Oral History of David Liddle

Interviewed by:
James L. Pelkey

Recorded: October 11, 1988
Metaphor Computer Systems
Mountain View, California

CHM Reference number: X5671.2010

James Pelkey: What were you doing before you got involved with Xerox PARC [Palo Alto Research Center]?

David Liddle: While I was still in graduate school, still working on my PhD, I worked for Owens Illinois in Toledo, a company that had a big contract—had had a big contract—and then themselves were making big investment in display technology of a certain kind, namely the plasma panel. The original patents for the plasma panel, the original fundamentals, were developed at the University of Illinois, and then Owens Illinois in Toledo sort of took that over as a development project. I had known—while I was at undergraduate school at Michigan—I had known some guys at ARPA [Advanced Research Projects Agency], and they sort of suggested to the people at OI that maybe I would be a good guy to come and do some of this display system work, because Owens Illinois was a major manufacturer of display devices—TV tubes and CRTs and stuff like that—and was anxious to be in this business, but had not done very much with graphics systems and that sort of thing.

Pelkey: What year was this?

Liddle: This was about '69, so I went to work there on this display project, and I got very much—we had an ARPA contract, and I got very much involved with the ARPA community. Steve Crocker asked me to be—no, Christ, it was Cordell Green—asked me to come back and talk with Larry Roberts, which I did, and they asked me to be on the ARPA Network Working Group, and specifically on the Network Graphics Working Group, and so I attended the first meeting of that. It was me and Charles Irby and Bob Metcalfe and lots and lots of guys. We all met at a conference held back at MIT, out at Endicott House. That's when I first met both Irby and Metcalfe, among others.

Pelkey: Do you recall what time frame that was?

Liddle: Oh, I guess that would have been early—it was warm, so it was spring or summer of 1970, I would say.

Pelkey: So that was your first involvement with the ARPA community?

Liddle: I had been involved with them for a few months before that, but that was when I first really met these guys in person and showed them what I had been working on, showed them a video of my product. I guess I had made a film of it. I couldn't bring the system to demonstrate it, but anyway, I showed them

this network workstation that I had built, which really—I don't mean to be so first person about it, but it was me, I had done it all by hand. It had this high-performance plasma panel and it was quite an intelligent terminal, but I had really bought, at that point, a computer from Data General and just stuck it in the middle of this terminal, and was interested in implementing network protocols directly in this terminal so that it could be connected and speak the right protocols more like a host, and at least take the graphics protocol burden off of whatever host there was and do all that translation locally. At that point, I was just zinging simple ARPANET Telnet stuff around, so there wasn't anything too exciting, but I met a bunch of those guys, including Eric Harslem, who later came to work with me on the Star at Xerox, and is now at Apple, of course. But Eric was down at Rand, and they had the big Rand Video Graphics System. Charles, of course, had gotten to know the ARPA guys because of the color stuff with Glen Culler, the OLS stuff down at Santa Barbara.

Pelkey: Was Charles with Glen Culler?

Liddle: Absolutely. Charles was one of Glen Culler's original...

Pelkey: That's where he knows <*inaudible*>.

Liddle: Absolutely. Charles was one of Glen Culler's original protégés, and by then Charles had left and gone to SRI [Stanford Research Institute] to work with Engelbart. I was really impressed by Metcalfe and by Charles and Peter Deutsch, who was then already at Xerox PARC, and I knew some of the guys at Xerox PARC because I had known Alan Kay a long time. Alan Kay was lover of plasma panels, because he saw them as potentially fitting into the Dynabook. Since we had both been in graduate school, I've known Alan Kay, and we had exchanged messages and thoughts and stuff on and off. When I finally finished my PhD, I knew a lot of these ARPANET guys, because I had worked on this Network Working Group and the Network Graphics Working Group. I guess we can come back to that later if it’s important, but that was where I met some of the other networking guys from Mitre and MIT and all those other places.

Pelkey: Where did you get your PhD?

Liddle: The University of Toledo, because it was right next door to Owens Illinois, where I was working.

Pelkey: And what year did you finish that?

Liddle: In '72.

Pelkey: It was in computer science?

Liddle: Affirmative. And so I came out to California; I called up Charles and said I was near to finishing my PhD and was sort of interested in and thinking about a job, and he said “Yes, come on out and talk to us." So I came out one time, because, I think, I was on a business trip and I dropped in and talked with Charles and with Engelbart and these other guys, and I just thought it was so wonderful, and Doug loaned me a copy of his film. So I went home and watched it and I made up my mind, that's where I want to go. I want to work with those guys at SRI. So then I made an appointment for a slightly more formal interview, and I came out about six weeks later, and it was really wonderful. I sat in on their whole process and everything, and I really liked it. Because I was out here, I just called up some of the guys I knew at Xerox to just drop in and say hi, and Alan Kay said: "Gee, if you're interviewing at SRI, why don't you interview here?" I said: "What, little me come to Xerox PARC?" Of course, it wasn't terribly old then anyway, but they said: "Sure, we're looking for people that have had good ideas or that are working on things that we think are interesting," and as it turned out, I knew Butler and I knew Dick Sharp and I knew a lot of these different guys because they had invited me to—Alan Kay had invited me to a Xerox-sponsored conference that had many, many people at it—the so called Alta Conference—a year or so before, so I knew a lot of those guys, and while I was there they kind of set me up to do some quick interviewing, and then they said: "Could you come back next week and describe my thesis research" and everything, and I did, and so they offered me a job. I reported for work about the same time Metcalfe did. I guess he had been, technically, on the payroll for a few months before that, although he was back and forth from Boston. That was October, anyway, of '72 when I started there. I worked on a lot of stuff. I have a memo that I just found the other day that I wrote in December of '72 called 'Debugging the Dynabook.' We were working on the design of the Alto in a very abstract way at that time, and I had been sent down to El Segundo to look at a micro-programmable machine that they were making down there, and I looked at that. Chuck Thacker was designing this microprocessor that we were thinking of using for what we later called the Alto, and anyway, one of the concerns was that if you build a complete micro-programmed machine, and it controls its own display and it kind of controls its own everything, how do you bootstrap it? That is, how do you know when it's first working? So, there was some little widget that I had designed that would sort of hover over this microprocessor and trap on certain instructions, and that kind of stuff. It was just a little piece of pick and shovel engineering, but it's just interesting to me because it's the first little piece of design of any significance that I did there, and I always thought that had a certain lyrical name, 'Debugging the Dynabook.' We didn't call it the Alto until quite a bit later when it was, among others things, clear that it was not going to be the size of a book. In fact, I remember when Chuck Thacker, in our little weekly discussion meeting—'Beat the Dealer' it was called—he said: "It's been decided that the Alto will be housed in a 30-inch tan featureless cube," which I always thought was a

