Interview of Kenneth T. “Ken” Pogran

Interviewed by:
James Pelkey

Recorded: October 27, 1988
Cambridge, MA

CHM Reference number: X5671.2010

James Pelkey: Thank you for your time. I look forward to this conversation. You were in the process of telling me where you were and how you came to this process, and since you know the period of time I'm interested in, why don't you just take off.

Ken Pogran: Let me begin by saying that I got involved in networking, really, starting in 1971 as a graduate student at MIT. It was then Project MAC, and became the Lab for Computer Science. I was involved with the Multics group at Project MAC. Of course, MIT was one the first sites involved in the Arpanet, so I became involved in implementing some of the early Arpanet host software for the Multics operating system, and in that context became part of what was then a relatively small group of people who were involved in the Arpanet and packet switching overall. Among other thing, I helped to implement -- or, not helped to implement but implemented the first file transfer protocol server on Multics. That was probably in late 1972 or early 1973 -- or, probably earlier than that. I can't quite remember what the dates were on it. No, that was probably right, late '72, early '73. I was around for what was really one of the first Arpanet demos, I guess, in I think it was July -- earlier than that -- July of '71 there was an early Arpanet demo, with a number of people who were participating in the Arpanet who were basically demonstrating, basically, that Telnet and NCP and all that stuff really did work.

Pelkey: The 'bake-off' was in the fall --

Pogran: The 'bake-off' -- that's right. Ok, that's what we're talking about.

Pelkey: October/November of '71.

Pogran: So it would have been earlier in that year that, as a graduate student, I started getting involved in the Arpanet stuff. I remained involved in the Arpanet aspect of Multics work at MIT through the early '70s. You probably already have, or I hope you have in the can, all the exciting stuff that happened at the first ICCC in '72 -

Pelkey: Right.

Pogran: - and the Arpanet special project demo.

Pelkey: Yes, the Arpanet public demonstration.

Pogran: Right, I was one of the people from MIT who were holding down the MIT -- was one of the MIT representatives to that, to that demo, to that project. So you probably have gotten some stories about how we walked into a completely empty small ballroom or large conference room at 8:00 one morning, and by 5:00 had a working Arpanet node with 35 terminals hooked up to a TIP, and everybody logging in and doing their thing at their home hosts.

Pelkey: That's right. I hear that -- people didn't want to leave at night, because anybody who knew anything about computer networking was in that room at that point.

Pogran: That's right, and that was very spectacular; just the notion of just walking into a completely empty room and setting up a functioning packet-switching network node and being able to log into host computers scattered around the country by 5:00 in the afternoon was pretty astonishing.

Pelkey: Yes.

Pogran: I have no idea how many union rules we broke, because we laid cable, built a false floor and all kinds of stuff -- you know, all these gung-ho young fellows from the universities that were involved in the Arpanet effort, so it was a lot of fun. To move a little bit more in the direction of local area networking and so forth, I took a look at some of the material that you already have, and I think that you quote late '73 as the time when the AI lab began thinking about --

Pelkey: '74.

Pogran: '74, ok -- began thinking about networking, and that really is about right. Certainly we had begun -- at MIT, we had begun to hear about Metcalfe's Ethernet stuff going on at Xerox, and it was probably in '75 --

Pelkey: Well, you would have seen his MAC report, his PhD dissertation, because it was published as MAC-1114 in '73.

Pogran: That's right, that's right, and then off he went to Xerox.

Pelkey: Actually, he was already at Xerox.

Pogran: Right, and we began to hear things about this network that he was implementing. Of course, we already knew about the Aloha Network, the radio experimental network in Hawaii, and we began to hear about the coaxial cable Aloha Network and thought: "Oh, isn't that a neat idea. We should learn some more about that." I think that the LCS interest in local area networking was really very, very comparable to the AI lab's interest, in the sense that we both -- both labs were beginning to have a large number of computers. We were all together beginning to use our corner of the Arpanet for a large volume of intra-MIT traffic, between the AI lab and LCS, and all of the computers that -- the major PDP-10 hosts and Multics hosts that LCS had and so forth. They were all on the Arpanet and we were using the Arpanet for our corner of local connectivity; even beginning to use TIP terminal ports as a means of terminal access -- you know, a terminal server, to get onto the variety of computers that we had. So there was a need for better connectivity, better than just using the Arpanet, for communication among the host computers that we had at the time; basically preceding, by at least five years, what's now typical, or became typical, in major universities and other organizations, as there began to be lots of good sized computers out there.

