Oral History of Robert Kahn

Interviewed by: Vinton Cerf

Recorded: September 30, 2006
McLean, Virginia

CHM Reference number: X3699.2007

Vinton (Vint) Cerf: It's 30 September 2006, and [I'm Vint] Cerf off camera, interviewing Robert E. Kahn on the occasion of the Computer History Museum Fellows Program. Bob will be inducted as a fellow later on in the year. This interview is for purposes of capturing his history in our Computer History Museum archives. The questions that I'm going to ask Bob initially are more or less a pro forma set of questions that we ask everyone, after which we will go into much more specific questions about Bob's history so the people who are interested will have some specifics. But let's start out, Bob, with some very simple straightforward things. Where did you grow up?

Robert (Bob) Kahn: I was born in Brooklyn, New York, I'm told. Don't recall the exact moment but I recall the history. I lived in Brooklyn until I was about 13 years old.

Cerf: Is there any reason to think that you weren't actually born in Brooklyn and you were only told that?

Kahn: No, I think I was born in Brooklyn.

Cerf: All right, fair enough.

Kahn: There are some people who can remember, they claim, back then. I sort of remember my youth from about the age of six or seven, and I have sort of vague recollections of things before that. But we lived in Brooklyn until I was about 13 years old, lived in a rented house, in my recollection, on 19th Street in the part of Brooklyn that I guess would be considered Flatbush. Before that we lived on Ocean Parkway and a few other places that I don't recall at all. I'm just told. We moved to Flushing, Long Island in 1952 or 1953 -- 1952 I think -- and lived there until I went off to graduate school.

Cerf: Well, we'll come back to that in due course. Can you tell us a little bit about what your father and your mother did? How were they employed, if they were, and what was that work?

Kahn: My dad was a high school administrator during all the time I was growing up. When I ended up moving out he was principal or assistant principal at a high school in Brooklyn -- Tilden High School. He had started there back in the depression teaching accounting, which is what his college degree was in, and he rose up through the ranks. I think he was Dean of students and he had a few other administrative roles, but he was basically a high school administrator. One of the interesting things that he had done was he teamed up with another gentlemen in the New York area that lived around the corner from us who owned about a dozen nursing homes in the New York area. My dad used to do business advising during weekends to sort of supplement the income.

Cerf: This is while he was still acting as a high school administrator?

Kahn: Yes. His background was in accounting and he was really good at analyzing the economic efficiencies of businesses. So he would supplement his income, which was not that large in the school system but was steady income, I guess, which was important during the depression. He used to supplement that by consulting with different businesses. A bakery that was marginal, and he'd explain how if you re-routed the trucks he could save money, deliver more bread. If you organized your accounting this way, if you made people consultants rather than employees, if you took out one plan versus another you could make it more efficient. He used to do that with some regularity. He teamed up with this gentleman, whose name was William Kellman, and my dad worked with him to set up a system to help him manage his nursing homes. He was having trouble even figuring out who his patients were, what supplies they had, provisioning for the future -- the sort of thing that you really need to do to stay on top of your business. He set up a paper forms system of managing nursing homes. They started to use it and it turned out to make a big difference. So, Bill Kellman, who was really the person who knew that business well, introduced that to some of his colleagues in Connecticut and New Jersey and Pennsylvania, some of the nearby states and it didn't take very long before they were in all 48 states of the United States. Of course later on it was all 50 states when Arizona and Hawaii came in late in the '50s. This was all in the late '40s, early '50 time frame that that happened. The company was called Macrobiotics, which was a name that didn't mean anything I believe back then. It later came to mean a significant amount, especially in the dietary regimes.

Cerf: Sure.

Kahn: But, it was a little business that they grew significantly. When Medicare came about in the early 1960s one of the things that they were in possession of was probably the most comprehensive list of nursing homes and names of administrators of those nursing homes, number of beds, and much of the information that the government wanted. They made that available to the government because they were the only ones that really had it. The government, of course, set up a system of requirements for reporting that applied to both hospitals and nursing homes. They were not in the hospital business. That was sort of a too big business. But the nursing home was a little niche business and they focused on that for their whole existence. It turned out to be a small but very profitable little business. I don't think they ever did more than $1 million worth of gross revenue, but the profitability was huge. My dad ran that business until the day he died.

Cerf: That's amazing.

Kahn: At age -- I think he was 89 or 90. He was still running it at that time. I've actually kept the business going. It's sort of a one-person business right now. It's really small but it's sort of like a government bond; it keeps producing a small percentage of revenue.

Cerf: That's fascinating. Some things have a life of their own, don't they? What about your mother?

Kahn: My mother was a very interesting woman, very interested in education, very socially minded. She was always trying to help folks. But in her 30s she had the first of what would have been, I guess, a total of eight heart attacks.

Cerf: Good grief!

Kahn: She probably had rheumatic fever as a child, and it really came to pass. The last of those heart attacks was so bad that the doctors told her that she would not survive another one, and they were coming at a frequency of about one a year. So she ended up going in for closed heart surgery in 1953, shortly after we had moved to Flushing. Closed heart surgery is a very different kind of surgery than open heart surgery, because with the open heart surgery they can keep the blood circulating in the body and they have more time to do the operation. In closed heart surgery, which they had to do back then because they didn't have those machines, they literally had to shut down the blood supply. So they stopped the heart, they stopped the blood flow. In five minutes you're dead and, in fact, I think they said after 90 seconds you start to suffer …

Cerf: Some side effects.

Kahn: some lossage in the brain in certain functions. So they literally set a budget of, like, 90 seconds or maybe 60 seconds, to literally do the whole surgery. So they could get you prepared for it but by the time they go in and do the surgery… This was valve replacement. Or I guess in those days it wasn't replacement, it was [that] the mitral valve had closed due to scar tissue so they had to go in and cut out the scar tissue.

Cerf: That's incredible they could do that in such a short amount of time.

Kahn: I didn't actually see the operation but I'm told they did it with little Exacto knives, those little things the model airplane guys use to cut balsa wood. Little knives on the end of their finger. They'd sort of go in, do the surgery, and then they have to get it all cleaned out because even a small piece of detritus in that system can cause blockages and strokes and the like.

Cerf: Sure.

Kahn: The mortality rate, I am told, the week of that operation was 50 percent.

Cerf: Good God!

Kahn: You had a one chance in two of surviving the surgery that week, and that the mortality rate was 95 percent in five years. So, you didn't go through the surgery unless those were good odds relative to the alternatives.

Cerf: There weren't many other choices.

