Computer History Museum

Interview of Jerome H. "Jerry" Seltzer

Interviewed by:
James L. Pelkey

Recorded March 7, 1988
Cambridge, MA

CHM Reference number: X5671.2010

© 2016 Computer History Museum

James Pelkey: When I began interviewing, I told the interviewees: "I will make sure that the historical reconstruction will be factual." It turns out if multiple people recall the same incident ten years ago, there will be likely as many different stories remembered, and they're all likely real, they just happen to be different versions of the truth. The challenge of recreating the one truth is more complex than I once thought.

Anyway, I greatly appreciate your time. Before you continue the thought you were on, know that my primary objective is to better understand the events - and your perspective – of the activities within the Computer Department of MIT from 1976 to 1979 regarding the history of local area networking and especially the technologies of token ring and Ethernet. Secondly, I would to better understand the MITRE Conference in 1979 and its impact. Lastly, I am interested in your understanding of how and why IBM chose token ring over Ethernet.

To the extent that you have views beyond those you will already have expressed, I anticipate an engaging conversation, sir.

Jerry Saltzer: Well, first of all, I was about to observe that one of the things that seems to drive people's decisions in the area of communications is their decision of what computer to purchase - at least up until now, or perhaps until very recently - has been driven by who manufactures the computer, rather than the primary concern in buying a computer being: What's the computer going to be used for? So if they're going to use it for administrative things, they go perhaps to IBM. If they're going to use it for engineering, they might go to DEC or Sun or whomever. The primary concern has become: What am I going to use the computer for, and therefore, they look at the software and whatnot that comes for their application.

Pelkey: Right.

Saltzer: The question of how well it communicates has always been a secondary concern. They buy the computer and it then comes in with something that is perhaps on the back end that does a little communications. If they buy an IBM computer, it comes with IBM's view of communications. If they buy a DEC computer, it comes with DECs view. People haven't generally purchased computers with a high priority on the question of how well it communicate and a lower priority on what am I going to do with the computer. And that has affected a lot in the way of what it is that happens by way of decisions here. This business of whether or not a -- which technology is appropriate or whatnot has very little to do with the quality of the technology, it has to do with the marketing and accidents of which companies happen to have picked various directions, and positions that companies have had at various times as to whether they could dominate by choosing something that's different, or can they dominate by choosing something that's standard. There are very different approaches that different people have taken to the same problem.

Pelkey: I agree with you completely.

Saltzer: So I think that these kinds of artifacts, almost, determine what's going on, and as a result, you can't really trace why something happened back to a question of this was better than that. Instead you if you succeed in tracing it at all, you may discover that it was an arbitrary marketing decision made by a bunch of salesmen somewhere or market specialists somewhere, or finance people.

Pelkey: Which may lead into that issue of MIT in the '78, '79 time frame picking token ring versus picking Ethernet.

Saltzer: Well, I'm not quite sure what you mean by picking token ring. We have always had all kinds of technologies here.

Pelkey: Right.

Saltzer: In 1976 we installed our first Ethernet. Sometime before that we installed Chaosnet. The reasons why MIT does things are sometimes not obvious; they're not immediate reasons that you would expect. They are usually based on very practical considerations. The Chaosnet was invented and built for only one reason. That is that Xerox had invented and built the first version of the Ethernet, but still considered it proprietary and would not allow anyone to use the internal knowledge. The fact that it worked was sufficient for someone else to say: "Well in that case, we'll build one too," so we built a Chaosnet. It's the only reason. It's essentially another net similar to it that's got slight differences in principle and, but the differences aren't important enough to worry about.

Pelkey: Ok. And that was done in the AI lab?

Saltzer: Yes, our group of people in the AI lab needed to connect a bunch of computers. At that time they were LISP machines. They wanted a local net technology. They looked around. They said: "Well, Ethernet works. We've seen it run out at Xerox. Let's build one." But they built one of their own without the benefit of being able to use, directly, the exact Ethernet stuff.

Pelkey: Right.

Saltzer: Now, I believe that about a year after that, maybe two years after that, Xerox finally opened up, and also began giving copies of their equipment to some universities, and we received some. So the next thing we knew we had both an Ethernet, a real Ethernet, it was called Experimental Ethernet at the time, it was the three-megabit version, and we also had Chaosnet running on the campus.

Pelkey: And this was in, what, '76?

Saltzer: I unfortunately would have to go back to my files, which are in another building, to get the years, but it was somewhere between '76 and '78 when we had both of those nets running.

Pelkey: Ok.

Saltzer: At about that time, we discovered that Dave Farber was busy working on a distributed computing project at the University of California at Irvine, and he had been experimenting with token ring as a method of interconnecting his distributed computing system. So we were getting interested in research on local area networks as contrasted with actually building one. We were sort of curious about some of the various aspects of what they were good for and what technologies were suitable, so we looked at the Ethernet and said: "That's interesting. It is full of nice properties and full of funny properties that worry me a little." Then we looked at the token ring technology that Farber was proposing to use and had actually tried to build some prototype equipment, and said: "That's interesting. That has some nice properties and it also has some funny properties," and decided that we wanted to do some experimentation in this area, and decided that the easiest way to experiment was to build something that was in contrast with the Ethernet that we already had running, just to see whether or not our initial intuition that certain pluses and minuses that applied to both sides might be born out. So we got together with Farber and imported his one-megabit token ring under the name of Version 1 Ring. What happened was that he had done a prototype and had also done a much more elaborate design and was working on the elaborate design.

So we imported the elaborate design and finished it here, and ran it for a few years, among a half a dozen machines as a one-megabit per second token ring. And while doing that it was apparent, once we looked at it carefully and began thinking, it was immediately apparent that he had elaborated the design in some areas that we didn't think were important, and that the key objection that people had to token ring, namely that everything had to be repeated by a whole bunch of repeaters on its way around, was actually very easily overcome by a simple topological trick of running the wire back to a wire center in between each of the nodes, and that give you a place to go to look for problems. In fact, that was our concern with the Ethernet, that there was no one place to go to look for problems. And we realized that with the token ring, there was an easy way to get maintainability. So that was really reason why we were exploring the token ring. It looked to us like it had a maintenance advantage over Ethernet. And so we built this one-megabit per second token ring using initially this design that had come from Irvine, and we finished off the debugging of the design. Ken Pogran was the key person who did most of that work, and while doing that we concluded that it was clear that the technology was a winner from our point of view. It was just straight: this is a good technology for doing the job.

Pelkey: Ok.

Saltzer: So we started up a second project to produce a second -generation version called Version 2 Token Ring, which is the thing that later became a Proteon product.

Pelkey: Now, when you got Version 1, did you do that just on your own at MIT? I know DARPA got involved in doing some funding of an analysis project between a ring and Ethernet somewhere along the line.

Saltzer: The group that I was leading, that did the work, was funded by ARPA. So that may be ARPA's description of the work we were doing. I don't remember. It sounds like it.

Pelkey: Now Bob Kahn in recollecting about that period of time, thought that you were interested in doing something. He wasn't interested in funding another local area network. You were interested in doing something in local area networking and then there was this kind of concept of a study, that is comparing the properties and the performance and the characteristics and so on, of Ethernet versus token ring. It was something that made sense to you and it made sense to him. Is that your recollection as well?

Saltzer: It's consistent. I mean, the purpose of this was to understand the difference between the properties, to see whether or not we had a theory that token ring was easier to maintain.

Pelkey: Ok.

Saltzer: -- and the only way to check that out, given that you can't do models of maintenance very well, without actually having one, is to go build one and see whether or not it was easier to maintain.

Pelkey: Right.

Saltzer: So we first tried with the stuff coming from Farber --

Pelkey: Right.

Saltzer: -- and we learned a little bit from that about this question and concluded that you really had to have one that was comparable in more dimensions, that is, higher data rates and that eliminated some of the features that we found, in the design that had been done by Farber, it had about three times as much circuitry as it needed, because it was also solving a bunch of other problems, having to do with broadcasting and multi-casting and a bunch of other kinds of things. They were not that important to the basic thing, but they tripled the cost and size of the interface, and it was apparent that by itself made it a somewhat uninteresting, in an economical sense. It made it uncompetitive in an economic sense, and it confused the issue on comparisons.

Pelkey: Now, do I take it that you led this group that was doing this work?

Saltzer: Yeah.

Pelkey: Did you initiate the contact with Farber, or was it someone else?

Saltzer: Either me, or probably Ken Pogran did.

Pelkey: Ok, Pogran was a graduate student at that point in time.

Saltzer: He switched from being a graduate student to being a staff member somewhere in there, but I think he was on the staff by that time.

Pelkey: Ok. And was Dave Clark a member of the team?

Saltzer: Yeah.

Pelkey: He was on staff at that point as well?

Saltzer: Yeah, he was a research associate.