Pelkey: Ok. So LCS got interested in '74 as well.

Pogran: Either in late '74 or early '75. I might be able to go back and look at some notes, but I don't know if I still have any of that stuff. So the motivation was pretty much the same. It wasn't just research motivation, it was 'hey, we need this for connectivity,' and since LCS didn't have quite the same hardware expertise that the AI lab did, we weren't about to go out and design something the way the AI lab folks did. I think you've correctly identified the competition aspects -- the aspects of the competitiveness between the LCS and the AI lab, and it wasn't clear that we just wanted to adopt the Chaosnet approach either.

Pelkey: It must have taken the AI lab some time to have created Chaosnet.

Pogran: Yes, and I can't quite remember what that time scale is. Certainly by 1975 I was asked to investigate what the Lab for Computer Science should do to build a local area network for LCS.

Pelkey: Do you recall whether it was early in the year or later in the year?

Pogran: I can't recall, off hand. I might have some stuff that I could look at, but again, I'm not sure. That was basically with Mike Dertouzos and Jerry Saltzer and Dave Clark that we decided to look into that.

Pelkey: And Michael was head of the lab at that point?

Pogran: Yes, Mike was the head of the lab. So I embarked on an effort to learn as much as I could about the Chaosnet, to learn as much as I could about the Ethernet, and also to learn about Dave Farber's experimental ring network technology. He already had a version of a ring interface running in the Distributed Computing System at UC Irvine at that time. It was really -- for Farber and Farber's group,

that ring was the means to an end, namely to explore DCS, rather than actually to produce a local network technology.

Pelkey: Right. Did you go to Xerox PARC?

Pogran: Yeah, I went out sometime during '75, and I can't remember when, I went out to Xerox PARC, and didn't recognize Bob Metcalfe when we came walking down the hall to meet me, because when he left MIT he was bearded and a tad on the plump side, and as I was standing there in the lobby -- this very impressive lobby of the Xerox Palo Alto Research Center -- in walks this slim, clean shaven, tan, very relaxed fellow, and that was Bob Metcalfe. So I learned all about the Ethernet, and some of the advantages and the disadvantages, the pitfalls, the motivation -- the very practical motivation that Xerox PARC had to develop the experimental Ethernet, which was based on the need to provide connectivity among the Alto workstations that they had there. Carrying cartridge disks wasn't a terribly effective means of sharing information among them.

Pelkey: In '75, sometime, they had like 100 nodes up and running.

Pogran: That's right. So, in any case, I spend the better part of a day with Bob and talked with a number of the guys there about the hardware, the software aspects of it, what they were doing -- a little bit of the electrical engineering behind it all. Then, I went on down to Irvine and met with Farber and some of his people and saw the original DCS ring, learned about their plans to, under DARPA sponsorship, design a more general purpose one- megabit ring network --

Pelkey: Ah, because the NSF had funded his original work.

Pogran: That's right, and so this was now a DARPA project, it was going to be done by graduate students, which should have told me something, because one of the rules of thumb at MIT LCS was 'If you have a project that you're trying to complete as, in effect, a deliverable under your DARPA contract, don't rely on graduate students.' That's why we had full-time research staff, to make sure these things actually got implemented, because the graduate student wants to write the thesis, and then do something else.

Pelkey: Yes, classes, etc.

Pogran: Right, so we learned about the ring network, and I came back to MIT, and we began to think about what we wanted to go ahead and do. I don't remember what the reasons were, precisely why we decided not to go along with the Chaosnet, and as I say, I really can't remember what many of the reasons were. I think we were probably worried about expandability; we were probably worried about the level of support. We really wanted something that was not going to be a toy, but that was going to be a real utility for us.

Pelkey: Someone, I think it was Dave that mentioned, that they didn't have a Unibus interface, which was critical to you.

Pogran: That's probably right.

Pelkey: In getting the Unibus, you would have to have depended upon them to do the Unibus interface, and that was a concern.

Pogran: That's right. We were starting to get involved with a lot of 11's and so forth at that point. Of course, the Chaosnet did have the PDP-10 interfaces, which were pretty important at that point too. So, we went through a little bit of analysis of what we wanted to do. We realized that, as you have in your manuscript already, Xerox wasn't sharing the Ethernet at that point, so it appeared that, if we weren't going to go Chaosnet, we couldn't go Ethernet, that the best thing to do would be to get into an effort with Farber at UC Irvine.

Pelkey: When do you think that happened?

