Interview of Leonard Kleinrock

Interviewed by:
James Pelkey

Recorded: April 27, 1988
Los Angeles, California

CHM Reference number: X5671.2010

James Pelkey: You were involved in communications in the mid '60s?

Leonard Kleinrock: Yeah, my doctoral dissertation, which I actually started in 1959, was devoted to the study of communication networks, and it basically laid out the analytic formulation for the performance evaluation of what we now call packet networks. It also looked into the design issues; how you do capacity assignment, routing, topology design, etc., of these networks. That thesis was completed in late '62, it was published as a book in '64 by McGraw-Hill, and lay fallow in the scientific literature because it was, in some sense, before it's time. Nobody believed these networks could be built. It drew upon some ideas from classical military communications networks. That's where the ideas came from. Then, in the mid '60s, as you say, '66 or so, there began to be an interest in data communications between computers. At that point, of course, Larry Roberts was beginning to run his experiments -- earlier than that, even -- between two machines.

Pelkey: Were you aware, at that point in time, of his experiments?

Kleinrock: Oh, yes. In fact, he had done them quite a bit earlier; just around the time I was -- I left Lincoln Lab in '63, and he was running some of those experiments just around that time: Q-32 to TX-2 kinds of things. You may want to correct me on the dates, but I was certainly -- because Larry and I classmates at MIT, along with Ivan Sutherland. The three of us took our doctorates on the same day. We shared some of the same professors: Minsky, and Shannon, and Arthurs and a couple of others there, and we all had to give a demonstration on the Lincoln Lab's TX-2. We all had -- a part of our dissertations was simulations or demonstrations; Ivan's sketch pad, Larry's three dimensional hidden line stuff, and my stuff on networks, had a big simulation. If you want to hear about that --

Pelkey: Yes.

Kleinrock: It was quite a day. We took our formal exam at the university. Then we had to demonstrate our theses, and Ivan went first, as I recall. He had sketch pad, which was a very nice graphics drawing package, where you could create an object, define it as an object, call up many of those, connect those, define that thing, call up many of those, and he built himself a computer with a few strokes of a light pen. It had, among other things, what's called 'constraint satisfaction.' He could draw what looked like a bridge truss, and put weights and stresses at various points, and it would adjust its position to minimize the energy, in order to -- so lines and things would lengthen. You could define -- draw four points and put some constraints. These three points are in a straight line, and the whole thing is a square, and you could put in contradictory constraints, and the system would do a sort of least mean square optimization of that. It did some very nice things. Then I believe Larry -- no, I guess Ivan was then -- he turned around from the computer and started answering questions, and Marvin Minsky went up to the screen. He defined his own structure and said: "Satisfy constraints," and it was a weird picture, and it couldn't satisfy it. The thing was exploding behind Ivan's head. Circles were coming and things were moving; the thing was going crazy, and I was over there. It never settled out. It was amazing. Anyway, the point is Larry and I knew each other since 1959, 1960.

Pelkey: And when was this dissertation done?

Kleinrock: November or December of '62.

Pelkey: Minsky was there and Shannon was there?

Kleinrock: Shannon was on my committee, and I think on Larry's; Minsky was on Ivan's. I forget the configuration, but Shannon was mine. Anyway, I knew Larry very well. He went to DARPA, and I went to UCLA. He stayed at Lincoln Labs for a while, and then he went to DARPA. Then he was busy -- these were the days of timesharing.

Pelkey: Let me ask you one question: When you were doing your work in '59, there were no such things as packets.

Kleinrock: Oh, no. We called them message communication networks. There was no idea of packetizing yet.

Pelkey: There were no examples of message networks yet.

Kleinrock: Yes there were. The military had messages -- there was a thing called UNICOM, some kind of Universal Communication Network they were planning. There were some small ones built. There were a few around.

Pelkey: When does Paul Baran's work --

Kleinrock: Paul Baran's work was in the early '60s, and he was talking about --

Pelkey: Military communications?

Kleinrock: Oh yeah, it was a military study he was doing. There was a guy named Reese Prosser who had studies networks of highly mobile terminals, and he analyzed the routing problems. He introduced - - he was one of the guys who contributed the idea of 'hot potato routing,' because the idea was that you don't follow people. You hope the packet finds them, but you can't track them. He had analyzed this, and I analyzed random routing also, which is what I called it in my dissertation, and showed the inefficiencies, in one sense, because of the delay. At least you get there. If you try to track people, you use the whole network up trying to track them. So there were a lot of studies going on at the time. There were some example networks around. So those were the days of timesharing, and what I was doing in those days, since there was no interest in networks, I was analyzing timesharing algorithms, using the same kinds of tools I used for network analysis. Larry was beginning to fund more and more research sites to do studies in computer science, and the first thing, when he went to a new university and a new researcher, the guy said: "You want me to do research on computers, buy me a computer."

Pelkey: So he would come to the universities and --

Kleinrock: Yeah, and the university would say: "You want me to do research on computers? Buy me a computer," and they were buying PDP10s all over the place. So a lot of good facilities got built up. Utah had excellent graphics facilities. Here at UCLA we began to develop some simulation packages. At SRI there was network storage and network information center kinds of things, before networks came about. They had some good ideas all over the place, and unique facilities. ILIAC was being built in Illinois. MAC and Multix were being started at MIT, and it became clear that you couldn't replicate all these capabilities at every site, in Larry's mind. So one of the things he did was say: "Ok, we've got to give people access to these unique resources so they can then go to those resources and use the facilities of those resources." So exactly who was first to have the idea that the Department of Defense should build such a network, I suspect it was Larry.