Pelkey: Ok. And who else was part of that team, if you will?

Saltzer: Those were the primary research leaders. There were two or three other individuals involved. Someone by the name of Larry Allen helped with the design. There was also a fellow by the name of Mel Kiapa.

Pelkey: Do you know where I can find either of those gentlemen?

Saltzer: Larry Allen works for Apollo. Mel Kiapa's a private consultant working somewhere out of Virginia, I believe.

Pelkey: Ok.

Saltzer: I'm trying to recall the extent to which those two might have been involved, you see, we were doing this work starting around 1976, and it went on to 1980, probably.

Pelkey: Ok.

Saltzer: Maybe '81. Well, in fact, in a sense it continued beyond that, because it continued right on through the time that we had Proteon set up and they were marketing this thing. I don't remember when that started, but it was fairly recent. So I just don't remember when they became involved. I know they were deeply involved. Both of them were deeply involved in the design and the implementation of the software for the second generation, the Version 2 Ring. I don't recall whether anyone other than Ken Pogran and Dave Clark were involved in the Version 1 Ring, although I suspect Mel Kiapa was around, and was involved in it somehow.

Pelkey: Ok.

Saltzer: The intellectual leadership I think came from Pogran, Clark and myself.

Pelkey: Ok. Now, when you started to do version 2, that's when you started to say: "Ok, we want to redo this. We want to do the electronics, the circuitry, the concepts involved in the ring to make it more robust, with higher data rates, to have a wire closet involved --

Saltzer: There was a wire closet in the Version 1 Ring.

Pelkey: Oh.

Saltzer: The basic concept that we figured out as soon as we looked at Farber's stuff was: "This could be made reliable, and the easy way is to make a wire center." It was just completely obvious. The minute we turned it on and began looking at it, we said: "Obviously the right way to run the wires is so that all the wires ran to a common point so you can get at them." At that point, we just built a passive wire center. It didn't have relays or anything in it. We just ran the wires to one point in the center and there was this little box and the wires all came in here and came out there. That's all there was.

Pelkey: How about that. And Farber hadn't done that?

Saltzer: No, as far as I know, this idea is one that originated with Ken Pogran and myself.

Pelkey: Ok, if I understand correctly, there was some issue about starting a net off, that you wanted everything to be the same, in terms of a sense of equality about the net up, a kind of philosophical approach to the problem.

Saltzer: Well, this has to do with some different designs people had as to how you might do rings.

Pelkey: Ok.

Saltzer: Whether one station is in charge of the ring and all the others are slaves, a master-slave kind of approach, or if all the machines are peers. So that is a question between two different ring design strategies.

Pelkey: Ok.

Saltzer: And there were some other ring design proposals floating around, and implementations, that used the other approach, the master-slave kind of approach.

Pelkey: You selected the peer?

Saltzer: Yeah.

Pelkey: Was it an arbitrary sort of decision or did it come down to --

Saltzer: Our assumption that it was fundamentally technically better because you don't have a single point of failure. You don't have to worry about what to do if the master dies. If everyone is equal, and all the algorithms are designed around an assumption that all machines are equally capable of talking to one another. The issue is the following: Do you use the same algorithms everywhere, or do you have one set of algorithms for one set of machines and another algorithm for a different machine for recovery fall and failures, and the master- slave approaches have the slaves respond to failures one way and the master responds to failure a different way, and it sort of doubles the number of things you have to think about.

Pelkey: Right.

Saltzer: That's all, a very simple idea.

Pelkey: Good. Now, if I understand correctly, somewhere around this tine, '77, '78, '79, DEC made an offer to MIT, to the Computer Labs, about putting VAXs in, but requiring that they be used as single user workstations, as opposed to being used in a multi- user environment. Does that ring a bell?

Saltzer: Yeah, yeah.

Pelkey: Did that come after this -- these concepts of the nets and so on?

Saltzer: Yeah. Well, the concept of the workstation was around well before Digital showed up, and it was the thing that drove the Chaosnet.

Pelkey: Right.

Saltzer: They were building LISP machines, which were basically the predecessors of engineering workstations.

Pelkey: Right. Victor Zue's group and the people were --

Saltzer: Excuse me.

Pelkey: Victor Zue, as I recall, he was doing speech research, and they had a LISP machine, one of the early LISP machines that they were doing speech work on over in the AI area.

Saltzer: Yeah, those were the users, but the people who were building them were years before that.

Pelkey: Yeah, ok.

Saltzer: The LISP -- the CHAOS Network Group was actually the group that did the LISP machines. It was really the LISP machine group, ok, and similarly Xerox pioneered the development of private workstations.

Pelkey: Altos.

Saltzer: -- the Alto, right. And that was the reason why they invented the Ethernet.

Pelkey: Right.

Saltzer: And we had Xerox Altos, and we had LISP machines and various other things, and the question then of whether or not the computer you're attached to is a mainframe timesharing system or workstation instantly becomes fairly uninteresting. All you care about is this is a computer and I want it to be on the net, so the fact that Digital came along sometime later with a bunch of VAXs and offered them to us to explore -- basically they were offering the opportunity to live in the future and pretend that you had workstation with the power of a VAX by giving us a bunch of VAXs and asking that we try to use them as workstation. So we did that.

Pelkey: So the VAX coming along was after all this token ring stuff.

Saltzer: It was the place that we applied the Version 2 Token Ring.

Pelkey: Was with the VAXs that were in house.

Saltzer: Correct.

Pelkey: Ok. When, now, Bob Metcalfe came to work at MIT as a consultant to Mike Dertouzos?

Saltzer: Yeah.

Pelkey: In the beginning of '79, as I recall.

Saltzer: Sounds about right, I don't remember the dates.

Pelkey: What was -- do you recollect what his role was and were you involved in that process with him coming on as a consultant?

Saltzer: He was basically looking for something -- a way to launch a company while -- and needed a place to live while being separated from Xerox.

Pelkey: Ok.

Saltzer: So we gave him a home.

Pelkey: And, was he arguing -- during this period of time, was there conversations or forums of, between his views about Ethernet and what was happening at Chaosnet and what you were trying to do with token ring, or was this kind of token ring experiment against Ethernet an independent event?

Saltzer: I would say it was largely independent. We didn't bring him in to tell us which one of those networks to use. His label consultant was an identification for purposes of making it clear he was not one of our regular staff members, he wasn't a professor or whatever it was. We needed a label for him.

Pelkey: Ok.

Saltzer: But basically, we gave him an office and a place to live, and a home while he was getting his feet on the ground trying to understand how to get going in the world, because he had separated himself from Xerox, and he knew he wanted to start a company, but wasn't quite sure exactly how to go about doing this yet --

Pelkey: Ok.

Saltzer: -- and he just needed a, you know, he was a friend. He was an old friend. He's been around for a long time. He did his thesis at MIT while he was graduating from Harvard, and so we invited him back, Dertouzos invited him back and said: "You can have a home in the lab while you figure out how you're going to proceed." So he lived in Boston for about a year, I think, a year and a half, something like that, and basically what he was doing, I think, was lining up venture capital and getting all the people -- all the ducks in order to be able to go out and jump in and launch a company.

Pelkey: Ok.

Saltzer: And, sure, we asked him questions, but we asked his opinion on anything he wanted to talk about because one of the reasons for bringing him into the lab was he had a lot of background and intellectual interests that matched a lot of our own.

Pelkey: Right. Now, during this period of time, do you recall some kind of a session in May of '79 in which, there was a forum in which some of these issues, I guess about the performance of token ring versus Ethernet were debated, or papers presented or any kind of a culmination of that sort?

Saltzer: Well, there was a conference over here in -- it was a local area networks conference sponsored by MITRE.

Pelkey: In May of '79?

Saltzer: I think. Sounds about right.

Pelkey: Copley Plaza or --

Saltzer: Yeah, that's it. That's probably it.

Pelkey: Can you talk about that? Did you participate?

Saltzer: We presented a paper there.

Pelkey: And the paper was on the performance of Ethernet versus token ring?

Saltzer: No, it was on the idea of star shaped rings.

Pelkey: Ok.

Saltzer: Pogran and I wrote the paper and presented it at that conference. That paper, as all papers are, the ideas in there were two years old, ok? That was the set of ideas that we figured out in '77.

Pelkey: Right.

Saltzer: -- and had implemented during the '77, '78 period by addition -- sort of a brute force addition to whatever the Farber ring stuff that we had, and by '79, I don't remember when it was that we commissioned Proteon to begin building something, but that paper in '79 described in some depth the design that we intended for the Version 2 Ring. I don't remember, for all I know we might actually have had it running by then, although I don't think so, I just don't remember the dates.

Pelkey: Ok. Do you recall that conference as having been particularly contentious or tense with a lot of debating about local area networking?

