Interview of Bruce Hunt

Interviewed by:
James L. Pelkey

Recorded: July 21, 1988
Apple Computers, Inc
Redwood City, California

CHM Reference number: X5671.2010

James L. Pelkey: Thank you for your time. I understand that the Ariel technology at Zilog [Zilog, Inc.] got there through you.

Bruce Hunt: Right.

Pelkey: Can you talk about where that came from before you got to Zilog? Where the ideas came from?

Hunt: Sure. The thing actually started because Joe Wells, who was then a director of the Stanford Center for Information Processing, wandered by my office one day and said, "Bruce, do you know how much it costs to connect a terminal to our computer system?" which was an IBM [International Business Machines, Inc.] 370 system at the time. I said, "No." He says, "About $3,000." I thought that was astonishing, since a terminal cost $2,000. I started to really think about how one might deal with that issue.

Pelkey: Around what year was this?

Hunt: As I recall, it was May of '77, and in the July issue of the ACM [Association for Computing Machinery], there appeared Bob Metcalfe's famous Ethernet paper, which I read, and went, "Wow, that's really amazing!" I started thinking about that, and just about that time, a fellow named Bernard Puteau, who did the architecture for the Z-8000 microprocessor, was working with a guy named Len Shustek, who went on to become a founder of Nestar as you may or may not know—Harry Saul. He'd be an interesting person for you to interview, too, later on.

Pelkey: His name is?

Hunt: Leonard Shustek. He was working on a PhD for Forrest Basket.

Pelkey: How do you spell Leonard's last name?

Hunt: S H U S T E K. He's now at Network General. He's a founder of Network General.

Pelkey: Leonard was at Stanford at this time with you?

Hunt: Right, he was a graduate student at Stanford University, and he had offices in the same building. He was essentially a graduate student working for Forrest Basket in the Computation Research Group at the Linear Accelerator. He had offices right around the way. We called him Lenny at the time. Lenny and Bernard Puteau were working on the instruction sets for microprocessors. In fact, it turned out that Bernard was working on the Z-8000 architecture, and he brought by a copy of the first Zilog SIO [serial input/output] spec sheet. I looked at that, and I said, "Wow, you put these things together and you could actually create a network." The thing that was interesting about it was that you could do it at very high speeds. The SIO was slated to run at 550 kilobits per second, which was astonishingly high speed for the time. I thought about it and said, "Wow, there's a way to do that," and I actually sat down and did the logic designs and thought through what the issues would be. The key issues then were: ”How could you get on and off the cable?” and the whole notion of how you could do appropriate collision detection; what kind of modulation could be done on the cable. These turned out to be major issues. A lot of thought went into that, and—

<break in audio>

Hunt: Basically, what it boiled down to was how we could use an appropriate technology to keep the line wiggling. I don't know if this makes sense, but the fundamental concept was that we wanted to go down a

shared cable. We wanted to be able to tap the cable and, if we were going to use some sort of a cheap connection, either through a capacitor or so on, we couldn't have a DC component on the line because the capacitor would stop it. A DC [direct current] component also wouldn't transfer through a transformer, so any way we wanted to connect and have isolation, so we had to think through that issue. It turned out that bit stuffing with an NRZI [non-return to zero, inverted] component turned out to be almost the right answer. It turned out that you need alternate mark encoding to complete it, so you could have a very, very low bandwidth, i.e. very close to the signaling rate is what you wanted to have the principal frequency components. That turned out to be the big issue—that, and how you do collision detection. I had sat and thought about that. Fundamentally, I guess, the issue was really raised because of what Joe Wells had said to me. I sat and spent a huge number of evenings just sitting down and thinking about that; thinking about how you could construct the hardware. I essentially came up with the basic scheme of how that could be done. About that time, Zilog was looking for people. I had spent about six years at Stanford and learned all there was that I thought was interesting about large mainframe computers. David Folger, who is an old buddy of mine, founder of Ridge Computer—I don't know if you knew about that, but anyway—called me up one day and said, "Why don't you come down and talk to us at Zilog?" So I did, and went to work for Dave at Zilog.

Pelkey: Dave was in what group?

Hunt: David worked for Joe Kennedy, and Joe Kennedy worked for Charley Bass. I went down and essentially brought two things to the table. One is an in-depth understanding of large-scale mainframe operating systems and how to build them and make them. I didn't know anything at all about the Z-80. It was an amazing jump, by the way. They set me up in a really nice office, and said, "Go to work. Help us design the next Z-8000 basic operating system stuff. I said, "Well, one of the things I'd like to look into is this idea of doing this communication thing."

Pelkey: Excuse me—when did you join Zilog?

Hunt: Seems to me it was—I have a picture of summer, and it was probably the summer of '77. In fact, May of '77.

Pelkey: May of '77 is when you thought Joe Wells approached you, and July of '77 was the Metcalfe issue.