Interruption in Interview

Kleinrock: As I say, I suspect the idea came from Larry, for networking, although he was working for Bob Taylor at the time, and it may have come from Bob Taylor. I don't know where the original thoughts came from. What I do know is that in 1967, he convened a group of us together. I was there, I think Doug Englebart, Larry of course.

Pelkey: Had you interacted with Larry before this meeting in '67?

Kleinrock: Oh, yeah, we had talked about the ideas of networking, yeah. He saw that what I had done was a model of what could be built, and he was aware of (unintelligible proper name) work, and Dave's work over at NPL, National Physical Laboratories. We had been talking about these ideas for a long time. So, again, the pressure from groups doing research for DARPA having unique resources put pressure on

DARPA to give access to those resources. So they convened this group. I believe there was another -- a guy named Doug McKay was there from IBM as well, and possibly someone from AT&T, I'm not sure.

Pelkey: Do you recall where that meeting was?

Kleinrock: Yeah, it was in Washington, in some DARPA office.

Pelkey: What time of year?

Kleinrock: For some reason I'm thinking April, but I could look it up somehow. Larry convened it. The idea was -- Doug Englebart was surely there. Oh yes, also I believe Herb Baskin was there. The idea was to write a spec for this network, which could then become an RFP and sent out to industry to get proposals. We sat down and we sketched out the ideas. I remember, among the things that came up, one was reliability. Reliability is a hard thing to define, so we took a very simple viewpoint, and said: "It should be that if any one thing in the network breaks, it doesn't stop anything else from functioning." So the idea of a two- connected network came up. Do you know what that means? Between any two nodes, there are two independent connections. Herb Baskin banged his hand on the table and said: "Look, if I have worse than a half-second response time, it's not interactive across the network. Packets move around with a round trip, taking less than half a second on average. Long ones take longer, shorter ones less; on average, half a second."

Pelkey: And it was clearly defined as a packet network at this point?

Kleinrock: Yes. The concept of packet switching was definitely there at that point. I was very much concerned about performance evaluation and measurement, and I said: "We have to build into these switches measurement capability, so we know what's going on. It doesn't have to be turned on all the time, but it's got to be there." So we said there's got to be measurement capability and specified the kinds of things you should be able to measure. And so we went.

Pelkey: It was a couple day meeting?

Kleinrock: One or two. I think it was one, but I'm not sure. I guess after that, Larry took those details and created the spec, and we may have interacted with him some after that. Something else happened at that time though. Oh yes. You probably remember the climate, at that time, in the Spring Joint Computer Conference, Fall Joint Computer Conference, from the mid '60s on until the '70s, there were always these panel discussions; the same panel discussion each time. Computer industry speakers here: communication industry speakers sitting there. Computer industry would get up and say: "Hey guys, we need good data communication." Larry was among those people saying that. Communication industry would say to these guys: "What do you mean. United States is a copper mine. We have wires all over the place. Use them." The computer guys would say: "No, you don't understand. You have a telephone network; it takes 25 seconds to dial-up, you charge us a minimum of a three-minute call. We want to send 20 milliseconds of data." Communication industry would say: "Little boy, go away. You're no revenue for us." And so these little boys went away and created their own solution, and so you have that picture. The overlap is that telecommunications is a little piece and the computer guys are the ones that did it, because they needed it. They used the facilities, of course, of the AT&T plant, but AT&T just said: "Go away." IBM, as I say, was involved in this in the early stage and they dropped out. They dropped out for reasons I don't have certain knowledge of, but it could have been that they wanted to bid on the network itself and couldn't be part of the spec, or --

Pelkey: They eventually ended up no bidding on it.

Kleinrock: That's right. They considered bidding it, but they no-bid it. Then the RFP went out. A number of companies bid the DDP- 316.

Pelkey: The Honeywell machine?

Kleinrock: Yeah, because that machine was a new machine then. It had 16 levels of interrupt, and that was the main thing that we needed to handle the switch. BBN bid it as well. I remember, I think it was the spring or fall of '68, the SJCC or the FJCC, where they had that machine demonstrated at that conference, and it was the hardened version for the military. It had a hook on the top of it, and they had it hanging from a cable, and it was running, and some guy with a sledgehammer was pounding on the son-of-a-gun, and it kept working.

Pelkey: Is this true?

Kleinrock: This is true. I saw it. I think that's the machine we got. Honest to God. Anyway, that's where the IMP is, inside.

Pelkey: Oh, God, that's a great story.

Kleinrock: We have the original IMP. It's now retired, and we're planning to put a case around it, but it's in our -- you can take a look at it if you want.

Pelkey: I'd love to see it.

Kleinrock: So what happened, the RFP went out in '68, early '68 I think.

Pelkey: How did the four sites get selected? The four sites were identified in the RFP, right?

Kleinrock: I don't recall. I think so. I'm not sure.

Pelkey: But it's clear everybody seemed to know this.

Kleinrock: Yeah, we decided UCLA would be site number one.

Pelkey: Why?

Kleinrock: Because I had the analytic capability here, I had the software people, and we had the measurement thinking here.

Pelkey: What software people did you have? Were there other groups that were being contracted with?

Kleinrock: No, no, it was the students, the graduate students, who had run simulations of networks by then.

Pelkey: But they were doing this host to host?

Kleinrock: Not yet. The idea of host protocol wasn't even around at the time.

Pelkey: But they were measuring -- computer network measurement

Kleinrock: Well, we were doing simulations of networks. At this point, I became very active again. I started producing papers in the area, and I had students that were actively interested, and Larry knew me and trusted me. It's surprising he did it across the country, but I suspect he wanted to get Utah attached real soon. Ivan was there, and that was an example of a resource that was useful to people. SRI had some real industrial capabilities at the time in system management. That may have been the reason there. Why they put on Santa Barbara I'll never know, except at that time --

Pelkey: Glen Culler.

