Interview of Daniel J. Warmenhoven

Interviewed by:
James Pelkey

Recorded March 9, 1988
Atherton, CA

CHM Reference number: X5671.2010

James Pelkey: As I mentioned, Dan, I will make a transcription of this and get that to you, and anything I use from this conversation I will also pass by you. I appreciate your time, and thanks for coming over.

Dan Warmenhoven: You're welcome. You asked for my background, well, let's see, let's start when I was at Princeton. I was an EE, and started getting involved in computer communications at Princeton, which is not really known as an engineering school. They graduated about 30 EEs in my year, and this is -- one of them, one of my classmates was Dave Boggs, the co-inventor of Ethernet, so it's kind of ironic that two guys out of a small class -- really two principals of building the competing technologies, but even before I got out of Princeton I started working with IBM, joined IBM four days after graduation at the Yorktown Heights facility, which was then the Advanced Systems Development Division, the ASDD.

Pelkey: What year was this?

Warmenhoven: 1972. Spent only one year there, and moved to Poughkeepsie, where I worked on all distributed systems and communications systems, all heavily communications based -- a lot of minicomputer orientation and so on, Series I, 3790, 8100 kind of stuff, which are all communications-based systems. I moved to Raleigh in '77, which was then the Communications Products Division. In '79, I became the development manager, working for Murray Bolt at the time -- actually, I guess I started in '80 -- I'm sorry, let me get the dates right here. In early '77, I went to work in the distributed systems organization and did the 8100 RJE package. Then, all that responsibility shifted to Kingston, which was an interesting problem for the Raleigh laboratory, the Raleigh organization. They even hired more people to move to Kingston. At the time, the 3705 multiplexer program, the mainframe interface, all had moved to (unintelligible), so there was no mission, if you will, for Raleigh. So the corporation decided to charter Raleigh to build communications products.

Pelkey: And this was '79?

Warmenhoven: This was '79. It may have spilled over a little bit to early '80 before all the decisions were done, but, anyway, in '79, late '79, the decision was made to -- did I get that right, the dates right?

Pelkey: '80 was IEEE.

Warmenhoven: Yeah, yeah, that's right, so it was early '79. That's all correct. By the end of '79, the corporation decided that Raleigh would focus on, basically, system independent communications products, meaning not 3705 multiplexers and the rest, and they listed various business alternatives for doing modems and various other things, but, in fact, at that time, an organization was put together called Network Products, which I became part of -- Murray was actually part of Network Products, the development manager -- and we build a product that ultimately turned out to be the 3710, 3708, and I was the engineering manager on that. The 3710 was basically an alternative to a statistical multiplexer. It was a combination protocol converter, line concentrator, and SNA spoofer. In fact, it's still the same architecture that Netlink uses today, Netlink base comes into our competitor. The 3710 could handle asynch and bisynch, and

the notion was to provide both the line concentration capability for economic justification, as well as a migration vehicle for RJE stations and 3270, sell into the world of SNA -- and never sold. The IBM sales force could never figure out how to sell it. The IBM problem with communications was twofold. The first one was lack of expertise in the field about communications in general. They knew SNA, but it was all mainframe based. They didn't really understand basic communications. And then the other piece was the, if you would, personal productivity issues built into the compensation plan. Here was a box that cost $25,000 or $30,000 that became difficult to sell because it's got so many technical aspects to it, and so on, and so it became time intensive for a salesperson, and you get very little return, so there's no commission structure. Everything was flat commission basis, so you couldn't get their attention, and that becomes a recurring problem, I think, at IBM.

Pelkey: Wait until you read the story about Codex introducing the statmux. They basically fired his sales force because they wouldn't sell it.

Warmenhoven: I'm not surprised. IBM wouldn't take the next step of really putting together a focused sales force, so that entry into that market, it was moderately profitable, but not something that was a real winner. But, anyway, I took that one to the point of introduction, and around that time I was asked to put together a product proposal for (ligible).

Pelkey: This would be late '79 then?

