Interview of Robert “Bob” Kahn

Interviewed by:
James L. Pelkey

Recorded February 22, 1988
Reston, VA

CHM Reference number: X5671.2010

James Pelkey: We are in Dr. Robert Kahn's office. Thank you in advance for making the time for our conversation. Maybe you could begin with your being at BBN in the '68, '69 timeframe, and working for Frank Heart[1]?

Robert Kahn: You'd like to go back and explore that whole period.

Pelkey: Yes.

Kahn: Okay, -- let me go back a little bit further.

Pelkey: I have a series of questions that I can fill back in later, but...

Kahn: But the BBN link is because you're interested in--

Pelkey: Arpanet.

Kahn: Arpanet, and packet switching?

Pelkey: Yes and onto TCP/IP and the LAN business. I'm interested in intellectual property, where ideas come from and who came in contact with whom and how these ideas grew.

Kahn: Okay. Now, if you manage to disentangle all of the Arpanet contributions, you'll have to understand at some fundamental level that you're dealing with a whole set of people who were involved in the evolution of the ideas of packet switching, that ranged from people actually involved in doing Arpanet to people who actually had no role in building Arpanet, but all contributed to the process. And it's going to turn out when you're all finished that there really is not a single, crisp answer to the question of who was responsible for it, because it was a compilation of a lot of contributions from a lot of people. I'll give you my perspective on that.

Pelkey: So you know, my first interview was with Paul Baran[2].

Kahn: Okay, well, Paul represents one end of the spectrum; Paul could properly say the ideas were his. And I can explain in what sense they could properly be viewed as his, and in what sense they could properly not be viewed as his, and so, that's again one of those anomalies.

Let me give you a little bit of my own background first. I got my Ph.D. in Electrical Engineering--there were no computer science departments in those days -- from Princeton in '64. During that whole period of time, I had also been working at Bell Telephone Laboratories as a member of the technical staff. I was primarily an applied mathematician and communication theorist at the time. After getting my Ph.D., I went to teach at MIT. I was on the faculty in the Electrical Engineering Department. I took a

---

[1] See Interview – Frank Heart
[2] See Interview – Paul Baran

leave of absence from MIT in the fall of 1966, primarily because I had been encouraged to get a little more practical experience with real systems. As a theorist without much practical experience, it's hard to know which problems are of real practical significance, and which ones are just intellectually interesting. The suggestion was made to me by some of my close colleagues that it would be beneficial to get some practical experience; and then maybe continue working on theory, or whatever. And so that's why I took a leave of absence at that point. I went to Bolt Beranek and Newman where I began to work on networking issues, since that seemed to be a very practical kind of thing. In retrospect, it was still rather conceptual work, but instead of just being a problem of intellectual interest that had been posed and formulated, I viewed this as a real national challenge to design some kind of a network to link computers and to figure out how it might actually work.

Pelkey: Now networks at that point in time were really primarily just modem links.

Kahn: There were no networks other than people dialing in through the telephone system. A 300-baud link was considered a good, solid network connection at that time. Timesharing had been around for a few years at that point; I had my first exposure to timesharing at MIT. And there were a number of people, including myself, who had access to timesharing systems through terminals that you dialed in from home, but they were generally the old Model 33 teletype-like terminals -- nothing like you have today.

I began thinking about networking while I was at BBN. I hadn't really started thinking seriously about that at MIT. I had encouragement from a fellow named Jerry Elkind who at that time was running half of -- let's see, the timing is not clear in my mind -- then or very shortly thereafter, he ran about half of BBN. Leo Beranek ran the other half. So all of the work in information processing was under Jerry. In the 1967-1968 timeframe, I thought I had a fairly complete view in my own mind of what networking might be like. I'd been writing technical memoranda, you know, how different parts of the system might work. And I remember at one point Jerry Elkind came to me and said, "There's a fellow who's just gone down to ARPA named Larry Roberts[3]. You really ought to go and meet him and give him some of your views of networking, 'because I think he'd be interested." And I remember writing a letter to Larry in 1967--I probably still have a copy in my archives--basically saying that I had been thinking about networking for a while, and had some ideas I thought he'd be interested in. If he wanted to take the time, I'd be happy to share them with him. And next thing I knew, I got a phone call back saying, "Come on down, let's chat." And I didn't know much about ARPA at the time. Its now called DARPA, but it was then called ARPA. Then I went down, we chatted, and Larry explained to me how, in fact, he had gone to DARPA to kind of make some kind of a computer network happen, and he'd be very interested in any ideas I could share with him. He had a number of people around the country who were helping him shape and form some of his ideas; people like Wes Clark[4], which is a name you may or may not have heard. Wes used to be at Lincoln Labs, but he was then a professor at Washington

[3] See Interview – Larry Roberts
[4] See Interview – Wesley Clark

University in St. Louis. Wes should probably be credited with the idea of placing the key network functions in a separate minicomputer rather than locating them in the mainframes. In the early days, they were thinking of linking a machine directly to another machine. The idea of store and forward had come out of Paul Baran's original work. The idea was that the mainframes would pass messages around. It was Wes, I think, who was primarily responsible for saying, "Hey, don't put these communication functions in the mainframes. Factor them out separately and build a separate kind of communication system based on minicomputers, whose resources are not being commandeered by other people's computations."

Pelkey: That's an important contribution.

Kahn: I think it was pretty fundamental. It was just an organizational paradigm.

Pelkey: Do you know where Wes is now?

Kahn: Yeah, he's in New York. I think he and his wife are doing consulting.

Pelkey: New York City?

Kahn: Yes. His wife's name is Maxine Rockoff. The consulting firm is probably called Clark, Rockoff and Associates.

Pelkey: Now, at this point in time, were you familiar with Paul's work?

Kahn: No, I wasn't. I knew that there was a Rand study on distributed communications, but I didn't have a copy of it and hadn’t read it.

Pelkey: And were you familiar at that time with work that Newhall had been doing at Bell Labs when you were there, or subsequent to that period of time?

Kahn: Not when I was there, but I subsequently learned about it

Pelkey: Given your work on networking?

Kahn: I was generally familiar with what had been going on in the communications field, but there wasn’t much work on networking. So, I had just started to think about the networking problem from scratch. At BBN, I was writing my own thoughts down. I didn't have any network reference points or reasons to search the literature (which I was generally aware of) because I wasn't yet trying to contrast and compare my work with anything. Shortly after I'd gone down to visit Larry and to share some of my ideas with him, he had asked me also if I was familiar with the Baran material, and I said-- I knew there was such a report, and he suggested I go get a copy of it, and I did look at it after that. But that was I would say--

Pelkey: In '67?

Kahn: Either then or early in '68. Anyway, I took a leave from MIT to go to BBN where I went to work for Jerry Elkind. When ARPA put out the RFP for its network, I was responsible for putting together the technical part of it together with a fellow named Severo Ornstein. You see, he was a kind of reality check for me on the engineering details. Severo had built a lot of equipment. He used to work for Wes Clark at one point at Lincoln. Severo was a hands-on engineer. Our conversations might go like this - I would say, "Look. Here's a malfunction and here's the way this has to work." And he'd say, "Well, this is the way you do it." He translated many of the ideas into fairly concrete realities. Together we put together a detailed technical plan that later came under review by several others including Frank Heart and Dave Walden. After the project was funded, Bill Crowther joined the effort in charge of software. Severo is a "seat-of-the-pants" engineer. His unique contribution was the pragmatics of the network implementation. If you were an architect, you would first design a building, and you would plan its implementation before building it. It was an important kind of learning experience for me at that point. To learn how you take ideas and map them into reality. It was also a very important reality check to know the ideas were actually fundamentally sound.

Pelkey: Yes. Now the spec- did RFP come out and specify a network?

Kahn: Yes.

Pelkey: How much was specified, and how much was left up to the responders as to what...?

Kahn: Well, you ought to get a copy of the original RFP if you can, but it basically said that ARPA wanted a four-node network that had the potential to expand. It had a notion that there would be these minicomputer switches called IMP's. And it indicated that one IMP would be at SRI, one IMP would be at the University of Utah, one node would be at UC-Santa Barbara, and one IMP would be at UCLA. In fact, they were installed one a month in the order of UCLA first, SRI second, UC Santa Barbara third and Utah fourth between (roughly) September of '69 and December of '69.

Pelkey: BBN wasn't a node.

Kahn: BBN was the fifth node on the net. It was specified at the time that there would be 50 Kilobit circuits connecting the nodes. The RFP stated that there would be computers attached at each of these sites, although I don't remember if it stated specifically what machines or not.

Pelkey: Do you know how those nodes were selected?

Kahn: No I don't. That's a question to ask Larry Roberts, because Larry made that selection. All the sites had significant research programs in computer science. UCLA was involved there because Len Kleinrock[5] was there. And Len was one of the people

[5] See Interview – Len Kleinrock

who had done the most thinking about networks up until that point in time. SRI was selected, in part, because there were several people that had helped Larry think through the networking issues. Elmer Shapiro was one name that I remember.

Pelkey: Right. Shapiro was at the meeting when Baran presented the concept of the Rand network to Roberts.

Kahn: Could be.

Pelkey: Larry said, "Wait a minute, I can use that for resource sharing on the Arpanet." He wanted to do a network.

Kahn: Could be. And all four of these places had major ARPA contracts at the time and they had a lot of computing resources Larry was trying to find some way to somehow get those computing resources to work together. Utah also had a major computer science research program at the time. Let's see, the principals at Utah were Tom Stockham, shortly later came Ivan Sutherland and the graphics effort so maybe Dave Evans was there. Englebart was at SRI at the time; they had a big AI lab at SRI. And at UCSB—

Pelkey: Ed White was there?

Kahn: I think that the key guy was Glen Culler, who was doing quite a bit of work...

Pelkey: Right, right. Mathematics.

Kahn: Yes. Something called the Culler-Fried System. In any event, that was one of the main reasons, then. That was the initial net that was specified, and...

Pelkey: So the basic concept was specified and 50 kilobit links were specified and...

Kahn: Right. And the notion of packets were specified, although I think at the time, it was called message switching.

Pelkey: Right. Block message switching.

Kahn: Well, Paul Baran had called his work something like—

Pelkey: Message block or block message—

Kahn: Block message address switching, or something like that. Same permutation.

Pelkey: And Donald Davies[6] was the one with the packet switching?

Kahn: Davies was the one who introduced the term packets and maybe packet switching.

---

[6] See Interview – Donald Davies

Pelkey: Yes.

Kahn: I think the term "packet switching" was really attributed to him. And if you look, in fact, in the 1970 paper that I wrote with the rest of the guys at BBN on the design of the IMP (which is what we actually built while we were there) you found that we used the term message switching and we didn’t use the term packet switching in that paper, because it hadn't quite—

Pelkey: Reached common knowledge.

Kahn: It hadn't quite gotten into common parlance. We talked about packets, but we called it message switching.

Pelkey: And you defined as a subnet at that point?

Kahn: We separated out the communications portion from the rest of the portions of the system at that time.

Pelkey: So how much freedom did you have to change the design?

Kahn: In what it did?

Pelkey: Contractual. I don't know who submitted else bids?

Kahn: I don't know. Larry would know that. I think several proposals came in, maybe half a dozen or more.

Pelkey: Right. Why did yours win?

Kahn: I think it was just probably stronger technically. I don't know for sure. I never saw the other proposals and so I don't know. I believe the one that was the closest competitor was from SDC. But I don't know. Ask Larry Roberts.

So the RFP specified that there would be IMPs, it specified that there would be these packets transported. And it basically outlined what the network was supposed to do, structurally, functionally. But it didn't say, really, how to do it. And so the real task was to figure out how to take this concept and make it work. Here’s my analogy-- "We want a machine that is 30 feet long. It should have wings on it to give it lift. It should use the principles of jet engines to get it going, and it should be able to hold up to 48 tons with a range of 5,000 miles," and the only thing you have left to do is figure out how to do that.

I think Larry did a fine job in integrating a lot of ideas from the community; much of that insight really came from him, how to structure it, how to get it going, and the politics of it all. He was very deeply involved in all of it, although he wasn't the only one involved. But he laid out what he wanted and how he wanted it done. And so it was fairly straightforward to see what it was he had in mind. The question was how to make it happen. You have to understand that it was necessary to understand both the telecommunications industry and also the computing industry. And the number of people who knew about both in the country--well, I mean—

Pelkey: There weren't a lot.

Kahn: I'll bet there weren't half a dozen in the whole United States.

Pelkey: Now my understanding—

Kahn: And I was probably one of the few. Larry was probably one of the few. There weren't many.

Pelkey: My understanding is that Rand and ARPA both asked AT&T to build this project. And they were turned down. Did you ever hear that?

Kahn: That's part of the folklore. I don't know. I've heard people say that, but I don't know that for a fact. Again, Larry would be able to tell you about that.

Pelkey: The-week before Christmas of '68 you were awarded the contract.

Kahn: And it started in January.

Pelkey: What was the reaction when you won the contract?

Kahn: I don't think anybody was terribly surprised, because I knew we knew we had a good proposal. We knew it was really solid. It was sort of straightforward that we would just make it happen. It didn't particularly seem like—

Pelkey: Were you excited when you won it?

Kahn: Yes. The excitement --

Pelkey: Was there anything exciting about it, was it a big thing, or...?

Kahn: I can't speak for everybody at BBN, because there was a number of people involved, but to me the excitement was thinking it through in the first place. Once we had finished thinking it through and writing it down, to me that was the height of the real excitement. The fact that we subsequently won the contract, frankly, I was, in fact, poised to go off and work on other things, because it seemed to me that it might be just an engineering job to go build it.

Pelkey: Right.

Kahn: But it turned out that a whole bunch of detailed issues came up as we started to get into it. I finally decided I really wanted to spend time overseeing the design, its evolution, rather than going off to work on other things. So that's what I did. But the excitement for me was finishing the job of thinking through the basic concept.

Pelkey: And what were some of the things that -- in terms of designing a network, since this was really the first network -- in your proposal led BBN to winning the RFP and got implemented in the subnet network. Things that were innovative, or had to be thought through for the first time, or something you felt particularly like you had really overcome something.

Kahn: Well, I think, first of all the whole issue of how to do routing in these networks was a really interesting kind of issue; that even when we wrote the proposal we hadn't firmly resolved it. I mean, we had a suggestion for what we'd do initially, which had to do with shortest-path routing, with certain metrics, but we left that open, because we knew it was an area in which we wanted to do more research. The implementation was sort of modular; you could plug in a new routing algorithm.

Pelkey: And who had influenced your thinking about that? Anyone?