Saltzer: It was --

Pelkey: Can (unintelligible) other people or --

Saltzer: It was probably, as many of these things are, there hadn't been any conferences on the subject of local nets before that that as I recall, or any significant ones. And this one, for some reason, did catch a lot of people's attention. I think maybe, is that the place that Tobagi and Hunt reported their initial work on the performance analysis of the Ethernet, I'm not sure, you'll have to look at the proceedings to check that. But any case, we reported this idea of how you could make rings into a useful net. I think that perhaps the reason why that paper got some attention is that there had been, since 1974, maybe even earlier, there had been papers proposing token rings --

Pelkey: Ok.

Saltzer: -- and every time a person proposed a token ring, someone responded by saying: "Yes, but its fragile," and as a result, people would say: "It's silly." Propose;

silly. By '79, this was the first time we published a paper saying: "Look, fragility is not a problem. All you have to do is run the wires the right place."

Pelkey: Ok.

Saltzer: All of a sudden that did produce a fairly significant reaction. We found people at that conference were very excited. On the other hand that was not where the idea had begun. We had been using it for some time. By then IBM was using it, in Zurich. Well I spent the year of '77, '78 working for IBM, and --

Pelkey: In Zurich?

Saltzer: I was working in White Plains, but I traveled a bit, quite a bit, and one of my stops was in Zurich, and I discovered that they were interested in token rings there. So we proceeded to compare notes, and I basically had supplied them with every piece of input that we had learned, and I felt that they were busy designing a very high quality ring lock-up mechanism that was very nice. So I encouraged them to proceed with that.

Pelkey: (Affirmative).

Saltzer: And so they were deeply into this by that time already. It's not as if the ideas had not been touched or picked up anywhere else.

Pelkey: Do you recall why they picked up the ring as opposed to, for example, Ethernet?

Saltzer: I don't remember hearing a specific reason, however, the laboratory doing it was in Zurich. At the time, this was1977, I think the only local area net in all of Europe, probably, was the Cambridge Ring.

Pelkey: Yes.

Saltzer: In other words, I have a feeling that, if the Zurich laboratory had as its charter to do communications, and it was trying to figure out could it do something useful in the area of local area network communications. Ethernet basically was something that was being done 5,000 miles farther away than a token ring, or a ring of any kind. So, they began by looking over the various things, and they came across the Cambridge Ring almost immediately because they saw lots of people from Cambridge. It's not very far away. And they began looking at that, and it was harder to look at the Ethernet because no one had published anything on it yet. You know, you would just hear, in conferences or you would hear in various places that the people at Xerox PARC, and Xerox PARC is a competing company, so you can't get inside to look at it. They hadn't published anything on it yet, so they began looking at rings, and that's probably the reason. I don't remember any specific reason given. On the other hand, what I found was when I got there that they agreed with me -- see it's my opinion, it always has been my opinion, that if you were to start off with a completely free slate, no history behind you of any kind, and ask yourself which is intrinsically better, the Ethernet or the token ring, I believe the answer is the token ring. And I think that because they happened to have begun working in that area, they quickly came to that same conclusion themselves, that it's got some fundamentally superior properties.

Pelkey: Ok.

Saltzer: I wrote a paper on this subject for the Mexico City Data Communication Symposium. It was presented, probably '81. I'm not sure.

Pelkey: Do you know where I could get a copy of that paper?

Saltzer: Well, it's an IEEE well known symposium. It ought to be in any library.

Pelkey: Ok.

Saltzer: I have them in my -- unfortunately, I have two offices and copies are in the other office.

Pelkey: Is there someone who I could speak -- and along that same line, this MITRE conference in May of '79, IEEE?

Saltzer: MITRE.

Pelkey: MITRE. So I'd have to get that proceeding straight from MITRE.

Saltzer: Probably that comes direct from MITRE.

Pelkey: Ok.

Saltzer: If you're interested in exploring this at all, you ought to have copies of those proceedings of the conferences at that time --

Pelkey: I agree.

Saltzer: -- because they do shed a lot of light.

Pelkey: Yes.

Saltzer: There's a series of conferences sponsored by the IEEE called Datacom. They're held every two years.

Pelkey: Right, yes.

Saltzer: And it was probably, I believe they're in odd numbered years, and I don't think it was in '79 because that would have been the same time, so it probably was in 1981, it was held in Mexico City, and the paper that I presented was entitled "Why a Ring?" Ending with a question mark.

Pelkey: Do you remember anyone who was in the Zurich group of IBM who I could try to contact to understand their thinking and their decision making process and where the intellectual property came for the work they were doing on rings? Because I suspect you're right, that it was the Cambridge Ring was the influencing factor, and it was because it was close by and so on, but that's just a guess.

Saltzer: Karl Kummerle, K A R L K U M M E R L E

Pelkey: Do you know if he's still at IBM?

Saltzer: Yeah. Zurich. I think he was leading the work at the time, but I'm not certain of that. I'm trying to remember the names of the engineers. That work was being driven largely by two very sharp German engineers who were -- they were phase locked loop designers, and they were working at that level, and I'm trying to remember their names. Again, he may have had a paper. I know he was at the Mexico City conference. It's possible he had a paper there, in which case, the co-authors on the paper may also be the other writing people.

Pelkey: Good.

Saltzer: I know they have published work on the IBM Token Ring. In fact, we're talking here about the initial work done by the laboratory at Zurich.

Pelkey: Right.

Saltzer: Because there is a whole later wave, having to do with IBM's sudden discovery that token rings are an interesting idea or that local networks are important, and their Raleigh laboratory decided: "Ok, we've got to do something," and they proceeded to import the technology from Zurich and redeveloped it in their own way, and there are a whole bunch of papers since then, but they're all derivative, in the sense of involving much more recent marketing decisions and whatnot and technical decisions not really related to the origins. I mean, they may be related to another part of your study, maybe, for which this sudden interest by Raleigh is something you might want to pick up on.

Pelkey: Yes.

Saltzer: But Zurich is where the technology developed at the time.

Pelkey: Good. I venture Zurich is a little more interesting.

Saltzer: And the names in Raleigh are easily derivable by studying the IBM Systems Journal and various other papers on token rings.

Pelkey: Right. Did you know about Soderblom during this '76, '79 period of time?

Saltzer: First time I ran across that name and the patent claims were sometime in, quite recent. I think, probably '81 or '82 is when I first heard about that.

Pelkey: Ok, so that whole concept was, I mean, you were just pursuing the technology and --

Saltzer: Yeah. I had never run across Soderblom. Since then I have received a copy of his patent and patent claim. I'm a little surprised that the patent was there and allowed, because it struck me when I looked through it, that this was a patent worded, what '76 or '77 I believe, for something that I thought was common knowledge in '74, so I didn't understand how this worked.

Pelkey: Maybe his argument would be the same thing as your argument about papers.

Saltzer: Oh, it could well be. I don't know what the dates -- what the priority dates are. I thought the priority -- the dates that I've seen on it seem to be well after the dates that I thought that this work was going on. The point is that Farber's work was all open and published.

Pelkey: Yes.

Saltzer: And therefore, by itself, not the basis of a patentable arrangement…

Pelkey: Ok. Did you -- were you aware of any of the work that Newhall and I guess it was Pearce did at AT&T on rings?

Saltzer: Yeah. That was even earlier.

Pelkey: That was in, what early '60s?

Saltzer: I don't -- I'm sorry, I completely lost there. Pearce proposed the idea of a ring a long, long time ago as a way of doing data, communications and telephone communications and all kinds of stuff. It was a far-reaching vision. And Newhall then picked up on, and there were several papers written and published in technical journals, and in fact I think it's possible that either our '79 paper or Farber's 1974 paper, I'm pretty sure, refer to the Newhall papers and the Pearce papers.

Pelkey: Ok.

Saltzer: I'm pretty sure that Pearce's work was theoretical and I don't recall whether Newhall actually built something. The immediate result of all of that was the next guy to pick up on that at Bell Labs was Sandy Fraser, who built a ring network named Spider.

Pelkey: When did he do that?

Saltzer: I'm sorry, but I believe that was probably available to be reported at that 1979 conference also.

Pelkey: Ok.

Saltzer: His was definitely done with a central monitoring station rather than relying on a peer kind of arrangement. That was the main characteristic that caused us to think that he probably was not quite on the right track.

Pelkey: And they never did anything with it?

Saltzer: He had a hard time selling his management.

Pelkey: Why do you think that was?

Saltzer: Don't know. It's very hard to sell local area communications in a world that didn't believe that this was important yet. AT&T was not the leader in the business of telecommunications for networks. It had never been the data communications leader.

Pelkey: Why do you think that was? I mean, at one point they were, because, I mean, the modem business was theirs, other than for select military, so I mean, in terms of data communications --

Saltzer: I think it's only been within the last year or two that it has suddenly occurred to them that more bits may move than voice.