Warmenhoven: I actually believe, now, this was '80. It had to be in the fall of '80. There was a task force commissioned, and, I remember, when I inquired about the objectives of the task force, I was told that -- this is an interesting little story. At the time, Xerox was heavily promoting their Ethernet, and there was a series of commercials airing heavily on Monday Night Football about 'My computer can talk to your computer,' and a guy drawing on the walls, and 'We can share a printer, we can share files,' and all the rest of this. Basically, it was the Star Systems, the 8000 series, the Xerox 8010 Information System. And the executives at IBM didn't understand anything about communications asked 'What's our answer?' And so they insisted they must have an answer, and so the task force was formed in the fall of '80 to compose what the IBM alternative would be. In fact, two proposals surfaced, and I remember the meeting well. One of the recommendations was we just adopt Ethernet, it’s just fine. And the answer came back 'We can't do that, because you can't be an industry leader by following somebody else's implementation.' And at the time it was already pretty clear that DEC was getting very closely aligned to -- it wouldn't be so bad if it was Xerox only, but having DEC in the fray, that was like a declaration of war. I mean, the Axis Powers had formed, and so it had to be IBM had to have a different solution, so the alternative was to pick the token ring. The token ring -- I'm not as familiar with the early stages, the mid to late '70s, but the Zurich lab had been working on -- and Zurich, incidentally, was part of the IBM research organization which has no product charter whatever, I mean, they are pure research, and it is closer to an academic environment than it is a business environment. There is no metrics for P&L or whatever.

Pelkey: Sounds like Bell Labs...

Warmenhoven: It is, it's very similar, it's a pure research environment, and one of the guys here

whose named Werner Bux -- in fact, Werner was given, I think, several large awards, and I forget the other --

Pelkey: B U C H

Warmenhoven: B U X. Werner Bux and, there was another guy whose name escapes me at the moment, had basically taken the concepts of the Cambridge Ring and various others, Saltzer's work, and so on, and had refined it to an architecture for a single ring network, one ring, and it was -- it had an architecture which was seductively simple, even though it was incredibly complex. I remember him describing this thing on one slide, it's really simple. You know, there's three diagrams of basically a ring with a little red token on it, and of course the token is free, and somebody wants to transmit. The next diagram shows they turn to busy, their message, it comes back around, they strip it off trying to (unintelligible). Very easy. It was incredibly complicated with all the token management issues, but you could describe to somebody in sales: “This is how you control it,” and it was -- it had the right cultural aspect.

Pelkey: It was like a polling ring.

Warmenhoven: And IBM had developed, in the late '70s, at least three different loop systems, which were all based -- all basically for the 3600, 3700 series of controllers, so it was a banking system, the B-Loop, and the banking system is still in use today, the supermarket retail systems with all the scanners and the rest of it. The supermarket and retail system is based on what's called an S-Loop. And another was the 3790, which was really focused on insurance and a number of other administrative kinds of organizations, and that was called the R-Loop. And all three of them had the same basic architecture, they had different performance characteristics and so on, but the same basic architecture of the controller, which was a microprocessor based thing that IBM had developed, that would issue a poll, and each station would, in series, append its message at the end of the then-existing message, so it would repeat whatever was on the ring, see the go ahead flag, generate its message, and then it would all come back to the controller, in which case it would acknowledge them and so on and so forth. So they already had the notion of this -- in fact, it was called an ORP, an Option Response Poll, ok? If you had a message you could broadcast it -- so the notion of the token ring was just the next, if you will, step in the evolution of an ORP, except now, instead of a single master station generating -- in fact, this is where Soderblom comes in: IBM reached the settlement with Soderblom, the license agreement with Soderblom, not because of token ring, IBM did it because the Soderblom patent really hit right at the heart of these loop systems, especially the R-Loop, the one used on the 3790 and 8100. If you look at his, it really was a master station generating a poll which multiple stations could respond to, and it had nothing to do with the token ring, ok? And IBM was accused of reaching an agreement with Soderblom and not worrying about the patent when they did the patent well before they even had any notion that token ring would ever exist.

Pelkey: Right.