Hunt: That was '76. July of '76 was when the paper was.

Pelkey: May of '77—

Hunt: I think it was about May of '77. Essentially, all of the concepts were done in the middle of the night, from May of '76 to July—I thought, "Gee, that's an interesting problem. Somebody ought to think about it." It was July that I really got the thing burning. Then, I worked on it right around through. Then, about May of '77 was when Zilog approached me and I decided to join them. I went down to Zilog and then literally built what was called ZNet at the time, over that summer with a guy named Ben Laws. I'm a software guy by trade and I'd taken a lot of hardware courses, but not done any real hardware design, and Ben Laws was helping. Ben Laws had come from Xerox PARC [Palo Alto Research Center] and he had joined Zilog, and was just a super guy to work with. I explained to him what I was up to, and he said, "Oh, yes, there are some great chips that we could use." We sat down and he did the hardware and looked at my simple designs and turned it into something that would work. Then we built the stuff, and then I did all the basic software that made it all work, and we had the thing up by the end of the summer. In fact, as I recall, I started spending about 60% of my time doing nothing but giving demos of how the system would work, because it really ran.

Pelkey: Now, was it called ZNet right from the beginning?

Hunt: No, it was called Ariel. In fact, the name of it was Ariel at the very beginning. It was actually named as a triple pun: it was named for—remember Ariel was the spirit or sprite in *Midsummer Night's Dream* [*The Tempest*], right supposed to be light and ethereal spirit, in the spirit of ether. My father's name was Ariel. It was a puny thing. Let's see, what else. Anyway, Judy [Estrin] was one of the early employees and was working for Joe Kennedy at the time. In fact, I think she and a guy named Craig something-or-other [Huffaker] were actually working on assemblers for the Z-8000 and Z-8—I guess [they] were the two new products that Zilog was about to introduce. The next interesting step was that Joe Kennedy asked me to talk to Charlie Bass. I sat down with Charlie and tried to convince him that LANs—local area networks—were a really good idea and that somebody ought to do something about it—so I guess he did.

Pelkey: Do you recall when that discussion might have happened?

Hunt: Oh, sometime during that summer. I don't remember the exact date.

Pelkey: Before you completed the Ariel network?

Hunt: Well, we had a working demo. He demo-ed it. Also Doug Broyles, who was the VP [Vice President] of System at the time, came by and I gave him demos. I gave demos to people from Exxon all the time.

Pelkey: What did Ariel do at that point?

Hunt: You could take two—I can't remember what they were, the ZDS's I think: Zilog Development Systems—and you could plug them onto a cable and they could talk to each other at 550 kilobits per second. You could actually type back and forth. We formed packets and shipped them out at half a megabit per second and captured them off again. We had complete collision detection, and we counted, and we knew when we were doing collisions. It really did work, and it was sufficiently cheap and fast to do that Zilog started to think about doing a product about it. It did all of the basic things: address recognition, all the stuff that you would have expected.

Pelkey: Did you have more than two development systems on it?

Hunt: Oh, yes, I think the most we ever got up was five. That was because we built all the cards by hand. They were all wire wraps. One of the wonderful things about Zilog, by the way, which I thought was just grand, was that—when I was up at Stanford, getting a technician to help you is really hard, but at Zilog, you could just take a circuit and two days later, somebody would come back with it wired. It was grand; it was just wonderful to work there.

Pelkey: Sounds like a phenomenal environment.

Hunt: Exactly right, exactly right. As a matter of fact, that was one of the things that was intellectually, extremely stimulating. There were just so many bright people there: Luis Trabb Pardo, founder of Imagen, was there working;Judy was there; Dave Folger was there; Ashuk[?] Suksino[?] was there; Joe Kennedy was there—

Pelkey: Suksino[?]?

Hunt: Ashuk[?] Suksino[?]. He was my compatriot in working on the operating system. He was really interested in operating systems, and ultimately I think he went off to work for Olivetti's Advanced Technology Group. I haven't seen him since he left Zilog, so I don't know what he's doing.

Pelkey: The people that founded XLAN were there too, right?

Hunt: Yes. In fact, there's a guy named Dale Laye[?]—I knew him just vaguely—who was Vice President of XLAN. I don't remember the other folks, but they were there, later on. In fact, just to complete the story—it was a really stimulating environment; there was all kinds of interesting things going on. I'm trying to remember the guy who did all the languages; PLZ, Tod Snook was the guy's name—really great guy. Anyway, they had done the PLZ language design and Charlie and Todd had worked really hard and did a really beautiful job and it turned out to be a nice development language. There was that whole group that was going on. Anyway, just lots of good intellectual discussions. The question was, how to get the communications going? About that time, Judy and her roommate had a party at their house. I can't remember her roommate's name; it's terrible. They had a party at their house, and one of the people that she had invited was Fouad Tobagi, and Fouad had just accepted the job at Stanford as assistant professor of electrical engineering. Judy introduced us and said, "Fouad, you ought to know Bruce. Bruce has built this little network." Fouad says, "Wow, I studied networks. That's what I did for my PhD for— " Oh, come on, the guy's famous for queuing theory—

Pelkey: Leonard Kleinrock.