Kahn: Not really. I just wrote down some ideas and they looked implementable and... I mean by that time I had read--by the time we had won the contract, I had read a lot of the other things that had been written, and it just seemed to me like our ideas were the right ones to try first. Paul Baran's contributions to routing were actually very interesting, because they were sort of the first dynamic discussions of how you do routing, but they were largely based on what he called "hot potato" routing. And you have to understand where Paul was coming from. He was thinking of this kind of network as a survivable network from the point of view of military communications. He was working in the era of thermo nuclear war. Those were the days, you know, when the military was really concerned about what to do about this new threat. Herman Kahn had just written that book “On Thermonuclear War” which was getting everybody all upset and concerned about, you know, how you deal with that whole area, and so the military was looking for ways to deal with communications in very violent kinds of threat environments. And this was one idea. I think Paul was motivated almost entirely by voice considerations. I mean, if you look at what he wrote, he was talking about switches that were low-cost electronics. The idea of putting powerful computers in these locations hadn't quite occurred to him as being cost effective. So the idea of computer switches was gone. It didn't show up in any of the writings. The whole notion of protocols didn't exist at that time. And the idea of computer-to-computer communication was really a secondary concern. He was mostly interested in voice. Go back and take a look at his papers.

What finally happened in networking was, in fact, a strong derivative of what Baran was talking about, but he had not really laid a clear template for what it was. Now the RFP that Larry put out from ARPA was a very clear template. I mean that was one-to-one with what actually happened in terms of function and structure.

Let's see, we're talking about issues here: routing was one.[7] The whole issue of congestion control was another issue that was very interesting to me. With all these packets, how are you going to keep the network from congesting? And there was another area that I was very concerned about and what we used to call: "deadlocks." That is, the possibly of actually having the network come to a grinding halt. And that was one of the more contentious issues that we dealt with, because, you know, this was just a notion in my head at the time that this could happen. It turned out that not only did it happen; it was one of the major things that one had to be concerned about in the design of computer networks. But it was impossible for me to convince anybody that that was real. And so we actually made it happen. The very first test I actually ran on the net was a deadlock test. We deadlocked it. And at that point, everybody else who was busy writing software and building hardware got—

Pelkey: Is that why you went to UCLA early on to be able to force some of these issues?

Kahn: Well I actually ran the first testing with Dave Walden, who is now the president of BBN labs. Dave worked with me on that. He was one of the guys on the team. And that was one of the first experiments we ever ran-- the deadlock experiment. I mean Dave got to see it right there. I think it was six months later before everybody else sort of gradually came along to really understand that. It was a very frustrating time, because I couldn't convey that notion to anybody else despite the results.

Pelkey: And the network, when it was first up, if I understand correctly, really got into some really serious deadlocks?

Kahn: No, you had to really force that. Most of the time the network worked just fine. And in fact the argument was that deadlocking would probably not be serious, because, it's like, arguing, for example, "Gee, if you got 15 people walking through the corridor, they might get stuck-- sandwiched." Well, you know, if you have a lot of things going on, statistically it was very unlikely, and, everybody said, you know, "Let's worry about that when we get our first deadlock." But in fact, it was very clear to me that you could just make it happen. So we did that.

Another concern was the whole issue of--how shall I say it-- I'm just going to call it independence as an issue. How do you build a system such that each of the nodes can sort of act independently of all the others and still talk to each other and not require global control. It's sort of the whole issue of distribution and interdependence. And so everything that we unplugged one of the nodes, the system should find out it got unplugged, reconfigured automatically, and got on about its business. And when you want to put a new node in, you just plug it in, the other nodes find out about it automatically, and keep going.

Pelkey: That must have been an awesome concept at that point in time.

[7] See Exhibit A – Scan of Notes from Interview Book

Kahn: I thought we got it really tight and implemented it just beautifully. But that was one of the critical design issues. And so you can see in putting this thing together, I mean, it was a very interesting mix between the issues of overall architecture and design, which I was principally responsible for, the issue of hardware development, literally building the interfaces, getting the machines up and configuring them, and making them work properly, and the issue of software design and development. Now these areas linked very tightly. The software design and the architecture and network design, sort of mapped one to one, because a net design was only as good as the software design that implemented it. There was a fellow named Bill Crowther that was involved with software and I was responsible for the design. This whole activity took place in Division 6 that was run by Frank Heart. Let me see if I can get this right. There was a Division 4 that Jerry Elkind ran, and there was a Division 3 that I think John Swets ran. And at some point or another, I think all three of these divisions reported to Jerry Elkind. Leo Beranek ran the other part of the company. Now I actually reported to Jerry Elkind, but the people that were going to implement the network like Ornstein, Crowther and Walden were all in Division 6. And what happened was I got loaned to Division 6 to help with the implementation because I was conceptually involved in putting the plan together. And I decided to get more directly involved in the actual project itself, because it was clear that what was going to go on critically depended on the evolution of the design. So I decided to get involved for a while and I was really on loan to Division 6.

Pelkey: Where 10-X was done?

Kahn: 10-X was done in Division 4. This is where the work on 10-X was done[8]; the work on LISP was done there, all the AI work, and much of the computing research. And the work done in Division 6 was largely Arpanet related (and medical). The original work that I had done before we actually wrote the proposal and even during it, I was located in Division 4. So, I would say the initial Arpanet design was done in Division 4, except that Division 6 got involved shortly before the point that the RFP came out on the street.

Shortly thereafter, while I was reporting to Jerry Elkind, Jerry left BBN under circumstances that I can't relate to you accurately, because I don't know the details. I think there was a conflict between him and the rest of the BBN management. And so he left and he went to MIT at the time. And he stayed at MIT for about one year, and then he went to Xerox Park, where he became the boss of Bob Taylor[9], who ran a lab called the Computer Systems Lab. I think Jerry maybe was brought in as head of the computer systems lab, and Bob was associate head or something, but in fact, Bob Taylor had been running it. Licklider[10] was the guy who started the Information Processing Techniques Office around 1962 roughly.

Pelkey: McNamara forced that issue.

[8] See Exhibit
[9] See Interview – Robert (Bob) Taylor
[10] See Interview – J. C. R. (Lick) Licklider

Kahn: Well I don't know what forced it exactly. ARPA had been created in '58. With the Russian's launch of Sputnik, it was the U.S.'s way to respond to. And Licklider, who was and still is one of the most interesting people you'd ever want to meet, had a lot of these ideas about linking machines together. If you asked: "Who's idea was this network of linked computers, interactive timesharing," you could go back and say it was Licklider. In 1964, Licklider brought Ivan Sutherland in as his replacement. Ivan clearly had in mind many of these ideas about networking that were part of the culture at the time. I think around 1965, if I had to guess, along came Bob Taylor, who later replaced Ivan Sutherland. Licklider went to MIT via IBM. Ivan went to Utah via Harvard. Bob Taylor hired Larry Roberts to ARPA.

Pelkey: From MIT.

Kahn: From Lincoln Lab.

Pelkey: From Lincoln Lab, excuse me.

Kahn: Larry became the office director in '69, but he had actually been there for two years before Bob Taylor left government. So here's another one of those little conundrums that you run into, because—

Pelkey: People think Taylor was really—

Kahn: Well, and I think that-- Bob Taylor was running the office at the time, and hired Larry in to manage the networking project. If I were in Bob Taylor's shoes, I would say, "Gee, that was my idea. I just hired Larry to go and do it for me." And in fact, it's very clear, you know, when you look at all the things that happened, that Taylor had the idea in his head of what this network was going to be like. And he and Licklider had written a paper many years ago about interactive computing and the role of networking. They both had a clear vision of networking, as perhaps did Ivan, although Ivan's interests were into so many other areas. He spent a lot of time on ILLIAC IV and graphics and things like that. Bob Taylor really gave Larry a lot of free rein, you know, to make the network happen. So you can see the confusion there. Take the case of the Arpanet activity at BBN. I mean, the Arpanet project was run in Frank Heart's division. So if you were Frank, wouldn't you say, "I did the Arpanet"? Okay? On the other hand, if you were Larry Roberts, you'd say, "Well, I did the Arpanet." And then if you were Bob Kahn, you'd say, "Well, I did the Arpanet." Likewise, if you were Paul Baran, you might say, "Well, it's one of my ideas from back then; they just simply updated it." Earlier, Donald Davies in England had actually started implementing a one-node packet switch. He could realistically say, "It was my idea. I mean, I've been calling it packet switching." If you go talk to a guy like Bill Crowther who wrote most of the software with a lot of help from Dave Walden, he could say, "Well that was all just a bunch of preliminaries. The real network was the software that ran in the nodes and that was really my doing." Right? That's the kind of problem you had. And I'll give you a couple more—

Pelkey: Good. Let me ask you a generic background question given that point, because I agree. Clearly in the intellectual environment -- I guess because of timesharing -- the concept of putting computers together was something that people were starting to be concerned with. So it took some time, although you say people were thinking about it for years, but it was really just the recognition now that, wait a minute, computers were going to be around, and we ought to start hooking them together.

Kahn: Not everybody believed that. A very small set of people believed that.

Pelkey: Right.

Kahn: Larry believed that.

Pelkey: But there was a set of people—

Kahn: Bob Taylor believed that, and I think Ivan believed it. Lick for sure believed it. All the people in BBN that were involved believed it. But if you went out to the research community at large, it's not clear that they would have agreed at all.

Pelkey: But at Lincoln Labs wasn't there some network that was being done around the same time?

Kahn: It was an experiment that Larry did in 1965, I think, that hooked a Lincoln Labs TX2 to a Q32-- it may have been TX0--I'm not sure -- to a Q32 at SDC. They literally took two computers, and connected them directly via a phone line. You know, slow speed lines. I forget exactly what they did, but these were brute force experiments to show that you could move a file from one machine to the other – but not very effectively. They had some interesting ideas, like the idea of modifying the operating system so that you could talk to at least two other sites. I mean it seems like a trivial idea, but if you can only talk to one other site, you can never build a network; at least not the kind that was being envisioned.

Pelkey: Right. It takes three to be net.

Kahn: At least. And in fact, you really want more than that, because if you can only talk to two other places, the most you can do is build a loop. So in any event, Roberts became director of the ARPA office and when Larry left, he was replaced by, lo and behold, Licklider, who came back again for a second turn around 1974. Actually, Larry left in September of '73. So there was like a three-month period where there was really only an acting head.

Pelkey: He went to Telenet.

Kahn: Right. I'll get back to that. That's another story.

Pelkey: And you were still at BBN during this period of time.

Kahn: No, Larry and I overlapped for slightly less than year.

Pelkey: At...?

Kahn: At ARPA. I'll get back to that too. In 1975, Dave Russell, who was the only non-computer scientist in the job, took over. He had a PhD in physics. He was an Army colonel. He didn't know much about the information-processing field and was put in by the then-director of DARPA, George Heilmeier, who wanted somebody running the IPTO office that would be more responsive to him. And so it became more of a military chain of command type of thing. And in 1979, I took over the IPTO office when Russell left. And then in late '85, Saul Amarel came in, and the name of the office then changed to ISTO, Information Science and Technology. And in 1987, Jack Schwartz came in, and he's the current director of that office.

Pelkey: In what year did it go from ARPA to DARPA?

Kahn: I don't know; somewhere around 1972. It could have been late '70, it could have been early '72, but nominally '71.

Pelkey: Okay. I'll come back to that question. But let's go back to where we were.

Kahn: So with the exception of Dave Russell, who was a physicist by background, everybody else was within the information-processing field. He was a rather competent scientific fellow, but from a—

Pelkey: Different perspective.

Kahn: --different field and the only military person in this whole chain. Although it may have been that Ivan at the time was in the ROTC. So it's possible, but he wasn't a career military, whereas Dave was.

Pelkey: So BBN gets a contract.

Kahn: BBN got the contract in early 1969. And one year later the actual installations numbered four. So it was like nine months to actually create that four-node net that I showed you and twelve to have all four nodes delivered. We did some testing and then from '70 to '71, then it kind of grew, '72, '73-- In late '72—

Pelkey: October 24th and 26th.

Kahn: I can't remember the dates. We did this demonstration at a conference that was called the ICCC. It took me over a year of planning and preparation to make that demonstration happen. And the reason was that planning and preparation actually made the Arpanet work. I mean if you looked at the net—

Pelkey: There were 15 nodes at that point.

Kahn: Um. If I had to guess, I would have said that there were about 20, maybe 30 nodes. You'd have to back and check the books.

Pelkey: Whose idea was the ICCC?

Kahn: Well you have to separate out two things. I had been pushing for a demonstration of the net as a way to get it more usable, but I had been suggesting that we try to do it either at the Spring, or Fall Joint Computer Conference. Larry Roberts was the one who had heard that there was going to be this new conference called the International Computer Communications Conference: the very first one. He was the one that came back with the counterproposal and suggested that we put on the demonstration there. I was asked if I would do it, and I said, "Sure."

Pelkey: So you were the taskmaster to get it done.

Kahn: Well, I mean, the idea of doing a demonstration had come from me. Larry may have thought of it independently.

Pelkey: But you were the one who was responsible for pulling it off.

Kahn: They asked me to do that, right. Now let me just show you a couple of other details, here. I actually got involved directly in running this project after I had moved back to Division 4.

Pelkey: Were you working closely with Elkind at this time?

Kahn: That's right, except Jerry left the company somewhere in that timeframe. At this point, I transferred back to Division 4. When I say, "Transferred," I mean I physically picked up my belongings and moved to another office--they were in different buildings. I physically moved back to the old building. Elkind was replaced as head of Division 4 by a fellow named Dan Bobrow who left BBN around '72 and went to Xerox Park. Bobrow was then replaced by Bert Sutherland; the brother of Ivan Sutherland. Bert left BBN in 1975 and also went to Xerox Park. There were several other people at BBN who also went to Xerox Park. Anyway, about this time, I moved back into Division 4 offices.

Pelkey: How much of your time were you spending on ICCC?

Kahn: I spent essentially a hundred percent of my time from late ‘71 to late ‘72 pulling this off and trying to help people figure out how to use the net. Because you had to understand at this point, while we had the net up, you couldn't do anything with it. That is, if you sat down at a terminal, and built a matrix of all the computers and their functions on the net – it was very sparse. Computer One, Two, all the way down to 28, or however many they were. "Here are the things that you could do with the net, like connect.”

Pelkey: Like remote login?

Kahn: No, get a herald. Connect usually means that you are protocol connected. Herald only means you got some characters back from the host. Login means you can actually now do a login. You run a program. If you tried to put X’s for functionality into this matrix, you might have a few X’s. So we really spent this year in time trying to make the network work.

Pelkey: And you had some people that worked for you during that period of time?