Kahn: So, she did and not only did she survive it, although the first five or ten years were really very difficult. She was bedridden almost the whole time. Of course, she had been bedridden for the eight years or so before. She finally came out of it. She actually lived until she was 80 years old. She died in 1994 after having gone through the third! She went through a second heart surgery in 1982 at New York Hospital. That was open heart surgery, by comparison a piece of cake. It's pretty invasive surgery, but that was actually valve replacement. Then she went through the same operation again in 1994. She actually came out of the surgery okay, but there was so much damage to the heart muscle it couldn't continue to pump on its own, so she never actually made it out of the hospital. That was unfortunate. But, I do remember her saying that-- or my dad relating to me -- that the first of her heart attacks occurred on April 12, 1945, which is a significant date because apparently she heard about Roosevelt's death on the radio and immediately went into heart failure after that. It must have been just enough to trigger her over the top. Growing up as a kid, my sister and I -- I have one sister, Diana, who lives in Cleveland with her husband Harris -- my parents really weren't able to deal with the house. My dad was busy with work and caring for my mother. My sister and I actually had the run of the household growing up as kids, so this was rather an unusual childhood. You're making dinner every night for your parents and things like that. I wouldn't say we did it all the time. Eventually my mother got into form where she could actually take over and do that stuff. But it was an interesting childhood. She was actually an extremely bright woman but I think she was required to go out and work when she was growing up to bring in income for the family, and so she never went to college. I think she always regretted that. In later life she actually went back to college to get degrees of one sort of another. I think the things that she was proudest of, other than her kids and her grandkids, was the fact that she successfully got through all the college courses later in life. Here she was with teenagers and here's a 70-year-old woman getting college degrees.

Cerf: That's a great story. Now, you mentioned a sister. How many brothers and sisters do you have, Bob?

Kahn: Just one sister, Diana.

Cerf: Is she younger or older?

Kahn: She's about two years younger. Diana was a mathematician in her earlier days. Things came pretty easy to me in school and I think they did for her too, but she was probably competing with me at some level. So every time I would go out and get a 98, she'd get a 99 in the same course. And every time I'd get a 99, she'd have to get 100, whatever it took. She would do that. And I don't think it was obsessive. I think it was just innate. She ended up getting one of the best records I've ever seen. She graduated valedictorian for her high school, class of-- same one I graduated from. I was probably eighth or something in the class -- I don't even remember anymore -- but she was first. When she went to college she went to Queen's College and she graduated first at Queen's College. Then she went to the University of Chicago and got a Ph.D. in mathematics, but she ended up getting more interested in health-related issues. She came down with diabetes in her 20s, Type 1 diabetes, which they claim might be…

Cerf: That's unusually late to show symptoms.

Kahn: For Type 1, yeah. This could be typically caused by a virus or something like that. Bottom line was somehow she lost interest in mathematics, although her son ended up getting a Ph.D. from MIT in math and so the legacy follows on. He married a woman who got a Ph.D. in mathematics, so she got two kids that are mathematicians.

Cerf: We'll probably come back to that for a couple of reasons, not the least being the research work that you did at Princeton, because that had a strong mathematical flavor to it too. But I wanted to ask you about some other relatives. You have another relative named Herman Kahn but I don't remember exactly what the relationship is. Can you tell us a little bit about that?

Kahn: I met Herman only twice in my life and I remember both times, because when I introduced myself to him, he sort of got very flustered and I wasn't sure why. Herman was at the Rand Corporation for many years, and he wrote some of those famous books; "Thermonuclear War," was one, I think "Thinking about the Unthinkable" was another one. Many people thought he was the caricature for Dr. Strangelove. Very heavy set guy, overweight. He was in my dad's generation and so that would make him probably some kind of second cousin once or twice removed. I don't remember exactly how they related, but I think his father was the brother of my father's father, if you can figure out the connection.

Cerf: Yes.

Kahn: So, he was sort of not in the main line of the family but one or two removed over to the side.

Cerf: It makes him a cousin of some kind or other and only women with the right DNA know how to figure out what that computation is, I think.

Kahn: I guess I can relate to that. But I only saw him twice in my life. I don't think I ever had a really serious conversation. The twice I saw him, once was when he gave a talk up at MIT. I was in the audience and I asked him a question. And the other one was when I was at DARPA and he showed up there one day and I shook his hand and mentioned that I was a distant relative. And he said, "Okay, nice to meet you" and then walked off. So I never had a substantive conversation with him.

Cerf: It sounds a lot like my interactions with Bennett Cerf.

Kahn: Your son.

Cerf: No, the other one. Let's move to a slightly different topic.

Kahn: Oh, there was actually another relative in my family...

Cerf: Oh, who did I miss?

Kahn: …who's got a very interesting career. His name was Moe Berg, and if you ever heard of the book by J.D. Salinger called "Catcher in the Rye."

Cerf: "Catcher in the Rye"? Yes.

Kahn: They wrote a book about Moe Berg. It was called “The Catcher Was a Spy." If you've been reading The Washington Post the last few weeks they had a comic series on him, so this is how well known he was. My wife Patrice has been clipping it out, keeping track of it. Moe Berg, he was never married, and [was a] very interesting and perhaps introverted kind of fellow. Again, I never met him. He's on my mother's side of the family, whereas Herman Kahn was on my dad's side. He ended up becoming a lawyer, and he worked to go through law school by being a catcher for the Brooklyn Dodgers, among other things.

Cerf: Oh!

Kahn: And I think he joined some other [ball team]; I think he may have been with the Washington Senators or some other teams. And he caught all the all star games. He was well known in baseball circles. I forget, he got a degree at either Columbia or the Sorbonne. He became fluent in languages. Anecdotally it's said -- since I never met him I couldn't verify this directly -- that he spoke, was fluent in something like 27 languages.

Cerf: Good heavens!

Kahn: Of which some number he was so fluent in you could not tell the difference between him and a native. He knew all the idioms and the expressions and the lore and the history and the like. He was recruited into the U.S. government during the Second World War to assist with some translation. There was a play that was recently put on that you may have seen called "Copenhagen."

Cerf: Uh huh.

Kahn: Very well known, was all about a meeting between Werner Heisenberg, who was then in charge of Hitler's atomic energy program, I guess, trying to build bombs for Hitler, and his thesis advisor who was Niels Bohr, the famous atomic physicist. The whole play is that -- just those three. It was Bohr, Heisenberg and Bohr's wife. He had gone there purportedly to find out something, or talk with him. And they gave you three quantum mechanical views of this little interaction, like: was he there to find out from Bohr what he knew about the Americans' plans? Was he there because he couldn't really make research progress without consulting his advisor? Was he there to try and convince Bohr that if the Americans would stop and they would stop… there were all these different competing views. And of course it leaves the end ambiguous, obviously. Well, Moe Berg was sent over -- and it’s reported in this book called "The Catcher Was a Spy" -- was sent over to Germany during the Second World War with the ostensible purpose, I think, of assassinating Heisenberg.