Warmenhoven: But, anyway, that's all background. So anyway, it became -- it was culturally acceptable, ok, with a controlled environment, you know, it was all these other things, it was the natural evolution of the ring systems they had, and so on and so forth, so it was very appealing

and it seemed very simple, and it was different than Ethernet.

Interruption in Interview

Warmenhoven: We're back on the air?

Pelkey: If I understand correctly, though, in December of '80, IBM was talking with whoever was at the IEEE meetings.

Warmenhoven: Yeah.

Pelkey: So by this time, it --

Warmenhoven: By December, yeah, but the key decisions were really made during November of '80, yeah, November of '80, that we would, in fact, proceed with the development of some kind of a LAN system, and Danny Sze, who now is working with Brooktree Corporation down in San Diego, Danny was the representative to the IEEE, and he had been going more in a monitoring capacity initially, but we were really cranking up to go propose the token ring system. And he basically took the Zurich architecture definition, highly distilled, if you will, and basically proposed that as a token ring. But you're right, at that time, there was no product. Zurich didn't even have a working implementation of the then-current architecture. They had been through several iterations, but the architecture submitted had no corresponding implementation. They were in the middle of building one at that time, but --

Pelkey: Sze?

Warmenhoven: I should have brought my little computer with me, it's got all my phone numbers, but it you're interested in talking to Danny, he can fill in some of the background of '80 as well. In fact, from a historical view, you might want to talk to another fellow by the name of Jim Markopf. Markopf had been given the assignment in Raleigh, at least a year before I got involved, of basically what IBM would term "advanced planning." It was kind of a "what if" kind of thing. He had a small group of people who became the interface between the Zurich guys and the Raleigh laboratory and the SNA organizations and all the rest. Markopf had been part of the SNA architecture crowd, so he was kind of in charge of putting in the first path proposal for what the IBM answer would be. Danny Sze worked for Jim Markopf, and so Jim really had his fingers in a number of different things trying to pull it together. From an early stage viewpoint, he can describe the transition from Zurich to Raleigh and how all that stuff occurred, and the IEEE, so he could fill in a lot of blanks.

Pelkey: Did you ever go over to Zurich during this period of time?

Warmenhoven: No, not in '80. Several of my people did, but I didn't.

Pelkey: And your role in network products, at that point, was?

Warmenhoven: Well, officially, I was still the engineering manager on the 3710, and I was

drafted into this task force to put together the product plan for the token ring. Coming out of that --actually, I didn't actually become part of the management team on token ring until the middle of '80. After the task force was over -- I'm sorry, yeah, after the task force was over in the end of '79 -- did I get this, no I'm sorry, middle of '81 I became part of the token ring. After the task force was over in the late '80 time frame, I went back to my normal responsibilities, and when the token ring organization got put together, I wasn't part of it. Murray had responsibility for it, and I was working for Murray at the time, so he had both the 3710 and the token ring development, but I was still on 3710. Then, a fellow by the name of -- geez that escapes me -- Cantwell, Frank Cantwell, was charged with putting together, basically, the system architecture to a macro-system plan for token ring. What I mean by that is not the details of token ring, but how it related to SNA, how you'd relate to other devices that were going to attach to it, getting the product organizations for those systems committed to incorporate the token ring into their -- You see, when IBM did a token ring, it did it a funny way. If you look at somebody like 3Com, they built boards to go into PCs and whatever, right? But in IBM, the token ring organization was set up only to build subassemblies, and each product organization was given the directive to integrate that subassembly into their individual products, ok, but that became a massive pushing on a rope problem. What product manager in the world wants to spend money on engineering for something that doesn't exist? So, anyway, Cantwell said: "This is stupid," and basically said: "I want out." He moved up north to Kingston, and they asked me to take on the job. So in, I guess it was roughly June or so of '81, I became the manager of what was called System Design and Architecture. And the other aspect of it was IBM was insistent that the token ring fit into SNA, had to have a migration path to SNA, and they were insistent that it be able to support the installed base of devices, which were dominated by dumb terminals, I mean, 3270 terminals. A 3278 has the IQ of a stone. It had no microprocessor in it. When you hit a key, it flowed back to the controller at the end of a coax, which would then reflect it on the screen, and the screen and the keyboard were logically disconnected, except for the umbilical chord, right. So they said: "We want to take these 3278 heads and patch them to the token ring," and I -- my eyes glazed over like yours did. I mean, this is -- they said: "If your system can't support this 3278, obviously you have the wrong system." So, anyway, it was --

Pelkey: They would have worked good on Ethernet.