Kahn: I basically recruited about 50 people; they volunteered to help put this demonstration together. They were people like John Postel[11], Vint Cerf, Bob Metcalfe[12], and Dick Watson who is now with Livermore. I mean, you go through the names of who's involved in networking, and a lot of them were involved in that particular project. Al Vezza. He’s now associate director of the Lab for Computer Science at MIT. He was very helpful back then. So was Jack Haverty, who is now with the BBN Computer Communications Corporation. A lot of really first-rate people were involved. The demo was to be public and visible. You really wanted to make it all work, and that meant making the network work. And that meant not only moving the packets around, but getting them into the machines, and getting the machines to compute over the net. I got something like 40 manufacturers of terminals to loan terminals to us, which we actually physically installed in the ballroom of the Washington Hilton. In fact I just gave a talk at a recent conference here in Washington, I showed some slides from that particular—

Pelkey: You mentioned that. And an audiotape, I gather it was.

Kahn: And we also made a film of the early days of networking, which I'm going to put on videotape make available for educational purposes.

Pelkey: Good for you. I look forward to seeing that. Now coming back to who created document 1822?

Kahn: Well, that's another one of those interesting things. Basically, I wrote the 1822 document, but we issued it as an unauthored document, simple because it was intended to be a reference spec. It was just called “Specifications for the Interconnection of a Host and an IMP.” And you see, at the time—

Pelkey: The reason I ask you this question is because a lot of the issues you were dealing with are host-to-host, which was outside BBN's responsibility.

Kahn: Yes.

Pelkey: BBN had responsibility for the subnet.

---

[11] See Interview –
[12] See Interview – Robert (Bob) Metcalfe

Kahn: Yes, but let me tell you that the 1822 document really got written--

Pelkey: Which was when?

Kahn: --the relative timing of things.

Pelkey: Fall of '69?

Kahn: I would say it was written sometime between April and June of '69. Because none of the other researchers could understand how to build an interface to an IMP until we told them how to interface to it. So they needed it well before the first delivery of any IMP. This was a complex interchange here, because in some sense, while I was trying to describe what the hardware was doing, the real issues were at the level of what the software would do. And the issue of what the software would do was a function of what the architecture and design specified. Well, at the time, Bill Crowther was busy with other parts of the software, and I couldn't wait for Bill to figure out how to write the document -- chicken and egg problem. So what I did was I wrote the document, and—

Pelkey: They designed to that, too.

Kahn: --then, well, fended off all of the issues and concerns that came from these quarters as to whether that was the best way to do it or not. Because we were trying to decide fundamental issues about how programs would work in some of these structures, and hadn't gotten around to it yet as a group.

Pelkey: That was a critical document.

Kahn: Yes. And we iterated on it. I mean, I wrote it, and then various people objected to parts of it, and so we changed parts of it to make them feel more comfortable. Other parts they decided to live with, because they didn't have time to go back and rewrite it. You know, that's sort of how things work. Now parts of what went into that document were based on where the software was at the time. If 34 percent of it was done, I'd take that 34 percent and invent the other 66 percent, or whatever it was. There was an interesting interchange going on here, but, you know, fundamentally we ended up with a document that everybody said they could live with. And that's what went out. And part of the reason it was unauthored is because when you're finished with a process like that, it's hard to sort out whose ideas are which. If you asked whose pen was on the paper, it was clearly mine, but I was just as happy to have that be an unauthored document. The question has never come up, by the way, in all the years that--I don't think there's anybody in the country that could answer that any more. Including the people at BBN.

Pelkey: And the Network Working Group came up as a consequence of that document. And Steve Crocker[13] at UCLA, I'm led to believe, said, "Wait a minute. We got to start talking about this host-to-host stuff."

[13] See Interview – Steve Crocker

Kahn: [long pause] I'm hesitating now because I can't remember the timing accurately to the year. Vint Cert could probably have told us that. Let me tell you what I remember. The Network Working Group had gotten formed at the request of ARPA. And I don't remember whether it was Larry Roberts or Barry Wessler (who worked for Larry at the time and still does). But they decided that somebody needed to worry about the problems of protocols from computer-to-computer: and Steve Crocker, who was at UCLA at the time, got the nod to lead the Network Working Group. In fact, Vint Cerf was at UCLA at the time, too. I'm not clear that Jon was at UCLA at the time.

Pelkey: According to Jon Postel, he was.

Kahn: He worked for Dave Farber[14] at Irvine on his PhD. He may have gone there briefly. That's right, he was there for a very short period.

Pelkey: And UCLA got these folks because of Kleinrock. It was really Howard Frank[15], you and Kleinrock, the three key guys underneath Roberts, it sounds like.

Kahn: Let’s see. Larry was worried about the grand design.

Pelkey: Right. The application-to-application level

Kahn: Yes, but I mean also that there should be IMPs that should be connected. But he never went one level lower: if we took the airplane analogy, Larry would not have been the builder of the airplane. He would have been the guy who said, "We need airplanes that can fly and we have got to have airports so let's build four airports and we need an air traffic control system, so let's get somebody to do that." But never getting into the details of any one of them, although I suspect he could have. Larry is a super web technologist.

Pelkey: In some sense, he was really the customer.

Kahn: That's right, that's right. What I really did was worry about how to make the--and let me call it the subnet, because that's really what it was, how to make that happen. Now you had a guy Howie Frank, and, you know--here, it's very easy. You draw some things, you draw some lines between them, but you really have a complicated system of hundreds or dozens of these locations. How do you hook them up? Because one of the other things you didn't want to do were string lines between every pair of computers. Because—

Pelkey: It's an n-squared problem.

Kahn: It's an n-squared problem. So Howie was the guy who took that problem and formulated it into a topological design paradigm. So he was really interested in the topology of a net and all of the techniques that you'd use to design it. And his principal

[14] See Interview – Dave Farber
[15] See Interview – Howard Frank

tool was simulation. But you had to know what you were simulating. I mean his simulations were not at the level of how the net would actually work. They were simulations of the properties of the net. Like what was the--for a Monte Carlo test--what's the reliability of this configuration, or how much throughput can you get. He'd do flow analyses. So he'd be able to, you know, produce--and you could get this kind of—

Pelkey: Throughput?

Kahn: This kind of a delay--throughput relationship. He'd do various analyses. Sometimes he'd use formulas that Len Kleinrock would generate. Now Len Kleinrock’s work was really analytical. Len would produce formulas for the behaviors of network parameters. You know, like the throughput was equal to the sum of the following terms. He'd use approximation techniques to arrive at a lot of interesting analytical results. He wrote an early book about networking.

Paul Baran had some of the early conceptual ideas. Kleinrock's analysis work was also done starting in the early '60s. He and Baran were in fact contemporaries. I believe Donald Davies in England was working on packet switching in the mid 1960’s. He had a lot of the ideas, but it was only one node he ever built.

Now there many people at BBN working with me on the Arpanet, and I've given you some of the key names. And then you have the people who weren't directly involved, but were extremely important, you know, like Licklider. In some sense Lick's ideas fueled all the early interest on this. Bob Taylor also did the same. Bob wasn't involved in building any piece of it or designing it, but in terms of a moving force and a conceptualizer, he played a very significant role. He hired Larry, for example.

Pelkey: Good decision.

Kahn: I think it was a great decision. Larry was the one who decided to give BBN the job of actually building it. I think that was a pretty good decision, too. Howie, Lenny and I actually wrote a paper on network design[16]. I don't know if you ever read it.

Pelkey: No.

Kahn: I think I have an extra copy for you. Anyway, Howie wrote this marvelous paragraph in the paper that says something like, "This paper has not been an easy paper to write." He said, "However, we have all benefitted from the interaction." He was clearly trying to express the fact that we were three people who just sat down and wrote this paper coming from completely different perspectives. My work had not involved any significant topology or simulation work. And I hadn't been doing any theoretical analysis. I was just trying to figure out how to make the network happen. Whereas Kleinrock had many insights about how the network ought to work, how many buffers were needed—much of that came out of his analytical work. Whether the results had any basis in practical design or not, you had to run some tests to find out what the real coupling was. Likewise, Howie Frank had been doing quite a bit of work on topology

---

[16] “Computer Communication Network Design- Experience with Theory and Practice,” Howard Frank, Robert Kahn, Leonard Kleinrock, Spring Joint Computer Conference, 1972

and simulation, but he was not involved in building the network equipment. So when we got together, we often found that we had come to the same conclusions for totally different reasons. We'd say things like, you know, "We don't really know how to build nets of more than about 60 nodes. Well, why is it that we don't know how to build nets of more than 60 nodes?". Howie would say, "Well, the simulation tables become bogged down." I'd say, "That's not the reason. The reason is you can't pass much traffic on the net if the lines are all filled with routing information." Lenny would say, "No. The reason is because the denominator of this equation goes to zero at that point.” And so we were trying to deal with wholly different conceptual formulations of the same problem, and putting them all together. So we'd resolve things as follows: "In simulation, here are the design issues. In analysis, these are the fundamental topics. And in the practical realities of building nets, these are the issues."

Pelkey: How often did you sit down and meet?

Kahn: Oh, I don't remember. Not often.

Tape side ends

Kahn: It was an interesting paper to actually put together and it involved very constructive and productive interactions: but they were somewhat difficult at times.

Pelkey: One other thing before we move on. Do you remember Bob Metcalfe coming to see you when he was a graduate student to discuss interconnecting the IMP issues?

Kahn: Mm-hmm. Only vaguely; that was back in when he was still at Harvard, I think?

Pelkey: He was at Harvard. However, MIT was going to let him be responsible for the project of connecting a MIT host to an IMP.

Kahn: Yes. Those were very interesting days back then.

Pelkey: Let's come back to ICCC, because that was one of the most important events during this timeframe.

Kahn: Well that's what made the Arpanet become viable. There are very few people who understand the importance of what happened there, because it wasn't a very glamorous activity. It involved getting vendors to participate, bringing people together and making all the network resources work. It was figuring out how to pull this demonstration off and then actually doing it.

The other thing that was rather interesting is that it was kind of a tour de force to actually make it happen. I forget which day of the week that the conference started, but we didn't actually have access to the room until something like 7:00 a.m. of the day it had to work. Now you have to imagine what it's like building a whole computer installation when you have something like five or six hours to actually do it. People were going to show up in the room to use the net that afternoon. And so we had to lay the false floor, and we had to do all the cabling, and we had to get an IMP installed, and we had to get the long-distance telephone lines, get all the terminals installed and do that all in a compact space of six hours.

Pelkey: It was a TIP that was installed, right?

Kahn: Yes. It was the TIP.

Pelkey: Now, the TIP wasn't originally designated in the spec?

Kahn: It was a subsequent addition, that's correct. The early spec called only for IMPs.

Pelkey: And the IMPs were only to be one node, i.e. one computer?

Kahn: That's right. That was another change that occurred later. The early spec called for one host connection per IMP. Very quickly, we found out that each site on the net had at least two and sometimes many more computers that they wanted to be able to hook up.

Pelkey: Did that cause problems for you?

Kahn: No, that was a relatively trivial change. The main issue there was how to deal with the distances, because-- I'll show you what the nature of the problem was. See Exhibit XX.

Pelkey: Historically, it turns out that intra-IMP traffic was a significant amount of traffic. It was in fact the first local area network.

Kahn: That's right. It hadn't quite been—

Pelkey: And no one really conceptualized that when you began, so I understand?

Kahn: I think that's probably accurate. The specification called for the host to be fairly close to the IMP. I don't remember what the number was, but let's say 50 feet. Some number like that. Well, at 50 feet, you don't get electrical differentials and potentials that are too big. The next computer might have been 1,000 feet away. What you should have done is put another IMP over there, but that was too expensive. Well, we ended up with three solutions. One was something called the local host, which is what BBN Report 1822[17] really specified. There was something called the distant host, which was really 1822 plus--let me just call it drive circuitry: twisted pairs, differential signaling, etc., the things that let you hook up over a longer distance. I think it went a few thousand feet, but I don't remember the numbers.

Pelkey: Was there a separate IMP next to these computers as well?

---

[17] BBN Report No. 1822, Interface Message Processor, Specifications for the interconnection of a Host and IMP, Bolt Beranek and Newman Inc.

Kahn: No.

Pelkey: It was just the electrical characteristics of the line that had to be ...

Kahn: It's just a way to make it work over longer distances. And then somewhat later--I don't remember exactly how much later--there was something called the Very Distant Host (VDH). This was a communications protocol, plus parts of 1822. It involved acknowledgements across the lines. In fact, you could have done VDH across the country if you wanted, but it would be terribly inefficient. It was designed to be efficient for distances of a few miles.

Pelkey: And you went to two bits so you could support four computers on one IMP.

Kahn: That's because we had an eight-bit field to designate the destination. And we decided that we would use six of them for the IMP and two of them for the host.

Pelkey: A practical solution?

Kahn: The pragmatics, right. We could address four hosts, but we couldn't have more than 64 IMPs, because with six bits, that's all you could address in terms of IMPs. Well later it changed, and we got to a bigger address space, but that caused quite a bit of change to the software.

Pelkey: How did the TIP come about, and was that a problem?

Kahn: The TIP was a big deal, because now suddenly here you had an IMP, and the question was, how could you hook up these terminals to it. And in fact what we did back then was a mistake necessitated by the economics of the situation. The first IMPs used 8k and then 16K of memory. But you could put 32K of memory in the machine. So we used the bottom 16K for the IMP, and we used the top 16K for the TIP. The IMP could handle four hosts, and the TIP was host 3.

Pelkey: Okay.

Kahn: Maybe it was zero, I forget what the order was, I think it was host 3 always. But the interconnection was made in software rather than hardware. At the front end of it was a multiplexer that would allow you to take all these terminals and MUX them into memory. So that's really what the TIP was. It was a bunch of software that got written in 16K of memory, plus this multiplexer which we called, at the time, the multi-line controller, also designed by Severo Ornstein.

Pelkey: And how many terminals could you connect? Do you recall?

Kahn: About sixty-four. I think it may have been minus one or two. I think 63 terminals were actually what you could connect, because one of the ports was a diagnostic or control port or something like that.

Pelkey: Okay. Do you recall when you did the TIP?

Kahn: During 1970-1971 I believe, but the development continued beyond that.

Pelkey: And that was because there were individuals who wanted access to the Arpanet?

Kahn: They wanted to get on the net and they didn't have their own machine. Remember, the original design of the net was supposed to allow people who had these machines to get on so they could interact with other machines. And then pretty soon we realized that there were a lot of people out there that just didn't have any machines. That's what made the Arpanet take off because people wanted access to the net. The people with the big machines had all the computational resources they needed locally. The users with the TIPs were absolutely dependent on the net to get to their computational resources.