Cerf: Oh, my gosh!

Kahn: Again I don't know the details firsthand, but apparently he managed to get close enough to him. I don't know whether it was a dinner, or whatever, and he was probably armed with a pistol and a cyanide tablet, or whatever. But he came back with the clear impression that Heisenberg was not a supporter of Hitler getting the bomb, for whatever reason.

Cerf: So, the speculation might be that Heisenberg actually was impeding progress somehow.

Kahn: Well, you could make that… I don't know what the rationale is. I've heard people say that it was reported that somebody who was a paranoid fanatic really shouldn't have access to atomic and nuclear weapons, things like that. And so, he made the decision not to carry out the threat and came back and reported why. Who knows what… Apparently it was a reasonable decision because he never really got those weapons. But that book was all about it. So that's another part of the family.

Cerf: That's fascinating. Let's move back to your childhood for a bit. Did you ever think much, that you can remember now, as a young person about what you would do when you grew up?

Kahn: You know, my childhood was a pretty happy childhood. We were pretty much in charge of some things, and my parents were very understanding. They were great people and I really related to them very well. I don't recall ever having really serious thoughts about the future when I was a kid growing up, because it was just a very happy and comfortable arrangement, despite the fact that my mother was not well. She was there. You could talk to her. They cared. I was sort of mainly doing what I liked to do, which was often go out and play ball with the kids in the street, and get on my bike and ride, and do the things that kids do. I had a very sort of normal childhood, I think, growing up. Although my dad would tell me I was sort of ornery at times. Like, they would go out for the evening sometimes and they'd come back and I would hide from the babysitter. They couldn't find me and they'd all go into a panic. I'd be hiding in the laundry hamper or some such thing.

Cerf: Well, you had mentioned that for a while you and your sister were preparing meals for your parents, which is an inversion.

Kahn: Well, for the family.

Cerf: For the family, I should say; an inversion. But I know that you have a great deal of interest in culinary things and you're very good at it. Do you think that stems in part from your exposure, early on, to cooking?

Kahn: I think probably not. I just think I developed a taste for food later on in life that probably did that. I think the mechanics of working in a kitchen, maybe, I kind of learned back then.

Cerf: And a certain comfort level with familiarity there.

Kahn: But you have to remember we didn't have Cuisinarts and all of the fancy gadgets that you might have today, convection microwave ovens. We just didn't have that. A frying pan and a pot. It was just pretty basic.

Cerf: Basic.

Kahn: Cook a hamburger, cook a chicken, boil a potato, that sort of stuff was what we did. So I wouldn't call this high cuisine back then. But who knows what things trigger you later in life?

Cerf: Did you have any idols when you were growing up? People that you wanted to emulate, or you looked up, to or thought… That somehow influenced your thinking?

Kahn: I was so into myself and the family that the main influences in my life growing up were my parents. My mother was clearly a spiritual influence in the family, because her spirit was very dominant. But my dad was sort of the pillar of the family, and he was just making the family work. I guess to the extent that there's anybody that I looked up to it was my dad. As a dean of students he'd come home at night sometimes and tell us about all the crazy things that happened at school -- this kid that threw the telephone through the window, and this kid that was plugging up the toilets and creating channels down the high school corridors. And I kept thinking, why do people do that? That's sort of silly. But I think as close to anybody it would have to be my dad.

Cerf: What about computers? Now computers weren't around all that much. I know you were born in 1938, so by the time you would have been 13 it would be 1951. There weren't too many things around, maybe Univac. When did you first have any exposure at all to computers?

Kahn: Well, my background actually was more applied mathematics and communications. So I really didn't think of myself as being in the computer or computing field until I started to work on computer networks. That was where I really felt close to computers for the first time. My first interaction with a computer was probably in graduate school at Princeton. That may be wrong. I'm trying to recall. It was one of those early IBM machines where you had to toggle the switches.

Cerf: Oh, my goodness.

Kahn: It was just sort of learning what a computer was, not any serious use. I guess the first really serious interaction that I had with computers was at Bell Laboratories. The first job I had after graduating from college was at Bell Labs, and I worked at the headquarters of Bell Labs. You'll probably want to get into this later but let me just fast forward since you asked the question now. I was in an office that had six people in it. It was a very big room with six desks. Pretty impressive people. One of the people I shared the office with was a woman named Jeannie Root, and her job was to build some of the computing infrastructure for all of Bell Labs.

Cerf: Wow!

Kahn: I was working in the headquarters building of Bell Labs, which in 1960 was at 463 West Street in New York, right on the Hudson River. It was before Holmdel had even been built. Most of the technical staff was at Murray Hill, but we were in New York at West Street. Jeannie was busily trying to figure out how to bring on to the IBM 704, which is what we were using -- this was not even a solid state machine as I recall, I think it was a vacuum tube machine -- the ability to program in Fortran as well as machine code.

Cerf: Machine language, yeah.

Kahn: Which I guess was FAP or something. It was some combination of machine code and Fortran so you could write…

Cerf: It might have been assembly language by that time.

Kahn: I think that FAP was the Fortran assembly code of some sort. I'm recollecting. It was her job, it wasn't mine. But that was my first experience with the availability of that. One of my tasks back then was to actually evaluate certain switching capabilities that were being developed at Bell Labs, including some ferrite core switches, looking at the materials properties of them and figure out what their failure rates were. I ended up writing Monte Carlo simulation programs, in some combination of Fortran and machine code, to try and do the analysis. So that was both my first experience. I never actually saw the machines except through glass windows because they were all big batch [processors].

Cerf: The fishbowl.

Kahn: And we, of course, were programming on pads and punch decks. And one day later you get the printout back and edit it, and go through all of the iterations that later were made unnecessary by moving to interactive timesharing kinds of systems.

Cerf: I'm sure people who grow up online on the Internet today have absolutely no concept of what it was like, with pigeon holes, with printouts and punch cards. Where did you go to college, Bob, and is there a particular reason why you chose one place or another?