Kleinrock: Glen Culler with his Culler-Freed stuff and all the rest.

Pelkey: Maybe also you certainly don't have a network without three nodes, and you really should go to more than three to have a network.

Kleinrock: Well, it was a weird network, you know. It violated the two-connected constraint, if you recall also, Utah was out on a spur. But we allowed the growing network to have that capability. That was Utah -- UCLA. We allowed that. So the story goes, as I recall I think the contract was awarded to BBN in January '69, and they were supposed to deliver the first IMP to UCLA in September '69, and that turned out to be a weekend.

Pelkey: Labor Day weekend.

Kleinrock: And that caused an interesting problem, as you'll see. I mentioned TTI [Technology Transfer Institute] before. You may not TTI before. You may not want to put this in the book -- TTI was formed in 1976, but in October of '68, there was a company called Linkabit formed. You probably never heard of it.

Pelkey: Right.

Kleinrock: You know them?

Pelkey: I sure do.

Kleinrock: M/A-COM Technology Solutions bought them out. I was the first president of that thing. There was myself, there was Andy Viterbi and there was Irwin Jacobs. Andy was a professor here at UCLA. We were buddy-buddy. Irwin was visiting at the time at JPL, I think. He was on the faculty at MIT. We formed this company. I was first president. Now in summer of '69, it was clear we were going to have a lot of work to do to simply learn how to take our SDS Sigma 7 and connect it to this thing called an IMP, whatever it is. We were going to be the first node. To be the first node on a network is almost a contradiction, but it's not. That's what you got to do: hardware, software, etc. Host to IMP level. It was clear we were going to have to write a lot of software and build a physical hardware and software interface in our machine. So I told these guys, come summer of '69, I'm going to have to quiet down here and I'll pick up in the fall again. They said: "Fine." (But comes the fall, I came back to UCLA. Then my 'buddies' at Linkabit said: "You're out." So I finally extracted from them a one-page document which I tackled in the office that night, which guaranteed that I had X% of the company, and then I went away. That X% became beaucoup dollars, when M/A-COM took them over, so I had a last laugh, in a sense.)

We worked the summer trying to get a spec on the IMP that BBN wouldn't release. They hadn't finished it yet, and yet we had do design against it. So it was a frantic time. Finally, late in the summer, they released the hardware and software specs, and we raced to build that interface -- and I had some people here doing it, while we were doing the research at the same time, and this thing was supposed to come out that weekend, and we were told: "Oh, it's labor day weekend and it's going to be a week late," or something. We say: "Whew. Wonderful! We need the time." Then they put the damn thing on an airplane and shipped it out, and it appeared shortly after, here, and they wheeled it in. Now, you've got to imagine who was there. Ok? Let's see, first of all the computer science department representatives were here; there were UCLA representatives from the administration, to see this great event; there were people from AT&T long lines who were going to connect, eventually, over some modems to these two; there was the local company, I think it was GTE; there were DARPA guys here; there were BBN guys here; there were Honeywell people here; and the SDS guys were here It's this grand event, and in walks this thing, and they were ready to point the finger if it didn't work "It's not my fault, it's your fault," right Well, to BBN's credit, they connected this thing up, and that day we had bits moving back and forth. I think the next day, we had packets moving back and forth, and messages worked. It came in, and the damned thing worked.

Pelkey: And it was your hardware that you had to put in the SDS machine that hooked to the serial link to the IMP?

Kleinrock: That's right. That's right, and the damned thing worked.

Pelkey: And the software that sat in the host you wrote?

Kleinrock: Sat in the host. We wrote it.

Pelkey: Now that was Steve Crocker and Jon Postel and Vint Cerf?

Kleinrock: That's right. Those guys were writing the software, basically matching the protocol, the host-IMP protocol, which was slowly being specified.

Pelkey: Was 1822 specified at this point in time?

Kleinrock: I don't think it had come out by then. 1822 was not. 1822 was the host to host.

Pelkey: Oh, right, but not host to IMP.

Kleinrock: So the host-IMP was an accrued document, notes passed back and forth. I actually somewhere have -- you may want this. I think its the first message that we sent; the first Teletype printout. I will ask my secretary to find it. Anyway, it was a great success.

Pelkey: It must have been a great relief. You must have been frantic.

Kleinrock: Oh, yeah. Well, I don't get frantic, but I was anxious, a little. I knew my guys had done a good job, but, if it didn't work the first time, it would work a few days later. They were going to sit here -- if I'm not mistaken, Bob Kahn came out also. He came out with a notebook, a Bell notebook, like an AT&T notebook, where you can't take the pages out, sort of yellow pages and graph paper in there; and he was taking notes. I remember that they were very careful, very scientific.

Pelkey: Did you know Bob at this point?

Kleinrock: Not well. He got to MIT after I did. I met him through this. I work, I take things, and I throw them away. He was VERY organized. The damned thing worked, so that was good. So by December this was up.

Pelkey: The second one was --

Kleinrock: The second one was -- ha, ha, I can tell. It's in my book. Sure, it was SRI. Definitely. I'll tell you why I know that.

Pelkey: Elmer Shapiro was the person at that facility.

Kleinrock: Yeah, I forgot his name, that's right.

Pelkey: And it was, what, the Network Information Center?

Kleinrock: Network Information Center. The reason I remember this is - - actually, there's this picture here (leafing through book).

Pelkey: So you recommended Howard to.

Kleinrock: To Larry as having a company that could do this kind of design. They were the world's experts at this time.