Pelkey: Do you recall how electronic mail came about?

Kahn: That was another big surprise, electronic mail.

Pelkey: Could you say, in retrospect, that Arpanet originally was supposed to be an experiment, but because of electronic mail it just became such an embedded part of the—

Kahn: Oh, I mean this was a complex issue of funding and sociology. I mean it was just a very important research tool back in those days. And I think that—

Pelkey: It was terminal traffic, i.e. E-mail traffic, that was really the dominant traffic in those days.

Kahn: There was a lot of early experimentation, moving graphics back and forth. I mean, Danny Cohen did some very interesting early graphics experiments between Harvard and MIT; aircraft carrier landings I believe. There was quite a bit of experimentation in moving files between machines. But in terms of total number of actions people took, there was probably more electronic mail activity, even though some of the other activities probably transmitted more bits. E-mail started as a hack, I think. My recollection is that it was Ray Tomlinson of BBN who decided to try and see if he could build a little program that would send a message to another machine and, you know, be able to retrieve it and read it coming back. His transmission program was something called SNDMSG, I think, that's what he called it. It enabled you to say who the message was going to, and it automatically filled in who it was coming from. You'd give it a subject and type in your text and send it. He created another matching program that was called--this is really going back to the early days--something called Readmail. There was one file to which any mail that came in to you was appended. Message.txt I think it was called. And so this program would just go into the message.txt file and print out everything that was there. Or print out everything that was there that was unseen. The trouble was if you had 10 messages and you turned on the thing, out it all came. Larry Roberts decided to do a little hack that you referred to earlier. I think it was called RD. It was really a TECO macro that enabled you to get a summary of the file first. It would tell you there were 16 messages and you'd say, "Well, I want to look at the sixth one." So it counted down to the sixth one, and printed that out. Or you could select the second one. You could delete it or file it or--really a neat little hack. For several years, that's what people used. I guess that some of the machines on the Arpanet probably still have the RD program, but I suspect the only one who might still use it is Larry.

Pelkey: Was E-mail part of the demonstration at ICCC? I know there were games, and...

Kahn: I'm absolutely sure it was. I wish I had my archives, I'm sure I have a scenario book.

Pelkey: I do.

Kahn: Do you have a copy of it?

Pelkey: Yes.

Kahn: Well, take a look at it, you ought to be able to find it.

Pelkey: Yes.

Kahn: Now, Bob Metcalfe put the scenario book together. And he meticulously went through and tested every single one of them and we debated which ones to include and which ones not to. Bob's did a nice job on it.

Pelkey: Was there a lot of contact between the two of you at that point in time, or did he have quite a lot of liberty on that?

Kahn: Oh, he had quite a lot of liberty, but there was also quite a lot of contact. Fundamentally, there was so much going on that you couldn't, you know, attend to every last detail personally.

Pelkey: How many people came to ICCC and saw the demo would you guess?

Kahn: I don't know. I would guess there were probably a thousand people that must have come by there.

Pelkey: Do you recall an AT&T group coming through?

Kahn: I think so.

Pelkey: Do you remember anything about it?

Kahn: No, were you there? Were you in that group?

Pelkey: Oh, no, no, no. I wish I had been, but I wasn't.

Kahn: The only thing that I can recall about it is that they were the grimmest looking group in the whole bunch, I mean, almost like they didn't want to believe what they were seeing. That demonstration was the demonstration that convinced people that this technology was real. All during this period, nothing that you could say would convince somebody—

Pelkey: Right. Because there were articles that were written about the Arpanet at the time, but it wasn't real until that conference.

Kahn: I think that conference had the effect of suddenly making people have to take this as a serious technology. And it really turned things around at that time. In the European community, I think, it really gave it some substance; also within the domestic industry. I mean, because there were people from all over the globe in attendance. This was an international conference, and it suddenly solidified the view that: “holy smoke, this stuff really does work.” There may still have been a lot of latent concern, at that point in time, that this was not an economical technology, or sooner or later something would go wrong, or that you could do it better some other way. But they couldn't deny the fact that it worked, because you could sit down at a terminal, you could connect to any one of the machines on the net, and you could get into the machine, you could run a program and then move to another one. You could analyze whether there were alternatives that were more cost effective or whether you really wanted to use these big, clunky mini-computers for switches, I mean micros didn't exist at that time.

Pelkey: That morning when you said you got the room at 7:00 a.m., and you had five or six hours to get it all together, that must have been a harrowing time for the group and for you.

Kahn: No, it worked like a charm. It was just all ready. And we just got in there, you know, put everything down, plugged it in—

Pelkey: How many of you were in the room responsible for putting it all together?

Kahn: I’m not sure exactly. There were two stages to it. One stage was what I want to call the set up. I mean there was a whole big part that was before that which was the planning. And that planning went on for almost a year.

Pelkey: And was there anyone other than you who was really in charge of planning?

Kahn: Al Vezza from MIT worked very closely with me on all aspects of the demonstration. He was super-helpful with all kinds of things. He really was an alter ego for this particular demonstration. The actual set up of the equipment was under Al's direction. He brought in people to help, and just made that part happen. That involved getting the terminals installed, working the union problems, checking out each terminal to make sure it worked, setting up the curtains, the false floors, the cabling, etc. Then there were the actual demonstrations, or demo/walkthroughs. A whole bunch of people helped make it a success. I think we called them facilitators. The TIP was right in the middle of the ballroom, and it was on a raised floor, and we had ramps all around it to get people in and out and to handle the cabling. It was structured so that people could walk through the ballroom and see all the demos. There were booths along the side as well as in the center.

We had all the AT&T equipment in the back. Each of the booths had somebody assigned to them and then we had people responsible for regions of the ballroom like Vint Cerf, Bob Metcalfe, or Jon Postel. They would kind of roam around and just help out wherever they were needed. If there were a technical problem in any area, they would go and attend to the problem, fix it, and if somebody were having trouble they would assist. I would guess there were at least 30 people doing that. We also had demos outside the ballroom, staff at the doors, and we had a separate room where we were showing the movie.

Pelkey: Quite a production.

Kahn: It was quite a production, that's right. It took almost a year just to make it happen.

Pelkey: I gather there was only one down during the three days.

Kahn: I think the TIP did go down once for a short period of time.

Pelkey: Ten or twenty seconds, so I was told..

Kahn: Some very small period. I think we were just fantastically fortunate, because the whole exhibition was hooked up with a connection from AT&T at 50 kilobits/second that we brought in.

Pelkey: Oh, I thought it was two lines.

Kahn: It would have to be two, because if it was only one line, it would have gotten into trouble maintaining connectivity. The error rates on these lines were such that if we had only one of them, there would have been a sufficient amount of service interruption. Every time you get long error bursts, the line would go down, it would stay down for a while, then go back up again. And every time it would happen, the site would get disconnected. All the connections would be broken.

Pelkey: At nights, did the group get together and have dinners?

Kahn: Not as a group, because you have to understand—

Pelkey: Where did the joy come from?

Kahn: You had to understand the sociology of this whole thing. For the first couple of days, when it was up and running before the conference started, all these network experts were trying to get their demos to work. When you've got a hundred or so of the best computer people in the country, and they're all in one room, any questions that you want to ask can probably be answered by somebody in that room. People weren't leaving the room until well after midnight, every night. I mean, nobody wanted to leave that room. There was no issue about going out as a group for dinner.

Pelkey: So even after the demos were done, which I presume ended about 6:00 or so—

Kahn: Well, remember they were hands-on demos by the attendees in most cases. They weren't generally people demonstrating their own programs. There were the Metcalfe scenario books. People would sit down at the terminals, and they'd pick a scenario, and they'd do it. It would sort of be like you playing—

Pelkey: Yes, yes, I've got you.

Kahn: --Pong, in the early 1970’s...

Pelkey: So they didn't end at six--traffic didn't stop at 6:00. Did traffic stop when there was no one to take them through?

Kahn: Well, people attending the conferences generally left around 6:00 or 7:00. But the people who were involved in the demos didn't leave. They were there all the time. I mean if we didn't decide to shut the doors and lock it up, they might have stayed there all night for a whole week. And it was just--

Pelkey: It must have been an incredibly intense, exciting period with all these people to come--interacting in the same room—

Kahn: It was very exciting.

Pelkey: Were there any incidents during that period that would highlight that kind of excitement or energy level or...?

Kahn: No, just that there's where you wanted to be.

Pelkey: So, when it ended, what happened?

Kahn: Just tore it down.

Pelkey: Just tore it down and went on? You're engineers and scientists!

Kahn: This was happening literally at the boundary between my leaving BBN and starting at DARPA. I think my official starting date at DARPA was the 29th of October.

Pelkey: Isn't that interesting? I believe it was Vint Cerf's last day before going on to Stanford.

Kahn: Anyway, literally that was just before I went to DARPA.

Pelkey: Okay, so, now, unless I haven't asked you something along those lines... The next step, as I understand, was when in'72, '73, you were asked to start thinking about interconnecting other networks.

Kahn: Let me back up one more time before we do that, because I want to tell you about Telenet. Now back in the '69, '70 timeframe, I had been pushing the BBN management to consider getting involved in packet networking on a more commercial basis. The technology was available, it was viable, and there was a possibility that DARPA might let somebody run the Arpanet or BBN could build one of its own. There was just no interest in proceeding with this at BBN as near as I could tell. A small committee of people was set up to look at it, they did, and they decided that at most a private network business might emerge but that was too small to be interesting, and the idea folded. And it always seemed to me during that whole period that it was a very interesting idea to keep pursuing. In 1971, I went back to the BBN management and proposed the idea in a slightly different context. I got no reaction this time. I came to the conclusion that despite the fact that this was an interesting technology that BBN was not the kind of place that was likely to choose to capitalize on it. Of course that changed dramatically in the 1970's.

Now during 1971-1972, Larry Roberts had invited me to come to DARPA. I had a lot of trouble convincing myself that that was a good thing for me to do. But, by the spring of 1972, I had come to the conclusion that I really ought to go do that and told Larry shortly thereafter. Well, in the '71 timeframe, Frank Heart, who was trying to figure out how he could capitalize on the IMP development, had hired a guy from the Pentagon named Lee Talbert. Lee Talbert turned out to be interested in just sort of doing this on his own terms. He had been at BBN for perhaps six months. I don't know if he tried to convince BBN to get into networking or not. One day he just quit BBN, and took with him several of the key people. None of them intimately knew the network details, but they were fairly knowledgeable about them. And he set up a little company called Packet Communications Incorporated (PCI). I don't know if you heard of this part of the history.

Pelkey: No.

Kahn: PCI was the name of Talbert's company and that was set up in 1972, and I believe it was May or June of '72, because the timing was unbelievably exquisite. A week or two before he quit I had finally decided that BBN would never get interested in commercializing the technology and I had decided to come down to DARPA. So Talbert left BBN. One day later, or certainly that week, Steve Levy is in my office, he says, "Bob, we've been looking at these memos that you've been sending up here, you know, indicating that you think this networking area has some potential. We've decided it's a good idea." And I said—

Pelkey: What was he doing at that point?

Kahn: He was sort of the corporate entrepreneur. Perhaps his title was VP for finance or something. But he was clearly a person who made new things happen.

Pelkey: And was he T-a-l-b-e-r-t?

Kahn: That’s how he spelled his name.

Pelkey: Talbert.

Kahn: Lee Talbert. I think BBN was responding exactly to this event. And I said to Steve, "Your timing couldn't have been any worse." I said, "If you had come to me a few weeks ago I would have jumped at the idea." I said, "But I've already made this commitment to go to DARPA and I think that's..." He said, "Look. Why don't you work with me on this thing, let's just sort of see where it goes, and maybe you'll change your mind." So I said, "Fine." And this was sort of the June timeframe. I didn't go to DARPA until the end of October. And all during this time, I was also in the final stages of planning the ICCC demo. Anyway, Steve and I put together the plan for Telenet at that point in time, and I was the one that in fact introduced him to the first two employees that we hired, namely, Stu Mathison and Phil Walker. Stu was at ADL and Phil was just finishing Georgetown Law School at the time. And one of the issues that I hadn't decided at that point was do I want to stay with Telenet and turn that into a business, or do I want to go to DARPA. And I finally said to Steve, "Look. I've already made this commitment. I think that's really what I wanted to do." And so Steve stayed on with Telenet as its first president. I think Telenet actually started in August of '72. That's when it was formally incorporated. And we spent quite a bit of time doing the planning.

Tape side ends

Kahn: --Roberts joined as the second president, so Steve Levy was the first president during this period.

Pelkey: Okay, so, September 30th?

Kahn: I think about the first of October, roughly the end of September.

Pelkey: Of '73?

Kahn: Of '73. Remember, I had joined DARPA in late ‘72.

Pelkey: Okay. And Larry was at DARPA—

Kahn: Larry was at DARPA, so I came in and Larry left and he became the second president of Telenet.

Pelkey: Do you recall why Larry left?

Kahn: I think he probably was looking for other things to do. I mean, DARPA is normally a place that turns over personnel quite frequently by design. I think Larry probably thought this was a tremendously exciting challenge to do, as I did. I just thought the DARPA thing for me was a much better choice at the time.

Pelkey: When did BBN spin out Telenet?

Kahn: Well, it was a subsidiary for a while, and they weren't making any money on it. Larry was very interested in the technology he kept pursuing more and more interesting things. Telenet became operational something like '75 or '76. But it continued to lose money, and BBN was anxious to get out from under at that point. I think they sold it to GTE something like '78 or ‘79. And Larry agreed to stay on with it for some period of time. In ’82, Larry went to DHL.

Pelkey: Where he is today.

Kahn: Well, he was actually president of DHL and he set up a subsidiary called Net Express. He's now chairman of the board of Net Express. That's what happened with him.

Pelkey: Do you recall, during this period of time when Telenet was going to be set up, that there were other people who wanted access to the technology, other than just PCI? Like NSS?

Kahn: Network Systems?

Pelkey: Yes, It was a timesharing system. There were some others, but for some reasons, BBN was relatively unwilling to share—

Kahn: I remember PCI wanting access to the technology.

Pelkey: And that there might have been some conceptual ideas about, "Wait a minute. If we just share all this technology, everybody will put up their own network, and if all of these networks eventually get combined, it's a much more robust network than if everybody kind of goes up and does their own and they're not compatible." Were those kinds of discussions ever taking place?