Hunt: Yes, Leonard Kleinrock, exactly. Anyway, so Fouad says, "Yes, you ought to take a look at some of my papers."

Pelkey: How do you spell Fouad's name?

Hunt: F O U A D T O B A G I. Fouad said, "You ought to go look at my papers," so I did. I went and read all of his papers and his PhD thesis, and what he had done, which it turns out it was something I learned later—Bob Metcalfe, in his PhD thesis, had proposed what was called the Linear Feedback Model, which was a way to analyze networks, but Bob didn't have the basic mathematical sophistication to carry it through. Fouad had a wonderful background in that area, and he carried it through for that carrier sense multiple access scheme, and did all of his work. In fact, he wrote four really famous papers. I don't know if you recall, but he and Leonard Kleinrock, of course, wrote these four very famous papers about how to treat CSMA [carrier sense multiple access]. When I read his papers, I said, "Gee, I can extend this to the collision detection case," so I did the derivations and added the extension of being able to handle collision detection, and I took it to Fouad. Literally, I think I did it in about 12 days. After you read this stuff, your head is full, so you've got to dump it. I did the derivations and I took it to Fouad—and it was about 12 pages—and I said, "Fouad, look at this." He says, "This is great, Bruce. Let's write a paper." I said, "Ok, great." We started talking about itI told him I was interested and that I had a bachelor's degree in mathematics and he said, "Why don't you think about coming back to Stanford and working as a graduate student?" So I did.

Pelkey: When did you leave Zilog.

Hunt: I left Zilog; I guess it was October or November, and I was—

Pelkey: Of '77.

Hunt: Yes, Zilog was really supportive of—by the way, one of the neat things about Zilog is that it was very supportive. They said, "Yes, you can go and take classes at Stanford." I guess there's another really important point, and that is that I was, of course, taking classes at Stanford while I was working up at the Linear Accelerator, because that was really easy to do. I like that sort of thing. One of the classes I was taking was Bob Metcalfe's. [He] was giving a seminar on local area networks. In fact, that summer, during one of the seminar periods, I gave a talk on Ariel. We actually brought it in and demonstrated it and showed to everybody.

Pelkey: When was this?

Hunt: I suppose it was in the fall—September? It was in the early part of the fall—September or October of that year.

Pelkey: Of '77?

Hunt: Yes.

Pelkey: When you left Zilog, you went to SLAC [Stanford Linear Accelerator Center]?

Hunt: No, I actually went and went to Stanford University as a graduate student.

Pelkey: But you said you were working at SLAC, or a conference was held at SLAC. Did you say "SLAC?"

Hunt: I probably confused you, because I'm jumping around. I worked at SLAC from about '72 until—well, it was five years and six months, or five years and nine months, I can't remember exactly, but it was '72. September of 1972, because I got married on December the 16th of that year. [I worked there] for five years and nine months, and then I left in May and went to Zilog and stayed until about October or November, and decided it would be much more fun to be a graduate student and work on the mathematics, because I really liked the mathematical side as well as the technical side. Then, I went back and was just a straight graduate student, and I was Fouad's graduate student, and went from an extraordinarily nice income at Zilog to nothing. I was having a great time, but Deb didn't think it was such a hot idea, as I recall. As a matter of fact, I tried to convince her to move into student housing, and it just wouldn't fly.

Pelkey: When was Metcalfe's course?

Hunt: I believe it was that fall, so it should have been September—

Pelkey: Before you left Zilog?

Hunt: Right, September, October, November time frame, I believe.

Pelkey: It was a course that he was teaching?

Hunt: Yes, it was a seminar in local computer networks. We can probably get the times; I don't remember exactly the date. I probably should go back and really try to set these things down if somebody really cares.

Pelkey: I care, so when you get this transcription, maybe you could—you were still at Zilog?

Hunt: Right, and in fact what I did was I took a leave of absence to see what I wanted to do. I didn't really leave, per se. Anyway, it was fun to be a graduate student. I really liked that; I worked well with Fouad, and we worked until—gosh, it seems to me it was like about June of the next year—and it was that May, of '78, that we went off to the MITRE conference. I believe it was in May. You can check that, too.

Pelkey: The first MITRE conference.

Hunt: Right, the first MITRE conference in local area networks.

Pelkey: Was it here or back at MITRE?

Hunt: It was actually at the—

Pelkey: There were three of them that year. They had three regional shows, as I understand.

Hunt: This one was in Boston, and it was called an LACN [Local Area Communications Network] symposium. I've still got a cup. They gave me a cup because they said we did the best paper.

Hunt: This isn't the one at the Copley Plaza?