Pelkey: And how did you know Howard?

Kleinrock: From 1966. Howie invited me to participate in a short course he was running up in Berkeley, where he was a faculty member, along with a guy named Ivan Frisch. Now they, together, wrote a very important book. Frisch and Frank. They wrote that book. They were well known in what you call 'network float theory;' shortest path, traveling salesman, maximum flow, topological design; and Howie was on faculty along with Ivan Frisch up at Berkeley. They invited me to talk there, which is a side story.

Pelkey: And this book is Communications, Transmission and Transportation Networks and --

Kleinrock: Yeah, but it's basically Graff Theoretically oriented, algorithm oriented, ok? It's not probabilistic. It's how you label things to find shortest paths, and all the rest. It's not probabilistic. Actually, they do some evaluation of reliability, but it's largely labeling techniques. But they were --

Pelkey: Impressive little book.

Kleinrock: They were also going to speak there, in the summer, for a one-day presentation. I said: "Fine," but that week I was up in Sequoia with my family, camping, and the idea was, I'd get up real early in the morning, get out of camp, and drive down from Sequoia to Fresno, fly to Berkeley, give my talk, and come back. Got up at five in the morning with my four day beard, a straw hat, jeans and a checkered shirt, and my clothes and notes in a suit-carrier, and I'm driving out of the park. And my brakes fail, downhill. I had a big old wagon. I step on the brakes, nothing happens. Put on the emergency brake, nothing happens. Put it into low, nothing happened. The transmission was not operating either. This thing was freewheeling down a straight- away, with a curve at the bottom -- I was doing 45 at this point. There was no way I was going to negotiate and almost right angle turn, left hand turn. This was a cliff, and this was sort of a rollway. I couldn't even bang against it. If I had gone up, I would have gone over. I didn't know what to do!" I opened the door, I looked out, I said: "No," put my belt back on, and I yelled: "Help! I actually did." What did you do?

Pelkey: You were alone?

Kleinrock: I was alone in the car. And son-of-a-gun, help came! Here I am, right. What happens is the rear left tire blew -- BOOM! And it stopped the car. It ground to a halt. I got out with my straw hat on, looking at this thing, and all the vital body fluids -- the brake fluids, the engine fluids, the oil, transmission fluid are coming out. I had no idea what happened. A ranger comes by, and I said: "Look, I got to get out of here. I've got to do lecture," because I knew now I've got to earn that money to fix this car. You know, I was poor. So he took me to a gas station, I said: "Look, there's a car back there. Tow it, fix it, I'll be back tonight," and I started hitchhiking. Same ranger comes by and says: "You can't hitchhike in the park." So he took me to the exit. I said: "You mean you're throwing me out?" He said: "Yeah." So now -- I've got to catch this plane, ok? It was leaving, I think, at eight o'clock in the morning. It's now approaching seven, and I'm 50 miles away down mountain roads from the airport. And, of course, nobody picks me up, because I look like a bum. So I went to the guy in the toll both and I said: "Look, explain to them I'm not a bum and that I've got to get to the airport." Next car picked me up, a quiet old couple. I go 'fast,' and he heads down. He's racing down this hill, making the turns, we get to the airport about a minute before the plane takes off, I thanked them, I rush through the airport with my suit carrier flying behind me with my beard and all -- get to the plane, and I get on. Just got on. Whew! Sit down, taxi out to the runway, stops at the edge of the runway, hesitates, and says: "Folks, we have a problem. Our brakes are not working." I said: "NO!" So they tow us back. I made the damned plane. So they tow us back, and I see another plane on this runway. Fresno is not a big place, so I rush back. The same people that watch me come back this way. I say: 'Where's that plane going to?' They say: 'LA.' 'Put me on,' so I hopped on, off to LA. From LA, at least I can get to San Francisco. From Fresno I can't get anywhere. I get up, go to the bathroom, shave, and change my clothes. They knock on the door, they said: "We're landing." It's too early. I got out, got off the plane, it's not LA. It was Bakersfield. I get back on; we fly to LA -- air traffic control. Delays. Late. Get down, I call Howie, I say: "Listen, I'm going to be late." So he says: "Don't worry, Paul Baran is here." He is also a guest speaker the day after me -- he wanted to hear me -- "so we'll put him on." So Paul did the first half-day, I got there for the afternoon. I took the second half-day, and I came back the next week to finish my other half-day. Then I go back to

Fresno, rent a car, drive up to the tent. My family is now asleep; it's late at night. I crawl in, my wife says: "How was you're day?"

Pelkey: When was this?

Kleinrock: '66, I think it was July or August '66. That's how I met Howie. It was funny. I'll never forget that.

Pelkey: That is great. Is that the first time you met Paul Baran as well?

Kleinrock: No, Paul I had known. I think I had met him in 1964 in a meeting in Orcas Island in Puget Sound; it was an ORSA meeting or something. I had met Paul before.

Pelkey: That's a great story. Thank you.

Kleinrock: So that's how I got to know Howie. I recommended Howie to Larry as well. That's what we have here. The problem of technology transfer -- and it was a negative on the Arpanet. Let's see what I have. (Leafing through some papers) I will lend you this thing. What's the date on this, October '79, that's the same thing I had before. Ok, here's an old Arpanet document. What the hell is this now? '69.

Pelkey: Oh, my goodness.

Kleinrock: That's before it all happened. We were talking about the host-host, how they would link up. You may want to borrow this - - but this is not the document I was looking for. Oh, here it is. Steve Crocker. He was thinking about host to host, in fact. Why don't you borrow that and give it back to me. There's not year on this. This is just some mail that I had. "Write a spec for the ARPA net," and we basically laid out a total outline for that, and I think Larry fleshed it out, and then the big RFP went out.