Pelkey: Right.

Saltzer: We received this --

Pelkey: They were still the market share holder in '65 or thereabouts, and I mean so --

Saltzer: Yes. They understood that you did have to move some data, but whenever we talked to them -- I mean we got into the data communications business at MIT in a big way back around 1962.

Pelkey: Ok.

Saltzer: When we began trying to connect up our terminals to out timesharing systems, and we discovered that there was no good way to do it, and we found that the only technology was these Bell System modems and dial up, and we discovered they didn't answer the question completely by themselves, because there were issues of concentration and holding time in the exchanges and things like this that were to directs. All that modems gave you was a point-to-point communication from some point, and from a system point of view, there were still some severe problems. We went to the telephone company and began trying to talk these problems through with them and they took an attitude then that they have taken, really, up until very recently, namely: data is an unimportant part of our business.

Pelkey: Right.

Saltzer: It's a piece of the business, but unimportant, and therefore they didn't put their best engineers on it. They didn't put any very good thinking on it, and so on. And that has colored their thinking, I think, very strongly, and certainly by 1976 or '79 or whenever Sandy Fraser was working on Spider, although HE was certainly convinced otherwise, I think that he was operating in the middle of an environment that still felt that data communications was a little tiny tail on the side of a very big voice communications dog.

Pelkey: Was Sandy at Bell Labs?

Saltzer: Yeah.

Pelkey: It is kind of strange that Bell Labs, I mean, the couple times I've been through it, I mean is an environment that strikes me as much like an MIT environment or a Xerox PARC in its heyday --

Saltzer: Sure.

Pelkey: Lots of very creative, energetic, resourceful people, and lots of computers around and why don't we connect these things to the other and do something? And why some of these technologies, in terms of data communication, just never emerged out of the Bell Labs environment, is an enigma to me.

Saltzer: Well, the comment I made at the beginning applies here. The quality of the technology and the quality of the thinking is not the prime thing that determines any other of these questions. What determines these questions are a couple of practical things. Who decides to go to market, and if you were to look over the various technologies, you might look at the Spider system. Sandy Fraser has since put together a essentially more interesting system called Datakit, which is a data communication local area network kind of thing. Suppose you were to sit down and just look over the technology issues and you were to decide Datakit is the right way to do this. They've really got the right idea. They understand the telecommunications technology and it solves maintenance problems better than anything else. You have, in your site, four VAXs, three IBM machines, a whole bunch of IBM PCs and two Suns. You say: "I want to put in a Datakit." Well, you discover you can't, because AT&T does not market interfaces for any of your computers. You have one 3B3 -- 3B1 over in the basement. Maybe you can get a Datakit interface for that, but when it comes to being practical, saying: "I want some communications," what you do is you go out and find someone who has decided to make the investment in the interfaces that you can buy.

Pelkey: Right.

Saltzer: The reason why PCs get linked with Ethernets or with Novell cards or with whatever they are is because you can get cards to do it. And so that dominates the decisions. No matter how much -- even if you think that Datakit is a whole lot better than Ethernet or Token Ring, you can't get one for your PCs and -- or actually you CAN now, but it turns out that up until I think a year ago that wasn't true, and that is the thing that, I believe by itself, it's marketing decisions by people who have got other interests than: is this the best technology, in mind.

Pelkey: Right. Let me generalize for the moment. I mentioned that I have read Kuhn's book, in terms of this paradigm shift.

Saltzer: Yeah.

Pelkey: This idea of these discontinuities, punctuated equilibrium, if you will --

Saltzer: Incidentally, there's another thing that, perhaps, is fairly significant here, before we get to that.

Pelkey: Right.

Saltzer: The question really, in comparing different technologies, has to partly be based on how different they are. How much better is this one than another one? If you have a dozen workstations, one mainframe, all within a building that is managed by one organization, and you say: "Which is better, Ethernet, Token Ring, Datakit, MAP, whatever it is the broadband protocol that comes from somewhere else?" If you do a very careful technical study of all of the things and you really do it right and you understand everything, you'll come to the conclusion that it doesn't make a whole lot of difference. All those technologies will work pretty well.

Pelkey: Right. Yes.

Saltzer: And you therefore make your decision on things like I was just describing.

Pelkey: Availability.

Saltzer: Who makes hardware interfaces for all the machines I happen to own? If there are three possible venders have the interfaces; who has the best delivery and the best price this week? Or questions like those. Or, I'm thinking of buying some machines in the future and they only use one of these technologies and that's all. So the decisions are made on that basis once you conclude that it's a wash on the technical issues. Now, the problem is the industry grows up by people starting --

Pelkey: Right.

Saltzer: -- with: "I have a dozen workstations," or: "I have five computers in my basement and I want to link them." Small scale. If you start looking at the upper end of the scale --

Pelkey: Right.

Saltzer: -- the thing that distinguished -- let's switch to a whole different paradigm for moment. The thing that distinguishes AT&T's approach to installing telephones is, always has been, that they think in terms of: "Anything we do has to fit in with a system that already has 19 million other telephones in it. You don't just casually go changing the standards because you've got to have the whole system worked together. They think of it from a system point of view, therefore a very large-scale point of view. Now, the problem is that on the computer side, people began thinking from a small-scale point of view. The technology differences, I think, show up as you raise the scale. If you have a thousand workstations to be scattered across -- the problem we were attacking, initially, is we looked ahead and said: "Look. MIT fine. So we have got 50 or 100 workstations today. Obviously in the future we're going to have thousands of workstations. They're going to be all over the campus. There's going to be as many as there are telephones, probably. PC on every desk, except it isn't a PC anymore, it's a big computer, it just happens to be packaged in the size of a PC, and it's priced like a PC, but it's a big computer and it wants to communicate. So what do we have to do? We have to think about the future where you don't have ones and twos and tens linking together on one floor, you've got thousands scattered across the campus."

Pelkey: Right.

Saltzer: If you evaluate that one carefully, that's where you begin to discover that technologies like token ring make a big difference, because maintainability, ease of finding troubles and what not begin to dominate your cost and your other considerations. So I think that also has an important effect here. Because the industry began by people linking small numbers together, at that level, most of the technology questions were a wash.

Pelkey: Ok.

Saltzer: And it didn't matter. It still is true. It doesn't matter.

Pelkey: So that --

Saltzer: Go ahead. You were going to --

Pelkey: In listening, I'm not sure I -- I'll have to get back to exactly the question I was asking, but it was about the kind of the paradigm shift. The paradigm shift at that time then was kind of local area networking, as opposed to specifically any given embodiment of it.

Saltzer: Oh, ok.

Pelkey: Somewhere along the line there came this idea of local area networking was a more powerful way of looking at connecting things together than it was token ring or Ethernet. Those are subsidiary issues. I mean, they came to play in terms of marketing --

Saltzer: Sure. Sure. Right. Right. Right.

Pelkey: -- and commercialization --

Saltzer: So, there was a more interesting question of whether you should be installing a local area net or should you be installing a bunch of modems? That was a different kind of question.

Pelkey: Yes, and the guys who were on the modem side, and the traditional data communications side, I mean completely missed out on this idea that there were going to be workstations and that there was a better way of connecting things together than modems.

Saltzer: Ok. Right.

Pelkey: Why do you think they missed that paradigm shift?

Saltzer: I think a lot of people missed the personal computer revolution.

Pelkey: Because they weren't believers in the use of computers?

Saltzer: It's a revolution that is so substantial. I mean, look who's making money on personal computers.

Pelkey: Board manufacturers.

Saltzer: Is it Burroughs? Is it Honeywell? Is it RCA? Is it Univac? I mean we had a dozen big computer vendors --

Pelkey: But mostly those guys missed out on minicomputers and just disappeared -- I mean they missed out on the minicomputer before they --

Saltzer: Fine. Let's look at the minicomputer vendors. Data General.

Pelkey: DEC.

Saltzer: Digital. Are they making any money in personal computers? No they missed it. The point is that almost everybody missed the personal computer revolution. IBM caught it by virtue of --

Pelkey: Luck.

Saltzer: -- its board chairman setting aside a little organization in Boca Raton and said: "Look, I think there's something interesting here. You are free of all the restraints of the corporation. Pretend the corporation doesn't exist." He bankrolled an independent company.

Pelkey: Right.

Saltzer: Almost, I mean, he was lucky.

Pelkey: Skunk works.

Saltzer: He bankrolled an independent company, practically the same as the Apple people were bankrolled as an independent company, so a lot of people missed the idea that you can put on your desk a computer bigger than the old mainframes, and that is the driving force behind the personal computer. Now, the people at Xerox PARC did not miss this. They're the ones who were among the first to realize that you might have something useful, and of course they immediately recognized the importance of communications. The other part of it then is, once you have a computer talking to a computer, it's apparent that higher data rates that can be useful.