Hunt: Yes, it was at the Copley Plaza.

Pelkey: That's '79.

Hunt: Was it '79? I'm off a year, somehow. That was the only one that I attended, was at the Copley Plaza.

Pelkey: It was '79 at the Copley Plaza.

Hunt: What did I do for a year? I don't know. I apologize. Oh, you're right, it was at the Copley Plaza. As a matter of fact—

Pelkey: <*leafing through some papers*> It's '79, it's May '79.

Hunt: I'm pretty sure it was May.

Pelkey: May of '79.

Hunt: We published the paper in Computer Networks. It turns out it's been a classic. They actually republished it. IEEE [Institute of Electrical and Electronics Engineers] asked for permission to publish it again. It's been published a couple of times.

Pelkey: This paper was originally published in—

Hunt: The basic results of that were first presented at that Copley Plaza LACN. I can't remember if it was the first or second.

Pelkey: Now was CSMA/CD [carrier sense multiple access/collision detection] a well known expression at that time?

Hunt: No, this was the paper that introduced the terminology. CSMA had been introduced by Fouad earlier in his PhD thesis and the papers that appeared in the Transactions on Communication.

Pelkey: Did the Ethernet that was being done at Xerox at this point have these concepts embodies in it and they just didn't call it this?

Hunt: Yes, this paper brought the first complete analysis of how all such systems worked. The purpose of this paper was to tell the world how collision detection works, and all of its facets. We looked at all the algorithms; in fact, we introduced what we called the ”One Persistent CSMA/CD” which is, in fact, the Ethernet model. There are some minor differences—the ”binary exponential back-off” was different than the ”one persistent” back-off scheme. Never the less, it has the same basic operation, and the ”non-persistent” CSMA/CD, which is a different style. Those were the two. Of course, you could do “P-persistent” which is something that we didn't bother to explore. In fact, Fouad had done it without collision detection, and it seemed like a fairly straight-forward arena. What this paper really brought to the party was the theory that allows you to take a look at collision detection as a scheme and determine how the system will perform. It brought in the basic parameters, which is number of stations, throughput, analysis and delay, and combined all of those things. Prior to this paper, there were only throughput analysis—essentially what's called the ”Infinite Population Model”—and there were no delay analysis of collision detection type schemes by that time. That's really what we did. It was an extension of the linear feedback model, to make that happen. The thing that was cute, of course, is that Bob Metcalfe had said, "That looks like the right model to use." Fouad had said, "Yes, it *is* the right model. Here's the way you can use it." I just added collision detection. "Here's the way you can use collision detection, or extend this so that you can now handle collision detection," which is kind of fun. It was very interesting. It's turned out to have been a far more important paper than I would have ever thought. It just keeps showing up. I guess thing that is fairly amusing to me is that it keeps showing up over and over again; people keep referencing it over and over again. It's been pretty much classic. Finally, last year or two years ago, IEEE came back and said they wanted to publish it in Classics of Communications. It's now in a book which is extremely expensive, $45.

Pelkey: How did the people at MITRE know about the work you were doing?

Hunt: Principally through the conference, when we stood up and gave the presentation.

Pelkey: At Metcalfe's conference? How did MITRE know about you to invite you to the conference in the first place?

Hunt: Probably because of Fouad. Fouad knew about the conference, and said, "This is the right place for us to go present." Of course, I was very interested in LANs [Local Area Networks]. As I recall, it was published in something like the ACM or someplace like that, so it seemed like it was the right, natural place to do it. Fouad might actually have a real good memory of exactly how we got in, since he took care of all the arrangements.

Pelkey: Do you remember anything about actually being at the conference?

Hunt: The most important thing that I remember about the conference that was fairly interesting, at least to me, was that I asked Bob Metcalfe, "What do you think?" He said, "I think we probably went overboard on binary exponential back-off," because one of the results of this paper was we tried to look at what the right back-off was, and it turns out linear back-off is the right one to optimize the delay. The binary exponential back-off guarantees stability, no problem, but it does that at the expense of introducing very long delays, and if you want to optimize—you still want the channel to be stable, i.e. not collapse under a heavy load, which is a problem with any collision detection scheme—then you need to monitor that and counter effectively, and you need to figure out what the right thing is to do. The answer is that linear back-off is the right thing, not exponential back-off. It introduces too much delay too fast. In fact, actually, there are some interesting experimental results. A guy named Guy Elms up at the University of Washington did a simulation study which showed the result that we had derived mathematically, which is kind of fun.

Pelkey: How about that. How many people do you remember as having been there?

Hunt: At the conference? It seems to me there were like 150, 200 folks.

Hunt: Was there a lot of energy there, a lot of excitement?