Warmenhoven: Right. But the other thing, you, back on the cabling system, you had asked about who invented the cabling system, another issue that IBM was getting beat on the head on was 'How do I wire a building?' Because the 3278 coax was different that the 3272 coax. There's two generations of 3270. The original generation used the technology called ANR, Alpha Numerical Replacement, and the other one used the later one, which was the 3274 controllers and the 78 heads and so on that use a different technique whose name escapes me at this point, but they used different electrical transmission schemes, and so they specified different cable. Now, it turns out that one can be reconciled, but then you had things like the 5250 Workstations and the System 34, and they used a different kind. It wasn't coax, it was basically a twinax system, and it was a multi-drop twinax, a little closer to what you'd find in those channel kind of structures, you know, and that uses a transmission rate of 1.5 megabits. I remember it was called the T1P1, T1 being the 1.5. Then there was all these loop systems, the R-Loops, the B-Loops, the S-Loops, and they all used a different kind of twisted pair structure, and on and on. I mean, it never ended. So they said: "We want a universal cabling system so we can tell our customers

how to cable a building," and, of course, once they got a cable for IBM stuff it's easier for us to plug our equipment in. So, anyway, one of the objectives for the, now, LAN, the IBM LAN, became a unified kind of cabling system. Let's see, let's back up to the task force in '79, excuse me, '80. This probably went from, basically, early October through, let's say, the Christmas period, which, incidentally, is IBM's annual planning meeting, so part of what we were trying to do was get a resource definition to be able to commit funding to build something, ok, but it was partly that they have a systems strategy. And I think Murray and I were probably the key architects of what I would consider to be three fundamental elements in this thing. The first one is that it had to be a structured system, so even though it was a ring, it was going to be a star wired system. It was going to model the phone systems, ok. In fact, part of the product plan was basically what turned out to be the SynOptics product line; smart concentrators. It started off with passive concentrators, but the whole cabling system and everything was kind of a star. The second one was that it would not be single ring, it would be a full mesh, ok, and it would be strictly peer to peer across the mess. Now, at the time, SNA was, you know, mainframe polled, and so here's your first anathema at IBM, that it's not mainframe centered. It is totally distributed. The focus was going to be on things like PCs, and intelligent workstations, and, at the time, the Displaywriter, if you recall, the predecessor of the PC, the 8088 or something. That was a hot product. It had facilities for printer sharing and very primitive resource sharing. So we basically said: "That's the model we want to generalize," and of course the PC was just starting to show up on the horizon. In fact, at this time, I don't think IBM had announced their first PC yet. We knew of it internally, but -- and Don Estridge was in place, but --

Pelkey: Because that was August '71.

Warmenhoven: '81.

Pelkey: '81, excuse me.

Warmenhoven: So the organization had formed and we knew it was off and running, but we all kept thinking the PC is the way things are going to go. It's going to be smart workstations and minis. So the notion it was going to be --

Pelkey: And was the third element peer to peer?

Warmenhoven: No, that's all -- SNA compatible, but strictly peer to peer, full mesh network. None of this multi-systems network crap from SNA and no mainframe requirements to run SNA, so the ring itself would be a full mesh network thing, ok? So it was a little bit more like an X.25 orientation. X.25 was in the architectural concepts included in this thing, and a lot of similarities --

Pelkey: And the third element was?