Pelkey: Anyway --

Kleinrock: Where were we? Oh, we were talking about network design. Finally it became apparent that we needed some help with the topology design. We were doing analysis, analytic design and measurement, but --

Pelkey: As I understand from Bob, he came out -- it must have been early '69 -- here to work with you. He thought that the network could get locked up.

Kleinrock: Yeah, there were some lock-up problems. We recognized, immediately, some lock-up problems.

Pelkey: And he came here to sort of test and experiment on the network to find out how it really performed, and whether or not you could lock it up --

Kleinrock: The idea was that we were going to create this four-node network, and then spend about three months shaking it down, and it was the end of December or early January that we started doing that.

Pelkey: Of '70?

Kleinrock: Yes, and I, at that time, had a student devoted to doing network measurement as well. We performed a number of measurements on a network using the code that we had forced BBN to write and we spec-ed that code. We told BBN: "We got to have measurement software in there," and they created a lot of software, which is described in this book.

Pelkey: So they put the hooks into the code that was in the IMP?

Kleinrock: That's right, and it's described in this book. The point is, between this network and this network is when Howie was hired. I can tell you that because we had been simulating the heck out of these networks as they came up, and we found that original line, the first line in the ARPA network, was not carrying much traffic, by simulation. Howie did his topological design stuff, and he also found it was not needed, and you notice this network doesn't have it. It's easy to add to nodes to a network. You can be a hero. You take something out and the network fails, it's your fault.

Pelkey: Absolutely.

Kleinrock: So we both agreed to take it out, and it was fine. Howie also, as one of his points in network design, he says: "Network design is really hard. A small mistake can cause some serious problems. As a result, you don't want to design a network on the back of an envelope." Look at that network. (Leafing through book) It looks like the back of an envelope. This is a joke.

Pelkey: It does look like a back of an envelope.

Kleinrock: In fact, that's just the way -- it's a logic design, of course.

Pelkey: That is a good joke.

Kleinrock: Anyway, and then it continued, and off it went.

Pelkey: Help me understand, on this network, when Howie got involved --

Kleinrock: '71 was it?

Pelkey: Yeah, 3/1/71.

Kleinrock: Now, he came in sometime between that and the April '72 date.

Pelkey: Now Bob Kahn and Howie and you used to get together?

Kleinrock: No, actually the three of us wrote a paper in '72, I believe. Jerry Cole did a dissertation on measuring the ARPA network performance using the tools.

Pelkey: And Fultz is using --

Kleinrock: Fultz is doing simulation. I got the dissertations here. They became very important works -- math model and all the rest, outputs, implementation -- we tried to make a Lincoln Wand, an invention of Larry Roberts' by the way. Fish-eye was an invention of mine. We even had -- (unintelligible) as to what would happen, and most of the focus was here, of course. Now, let's see if we can find the host-to-host being finished. Where are we? '69, I guess. Some host-to-host -- this was supposed to be done, you see, in the fall of '69. It didn't happen -- the full implementation didn't happen until October '71. They just dragged and dragged.

Pelkey: Now, to Larry, this host-to-host was a source of great frustration.

Kleinrock: Absolutely. What happened is, we made a mistake, like everybody else does. We, first of all, underestimated the level of effort needed for the software project. Secondly, we gave it to a bunch of bright -- in some sense too creative -- graduate students who don't understand a schedule. They've got their own studies to be doing during the academic year, so it sort of quieted down during the year, in the summer it got going, and they dragged in a lot of people from other facilities, which is great, but it took two years to get it -- well, two years before it was fully implemented. What happened was, they had to create a spec. After the spec, you have to get the spec approved, and then installed. Every site had to

do it's own installation, with a different implementation. They just gave us a high-level spec, and that slowed down --

Pelkey: Now, at some level, these people reported to you -- Steve Crocker presumably --

Kleinrock: But I'm not a software engineer, so I depended upon --

Pelkey: Right, and he said that you interacted a little bit, but he wasn't managed at all.

Kleinrock: That's right.

Pelkey: I asked him about the interaction with Larry, and he said: 'Two or three times a year – ' I said: 'But Larry expressed that there was a great deal of frustration. Wasn't there more interaction.' He said: 'No, we knew we had to get it done, but it wasn't that tightly managed.'

Kleinrock: You've got to understand, that Steve, in particular, his personality is such that nothing ever gets done. He's one of these guys who adds bells and whistles all the time; no to fault him, but he's just too creative, and he never gets closure. It was drag, and drag. These guys were bright! Steve is a great talker, with great ideas, but doesn't ever carry out, and it permeated the whole group. It was not managed, and I couldn't supervise it, because that's not -- I'm not a software designer.

Pelkey: Did you and Larry talk about this issue, the host-to-host?

Kleinrock: Yeah, but -- it was a pain in the neck, but it wasn't a first order urgent problem as presented to us. The network was growing. Who was using the network at that time? People who had learned to use machines -- let's think about the Utah machine. They would then take a job at Stanford Research Institute, wanted to continue to use the old machine, and they knew how to use it. They could basically get through the network without a host-to-host protocol, because they knew how to get into that host.

Pelkey: So they had remote login --

Kleinrock: Exactly. Remote login and off they go. That was what was going on. That was the main use of the early network.

Pelkey: Was there E-mail up at this point in time?

Kleinrock: I forget exactly when it came up. I saw in your notes that Larry designed the first E-mail. As I recall, BBN did.

Pelkey: The guys who did 10-X --

Kleinrock: It was an add-on. A little add-on, and it took off like that.

Pelkey: Larry did a thing called TECO-Hack.

Kleinrock: Ok.