Hunt: I remember just being really thrilled. It was the first time I had any been at a professional conference of any magnitude, so I was really excited about that. I had a grand time. I had spent very little time in hotels in my life prior to that I was in ”cultural overload” is what you might call it, so I had a great time. I was worried about making sure that I didn't stick my fork in my foot or something like that and embarrass myself. I remember just being amazed at how many people were really interested in local area networks. I don't know any researcher who secretly believes that what he's doing is really important—

Pelkey: Or that a lot of people are interested in it?

Hunt: Right, a lot of people are interested in it, I was just surprised that people really were interested in it.

Pelkey: Do you remember any discussion of token ring versus Ethernet at that conference?

Hunt: Very little about token ring, as a matter of fact. I had a chance to meet Dave Farber later; the guy who really invented token ring. I don't know if you've talked to Dave; dynamite guy.

Pelkey: I interviewed him in Newark, DE.

Hunt: Right, he's at the University of Delaware, as I recall.

Pelkey: No, University of Pennsylvania.

Hunt: He was at University of Delaware. Great guy. I was really excited to meet him. I remember that as a lot of fun, to meet some of these people who had really done it. There was a guy named Tony who was at Queens College in London—I can't remember his last name—who I had actually done some correspondence with, because he was interested in high speed LANs. We talked about it—sort of communicated—

Pelkey: When you finished the conference, what happened?

Hunt: I went back home, and about that time I had discovered that I couldn't live in the style to which I have become accustomed. I had exhausted my savings after having been a couple of quarters at Stanford. I decided that I really liked working with Fouad, but I had to make some money. Fouad suggested that I go work at SRI [Stanford Research Institute], so I did. I went over and visited Don Neilson, and Don said, "Gee, would you like to come on?" I said, "Sure." Part of the deal was that I got to continue my work at Stanford, and that turned out to be a lot of fun. I spent the next—gosh, '79 to '82, I guess, if we can find the lost year—from '79 to '82, I essentially was at SRI, where I wrote the rest of the papers; Unidirectional Broadcast System with Ralph Urone[?]. Nachoum Shacham and I sat down and tried to study the “one-persistent” CSMA/CD model and try to understand how it worked. We came up with a new derivation that gave us a real good insight. From that we were able to predict a thing called “channel capture.” It turned out, there's some real physical evidence that it really occurred. Ethernet-like channels can exhibit a behavior called “channel capture,” where one station will grad the channel and hang on to it to the exclusion of all others. It turns out it's a completely statistical property of the way it works, and you can adjust it. We showed how it worked, and that was a really nice paper. In fact *<handing over a paper>* this is it. It's called "Performance Evaluation of a CSMA/CD One-Persistent." You're more than welcome to take that with you. I have a couple of them. Anyway, we did that paper, and Ralphy and I were looking at other schemes. Ralphy had come up with this ”unidirectional broadcast

bus" idea. I sat down and said, "Wow, I've got all these great mathematical tools." I did a fast analysis of it, and it turned out to be a really good idea. We did this analysis together of what we called the "Unidirectional Broadcast Scheme." That got me to thinking about why it had to be unidirectional? That is where the Metapath ideas came from. This was a real early example. I actually never published anything about the Metapath stuff. I got really interested in making it real. One of the characteristics that I have is that I like the practice as much as the theory. I like the combo. I got to thinking about that and came up with a way in which we could extend the unidirectional broadcast bus so that it would work without having to have unidirectionality that you could still get it. It turns out that the fact that the speed of light is constant in the universe turns out to me that there's always a way to order events down a single cable. By definition, when you start, it's going to be—an event will occur earlier and to someone else later, and so it does this automatic ordering, and that's exactly what we took advantage of in the Metapath scheme. Then, in order to make it cheap, we had to leave a little slot so that you could detect when it was your turn, but it turned out to work just beautifully in practice. You could get really high performance. The thing that was interesting is that scheme performed like an MD1Q, which is wonderful. Literally M exponential arrivals can come to a deterministic queue, which is your single channel. It's the best theoretically that you can do. If you can't predict the arrival rate, and you can deterministically handle it. It turns out it works exactly like that in practice, and that's what was very neat. There was a guy named Michael Sholl who went and derived a whole series of things that show that all those schemes worked, and actually published papers. He was one of Leonard Kleinrock's students. The Metapath technology was based on exactly that. We went and got a patent on that idea, and I decided I wanted to try to start a company in Silicon Valley. A guy named Tom Lunzer and I used to meet in my garage, or my den, every night.

Pelkey: Starting when?

Hunt: Gosh, I guess it was probably '80, '81. It was about 1980. The company was founded in August of '82. That's when we did Metapath, when we actually sat down with—

Pelkey: It was still Ariel Systems at that point?

Hunt: No, I had just talked to my brother and let him fund us. I went and said, "Hey, Mark, I want to build a chip." About that time, I was very interested in VLSI [very-large-scale integration] because of what I was doing at SRI at the time.

Pelkey: Did you get your first funding from—

Hunt: Norm Dion.

Pelkey: Norm Dion.

Hunt: March of '83, I believe.