Kahn: Well, I started at Queen's College in New York. I had actually applied to about a half a dozen other schools, and I believe I was accepted at most of them if not all of them. But the reason I went to Queen's College was because of family issues. The situation with my mother still hadn't been resolved. We didn't have much money in the family at the time. The City College -- Queen's College system hadn't been merged yet -- but it was basically free by comparison. So for that combination of factors I just decided to stay local, be with the family, and it was economic for us as well. They had a program, an engineering program, where you could spend two years at Queen's College taking all of the liberal arts, mathematics type of stuff, and then shifting over to City College which was up at 139th Street in Manhattan. And so I spent two years at Queen's College, and then shifted over and finished the rest of the program. It was, I think, a five year program. So it was two years followed by three. I graduated in four and a half, so I graduated actually in January of 1960. I had the coursework finished in December of '59.

Cerf: And what degree did you have? After this five year program this was a Bachelor's degree in mathematics?

Kahn: Electrical engineering. The degree was actually called a BEE degree and it was a fairly broad-ranging program. City College, the whole system, had really a wonderful educational program that was geared toward people in the city, mostly middle class kids, which is what we were growing up. But it was a wonderful educational program. How many people do you know that have ever had experience in large power equipment? We used to go in on, I think it was Thursdays, in the grubbiest clothes we could muster because we had to hook up these big power engines to power consumers. We were carrying cables on our shoulders that must have been three or four inches in diameter and you can barely drag the cables. It felt like elephant trunks that we were carrying around. You'd hook them up and you'd have to be careful because it would be like 1000 amps or huge amounts of current flowing. And you'd learn how to characterize the performance of a squirrel cage motor and a Prony break and all that stuff. Today you don't have any need for it, but back then this was what the field was like. It was a really good background. We learned about transistors by testing them one at a time. Nobody today would ever even be able to get a transistor to test, or think that it was necessary to learn, but back then this was brand new stuff. The whole semiconductor world had only been created in the late '40s at Bell Labs and so by the middle '50s it would be these single case transistors with the three spidery legs coming down, and you'd actually go through all the mapping of the currents and the voltages and the biases. It was really a very good education.

Cerf: These were discrete components in effect.

Kahn: Discrete components. The experience there was about what you'd expect. Most people commuted by train; it wasn't a big dormitory system. It wasn't like a college campus of the type you might find at a place like Stanford, where you spent some time, or any typical university that's a campus-type setting. There was a campus there but it was more for commuting students. There wasn't much of a social life there after hours, and certainly not before hours. Bottom line: it was a very functional education. Well trained people. I think they had some really good folks come out of there. In fact, I've been up there in recent years, and one of the people who is there a lot is Colin Powell, who explains how that's how he got his education.

Cerf: Now that you mentioned Colin Powell, were there other contemporaries that you recall that might be notable in our field?

Kahn: Well, in our field I know Len Kleinrock graduated from there a number of years before, I don't remember exactly how many. In my class, and again I'm talking about the class of 1960, people you might know there's a fellow named Andy Grove, who was a refugee from Hungary, I think.

Cerf: Hungary, if I remember right, yeah.

Kahn: This thick, syrupy Hungarian accent. I think he was a chemical engineer at the time. I guess they just named a school after him because he made a sizable donation. There were a number of really well known people that came out of that because the education was really first rate I think.

Cerf: After you finished your undergraduate work at CCNY did you go immediately to graduate work and, if you did, where?

Kahn: No, I didn't. I was sort of young going through the college. I think I went into college at age 16 and I got out when I was 20, roughly, plus or minus a year one way or another. I thought maybe I ought to go out and get some job experience first -- find out what a real world setting was like. So I applied to a number of places and ended up taking a job at Bell Labs, which was in the New York area, as I mentioned. I ended up choosing a job... there were several opportunities. Most of the people in my field went into what were considered the sexier parts of the field. Like color television, which was just emerging. Some people went into modulation, kind of theory work. Some went into advanced materials kinds of stuff. I opted to join what was essentially an operations research department, which in part was sort of mathematically oriented. This was a group that was responsible for the whole traffic engineering of the Bell system. Where do you put the lines and how do you decide what switches go where? It wasn't about the hardcore implementation of any piece of equipment, but it was: how do you system engineer this thing to meet customer needs? The gentleman who ran that division was one of the nicest people I've ever met in my life. He was a very consummate professional. He reminded me of Walter Cronkite.

Cerf: Wow.

Kahn: He came in with cardigan jackets and he smoked a pipe. He was a mathematician of sorts, but he would draw curves and intersections and figure out, "Well I think we need about this many of that there because this is where the intersection goes." He was figuring out the curves and he really developed all the total traffic engineering theories as far as I know for the Bell system. He was probably the last person that understood where everything was in the Bell system. His name was Roger Wilkinson. He lived somewhere up in the Larchmont area, and I know he would go sail his boats every weekend. He was just a very stylish gentleman. He ran a group of practitioners, mainly mathematicians. You may have heard some of them. John Reardon was a well known mathematician…

Cerf: Yes.

Kahn: …wrote some books on combinatorial theory, was in that group, worked for him. Walter Helley, who was well known for total traffic engineering theories, [did] some interesting analyses of traffic through the tunnels based on queuing theory. We had a fellow named Fred DeClue [ph?] there. Fred later went to Bellcore after the breakup and was one of their well-known mathematical types. I actually worked most closely with a fellow named Joe Webber there. Joe had a long and distinguished career in the Bell system. One of the people that I got to know very well in that time was a young lady whose maiden name I can't remember but her name was Gigi, and she later married Dave Farber.

Cerf: Oh, that Gigi! I'll be darned.

Kahn: People think that because I've had a long relationship with Dave over many years that I knew Gigi because she's a spouse but the reality is I knew Gigi before I knew Dave.

Cerf: I didn't realize that. That's fascinating. Wasn't there a fellow at Bell Labs who was a connoisseur of wine, and didn't you get introduced to some pretty spectacular vintages while you were there?

Kahn: Yes, you know me well. We had a little luncheon group. Because we were working at West Street it was right on the boundary of Greenwich Village, and there are all of these little boutique restaurants. These were not big restaurants ala today's mega style restaurants. These were little excellent holes in the wall, places you had to know where they were. You go down some three stairs, down in the back. We used to go out for lunch every day. It was a group that consisted of myself, Joe Webber, a fellow named Dick Priest, who was the wine connoisseur, and sometimes the fourth one Barry, but Peter Linhart [ph?] was principally the fourth. Sometimes Gigi would join us, and so forth. Dick Priest was a very unusual fellow. I have not seen him. I didn't see him before, and I haven't seen him since that experience. He seemed to have three main interests in life, one of which was wine collecting, and one of which was first editions -- and I'm not sure of what exactly -- and the third was opera. Those were the three things he seemed to care about.

Cerf: I hope he was interested in wine consuming as well as collecting.