Pelkey: There was a library that contained messages. Actually, Larry's perspective on E-mail is very different than anybody else's. Everybody else remembers E-mail as not having been part of the original specification.

Kleinrock: That's right.

Pelkey: And it came as an accident.

Kleinrock: Yeah.

Pelkey: Larry says: "Well, obviously that was there. The intent was E-mail."

Kleinrock: It may have been in his head, it was not explicit.

Pelkey: Because it just kind of happened.

Kleinrock: That's right.

Pelkey: So E-mail was a great surprise to the network.

Kleinrock: Oh, yes, absolutely.

Pelkey: Like Paul Baran says: "The experiment that went bad."

Kleinrock: That went good.

Pelkey: Right, but it was meant just to be an experiment, but it became so important that he couldn't get rid of this experiment, and it took on a life of its own. It was too successful an experiment.

Kleinrock: (Leafing through a book) What I'm looking for here is the growth of traffic in the Arpanet book.

Pelkey: Howie got involved in '70?

Kleinrock: '71.

Pelkey: One of the things that Bob was sharing with me was a different perspective that Howie brought, you brought, and Bob brought to what a network was. You're heritage, your intellectual understanding and all these are networking from a very different perspective, and reconciling these, which -- some of it was in this paper that the three of you wrote.

Kleinrock: That's right. That was a hard paper to write.

Pelkey: Yes, he said it was a very difficult paper for the three of you to write.

Kleinrock: Well, Bob is a stickler, also. He's a hard guy to work with, Bob. He's very thoughtful, but he can't stay focused. Things interwove, but there's an issue here and an issue there. Hey! So that's the only thing that held us up. I was going to try to get you a map of the Arpanet traffic that is revealing, here. (Still leafing through book) All these are measurements I'm passing through here. There it is. Now, you notice, the traffic pictures begin in October '71, when the host-to-host was installed, and then it was only 100,000 packets a day. It was really quite small before that. Then it grew exponentially, log scale here. This blip was Tinker and McClellan doing their own host-to-host protocol, basically, in this interval. This was October '72, which is the ICCC -- a lot of free accounts were let out, and weren't removed for a month or two, so people kept using them. Then it began to saturate. The saturation was due to the fact that the hosts couldn't offer more service. The IMPs had the capability, but the hosts couldn't. This picture stops here when the book was published in '76, and then when they released the requirement of no more than 64 IMPs, it began to take off again.

Pelkey: So it's about five times ten to the sixth.

Kleinrock: Yeah, now it's -- do you get the printout of the traffic statistics? I've got lots of them.

Pelkey: Maybe I'll get back to you for that. In light of having twenty minutes, now, you and Howie and Bob dealt with the issue of the topology and designed networks and traffic --

Kleinrock: We worked some together. We developed a design algorithm here, which is reported in my book -- they had their own design - - topologic was an algorithm -- very different, leading to very different topologies, both with the same performance. So we were doing analytic algorithmic studies based on delay and queuing theory, they were doing it based on topology, but theirs was the one that was being used. They were hired to do it. We were interacting a lot, Howie and I, Kahn and I, not so much, Larry and I a fair amount. Larry and I were very good buddies, by the way.

Pelkey: Give me, if you would, some examples of the perspective that you brought to this process --

Kleinrock: Measurement being key. Nobody was pushing that but me. I got it, and it turned out to be very important. It's from that - - see, we became the Network Measurement Center. I forget when it was, but we stopped in '75. I can't remember when we started, but early on. It was our job to not monitor what was going on, that was BBN's function as a control function, but to experiment with the network, which in our mind meant get performance profiles, get limits of performance, and push it as hard as we can in weird configurations of traffic. We had artificial traffic generators, both in the IMPs and on the hosts. So we began to break the network any time we could to find weaknesses. BBN hated us for that. Frank Heart was the manager, and he was -- he's an interesting guy, but he was trying to protect his position running the network, unlike Dave Walden and McKenzie and McClellan; they were a little friendlier to us. We would find a fault with the network, but we were not given the IMP code. They wouldn't release it in the early days. They considered it proprietary. You know that story?

Pelkey: About PCI and --

Kleinrock: PCI? PCI was the first network that went bankrupt.

Pelkey: They're the ones that tried to --

Kleinrock: To get it, to release it, right. Finally, it got released. Until then, here we are measuring the damned network, observing problems from a measurement point of view, telling BBN: "Look, there's a problem here. We think it's this. Fix it." The first response would be: "Oh no, that's a six month fix. Go away." So then one day they release the code, we told them, now, not only what it was, but how to fix it. We wrote the code for them. It still took six months, because their point, more or less was, if the network isn't broken, don't fix it, and some of these things we found were potential things.

Pelkey: They became the new AT&T?

Kleinrock: Exactly! Hard to deal with; it was awful. We told them, in the early days, when they came out with their proposal, we looked in their routing algorithm, which we were very interested in, and said: "This thing has loops. A packet can go back and forth, or in a circle," and they said: "No, there are no loops." There were LOOPS. They said: "No." So we simulated and we showed them and said: "Look, there's loops." They said: "There are no loops." Then this (pointing out something) network came along, and we showed them loops, and the answer was: "Yes, there are loops."

Pelkey: So measurement was something that you brought to this.

Kleinrock: Absolutely, and simulation, and analysis, and our own design versus Howie's design.

Pelkey: What have been the implications beyond Arpanet of those concepts?