Pogran: It's got to be in '75 -- '75 or '76, because I certainly remember that, as we began, at LCS, began planning to have a network, Mike Dertouzos began to talk about it as 'the 76-Net.' We were going to have a local net up in '76. So it would have had to have been '75 when we began to embark on this. I've probably got documents that I could look at that can sharpen that --

Pelkey: If you get a chance, it would be helpful. Do you recall there ever being an experimental Ethernet?

Pogran: At MIT? Yes, but later. I'll get into that as we go through the chronology. So we began to work with the Irvine folks, and they were developing their ring network interface as a graduate student project, essentially, and I can probably pin down these dates after I look at some stuff, but I can't now. I made several trips out to Irvine with the objective of helping them begin to test and debug the ring interface. We began by trying to get the Unibus interface working, so that we could run the damn thing and bring it through more debugging. There were a lot of problems and issues that came up. One was that, although Farber and his students had committed to doing this for PDP- 11's, Farber's group itself did not have much in the way of PDP- 11 resources, so the very first debugging that we did -- we went up to Jerry Popek 's shop at UCLA and used one of his 11's as the very first test site. It was clear that that wasn't going to be very good either, so we decided to move the debugging venue to MIT. So we rather -- I, rather naively, said: "Sure. We'll get together -- you guys from UC Irvine can come out and we'll plug this into our 11 and we'll get it going lickety split. Everything will be just fine." My background -- I should digress a little bit and point out that my background was in both electrical engineering and computer science, and probably as late as the end of my junior year at MIT, I wasn't sure whether I really wanted to go into software or hardware stuff -- digital logic or programming, but I did spend the first few years of my professional career writing network software. In early -- it must have been early 1976, I suddenly discovered that I was into debugging hardware in a big way, because really, aside from frequent trips out by the primary UCI grad student designer, I was the guy who was trying to debug the local network interface, the LNI, pretty much by myself. A lot of telephone consulting from Irvine and all that, and I was learning all about families of digital logic that hadn't quite existed when I did the undergraduate work, and so it really was -- and I completely underestimated the magnitude of the tasks, to debug a device which basically consisted of about 300 LSI and MSI TTL IC's, and a mass of wires that I hadn't designed, and to discover, sort of a step at time, what' s going on here, was an incredible challenge, and also bad for my reputation at LCS, because it was taking forever, and I always thought it would be done soon. So the bottom line is that we really didn't have a '76-Net up at the lab for computer science. Again, I can probably figure out when we did start to have things, but it was a little bit later on.

Pelkey: So it was in the first half of '77?

Pogran: Probably would be. I can check that. I'm sure that I've still got some stuff on it. So we basically got over the hump. We began to get a few more copies of the local network interface built and in, continually making some improvements and so forth.

Pelkey: Now, the ones that you received initially, are these the famous ten prototypes that you received?

Pogran: I don't think that we ever received ten prototypes. I saw that in your manuscript, and I don't recall that being the case. I believe that we got maybe the first one, or the first two -- let me see. The very first prototype ring interface certainly was given to us by UC Irvine; stuff that they bought under their contract. Mike Wile, who is the lead guy from UC Irvine, and I kind of very quickly sat down and went through some catalogs and did some repackaging of it, and I think they switched vendors as to who was doing the wire wrapping and so forth, and basically went with a fabrication house out in southern California that could build something -- easily build something of the scale that we needed with the technology that we wanted. I recall that all the subsequent prototypes, all the subsequent units, were things that we at MIT actually bought off our DARPA contract. Sometimes it was their DARPA contract or out DARPA contract; it was the government's money anyway, and we, at that point, really began to control what was happening. We probably had seven or eight of nine or ten -- I don't quite remember how

many, the highest number that I remember was seven -- local network interfaces, V-1 LNI's, of which as small number went out to Jerry Popek at UCLA, probably two or three or four units were out at UCLA. So we got to the point where they were fairly reliable and actually useful for communications among the PDP-11's that we had at the lab. It was an interesting time, because we were trying to make this new technology happen, and really, although we didn't think terribly hard about it -- inventing some of the principles of local area networking in the process.

Pelkey: What kind of software did you run at that point?