great way of describing it. That's exactly what it was: a featureless cube. That was just sort of funny. Here was this thing that had this wonderful display, and could do all these spectacular things, but the rest of it was a featureless cube. To try to talk about networking, I suppose, we actually did a lot of different kinds of networking.

Pelkey: You had a Data General network.

Liddle: Right. We were on the ARPANET.

Pelkey: <*inaudible*>. The first thing was to get the Data General machines you had on the ARPANET, as I recall.

Liddle: Right. I did some of the monkeying around with the MCA. Other people did the real work of putting that network together and all that kind of stuff, but there were some programs—Bob…nevermind, I won't use any names—there was one really incredibly brilliant and capable guy who worked for us who only could work part time because he also had a lot of other commitments, so he came in and worked for us for a while, then he was back in graduate school, then he came to work for us, then he went back to graduate school. He was, in some respects, one of the most productive people that ever worked at Xerox PARC, but he had written a program that used the MCA exactly as if it were part of the ARPANET, and it was a very neat program. You could sit at any terminal on it and log in and pass through to any other machine and all that stuff—it was really very neat—but it just wasn't reliable. For some reason, nobody really quite understood that, and at that point I was very excited about operating system stuff. I already had done my PhD, but I had sat through Butler's operating system course at Stanford. I was doing post-doc stuff at Stanford, and I took the same course over and over when it was taught by Forrest Basket and by lots of people. I was very interested in low-level issues of synchronization and contention and process management in communication environments, and so I monkeyed around with this program for a long time and just discovered some fundamental critical sections in it and sort of fixed them up. So I sort of got grafted onto some of the networking activities for a while.

Liddle: Charles Simonyi, who is now at Microsoft, of course, but at the time was still finishing—in fact, is still only part way through graduate school, I guess at Berkeley and was about to move over to Stanford and work on his PhD—but for a while he thought he would do a network. It was going to be called SigNet, which he said was an acronym for Simonyi's Infinitely Glorious Network, and SigNet was going to have little IMPs, called micro-IMPs, that were exactly like the ARPANET IMPs, little message processors attached to each little local host, even though they were going to be Altos and Novas and stuff. That was Charles' notion, and one of the problems was that it would have to be rather high speed, because he

wanted the SigNet to run extremely fast, that is it was going to be a very, very high bandwidth network. Well, Charles was basically a software guy, and the hardware guys were interested in this, but they had other fish to fry. The hardware guys who could build this hot shit high-speed local network were actually working on the Alto and working on other processor work.

Pelkey: Which would seem to make more <*inaudible*>.

Liddle: That's right, and there wasn't quite access for somebody to go and put a lot of effort into that. By the way, he himself was doing all kinds of other worthwhile stuff and working on a doctorate and so on, so this thing didn't rush along real fast. In the meanwhile, Metcalfe was finishing his—he and I were both, actually, finishing writing our theses while we had already come to Xerox, and he was thinking a lot about packet communications issues anyway. He had studied ALOHA and other things like that. Now, I'm going to tell you a story which I believe to be true, but I am willing to accept the fact that it may be apocryphal and I have convinced myself of it. Do you understand what I'm saying?

Pelkey: Absolutely. I've heard many of these.

Liddle: Yes, and I'm not sure that this one will be quite correct, but this is as I remember it. I do remember that there was talk about making this ALOHAnet right in Xerox, right in the building there at PARC, this little packet radio thing, and there was concern about whether or not we could make those little transmitters and receivers such that they didn't also screw up all the other computing devices that we had around, because RF wasn't our bag, it wasn't our strong suit. So this observation was made, I think, by Chuck Thacker, that you could just do all this with coax, because coax was just nothing but captive ether—as in there luminiferous ether—and I remember Butler talking about it a couple of days later at 'Dealer' saying: "Yes, we're going to put coax all over the building to take the place of the luminiferous ether, and thus we won't have to worry about whether we've done a great design. We can just use sort of digital components and just smack the thing with five volts, and not worry about delicate little 'Fancy Dan' modulation that makes all that shit so expensive." So this was about to be ok, and I remember Metcalfe writing a memo—I read a memo that Metcalfe wrote—about a brancher, and the memo said something like 'in search of a passive brancher.' He always wrote these kind of funny, scribbled up memos, and he was back to his early electrical engineering days, sort of drawing little circuits with resistors and all that sort of stuff.