Pelkey: Did you call it Metapath then?

Hunt: No, it was Hunt Forth. Actually, it was August of '82 when we incorporated the company. Craig Johnson at Wilson Sonsini—what had happened was, the guy that lived right across the street from me was a guy named Jim Pardee. Jim was Chief Financial Officer of Applied Biosystems at the time. We got to talking one day and he says, "What are you doing? I keep seeing all these boxes go into your garage." I said, "I had this idea to start a company." He had an MBA [Masters in Business Administration] from Harvard and was interested. I asked him if he'd like to come take a look and help us get going, and so he did. Jim helped us a lot with the finances and financial aspects of the company. It was cute. We would talk about it, and we started visiting venture firms. Literally, we decided about the August time frame that it was time to get really serious. We had, by this time, produced 20 Version I Robins. We had our custom

integrated circuit complete. Mark, my brother, had paid for it, and out of what Ariel Systems had done, he loaned me—or actually gave me, literally—the $20,000 it cost to build the thing. We bread-boarded the sucker up—discrete breadboard. It was really amazing; a real mess.

Pelkey: There's got to be a better way.

Hunt: Absolutely, and there now is. It's just amazing. Anyway, we bread-boarded up the chip and built that. Then we put together the 20 Robins and we demo-ed it to—in fact, the exact way that occurred was that we had gone and talked to a lot of people. I'll never forget this; when we talked to Craig Taylor at—oh dear. It was actually a venture fund, and he was really interested. Then, he asked for the—oh, I'll remember who the people were, but—we talked with them for a long time. It took me a long, long time, because I was extremely naive, to figure out that when those guys were being very polite and asking us questions, they were trying to say no. They didn't know how to say no. I wish they would have just said, "Hey, Bruce, you don't understand how to run a business yet. Here's what you need to do." We just persevered, and what occurred was that my brother had a friend who was working for one of Norm Dion's start-ups, and they were interested in data communications. We had our little presentation for Ariel, and we zipped on down. We called it Ariel at the time, too—one name, use it over and over again. We went down to the company—I can't remember the name of the company, which is terrible. They had an interesting technology. They had seven floppy disks that they were going to stack together.

Pelkey: <*inaudible*>

Hunt: Yes, probably, if that was what it was. It was something like that; sounds right. They had these seven floppy disks, and they were in Saratoga over a garage. We went down there and we gave a presentation of the technology. One of the people sitting in the audience was Bill Harry, who was Vice President of Development at Dysan. Bill Harry asked us a bunch of very probing questions, and says, "Who's funding you guys?" I said, "Nobody. Got some help from my brother, and we're in my den right now." He says, "I'd like you to meet somebody. I'd like you to meet a guy named Norm Dion." I said, "Great." That night, I went home and I called Craig Johnson up, and said, "Have you ever heard of a man named Norm Dion?" He says, "Where did you hear about him? What do you know about him?" I said, "We're going to have a meeting with him Monday," and Craig said, "Bruce, you're in the big time. Don't blow it." That's exactly what he said. He said, "You guys need to be really prepared." He explained to me who Norm Dion was. He's the founder of Dysan and chairman and CEO [Chief Executive Officer], so if you're going to see this guy, you better be prepared. We worked our asses off. We went into overdrive. Literally, I think we slept two hours until that Monday. We did a complete presentation of the business plan. We had foils and everything all set up, all the papers I had ever written. Tom Lunzer and I worked like dogs. He had what he called a standing meeting at, I guess, one o'clock in the afternoon. He had a Monday afternoon staff meeting, so we walk in there. We went to the Wedge. I don't know if you've ever been down to the Wedge, but boy, what a beautiful building. I looked around; gosh. His administrative assistant came down to get us—a really neat lady. She came down and was very pleasant. She took us up to this room and walked us back in there. It was this giant room, and there were all these little place cards for all these people—George Ferinski[?], Chief Financial Officer; Dennis Murray, Vice President of Engineering; Bill Harry, Vice President of Development—just down the row, bottled brains. We sat there, and Norm comes in. The meeting, as I recall, was supposed to start at one o'clock, and Norm wanders in about 1:45. He sits down and everybody immediately straightens their—it was very impressive. I have this really great picture. He said, "Hello," introduced himself, and we gave him our presentation. He said, "Thank you very much, that was very interesting. We'll talk to you later," so we walked away. About Thursday, Bill Harry called and said, "Could you guys come back in next Monday?" We said, "Sure." Then on Friday, he called back and cancelled. There was two weeks of horrible "Oh my God, they're not interested in us.". About two weeks later, they finally asked us back, and we came in on Monday; same room, but this time all of my papers were open and on the desk with our business plan. Somebody had actually read it. You could see notations by the side and yellow highlighter. Somebody had really looked at this stuff, so I figured that means they're at least interested. That was a great sign. We sat there. Norm started to really probe me. Boy, probably the toughest interview I ever went through in my life, was Norm asking me questions. I was in such a state of excitement that my brain was racing, absolutely

crazy. I noticed that he would ask me an easy question, and follow up with a really hard one. It was like he would disarm you with this one, and then just hit you with a whammy.

