



Interview of Yogen Dalal

Interviewed by:
James L. Pelkey

Recorded August 2, 1988
Santa Clara, CA

CHM Reference number: X5671.2010

© 2010 James L. Pelkey/Computer History Museum

James Pelkey: You mentioned that you were with Vint Cerf in 1972, '73 when he came to Stanford. Were you doing graduate work at Stanford?

Yogen Dalal: I joined Stanford in the fall of '72, getting ready to go into the masters and PhD program, and Vint had joined Stanford in the fall of '72 as an assistant professor, and the area of research that I wanted to do was data communications, so I was pestering Vint to make me one of his research assistants.

Pelkey: How did you know of him to pester him?

Dalal: I was taking courses that Vint was teaching -- systems programming courses -- and Vint was trying to get an ARPA contract from Bob Kahn in those days, so he asked me to be a teaching assistant for a couple of his courses to sort of check me out and vice versa, and in the summer of '73, I formally entered the PhD program and became one of Vint's research assistants. It was sort of an interesting summer, because at that time, Vint and Bob Kahn were trying to formulate an internetworking protocol, which eventually became TCP, and was based on some of the work that Louis Pouzin was doing on Cyclades, so for the year prior to that, Bob Kahn and Vint and Louis were all chatting about how you connect networks together, and that summer we had a number of brainstorming sessions to figure out what to do. That's when I first met Bob Metcalfe. He had started at PARC and was planning on -- he wouldn't tell us what he was doing with Ethernet in those days, but he was participating in the TCP design, and in parallel, had taken some of those ideas and incorporated them at Xerox PARC in the form of the PUP protocols. It was at that point that Metcalfe provided a lot of input that didn't necessarily implement PUP the same way that TCP was being implemented. There were valid reason for that; the main one was that Vint and his group were concentrating on how you interconnect large networks together, Arpanet-like networks together, so they were on a particular design point, while Bob was busy building Ethernet there, and of course the computers. He was operating on a different spectrum. He was thinking of small machines that communicate very fast, so some of the design parameters were different, so he, David Boggs and Ed Taft developed PUP, while Vint, Carl Sunshine, myself, Dick Carr, some of the people at BBN -- Ray Tomlinson in particular -- were trying to refine and create life to TCP.

Pelkey: Was it your recollection, relative to Ray Tomlinson back then, was that he created the first network version of E-mail when they distributed TENEX?

Dalal: He was one of the main authors of TENEX, so I'm not sure whether he implemented any aspect of MSG, which is what the product was on E-mail. I know that he was responsible for a large chunk of TENEX, and was sort of considered a guru of systems design there. Another person, Bill Plummer, was also involved with Ray Tomlinson. So the group of us were busy trying to create a specification for --

Pelkey: I was led to believe that there were meetings in some small office somewhere in some building in Stanford?

Dalal: Well, at Stanford, we, the electrical engineering department, had a number of small buildings and we were in one of the smaller ones, which was sort of a post-war building; a very simple looking, unpretentious building, which is where graduate students and their research advisors hung out. So it was very unpretentious and we would just hang out there and try to figure out what TCP was about.

Pelkey: When you say 'hang out,' were there formal meetings or did you just have offices there together?

Dalal: Well, each one of us in Vint's group was a graduate student, except for Judy Estrin and Jim Mathis, who were -- Jim Mathis was an undergraduate and Judy was there for her masters.

Pelkey: But Judy was there in, what, '76?

Dalal: Which was later, she came in, but each one of us was taking courses. We were trying to help Vint on the research projects that were being funded by ARPA, so the meetings would be in offices, or in sort of -- in fact, we never had our own offices much. Many of us would just share a large office with two or

three desks in that, so we'd get through our coursework and take our finals and midterms, but also really work on this. The eventual goal was that we would stumble upon a problem that could be big enough to wrap a PhD thesis around it. The interesting thing is that only one PhD thesis came out that was related to TCP, and that was Carl Sunshine's thesis. Each one of us sort of tripped upon another problem that took us in a completely different world that had nothing to do with TCP, but that's the way life was. So there was a lot of very energetic -

Pelkey: So the summer of '73 was a very active period of time for the group?

Dalal: It was a very active period of time for that group. It was over the next 12 months or so that the definition of TCP formed, with participation primarily from Bob Kahn and the people at BBN. I recall that it was sometime towards maybe late -- well, it was sort of the end of '74 when we built the first prototype. We built one at Stanford and Bill Plummer and Ray Tomlinson built one at BBN, so this was a very exciting day when we sent the first TCP packets back and forth between the two machines. It was a very exciting period of time there -- a lot of high energy. Jon Postel got sucked into it. He was coming out of the work that he was doing at SRI on the older ARPA protocols -- NCP and things like that -- and he decided to go down to southern California and work with Steve Crocker, so Jon Postel started coming into the TCP/IP discussions.