Liddle: Well, when I was a struggling undergraduate, and feeding <*inaudible*> and little kids and everything to take care of, I worked for the Detroit Board of Education in their Engineering Department, while I was in school, and one of the things that I did was to do the analysis and layout and design and so

on of cable TV systems that were being installed in the schools under a big federal grant. That was all good fun stuff, and I went up on roofs and measured microwave signal densities and did stress analysis and made drawing of how tall this tower would be and how strong the cables had to be; all that pick and shovel stuff. It was good fun, I made money, I could do it at times of day that also allowed me to go to school. Well, not to spiral in on that, but at that time, the cable TV technology was quite new, and the only places that it was really used were things like motels or things like that, and for these educational TV sorts of situations. So, you had a specialized antenna and amplifiers that amplified the appropriate channels and sent the signal down a parent coax, and this coax was tapped as it went by the rooms, with a more or less impedance matched tap. So, I knew all about this. Nobody else at Xerox had ever had anything to do with anything so plebian as that, but I went to Metcalfe and said: "We don't—this is not what you want to do. You don't want to try to build an elaborate branching circuit. What you want to do is to make a tap that's impedance matched, and then it'll be very easy to install these things and move them around and all that kind of stuff." He said: "What the hell do I have to do to make one of those?" And I said: "No, this company Gerald was that company that—it was its name on all this stuff that the Board of Education bought with this federal money. Why don't you call them and they'll send you a catalog and so on," which he did. So we started off by using the so-called Gerald taps. Now, although they weren't great, they got us instantly started, and in no time at all, these guys had built a workable network using these stinger-type taps. Now, after that, of course, they learned how to make real Ethernet taps that had various other—were better impedance matches and had various other wonderful properties, but these, of course, were early days and it was a low-speed network and…

Pelkey: You needed to get some kind of <*inaudible*> to get it to work.

Liddle: Exactly right, and that was how all that proceeded, and of course it snowballed very quickly. I mean, the damn thing just worked. It was a very easy concept, and inside and outside Xerox both, there were all kinds of people who were very telecommunications oriented who thought our next thing was we should rush off to Collins Radio and get them to design some complicated goddamned thing, and of course that would have meant you had a modem or a CB radio or something on every desk. I said: "No. I mean, we know these signals are sloppy, but it doesn't matter, they come out right at the other end, and by the way, if they get fucked up, they just get fucked up. You just retransmit them. It's not that big a deal." So we stayed with this very loose design, and it worked out pretty well. So, apart from that tap business, at that point, I was just a beneficiary of the Ethernet.

Pelkey: But, now May of '73, when Bob sent a memo around where Ethernet was used, there was a point in time at which Bob needing some help to do this project, and he said he came to you and you suggested that he go to this guy…

Liddle: David Boggs. David Boggs was, at that point, working for me. I was working on this big office system we were trying to build.

Pelkey: This was a combination of Altos and Novas?

Liddle: In fact, there weren't any Altos yet. There were only a few stand-alone Altos at that time. Of course, then there wasn't any Ethernet, but we did have the Novas on the MCA, so it was originally designed to be all Novas; a pool of Novas on a network, and a very elaborate system that—the Video Terminal System, it was called—and it was quite neat. You had this really gorgeous, wonderful, high resolution terminal on your desk. Down in the bowels of the building, there were a pool of character generators—very elaborate, but still fixed function—and they could generate different fonts and stuff like that to go on this screen. I had worked on the electronics of this display even and stuff like that. There was this big rack full of character generators, and you were patched into one of those to the terminal on your desk, but then you connected—there was one or two Novas that served all those, but they were really just gateways to connect you to the Nova that was going to serve you, and there was a big pool of all these Novas. This was how—in other words, this was a distributed processing thing. Instead of a single host, there were all these different Novas that would help you, and you could dial out through them onto the ARPANET or to other hosts in other places, or you could do your computing right in that machine. The MCA was the backbone of this, and Bill Duval had done very beautiful work in building the operating system for Polos and the communications software and so on.

Liddle: Anyway, I had this guy, David Boggs, who worked with me, among others, who was still working on a PhD at Stanford—in fact had just lately come from Princeton, I think—but he was a really, really clever guy who could do a wide range of stuff. He was really—very soft spoken but tremendously impressive engineering guy, and so I knew that he had been a ham radio guy, and I knew that he was willing to take electrical engineering type measurements—that he wasn't, like Metcalfe, a software-oriented guy, particularly. So I said: "Look, why don't you borrow David Boggs. We'd all like to have this network you're working on, and it would be worthwhile," so David Boggs came up with some extra time to go off and work with Metcalfe. Of course, Boggs is a tremendously powerful thinker. In a sense, this was like introducing Ramnijan to Hardy. The effects of Boggs and Metcalfe working together were very significant. By the way, because this was a big deal that would affect the whole laboratory once they got started on it—Chuck Thacker was building the Alto and Butler Lampson, who was sort of essentially an intellectual parent figure to many ideas that were there, both also sort of got involved with this and made very worthwhile contributions to that design process.

Pelkey: It's my understanding that Boggs got involved subsequent to the proposal of using the coax, and the labeling of what was then 'Ether Net,' and that David was critical and essential to the process of the implementation, but the process that created this—of the ALOHAnet and the ARPANET and the coax—those things, in some measure, were in place conceptually. There was still the business of reducing it to practice in a way that was…

Liddle: Everything you said is true, but the one problem that I have with it is it's very difficult to draw the line as to where things were implementation and where they were creative. For example, I am not sure who developed the binary exponential back-off algorithm—whether that was Boggs or Metcalfe. The idea that, upon a collision, you retransmitted—you randomized the retransmission interval but you made the average value of that random number increase in a binary fashion—I don't even know which of those two guys did that, and I'm not sure that it's very separable. In my view, that was a highly creative thing that they accomplished. That's one of the things that makes the network work.

Pelkey: There's also a professor at Stanford, whose name I'm not going to remember now because it's a complicated name, who did some work around and, I'm told, helped create the label CSMA/CD [Carrier sense multiple access with collision detection].

Liddle: Oh, yeah, the CSMA/CD didn't come 'til years later.

Pelkey: But this guy was a Stanford professor who did work on this, in terms of—because the first back-off algorithms were relatively simplistic, and over time, as the phenomenon got known more, had to evolve and become more complex.