Warmenhoven: The third element, and this is really an interesting wrinkle, it had to be available for non-IBM device attachment, and this was another anathema at IBM. This leads to several conclusions: The first conclusion is, you want the chip set available in the semiconductor merchant market, and IBM had basically had a business strategy that said they were not going to

enter that market. So here's a universal local network, right, that had no way for DEC or Xerox or anybody else to attach, and clearly, that's a problem. That's got to be a problem. So Murray, basically, I think, more than anybody else, made the decision that he was going get the chip set designed by one of the major semiconductor manufacturers. The three big candidates were Intel, Motorola and TI, and he spent most of 1981 negotiating with those guys about how to do it.

Interruption in Interview

Warmenhoven: Yeah, the corollary to the third point about the chip set being in the semiconductor merchant market, is you had to publish the architecture. In fact, that's what lead to the 'Let's go contribute to the IEEE.' And the system had to, I use the word 'appear,' even though it's kind of an understatement, appear non-proprietary. That really flew in the face of everything that IBM traditionally believed, right, that if you're going to be successful, you've got to have proprietary technology, but the problem here, clearly, was to be successful in the LAN world, you had to be able to support the attachment of everything else, and to be able to support that attachment, you had to publish it. So success required non-proprietary, and I can remember arguing, probably, a hundred times with various executives as to why we had to give away the technology, at least in a paper form, an architectural form, and they used to think I was a stupid kid, you know. I mean, honest to God, they'd go: "You don't understand," and I'd go away shaking my head about: "No, you don't understand." There's a hundred computer companies at this time, and the one thing customers always kept saying was they don't want to be locked into IBM. So how do you come out with a new system in the early '80s that's IBM only and only IBM stuff can attach to it? Nobody will buy that. So, you know, it was partially driven by the notion that IBM sold systems. They didn't sell communications, they sold systems, and communications were only there to hook up their systems, data connect them, and as long as that was your marketing approach, a system-independent local network didn't make any sense to them. That's a non sequitur. But the problem was that we were asked to go build a local network. We weren't asked to build an interconnect for systems, we were asked to go build a local network to compete with the DIX consortium and Xerox and the rest, so it became a major problem, conflicting objectives. It was bizarre.

Pelkey: Unbelievable.

Warmenhoven: Yeah, it really was. So, anyway, Murray spent a year, basically, during 1981, negotiating with semiconductor companies, and got to the point where TI was selected, and I think that was circa the '81-'82 annual boundary, and then spent another year negotiating a contract trying to do it. Of course, development is already underway, and I think they -- as I recall it, the contract got signed in roughly the first of December of '82. Now, I remember the meeting. TI came in probably ten days after the contract was signed, and said they were off the schedule by a year. I mean, they may have started the discussion 'We're off by up to six months,' but by the time you get into it, it's off by a year. And, from their perspective, they felt as though the specs had changed, all kinds of other things had gone on, but from our perspective, we had signed a contract ten days ago -- I'm not sure it was exactly ten, but some very short period ago -- and all of a sudden you realize that the dates you committed to, within less than a full moon, are off by 50%, and you're going 'Geez, what's wrong with this picture?' So, that really set that one at odds. Anyway, in '81 is when I got into the system architecture game, and my challenge was to

go get all the other product managers in the company to commit to putting their systems on this. I got into every issue that related to SNA and interface to cabling systems, and all the rest of it. The cabling system was a big issue, because that meant they had to go test all their devices for distance and noise and all the rest of that good stuff, so they had to do most of that work too. A lot of time and energy was spent on that cabling system, and of course the other one that, in the cabling system, that became a big issue was which fiber. AT&T was pushing premises wiring structure, I think they called it, and they were recommending a 50/125 -- 50 micron, I guess it is -- interior diameter and a 125 exterior, outside diameter, and we had already picked 100/140, so now you've got two conflicting cable alternatives. I'm kind of not sure 100/140 is used much anymore, is it?

Pelkey: I don't think so.

Warmenhoven: I mean, the token ring could run on either one, it really didn't make a lot of difference. It has different attenuation characteristics and all the rest. The 100/140, I think, is what we had already submitted to the IEEE that got written into the standard, and so, what do you do then? So, anyway, back to the cabling issue, we had to pick a cable that would support 3270s and all the rest of the stuff, so part of the definition of the cable was not wrapped around token ring, partly it was wrapped around 'what do you need to support all this other junk?'