Pelkey: This was toward the end of '74?

Dalal: I would say probably a little later, even -- maybe in the '75-ish period.

Pelkey: Had he gone back to southern California?

Dalal: By this time he had gone back to southern California. The whole notion of TCP became sort of a reality. We had a couple more implementations working. John Shoch, who is now at Asset Management, who was at Xerox at that time, sort of participated toward late '76, I think, in TCP discussions. Just sort of tangentially, roughly in the '74 time frame, TCP and Bob Metcalfe started diverging, in that Bob spent less and less time on the actual evolution of TCP and spent more and more time in making his prototypes at PARC work, so it was roughly in the '75, '76 time frame that PUP started evolving, real implementations, real Ethernets.

Pelkey: Was Bob pretty active during the summer of '73?

Dalal: In the summer of '73, Bob was extremely active in a lot of the discussions that we had. Of course we didn't know what was going on at Xerox, so it was hard for us, sometimes, now in retrospect, to understand why Bob would argue for a different perspective. Xerox was extremely hush-hush about what they were doing with the Alto and the Ethernet. I went to Xerox in early '77, and by late '76 things were beginning to be a little more open, especially if you were in the inside crowd.

Pelkey: So Bob, in '74, started to reduce his participation, and then in '75 and '76 PUP started to be developed on its own as a protocol for connecting Alto machines?

Dalal: Connecting Alto machines and a few of the early laser printers at PARC.

Pelkey: you were at Stanford throughout this whole period?

Dalal: Through the whole period. Through early '77 I was at Stanford.

Pelkey: And Judy Estrin joined the group in '76 as I recall?

Dalal: I would probably say '76 -- no, '75. She was there from the fall of '75 through the fall of '76, for about 12 months getting her masters, and so sort of came in late into the TCP/IP development. Because the world is full of connections; Jerry Estrin, Judy's father, was Vint's thesis advisor at UCLA, so she was part of the research project, and she was participating in some of the prototyping, and that was very hard

for somebody who was only there for a year or 15 months to have a big impact, but it certainly had an important enough impact on her, because after graduating from Stanford she went with Zilog where she tried to create Ariel, which was Zilog's local area network. The other person that was involved was Jim Mathis, as an undergraduate, a super guy who until recently was at SRI. He was at SRI for roughly ten years, and now he's running a major development effort at Apple, getting Apple into communications standards.

Pelkey: John Shoch was at Xerox when he was interacting with the group of you, or was he at Stanford?

Dalal: Well, he was at Xerox. He started interacting with us really later in the '76 time frame. That's when I first met John. He was at Stanford getting ready -- he was at Xerox, but getting ready to reenter the PhD program, just as some of us were actually getting ready to leave. That was because even though we were contemporaries, he took a couple of years off doing research at Xerox while some of us were going straight through school. So he started interacting with Vint on networking. Things were a little more open at Xerox, though still quite hush-hush. The interesting thing is that TCP and PUP sort of developed somewhat in parallel, and towards the end of '76, TCP gained a fairly wide momentum behind it. People at UCLA, ISI, Stanford, MIT started getting involved with Dave Clark and Dave Reed, who is now at Lotus, and the whole notion of --

Pelkey: But they were still running NCP on the Arpanet?

Dalal: They were still running NCP on the net --

Pelkey: But the sites that had brought up TCP could talk to each other using TCP, and the group of you were talking to each other trying to build this more robust protocol.

Dalal: This more robust protocol, and the goal was, eventually, to get NCP out of the way and TCP sort of well entrenched. A couple of things that happened along the way that caused the boat to get rocked; one was my going to Xerox. That itself didn't rock the boat. I went off and started getting involved in the next generation of protocols, which eventually became XNS, but in the development of XNS I studied PUP, and it was clear to me, after I had left Stanford to go to Xerox, the impact that local area networks would have on internetworking, and that while the theoretical problems associated with internetworking had been solved in the ARPA context, new light was being shed on what personal computers might want of an internetworking protocol. So we carefully tried to influence Vint and his people that TCP should change somewhat. In fact, TCP changed to become TCP/IP. The concept of the datagram evolved out of some gentle hints that were coming out of John Shoch, myself, David Boggs, who were working now on PUP and XNS. PUP was then beginning to be disclosed, but even before PUP had been disclosed we tried to convince Vint and Jon Postel and others, that breaking things into a datagram and a session protocol was very, very important, primarily because datagrams are useful in local area networking contexts. It was Vint and Dave Clark and Jon Postel who saw that immediately, and sort of slowly started modifying TCP to be a TCP and an IP.

Pelkey: I want to discuss that some more, but Dave Clark's recollection was that there is a gentleman whose name I forget that was at MIT --

Dalal: Jerry Saltzer?