Kahn: Well, I can certainly vouch for the fact that he consumed it at lunch, but so did we, with him. He had this big cellar that he kept. I believe it wasn't in his house. It was at some wine store on Madison Avenue that purchased it and apparently kept it for him. He had this just enormous collection, apparently, of wine. What he offered to do was to bring in a different bottle of wine every day. These were all French wines. We never… It was nothing but French. It was thoroughbred French. Every day he'd bring in a different bottle of wine, and we'd reimburse him for his cost of the wine.

Cerf: Cost.

Kahn: The cost -- what he paid for it when he first bought it. Who do you know, in a 20 or a 30-day run of lunches, would start out with Lafite Rothschild 1928, followed by Lafite Rothschild 1929, followed by Lafite Rothschild 1930, followed by… you know. We'd do a run on that brand. And he'd come in with running commentaries. That's where I learned all of the words -- about a nose, and bouquet, and nutty, and fruity -- all of the words that would apply. He'd give us running commentary and I was sort of making notes. But it's very hard to go out and buy a good bottle of 1931 Lafite Rothschild these days, so it doesn't do me much good. Then we'd go on to the Chateau Margaux's and the Chateau Beychevelle's, and whatever it was. We just went through all of them and it was a marvelous experience in learning about wine. It was only one bottle and there were four of us, so it wasn't like we couldn't have functioned in the afternoon. It was a really good experience.

Cerf: What an incredible experience at a rather early age.

Kahn: One of the restaurants I really like the best was a little one called the Beatrice Inn, which was-- you probably wouldn't normally think about it, but it was a company called the Beatrice Foods Company that has all kinds of brands underneath it. I can't even name what they are but you'd probably recognize all of them as household names. Well apparently their family owned a little tiny restaurant in the area called the Beatrice Inn. That was my mother's name, so I kind of had an affinity for it.

Cerf: Fascinating.

Cerf: Bob, let's go back now. We were just talking about Bell Labs, but let's go back now for a moment to your graduate work. After Bell Labs I take it that you went to Princeton. And I'm interested to know a little bit more about what you chose to study there and why?

Kahn: While I was at Bell Labs I had applied for an NSF fellowship. It came in sometime that summer and it was applicable I think at any school of my choice. I had applied to Princeton and I'd gotten accepted there. I chose to go to Princeton for several reasons. One was it was reasonably close to home, so I could still imagine driving back home on weekends, or at least monthly, so I didn't get too detached from my family because they still were in need of some kind of iteration and help. Sometimes my dad would just need me there for some reason -- helping my mother or something. Second, because it represented a change of style that was something that appealed to me. I mean, Princeton was sort of the inverse of City College in just about every way you can imagine. Wasn't in a big city. It was in a little town. It had a nice campus. It also had a nice golf course right beside it, which is the whole of the story I'll mention in a moment. It was a much more social school and I felt like I could have used a little more training in the social graces. I wasn't used to wearing tuxedos all the time and just doing the kind of things that a good social environment could sometimes bring out. So while I was at Princeton I actually volunteered and became in charge of all the social affairs at the graduate school. I would set up all of the cultural events, and would set up all of the social events. I used to invite all the girls' schools to come down from whether it was Wellesley or Vassar or Bryn Mawr or whatever, and I'd be the host. There would be a counterpart on the women's side and we'd arrange all the social events that took place. There was no counterpart for that at City College at all. It just didn't happen.

<crew talk>

Cerf: Let's continue now with the discussion about your time at Princeton. You were talking about being in charge of the social and cultural events. Do I remember correctly, Bob, that you were required to wear academic robes at certain times while you were at school?

Kahn: Dinner in the graduate school was served in Proctor Hall, which was a very long English-looking venue and, yes, we had to wear academic robes to dinner, and we had to say the appropriate graces. That's where I learned how to say appropriate Latin graces because sometimes I had to actually say it, you know, <speaking Latin> and so you--

Cerf: You weren't required to go through the <speaking Latin>.

Kahn: No, this wasn't religious. This was bread on the table and give thanks for the food. It was a very interesting event every night. Of course, most of the graduate students were there in jeans and tee shirts underneath these academic robes. The big problem was these robes had big, long tails on the sides and everything they seemed to serve for dinner had gravy, and so the biggest problem was getting gravy on your sleeves. So you'd learn to learn to reach for things like that. [Demonstrating arm motions.] But those were interesting times. I ran a pretty frenetic kind of existence compared to most students because I kind of learned to function for part of my time there on three chunks of two hour sleep. I was getting six hours sleep a night in three two-hour chunks. I'd get up at 6:00 in the morning and I'd head over to Baker Rink to do my figure skating patches. Then I'd come back and get showered, go off to school. At noon, I'd come back and we'd have lunch. I'd go to sleep for two hours, like from 12:30 to 2:30. Get up at 2:30, head back to the engineering school, work until 6:30, come back and dinner was like 7:00 to 7:30. Then I'd go to sleep from 7:30 to 9:30. Then get up and work until 4:00 in the morning, go to sleep for two more hours and get up at 6:00 and go do my figure skating patches. The funny part of it was that you'd think you'd be tired all the time, but I wasn't. I was revved up all the time. It was sort of like you don't feel like you've been out for eight hours and you got to get started up and get coffee into the system. You wake up and you're going full bore. It was really a very good existence. I can't do it anymore.

Cerf: I remember going through a similar period of time when it was the only way to get a lot done. You got the same number of hours of sleep but you broke it up so you could almost work 24 hours a day. It's terrific practice. You mentioned ice skating, figure skating. Is this when you picked up an interest in golf as well?