Pelkey: Precisely, junk.

Warmenhoven: I use the term loosely, but it had to be a universal cable, and a lot of the stuff that came in the shielded twisted pair was just for that reason. The other thing that IBM was paranoid about was any kind of regulatory issue. A lot of LANs, at that time, radiated like crazy. IBM was going to enforce all of the system standards, and I mean really enforce it with lots of margin, right? The IBM requirements were worse than the FCC, because they wanted make sure they weren't in trouble. And we had this plan for a structured wiring which had a big wiring ring, and the ring is all clocked identically -- I mean, every lobe is looking at the same clock cycle, so you get the aggregate of all these things added together -- so a little bit of radiation, multiplied by, let's say, 100 connectors, and all of a sudden you've got a huge antenna. So a lot of the shielding and the connector design and all the rest of that went to just minimizing that at the wiring panel, and terminating the shield into the connector easily. I mean, there was all kinds of work involved. The other one that cracked me up, I remember, on the cable was that one of the IBM requirements was that the cable be able to withstand temperature cycling down to minus 40. Now minus 40 C and F are the same. I mean, this is cold. Minus 40 C and minus 40 F intersect. They would run tests on temperature cycling, and of course some of the fiber at that time, fiber was a fairly new technology, it would crack. It would get stress fractures or whatever, and we went through hell trying to get Siemens and Siecor, I guess it was, actually, the Siemens Corning joint venture, and Sumitomo Electric and a few others, to finally produce something that worked. The third thing that was going on at the time was that the EPA, or somebody -- I have to remember who it was, OSHA, OSHA, I guess -- came out with this toxic smoke emissions for stuff left in risers, so if it burned, right, it had to be Teflon coated, and, by the time we got done conforming all this stuff, we had a cable. I mean, I described this thing to senior management one day, Jim Boyle. I'll never forget, I described all these characteristics of this cable, and he said: "I don't believe this." He says: "You've got a cable that's bulletproof, rat-proof. I mean,

you know, rat's bit it, they're going to die." It was unbelievable. Anyway, back to the token ring story, Murray left at Christmas of '82, and went to Codex. He and I were kind of peers. He had the engineering organization, I had the systems design organization, both working for the same product manager, anyway, both working on token ring. He had other products under development in his organization, but what they decided to do around this time was, when Murray went to Codex, was combine the two and I took on the whole thing. So, by January 1 of '83, I was running a combined organization, and the marketing, product planning kind of stuff was under a guy by the name of Tony Terrel. So there's a business planning function. This is, obviously, the time where TI had just walked in and said: "We can't make it." So I climb into the seat, the engineering seat, not having had a great deal of involvement with the TI contract, Murray did it, basically, and the first thing I got was 'We're a year off.' Now, the contract really had just been signed. I can probably get the exact dates if you're interested. It was within a matter of weeks, and of course there was a lot of speculation that Murray left because he knew that the plan was going to slip. Personally, I don't think that was the case. I think he got a better deal up at Codex, but that became a real problem. So, all of a sudden, IBM's irate. They didn't want to give the stuff to TI in the first place. They didn't want to have a third party. 'How come our technology's not good enough?' the Burlington guys: 'How come we can't go it in house?' And the whole issue had nothing to do with technology, it had to do with openness, system openness. How were you going to sell it? We can't tell engineering that. So half of the IBM power structure came descending upon us. I was -- I mean, it was -- in corporate relations, I was -- it was a bloodbath.

Pelkey: I can imagine.

Warmenhoven: And from that point on, I spent -- I remember in 1983, I racked up 300,000 air miles and never left the United States. I went from Raleigh to Harrison, White Plains, and back, and out to Dallas, and back, you know, TI. I spent --was making, on average, two trips a week up to the headquarters area in White Plains, and it was -- actually, it was Harrison. It was across the street -- literally, across 287. It was unbelievable.