Liddle: Well, there were papers written proving that there were various other algorithms that were slightly better, but the point was you needed a fucking computer for those. You need a shift register to do binary back-off, so Boggs and Metcalfe had the sense and good taste to say: "Never mind. Let's not make a silk purse out of a sow's ear. This is a simple thing that we can make with a one dollar part."

Pelkey: Were you involved with PUP [PARC Universal Packet] at all at this point in time? Were you aware of what was happening over at Stanford with Vint Cerf and—or were you largely just on the display side? You guys were involved in a zillion things.

Liddle: I wasn't on the display side much anymore after that. I was doing more system-level stuff, although also you've got to remember that we actually started—we started SDD [Xerox's Systems

Development Department] in January of '75. I think in '74 was when the PUP stuff first began to raise its head. The PUP…

Pelkey: And it was working in '75.

Liddle: Yes, but the PUP stuff was all done in PARC, it's just that Stanford was the one place on planet earth where they so deeply 'instantiated' it that it became a major part of their network architecture. Nobody else anywhere uses PUP, but they do use it at Stanford very heavily. I remember that…

Pelkey: But there's a real intellectual digression between PUP and TCP.

Liddle: Oh, sure.

Pelkey: And the group of you at Xerox, although there were interactions in later years between yourselves and Zilog, and there were interactions between yourselves and Stanford—witness your going to classes, and there's a lot of interactions between things that were happening there—but Metcalfe wasn't under—as people who described the meetings that were happening at Stanford in that period, '73 to '74—in fact, Metcalfe had to eventually drop out of that process—was that Metcalfe had this hidden agenda, which was a local area networking scheme, whereas the TCP was really thinking about 50 kilobit lines and radio networks and all kinds of internetworking.

INTERRPUTION IN INTERVIEW

Pelkey: So we were talking about PUP and SDD and that there was this secrecy, so there's this digression on TCP and PUP, and PUP was much more confined—it was coax, it was workstations, specifically only a few types of workstations, and it was meant to be up and running and working because you guys needed something. You didn't need to solve all problems, you needed to solve your problem with that. Then, SDD got developed in '75, and Bob left and went to…

Liddle: Went to TTI, and then I got him to come back…

Pelkey: In '76.

Liddle: Right.

Pelkey: He came to work with you at SDD. What was the reporting hierarchy at that point in SDD?

Liddle: At the point when he came back?

Pelkey: Because he was part of your team, right?

Liddle: Well, in SDD, it was started in January of '75 and headed by a man named Harold Hall—a wonderful guy who just retired from Xerox—and he was the boss of it for six months. Then, in June or July of '75, it was headed by a guy named Steve Luchesik, who had been the director of ARPA, and he held that job for one year, until July of '76. Now, when Metcalfe came, I'm not precisely sure when he came back, but—do you have that?

Pelkey: [*Leafing through notebook*] I probably do. While you're talking I'll look it up. Actually, I'm not sure.

Liddle: I don't know that that matters a lot, but I'm just trying to give you the reporting structure. Well, to make a long story short, anyway, SDD—Harold Hall was the first employee and I was the second employee, and I always managed all of the Palo Alto activities of SDD, and as I say, Luchesik was in that job for a year, and then in mid '76 was replaced by Bob Spinrad, who held that job for two years. By the fall of—shortly after he took over—by the fall of '76, I then had all of the development aspects of SDD, everything reported to me. A fellow named Hal Lazar was sort of the program manager/interface to the rest of the company kind of guy. Harold Hall was responsible for—he had sort of a senior staff role, and he attacked many staff problems working for Spinrad, but I essentially had all of the line management of it under me, and that's about the time Metcalfe came back, so he worked right for me. I gave him increasing amounts of responsibility. At first he had just the communications stuff, but fairly quickly we made him responsible for all the—well, he had the system architecture group and he had guys who did the diagnostic stuff and guys who did various things like that. I even think he managed Charles Simonyi for a while. So there was a different group that worried about the operating system kernel and the programming environment. There was another group, located in LA for historical reasons that reported to me that was writing the user level software. Metcalfe had much of the other system software and system architecture kind of stuff.

Pelkey: Your objective was to take the Altos and make it become the Star, a system level product.

Liddle: Yes. I mean, it's more—it's a much more complicated story than that really, because the Alto wasn't really—we re-engineered it a couple of times. We used it very heavily, but it was not the kind of

product upon which you could make a lot of money, because its manufacturing costs were sufficiently high that, to really sell it and make any gross margins, you would have had to sell it for forty or $50,000, so we were redesigning—we were designing Star. We really designed a processor family called the 8000; a network architecture, an operating system, and a programming language that was supposed to be the basis for about ten years of Xerox products—servers, workstations, electronic printers, gateways, everything. So mostly I was pushing that system-level activity, and then there were two product families to come out of that; one was Star and the office related stuff, one was the electronic printing kind of stuff, and all those things did come out, too. The most important thing, from a communications point of view, that happened during that period of time—I won't say the most important thing, because there were several, some of which were very controversial—one was we went from three megabits to ten megabits. The people at PARC complained about this a lot, and said: "No, this is 'biggerism.' Why does it have to be ten megabits? Why isn't three megabits good enough?" We simply said: "We didn't think three megabits would be good enough over this time horizon," because we thought there would be more movement of big files and databases and printing big images, and the product was really virtually playing almost the role of a bus, not an old-fashioned 'beep beep beep' communication line. This was somewhat controversial because, of course, it increased the cost and subtlety of designing some of the components. On the other hand, Xerox owned some disk drive companies then and we knew that read/write amplifiers could be made that were pretty cheap, if you knew what you were doing.

Pelkey: So the concept of going to ten megabits came well before any involvement with Intel?

Liddle: Well, it was right around that time.