Pelkey: His pecking theory.

Hunt: Exactly, so pretty soon I got into the style of it. ”This is the easy one, so take it easy, and then really focus on the next one, because that's where it's going to be.” Finally, he stopped the meeting. I think it was about 4:00—as I recall it was in December—and he just stopped the meeting, and he says, "Well, thank you, I think I've had enough." Then he sat there, and he literally put his head down and he started to think. I'm sitting there and I'm saying to myself, "My God, should I say anything? Jiminy Christmas, there's this whole row of people." I had this impression of all this incredible amount of valuable time was going to waste. We must not waste time, but you think I'm going to say anything now? No way! It got to the point where we were sitting there and fifteen minutes goes by; 20 minutes goes by; 30 minutes goes by, and I'm going, "Oh no." Finally, at about 35 minutes, I made peace with myself. I said, "Look, he's got a tough problem, and if I jump in now, he's not going to—he's got to come to his own personal conclusion. If it's going to take all night, I'm going to sit here and wait." I just sat there. At 45 minutes, he looks up and says, "I think I'm going to do the deal," and that was it. From then on, it was <*claps*>—I don't know if you know—one of the things I really like most about the man is that if he gave you his word, you could bank it. He told you what he was going to do, and he would do it. You'd never have to worry about it. He said, "Let's get going. Can you pull out a couple of those letters of intent that we did before?" He had several of them. One of which was the one he did with Al Shuggart to start Seagate. What a deal! Really exciting stuff. Anyway, Motornetics—the company that had the technology where you could get just unbelievable control over analog motors, some company up in Marin. I don't know what ever happened to them—and the company that was doing the 3 inch disk drives, <*inaudible*>., I had all the letters of intent that I had signed with the founders. He says, "Now you go off and you figure out—." We talked about what it was going to take to do it. I had my little budget, and he came back and he says, "No, I want to see proof of concept, so I want you to estimate what it will take to get you to this stage, and that's what we'll do." We said, "Fine." We went off, and that's where the $640,000 came from. It was exactly like that. He looked at it and said, "Yes, that looks great." Then he said, "Great." Three days later, a $190,000 check shows up to me from Norm Dion. That's the way he started the company.

END OF TAPE 1

START OF TAPE 2

Pelkey: We've come through '83. 3Com is going and Bridge is going and Ungermann-Bass is going and Sci-Tek is going. There's an awful lot of activity that has come out of this intellectual heritage of computer networks, which became local area networks. I haven't pinned down when it went from local computer to local area.

Hunt: It's probably because you have to say L C N [Local Computer Network], and you can say ”LAN” for L A N. I suspect it's something as simple and trivial as that.

Pelkey: Probably so. You were at SRI, where Engelbart had been, and there were lots of ideas floating around. Can you characterize this time period, from your perspective?

Hunt: Fundamentally, there was a perception that we were on the edge of something really astonishing. In fact, I've said to people for probably the last ten years that this is probably the most exciting time to be alive or any time I can imagine, for somebody who is interested in technology. The reason for it is that the emergence of integrated circuits and the ability to pack incredible amounts of logic on a machine, the fact that those machines can manage and control themselves, the storage technology is moving up, the communications issues are arising, that graphics is coming in. You could just smell it. You could sense

in the air that major, major things were happening. The Z-80 was going like crazy at Zilog. There was all that ferment—what Intel was doing, what Motorola was doing—really exciting times. For example, when I was a senior in high school I went down and toured a couple of electronics firms, and they gave us a handful of field- effect transistors, the very first MOS FETs [metal-oxide semiconductor field-effect transistor], which had the incredibly high input impedance, and changed—we were transitioning out of the vacuum tube era into the transistor era with the field-effect transistors. They acted like tubes but had these incredible input impedances and all these great characteristics. You had this perception of being at the right place at the right time. The other thing for those of us that had a reasonably good background in computer science and mathematics, was the sense that there was a true embodiment of essentially much of these abstract ideas. That was really exciting. Just *beaucoup* exciting. The other thing was a perception that, if you had a really good idea, you could turn it into a business, which is very unusual, as I have come to understand. Very few people think of it in those terms; "How could we turn this into a product that somebody else would use, and make money off the deal?"

Pelkey: Where did that stimulation come from? You had this issue about theory and practice, but was it also the climate you were in, the culture which you were in, which was creating companies?

Hunt: That's exactly right. As a matter of fact—

Pelkey: If you had been in Chicago, you might still have had this interest in the theory and the practice, but it might have surfaced in a more academic environment.