Pelkey: You must have felt very productive.

Warmenhoven: Yeah, I figured I spent 40 hours a week holding the thing together, either with TI or the IBM executive set, and then I would work weekends and evenings to try to keep the engineering program running. It was unbelievable, and I went that way all the way through 1983. And, of course, every product manager we had signed up all of a sudden said: "Well, shit, you guys aren't for real. You're not going to ship." They cut it out of their plans, so there was all this putting it back together again. And the SNA guys were against it, because it killed their architecture, right? So the token ring became a target for a dozen different organizations within IBM. It was unbelievable. I look back on it thinking that it was almost impossible to make it through. People ask me what it was like, and why did I leave IBM, and I tell them the reason was that I did this token ring thing and I felt like we really achieved something of significance, but it was so goddamn hard I felt like a salmon swimming upstream to spawn. And if you could finally make it over all falls and all the rest of the stuff, you're going to die anyway.

Pelkey: You're dead. You had no future in the IBM organization.

Warmenhoven: Right, so, anyway.

Pelkey: Now, did you do Sytek?

Warmenhoven: Well, yes and no. That's another interesting story.

Pelkey: Because Sytek must have been happening in '83.

Warmenhoven: Well, yeah, Sytek was happening in '83, and I was part of the due diligence team that went out there to look at 'should we acquire them?'

Pelkey: And did you have anything to do with the development contract with UB?

Warmenhoven: Oh, Murray started that and I finished it, yes. Murray had started the discussion with UB, and I put the deal together.

Pelkey: And that deal was about what?

Warmenhoven: UB was -- well, the token ring plan was a very, very aggressive plan to build an entire infrastructure for the network, bridges -- there was the notion -- you'll get a kick out of this -- bridges and tunnels. A tunnel was a two-ring bridge, so ring to ring only, and a bridge was viewed to be a multi-port notion, so you could have either four or eight so interconnected, and it was more of a large device for rings coming in -- and then there was the wiring cabinets, and the wiring cabinets had both passive and active, and so anyway, it was a very aggressive plan for, really, a lot of components, some of which were very sophisticated. I mean, hubs, I mean all the definition of hubs, and how to isolate a lobe, and diagnostics, and basically what was the ring version of Sniffer was in there, as well. So you could think about what is now Madge Networks, or SynOptics, to a degree, the sniffer guys from Network General. Now, all of those were basically Ethernet oriented when they started, but then this definition, the whole list was defined in token ring in 1980, '81. But anyway, it was a very aggressive plan, and the marketing guys wanted the whole damned thing, because they saw in this a way to catch up and pass Ethernet. It would be a highly managed network, it had a complete solution.

Pelkey: Yes. We want it all immediately.

Warmenhoven: Right, and so it became a huge engineering bill. I mean, I put together an engineering plan for 300 people. Of course, the cable system was in there, right, and all the bounds, or whatever you call them, the interface devices between the old devices and this cable. I mean, 'How do you make a 3278 head designed for coax run on a twisted pair?' Design an interface that works. So, there was all of those components, all the wiring components, the ring adapters, the network management software, all these wiring components, bridges, and the little gateways, basically micro-routers, this split bridge notion, and this was a very extensive set. So the first thing they said is, you know: "We can't afford this," not surprisingly. Actually, when I left, in 1983 as I recall, or maybe it was '84, my engineering budget was $30 million for this. It was a BIG undertaking, and the organization, in fact, grew to about 250 people, but even so, we

couldn't afford to do it all, and Murray recognized that, so he started discussions with guys like Ungermann-Bass to see if they'd be interested in building some of these parts. The appeal to them was they had early access to the TI chip set, so they could be first to market, there was a time to market advantage, plus they'd have IBM as a distributor of whatever they built. So, there were a number of, basically, card level products; one for the PC.

Pelkey: Which was Sytek.

Warmenhoven: No, I'll come back to Sytek.

Pelkey: Ok.

Warmenhoven: And the bridge product, the bridge-tunnel, both of which were given to Ungermann-Bass.