Kahn: Well, I don't remember it in that generality. I do know that after Lee Talbert left, PCI went to BBN and asked for the software. Or maybe they took the software with them, for all I know. They could have just run off a copy. But somehow they had to get

BBN to agree, or there might have been a lawsuit, I suppose. PCI filed for their service--they were the first one to file with the FCC. And they got the first approval. But they subsequently went bankrupt. Perhaps Talbert was unwilling to share enough of the company with the venture capitalists. He must have been unrealistic in his assessments. He couldn't raise the money and so they finally went bankrupt. But during that period, he was pushing to get the technology, and ultimately having no luck in direct dealings with BBN. So there was no way Talbert was going to get that software from BBN. So he went to DARPA. And said, "Look. This was funded at government expense. I would like a copy." And the way that was finally resolved, DARPA finally went to BBN and said, "You've got to make this software available." BBN at that time thought there was more value in the 10-X software than there was in the Arpanet software, because they weren't thinking about the fact that other people would want nets. But they were thinking that lots of people might want to buy 10-X systems. So the quid pro quo was the government let them keep the rights to 10-X, but forced them to give the Arpanet software away. So PCI got it. The fact of the matter is it wasn't very useful to PCI.

Pelkey: Do you remember when that compromise was reached?

Kahn: In '73 or perhaps '74.

Pelkey: Sometime in 1973. And PCI went bankrupt in...?

Kahn: Probably in '74, '75.

Pelkey: Not long after.

Kahn: Yeah. They never got their financing. But the fact of the matter is having the Arpanet software was almost irrelevant to them. Because the concepts embodied in the code were everything. And even if you had the software, that didn't necessarily guarantee that you knew why the software was the way it was. There was nothing that required BBN to work with those guys. I mean, even if they gave them the software, they could say, "Fine. Here it is. It's all yours. Good bye." And that was the end of that discussion. So PCI fundamentally couldn't capitalize on it.

Pelkey: A value question. Should BBN at that time--given that it was funded publicly--should it have been more accommodative and supportive of PCI and anyone else who wanted to do a network?

Kahn: What position would you like me to take? Do you want me to put a BBN hat on? Or put a DARPA hat on? What hat would you like me to wear?

Pelkey: BBN's perspective.

Kahn: From the BBN perspective, they were trying to be as proprietary as they could of what they thought might be a potentially valuable product. Now that view was not shared unanimously throughout BBN, but the people who had a say in the decision all felt that this was potentially something valuable. Therefore, let's keep it to ourselves.

Pelkey: Because of some network issues, the Arpanet code and design never really did get out into the public as much as the host-to-host code did. BBN did keep that close to their vest.

Kahn: Well, it depends on what you mean by that. The software was available to anybody who wanted it. How much time should BBN spent trying to explain what it did? I mean, I was writing documents, but unless you went through the code literally line by line and explained how it all worked, it could be a lifetime proposition. In fact, in that early Arpanet paper that we wrote, we went through a lot of those details: how the software was structured, why it was structured the way it was, all those issues. I mean there was probably more written about that network than any other network in history, including the Bell System.

Pelkey: Gotcha.

Kahn: So that was Telenet. And that was really the precursor of a lot of other commercial developments.

Pelkey: Now Tymshare developed separately.

Kahn: Tymshare was a provider of timesharing services and software. Users originally dialed up through the telephone system. Tymshare basically concluded, some time in the 1960s, I don't know when exactly, that it would be more cost-effective if they built a network and let users dial in locally. They then would have shared lines across the country and that would be more cost-effective for the company and for the customers. So the TYMNET network was built around that premise. It had a lot of properties that were similar to Arpanet, except that the whole goal was to get this customer base to a set of centralized machines in California.

Pelkey: Right. As opposed to distributed machines.

Kahn: There was no notion really of distributed computing. The routing was fixed. I mean, when you got logged into the TYMNET system, they would create a path and all your characters would go over that path. With Arpanet, whatever the routing operator said for a packet, that's what happened for that packet. So that packets might arrive out of order, but that never happened with TYMNET. Furthermore, TYMNET was organized around character interactions. An Arpanet packet could have anything in it. It would have a little header on the top that said where the whole packet was supposed to go. In the case of TYMNET, every 16-bit word was broken into two 8-bit bytes, one of which had a character and the other of which had an identifier for that character, which would tell something to the next switch about what to do with that character, what line to put it out on. In a single TYMNET packet, one might go one place, and the next character go to a different place. So each of these packets were split apart at each node and reconstituted, depending on how many lines were coming in, so that all the characters that had to go out that line were repackaged with new IDs. The IDs changed because they were node and line-specific IDs. This approach couldn't handle a whole lot of throughput, and the user was basically lucky to get 300 bps, and it was totally focused around terminal to computer interactions at low speed. Now TYMNET would probably, on defendable grounds, say that it was a form of packet switching, but I would have called it low-speed character switching, which is really what it was. They later changed and made it more Arpanet-like in the mid-1970s, and in fact later, it became a carrier. So that instead of just providing access to their own machines, it also made access available to others. I think Telenet was actively involved in trying to get Tymshare to open up and obey the common carriage rules. Because Telenet, at that time, thought and probably were obliged to behave under the rules of common carriage that the FCC set down. But since TYMNET was offering a similar kind of service as a private thing, they weren't. And ultimately TYMNET filed with the FCC and came under the same rules and regulations. So that was a bit of a hassle.

Pelkey: There was a thing called NCP that existed when you went to DARPA on the network, right? Network Control Protocol?

Kahn: Yes. The Network Working Group got set up to deal with these protocol issues. Now my recollection is that that group actually started in the '69 timeframe, but I'm not positive of that.

Pelkey: I have the first RFC.

Kahn: I'll bet if you look at those numbers, they probably—

Pelkey: I think April or May of '69 was the first one.

Kahn: That’s about right. Barry Wessler really got that whole thing going from within DARPA. The working group designed a protocol to move packets from one host to another. And it was called NCP. It was really designed collectively by a lot of different groups, all working together. DARPA was part of the process. And Crocker spearheaded it. And that was another sociological phenomenon. Because all the machines were different, the operating systems were different, and this whole notion of protocols was kind of a new invention at the time. And NCP was first implemented in the—

Pelkey: Late '69 timeframe?

Kahn: --late '69 timeframe, because when I went out to do the Arpanet tests on the four node net, I'm sure that we did some interactions that were over the net with multiple machines.

Pelkey: Did those engineering issues attract you at all or fascinate you, or were you just so consumed with subnet issues and so on that—

Kahn: They were quite interesting. Several people were helping me run the test at UCLA. Dave Walden was there to help our planning and executing the subnet experiments. I might sit down and say, "Dave, here's an experiment I’d like to run." We’d lay it out together. If I would say "I'd like to patch that machine so every time a packet goes by, we can see it and I want to blink on that light." He'd say, "I got it." He'd figure out a four-line patch and he'd put it in. Vint Cerf was also very actively involved. Vint was a graduate student at UCLA at the time. He wrote a measurement program that ran on the UCLA machine that let us collect all the statistics of what was going on. Len Kleinrock actually had a contract to make something called the Network Measurement Center happen.

Pelkey: Right. From ARPA, which was really to be measurement of computer-computer statistics, and then got modified to be for computer-to-network.

Kahn: Something like that. I don't actually remember all the details. But anyway, Vint was a key guy, and he wrote the software for the UCLA machine that we used for the measurements. And there was quite a bit of iteration with the measurement software that allowed us to check out the IMP software. We made thousands of modifications over the next several years to the IMP software based on the tests we were doing on a real net, plus the tests that we were doing back at BBN—back room experiments. The NCP was the protocol for host-host communication, and there were specs for other kinds of protocols. The telnet protocol handles terminal to computer interaction. FTP handles file transfers. Some interacted with the IBM protocols, although there weren't many IBM machines on the net.

Pelkey: So the file transfer and remote log in were critical ones.

Kahn: Well, let's see. FTP, that was pretty crucial, Telnet, NCP, something called ICP, Initial Connection Protocol (that was sort of almost a part of NCP). You had Remote Job Entry—

Pelkey: By the way, do you know where the Telnet name came from?

Kahn: I don't know who actually generated it. Maybe Crocker and Postel. It stood for something like Teletype network. Teletypes were still popular back then. In fact, Telenet was just a perturbation of Telnet. I think Steve Levy actually came up with Telenet as a name, or he had some firm do it. A lot of these names just sort of emerged.

Pelkey: But then you got to DARPA, you wanted to do something about all this, right?

Kahn: I basically hadn't been too involved in this part of it, because I had my hands full just with the communications part of it. I had to use all the protocols in testing, and I knew what they were about. From a communication subnet perspective, host protocols are a higher-level issue. At times, it got deep into the operating system side of things. I had a feeling that we could have done a little better with the early protocol designs, but basically it didn't seem to be a terribly high priority to optimize them at the time. Many of the Telnet issues were terminal-specific. And I knew that was going to evolve over time, and so we might as well let it evolve. Now remember, to develop a protocol, you had to get dozens of sites to agree. And operating systems were all different. It was just too important that something happened, not that it happened perfectly.

Pelkey: This is NCP and ICP.

Kahn: ICP was just the Initial Connection Protocol. And NCP was what happened. Now let me tell you a little bit about NCP, just to show you why it had to be replaced eventually. You remember, the Arpanet was built around the notion of a message that could have up to eight packets. A message could be up to 8,192 bits long, and you could break that into at most eight packets. The packets would all get reassembled at the other end, put back into the whole message, and delivered as a single chunk. So that was the idea of how the network was going to work. There was not a whole lot of thinking about a pure datagram service. Somehow the idea was that hosts basically want to send big chunks and the network would have to break them up into packets. The hosts wouldn't want to deal with packets; they'd want to deal with the bigger messages. Which was why the original papers referred to message switching, because the host was sending messages to the IMP, and the IMP was breaking the messages into packets. A lot of people recognized that this was not such a good idea later on, but that's the way it started. We wanted to insulate the host from having to know about packets. On the other hand, you couldn't allow the host to send arbitrarily long messages, so we said, okay, they can send up to about 8,000 bits (which seemed big at the time) and the net will break it up into packets. A host could send one of these messages, and then it had to wait for something called the “Request for Next Message”(RFNM) before it could send the next one. This came back from the destination to the source. Now the problem with that strategy is the further away the destination is, the longer is takes for the RFNM to come back, and so the time between transmissions of the same amount of data gets longer and longer, and so the effective data rate goes down. So you couldn't sustain high throughput over long distances. That was problem number one. Problem number two was that NCP assumed the net was perfect. That is, the NCP protocol viewed the net as if it were a disk or a tape drive, in other words, some perfect peripheral out there that if you sent it a message, it would show up at the other end perfectly. So if any messages were lost you were in trouble.

Pelkey: Because the transmission wasn’t perfect.

Kahn: Didn't have a recovery mechanism. Problem number three had to do with addressing. The Arpanet was originally constructed as the pioneering network. You sent something in here, and it popped out there. The host specified at the beginning, in a couple of fields, essentially where the message was going to. And since only one message could be out at a time, it didn't even have to sequence them. There was also this “link field” which was intended for the host to use, but they could never quite figure out what to do with it. When a message arrived at the destination, you didn't have to tell the recipient computer who it was going to, it was going to that computer. By the time a message was delivered, the “to-field” got changed into the “from-field.” And the idea that you needed to know the “to-field” didn't exist, because some place in the machine, there would be state information about this connection. It was all in the tables of the NCP. The NCP built these tables, and when you'd open up this connection via the Initial Connection Protocol, it would assign you a “socket”.

Pelkey: The virtual circuit would be created as long as the circuit remained open, and presumably it was good.

Kahn: That's right. It was all in state information. Well, the thing that really caused me to take a harder look at this was the packet radio network. We were generating this packet radio technology all based on these little radios that had micros--the first systems application of microcomputers in the military. We're doing this in the ’72 and '73 timeframe.

Pelkey: But the real first use of the microcomputer was the calculator, right?

Kahn: I'm talking about a serious system application of it. I mean, people were using chips for a lot of different things. But even the HP pocket calculators I think first came out like seventy—

Pelkey: Two.

Kahn: Maybe late '72.

Pelkey: Now was your packet radio application influenced by Aloha?

Kahn: The Aloha network was a precursor. And that was funded by DARPA and AFOSR[18]. The Aloha system was to packet radio like the original Tymshare network was to Arpanet. That is, the Aloha system was a way for users at terminals to get their characters into a central computer system at the University of Hawaii. It was a centralized system: it was also a STAR; it didn't have multiple hops, it didn't have routing. I mean, if you were at home at a terminal, you could somehow broadcast on a radio, maybe you could or could not get it to the central computer.

Pelkey: That was Norm Abramson[19], right?

Kahn: Yes, it was probably his idea. But he and Frank Kuo jointly directed the effort.

Pelkey: Do you remember what year it was?

Kahn: I would say it was worked on largely between 1969 and '74. And many of the ideas really led to packet radio. It was predicated on the success that they had. But there was no way that you could turn the Aloha system, as it was, into a multi-node system.

[18] Air Force Office of Scientific Research
[19] See Interview – Norm Abramson

User radios transmitted on one frequency, and received on a different one. So now think about how you would talk to a neighbor down the street. He's transmitting on the same frequency that you're transmitting on, and you're receiving on the same frequency he's receiving on, so you can't talk to each other.

Pelkey: Right. You need a lot more frequencies.

Kahn: It's a different system design.

Pelkey: I think I understand.

Kahn: It was not designed as a network in the classical sense. But anyway, packet radio was a new problem in many dimensions, a lot harder than the Arpanet technology was from a technical point of view. It involved spread spectrum, microcomputers, mobile units, dynamic reorganization, all the things that you could take for granted in the Arpanet suddenly were open variables in packet radio. And so having now planned the first of these packet radios—

Pelkey: And what year was that?

Kahn: '72, '73.

Pelkey: That was when you first joined DARPA?

Kahn: That's right.

Pelkey: Did you start the program?

Kahn: The plans for moving into that area were in place. The program hadn't started. I mean, in some ways I was probably to that what Larry was to the Arpanet development. So the plans for doing that were already in place, and Taylor clearly hired Larry to do it. But they really hadn't started the program yet. I mean, I called the first meeting of the working group.

Pelkey: Now you needed a protocol to deal with lots of users sharing the common channel.

Kahn: I was concerned with getting many computers on a mobile radio net. You know, in 1973, the smallest computers you could get were minicomputers, and they weren't terribly powerful. And Arpanet had like 50 or 100 hosts at that time, so my idea was to connect the packet radio net to the Arpanet so that I could test the ability of the radio net with some real computers. A 10-X machine would take up a room. You weren't going to put it in a van and move it around. But you could put a user with a terminal in a van and try if the connectivity was there. So I needed a way to hook these two networks together. We came up with the idea of a gateway between the nets, simply because these nets were totally different. I mean, the packet radio net ran at 100 kilobits or 400 kilobits and the

Arpanet ran at line speeds of 50 kilobits. We had to find some way to get a computer on one net to talk to a machine on the other net. Well, it was clear that NCP had no chance of working in this situation.