Pelkey: Because Al's recollection is that they really wanted to go to 20 megahertz, but the state of the art…

Liddle: We had already decided we were going to do it. Within Xerox, Ethernet meant the three megabit system, and so some other guys somewhere else in the corporation said: "Well, this is a good idea. There has to be a Xerox wire, and so for a while, the ten megabit network was called 'The Xerox Wire,' or 'The X Wire.' That was simply to separate the ten megabit from the three megabit version, which the bore the name Ethernet. Well, we always wanted to go ten megabits. The DEC people were very pleased to find that we wanted to go at ten. The Intel people thought that it should be able to go even faster. Almost everybody else in the industry thought that this was a mistake and we were screwing everyone.

Pelkey: Getting it too high cost.

Liddle: Right, exactly. The Zilog people just howled about this, as did many people. They said that: "You guys are just trying to drive this and make a standard so that all the little guys will be somehow squeezed out." It was all rubbish, but a lot of people were concerned about that. So anyway, a couple of important things happened, but in '76, '77, it was a ten megabit product. There was no question about that; that was the tack we were on. We were working our tails off to design controllers that would work and transceivers that would work and everything. The opportunity, later on, to work with other people—well, one important thing happened. We went from three to ten megabits.

Pelkey: And the second one was PUP going from…

Liddle: Well, it was doing XNS [Xerox Network Services], right, because there wasn't anything wrong with PUP for what it was used for, but PUP wasn't nearly a general enough specification for all the kinds of things that we had to be able to do. Now, Yogen, of course, had been a graduate student. Metcalfe knew him because he was in one of Metcalfe's classes, and Yogen [Dalal] had worked along with Carl Sonnenschein and Vint Cerf and all these guys on TCP/IP, and Yogen knew a lot about what made network protocols work, but in particular he also knew the strengths and weaknesses of those approaches. Yogen and Jim White—who by that time had come over to work for us—and Metcalfe and Larry Garlick, all those guys thought through a lot of those issues very carefully, associated with a wider networking environment. Jim White, of course, had always been this remote procedure call protocol jock, and so we designed the Courier protocol to do that kind of stuff. There were a number of important levels of the XNS protocol, I won't bore you with that now, but there wasn't any other complete protocol like that. OSI [Open Systems Interconnection] was complete bullshit. It was nothing but viewgraphs. It was just all these layers and lots of multilingual bullshit artists were all talking about what a dandy thing this will be in the golden future, the sort of classic European inaugural tea-sipping standards committee thing. The XNS stuff, on the other hand, was just sort of real good pick and shovel engineering and it was made quite real fairly quickly. Well, you understand what I'm talking about.

Pelkey: I don't understand the contrast there. Maybe you can make that clearer.

Liddle: Anyway, we just cranked ahead and did XNS because we were going to do XNS. We needed to have it for our products.

Pelkey: So you really didn't participate at all in this OSI stuff. That just seemed like that was Fantasyland.

Liddle: Well, it was way behind, and all the guys that were involved in it were—the biggest problem with it was it was terminal/host oriented. It was just like X.25. Those guys were all still stuck in the ARPA way of looking at things. It wasn't about terminals and hosts anymore.

Pelkey: <*inaudible*> connectionless.

Liddle: Yes, exactly so, and in fact, the connections were not peer-to-peer connections at all, and so we just forgot all that and said: "Well, this is bullshit. What do you mean, the 'session layer?' What the hell does that mean? Why are there fucking sessions? Send me a datagram, get the hell out of the way. Grab the loot and run for the train, you know what I mean? This is a busy network, goddamn it. I've got code to execute, disk pages to fetch. If I want an acknowledgement I'll ask for one." See what I mean? Well, that wasn't the idea. They all wanted to do this delicate ballet of terminals talking to hosts, so we didn't pay much attention to that stuff, and we banged ahead and did XNS, and had already done the engineering to get the network to a reasonable state. Well, Metcalfe wanted to go off and do other things, just simply because—I mean, I had promoted him as high and done as much with him in Xerox as we reasonably could, and he just had the interest to go do other things. Reasonably enough, he wasn't sure at all what he was going to do; he just left and started consulting and waited for the phone to ring. IBM called him up and wanted to talk to him about what Ethernet was and different stuff like that.

Pelkey: But he was precluded from going off and really pushing Ethernet at that point, correct? He was under some confidentiality agreements and some non-compete agreements and so on.

Liddle: He wasn't under any non-compete agreements. Non-compete agreements don't mean anything anyway. There was nothing like that. It was simply the fact that he did not own the Ethernet at all, and all the details of it were, as we say, 'Xerox private data.' He didn't have any right to them, and certainly not to give them to our competitors or anybody else. Now, before he left, he had published some papers, and everything that was contained in those papers was—it was made very clear to him that that was not restricted, but otherwise he didn't have any specific rights or ownership in the Ethernet, or any rights to tell any of these other companies the details. I'm absolutely sure that he never did. He was very scrupulous about that. He would get a little annoyed with me once in a while, because he…

Pelkey: He felt something should be in the public domain and was covered by what already was out there.

Liddle: Well, or at least, he wanted to somehow stimulate things happening out there. The problem with that—and that was ok with me but—the trouble was, I didn't mind that happening, provided Xerox got

something out of it. I didn't want him to just go do it to make a big, blooming consulting business, and have Xerox not benefit from it. So he sort of wanted to be the marriage broker of getting Xerox to open up the product to lots of other customers, lots of other companies. My position was: "No, but a few selected companies that could make really strong partners so we can make people see this as a de facto standard, then that's fine." Of course, there were all these other guys applying to us for Ethernet licenses. Somebody wanted to license GE and all these other guys, I didn't care about them, but…

Pelkey: And Bob knew that you would be interested in working with other parties if, in fact, other parties came to Xerox to work with them?