Pelkey: Yes.

Saltzer: The data rates that you get out of a modem, that is the traditional modem industry of 9.6 or 2.5, whatever they are, things under 20 kilobits per second or under 60 kilobits per second, are suited for an interactive mode in which there is a human being at one end of the wire and the computer is at the other end of the wire. The minute you have a computer at both ends, then it's possible for the computer to take significant advantage of being able to use higher data rates. And that is the driving thing that caused -- people who figured out that they wanted a computer on their desk, also figured out that they wanted higher data rates. People who didn't understand they needed a computer, couldn't see any reason on earth for a data rate higher than 19.6 kilobits, cause all that would happen is the screen would paint a little faster, not enough to make a difference.

Pelkey: Now, what was there about the engineering environment that drove, or if you will, what was the application that drove the need for higher data rates if there was a recognition: "I want to speed higher data rates around." What was it that they wanted to speed around, in terms of moving more data, that caused it to be accepted in the engineering environment?

Saltzer: Why are you moving data to begin with?

Pelkey: Certainly, one are files.

Saltzer: No, no. The reason you used to move data was to move it from the computer to the screen.

Pelkey: Ok.

Saltzer: Now, when the computer is on my desk, the path from the computer to the screen does not go through a modem or a telephone wire or a local area net. It's directly wired.

Pelkey: Right.

Saltzer: So, in other words, the point is a desktop computer leads you to a whole different architecture of what you're doing.

Pelkey: Right.

Saltzer: It's computers talking to computers, and therefore you have to ask the question -- you don't say: "Why did it change? Or why is it higher than it was?" Because it isn't higher than it was. Computers didn't ever talk to one another at low rates, because back at the time -- at that time computers weren't talking to one another.

Pelkey: Yes. Gotcha.

Saltzer: That's not completely true. That's an extreme version of the state, but I think to a first approximation it's pretty close. Yeah, we've always had computers on the east coast and computers on the west coast sending stuff via whatever speeds they could get sending the payroll information from one end to the other, and generally, those computers actually had trouble because you have a 9.6 kilobit line and it takes all night to send the payroll. It was a big hassle to do so.

Pelkey: But I think your point is obviously well made relative to that the computer to the screen was always there. It's just that when you have lots of computers, you'd want to have a computer other than one that was sitting immediately in front of you spatially to be able to put something to your screen, and that's a lot of data. But it was kind of really the kind of a productivity tools that strikes me, the CAD or CAE types of applications --

Saltzer: That's cause you're still picking this up from the wrong side. The fundamentals are that the computer . . .

Tape Side Ends

Saltzer: The revolution in which computers were replacing workstations coincides with the conversion, or the change in interest from modem, low-speed data communications to local area net communications, and so what you’re really asking is a different question, namely: "Why did people want personal computers?" That's the fundamental question you're asking. If I use a spreadsheet on a timesharing system with an old fashioned terminal on it, I can use it at 9600 baud.

Pelkey: Right.

Saltzer: I do not need high data rates.

Pelkey: Right, but --

Saltzer: So it isn't a personal productivity tool, that has very little to do with it. You asked the question: "Why do I have a personal computer?" That's the question you have to ask in order to be able to --

Pelkey: Yeah, but that was -- wasn't that really a derivative of a timesharing environment? You wanted more response time. You wanted to be more interactive. You wanted a faster response if you will, and a personal computer was that -- you had control over your environment and you've got fast response time.

Saltzer: (Affirmative).

Pelkey: Now the concept of connecting it to some other computer was a separate issue from the personal computer.

Saltzer: Right.

Pelkey: And the personal computer was an embodiment of, I guess out of the timesharing environment, people started getting used to interactive applications --

Saltzer: Right. So now I want to share information. That's fundamentally what the network is for, for sharing information.

Pelkey: But until there was something that I wanted to share -- I mean back when personal computers first came on the scene, there wasn't any -- I mean, what were you going to share with each other?

Saltzer: Software.

Pelkey: But you didn't need a local area network to share software. You could have done that over --

Saltzer: Your response is a little bit like the response of the London business community to the idea of telephones. I mean, they said: "We have messenger boys, and that's sufficient." You do want to communicate. Whether you can do it with a primitive method or a fast method is a question of how important the communication is to you and how cheap it is. So I --

Pelkey: I agree.

Saltzer: You really DO want to communicate, and the local area network is a way of doing it smoothly and fast and easily.

Pelkey: Yes.

Saltzer: And the things you want to communicate are data. You want to share information.

Pelkey: But why did that happen to the engineering environment. I mean, it was in the engineering environment that local area networking found its home, in university settings and the engineering environments within the universities. Was it something unique about that or was it just because it was computers and it was computer engineers and they wanted to connect them and it was --

Saltzer: Well, let's see. In the AI lab, for example, they built a bunch of LISP machines. It was very apparent to them that, it's always apparent to anyone who builds software that there isn't enough time for you to develop all the software that you would possibly want to use.

Pelkey: Right.

Saltzer: Ok, nobody, since 1960 perhaps, no one has ever taken an approach that said that I'm going to write every piece of software that runs on this computer.

Pelkey: Right.

Saltzer: You bring in compilers, you bring in spreadsheets, you bring in word processors, you bring in database systems, you bring in all kinds of -- most of what you run, someone else wrote, so you’re sharing programs.

Pelkey: Ok.

Saltzer: And I think they recognized immediately that the main activity that they were involved in was sharing programs. I think the people at Xerox PARC, during development time, realized that that's what they were doing. They also realized their future users were going to share information, because the timesharing systems had proven -- see the initial timesharing systems were giving each user a share, and independent share of a big computer. At the time they were considered big computers. But they were quite independent and couldn't communicate. Then it was very quickly discovered that one of the main things that you can get out of a timesharing system is, if you've got two users, if they can exchange files, it's a big win.

Pelkey: Right.

Saltzer: So the concept of the value of sharing had been established very fast, sharing information --

Pelkey: Right.

Saltzer: -- had been established very quickly within the timesharing environments, because timesharing initially didn't have that as one of its goals, and that got added very soon after it was invented, because it was discovered to be so powerful. So once you had independent workstations, immediately people said: "Well look, these are independent. We've got to get the sharing back," and the local area network is the answer.

Pelkey: Ok. Now, let me ask you a question of a different kind. When ARPA created its original spec, it did not include electronic mail.

Saltzer: You mean for the Arpanet?

Pelkey: For the Arpanet. It was to be a sharing of resources --

Saltzer: Resource sharing, yeah, right.

Pelkey: Resource sharing.

Saltzer: Yeah, but it had a problem, that every time that someone got an ARPA contract, at another site, they had to go buy a computer, first thing, and those were expensive, so they figured this might be a way to allow a site to come on board, take on a contract, and get nothing but a communication system and use one of the computers somewhere else.

Pelkey: Actually, I think that's not even technically correct, cause the TIP wasn't even part of the original spec.

Saltzer: Ok.

Pelkey: The original spec was on the IMP. And then started to realize: "Wait a minute, I've got to be able to make a TIP available," so it was really meant for the MITs and the Stanford’s and the UCLA’s that were the dominant sources, and in fact the original spec only had for one computer per IMP. Then they realized that: "Wait a minute. Everybody has multiple computers," and it turns out that most of the IMP traffic was local.

Saltzer: Right. Yeah.

Pelkey: Different than being on the net.

Saltzer: Right.

Pelkey: It never got to the subnet. Now, given your line of thought, given that E-mail was never part of the original spec and was kind of thrown in by some guys at BBN who were doing 10X, some years later. They learned that they couldn't kill the Arpanet experiment because most of the transactions going over it were E-mail. The total amount of traffic because you were passing big files and you were doing experiments, but if you look at data, you could still say it was files, but in terms of the number of transactions, E-mail swamped out file traffic and no one ever expected that.

Saltzer: Ok.

Pelkey: Now, why did that happen, if in fact the real concept was resource sharing, but it became a really a people-to-people, as opposed to a computer-to-computer.

Saltzer: It was timing. When did that occur?

Pelkey: It started '69 to '71

Saltzer: '69. That's the point. In 1969, people still thought timesharing systems were to share the resources of the machines.

Pelkey: Ok.

Saltzer: As of 1969, the idea that a primary goal of a timesharing system -- look at it this way. Timesharing was invented in the sense of just being available, roughly 1959, and it wasn't really used -- it was only in a research environment at that point. It was available to people in a sense of people other than the creators of the timesharing systems, roughly around '61, '62, '63.

Pelkey: Right.