Pelkey: Why?

Kahn: Well, just take example of the perfect net. I mean, when you're dealing with a radio system: static in the air can zap packets, or you can go into a tunnel, and packets can collide. Packets didn't collide in the Arpanet. So first of all, this perfect network assumption went away. Second, let's take addressing. A packet is sent from a computer through the Arpanet into a gateway to the packet radio net. The gateway says, "Gee. I know who it's from, it's from that computer on the Arpanet, but I have no idea who it’s going to." With no state information in the gateway the NCP addressing mechanism couldn't work. We also wanted to avoid logging into gateways all along the way. Even if the message reached its destination, the destination wouldn't know where to respond, because all he could do is send it back to the gateway, probably, and so it just—

Pelkey: So it completely broke down.

Kahn: This broke down, too. In fact, we never bothered to try it. Then finally, none of these other nets used RFNMs, so how would this computer ever get a RFNM. They were unique to Arpanet and it was generally understood that was a bad idea by then. Even if you could fake a RFNM somehow, it would either lose its original control function or it would be going through longer and longer paths every time, and the throughput would—

Pelkey: No throughput.

Kahn: Right. So all three of those things said that whole original NCP strategy just couldn't work. And that was part of my original intuition that this was not the right thing to do, but I couldn't articulate exactly why, because we didn't have a credible example, and I wasn't working on the problem. So anyway, with this in mind, Vint Cerf who was then at Stanford and I started to think about some of these problems. And Vint and I probably had talked a little bit from early '73 through early spring—

Tape side ends

Pelkey: So in early '73...

Kahn: I had some fairly clear ideas about how I thought the new protocol ought to work. And remember, Vint had been actively involved in making the original NCP happen. So it's again a match of two talents that were sort of necessary to deal with it. And together we sat down and formulated a protocol for network interconnection that we then wrote in a paper. And that paper was actually written from probably May through the summer. But by September '83, we actually had it documented and distributed in draft form at a meeting in Sussex, England; it was an International Network Working Group (INWG)[20] meeting. The IEEE published the paper in May 1974.

Pelkey: In your written notes?

Kahn: Yeah, here's the internetworking experiment, the major technical progress--groups we have--that's an interesting one--here's the thing I was saying about BBN, carried under-- --continued to advance more than any other individual with the possible exception of Roberts, himself. They couldn't quite make it any stronger than that. Oh, here's a--you can, you can tell how many were there. September of seventy--no, this was two months before the network. So you can count the number of nodes: One, two, three… 27, 28, and 29. What do you know?

Pelkey: You then wrote this joint paper?

Kahn: So then I wrote this joint paper, and that was basically called A Protocol for Packet Network Interconnection, or something like that.

Pelkey: I'll get that paper.

Kahn: And that paper basically described the Internet architecture. I mean, it laid it out conceptually. This was an architectural design. But somebody still needed to do a software design, and actually do the implementation. That's what Vint and his students then did. He and a group of his students at Stanford did the first implementation of TCP/IP. We described the process in a composite way, and later it was decomposed into a part that could go from gateway to gateway-- that was the part that was needed by the gateways (IP), and another part that was needed only at the ends (TCP). The local area network hadn't been invented at the time. And once you got into LANs, then you got the idea of connections from another gateway to a local area network. In that mode, you really wanted the gateway-to-gateway protocols separable from everything else. And in fact, you know, the Packet Radio Working Group that I put together included Bob Metcalfe. And partway through that development he came up to me, and he said, "Bob, I'm going to have to drop out of this working group because of a possible conflict that I have at Park." And that's when Ethernet started to happen, I believe.

Pelkey: Well actually that was XNS.

Kahn: XNS was much later.

Pelkey: I mean PUP.

Kahn: PUP was much later, too.

Pelkey: Oh, it was.

[20] At this time the INWG had been reorganized under the IFIP as Working Group 6.1 (WG 6.1) on Network Interconnection.

Kahn: Yes. In fact the TCP that we originally just called it a TCP. I'll get back to PUP in a second. It was developed by Vint Cerf and his students at Stanford, Peter Kirstein and his students at University College in London, and a group at BBN, led by Ray Tomlinson, the guy who did the original SNDMSG stuff, and some of his people at BBN. There may have been others. Well, anyway, those three did separate implementations on different machines. You really needed different people to do it to show the robustness. And then they went through something like four different major versions until it became the DoD standard in 1980.

Pelkey: Did you have anything to do with the DoD standard?

Kahn: We worked very hard to get it to become a DoD standard.

Pelkey: Was it the DCA that said, "Create a policy statement"? I am guessing in '83?

Kahn: Well I think it was originally created as a standard in 1980. And then it became mandatory for the DDN in—

Pelkey: January 1, 1983.

Kahn: 1983. Now you have to remember that Vint actually left Stanford in 1976 to come work at DARPA, and so he managed the evolution of this Internet technology while he was at DARPA. He left just before we went through the Internet transition on the Arpanet. He went to MCI to build the MCI Mail system. And he was really the designer of the MCI Mail system, working for a guy named Bob Hancharik, who was the president of TYMNET before that. It's really a small world. And then Vint came back to work for me shortly after I set up the Corporation for National Research Initiatives in June of '86.

Pelkey: June of '86?

Kahn: Okay, so, to put this in perspective, the packet radio work was really a follow on to the Aloha work. And while we were doing this, Bob Metcalfe started doing Ethernet. Ethernet was really the Aloha system on a wire. I mean, not exactly, but close. It went on a wire rather than over the air. Bob was superb in making something like that work simply and effectively. That was his brilliance. It was an engineering coup. But the key ideas were all generated in the Aloha project. But making it happen commercially was really the key, and he was the first guy to do that. Not only did he make Ethernet happen, he got the industry to agree on it as a sort of pseudo-standard. He sold DEC and Intel and eventually built a business around that. But it happened while we were doing packet radio, and that, of course, was a much more grandiose kind of idea. Now TCP/IP was really TCP that became TCP slash IP, was really generated because of the need for something that would work across the multiple nets. But remember the LANs didn't sort of show up until later.

Pelkey: I guess '74, '75 was when there were LANs, at least at PARC.

Kahn: Something like that. Xerox initially decided that they would participate in the TCP development. But the problem was we were trying to make this be something that would work across--

Pelkey: Multiple hosts?

Kahn: --multiple hosts and multiple organizations.

Pelkey: And they were worried about the implications?

Kahn: And it turned out to be, you know, a four-version, six-year standardization process to get there. So we didn't actually get something that was a quasi-standard until sometime around 1980. Meanwhile, they wanted something in the mid '70s, and they didn't want to be in series with all this consensus building that goes along with standardization.

Pelkey: Was there serious discussion with Bob about the paper at that point in time?

Kahn: He hadn't been involved, but he understood the importance of it. I mean, Metcalfe instantly understood the importance of what we had done, and probably would have contributed greatly to, you know, elaborating it. I'm sure he had dozens of good ideas, for example, the whole notion of local area nets.

Pelkey: That had not come into your—

Kahn: I mean, I'm sure it was in Bob's head at that time. It wasn't in our descriptive psyche. As we were describing the things, we weren't writing about local area nets in all we were doing.

Pelkey: But that was in his psyche.

Kahn: It was in his psyche, and it was subconsciously in ours. I mean because the protocols that we derived were perfectly good for that case. It was like the Arpanet worked for four hosts as opposed to one with a trivial modification because we had sort of done it right. But Bob would have contributed to the paper if we'd done that. We were happy with that paper as a tangible contribution, and we wanted to move on. And so, Xerox developed something called PUP, the PARC Universal Packet, or Protocol, or whatever it was, and that was their version of TCP/IP. And it was essentially the same thing with minor variations. I mean, they used sockets and we didn't, but to first order, the architectures were isomorphic. And PUP later led to—

Pelkey: XNS.

Kahn: They used PUP for quite a while. I don't know how long they used it, then they had something called the OIS protocols, the Office Information Systems protocols, and that led to XNS.

Pelkey: Now TCP/IP. Were you involved in the issue of creating Berkeley UNIX?

Kahn: Yes.

Pelkey: Because it is my understanding that it wasn't until the TCP/IP code got distributed with Berkeley UNIX that you really started getting TCP/IP out to the community. Rather, it was really the recoding of Berkeley UNIX and TCP/IP by Bill Joy that made become a standard in the commercial sense.

Kahn: Yes and no. Let me-- you see, all during the early part of the 1970s, there was a focus on-- I'm trying to think of the name of the operating system that ran on the DEC minicomputers back then.

Pelkey: RSTS or RSX-11?

Kahn: Something like that. There was also a system that some people in Santa Barbara had built called ELF, which was German for eleven. It ran on a PDP-11. At about the same time, UNIX had come out of Bell Labs.

Pelkey: But UNIX came out later, right? UNIX came after 10x?

Kahn: 10x was created around 1967 or '8, let me call it '68. I mean there were only a few timesharing systems back then, one Multix[21] at MIT, 10-X at BBN,…

Pelkey: And Multix was John McCarthy?

Kahn: Multix was [Fernando] Corbato at MIT. McCarthy was credited with the idea of timesharing, but I don't know if he ever built one himself. Multix and 10x were two of the most significant systems. There was also Project Genie at Berkeley, where they built a timesharing system for the SDS 940 that was very influential in the 10x development.

Pelkey: Now there was already commercial timesharing when these systems came out, right? An IBM environment sort of?

Kahn: IBM eventually came around to the timesharing idea, but they originally thought it was a bad idea-- probably for marketing reasons.

Pelkey: Okay. So these are the original timesharing operating systems.

Kahn: There was a timesharing system on the PDP-1, and there was something called CTSS at MIT, and... There were a number of precursors, but 10x and Multix were two

[21] MIT originally released the code as MULTICS – the father of UNIX.

big efforts, plus the SDS 940 work at Berkeley that developed a lot of the early ideas. And a lot of those people ended up at Xerox PARC. Butler Lampson, Peter Deutsch, Daniel Bobrow and people like that. ELF was the operating system that was originally used in the packet radio stations. In fact, all the early gateways used ELF. But UNIX began to emerge as the operating system of choice for minicomputers. AT&T was pushing it. Now along around 1979 came DEC with the VAX. And the thing about the VAX was that it had a 32-bit address space. So in fact you could run much larger programs on the VAX than you could run on a PDP 11 or other 16-bit minicomputers.

A large address space machine was very desirable, in the case of DARPA, for two reasons. One was our image understanding work, dealing with images; you encounter very large files and programs for working on them. The PDP-11 was just not sufficient. The second was VLSI design. We were just starting up serious work in both of those. We came to the conclusion, with a lot of input from the community, I might add, that what we really ought to do was go with the VAX, since it was the only appropriate machine around, and unfortunately, they weren't happy with VMS or UNIX. Neither one had all the features that they needed. So we gave a contract to Berkeley-- Bill Joy-- to take the UNIX system, put it on the VAX, and add all these other features, which he did. And then shortly thereafter, you know, we arranged for BBN to provide a TCP/IP implementation that would work with—

Pelkey: Bill's version of UNIX.

Kahn: Bill's version of UNIX. Remember, BBN had by that time five years' experience with internetting[22]. Berkeley had none. The problem was that Bill, being properly covetous of his design-- he was probably the only one who understood it fully—wanted an efficient implementation of TCP/IP for the BSD UNIX. BBN built it to interface in standard ways. Bill Joy just didn't feel like this was as efficient as he could do if he did it himself. And so Joy just rewrote it.

Pelkey: Was there some debate about that?

Kahn: There was no debate at all. He just unilaterally redid it. And all of a sudden, he got it out to everybody. I mean this was one of the classic standoffs. Because when Bill did it, he didn't know much about the Internet architectures. He was trying to make the BBN code work more efficiently. And so he would do things that they felt weren’t quite right, and so the BBN guys would resist. And he would say, "Theirs doesn't work well." The ultimate resolution was we went with the Berkeley code and the BBN code. The guy who was managing this at the time was Duane Adams, who is now at Carnegie-Mellon University. Any recipient of the Berkeley UNIX software could designate the BBN code if they wanted it. Initially there were more features in the BBN code than there were in the Berkeley code, except the latter worked faster and was more tightly integrated. In the final analysis, the Berkeley UNIX supported both sets of software. It supported the Berkeley code and the BBN code and it was up to the user to specify which they wanted. In the early 1980s DEC agreed to make available to the universities something like four

[22] To become known as internetworking

or five 11/750 VAXs and a file server for approximately $180K. A single VAX in those days used to cost several hundred thousand dollars. So suddenly, you were getting five VAXs for less than the price of one 11/780. And that helped to promote the Internet as well, because once you got the VAXs and the local area network, it made sense to install a gateway and get into the rest of the Internet. Before then, people couldn't afford -- I mean, by the late 1970s, even though the Ethernet had been propagated, the research community couldn't afford enough workstations. What could you get?

Pelkey: How about the Altos workstations?

Kahn: The Altos were not powerful enough.

Pelkey: So MIT's Laboratory for Computer Science (LCS) ordered the DEC VAX 750s that, if I understand correctly, had to be used as single user workstations.

Kahn: That was part of the deal that was negotiated with DEC.

Pelkey: I guess you convinced DEC that this was the trend and they ought to understand and create an environment for them, or was it just an economic issue? Was there more strategy?

Kahn: I don't know what DEC was thinking, exactly, but, basically Gordon Bell and Sam Fuller-- Sam worked for Gordon at the time—decided that if they made these VAXs available to the research community, people would explore interesting ways to do distributed computing on them, and that was going to be the wave of the future. DEC said, "Okay, we'll do that, provided that these machines are not used as stand-alone timesharing systems."

Pelkey: Around that same time, LCS’s Michael Dertouzos was debating how to put a local area network into the MIT environment. And Bob Metcalfe came back as a consultant and argued for Ethernet, and others, e.g. Jerry Saltzer, favored token ring. Were you involved in any of that?

Kahn: Yes, in a way. I think it's probably fair to say that there were more good systems ideas and technology coming out of Xerox PARC at that time than anywhere else. MIT was just sort of getting over the Multics experience, and beginning to work on distributed processing. Mike Dertouzos was doing a very good job reorienting the group. By the '80s, MIT was starting to pick up some steam and what they didn't want to do was use only Ethernet because that made them look like they were the follower. Instead they explored the pros and cons of the Ethernet and token ring technologies.