Liddle: Absolutely, in fact he pushed that point of view very hard and he published his little newsletter which he sent to lot of people saying: "I have been encouraging Xerox to do X and Y and Z," and finally I either wrote him or called him up and said: "Bobby, come on, don't do that. Don't make it so that, if I go do something—it isn't this fucking easy to let me get the corporation to let me do something like this. Don't make it look like you dragged us kicking and screaming into it. You'll get enough credit for it. Don't do that." He told me in later years he thought of that was 'the threatening phone call.' It wasn't threatening; I just called him up and told him that it wasn't nice. It was making my life difficult. Anyway, then he was consulting to DEC.

Pelkey: Did you know it before you received the letter?

Liddle: Oh sure, absolutely.

Pelkey: You and he were friends?

Liddle: Yes, and he was back in Boston almost all the time, then. I hardly ever saw him out here. He was back there consulting to DEC, basically, among others, and I knew that he was working with them and, in fact, I knew that he was also being very careful about what he did and didn't tell them, everything like that, but he encouraged them to get ahold of us. I don't know, they wrote to the wrong guy, somebody at some strange spot in the hierarchy, but eventually, anyway—some communication planning guy someplace, or whatever—but anyway, eventually I got to go talk to Gordon Bell and just…

Pelkey: Do you recall a letter from Gordon to you?

Liddle: Yes.

Pelkey: In February, or so, of '79?

Liddle: Early '79, yeah.

Pelkey: I'm led to believe that Bob drafted that letter for Gordon's signature.

Liddle: Well, the thing that was funny—that could very well be, because he certainly was consulting to them and trying to tell them how to make this happen, and I knew that he was also doing that—there's things here that I don't know if I want to see them in a book.

Pelkey: I understand. Had Intel contacted you beforehand? Was there any contact between Intel and yourself at this point, between Intel and Xerox at this point? Intel was interested in…

Liddle: Let me just, off the record, tell you something for a minute here.

TAPE MACHINE SHUT OFF

Liddle: ... So, anyway, then we went and met with…

Pelkey: First, there was some bilateral meetings between yourself and DEC.

Liddle: Yes.

Pelkey: It was you and a planner, and on the other side it was Gordon and…

Liddle: And a cast of thousands. Jesus Christ, all these bushy-tailed guys, sort of a pain in the ass. Anyway, after that, I got rid of the planner and he settled on two or three guys who could work with me, one of whom you should probably talk to. It's Dave Rodgers, and he's, I think, engineering manager at Sequent now, or at least he was for a while. He's a very good young guy. I have hanging on my wall someplace, I guess it's at home, I have a printed message which says: "From Dave at DEC, Node Blue." This is from the 1980 NEC where we strung a gigantic cable across the McCormick Place, and the first DEC to Xerox Ethernet electronic mail message was sent by that pathway. Anyway, that guy Dave Rodgers, who sent it to me, was the guy who coordinated their end and pushed things along. He was a

good fellow. He worked for Bernie Lacroute, who now, of course, is the Chief Operating Officer at Sun, and anyway, we had those series of conversations with them. Around the same time, we were independently contacted by Intel. Intel was interested in making a networking part primarily for IBM. They wanted to make a single part that could do both SDLC and Ethernet, and they wanted it to be variable speed and able to run real fast and all this kind of stuff. They, as you say, were anxious that it should be at least ten megabits, maybe faster. Charlie Bass was sort of swirling around all of this; Ungermann-Bass had just sort of been, or he was at least thinking of starting Ungermann-Bass.

Pelkey: Right. He was poking around, looking for something to do after he left Zilog.

Liddle: That's right, and he kept saying…

Pelkey: He knew, because of the social interaction with Zilog, they were doing the Ariel network, and what you were doing, and when he left, he was looking for something to do, and started poking around in this area of networking.

Liddle: That's right, and in particular, he thought the standards thing was a great idea, that is 'let's get together and do something.' He sort of wanted in on it, or he wanted to know all about, and I wasn't in a position to tell him very much about it. Anyway, we went to—I went to DEC and said: "Intel's interested in working with us, too. What do you say we make a three-way arrangement," and—I'll tell you my reasoning for it. It was very simple. Doctor Wang and IBM and all these people were after us, because we had been running these very successful TV ads and were scaring the shit out of everybody, and the claims were getting made that: "Oh yeah, but Christ, this is really costly and it always will be." So I said: "If I want to get stuff that's *<inaudible>*, first of all I've got to show that I've got a semiconductor manufacturer who says: 'We'll make all the components so they'll be cheap,' and secondly I need down and dirty equipment vendor who makes modems and terminals and throw-away stuff, which was certainly Digital's reputation at that time." They were just swarming over the terminal business then, making very low-end products of all kinds.

Pelkey: DEC 100s.

Liddle: VT-100s, and all that stuff. So, I thought this will get away from the image that this is a premium high-end expensive product, and it'll encourage it—I wanted to start an industry at that moment. That was the only reason I did it. The whole Blue Book thing was intended to start an industry, not an industry for Xerox to make a lot of money in, a feeder industry to drop the whole overall cost of communications so low that it would just be a component of these systems. I didn't want any proprietary network. I wanted to

be able to say: "Look, I can hook to all this other gear that other people make," and I didn't want to have a post hoc networking industry develop. I always thought that was wrong, that there should be companies whose job was to go in and charge you a whole lot of money to hook up your equipment. I took the view that it should be just one accessory of the equipment, that of course it goes on a network. So anyway, I knew Xerox wasn't going to sit around and make modems for the world. The idea was 'how to make this stable enough as a standard that little guys would make barrel connectors and cable and…

Pelkey: Anything and everything.