Pogran: Real simple stuff, at first. Just, kind of like repeating the early days of the Arpanet: device drivers, just moving packets from here to there and making use of them. Noel Chiappa was, I think, around at that point and writing some of the early software to run with these things. There were early versions of TCP that were coming out of the protocol research. I don't think it's fair to say that DARPA gave MIT a copy of TCP, but I think rather, that people at MIT were part of the working groups that were developing TCP protocol specs, and were among some of the early implementers of some of the early versions of TCP. So some of them were built to run on the 11's with the ring interface. At that point, I was enough -- I had sort of done a 180 and I was deeply into the hardware of it and not really paying close attention to what was going on with software.

Pelkey: Ok. So now we're into sometime in '77; the network is up and running, and it must be -- you continued to expand it, and you must be into '78 some time.

Pogran: Probably right. In fact, we'd have to look at some dates on some papers now, because Dave Clark and Dave Reed and I wrote the Introduction to Local Area Networks paper which appeared in Bob Kahn's November, '78 issue of the IEEE Transactions on Communications -- no, Proceedings of the IEEE, November '78. The Introduction to Local Area Networks paper, which I think still gets -- I cringe every time, but it still gets reprinted today. So, we did that. Not too long after that, of course, Ethernet was receiving a lot of publicity -- the three megabit Ethernet was receiving a lot of publicity and Saltzer and Clark and I felt that we need to equally make the case for ring networks, so we put together a paper that was entitled Why a Ring?, which had also been reprinted a fair amount, and as purely a pragmatic matter, from the point of view of trying to cope with some of the reliability issues of ring networks, we developed the concept of the star-shaped ring and the wire center for a ring network, and Saltzer and I did a paper on that.

Pelkey: Now, did those events happen in the '78 - '79 time frame, or did they happen real early on, in terms of the ring. Did you basically import Farber's stuff and get it working and then say: "Wait a minute, there's a better way to do this"?

Pogran: Yeah, well, we imported Farber's stuff and got it working, and we really -- a lot of the final stuff in that ring interface was ours -- you know, mine, rather than Farber's, but the wire center idea came to us as sort of a pragmatic 'how do we cope with this,' rather than a deliberate -- any deliberate design. You know, we said: "How do we cope with the ring? How do you expand it? How do you avoid pulling wires in and out all the time? ' and so forth, and we came up with the wire center idea. It was actually quite a while before we implemented it. We basically built the wire center. We built a star shape ring and cabled it that way with manual connectors and so forth, kind of to patch the ring together. Let's see -- some time in, probably '78, and again, I probably have some documents, Xerox, of course, discovered that a lot of colleges were clamoring for Ethernets and Altos and all these neat things, and they established a University Gift Program, or I forget exactly what it was called, and at this point I was managing a small group at the lab. I basically had a technician and a support type guy working for me, and MIT, Carnegie-Mellon and, I think, Stanford, although I'm not sure, were named as recipients of stuff from Xerox under the University Gift Program, so my two guys went out to Xerox to get some training on the care and feeding of these beasties, and sometime, probably in '78 - '79, although I can't remember when but we can figure it out, we got some Alto workstations from Xerox, we got a Xerox file server, we got a Dover printer from Xerox, and at that point, we had experimental three megabit Ethernet technology at the MIT Lab for Computer Science, and that was the point at which we had that.

Pelkey: Ah, good, good. That's helpful. Now, somewhere along the line, you also decided to do a second version of the ring.

Pogran: Right, again, the chronology doesn't come to mind exactly, but we can figure it out. Mike Dertouzos was worried about the process that we went through to get the Version I Ring working. It wasn't what he had in mind, and it wasn't what I had had in mind either, but that's the way it turned out. So he was concerned about getting a new version of the local network that would be easy to obtain, and from his perspective, that meant working with a company that could build the stuff for us and hopeful, if it was successful, actually have it be a product that we could buy, but certainly the notion was: "Let's get a company involved who can do this hard stuff for us, because we don't seem to be terribly good at it," and we already had a lot of ideas about how to improve on the ring network, and so forth. So he knew some fellows at a little company called Proteon by the name of Howard Salwen, who had been a friend of his from --