Kleinrock: Well, first of all, we were able to tell, then, for example, what the phenomenological behavior of a network is. You give me a network, I can analyze it's performance, and in a very simple way. We developed exact mathematical formulas, we've got some approximations that are easy to use, we've got design procedures which didn't really impact the field, because Howie's company was doing the same kind of thing but different, but we learned things from this. Measurement was very important. We were able to demonstrate deadlocks and degradations, explain them -- they're all detailed in this book here --

what was causing them. Then we found that the whole issue of flow control was a key issue. For example, sequencing, keeping things in order, causes dead- locks. See, flow control; the function of flow control is to make things move smoothly. If you want to do that, you put in controls. Once you put controls on, that's called a constraint. If the network can't meet the constraint, it crashes. So the thing you put in to help it out is the thing that kills it, or degrades it, so we have a whole catalog of degradations and deadlocks that they eventually fixed. Those things are still present in every network today.

Pelkey: And in the late '70s and '80s, what are some of the things that have come about as a consequence of the work that you brought to this?

Kleinrock: One of the things that came out is a kind of a left-handed by-product. It was a whole army of my graduate students. In fact, they're listed on the board behind you. I have now graduated 26, 27 students. Those students formed a cadre of people out there who knew about packet networks. Some went to industry; some went to universities who produced yet more students. It created a cadre of excellence in network design, performance evaluation, measurement, innovation, etc., and the focus on doing analytic modeling of these things before you spend the million dollars, and it helps a lot, in topology and capacity assignment --

Pelkey: After '71, '72, what did you do then?

Tape Side Ends

Kleinrock: . . . Notice I'm being supported by ARPA since 1969. They got interested in satellite stuff, started putting my research in that; packet radio, a bunch of students did that; created some very fine people who are out there in the universities now; that's the time Metcalfe was sort of involved with that early point. Some of his ideas came out of these broadcast access methods.

Pelkey: When did you first meet Bob?

Kleinrock: I can't tell you exactly; it was probably '69, '70, in that range.

Pelkey: Did you meet him in conjunction with ARPA?

Kleinrock: Yes, and also part of Xerox PARC. We were interacting with them because of the network, but I met a lot of the other guys. I met Allan Kay up there, Englebart, that whole crew, and he came with the idea of CSMACD on a wire --

Pelkey: When did you first hear about that?

Kleinrock: Trust Bob on that. It was early '70s. I can't tell you the date. He recollects meeting you at Washington Airport. We met there; we also discussed it at the ICCC, the first conference in '72.

Pelkey: He remembers you as being somewhat critical.

Kleinrock: The problem was it wasn't clear that -- we were doing satellite experiments at the time, satellite thing, and he was suggesting it for an environment where the satellite -- if this is the propagation path, a packet is that big. It goes up and comes down. Listening to a packet is ancient history, and it wasn't clear, because he hadn't specified the parameters of his system -- he mentioned something like a ten-megabit kind of thing. He had in mind a packet where only a tenth of it would be on the pipe. You start pumping it in, first it comes out before most of it's got in; so what you listen to is current. Now he didn't give me the parameters of the problem for me to realize that and switch from this kind of a picture to that kind of a picture, where the packet's this big instead of this big. Now, it's got to be this big to be (unintelligible). In fact, 'carrier sense' still suffers from that problem. If the packet gets too small, effectively, if, namely, a twentieth of it is in, can fill the pipe, you're already in trouble. It's got to be less

than a twentieth, roughly, and Ethernet is doomed to failure, because as you go to fiber optics, the packet's going to look this big again. When you go to longer cable lengths --

Pelkey: Yeah, Saltzer and Clark are clearly into the token ring and the deterministic network as opposed to probabilistic statistical networks.

Kleinrock: Well, I love the statistical part, because that's the environment I live in, I just love that. Anyway, I didn't say no to it, certainly, I said: "You've got to make sure what you listen to is current."

Pelkey: Do you recall that meeting in Washington?

Kleinrock: I don't remember the airport meeting in Washington, frankly. I know we discussed this at ICCC '72, and I remember meeting him in Paris one time when my plane got detoured -- I forget when that was. It was around then also, but the Washington meeting I don't recall. I remember the conversations, more than one about it. So what did I do? Then I went to packet radio, and then to local area networks, as each of these -- now, local area networks is not really something that DARPA, per se, was supporting, but my research naturally applied there, and I did a lot of work in that.

Pelkey: This is, what, early '80s?

Kleinrock: '77 I had a key paper in this area, so I was already doing local networks by then, and since then we're doing that. Now I've moved completely over to distributed processing. I'm no longer doing networking, per se. I sort of recognized, just before DARPA stopped supporting networks, I recognized the field -- see, whenever a field gets such that there's more papers in the field than I have time to read, it's time to get out. So I keep trying to stay at the front tier instead of -- Some people want to stay behind and fill in the holes. I want to get out there and do new stuff. So anyway, let's see if I can. Here's ARPA monthly report, October '72. Spent the bulk of my time working on ICC arrangements. This is for Steve. I made everybody write these things up. I don't know it that's important to you or not. September '72 -- Network measurement supervision. Vint was helping me with network measurement at the time. ICC planning -- that was the month before. Ari was my administrative head. Bill Niello was the guy doing network measurements. We were doing a whole bunch of measurement, a lot of stuff under that.

Pelkey: What is this Tinker The Air Force base?

Kleinrock: Tinker McClellan paper describing the traffic that they measured on that. I don't know if you want this or not. I'm just trying to give you a feel of what was going on at that time.

Pelkey: Is there anything that we haven't touched on during this period of time that strikes you as relevant? We haven't talked about the layering.

Kleinrock: Layering was a product of these software guys, no question. It was a natural evolution on their part. The comment that AT&T didn't get back into this game is accurate. They're still not big, such as Net-1000, that billion-dollar disaster kind of thing. In April of '78 they were going to come out with the Bell Data Network, they never did that. Advanced Communication, ACS didn't happen, so they really missed the boat. IBM sort of crept in --

Pelkey: And they missed the boat because they were circuit switching?