Pelkey: Gotcha. And there was already Chaosnet, which was a derivative of Ethernet in the AI lab.

Kahn: That's right. Mike ran the Lab for Computer Science and Pat Winston ran the AI Lab. They were not of like minds, either. The Chaosnet guys did their own thing (which wasn't Ethernet), and it subsequently vanished. But LCS wasn't about to use what the AI guys were using, either.

Pelkey: So the LCS chose not to use the local area networks of either the AI Lab and/or Xerox PARC. So what did they do?

Kahn: Dave Farber was at Irvine at the time, and he had an NSF-supported project to demonstrate a ring network. Farber is a really very interesting idea person who also likes to experiment with building things; I think he is much better with ideas than he is with implementations, although he's not bad with implementation. That's not his forte. Well he had just developed a system called the Distributed Computing System (DCS) that involved multiple computers on a ring network. Dave had gotten an offer to go to Delaware and MIT was contemplating what to do in local area nets. I remember having several interactions with them about what to do, and they decided to work on ring networks and to compare them with Ethernets. Dave Farber worked with Clark and Saltzer to import the ring network to MIT. They decided to bring in the ring technology and the Ethernet technology and do a cross-comparison between the two. It was a toss-up. They really couldn't determine whether one was better or not at those speeds. It was very clear that as the speed went up, eventually the Ethernet would degrade in performance. Because at higher and higher speeds the propagation delay on the net was getting larger compared to the time it took to put a packet on the net, and so the probability of getting collisions on the Ethernet went up as the speed went up. There were no collisions on the ring.

Pelkey: Some believed, around that time, that an issue of Ethernet was it did not handle fiber very well.

Kahn: I'm sure you could have made an Ethernet equivalent with two fibers, each propagating in opposite directions.

Pelkey: The collision problem has to be—

Kahn: The ring was more controlled. This would have required two fibers. The ring was better suited to high-speed transmission. The research project was the cross-comparisons. The cross-comparison turned out to be a wash-- six of one, half a dozen of another, but they ended up redeveloping and refining the ring technology that was commercialized by Proteon.

Pelkey: Proteon.

Kahn: Proteon called the system ProNET.

Pelkey: ProNET. Were you involved in that process at all?

Kahn: MIT decided to get Proteon to work with them and actually do the implementation. That's how the ring network started out there.

Pelkey: And then there was a Symposium in the Spring of '79 when there were a bunch of papers presented, and a public discussion, about Ethernet versus token ring. Does that ring a bell?

Kahn: I remember that there was such a Symposium, but I don't recall that I tracked it very closely.

Pelkey: And then IBM shortly thereafter picked token ring.

Kahn: I think what MIT was doing and what IBM did weren't exactly the same. But IBM entered the token ring business about that time. I don't remember the exact dates. You have to remember that from the DARPA perspective, the whole area of local area nets appeared largely as an industry initiative with some support from NSF. Xerox PARC with the Ethernet-- that was an industry initiative. The original support for the ring network at Irvine-- that was an NSF initiative. Bell Labs also did some interesting work on ring networks.

Pelkey: Can you elaborate?

Kahn: Well, in reality what MIT did was make the ring network real. That was a significant contribution. On the other hand, it's a research project to do the comparison, and it was useful to have somebody look at it. And it wasn't an expensive project anyway to do that, so, it didn't make much difference.

Pelkey: Well how does X.25 protocol fit into this process? Were you involved in that at all?

Kahn: My recollection is that Larry Roberts, at Telenet, spearheaded the X25 protocol efforts. He wanted to get international agreement on an interface standard.

Tape side ends

Kahn: Roberts got several groups together-- I don't remember all of them-- but it was Bell Canada, the U.K. was involved, and one other. They made X25 happen, and X25 was really the first several layers of the OSI protocols. There had been a paper written some years before by Vint Cerf, Hubert Zimmerman[23], Alex McKenzie and Roger Scantlebury. It was an attempt to describe an Internet protocol that they could all agree upon that was largely derived from TCP/IP. For a variety of reasons, they decided to make some changes. Instead of doing control of TCP/IP connections on an octet basis, one byte at a time, they decided to do it on a segment basis. They preferred larger entities for control. What they came up with was very close to TCP/IP but not quite identical. There was a lot of sentiment in the European community to develop something that was different from TCP/IP. This work eventually led to the OSI suite of protocols. The OSI suite was directly influenced by what we did on TCP/IP. I mean, TCP really translates

[23] See Interview – Hubert Zimmerman

into something now called TP4. OSI also incorporated an IP protocol that they didn't have originally. And that IP is now the same as our IP.

Pelkey: Okay. And there were some layers on top of that.

Kahn: There were seven layers in all.

Pelkey: Going back to the earlier years, I hear very little that came out of either the University of Utah or UCSB. Most of what came out from the network experiment was really out of UCLA versus either of those other two university nodes.

Kahn: Came out in what sense?

Pelkey: In the sense of influence or people who impacted things or...

Kahn: Well, out of Utah came people like Alan Kay. Alan Kay has been a major contributor - he created the Smalltalk language.

Pelkey: Right, but in terms of the network. In terms of networking, and Arpanet, and...

Kahn: That's probably true. Utah had more of a computer science orientation. Tony Hearn was working on REDUCE for symbolic manipulation of algebraic expressions. Tom Stockham was working on signal processing (applied visually). Ivan Sutherland was working on graphics. Networking never was a major interest for them.

Pelkey: What about at UCSB?

Kahn: My impression is that was a one-man show for quite a while and when Glen Culler left, the network effort at UCSB sort of stopped. Actually, Glen Culler was very active in some of the early work on packetized speech along with Lincoln Labs and USC/ISI. In the '74, '75 timeframe we demonstrated that you could actually send speech over a packet network. Glen Culler was very active at the time; he was at Culler-Harrison by that time. His network connection was through the UCSB node. But nobody from UCSB was deeply involved in that experiment, now that I think about it.

Pelkey: Although there were a lot of spinouts as it turned out in the Santa Barbara area of people who worked in that project early on. A number of companies were created from people who were associated with Glen and/or that node.

Now in the early '70s, was it true that ARPA/DARPA, when you were there, really tried to get AT&T to take over the management of the network before it was ever handed off to DCA in '75? That there was an effort to get AT&T to take it over, and AT&T wouldn't?

Kahn: I can tell you what I remember about that. That's a question you ought to ask Larry Roberts, in particular. My recollection is that there was a study done by Paul

Baran, who was then at a company called Cabledata Associates--in fact, Vint Cerf worked with him, Vint was at Stanford at the time--to figure out what the commercial applications of that might be and what a spinoff strategy might look like. And Paul came up with some recommendations in a three-volume report and a spinoff strategy in another. The report described a strategy to devolve it to the private sector and the second report was what such a net might be most useful for. He became fascinated with the idea of using the network for automating the procurement process. He had determined that's where most of the leverage was. As for the spinoff, the auditors from DoD had decided that there was no possibility that they were going to support a decision to just get the net out of defense until the DoD had taken a clear look at whether it had any utility within defense. They recommended that we at least find out whether there was some utility for the Arpanet within defense before giving it away. The conclusion that came out of that internal look was, yes, it had applications within defense. In fact, it later became the basis for the DDN command and control backbone of the defense communications system. DCA agreed to manage the Arpanet. My view was that DCA had mixed feelings at that time, because it was deeply involved with circuit switching. Packet switching was a new notion for them. But it turned out to become more and more useful and effective and they later decided to build a net of their own, called AUTODIN II, which they actually contracted for, got Western Union to build, and canceled like the day after they accepted it. It was not as effective for DOD as the Arpanet.

Pelkey: Did it work?

Kahn: It satisfied the specifications, I believe. The question was, did you get acceptable performance for the cost? It turned out to be more costly than the Arpanet but it had features the Arpanet didn’t have. DOD decided it wasn’t cost effective, so they finally went back and started to use Arpanet instead.

Pelkey: But then Arpanet was really still a network for the research community as opposed to DDN that was a command and control network.

Kahn: That's right.

Pelkey: It is a much higher speed network as well, right?

Kahn: No, DDN is—

Pelkey: Is it a clone?

Kahn: The technology is the same. In 1983, when we cut over to TCP/IP, a part of the motivation was to allow the Arpanet to split into two separate networks. The Milnet was the military side of it, and the remainder, renamed Arpanet, was retained as the university part largely.

Pelkey: So there are at least three clones of that. MILNET, Arpanet and DDN.

Kahn: No. The DDN is the generic name for the Defense Data Network, which actually, I think, encompasses lots of different nets, of which MILNET is one.

Pelkey: Okay. Another question? Having been at Bell Labs -- Bell Labs was kind of a repository of communications technologies in the early to mid '60s. Then the modem guys started to come alive in the '60s, Codex, Milgo and Vadic. I guess in the '60s, there were a lot of startup modem guys, and a lot of that came out of telemetry applications and building modems for military applications and so on. And none of those companies became factors or participated in this process that we've been speaking of.

Kahn: They supplied modems.

Pelkey: Right.

Kahn: I mean, if you wanted to dial in from one place to another, you'd need the modems, so—

Pelkey: Right. And that was it originally. They were vendors of modems.

Kahn: Vendors of modems, that's right. Well until probably the mid 1970s when they started to provide more network kinds of services. Today, Codex provides network management devices and controllers and things that will help you organize a network, and they have a lot of technology that derives from what we did. But you're right. They played no essential role in that early development.

Pelkey: Why?

Kahn: I can only conjecture.

Pelkey: Well I'd love your conjecture, because you must know--there must have been interactions.

Kahn: Well, first of all, the guy who built the first high-speed modem was Jerry Holsinger[24]. He did the original work while he was at Codex, and then he left to start his own company. They were fixated at the time on the ability to get 9600 bits through a phone line that defined their horizon at that point. If you grew up in that discipline area, you would realize it was a monumental achievement to do that kind of thing. Furthermore, the talents that it took to build modems--is that not working? [Tape pause] Is it recording now?

Pelkey: Yes.

Kahn: Okay.

Pelkey: My battery must be going down a bit, but we're okay.

[24] See Interview – Jerry Holsinger

Kahn: Okay. The kind of talent that it took to build a modem like that required strong analytical capabilities and knowledge of analog circuitry. Effective digital componentry was not really available in those days. A different set of talents were involved in the networking task.

Pelkey: There was an element of math--mathematicians were involved.

Kahn: Mathematics was involved in analyzing things like spectra, analog circuitry, filtering, bit synchronization, “eye” patterns signal design... I mean, these were key issues in modem technology. The key things in the network were knowledge of operating systems, network management and real-time software. Software design didn't show up in the modem activity at all. About the only point of overlap between the two was something called ARQ that was essentially error detection and retransmission systems or, more generally, error-control. That was the only overlap between those two areas, and maybe use of circuits and modems. Now the network engineers viewed circuits and modems as wires. So they had a very simple model. These modem manufacturers could spend years designing modems. There was also another area: system design in the large, which the modem manufacturers could have done, but didn't have the framework in which to get into it then.

Pelkey: Right. But in the early '70s, Codex was delivering big statistical multiplexers. Worldwide link networks—

Kahn: Right. It was one piece of the problem. They got into it-- they started to bite off pieces incrementally. MIT was to communications in the 1960s as Xerox PARC was to LANs and personal computers in the 1970s.

I mean, some of the key people that have populated the field were there and doing their most active work. The people at MIT working on communications at the time were not focused on networking. It was not analytical. The work at MIT was heavily influenced by the information theory, which was highly mathematical, formulistic, analytical... The early networking efforts were largely seat of the pants. The two approaches just didn't mix well at first.

Pelkey: Okay, now, fit AT&T into this.

Kahn: Well, let me fit myself into this just for a moment, because I was sort of split across these camps. I mean, I was quite involved in communications-- remember that I was also at Bell Laboratories, and I had worked in the communications group at MIT. But at the same time, I had a lot of friends who were heavily into computer science and I followed that area closely. When I went to BBN, this was the primary expertise that existed within BBN (along with acoustics, of course). BBN, at that time, was not into communications at all.

Pelkey: Yeah, I was going to say, when you joined BBN, I mean, they weren't in communications until they got the Arpanet contract, right?

Kahn: That's basically right. That was the first significant effort they had that was in the communications area. I like to think that I played a key role in getting BBN into that area; other people at BBN were also involved. For the first time, I was in a place where I could synthesize a program in computers and communications. I felt I could somehow bring the two areas together.

Pelkey: And what competence did BBN have in operating systems and such when you joined it?

Kahn: Well, the very first operating system for the PDP 1 had been built there in the early 1960s, and they had run the first commercial timesharing system in the mid-60s called TELCOMP. They were experimenting with some of the most advanced operating systems.

Pelkey: Now how did this, your information theory background and knowledge impact your views, concepts, and designs when you joined BBN?

Kahn: Just as—

Pelkey: What was it that you had that someone else who hadn't been exposed to your background wouldn't have brought to designing a network?

Kahn: What can I say, I don't know? It seemed like something I could do.

Pelkey: And were their specific things about your Bell Lab background that—

Kahn: I remember thinking about all of these things. It was very natural to me. If I asked you to think about building a superconducting chip for my new levitation reactor, how comfortable would you feel thinking about that?

Pelkey: Not particularly.

Kahn: I mean you'd have to somehow get into it first.

Pelkey: I would just try to convince you to give me another problem.

Kahn: There were probably less than a dozen people in the country who knew about operating systems. There was no easy way to learn.

Pelkey: But by going and doing.

Kahn: I just happened to be there and close enough to people creating operating systems to see what was going on and to understand it well enough to be able to go and tackle that as a joint problem. Larry Roberts could have done it, but I don't think there were half a dozen people in the whole country that had enough real deep insights into the fields to bring them together.

Pelkey: Why didn't AT&T play a role?

Kahn: They provided the circuits. But they had a vested interested in circuit switching at the time. I mean, there was no market for computer-to-computer communications at the time.

Pelkey: I mean, you can tell me about Bell Labs. How Bell Labs was viewed as such a national resource, science center and so on: they were science. Why were left behind? Why did Bell Labs not do something?

Kahn: Bell Labs could have done research on this.

Pelkey: But could they have connected their computers inside their shops, and—

Kahn: But you have to remember that in those days, AT&T was precluded from being in the computing business. To the extent that they used computers, it was for large-scale job computations or experimental research. They weren't in the computer business.

Pelkey: And the Newhall Network, when was that, the late '50s, early '60s?

Kahn: They did some experiments with ring networks in those days. In fact, you know, even though ARPA did all its work on packet switching in the late '60s and early '70s, AT&T filed for a patent in packet switching in 1978.