Pelkey: School days --

Pogran: -- from before, right. He brought the Proteon guys in and said: "Let's see if we can work together to develop some sort of local network that we could build." Well, the Proteon guys were more comfortable working in RF types of stuff, and they were more interested in possibly something like a MITRENET type of thing -- a broadband contention network on CATV. So Dave Clark and I and Jerry Saltzer worked real hard -- basically did a selling job on the Proteon guys and said: "No, no, no, ring networks are really better. Here are the reasons why. Here's what local networking is all about," and quite literally, I think we taught -- the three of us -- talked Proteon everything that they knew about local area networks, and we eventually convinced them that what we really needed to do was to do a much faster, much more reliable version of the ring network. Now, we already had a lot of ideas sketched out on paper for how to do that, how to simplify it, how to make it better, and that became what we call the V-2 LNI, the Version II Local Network Interface. The objective was a 10 Megabit/sec. ring. We used Proteon's expertise to handle things like clock recovery and phase-locked loops, clock stability and all those issues. We put in our ideas about how the controller should work; the notion that there should be a network interface board and a host-specific board, a controller and a host-specific board, modular so that you could have different kinds of host interfaces. I had some ideas about how to do a Unibus DMA controller that would be more reliable than what we had on the Version I LNI, and really, I think that a lot of the MIT ideas got directly implemented in the very first versions of the Proteon ProNet ring interface.

Pelkey: Now Proteon got involved in 1979 sometime --

Pogran: Ok, that's probably right.

Pelkey: -- and Metcalfe was back, now, acting as a consultant to Dertouzos in what was happening.

Pogran: That's right.

Pelkey: Do you recall any discussions relative to -- I presume that Metcalfe was really arguing that you should use Ethernet versus persisting with your ring.

Pogran: Well --

Pelkey: Or wasn't it that way.

Pogran: I can't recall him arguing very hard. I remember, certainly, Metcalfe came back at one point and began extolling the virtues of fiber optics for local networks, and I remember one discussion with him in which he basically was saying that: "Hey, fiber is mostly a point-to-point medium. It's going to be a long time before we get passive taps for fiber and things like that," and fiber, once again, motivated a ring network more than a bus.

Pelkey: Right, absolutely.

Pogran: So I think Metcalfe, in his role of consulting to MIT at that point, was probably a lot more objective than he would have been publicly willing to admit at that time.

Pelkey: Ok. Do you remember the Copley Plaza meeting?

Pogran: Yeah.

Pelkey: Was that an important meeting?

Pogran: I think it was an important meeting because it brought a lot of people working in the local network field together, and there was a lot of idea sharing going on, and notions of what directions things would take. It was probably a lot more important to people who weren't as up to speed as we were.

Pelkey: Right.

Pogran: I thought it was a good time.

Pelkey: I would think, yeah. It was very different for you than it was for others.

Pogran: I think that's right. I think more important to me was the seminar -- was a workshop that NBS held a year or two before, which a lot of people who were working in -- see, at MIT we were using the phrase 'local networks,' and that NBS workshop was where I think 'local area network,' that phrase, came to be used. I think that's a conference at which I developed a lot of appreciation of what other people might want to do with local area networks and what the technology might be --

Pelkey: You were deeply embedded in getting something up and running --

Pogran: That's right, and most of the people who were at that meeting in '77 were doing something in the field; they either wanted to learn about it or they had something to bring to the table.

Pelkey: Do you recall whether or not your network was completed when you went to that conference in '77?

Pogran: I don't recall. It's hard to say.

Pelkey: We've run out of time, but I will, if you don't mind, give you a call and update --

Pogran: Well, I think we can continue until the knock at the door when the fellow who was supposed to see me at 1:30 is here.

Pelkey: Ok.

Pogran: So that's --

Pelkey: So Proteon got involved sometime in '79, and so Version II started to happen.

Pogran: That's right. I left MIT in -- at the end of April of 1980, with the Proteon implementation effort well under way. I think we had begun to see bits and pieces of prototype hardware and Proteon guys were at MIT trying to debug their Unibus interface, and that sort of stuff.

Pelkey: But their process took a lot longer than anybody expected as well.

Pogran: Yes.

Pelkey: Small consolation.

Pogran: That's exactly right. I'm not sure why it was Mike Dertouzos thought that having somebody professional do it would be magic, but --

Pelkey: Where did you go when you left MIT?

Pogran: I came here [BBN].

Pelkey: Do you remember what was happening in the TCP/IP front with Dave Reed and Dave Clark during this period of time?

Pogran: Not very clearly, because I was really paralyzed around the (unintelligible) net hardware at that point. In fact --

Pelkey: I can imagine you poking your head up and someone knocking it back down.

Pogran: I know in your manuscript you mention professor Reed, and I don't quite remember exactly when he went from graduate student to professor, so you might want to check on that.

Pelkey: Do you know where he is now?

Pogran: Ooh, no. The last that I knew of Dave, he had gone . . .

Tape Side Ends

Pogran: . . . '71, '72 at MIT, so Dave eventually went off to Software Arts, as I said. Where he is now, I don't know.