Pelkey: No, another gentleman who wrote a paper about this issue that he felt was critical -- I apologize for not having his name. I'll look through my notes while we talk. Your recollection is that the experience came out of the group of you being exposed to what was happening in local area networking and PUP and the need for datagram type of services because of the connectionless orientation of Ethernet, in contrast to this kind of virtual circuit view of the world that existed in TCP.

Dalal: Even though the Arpanet itself was a datagram protocol, people hadn't seriously considered having computers use datagram protocols themselves, because reliability was a big issue in the Arpanet context, while reliability is less of an issue in local area network contexts because speed was so high and

you could achieve the reliability because local area networks inherently are more reliable. In local area network contexts, you needed protocols that didn't need connections. For example, if you ask for the time and you don't get a response, you can ask for it again. You don't have to open a connection to the time server to get the time. So we started discovering that there were lots of applications that didn't require connections in a local area context, and therefore, if there was going to be an internetworking protocol, it should allow the user of that protocol to either interface at the datagram level or to interface at the session level of TCP.

Pelkey: It strikes me that, in the history of how this technology evolved, that became a very critical point that is still unresolved.

Dalal: It's still unresolved.

Pelkey: And as we look forward to fast-packet switching, we're going to revisit all these issues again about connection/connectionless. At that juncture of what was happening with Ethernet, in the introduction of the mainstream world to IP, really was something that was very important.

Dalal: Very important, and I think that that's what caused TCP to be completely revised. In fact, it's lucky that it happened when it did, because it got into the mainstream of TCP/IP development - - in fact, at that time, Jon Postel took over the management of the entire protocol handbook -- and it happened in advance, so that when the entire Arpanet sites converted to TCP, they converted to the new TCP/IP, otherwise they would have gone through two evolutions. The one thing that didn't get into TCP, which is still a sore point, I think, was the notion that there would be thousands and thousands of personal computers, and therefore the address spaces aren't big enough in TCP, and at Xerox we were building XNS assuming that there was going to be a high-speed new Ethernet, the one that we call the Blue Book, and so we chose to use large address spaces that were 48 bits wide, so that we anticipated the need for thousands and millions of computers connected, each with IDs that would get stamped into them when they were created. So the concept of a universal ID per machine was another very important concept that came out of Xerox, and now you have to do some standing on your head to go between the short addresses and the long addresses, that sort of thing. An engineer's thrill, and it didn't really matter, but --

Pelkey: This period of time when IP started to get introduced was in early --

Dalal: Second half of '77, because it was certainly after I had joined Xerox.

Pelkey: Do you recall when IP came to be formally viewed as a separate layer?

Dalal: Probably around '78. I would say roughly in '78, roughly a year later.

Pelkey: And were you -- you weren't involved in any of this activity other than for having influence because you knew the people, and at this point in time, Vint's back at ARPA?

Dalal: Vint's back at ARPA.

Pelkey: With the responsibility to get this TCP/IP completed and out and used, and getting off NCP.

Dalal: Right. In fact, for my first six months at Xerox, I would attend the TCP meetings, but that was sort of the pull from the past, and the pull of the future soon made that impossible. There just weren't enough hours in the day to go to TCP meetings as well as to do development work at Xerox. Minor points of TCP have changed over the last four or five years, but roughly TCP/IP, as defined in late '78, early '79 --

Pelkey: This being version II -- there's four. Version four, as I recall, is the one that got implemented.

Dalal: But the changes from '78 to version three are all microscopic changes. The big building blocks and the major architectural decisions had been set in concrete by then.

Pelkey: Did you start working on XNS?

Dalal: Yes. As a matter of fact, that was the first thing that Bob Metcalfe asked me to do. I went to work with him there.

Pelkey: Were you formally working for Bob at this point?

Dalal: Formally working for Bob, so when I graduated in early '77 and I was trying to figure out whether to join Steve Crocker and Jon Postel -- well, Jon was still at SRI in '77 -- but whether I should go work with Steve Crocker down at ISI or work with Bob Metcalfe at Xerox or go all the way back to the east coast and work with Bob Thomas and his team at BBN. The pull to PARC was extremely strong. Once I had seen the Alto, there is very little that could convince you to go anywhere else. You had seen the future and you knew that this was it, and I considered going and working with Bob Taylor's group in the research labs, and chose, instead, to go work with Dave Liddle and Bob Metcalfe in the product world, because I saw a lot of things that I had done impacting the future of Xerox. So, at that time, when we started in early '77, the goal was to build a 20 megabit per second Ethernet, which did come to fruition, and to build a brand new set of protocols based on the experiences of PUP.

Pelkey: PUP was out in the marketplace at this point in time in one megabit or three megabit --

Dalal: On the three megabit Ethernet in about '78, '79, Jerry Elkind created the Advanced Systems Group and they tried to seed the White House and a couple of commercial places with the Alto three megabit per second network. The three megabit per second Ethernet never really made it into the outside world. It made it into corporate Xerox back at Stanford and places like that. Our goal was to build the next generation.