Kahn: Actually I had never played golf until I was in high school. My dad had an old set of clubs at home and he used to go out occasionally to play. Those were the wooden shaft clubs, Hickory shafts. After we moved to Flushing, I had a friend, Jeffrey Rothman, and he lived a few blocks away. We were probably a mile or two from school, so we used to walk it every morning and come back at night, just two of us high school kids. I was, I guess, a junior at the time. I started in high school in Brooklyn at Madison and then my junior and senior year we had moved out to Flushing, Long Island. It turns of Jeff was a golfer. He was actually a pretty good golfer. I sort of knew that, but it didn't really get into my inner thinking because it wasn't something I did. Then one day he said he was going to be a little late, could I meet him at such and such a room, which I did. It turned out it was a meeting of the golf team. He came out a few minutes later and basically said to me could I join them inside? Well, it turned out that they had only four people that had shown up and you needed five people to constitute the golf team. So he asked if they could use my name on the golf team in order to just have a roster. He said, "We'll get somebody else." I said, "You have to realize I've never played the sport, don't know anything about it." He said, "It's all right. We'll find someone else." Bottom line, they didn't find anybody else. It came the first match and I was the fifth man on the team. And believe it or not, I actually had something like a 1-8 record that first year. I actually won a match because some other golf team had somebody else in the same situation. The second year I liked it enough that I went out for it again. This time I made the golf team on merit, I guess, although maybe only five showed up again. I'm not really sure. I think I had a 5-4 record the second year. Then when I got into college I went out for the varsity golf team and actually made the college golf team. I was getting pretty good at that time and we played some of our matches on the Black Course at Bethpage where they just held the U.S. Open in 2002. We played a lot of our golf at some of the area courses. I remember playing a match against Iona College at Winged Foot where they played the U.S. Open this last year. They played it in 1959. I think at that time maybe Billy Casper won the Open that year. But we played it on the Monday right after the Open ended, so it was literally the U.S. Open set up. They hadn't changed the pins, the rough was even a little higher. If you've ever played U.S. Open style golf you know that it's not like regular golf. Because in regular golf if you're a little off the fairway, you just play the ball out of the rough. But for the Open, they put a premium on accuracy. So when the rough is that high and your ball gets into it, it's like hitting out of steel wool. You're lucky if you can get it back on the fairway. You can be two inches off the fairway and it's essentially half a stroke or a stroke penalty. I shot one of the best rounds of golf I had ever played in my life. I was hitting everything as well as I knew how. I wasn't that accurate and I shot a 92. And I said, "If that's the best I can do, top of my game, these guys are pretty good." So that was my one calibration with premiere golf. It put me in my place. I knew where I was. Just being good isn't enough to compete in the league. I began to appreciate golfers like Jack Nicklaus and others in later years much more than I would have thought I would as a young kid.

Cerf: Well, I know you've been a great and avid golf player and I thought it was pretty neat when you got the Presidential Medal of Freedom and you got to have a picture with Jack Nicklaus. I know that meant a lot to you and it was fun to watch the enthusiasm. He's a really nice guy isn't he?

Kahn: Well, I enjoyed meeting him. I think he was... well, you got a chance to meet him too. I did get a picture of him; in fact, I have it up on the wall. One of the things that really struck me was that he was beginning to use the Internet, so it seems like a fair trade of some sort.

Cerf: What would be nice, I suppose, would be to have a chance to play a game of golf or a round of golf with him. Maybe you should think about that.

Kahn: Well, I don't know what would be in it for him exactly, so I think I won't propose it. But, if I could shoot like… My one dream in life was to get on a golf course and shoot 18 birdies, a birdie every hole. If I could figure out how to do that sometime I'd be happy to play a round of golf with him.

Cerf: It sounds like you need a GPS receiver on the golf ball plus a gyroscopic control and--

Kahn: And a computer program, yeah. It's one of those fantasies. You play a really hard course in your mind and you say, oh, drive down the fairway two feet from the hole, knock the putt in, do the same and 18 holes later you've got 18 under par, no problem. It doesn't happen in practice.

Cerf: When we talk about CNRI we'll talk about microelectronic mechanical systems and maybe there's a possibility of an 18-hole birdie. Let's come back for a moment to Princeton. Were there people who -- either at Princeton or at CCNY or Queens -- were there people that particularly influenced your thinking and your career while you were in college?

Kahn: I would have to say yes. At City College, the person who most affected me and influenced me was a professor by the name of Egon Brenner, who is now retired and I think he's in Florida. He actually lived not too far from us in Queens and would often drive me either to school or back, so it was a very interesting opportunity for a student to get to know a faculty member. I was just an undergraduate at the time. He later went on to become dean or provost at some of the other schools, so he had a career of his own. He wrote one of the outstanding textbooks in circuit theory with a fellow named Javid Monsour. But I remember in the final year at City College what Egon said. It was very observant of him. I was a student that had never really been motivated by school that much, because the learning was just something I knew I could do, and homework was something you just had to do because they required it. Oftentimes I could just look at it and say, "I know how to do all these projects and answer all these questions, but why do it? It's like doing a really easy crossword puzzle." I know I can fill in all the answers but do I have to bother? So I was not that motivated. He came to me in my senior year and he said, "You know, we have a class called advanced studies and it can be anything we like." And he set that up, and the class consisted of the two of us. I was the only student. The thing was really focused on trying to get me to think about doing research. He would pose challenging problems for me to think about, maybe even unsolvable ones, I don't know. I can't tell you any specific thing we did, but I remember that whole experience was one of, for the first time, really forcing me to dig deeper and try and figure out how to solve problems for which there were no textbooks.

Cerf: Yeah, there were no answers in the back of the book.

Kahn: And no guidance as how to even deal with that kind of a problem. I think it stood me in very good stead. I think he probably knew what was going on at the time better than I did, but I think in retrospect I could really see that that made a big difference to me in just sort of finding the inner core of who I was at a scientific basis. So that's at City. At Princeton, I had really good collegial relationships with some of the students, two of the faculty in particular. One was a fellow named Ken Stieglitz, who had just joined the faculty. In fact, I think Ken is still a professor at Princeton to this day. Very smart, very energized. Really clever fellow, but a nice fellow too. I really enjoyed those interactions. And, another professor who I believe is also still on the faculty, named B.D. Liu, who had also joined. But those were more collegial even though they were faculty, and I was sort of a senior graduate student at the time. My advisor was a gentleman named John Thomas. John was more distant from me, and I don't think I really appreciated what John's contribution was at the time. I do now. Nor did I think he understood me very well because I think I was probably the biggest surprise they ever had at Princeton. They didn't know what I might amount to, because I was a pretty independent spirit. I'd go off and do my own thing and I was not part of the normal milieu of people who paid due deference to-- I was pretty much of an independent spirit, is the bottom line. John's contributions were very structural. I believed at the time, and even more so now, that he could have advised anybody on any topic in any field. And that's not to say he didn't know a lot about the field because he did. But he didn't inveigle himself into the middle of it. He would say, "Okay, tell me about your progress in the last" whatever period of time it was that he described. And he said, "And what problems did you encounter? What's your approach to solving them? How are you going to deal with this thing that you're not sure how to deal with? What things are you going to document?" He would, by virtue of asking generic questions, cause you to form the plan for how you're going to deal with your stuff. Never specifics. He would never say, "Go find a solution to this and get the answer to that and calculate this Eigenvalue." It would all be [to] have you tell him what your plan was for dealing with it. It was really pretty effective. John retired from Princeton sometime in the late 1990s, maybe '96 or something like that. I forget the exact date. I went to his retirement party on the campus. He had a number of students, too, that were fairly well known. You probably know some of them. Eugene Wong was I think his first student.