Pelkey: During this period of time, there was a lot of other activity going on in networking. There was, as you saw in my manuscript, there was activity going on at Zilog, at Prime and at DEC -- less so DEC. DEC was more software. There were a number of other hardware sorts of things being attempted as well. Were you, since you were deeply embedded, were you aware of these other activities?

Pogran: Yeah, I think we tried to stay on top of it. I remember one meeting that we had, Saltzer, Clark and I, or maybe just Clark and I, I'm not sure, with the guys at Prime. So we did talk with Prime, and we learned about some of the things they were doing, some of the advantages and disadvantages of their ring net approach.

Pelkey: They were very much influenced by the Multics experience. [William] Poduska had come out of a Multics environment as well.

Pogran: That's right. So we did spend a lot of time with Prime. I had been out talking with MITRE about their broadband cable stuff; got to know Greg Hopkins pretty well at that point. So yeah, we tried to stay pretty much on top of it and involved with it.

Pelkey: You were making the comment -- I presume the protocol is they'll knock.

Pogran: Yes.

Pelkey: I presume -- you were making the comment about internetworking and the importance of internetworking. If you don't mind, would you take a few moments and speak to that issue.

Pogran: This gets a little bit more into the current decade, rather than the previous decade, but I guess, from a technology perspective, from the perspective of the evolution of networking technology, in some

sense you can say that the research that was done in the interconnection of networks, or internetworking, and the development of the internet protocol and TCP is what we now call transport layer protocol, made it possible for some of the early local networking to get to be publicized. Local networks started to come into existence and they started to be gatewayed to the Arpanet and to other wide area networks, and I think made it easier for local networks to be publicized. By the same token, I think that the proliferation of local area network technology really created an even greater demand for work to be done in the area of interconnection of networks, in internetworking, because as local networks got out there more and more, the need for them to communicate with one another became greater and greater. So there's really a lot of synergy going on between the continued availability of wide are network technology, DARPA sponsorship of internetworking research, the burgeoning local area network technology, and eventually, the advent of things like diskless workstations. In the early '80s, those of us who were in the local network field talked about diskless workstations as a possibility, and I spent -- in fact, I spent a few years as an ACM lecturer, giving talks about the challenge of local area network technology, and talking about applications as -- one of the applications for it being diskless workstations, doing your memory swapping and so forth over a network, rather than off disk, and the audience in that area was generally incredulous. Of course, we have it today, and what's even more important, we have it in the context of internetworking and the protocols available for internetworking, and so it all really plays together very well. So there's been an awful lot of synergy between wide area network technology, internet research and development, local area network technology and all of that sort of thing, plus I think that the advent of the Defense Data Network as a very large network and a big customer for host software and hardware has also served as a demand pull agent in internetworking. So I think that internetworking is being sort of the way of bringing order out of potential chaos, of all of those networks out there, is really key, is really very crucial to the continued evolution of data communication. In many -- I think it's well recognized now that an internet, in some sense, is the next higher order, or next higher level, of packet switching network. In a wide area packet switching network, you have switches, you have your PSNs, or your IMPs. They use leased lines as their transmission medium. In the internet, you have gateways, or routers, and they use networks as their transmission medium, but it's really -- there's really a mapping between two. The internet is the next higher order of packet switching network, and in fact we're getting to the point where in some implementation, there are beginning to be blurs between packet switching and gateway types of functions. DARPA has plans for the Defense Research Internet, the SDI community has plans for the SDI Network Interfaced Processing Element, SNIPE, and both of those represent, really, a blurring of the distinction between high bandwidth advanced packet switching, if you will, and gateway and internet technology. Basically, we're moving into an era in which the way of connecting to a wide area network is not through an X.25 interface from a host, but rather through a port to somebody else's network, to a local area network, for example.

Pelkey: Yes.

Pogran: So you've got combined gateway and switch functionality in that. So I think it's really a key element of where the technology is going, where the community is going. It's the thing that --

Pelkey: As well as packetized voice, that will start to come into this as you get the bandwidth that will be able to handle it. It will be truly digital, in the sense that a packet will be a packet, and what's in it is 'who knows.' You packetize voice now for compression purposes, but you don't packetize it in terms of end-to-end movement, or link movement.

Pogran: Well, not necessarily true. I don't know if you're familiar with the broadband network project that's gone on over in BBN Labs, but they've been dealing with voice and video for quite a while there.

END OF INTERVIEW