Pelkey: Who else was on the team?

Dalal: In addition to myself at the XNS bottom levels, there was Will Crowther who was an ARPA/BBN person who had come to Xerox who is now back on the east coast with BBN, and Hal Murray, who is again at Computer Corporation of America and had known Bob Metcalfe and had just joined Xerox.

Pelkey: Did he work with Tom Merrill?

Dalal: Probably.

Pelkey: He had worked with Larry Roberts on the experiments in the mid '60s.

Dalal: So Hal was the engineer's engineer, so he was doing more of the experimentation, the coding, the prototyping. I was doing some of that, I was doing all the spec-ing, and actually the design work, with some of Crowther's assistance, and then Crowther went back to the research lab. He decided that working with Liddle and Metcalfe and me was just not researchy enough. Too many decisions of everyday -- so he went back, and I basically took over the XNS lead and tried to define XNS as it exists today.

Pelkey: And this would have been in '78?

Dalal: This would have been in the '78 time frame. By the end of '77, early '78, XNS as we know it today had more or less become defined. Some changes were impacted by our relationship with DEC and Intel, and we started with -- the 20 megabit per second Ethernet that we were trying to develop had cost problems and it had --

Pelkey: But '78, it wasn't until '79 that you got involved with --

Dalal: With DEC and Intel. So XNS as we know it today was more or less defined by the end of '78, but used 32-bit IDs, not 48 bit IDs.

Pelkey: End of '78 or '77?

Dalal: By the end of '78.

Pelkey: And you took over XNS?

Dalal: In early '78, and those people were working on Ethernet at that time and were making sure that this architecture that we were developing -- in many ways, the entire ISO protocol stack -- was self consistent, so if we used, at that time, 32-bit host IDs at the Ethernet, we used 32-bit host IDs in XNS, and Jim White was working above XNS on defining Courier and some of the higher level protocols.

Pelkey: And Jim was in a group related to yours?

Dalal: Well, Jim and I and everybody reported to Metcalfe. Then, with time, Jim had a lot more management experience so he became the manager of the group -- the whole communication software group -- while Metcalfe just took on bigger and bigger responsibilities at Xerox until he left in '79 to go start nothing, but what became 3Com. So in '79, Metcalfe left.

Pelkey: Actually, he left at the end of '78, because for five months in '79 he was at MIT.

Dalal: That's right, so he left at the end of '78, and that's when I took over most of the management of the networking group. Jim worried about the higher-level protocols and some of the definitions, and I worried about internetworking.

Pelkey: You were aware of the ISO model?

Dalal: We were aware of the concept of the Open Systems Interconnection reference model. Hubert Zimmerman and those people were trying to define that. We tried to influence many of those thoughts.

Pelkey: When you say 'tried to influence' how do you mean that?

Dalal: The OSI reference model had a datagram level to it and we, I guess, were a little bit more applied people, so while the reference model has a lot of formal verbiage, we tried to inject the layman's translation of that, or what a programmer would think. So we were aware of that. We considered it somewhat academic, that it was a reference model that attempted to formalize what all of us knew, and the reference model, as most models did, attempted to concentrate on the bottom layers because people had experience with the bottom layers, and the higher levels became sort of

Pelkey: Fuzzy.

Dalal: Fuzzy, because you hadn't got to that yet.

Pelkey: The divergence between the OSI reference model and TCP/IP, do you have any views why TCP/IP, where the resident knowledge about networking protocols really existed in the early '70s, why that community didn't impact OSI?

Dalal: Well, I think for a couple of reasons. One, the history of TCP/IP and Ethernet is interesting in that it is really computer scientists or computer systems people who defined a communications standard, as opposed to classical communication people, and while a lot of those people were theoretical and had very formal backgrounds, they were also hands-on people -- people like Metcalfe and Boggs -- so they understood reality, they understood what made prototypes work. So, while the reference model was exciting to them, we found some of the excessive abstractions to be academic exercises, and we were more excited by building. The influence was nominal, primarily because Xerox made it difficult for us to talk about some of our experiences, and many of our differences on how we viewed the model had to do with the higher-level protocols. For example, everybody talks about the presentation layer, and the

presentation layer is one of those interesting layers that nobody quite understands what it means. It makes a lot of sense if your coming through a terminal-oriented network, but when you're working on a computer to computer, or at least a personal computer network, the concept of a presentation layer doesn't exist very well, and I think what's happened in the last couple of years with the cooperation of Open Systems and ISO protocols being defined with NBS, they have sort of come to peace with what each one of these layers means.