Cerf: Indeed, yes.

Kahn: You probably know Steve Wolf was another one of his students.

Cerf: Wow.

Kahn: So he had a number of students that became fairly well known in our field. I was, like, student number six. John decided to introduce all of the students who were there one by one, with a little bit of background on them. When he got up to me he said "Bob was always a challenge for us." He said, "Because we've always"-- I forget exactly how he phrased it but, a very nice introduction. But then he went on to say, "And he probably doesn't know this, but he's the only graduate of our department that ever wrote two Ph.D. theses."

Cerf: And you didn't know that either, yeah.

Kahn: Well, I didn't know that either. What he meant by that was: my Ph.D. thesis was actually an A+B thesis. It had sort of an A part on modulation and coding kinds of things, and it had a second part on sampling theory. The reason it did was they had come to the conclusion, according to him, that I had done enough work for the Ph.D. thesis in a few months and that I just hadn't spent enough time in the program. So instead of having me continue down the path that I started out, because I had pretty much dealt with that problem, they reassigned me to work with B.D. Liu, who had just joined the faculty. That relationship actually worked out pretty well and together it was more collegial, as I say. We ended up pursuing a whole other area in the sampling theory. So that's the reason why when it got packaged up together it was sort of viewed as an A+B package. Very interesting but very different kinds of things. For example, on the modulation problem I was able to show that… The most general form of modulation involves modulating the amplitude and the phase -- amplitude and angle -- of signals. I was able to show that if you specified one, the solution to the other was actually the solution to the Schrödinger wave equation. If you want to compact the signal into the minimum bandwidth, the question I kept asking myself is why would that come to play here? What is there about quantum mechanical things that would play in bandwidth compression and signal processing? But that was the end result of that analysis. We did a similar thing on the sampling side of the world to try and deal with what happens if you take multiple samples. Could you take a signal that was wiggling really very fast and sample it a lot slower and still be able to reproduce it, because you don't have the technology that will sample it fast enough? I was able to show that if you actually sampled it a lot slower but you independently processed the samples, like you processed the signal and its derivatives and different combinations of the signals that you could derive, you took the signal and put it through a lot of different filters, let's say, and then sampled slowly the output of all those different filters, you could then for an appropriate choice of filters, namely they couldn't be all dependent linearly on each other -- like they couldn't be identical, that wouldn't help -- that you could actually exactly reconstruct the signal if it were appropriately band limited. It was an interesting result, which I think probably has many applications. Like if you're trying to build a high speed oscilloscope today, and you can't make the electronics work fast enough. You could actually make the electronics work 1,000 times slower and still be able to figure out what these signals looked like. It was another interesting contribution that I didn't really do more than just show the theory of. The experience at Princeton was mainly a theoretical one. It was mainly into physics and math and things like that. So I just tucked away in the back of my mind that I had written two Ph.D. theses. Okay, didn't know it, didn't even know he was thinking about that. In 1998, Harold Shapiro, who was then the president, awarded me an honorary degree at Princeton. And when I first got the notification I said they made a mistake. They forgot I already have a Ph.D. from Princeton.

Cerf: Yeah, I already earned one.

Kahn: They forgot, because I didn't think they actually gave honorary degrees to their own graduates. They certainly rarely gave them in engineering in any event. It might have been one of the only ones they ever gave. So, when he had his dinner the night before, he asked to say certain things about your experience at Princeton. So I said, well when I first heard about the honorary degree, I thought there must be a mistake. But then I remembered John Thomas saying that I had written two Ph.D. theses so I thought it was only fitting that I got an honorary degree too.

Cerf: It wasn't honorary in fact. It was in recognition of the second one that you did.

Kahn: No, I think it was honorary. I think it was more in recognition of the Internet as being a social contribution to society, not for the technical parts of it but more for the impact that it had. The other thing that was really interesting at that particular dinner was I recounted my experience because on the golf course… They had this wonderful golf course at Princeton, Springdale. It was right adjacent to the graduate school so I could literally walk out. I could have played in academic robes if I wanted. The summer after finishing all of the required coursework, classwork… I took very few classes at Princeton because I was mainly in that library that "A Beautiful Mind" recorded John Nash is doing his stuff in. I was in there all the time reading books and doing a lot of studying. I took some classes, but a lot of work on my own. I took something like four or five months off and did nothing but play golf, just to transition between the infusion of knowledge and the writing of the thesis. So from roughly May of 1963, I guess, through maybe October of '63, I did nothing but play golf. Four rounds a day.

Cerf: Good grief.

Kahn: Which is a lot, because a typical round today on many courses is like four or five hours. I played with this other fellow student by the name of Stu Smith and we could walk the course in parallel; we played ready golf in parallel, and we basically played the course about two and a half hours. We'd be on the course at 7:00 in the morning. Nine-thirty we'd be finished with the first round, back on the first tee, 12:00 we're finished with round two and so we'd have a sandwich and Coke. Stu was as gung-ho as I was and we'd be back on the first tee at 12:30. Three o'clock we're done with round three. Now we're tired because we're walking, carrying around a bag. We lived that summer in a house on Battle Road we had rented for the summer with some other students. It was right near the golf course, so went back there, went to sleep, took a shower, went to sleep for an hour or two. Five o'clock we're back on the first tee again, 7:30 we're done because it was summer hours. We'd go out and do barbeque in the back or something, grill a steak, go to sleep and repeat it. And we did that. That's when my game really got really good.

Cerf: Wow!

Kahn: That's where I really started to figure out how to get it down. So, I was recounting at the dinner this particular story and I said, "I know it's hard to believe but that's what we did and I've not seen or heard of Stu since then. He's the only one that could vouch for that." At which point this hand goes up in the back of the audience and he says, "Hi, Bob" and it turns out it's Stu Smith, who at that point was chairman of the physics department at Princeton.

Cerf: Oh, my gosh.

Kahn: So he was able to vouch for what had actually transpired.

Cerf: That's absolutely terrific. What a great story. How many people have gotten to play golf for four or five months straight? So you sort of retired early and then went back to work I would have to say. We touched a little bit on your first job out of school, and that was at Bell Labs. After you finished at Princeton, if I'm remembering right your thesis was on rate distortion theory, is that the correct term?

Kahn: Well, it was called "Some problems in the sampling and modulation of signals". It was about sampling signals and reconstructing them from samples. This was a well-known theory. Claude Shannon and Nyquist had worked on similar things. They showed if you sample a signal at twice its highest frequency…

Cerf: You could always reproduce it.