Kleinrock: They didn't believe there was any money to be made in this game, and they didn't devote any effort into packet switching. Then, when they tried to do it, I don't know why they couldn't pull it off. They should have been able to just try carbon copying. In early '80, they tried to create AUTODIN II.

Pelkey: That was Western Union.

Kleinrock: I was consulting for Western Union at the time.

Pelkey: That was two years, and then closed it down.

Kleinrock: That's right. It didn't meet specs; one of the reasons was the government changed the specs on them. They said: "All packets will be full length," and it completely threw the statistics off, and I warned them about that.

Pelkey: Then they went ahead and implemented a replication of the Arpanet. During this period of time, ARPA went from ARPA to DARPA. This is a research question. ARPA went to DARPA, which slowly, but importantly, has reduced the amount of research dollars that are falling --

Kleinrock: Significantly, and to the mentality --

Pelkey: And the mentality. While NSF was supposed to pick it up, NSF is for the solo practitioner, as opposed to --

Kleinrock: That, and equipment grants, the big CER grants, experimental computer science, where they give millions of dollars in a lump, for five years kind of thing, but it's typically --

Pelkey: But they're not project oriented. One of the things about ARPA, and the Arpanet, to hold that out, was the connecting of the community in a much larger project. One might argue that parts of it were poorly managed. It brought an incredible number of people together with different disciplines --

Kleinrock: With unpredictable by-products, which were enormous.

Pelkey: Right, and there aren't a lot of those going on, to the best of my knowledge. The second thing that happened to the communications business was, after deregulation, the charter of Bell Labs change. So between Bell Labs and ARPA, really ARPA funding, for example it has been asserted that 60% of computer research in the early '60s came from ARPA. In addition, Bell Labs has always been a major source of research funding, how effectively spent one could argue, but at least research. If in fact, research leads a decade or two or three later to economic products and growth by society, where are we, as a country, investing in fundamental research that's going to lead to a source of economic growth twenty or thirty years from now?

Kleinrock: Well, this National Research Network, this thing I mentioned before that I just finished the review of, does has a plan to get us to gigabit network speeds by the end of the century, and that speed of networks is a totally new research slate, which has not been accepted or funded, but that could be a source.

Pelkey: Now, I've argued -- I think that SDI Test Bed project that they let to Martin-Marietta, they should have put that to the university community, as opposed to burying within the industrial sector.

Kleinrock: But SDI is not easily accepted by the university community unless than can partition off any taint of SDI --

Pelkey: But it would seem that the building of the network and the applications they're running on the network are two different issues.

Kleinrock: That's true.

Pelkey: At some point there is a link.

Kleinrock: Somehow launder the money. Launder the process.

Pelkey: Yes, so the concept is acceptable to the universities.

Kleinrock: There's not much. DARPA is making it hard. DARPA stopped with research.

Pelkey: Super-conductivity is --

Kleinrock: Did you see the paper this morning?

Pelkey: No.

Kleinrock: The head of the National Academy of Sciences, Frank Press, came out and said: "We in the scientific community are doing a disservice to Congress by not telling them what the priorities should be. We've got a $16 billion space station, a $4 billion super-collider, a $3 billion genetics engineering. We're not telling them where to spend their money."

Pelkey: Absolutely, we're not.

Kleinrock: And he said: "It's ok if those things get postponed, and in fact this network that we're proposing which is billions of dollars a year and we're hoping will in fact get in there, because that serves the entire community. It's an infrastructure that the Japanese can't copy. Once you get that, you've got a communications network that scientists need, you've got the --

Pelkey: Infrastructure building. I agree with Bob Kahn, with what he's trying to do in that regard.

Kleinrock: And he may, in fact, be an organization which helps implement this. He's got a conflict of interest, being on the committee as well. He's also on my review committee.

Pelkey: But this is a profound issue for us as a country.

Kleinrock: I agree, and it's not being managed. No one is worrying about it, totally, and this CSTB group that we have is worrying about it, but can we get the ear of Congress? It's not clear. That's where I'm going this afternoon. I'm catching a plane to Washington D.C. I am worried about it, very much.

Pelkey: I am too.

Kleinrock: The one thing that we have that the Japanese don't is that innovation engine. Their universities are weak. Their entrepreneurial capability is poor. Their taking it from us -- so what they do is they come here and they take advantage of our stuff, and they take it back.

Pelkey: That issue also, how many of our graduate students are foreign students versus domestic students?

Kleinrock: No, we win that one. Look at this list of students. I would say 60% of them are foreign students who have been captured here in the United States and contributing in the first ranks. Some go back. You get very few Japanese students. I have one Japanese student. He went back. He's working for IBM Japan. He's very good. Very few Japanese students come to America. You get the Chinese, Koreans, Vietnamese, but you don't get many Japanese, which is good. Let them stay there.

Pelkey: And you agree, then, that when ARPA went to DARPA, with the Mansfield Amendment, over time, while people got around it for a while --

Kleinrock: It finally killed it. We had such a hard time in the last few years. We don't know what's going to happen now that [Jacob T.] Schwartz is managing the IPTO effort. I don't think it's going to be good, but with Saul Amarel in there, it was nice, and with [Robert] Cooper in there, well he killed networking. That's why Bob left. You know the whole story.

Pelkey: I don't know the story, but I know he left. It's almost time to conclude our interview. Thank you very, very much for your time.

Kleinrock: If I have any other thoughts, maybe I'll give you a call.

END OF INTERVIEW