Pelkey: You continue to communicate within that world. You stopped going to the TCP meetings and the communication slowed down, but you were on Arpanet and you still had an interest. Vint, who was responsible for this international networking group, which was where the culmination of this conflict between TCP and the ISO reference model really came to a head and then got resolved, was it not?

xxxxx

Dalal: The (unintelligible) got resolved --

Pelkey: And in the United States community, Vint resigned his chairmanship of that, and the TCP community started to lose influence over those events.

Dalal: That's right, and that's primarily because TCP was still part of the Department of Defense, and ARPA became more officially known as DARPA in an attempt to do that, and at one point the management of ARPA was taken over by DCA, the Defense Communications Agency, and became a lot more constrained. It was at that time that MIT got involved with putting TCP onto local area networks. Basically, Dave Clark's group and a certain amount of work started happening at Carnegie-Mellon, and I think it was roughly at that time that Jim Morris might have gone to Carnegie-Mellon from Xerox, and the whole concept of integrating local area networks into TCP/IP began to happen. I think that's when the ISO reference model and TCP tried to sort of fit back together. Once Ethernet started making it's way out, NBS got involved. This was Rosenthal's group at NBS that started getting involved. To be honest with you, there was a certain amount of intellectual snobbery that was going on on both sides. I think some of the formalism people wanted to do everything precisely, and some of the hands-on people said: "Here we go again; standards being defined by people who have never built anything."

Pelkey: It is very constructive to understand, because an awful lot of planning energy went into these things, and here now in the world we have TCP/IP which is out there working, but now we have OSI that's going to become dominant, and for some period of time we're going to be very involved in gateways and so on, but presumably, eventually, TCP will disappear. It seems to me that a real opportunity was lost at that point in time that has now created confusion in the marketplace because of an inability of Vint to sell something, and some of it was snobbery.

Dalal: Some of it was intellectual snobbery, and then sort of conversely, many of the people that were attending the ISO meetings -- not people like Zimmerman who were intellectually --

Pelkey: Computer science oriented.

Dalal: But you started getting people from some of the ANSI groups that were responsible for ADCCP, which is a low level protocol, and you got large corporation backing people who were professional standards makers, and that's all intellectual snobbery, that's how it showed through. They attended these meetings that were always in fancy places -- hotels or clubs -- and so sometimes we thought that people were attending those meetings more as a way to break out from their normal drudgery. In retrospect, it did hurt some of those efforts, but the interesting thing is that TCP wouldn't have made it into being the pseudo-standard it is today if it hadn't been for a very fortuitous set of events that had to do with the fact that Bob Kahn funded UC Berkeley to put Unix onto PDP-11s and the VAX.

Pelkey: The gentlemen there, Bill Joy, thought he knew a better way of doing TCP than you did and recoded it.

Dalal: He recoded it, and that was fine, but --

Pelkey: Then it got out in the source license, and all of a sudden TCP started showing up everywhere.

Dalal: Well, what happened is every minicomputer company got into networking without having a networking guru because they got TCP/IP for free, so in many ways TCP/IP's success is based on the fact that it was available.

Pelkey: And that happened in '83?

Dalal: Roughly in '83.

Pelkey: Was that the 4.3 release? How important did you consider MIT's port of TCP to the PC?

Dalal: I think that Dave Clark's group did a lot of work in bringing TCP to local area networks and to the Alto and to some of the PC market, and I think he had lots of case studies of TCP onto the PC -- you mean the IBM PC?

Pelkey: Correct.

Dalal: I think they did a lot of groundbreaking work, but I think TCP on the PC happened a lot because of Interlan, a company in San Jose. TCP/IP was sort of floundering, and 3Com and those people were building Ethernet boards for the PC, but Excelan decided they were going to build --

Pelkey: Which company?

Dalal: Excelan -- the one here -- to built Ethernet boards for the PC with TCP/IP on it, so various people, like Excelan, started building boards that had TCP/IP on it. That did a lot to bring TCP into the PC world.

Pelkey: I want to come back to that in a minute, but it's Dave Clark's view of the world that getting TCP on the PC, and having a standard version out there to talk to other systems, was maybe not as important as Bill Joy and the Unix implementation, but was very, very significant.

Dalal: You indicate -- the criticism of TCP/IP is that it's these big, monolithic --

Pelkey: -- and all of a sudden now it became available, retail stores were selling it.

Dalal: So he demonstrated the fact that you could have a real lean, mean implementation of TCP. He didn't commercialize it enough; that is, commercial companies like Interlan, somewhat in Massachusetts, but mainly Excelan here that brought it into everyday life. ARPA put it into Unix through Bill Joy and into IBM PCs through Dave Clark, and in my opinion, the greater of those impacts that made TCP/IP succeed was the Unix.