Hunt: Exactly. I was at Stanford and SRI, lots of academia. If you look at my whole life, my dad loved school, and lots of academic stuff, so it was real natural for me. If I really think about it, it was the fact that I grew up in the area, went to high school in the area and watched companies form. Intel come out of nowhere, literally, and start to dominate. Zilog [was] even more impressive. They start the company, they do the Z-80 and it works, literally, on the first pass. Shima did just an unbelievable tour de force. Federico Shima and I guess Bernard. I'm not sure if he was on that early. Anyway, those people just "BAM!" and they had it up working. They were shipping, and it was an order of magnitude better than—for the time—the Z-80 was very much better part than the 8080. First of all, single power supply, completely identical, a much better instruction set for the time. There was that sort of thing. Then the other thing that I think happened, was that a number of people in Zilog went away and formed their own companies. I watched that happen. A guy named Bob Marsh went off to form Onyx Systems. I remember the day Charlie Bass fired him. I don't remember what the deal was, but Bob got fired and went off and founded a company to do Z-8000 boards and Unix systems. Then, Doug Broyles left to do IMI [Intermolecular, Inc.], the disk manufacturing start-up. I watched that sort of thing happen. It seemed a very real thing to go start a company. I had absolutely zero perception of how hard it is, what a lot of work it is, but anybody could do it, and it seemed to be that way. Furthermore, there was a perception that if you had a good enough idea, that you could find someone who would fund you until you could fund yourself, which was a really important characteristic. There was enough venture funding around, even back in those days, that you could really seriously think to do that. It turned out that we were fortunate. There was the economic basis—i.e., there was a technology that was there, there was somebody interested in funding the technology, and there was a real demand. You put those things together and bingo, it happens. It's just the right set of ingredients. Boil, stir well and watch what pops out. In fact, I would argue that it's still happening today. I can see more opportunities. You could start a company a week, literally, and there're really great ideas that probably could make money.

Pelkey: That party that you went to at Judith's house, how many people were there? I understand that was something that happened periodically.

Hunt: Oh, fairly regularly.

Pelkey: Did you go to a number of them?

Hunt: Well, that one really stands out in my mind. I think I was at about three or four of them.

Pelkey: How many people would be at these?

Hunt: Twenty-five to 30. They were the Zilog-ers; the Zilog crew.

Pelkey: Xerox PARC people would come as well?

Hunt: Yes. Judy's roommate was a Xerox PARC- er, and there was a lot of bi-play. Ben Laws had come from Xerox PARC. Oh gosh, I wish I could remember—the lady who went on to do a lot of stuff at Grid, Carol Hankins, was there a lot. She had come to Zilog and then gone off to Grid. Dave Folger was there. Dave was real interesting.

Pelkey: Those must have been exciting parties, in terms of the fact that there was an incredible amount of energy going on at both places at this point in time.

Hunt: Absolutely, and people were talking about—gee, I wish I could remember—the guy who did the Bravo text editor had come from Xerox to Zilog, and was building some really nice text editing facilities. On almost all levels, you could look at, there were real hard cored engineer types, logic designers.

Pelkey: Were the conversations technical product conversations or—?

Hunt: Much more idea oriented. 'How can we build it?' For example, most of the conversations at Zilog were based around the idea of doing the Summit System, which was the operating system. There were lots of discussions. The basic concept behind Summit was that you had a series of cooperating Z-8000 processors who could talk to each other and literally share processing. Communications was a major part; "How can we get this part to communicate?" I had a solution to that part of the problem, of course. We talked about how we would do that, how you could attach and detach yourself from the network, what would you do about security, and how could you handle what I called the shared memory problems. We had a big debate. II remember a debate about how you could share memory across a local area network. We talked about how we could make the network run ten times faster. We didn't want it to be at 500 kilobits or anything like that. We wanted it to be at 5 megabits, right. That would have been just absolutely grand, and it appeared that it was technologically reasonably do-able. Those issues went on; how we would do that. The thing that was interesting was that it was almost a given that we could do reliable communications. The real issues were issues that still plague a lot of people. That is this "name management." How do we manage the name space of the system? How do we manage resource allocation? We have to have this cooperating set of processes, any one of which can die at any moment. How do you guarantee reliability? The other issue that turned out to occupy a huge amount of time was the issue of; suppose you wanted to disconnect your system from the network, take it home, work on it, and bring it back and reconnect it. How do we know it's really you? How do we know that you're trusted to get back into the files that you have? A lot of those issues still are not well understood today. Dave Folger, Ashuk and I had huge conversations that lasted through parties—lots of conversations in that arena; how we could better improve the basic algorithms for the network. In fact, Dave invented a couple of really cute algorithms which nobody has really taken advantage of to this day.

Pelkey: What an exciting period of time.

Hunt: What can I tell you about it, other than just that that's the way it was? We just did it.

Pelkey: I think you've filled me in on the things that I was looking for. You've been extremely helpful. Thank you.

END OF INTERVIEW