Liddle: That was who I cared about. I didn't care about the creating industry of networking companies, because I thought that was sort of wrong. I just wanted to make it so that these things were just—well, Gordon Bell used a great phrase. In the first meeting I met with him he said…

TAPE SIDE ENDS

Liddle: ... You get these smaller companies committed to just making components and pieces and driving the price down, and if you can do that, if you can create those little ball bearings, then the industry can roll on those. So that was the sole reason that I wanted to do it. It was the sole reason for the Blue Book activity. It wasn't to dazzle the eyes of Fortune 500 companies by saying: "Ooh, now that it's blessed by Intel and DEC, that means Xerox has something wonderful." It had a little bit of that effect, but that was not it at all. It was to just open up the industry, and have an excuse to charge a very low licensing fee in exchange for people agreeing to do things in a standard way. See, instead of having the world make a million different, incompatible things, I wanted there to be one, so we said: "Look, for a thousand bucks, we'll give you a block of addresses, all the specs, and everything else that you need. This is an open thing. We'll share all the engineering data that we've got…"

Pelkey: Which was radically different than what Xerox had been doing before, in terms of everything closed in, and a closed environment…

Liddle: But Xerox was a company founded on proprietary technology. That's all Xerox was, and it made that company from a little jerkwater film company to one of the most profitable and successful companies in American industry. Quite understandably, they had a strong bent towards doing that. It's almost never the right thing to open up your technology to anybody else, almost never.

Pelkey: But that was a source of big battles internally, in terms of opening it up and having that view prevail.

Liddle: That isn't true. It wasn't bad at all, and I did it. Nobody else knows that because nobody else was there watching. It was a snap.

Pelkey: Because the way you positioned it wasn't perceived as being core technology, and it opened up the network so that they didn't have to worry about being in left field or right field. They were going to be…

Liddle: That's right. It made us mainstream, and the other guys were going to take on the hard work.

Pelkey: You couldn't take on everybody. You couldn't build all the big computers and all these other things that you needed to connect to, so therefore they did it.

Liddle: That's right, they liked it. In fact, Peter McCullough, who usually turns out to be so wise—when you really pick apart all the things he did, one of the things he said was: "Look, we've got a portfolio of patents, we should use them to acquire other technology and other commitments. Instead of sitting on all of them, we should use them creatively," and he loved this idea. He said: "God, what a great idea, trading the rights to this patent for improvements and all that, and their commitment to make parts, and then use it as a standard. What a great idea." So it wasn't a hard sell at all. I sold the damn thing internally in less than a month. It was harder to get Intel and DEC to stop squabbling with one another. They had had a bad falling out over the licensing of some semiconductor process—I forget what it was—and so there was a lot of—they would do the equivalent of sending guys to their meetings who were either attorneys or managers who had been involved in the other squabble, so sometimes it was a little tense. Then we wrote the press release, and Intel put some stuff in the press release about how successful they were in the microcomputer arena, so then DEC put something in about how they had shipped 9 zillion LSI-11s and all of that, and I said: "Guys, wait a minute. This isn't what this is about." So, it was a…

Pelkey: Now, Bob's perspective was that there were lawyers involved because of antitrust reasons, and that the issue of trying to get Xerox and Intel and DEC to be able to talk to each other. Not true?

Liddle: Not true. That's not a bad theory, but he was outside of it and I was inside of it, and there was no antitrust discussion about it at all; absolutely none. There might have been at DEC and Intel, but I never heard it. I never heard it once.

Pelkey: The lawyers being involved…

Liddle: They were fucking intellectual property lawyers. They were patent attorneys, that's all.

Pelkey: And the patent attorneys were there strictly because you had intellectual property rights, and how you transferred these and what you got in return for them.

Liddle: Right, and we made a relationship which was very simple. Number one, there was a three-way agreement for a long period of time—I forget how many years, but it was a long time—to exchange all research in local area networking, to mutually disclose and exchange.

Pelkey: And it now had to be negotiated during the course of these discussions.

Liddle: That's right, not that it was a terribly big deal, but it had to be negotiated. There was that. Then, second, all parties agreed that we would all do this provided that—we would exchange it three ways—provided the following three things were done: all three parties had to agree to implement this product as a standard for their future products, it didn't have to be their only network they used, but they all had to agree to implement it and support it; number two, Xerox had to agree—Xerox proposed—to license, for $1,000 each, all comers to practice the art in a standard way and give to them, reserve and give to them a block of addresses and so forth; Intel had to agree to manufacture an integrated circuit chip set to implement the protocols and control; Digital had to agree to manufacture, to design and manufacture, a low-cost transceiver.

Pelkey: Transceivers were a hairy…

Liddle: Yes, and they had volume for that; they did a lot of that stuff. They had a telecom plant up in Nashua, NH, and stuff. We didn't do that and neither did Intel. So their job was to make a transceiver, Intel's job was to make a chip, our job was to freely license any and all comers. That was it. That's all there was to it. The contract was a simple contract. It wasn't any big deal. Shit, the press release was harder than the contract, I'm not kidding you.

Pelkey: So the Blue Book was announced in the spring of '79, and then…

Liddle: We gave the Blue Book to the IEEE.

Pelkey: But you gave that in September of ‘80—was it September? You worked on it from—it was September of '80 when the Blue Book finally came out, September 30th of '80.

Liddle: Came out to where?

Pelkey: It was published. What was known as the Blue Book got published in September of '80. It got announced that you were going to start on it in the spring of '79.

Liddle: Right, that's correct, and then we worked together…

Pelkey: For a year to flesh the things out…

Liddle: Well, no—the details of it, right. For example, Intel raised questions about—we tended to use ground and plus three, or something, and they said: "No, it should be balanced line drivers," and we did a lot of screwing around with that. Digital wanted to have a heartbeat that checked the transceivers to make sure that they were working, as a reliability technique, so we spent time on that. There were questions about the address space. There were questions about things like that, and there was engineering data and there was a lot of—the original Blue Book spec, as you know, was ultra-conservative. All these things like 'CheaperNet,' that works, and in-the-wall, this shit I've got in here works and all that stuff, but at the time, it all had to be…

Pelkey: Mil spec.