Saltzer: It became something that a non-specialist could get his hands on and wanted to try it out. It was the arrival of the Multics system around '65, that is, the design was specified in '65 but it wasn't really available to anyone to use until '69 or '70, that the idea that sharing of information across the users of timesharing systems, with things like electronic mail and communicating software, was the right way to design things, to make the central design. The Arpanet was created at a time when it was still thought that the main issue was sharing the hardware. Of course, today we are able to buy computers approximately a thousand times cheaper than they were in 1960. If you think about a factor of a thousand in the price, if someone were to tell you that your automobiles 30 years ago cost a thousand times as much, you would not be surprised to find that people were talking a lot about methods of sharing the automobile. Ok? That's fundamentally all that's gone on now. Very simply. The price of computing cycles has dropped to the point where sharing the resources in the sense of: "We've got this expensive computer and we can't possible afford to pay for it unless we distribute it among 30 people or 100 people," that consideration has just vanished from people's thinking by a factor of 1000. All it takes is a factor of 1000 to make people think differently about little things like that, but that's why the timesharing people back in 1961 were focused on distributing the processor. That's why the Arpanet people when the put it together were focused on sharing the big computers at MIT and Stanford and Berkeley.

Pelkey: Yes, right.

Saltzer: You see, the slope had already started to come down. In 1961, computers the size of my PC here cost a million dollars. But by 1971, they had come way way down. They were already a factor of ten down that scale, but it takes a while for changes like that to get through everyone's thinking. When you make a big ground rule change, it does not instantly cause everyone to realize that the world has changed around them. When you do a factor of ten three times in a row, in three decades, it finally begins to get to the point where you just can't avoid it anymore. I think that's what really started happening.

Pelkey: Ok. I want to come back to that later, but first: You made a comment in the sixties that MIT was a communications center, in terms that Shannon was here,

Saltzer: I didn't make that comment. I think someone else did.

Pelkey: I thought you said early in the '60s that MIT was very strong in communications. In fact you went to AT&T. We got to talking about AT&T going on about modems --

Saltzer: I may have misstated that. What I was trying to say is the following: In 1961, we began to realize that the basic problem of building a timesharing system was going to have to get itself involved with communications because we wanted to put the terminals someplace other than the computer room, so we suddenly began having to interact with the telephone company.

Pelkey: Gotcha.

Saltzer: This is completely independent with MIT's position with respect to research in telecommunications.

Pelkey: Ok.

Saltzer: At the same time, there happened to be world famous communication theorists around who were doing information theory that had absolutely nothing to do with this.

Pelkey: And did those of you in the computer science side ever go access and deal with on an intellectual development, the talent that was over in the telecommunications side or was it different universes?

Saltzer: Well, I think everyone was aware of one another and there was considerable contact, but they were working on another problem. They were working on coding.

Pelkey: Yeah.

Saltzer: And coding is directed toward a different class of issue. Coding is trying to get the bits through in a noisy environment, and you're fighting the elements as hard as you can and what-not, and in general, actually, MIT did not have much by way of presence in telephone communications more ordinary kinds of things. There was the stuff that was the basis for trying to develop computer and data communications. So we actually had to scramble. When the time came to configure the telephone system here, we indeed found someone on the faculty who knew how to configure telephone systems because he had designed them back in the '30s. That was the last time that MIT had really been into that area.

Pelkey: Gotcha. Did you go to the scenarios, the ARPA scenarios in '72 when they demonstrated Arpanet?

Saltzer: When they demonstrated Arpanet.

Pelkey: The Arpanet for the first time. They had this conference down in Washington, D. C., at the Sheraton or the Hilton.

Saltzer: Oh. No, Dave Clark was at some of those. I never got involved. I ran the group, but the people in the group participated at that level.

Pelkey: That was, from what I'm hearing about looking at seminal events where people got together and there was a lot of energy and important in terms of exchange of ideas because a lot of people came together intensely --

Saltzer: Yeah.

Pelkey: That was credible. I mean Arpanet became real, as well as you had all these people who were, presumably, many of the leading experts on computer to computer communications and so on, all in the same room in the same time.

Saltzer: Ok, although, in a sense, it's hard for me to tell whether or not the media event is more important than the actual breakthrough, which occurred probably a couple of years before then.

Pelkey: Right.

Saltzer: The thing that I found interesting was when we first got the first three or four computers running on the Arpanet, and we were able to send messages from a computer at MIT to a computer in California, at Stanford or wherever it was, and I thought that was where something significant had happened, and in order for that to happen, many of the same people had to have also gotten in the same room several times, but they were smaller groups. They were just the groups that were really doing it.

Pelkey: Right, right.

Saltzer: So in that sense, that activity had been going on, and I think we got that going in 1970.

Pelkey: Yes.

Saltzer: Maybe there were three nodes on the Arpanet in 1970, which allowed some communications to begin. And by the time that demonstration was being given, what was really going there was staging it, in order to get attention --

Pelkey: Right.

Saltzer: -- in order to develop the political support for getting more funds into it, which is fine. And it also has the side effect of bringing a lot of other people in and other experts showing up and saying: "This is a good thing," and all that kind of stuff.

Pelkey: Now, you had similarly characterized the conference at the Copley Plaza in '79 as exactly the same thing.

Saltzer: Right, right. The neat things that happened -- Well, there were two things. The neat things that happened were -- people who needed to know about LANs already knew about them.

Pelkey: Right.

Saltzer: The communications, in a technical sense, among the key workers in the field, had already taken place, generally, and the conference was a place where, yeah, people got together and exchanged -- you know, shook hands and friendships began. A lot of people outside the central --

Pelkey: Circle.

Saltzer: -- dozen, circle, at that time, were able to catch up on the events, and that's perhaps its importance.

Pelkey: Right, cause that was -- I mean, going back to that event, I characterize that '79 conference at the Copley Plaza as being a really important -- when I look back upon this last 20 years, there's only so many times I can find where the was a really key inner circle coming in contact with the much broader circle and it was done in a way in which there was lots of energy and it really set the tone for things, and it strikes me as that '79 conference was one of those, was it?

Saltzer: I don't have a whole lot of comment on that, cause I'm not sure I'm in a position to judge the fallout.

Pelkey: Why did you pick Proteon?

Saltzer: We needed a contractor that knew a little bit about getting bits down wires, and they had done some contracting for the Navy for various little communications systems. Mike Dertouzos happened to know the head of the company, and brought him in, and we looked over the problem, and we concluded they would be able to do the design as well as anybody.

Pelkey: The fact that you had been exposed earlier to these smart German engineers who were phase lock loop engineers and that that, in fact, was Howard Salween's, that was an expertise that they had, from your perspective, did that run into any conflicts?

Saltzer: It wasn't earlier.

Pelkey: But Proteon wasn't engaged until after you came back from IBM, right?

Saltzer: I'd have to go back and check my notes on that. I'm not convinced of that fact. I don't know when we (unintelligible).

Pelkey: Unfortunately, I've been doing all this interviewing, but I haven't gotten the transcripts, I apologize for being a little unsure myself.

Saltzer: Unfortunately, I don't keep the dates like that carefully in my mind because I never worried that much about it. It is possible that we engaged them after I had encountered the people from IBM. I'm not sure, but that was not the primary consideration. The primary consideration was one of simply -- they seemed to have -- it struck me that -- you see, the following thing struck me, that building an Ethernet, especially building it to the specs that Digital, Intel and Xerox had made for the 10 megabit Ethernet, was apparent that it was very difficult to do. You had to have a real top-notch engineer to do that. I mean, it was hard, because they pushed the specs a little too far, I think. I wrote a paper on the subject shortly after they brought the specs out and sent it to my friends at Digital and Xerox because I thought that they had probably made a mistake. I still think they made a mistake in pushing the specs too far, because the result was that you could not build an Ethernet interface cheaply.

Pelkey: If they had stayed at three or five or something like that - -

Saltzer: It they had stayed at five, or even three, would even have been better. With three megabits per second, they could have put the whole thing on a chip.

Pelkey: Yeah.

Saltzer: Single chip, very simple, anyone could have done it.

Pelkey: Right.

Saltzer: And the result would have been ubiquity, I think, it would have been an unstoppable revolution that would have followed that, because they would have been everywhere.

Pelkey: Yeah, it would have been a hundred dollar, two hundred dollar per port --

Saltzer: Much earlier. At this point people are finally figuring out how to get it down there. The basic trouble was that, the minute you put the ten megabit spec on there, some circuits had to run at 20 megahertz. 20 megahertz means you have to have ECL. It's a hairy logic.

Pelkey: Yeah, it’s different.

Saltzer: So, it uses up power and it's got different voltages associated with it and everything else, but the fact that there had to be a few ECL chips to make it fly, it meant that you could not just make a chip of the whole thing neatly, and so I think that was the problem.

Pelkey: That's an interesting perspective. That's the first time I've heard that expressed, by the way.

Saltzer: Oh.

Pelkey: I agree with you.