Pelkey: Absolutely. No doubt about it. The other thing, there was a negative in terms of the local area networking aspect of TCP, which is the more dominant at least at the moment in the market, was that XNS -- Xerox couldn't get off the dime, in terms of making decisions about whether or not they wanted to open this up. They started the process of making the protocols available, then somewhere along the line they said: "Wait a minute. We can't do this. We're giving the shop away," and people like 3Com and Bridge and other people -- everybody universally says for local area networking XNS is far superior to TCP. I think there may be one or two people who challenge that, but they get drowned out pretty easily.

Dalal: That's definitely true. In fact, I remember working with Dave Liddle. I was on the team that was responsible for Ethernet once Metcalfe had left, with DEC and Intel, while Liddle worked on the many higher strategic issues of pushing that through at DEC and Intel, but many of the people that worked with me and were responsible for making the engineering decisions happen, and when we talk about it David and I just sort of marvel at how hard it was for him to make that happen.

Pelkey: The Blue Book?

Dalal: The Blue Book, just getting Xerox Corporation to make this into a standard. Even though we talked so many times about how Xerox had two standards -- had two products: they had the Star System which is a turnkey, you name it; and they had a network architecture -- and yet, and I kept telling David: "Why can't we get Xerox to sell or to standardize its network architecture product in addition to its commercial product?" It was just, again, a feeling that Xerox would conquer the world. In many ways, those of us who were on the Star team felt very much like Steve Jobs and the guys on the Mac team. We had something, we knew it was cat's meow, and important revolution, and we were going to make it succeed. Other parts of the company had similar visions of grandeur, but were afraid that if they gave the standard out, they would lose control. That was the biggest mistake Xerox ever made. Over and over again, we tried to make it happen, and by the time it happened, when I left Xerox to go to Metaphor, Liddle and I asked Xerox for Xerox's protocols so that Metaphor could implement them and make products compatible with the Xerox system. They said no, so Metcalfe, Metaphor, all of us used the lower levels of XNS and we never used the higher levels of XNS, and by then Unix had begun to take off in the mini world, before Sun even got there, and Sun just sort of completed the acceptance with workstations having TCP in them.

Pelkey: And once that happened it was just -- it was a negative act on the part of Xerox, of not getting XNS, because if they had put it in the market, who knows how history would have been written? The fact is that XNS, for local area networking was superior, and you had people out there who had come through Xerox and knew it - - Metcalfe off at 3Com -- the local network companies, at the time, would have gone to XNS.

Dalal: In fact, 3Com would have gone to XNS, Metaphor would have gone to XNS, Bridge had already gone to XNS, it would have happened. Even as late as '82, '83, when Metaphor started Xerox had a slim chance, but by the time -- XNS would have been here, and I think, without letting my vanity show through here, the transition from XNS to ISO would have been a lot easier and simpler, and I don't think people would have stood on their heads as much. Right now, TCP makes it a little awkward. There are gurus out there, and any awkward situation can be resolved with a few more changes here and there. It's engineering. It's just a tragedy that something that was architected to be designed will never make it. Many of the ideas in XNS that Jim White had with Courier, or I implemented a thing called clearinghouse on top of it which was an electronic name look-up mechanism very much like the White Pages and the Yellow Pages, those ideas are all getting out, and everything that Xerox created is being replicated and duplicated, and sometimes for the better. It's a shame that something that was so important just sort of fell away.

Pelkey: There was a level of intensity around that period of time. The group of you now had to work on this thing for a number of years. There was a group of you who were incredibly competent, and there was a sense that you were doing something that was important. I'm sure that the group of you were having fun, so it was very, very intense. That's not to say that people don't work intensely in the world today, but that was a unique period of time, in terms of what was happening. I can only imagine that, in 1979, when MITRE held these conferences, the group of you who were into LANs and knew what was happening, that must have been a very exciting period to see that the rest of the world was starting to come to recognition that this stuff had importance.

Dalal: Greg Hopkins was at MITRE in those days, and then he subsequently went to Ungermann-Bass, and I remember the days that Charley Bass and Ralph Ungermann left Zilog and hung around Xerox trying to figure out what they were going to do next. It's funny, when you look back on those years at Xerox, everything we did has made its way out. The only one thing that made its way out intact was Ethernet. Everything else is a variant. PostScript is a variant on Interpress; Mac and Windows and Open Look is a variation on Star; TCP or ISO is a variation on XNS; there's an analog to everything. The only thing that made it out intact was Ethernet.

Pelkey: Now, during this period when you were at Xerox, was there a contact between yourself and Zilog? Judy Estrin was over at Zilog --

Dalal: There was a lot of social contact because of the fact that a lot of Zilog people were ex-Stanford people, so there was just a lot of -- we were all in our late 20s, and there was a lot of camaraderie, but very little technology transfer. Zilog was trying, in those days with Charley Bass and others, they sort of thought of Zilog as somebody that would commercialize products in the same way that Xerox would, so they were trying to build on the Z-8000 platform a workstation. Again, that was a heyday of very exciting people, and Zilog just sort of collapsed under its own weight, I think.