Liddle: That's exactly right, so you know, 'big, yellow, Teflon cable' and all that stuff, but that was alright. We thought that the Xerox design was conservative and then those guys jacked it further up, but that was alright, because they wanted it to be reliable. So it was a very tight piece of engineering that we had already done and it took us a year to publish that. Then we gave it to the howling mob at the IEEE.

Pelkey: Were you involved with the howling mob?

Liddle: I tried to stay away from it as much as I could, because I didn't want to escalate—I was a Vice President of Xerox and everything, and I didn't want to escalate it all by…

Pelkey: It being a big, corporate issue.

Liddle: That's right, and everybody dragging their executives in, you know what I mean? But it was very messy and very small, chicken shit stuff. It was real bad. Mostly, three companies really stand out as major obstructionists:

Pelkey: IBM.

Liddle: HP, Data General, and IBM.

Pelkey: Now, HP perceives themselves as—this is through…

Liddle: Tell me who and I'll tell you how reasonable it is.

TAPE MACHINE SHUT OFF

Liddle: ... Loughry was great, ok. He was completely sensible, reasonable, a wonderful guy. He understood what was going on. Paul Ely, Jesus Christ, the most bullshit, anti-competitive stuff—competitive stuff—you could think of. He said: "Well, we think that Ethernet has dangerous grounding and it also has finite addressing." Well, it does, right, you can only label every three kilograms of the mass of the universe into one of its addresses. It was just 'we're not saying anything good about a competitor, goddamn it.' They wanted to screw it up, because we already had an installed base of products and Digital had already done a design. Data General was the same way; they just wanted to erect some trip wires, rather than create what we three were already doing as a standard. IBM threw up a lot of smoke, and all this Token Ring stuff. That's well documented elsewhere. Anyway, all I'm saying is those three, plus a faintly psychotic guy named Jerry Clancey who worked for Honeywell, were really the big problem. Of course, the way the problem got solved, it had nothing to do with it being reasoned out by reasonable men—and the IEEE wasn't at all. You probably know the story about how it got done. We had, by that time, begun working with Siemens. I was trying to OEM the Star, which I ultimately did, successfully, through Siemens, and there was a guy who worked for Siemens. He would come and visit me occasionally, and he also came and visited Metcalfe occasionally, and they were very anxious to see this Ethernet thing get blessed. They could see how they were going to hook it to their stuff in Europe. Standards committees are different in Europe. They're not parliaments, you know what I mean, especially in the good old Federal Republic of Germany.

Pelkey: Siemens, what do you want to do?

Liddle: Right, and what Fritz did…

Pelkey: Fritz?

Liddle: Roscheisen, and Fritz Roscheisen and Frau Ingrid Fraum, who worked for him, simply went to the European—what the hell is it called, whatever the hell it is, it wasn't CCITT, that's a tea-sipping society—but you know the one I mean?

Pelkey: ECMA [European Computer Manufacturers Association].

Liddle: ECMA, right. So they went to ECMA and they said: "We don't know why the Americaner—" Getting on with it, we think it should be the regular Ethernet, just like Xerox, DEC and Intel are using, so what they did was write a letter to the IEEE basically saying: "We have decided that we're going to ratify the DEC/Intel/Xerox Blue Book just as it is. We just wanted to let you know." Of course, IEEE said: "Wait a minute. We're not being upstaged by these guys," and so they turned around and said: "Bang, this is it, 802.3." The only reason that they ever knocked off the bullshit is that they simply got stampeded by the Europeans, and that was simply done by Siemens' pointing out what was going on to the other members of the committee, and they said: "We ought to kick those guys off <*inaudible*>."

Pelkey: That's the first time I've heard that story, and I've talked to a lot of people involved in that.

Liddle: Well, you can talk to Dr. Metcalfe about that. I think he knows about it, but the best thing to do would be to contact Fritz Friedrich Röscheisen or Ingrid Fro.

Pelkey: I think I have Ingrid Fromm's card. I went to an 802…

Liddle: Call up Bob—did you talk to Bob Prentice?

Pelkey: No.

Liddle: Oh, you need to talk to Bob Prentice—Bill Lynch and Bob Prentice at Xerox. By all means, you need to talk to both of those guys. Bob Prentice and Bill Lynch. Bob Prentice used to work for Bill Lynch, and Bob Prentice is a really good guy. He's a very calm, composed black guy. He was our torpedo on the 802 committee, and he worked for Bill Lynch, and they'll tell you that whole wonderful story.

Pelkey: So now, it went to 802, and at that point in time, your involvement, relative to the data communications side, you were just—you were involved in being in the front wave of it, but you went on to other things.

Liddle: Well, I was primarily, at that point, really worried about—by that time we were close to launching Star, so then, after that, it was all about workstations and user interface and getting all that stuff done right. It's also true that I considered the datalink problem to be solved, and just went up one level and said: "By the way, guys, here's our—" Then we published the XNS protocols and trampled all over—then we had to fight WangNet and all that kind of stuff, so then we began worrying about that sort of stuff. I gave a speech written by Bill Lynch—I wish I could remember now when it was given—in New York to fight a wave of WangNet propaganda. Boy, it was a real sizzler. We called a press conference and I basically told them all the things that were wrong with broadband, and particularly with the WangNet type approach, and all that sort of stuff, showed them how XNS worked and all that sort of stuff, but I don't remember the dates of that. Anyway, then it was all higher-level issues, so that was the end of the trail for me.

Pelkey: Yogen Dalal, going back to that period of time, he was really the person, in terms of XNS.

Liddle: Absolutely, he was XNS. Absolutely, Metcalfe was gone. Yogen did it.

Pelkey: I think I've had my questions answered. After I put all this together, I may have some additional questions, but thank you very much for your time. I greatly appreciate it.

Liddle: That's ok, it was a pleasure.

INTERVIEW ENDS