Saltzer: Several of the guys you talked to should have received copies of the paper that I wrote back in '79 or so.

Pelkey: I'll have to see about my getting --

Saltzer: -- whenever it was that they came out with the spec, within a couple of months after that, I wrote a memo and sent it back to both Digital and Xerox, saying: "Gee, you guys have pushed over the limit here. Good idea but, you can't push it that far without -"

Pelkey: Gordon Bell’s agreed to sit with me.

Saltzer: Yeah.

Pelkey: After he gets his product here launched later this month.

Saltzer: Ok. He might have a copy of that memo, although he may not be able to find it any faster than I can at this point, its ten years old. Anyway, the point is that, compared with doing an Ethernet transceiver, or doing a ring repeater, any engineer ought to be able to do it, and therefore I didn't consider the choice of finding a super engineer to be an issue. It was just to get someone who could crank out the card.

Pelkey: Someone who was convenient, that was there, had the experience.

Saltzer: They were a contractor who had had some experience in doing contracts and turning things out for the government, and therefore knew how to meet specs and so on. We figured, just contract (unintelligible) work. That's all there was to it.

Pelkey: After that event, the period of time after May of '79, did you move away from this field and moved into something else? You were still in contact with it, but in terms of your focus of activities, that is local area networking, you moved away from this.

Saltzer: Oh, I'd have to look back to see what I was doing there. I don't think I moved that far away from it. We continued a lot of network research at the time.

Pelkey: Yeah. Let me ask the question differently. Did anything after May '79, did anything else happen within your group that, I have no visibility of it, that would be important to the story that I'm trying to communicate.

Saltzer: The prime things that we were doing then was we exploring a set of closely related issues, which is interconnection of local nets, to get them to the next larger size of scale.

Pelkey: Right, inter-networking.

Saltzer: Because the individual local nets are good for a single building and a few hundred workstations, if you have a few thousand workstations spread across the campus, you by definition have a whole bunch of local nets and you have to worry about how to link them, and that's the area we were beginning to work on. We started that work, actually, even before that. We started that work in coincidence with the time we started the ring work, and the spine network, the main network that visits the main buildings at MIT is a token ring. The MIT campus net consists of a 10 megabit Proteon fiber token ring that visits 30 different places on campus, and there's a gateway in each of those places to a local net in that building, where all the local computers are connected, with probably, at this point, a couple thousand computers on our campus net connected up that way. So we were exploring that technology, the issues of how to design a gateway, how to make a gateway fast. The issues of the protocols in order to make use of things without beating the gateways down and so on and so forth. So it was that set of ideas we were into.

Pelkey: In '72, '73 at Stanford, Vint Cerf started to work on TCP/IP that was driven by the need to --

Saltzer: They wanted to connect the radio network --

Pelkey: Right, to a local area network, and then you had all kinds of problems because NCP wasn't meant to do that, and in fact, XNS, if I understand it, is a much better protocol for local area networking in terms of speed. It was designed for a coax cable, as opposed to TCP/IP that was designed for a 50 kilobit line.

Saltzer: That's a curious way of putting it. I don't think that there is a significant difference between XNS and TCP technically.

Pelkey: It's amazing how many people believe, or have the view that XNS is superior. It's just that Xerox didn't let it out into the open environment.

Saltzer: It was a real, real problem that they had in marketing that. They really screwed up.

Pelkey: And then TCP/IP, it was meant to be an internetworking protocol.

Saltzer: So is XNS.

Pelkey: But for local area networking only, where as TCP/IP was meant to --

Saltzer: This is just -- partisans, I think, are able to put a whole lot of favorable spin on these things, but the differences, if you stand back ten feet from TCP and XNS, and ask what are the differences, its almost impossible to tell the difference. I mean they technically are solving the same problems, they do it almost the same way. The differences are at the lower level when you look, it's just like saying: "Do we have stainless steel or copper plumbing in this building." And when you flush the toilet you really don't care. When you get down underneath the floor, you discover that it's harder to drill thought the stainless steel and the copper is easier to bend, but, you know, that's what's going on here. TCP/IP and XNS are so similar, when compared with all the protocols in the world, that they should be viewed as identical. It's just an arbitrary decision.

Pelkey: What happened, in your group then? Were there protocols on top of some nets here? Did you just take what was available elsewhere, or did you develop your own protocols and then you replaced them later?

Saltzer: Well, several things happened. The AI group needed something to communicate among their LISP machines, so they developed their protocol to go with the Chaosnet know as the CHAOS protocol, and that actually became a dominant protocol because they had the most workstations in the early internetworking stuff at MIT. At the same time we got involved with the TCP/IP stuff, and Dave Clark is the TCP network architect, and has been for several years. Back in 1974, when, wasn’t it Vint Cerf published the first paper on TCP?

Pelkey: Right.

Saltzer: We looked at that, saw it coming, and a fellow by the name of Dave Reed, who is now at Lotus, who was in my group at the time, he was an assistant professor at the time, looked over the TCP spec and said: "This has a major flaw in it. The major flaw is that it does not really handle separate networks properly," and he proposed an alternate protocol to be called DSP, which basically was to separate the TCP and IP into two layers. And he made this kind of proposal. We went back to the ARPA community and said: "You've got this fundamental flaw in TCP, that, unless you fix it, we won't be able to use it at MIT cause we see we've got lots of local nets that we have to link together." The result of that was that TCP suddenly sprang two layers: TCP and IP, and it became the standard, and the minute that happened, we adopted it again.

Pelkey: Gotcha. Now, in TCP/IP, although it was out and about, it really didn't become dominant until it got distributed in Berkeley 4.1.

Saltzer: Right, but that's an interesting kind of thing in which there were just too many network protocols out there. Everyone had their own network protocol, and if you were on the west coast, it looked like XNS was the big winner. If you were on the east coast, it looked like TCP was important. If you were in Europe, people thought that ISO was actually going to get there before anybody else. If you worked for IBM, it was obvious that LU 6.1 and LU 6.2 were much more important than anything else, and none of these had any particular relevance because no one had successfully managed to get anyone's attention.

Pelkey: Right.

Saltzer: So, when was exactly when TCP/IP came out, I'm not sure. There were two events, one of which we were involved in, and one of which Berkeley was involved in, back in 1979 or 1980? First was Berkeley 4.1, which started making TCP/IP available as a much more widespread event, and the other is that, roughly about that time, 1981 I think, 3Com came up with an Ethernet card for the PC. Within a few months after that, we had TCP/IP running on the PC, and we made a public domain implementation available for that. And the thing that that did was kind of curious. PCs by themselves were not very important, but because Dave Clark was, at the time, on top of the TCP working group, therefore the PC version of that program was the current protocol. If you wanted to know whether or not you had a working TCP/IP on your old whatever machine, about the best way to deal with it is to connect it to an Ethernet and see if you could talk to the PC, and if you talked successfully to the PC, that meant you had an interworking implementation. The fact that you talked to Berkeley didn't tell you, because Berkeley actually had done a number of things that were not quite legitimate with respect to TCP, but it worked fine itself.

Pelkey: Right.

Saltzer: So, the TCP for the PC became a reference implementation, and what happened was that people suddenly discovered two things. If they got themselves a network, an Ethernet, they could run UNIX, they could run their VAXs and something from Berkeley. From MIT, they could get something to run on their PCs, and almost at the same time from Sun, they could get another version of the Berkeley UNIX, and the ability to do heterogeneous communication suddenly became a property of TCP/IP.

Pelkey: Right. So it was a culmination of those three events --

Saltzer: I think the key was heterogeneity, and the heterogeneity of the Sun workstation and the VAX and the PC is what caused people to begin to push people over the edge and say: "Gee, if I want a heterogeneous net, I better do this. If I have a homogeneous net, I can use IBM, or I can use Xerox, or I can use whatever, but heterogeneous, this is the only possibility," and there were so many people with heterogeneous nets, that they began, one by one, to move to TCP/IP, and every time one more moved onto it, every time one more vendor got his software working on it, it increased the momentum.

Pelkey: Right, and pretty soon it just became the standard. So you worked on inter-networking and that was one thing, and was the second one the TCP/IP on the PC? You mention that you were working on two problems.

Saltzer: That was one of them. I mentioned that TCP/IP was a major project that we were working on, both networking -- the building of gateways and the debugging of protocol strategies for use across gateways, and the development of TCP/IP and the specific development of the PC version of it, all were activities of my group. So those were the things we were doing in the time up through around 1983.

Pelkey: Which is when you -- then you went off and started work on - -

Saltzer: I then moved to Athena in 1984.