Pelkey: Trying to do too much.

Dalal: Too early.

Pelkey: Helter Skelter.

Dalal: We at Xerox were sort of naive people that had never built a product, but so were the people at Zilog. We just had tremendous technology resources and some of the brains at Xerox were just -- are legends now.

Pelkey: Do you recall Xerox's decision to open up XNS -- sell it or license it -- and then, at the higher levels, they just said: "No, we're not going to do it anymore."

Dalal: What happened was, with Ethernet being up there, they felt that you had to give the next levels out, so they got XNS, in terms of the datagram and the session protocol and Jim White's Courier, which is a way to exchange procedure calls back and forth, but they were hesitant to send out the filing protocol, the printing protocol, the name look-up protocol, the electronic mail protocol; all the protocols that you really needed to do something. Again, I think they felt the lower level protocols it was ok to send out so that the Bridge's and the Ungermann-Bass's could build connectivity hardware, but nobody could really build servers that would compete with Star or the filing systems that Xerox was producing. We, even at Metaphor, Liddle and I sort of hammered at Xerox, as late as the summer of '83 saying: "Look, guys, this is insane. It's the summer of '83. So much has happened. This is water under the bridge. Forget the fact that we're no longer at Xerox. We're trying to popularize what you've got," and it was: "Oh yeah, well, we've got to do this, we've got to do that," and by the time it was published, I believe in '84, '85, it started making its way out. Jerry Elkind and his group were responsible for creating the real standards. Bill Lynch was there with Liddle and Metcalfe -- they finally got the protocols out. By then it became such a marketing battle of convincing anybody to use these protocols. By then the Mac was out. Xerox had lost it's --

Pelkey: It's prime position -- it's uniqueness.

Dalal: The world didn't think it important to implement the Xerox protocols, because it was like: "Why should we?" There was a certain amount of resistance on other people's behalf, which was: "Why should we implement Xerox's? We can do it better ourselves, or we can do it slightly different."

Pelkey: You went to Metaphor because of the challenge of what was to become Metaphor, and I suspect some level of frustration of being with Xerox and not seeing anything commercially being successful?

Dalal: Right. For a long time Metcalfe tried to get me to leave. I could have been employee number two or three at 3Com, and I resisted for what were great reasons then, which was, I kept telling Bob that I had unfinished business. If I had left, XNS would never have happened; a lot of other things wouldn't have happened. On the other hand, if I had joined Bob, who knows what we could have done at 3Com? I, in retrospect, felt it was the right decision, but towards the middle of '82 I was growing frustrated and feeling like my hands were tied, as were Massaro and Liddle, and so when those two decided "well, let's try to do it on our own," it was the perfect opportunity.

Pelkey: Were you working with Charles at that time?

Dalal: Yes, my last . . .

Tape Side Ends

Dalal: In my last year at Xerox, I decided that I had spent a long time doing data communications, and maybe I should do something different. In fact, I think that decision may be why I have the job I have today, which was the connectivity, the plumbing; XNS, Courier, Clearinghouse is all being handled, so I decided to work with Charles in the Advanced Development Group, to figure out what new, exciting services could you build on top of this network foundation. I'm not saying filing was passé, but in the Xerox sense of the word, the problem had been solved. There were implementations, electronic mail existed, all of this stuff existed, so I was working with Charles to figure out what would voice be like on a local area network, or what would scanned images look like -- getting more and more into sophisticated distributed applications.

Pelkey: I have no other questions. I think we have touched upon --

Dalal: I think we have touched upon it. I just tried to make some notes myself, just to understand what happened. I think XNS played a very important role in defining TCP/IP, and very few people realize that TCP/IP as it exists today -- is based on so much the PUP and XNS experiences.

Pelkey: Yes, and PUP was done largely by --

Dalal: David Boggs, and Ed Taft.

Pelkey: Ed Taft, and when you went over in the end of '74 --

Dalal: Early '77.

Pelkey: -- PUP was working -- only internally, but it was working -- and when you were given charge of developing the next generation networking protocol, seeing what was happening with the Ethernet and being exposed to PUP, finally after not having been able to get access to it because of Xerox policy, you started to go back and influence, and develop this datagram orientation as well, which led to the IP?

Dalal: I think that is very fundamental. I think TCP without IP would have been a disaster today. I think it would have caused a fair amount of agony by the community that's trying to build sophisticated stuff. Sun would be hurting. They would have invented their own IP because they wouldn't have been able to live with just the TCP, but I think TCP is where it is today primarily because of Bob Kahn's efforts in seeding UC Berkeley through the Unix. It's very ironic, but --

Pelkey: Let me ask you a question of a different kind: during this period of time, the United States funded, through ARPA, some major systems development. The Arpanet, which was graduate students working on Arpanet in the late '60s and early '70s, and then we had Xerox PARC, with the whole concept of Star and the office environment. At some level you had Zilog, which was different, but there was a lot of money being thrown at pursuing the big vision. It was these major interdisciplinary projects that were really critical to advancing the state of technology, in contrast to any given individual, no matter how great his contribution, it was the systems orientation towards bringing something forth which was unique about those projects.