Pelkey: Really?

Kahn: It seemed strange, but they claimed they'd been working on it all along.

Pelkey: Because of Newhall?

Kahn: Newhall was there, as was John Pierce.

Pelkey: That's where Farber picked the ring technology up. So it existed within Bell Labs, I mean there were some concepts of local area networks, i.e. ring networks.

Kahn: In the time frame that this work started, timesharing was largely a product of DARPA. I mean, DARPA had supported the early work on timesharing, and there were only a few sites in the country that had working timesharing systems. And I'm not sure if AT&T was one of them at that point in time. If you think about their motivations, they just didn't have the same user community out there. But even if they had succeeded, there was no market for computer communications right then. Besides, AT&T already had a big commitment to its own existing circuit switching system. We held a meeting with AT&T around ‘69 at their headquarters in 195 Broadway. The view was expressed that packet switching was probably not workable as a technology. If you wanted to switch rapidly between machines at 50 kilobits, they would have such a system some time in the future.

Pelkey: Now was that presentation made before BBN was awarded the Arpanet contract?

Kahn: I think it was just after, but I'm not positive. It could have turned out to be just before, but it was certainly right around then. Because I remember one of the concerns that we had was that if you hook up one of the IMPs to a modem, and the interface didn't work, then what? So we wanted to find out what was in their modems. They kept telling us it was idiot-proof. "Don't worry, nothing will happen, it will all work." We never did succeed in getting much information at that stage, I think we later got a lot more details, but it turned out to be irrelevant, because they were right--they just worked.

Pelkey: Why did IBM pick token ring?

Kahn: I don't know.

Pelkey: E-mail. Why—

Kahn: Let me back up again. Just for me to be able to go into this thing was actually a major decision on my part, because—

Pelkey: You had jumped out of this field of communications knowledge, and—

Kahn: --you see, I had come out of the communications field. If you think about the number of people who knew about computers and communications, they were either at Bell Labs or maybe there were maybe six to eight people who really knew about both in any real detail. Perhaps a dozen, I don't really know.

Pelkey: But there was a growing number at MIT?

Kahn: For me to now move to BBN and work on networking was a bit heretical.

Pelkey: Did people think you were crazy when you made that jump?

Kahn: Many people thought that I had gone a bit off the deep end. If I had not been successful with the Arpanet project, my reentry paths would have been difficult.

Pelkey: Were you aware of that, or did you ever think about that?

Kahn: I knew about that. It just seemed to me like the right thing to do at that point. I mean, it had all the wrong properties. I don't know if I can even explain it to you--it's sort of like somebody trained in good etiquette suddenly becoming a brawler in the streets.

Pelkey: Yes. This is all nice and formal—

Kahn: This is all nice and formal, and suddenly I was jumping into something that had none of those attributes. And as I say, if that had not been successful, there was probably no good reentry path. But, as it turned—

Pelkey: Was it a big decision for you? Or it just kind of came naturally. Was it intuitive?

Kahn: Well, it was sort of intuitive and a big decision at the same time. I mean it was sort of something that you know you have to do, 'cause I really felt that that was the right thing to do. I knew it didn't obey any of these properties, so I figured, well, "I think I can make this happen, so I'll just make the commitment and go do it." Maybe it's like, you know, jumping into a presidential race when nobody thinks you should. If you win great, but if you don't, you're probably nowhere. I think I made the right decision, and it turned out well. I then went into DARPA and ran the program there for almost 15 years, so it really worked out okay for me. But it could have been a disaster in my career.

Pelkey: Yes. And as you say, there weren't very many people who could have made this jump, because of the limited universe of people.

Kahn: And make it happen.

Pelkey: Yeah. I mean it's not clear to me, other than--at least, you're not able to express what it was from this more formal, i.e. the modem side of things that you brought to the this side, that influenced your thinking that caused you to be successful over here.

Kahn: Well, I wouldn't say it was the modem side, I really think it was the communications side.

Pelkey: The communications side.

Kahn: It was the ability to think about communications as more than just a bunch of wires.

Pelkey: Maybe you can explain while you're here? Deadlocks? Routing, or any of those things come out of your communications intuition?

Kahn: Well, it came out later, I mean, deadlocks was something that emerged.

Pelkey: But it would never have come out of your thinking about networks?

Kahn: I think anybody would have trouble with those notions.

Pelkey: Routing would have come up from communications

Kahn: That could have come up from either side. How to implement network software was more likely to come up from the computer side. How to think about good routing was more likely to come up from the communications side. One culture would have analyzed the problem mathematically, optimized the flow patterns, picked exactly the right metrics, whereas the other culture would have said, "Well, let's ship it to someone who can get it there. Anybody who is closer and who'll take it, let's give to them. Maybe they can get it there." It's sort of a real philosophical difference between those two points of view.

Pelkey: This process of intellectual development is really important to me in terms of my thinking about the book that I'm writing. How ideas develop, and who were the carriers of the ideas?

Kahn: Well, my experience is that you start with a belief, a philosophy, an approach and until you prove it, it's suspect. And after you prove it, then everybody says, "Well, of course." Take what we're doing right now. There is not yet a groundswell of support that says what we are now doing in this organization is the right long-term approach, but I think if you come back and look 10 years from now, assuming we're successful, everybody will say, "Well, of course. Obviously that was the way."

But it's so counter-culture right now that it's hard to even explain it.

Tape side ends

Kahn: Here's a note. Do you know Keith Uncapher?

Pelkey: No I don't. I only know him by name.

Kahn: Here's a note Keith sent me today--comes on the network. [Pause] That's an example, and that doesn't tell much about what we're doing, but there's an idea in there that's as powerful as the whole DARPA idea has been over the past 20 years.

Pelkey: There was a social behavior that came out of E-mail. It was a significant one. Why wasn't it thought of beforehand? Why did it arise serendipitously, and why was it so successful?

Kahn: Well, I think, it turned out that there was a need that it satisfied when it was developed that people didn't really foresee. I mean in thinking about a motivation for doing a project like Arpanet, you don't focus on the need for a person to communicate with another person, because you have the telephone. If you want to talk to somebody, you can talk to him on the telephone. Or write them a letter. Send them a telex. All right? The value of Arpanet was these grand ideas, the grand challenges? You could somehow manage to get computers to talk to computers. This was not people talking to people; there were ways for people to talk to people already. The motivation was based on what these computers could do if they were interacting. 'Well, this computer could run those programs over there, or it could ship its files.' I mean, they were thinking about things that couldn't be done. I mean that's the grand challenge of a DARPA. Not to figure out what you now could do a little bit better. But sometimes doing what you now can do a little bit better can make a radical difference in how you carry out your functions. And I think that's what happened there. It was just a matter of recognizing that with this network in place, who's motivation was totally different, it was now possible to move a message electronically from one place to another. And then suddenly people realized, ‘Hey, that's really neat!’ It didn't cost them anything, so economics wasn't an issue. Suddenly they could type a message and send it to a hundred people simultaneously, get 14 interchanges a day, they could do it at deferred times, they didn't have to get the other person on the phone. They could send a message at three in the morning, it could be read at six in the morning and by nine o'clock the next morning have five interchanges. The minute the capability was there, they could see what to do with it.

Pelkey: When you live in a two-dimensional world, the third dimension isn't obvious.

Kahn: Yes, I mean, it gets into this business of infrastructure in a big way, which is really what we're focusing on in this whole organization. Suppose somebody said to you in 1876, "What would you do if you had electricity in your house?"

Pelkey: Right. I agree.

Kahn: I mean, you might have said, ‘Well, I'll replace my kerosene lamps.’ You probably wouldn't have said, ‘Well, I'd have a Mixmaster, an electric fan...’ Just somehow that whole thing would have passed you by. For some reason or other, people weren't thinking about electronic mail just because it was too easy, this connectivity and ubiquity, which accommodates trial and error, new things happening, sometimes critical mass building. It's kind of an experimental primordial soup, if you will. Now you can call that learning, and I'm comfortable with that notion. In the case of the Arpanet, a lot of those things happened because it was there. And there was no way to plan it-- just the presence. It's like the chemistry between two people. You put them in the right place, and maybe something will happen. If they never get into the right place, nothing will happen. And there's no way that you can think it through, because if it was a chemical reaction or just a spontaneous happening, they had to be there. We [CNRI] are setting ourselves up in such a way that this organization is a piece of infrastructure. We set ourselves up to be in a position to essentially provide the leadership to create infrastructure. There's never been an organization that's really tried to do that other than perhaps Bell Laboratories and AT&T before divestiture. But they were driven by market forces rather than by just plain research interests.

Pelkey: I want to ask you about protocols. Clearly the concept of protocols existed prior to this period of time.

Kahn: Well, there were protocols before, but they weren't recognized as being in the province of computer networking. Remember I mentioned ARQ protocols. These were techniques for moving data across a circuit. You send something, you acknowledge it, you send the next thing; you acknowledge it; or send five things and acknowledge all five. The notion of a protocol existed, but nobody working on operating systems in the early 1960s thought about layered communication protocols in their machines.

Pelkey: Right. Where did the concept of protocols, particularly the concept of layered protocols, come from?

Kahn: I think it came out naturally from the DARPA work. I mean, having the network there, layering was just sort of the natural way that people thought about it, because you can incorporate them one at a time. It wasn't necessarily the best choice. I mean, there were many reasons why you might want to do things slightly differently. In fact, these layerings usually don't allow you to poke through the layers. Sometimes you'd like to go from the top layer right to the bottom and say, "Look, do something down there, 'cause you're slowing me up." Except you have to work from layer to layer to layer and you can't quite control the whole process. So it's not clear that layered protocols per se--the way they were done--are necessarily right. On the other hand, the notion of layers of protocols is a natural way to think about the problem, at least in natural systems compositions, but you've got to think about it in terms of the overall performance of that collection as a system.

Pelkey: Jon Postel holds the view that the reason layering came about was because of the people who were working in the Network Working Group all knew about operating systems. And operating systems were layered structures. And that seemed the most natural thing way to do networking protocols.

Kahn: Jon's in a better position to know. I would respect his opinion.

Pelkey: I'm trying to understand how industry grows in terms of its intellectual complexity and robustness.

Kahn: But you see, I think that there weren't many other alternatives, simply because, you know, people first took on this problem. I mean, the way you think about programs interacting with programs, you know, they call them, right? That's what they do. And, I mean, you can think of subroutines as layered in some sense if you want, although it's not necessarily a one-dimensional layering. The notion is data moving from one place to another. And therefore these programs that are calling each other could be viewed as layers. But you didn't have to think about it that way. Frankly, I don't think about protocols as layered structures myself. But there's always room for multiple interpretations. When Vint and I wrote the original Internet paper, he and I had somewhat different views of how TCP would actually function. I mean there was a time when Dave Reed from MIT, and maybe Jerry Saltzer[25], had come up with another protocol because they didn't like TCP. It was called DSP (Data Stream Protocol). I sat down with them, and they described the Data Stream Protocol, and I said, "But that's what TCP is like, with one minor difference." And they said, "No it isn't, because read what Vint says in this report." And I said, "That's just the way he interpreted it. Go back to our original paper, and you'll see you could interpret it your way, too." So even that

[25] See Interview – Jerry Saltzer

original paper that we wrote was subject to multiple interpretations about how you actually do the implementation. I always thought about protocols as collections of functions which, according to some sequencing, probably driven more by the needs of the data or the machine, would be called into play at the right time, rather than having a pipeline arrangement where first you do this function then you do this function then you do this function, as if they were all there in layered sequence. That's a matter of how you think about it. I'm sure Jon accurately described how he thinks about it.

Pelkey: One last thing? Do you think that when ARPA became DARPA, that this center of science research funding changed in that DARPA will never be what ARPA was in the '60s because of its reduced mission, or its mission focus on defense projects, rather than on funding long-term research?

Kahn: Well, let me say that I was with that organization for almost 15 years starting in '72; that was after ARPA became DARPA. This was the period of the Vietnam War anti-military syndrome during the '70s. During that whole period, DARPA never really significantly lost its view of its role as supporting basic research. And in fact, you know, when I took over the office, it had an annual budget that was somewhere in the range of 40 to 50 million dollars, and it went up to almost a quarter of a billion dollars during my tenure, most of which was spent on basic research. I mean there's always concern for the fact that an organization like DARPA can turn from what it's been doing in the past and, in recent years, that concern has amplified significantly. I think DARPA is still, more than any other place, a supporter of basic research in computing. Whether it will change or not, I don't know. But I don't think the placing of the D in front of its name made any difference. DARPA is an organization of people. And if you put in people who are more concerned with direct military applications mostly, that's the way the organization will be. And if you have people in there that are more concerned about basic technology, which might have military or commercial impact, then that's the way it will be. It's a function of the people.

Pelkey: Are you concerned about the level of science being done in this country?

Kahn: You mean, the amount of it or the quality of it?

Pelkey: The quality. It's similar, in that, I think you have to do an amount to have quality.

Kahn: It's unpredictable. I think I'm mostly concerned about the fact that the quality of the work that's funded by the government is affected in ways, better or worse, by the quality of the people that are managing the programs, whether they're in DARPA or anywhere else. And those are people who are making decisions about what to fund, both in terms of programs and proposals. They're the ones that are putting together the RFPs, they're the ones that are making the evaluations. With the best people in those positions, you're going to have a much greater chance of success for any level of funding than if you have people who are less capable. Of course, good researchers are necessary to keep the quality up. To have a situation which allows the government to pay half or a third of what the best people could earn in private industry guarantees that you're not going to attract the best people into the government. That doesn't mean that there might not be a few people who say, "Well, look, I don't care, I'll take a 50% cut in salary, because what I'm doing is so important."

Pelkey: But unlikely.

Kahn: Yeah. I think that that's my most serious concern right now. The government really cannot attract the best people, and by not attracting the best people, it means the quality of oversight and management is bound to be less than the best.. So I'm very concerned about it from that point of view.

Pelkey: That raises the question of the educational system that is a major concern, as well. Before we finish and we go feed our hungry stomachs, I want to express my immense gratitude. This has been a fascinating interview and more than I had hoped for. I very much appreciate your taking valuable time in your busy schedule.

Interview ends

Interviewees Mentioned in the above Robert Kahn Interview

Norm Abramson	Paul Baran	Wesley Clark
Vint Cerf	Steve Crocker	Donald Davies
Dave Farber	Howard Frank	Jerry Holsinger
Leonard Kleinrock	J.C.R. Licklider	Robert Metcalfe
Jon Postel	Larry Roberts	Jerry Saltzer
Robert Taylor	Hubert Zimmerman	