Pelkey: Ok. Going back. Ethernet being put out into the public, i.e. being forced into the standards committee was an unusual event. I mean, it was kind of forced, and then IBM came out and started to sponsor token ring as being a standard, so I understand. Do you have any reflections on that? Was that good? Was that helpful? Was it appropriate? It was done, if I understand correctly, only because Xerox and Intel and DEC couldn't get together to cooperate because of anti-trust reasons, unless they had in their own mindset that, by putting out to the public domain and making a standard of it, then they could get together and cooperate, and it wouldn't be perceived as anti-trust. And that's because IEEE, I gather up until that point in time, really hadn't dealt with standards and if there was anything, it might be one standard, as opposed to what became this kind of multiple standard approach to networking.

Saltzer: It might be useful to distinguish the formal setting of standards by standards bodies, in the sense of IEEE or ANSI or whatever, from the de facto standards that are created by computer companies building things, because the event that Xerox and Digital and Intel created by publishing the Ethernet standard, was of the variety of company setting a de facto standard.

Pelkey: Right.

Saltzer: It had nothing to do with any standards bodies anywhere. I think. I'm not sure if they were involved in their own internal arguments and discussions, or their legal arguments, or whatever as to how to go about doing it. The idea that it was made public domain in order to reduce the risk of an anti-trust concern probably sounds plausible, but I don't know. But it had absolutely nothing to do, as far as I can tell, with any standards setting, in the sense of ANSI or IEEE or whatever. Those people came along much, much later.

Pelkey: Ok.

Saltzer: I mean there have always been, from the moment that anyone comes up with an idea, there will be an ANSI committee usually formulated to go off and standardize it, but that doesn't mean it's of any significance.

Pelkey: But local area networks strike me as a little unique, because the goal was to have communications between different computer vendors, and that begged the issue of standards. When a computer vendor controlled their own environment, there was no need for standards. When you had a mixed vendor environment and you wanted to communicate with each other, standards were important. The tradition of the modems…

Saltzer: There are computer standards. There's standard FORTRAN. There's a standard APL. There's a standard PL1.

Pelkey: That's true. Maybe I better go back to the drawing board on that one. Because, I mean --

Saltzer: They have to do with communications.

Pelkey: Yeah.

Saltzer: FORTRAN is standard because you want to be able to move a program from one place to another.

Pelkey: And that is communication.

Saltzer: So it's a very similar problem. The point is I think there are several different things going on here. First of all, computer standards and communications standards have historically, up until 20 years ago, maybe 15 years ago, there was something very different driving the industry. If AT&T said this is the communication standard --

Pelkey: That was it.

Saltzer: ANSI would say: "Ok, I guess that's what it is." If Digital walks in an says: "This is the standard way of running VAXs," you know these standards can be (unintelligible), and ANSI would say: "Well that's a Digital proprietary thing. Why should we care?" So the fact that there was a single company setting the standard in the communications business, I think is probably fairly significant, and I don't know, probably you know more than I do about what happened when Vadic came in with a different modem standard. Ultimately that did become a standard of sorts, although I don't know when, relative to their arrival with it and the later standardization by the standards committee, I don't know how long (unintelligible) was there or how much disbelief there was of the concept or whatever (unintelligible)

Pelkey: Or Holsinger, who came out of MIT and went off and did a 9600-baud modem. David Forney took on the problem of performance and developed a better way of dealing with some of the issues that then got adopted in the standard.

Saltzer: So the process, on the whole, standards follow practice, I think it's normally the case that standards follow a long time after initial implementations, even in the case of de facto standards. After initial implementations, you finally have people honing in on something with de facto standards, and only much later do you get real standards. I think that's the normal process.

Pelkey: Although there's exceptions. I mean V.32 is a big exception. 9600-baud modem standard, I mean that has been in the standards committees for years and they tried to create a standard earlier, in fact, there wasn't any practical implementations of it. You're just starting to see V.32 modems now, long after the standard.

Saltzer: I believe that on the whole, that standards that people try to pull together that way are failures, usually.

Pelkey: I think so too.

Saltzer: And the reason is that the standards committees can get in to begin debate at about the point when people get to the point that they believe that the only things left to argue about are taste.

Pelkey: Well, not only that, but it goes back to your earlier comment, I think, which is, it becomes a political process, as opposed to a technical process.

Saltzer: Yeah. I don't believe that standards committees ought to be out there trying to solve hard technical problems, because those have to be resolved other ways, in the marketplace, and then the standards committees can come along and decide what to do. So, for the last 20 years, you've seen people attempting to set standards both ways, and on the whole the people who have tried to do it by standing up and saying: "We're going to invent a better local area network and make it the standard," they fail.

Pelkey: Yeah, right.

Saltzer: That was where the standards committee, actually standards committees started. They started off and said: "Oh, gee, there are all these people with local nets, but they're proprietary and this and that. We're going to do a better one, and it'll be the standard." And they failed utterly. Ultimately they ended up embracing Ethernet.

Pelkey: Right.

Saltzer: And the reason was that they couldn't make it, and we've seen this in a dozen different areas that I think --

Pelkey: That's an important issue.

Saltzer: Oh, a very important one. Have you talked to -- let me suggest another person to talk to, if I can think of his name? Carnegie- Mellon University -- Marvin Sirbu.

Pelkey: I tried to last Thursday night.

Saltzer: Ok fine. He has definitely explored this question of where do the standards come from and why do some succeed and why not.

Pelkey: One other question of different kind -- from my perspective, which is obviously from a great distance, in the 1960s MIT was a dominant source of research and ideas, I mean, just in general, in terms of being the dominant place of idea generation. I mean, timesharing --

Saltzer: You're talking about computers and communications.

Pelkey: Computers and communications. In the '60s MIT was dominant, and then in the '70s was Xerox PARC. Then in the late '70s, so I gather, there was some conscious effort for MIT to reestablish itself. That may be too strong a word. It makes it look political, but there is something to be gained for an institution of the quality of MIT to be perceived as being where exciting things are going on and therefore its a magnet to attract competent people and good students and so on. Why did Xerox PARC, from your perspective, become such a hot area, particularly when you had a corporation that didn’t know what it was doing?

Saltzer: Well, it's an interesting set of premises that you set out. Having been at MIT during this period of time, I'm not sure I would have described the scenario quite that way, because what happens, I think, is that, at any given time, we have a lot of different people at MIT. We've got a hundred faculty members working on electrical engineering and computer science, 35 in computer science. We're doing lots and lots of different things, and if you look at one narrow area, such as timesharing, we're talking about three of the 35 faculty members. You talk about local area networking we're talking about one or two of the 35 faculty members in computer science. Ok, you talk about knowledge based systems, we're talking about four of the faculty in computer science. If you talk about robotics, we're talking about three of the faculty in computer science.

Pelkey: Right.

Saltzer: Again, we have got 35 faculty members in computer science. Some years some of them appear to be big on the scene and other years they don't. Part of it has to do with a hype that happens to go with “is this a fashionable thing this year?” Part of it happens to do with whether or not they were lucky enough to have a breakthrough. The point is breakthroughs aren't something you schedule or you commission. They are something you set people out, against a hard problem, and some of them will make it. So a substantial amount of luck or risk or whatever is involved, and I think that's part of what's going on here.

Pelkey: Right.

Saltzer: Now, the interesting thing about all of this is that if you get very good people together with vision, this doesn't guarantee that you get a big winning maneuver out of the situation. It simply increases the probabilities. If you get a dozen such groups going around the country, one of tow of them may happen to have hit on the right idea, and some of the others probably won't, and as a result you will then discover that Xerox PARC seems to have a big win in this five year period. Someone else at another research laboratory doesn't seem to. It's not clear -- I mean, you can always go back and try to offer excuses having, or arguments having to do with: "Oh, there people weren't as smart or they somehow got off in the wrong direction," but I think that it's -- there's a lot of flakiness here. There's a fair amount of randomness and luck that's involved in that, you know, standing in the right place at the right time. "I'm going to try this idea. That guys going to try that idea. It turns out this idea is better," but you couldn't have told that in advance. Or another aspect of it is you try this idea and it turns out that, accidentally, a lot of other people got excited about it at the same time. Someone made a breakthrough, someone else that happens to mesh very well with your breakthrough, and the two together have got a huge amount of symbiotic effect, and you're able to take off. I think effects like that are probably what we're talking about here, more than anything else. The Xerox PARC group had everything going for it. They had a top notch leader and a bunch of very good people, and --

Pelkey: And a hard problem.

Saltzer: Well, I was going say that the thing that Xerox PARC really had going for them was they recognized early that the personal computer was going to be a major idea. They were one of the first labs to really pick up on that, but at the same time, we had a laboratory here, the AI lab was going the same thing. They were building personal computers too, only a slightly different design. Their goal was -- they were trying to solve the AI problem, so they were building LISP machines. The people out at Xerox PARC were trying to solve office technology problems so they built a slightly different kind of machine.

Pelkey: Thank you. You've been very kind with your time and very helpful. I greatly appreciate our conversation.

Interview Ends.