Dalal: I agree with you 100%. I think people like Bob Metcalfe deserve the glory and the fame that they have, because they were powerful personalities that stood up for what they believed in, but even Bob couldn't have achieved his vision and his goals if it hadn't been for the systems research, or emphasis, whether it was by Xerox or by ARPA; whether it was trying to get DEC and Intel together. So, I think Ethernet exists today because of the massive understanding by people in the semiconductor business,

people in the personal computer -- meaning Xerox in those days -- and DEC in systems integration that this could, in fact, change the world and impact it.

Pelkey: And today we, unless I'm unaware -- I mean, there aren't these major systems undertakings. Bob Taylor's trying to do something at DEC now, but there really is a lacking of this kind of efforts. When ARPA went to DARPA, that profoundly changed what was happening there, in terms of more immediate impact on defense systems. On another front, we lost Bell Labs after divestiture in terms of being a major research center. It's a different kind of research. In terms of us getting back into research with this kind of systems view, which has been so important to the economic growth you see in these technologies, we're just not doing that very effectively.

Dalal: We're not doing that very effectively. It's interesting in that one of the reasons why I think ARPA was as successful as it was, was all of us that worked on the Arpanet never thought about making millions of dollars as entrepreneurs. We were in it for the challenge, for the excitement, for making an impact, for changing the world in one way or another. Those were our ideals. We were children of the '60s and '70s, I guess, and this technology that we have unleashed on the world is not commercialized. There are so many more players involved -- people trying to make money, and that's good. In fact, that's ARPA's ultimate goal, to insure that this technology makes its way into the world, but will Bob Taylor achieve the impact that he achieved at Xerox PARC? Probably not, because Xerox was on its own. There weren't any competitive pressures pushing it in any direction, and the interesting thing in my opinion is how will NRI, this new thing that Vint and Bob are on, succeed? How can they recreate that kind of mystique, almost? How can they set their sites far enough in advance, or at least ahead of the commercialization potentials so that you're not driven down into profit and compromise and quick turn-around?

Pelkey: A question of another kind. I shared with you earlier my view that networking is going to change the way we run our organizations and manage them, do you subscribe to that point of view?

Dalal: Absolutely, in that just in the way Xerox did it's business, the most useful thing of the Arpanet was electronic mail. At Xerox, mail was responsible for keeping Liddle's El Segundo Division and its Palo Alto division going. It was a very bad situation, where you had these two groups so far apart working on such a state of the art technology. Without electronic mail, that would have been a complete disaster. One of the things that I believe, and what you said earlier, is that hierarchy is necessary in every organization. It's sort of the natural law. I think that over many years of evolution, we always think of hierarchy as --

Pelkey: Part of the genetic code.

Dalal: The genetic code -- but when people talk about third wave companies or the companies of the 21st century, a lot of stuff happens through the informal channels of communication, and if you make those informal channels of communication be high bandwidth, then the right people get together to solve the right problem. It doesn't always have to go to the top. It's a little scary for people to conceive of organizational dynamics changing rapidly where the hierarchical structure isn't as much in control as it has been in the past, but never the less, if that change happens, piecemeal, it will happen, at least the way I view my role at Claris, and I hope that my peers view their roles elsewhere, is we're products of this information technology, having been around when it happened. Now that we're managing executives and fancy titled people, how do we run our organizations? Do we allow that informal communication channel that exists because of technology to impact the way the organization changes?

Pelkey: I view that as a great challenge, one I hope you're successful with.

Dalal: One of the things that I try to think about -- I think that this country really needs this whole concept of data highways, and whether it's large organizations like IBM or AT&T or somebody like that and forcing them on us, or its evolution through things like Corporation of Open Systems and commercial forces. In fact, that's how AT&T was regulated. They provided us the telephone network, but on the other hand, the television network sort of came about through competitive juices. There was regulation, but sort of

pseudo-regulation. Now that we're into telecommunications and local area networking, how do we ensure that --

Pelkey: Clearly the bandwidth is important because video is the bandwidth that's got it, and fiber and the local loop and video is going to be -- the world doesn't know what it's -- it's the future. It's not ISDN.

Dalal: No, no, ISDN will be just a stepping stone. I think, unfortunately, this is literally too little too late, but it will play its role in that it will replace modems for the most part, and we'll all forget what modems are about and move one step further.

Pelkey: I agree, thank you so very much for your time.

Dalal: My pleasure, I really enjoyed it.

END OF THE INTERVIEW