Pelkey: So the PC was given to them as well.

Warmenhoven: Yeah, a PC card, for token ring PC. That's another interesting one. There's another token ring chip set running around that a lot of people were never aware of. Ungermann-Bass did theirs. They did it for this card, this PC card, with Toshiba.

Pelkey: No, it was Fujitsu.

Warmenhoven: Fujitsu, excuse me. They did it with LSI logic, which was using the Toshiba gate arrays, but the real foundry source was Fujitsu, and Charlie Bass was the counterpart on that. He can fill in a lot of that.

Pelkey: I'm meeting with him today.

Warmenhoven: Is that right? Well, say hello to him for me. Charlie and I are still good friends. We met in 1981 or whatever, just at the time of this.

Pelkey: So, this Ungermann-Bass thing was, what, '82?

Warmenhoven: It started, really, in the early part of '82 and ran all the way through the time I left in '84. That was another thing that --

Pelkey: Did you have other companies in addition to UB working on this?

Warmenhoven: It was UB primarily. We looked at a number of others. We looked at a number of others. We went to Proteon, we looked at -- there was one down in, it's a company that, eventually, Ungermann-Bass bought --

Pelkey: Oh, yeah, Amdax.

Warmenhoven: Yeah, probably Amdax. They had a 50 megabit ring.

Pelkey: It was a bunch of IBM guys.

Warmenhoven: Right. And then we looked at -- PC, how did the PC come about, the Sytek deal? The PC guys, when we took this one year slip, there were two issues for the PC.

Pelkey: Because UB merged Amdax into them on January 1, '83.

Warmenhoven: Right. We had looked at Amdax prior to that, and then UB bought them. The personal computer guys, first of all, just as background, were exempt from all corporate bullshit, to use a term. They were really set up almost like a wholly owned subsidiary or something. They had no requirement to follow IBM procedure or anything else. They were basically on their own.

Pelkey: That must have been interesting for you, given what you were trying to do with cabling and connecting everything else, and then to have these guys be able to go wherever they wanted to go.

Warmenhoven: Right. In fact, I remember we got into a big shoot-out over this PC network thing. Don Estridge, who's a hell of a guy, he was -- you'll never get a chance to meet him because he passed away in a plane crash --

Pelkey: No, I met him.

Warmenhoven: Oh, you did? He was a world class guy.

Pelkey: I was president of Sorcim.

Warmenhoven: President of?

Pelkey: Sorcim. We did SuperCalc.

Warmenhoven: Well, Estridge was pretty savvy in a lot of ways, but, anyway, background, his organization concluded that whatever network they came up with, first of all, they didn't want any dependency on anybody else outside, they were going to do their own. Number two, they wanted to be able to support video. They wanted to have a response to what was then Apple with educational systems, and a lot of the educational systems were video based. So they wanted a network where they could basically control video, and run video down the wire, and we could do that, but only at a great expense. So they concluded 'We want a broadband system. We want a cable TV system,' ok, and that's what went up to Sytek. There were two at that time, Sytek and I forget who the other company was, they eventually paved their way, but these guys had lined up Sytek to build a version of the Sytek broadband network for an IBM PC. That went up to the highest levels of the company, 'why are we doing two local networks?' and the question to us was: "If your network is supposed to support PCs and the PC guys don't want it, what good are you?"

Pelkey: Now, this is in '83?

Warmenhoven: Right, this is '83 now.

Pelkey: Because Sytek met IBM, conversations began in '83 after a conference at Brown University which Sytek had wired.

Warmenhoven: Right, that's correct. And all their work was done between Sytek and the IBM PC group down in Boca. The product manager down there was a guy by the name of Jim Turner, who was responsible for what turned out to be the AT. I forget the code name internally, but the AT was supposed to be the introduction vehicle for the LAN, the IBM PC Network. At one time, the PC guys got an edict from the corporation that said: "No, you will use token ring." We slipped in early '83, and they said: "We can't hold up the AT announcement." So that's what gave them the go ahead to go do their own.

END OF THE INTERVIEW