Kahn: …you can reconstruct it exactly if it was band limited. I generalized it to a general theory of how you could sample signals with multiple samples at different rates and the like. Then the other part was on the modulation one I described earlier. So, I began to think about rate distortion theory when I was up at MIT and I actually wrote one paper on the subject. But I really didn't stay with that very long, as you well know, because I got into networking fairly soon after.

Cerf: Actually let's follow that. Now after you finished at Princeton did you then go straight to MIT to teach?

Kahn: Yeah, I did. I originally thought I might like to go to a small school. Again, my model was I don't function very well as a small fish in a very big pond. I'm too individualistic, I guess, or something like that. So I thought maybe a real small school would be the first choice, and I sent letters around, interviewed with some of them. It turned out the more I looked at it the more I realized this was a bad choice for me. There wasn't enough support structure. I'd be literally all on my own, and it would be hard to really do anything. So I changed gears and I think I sent two letters out very late in the game. This was probably in May or June of '64, when I was graduating. I sent one to Stanford, this little school on the West Coast, and one to MIT, this little school on the East Coast. I got a polite letter back from Stanford basically saying it's sort of late in the game. Everything has really been pinned down for the next year but my application looked very interesting and they'd be happy to consider me for the following year.

Cerf: To which department did you submit it at Stanford?

Kahn: That was probably electrical engineering.

Cerf: Okay.

Kahn: Remember this was '64.

Cerf: That was probably… John Linvill would have been the chairman at that time.

Kahn: It could have been John, yeah. That's probably, in fact, exactly who it was. And MIT sent me back a very interesting letter. This was from Peter Elias who was then chair of the department, and he said virtually the same thing. It's late in the year. Basically all their slots are filled. But, on the other hand, they would be more than happy to invite me to come up and meet with some of the faculty and just get to know them, maybe give a seminar. Which I did. I came back, and one week later or two weeks later there was a letter in the mail saying they had found some additional funding, and were willing to offer me a post doctoral fellowship and assistant professor up at MIT. So I decided to go do that. I don't know what I would have done if they had said no.

Cerf: So this was electrical engineering also.

Kahn: Electrical engineering. It was an interesting group because that incoming class included people like Mike Dertouzos, who you probably know.

Cerf: Yes.

Kahn: I think Bob Cooper might have been in that class.

Cerf: Oh, my gosh, amazing.

Kahn: A few other people that we know well.

Cerf: Mike ended up running the Lab for Computer Science.

Kahn: And recently passed away.

Cerf: Yes, and Cooper ran DARPA for a time.

Kahn: Became head of DARPA for a while.

Cerf: That's fascinating. Now, you were one of the youngest professors in the department if I remember right.

Kahn: Probably the youngest. I can't vouch for that but I think it's probably the case. It was rather interesting because I had basically been trained, even though it was electrical engineering in my background and even though I got a pretty good training in the pragmatics of it at City College, my work at Princeton was largely theoretical. I took a lot of courses in applied math and physics. Most of the electrical engineering stuff I pretty much learned myself. I took some courses in it, but I was really an applied mathematician up there and the courses that I was teaching… I taught the course in information theory for a while. I taught some of the courses in circuits and communications theory, things like that. But they were basically pretty mathematical, which is why I found electrical engineering so interesting in the first place. When I first went into college I thought maybe industrial engineering was what I would go into and I very quickly decided that was not the right choice for me for a variety of reasons and, in fact, there was no program for that even. So I switched to chemical engineering and I was in that for all of about two or three weeks when I decided I didn't like the lab work. So I switched to electrical and that kind of stuck. That's principally what I was doing at MIT. The faculty at MIT was a pretty amazing faculty. They were all people who not only had established reputations in the field but they seemed to have insights into things that I didn't have. It was really brought home to me one day when… I had periodically been meeting with the gentleman who ran the group. His name was Jack Wozencraft. Jack is now retired, living out in California or Oregon, someplace on the West Coast right now. I really admired Jack Wozencraft. I really looked up to him. He just seemed to be smart about so many things. We had a very good group at that time. Irwin Jacobs had one of the offices on the faculty. Claude Shannon was on the faculty. Bob Gallagher, Bob Fano, Harry Van Trees. There were a lot of very good people that were there and on that faculty at the time. I was busily working my mathematical problems. I'd go into Jack with them and I'd be working them on the board trying to get his insights or whatever. I would say every now and then he'd get animated, he'd get up on the board and start working with me. But more often than not he would be writing his papers, reading his book. He'd listen, very politely tolerate me being there. Finally one day I just said, "Do you mind if I ask you a personal question?" I said, "Why is it that sometimes you get up on the board and you're really animated and the other times I kind of feel like I'm interrupting?" He says, "No, you're not interrupting. I'm listening to you." But, he said, "The reason that I get animated is when I can see what the value is of what you're doing. If I know what to do with the results of your problem I get real interested in it." He says, "If I can't figure out -- so you get the Eigenvalue to this equation -- I don't know what to do with it. If I told you I knew the answer, what would you do with it?" I said, "Well, I'm just trying to find out what it is." He says, "Well, that's the problem. I want to go a step further. I want to know why I'm trying to get the answer to the problem. If I can't -- if I don't know why I'm trying to solve the problem, other than it's an interesting problem -- I just don't get that interested." All right, well that was a real insight for me because that had never been part of my formulation of problems.

Cerf: Problems were interesting, and if they were interesting you attacked them.

Kahn: It's like if somebody gave you a crossword puzzle to solve and you started to work on it and I said to you, "Why are you working on the crossword puzzle?" Probably the answer would be “Because I feel like it”, or ”Just to occupy myself for awhile”. But there was no larger answer for most of the problems that I was working on. And he said, "It's basically experience. If you really want to understand how to deal with problems, you need to get some experience." He said, "I'll give you a list of people. If I were you, I'd go take a year, two, take a leave of absence, apprentice yourself. Figure out how to build something and then come on back." So, that's what I ended up doing. One of the places he suggested I go visit was Lincoln Lab. My friend Paul Rosen was then running the communications division and Irwin Lebow was his deputy, and I knew them both very well. In fact, Paul had been in my class as a student when I taught information theory, which was kind of interesting because Paul was a real good practitioner of the field but he wasn't really a theoretician. Every now and then when somebody would ask that question about why would you want to do that, Paul would jump up and he'd have all the answers. It was really very interesting. But Lincoln had, right after I had applied, been put on a freeze by the DOD. They literally shut off all hiring. So I got this -- I forget what it's called -- a letter from Irwin who basically said, "Look, we'd love to have you come, but right now we just can't hire anybody. We don't have any discretion on that." Well, the other place that Jack had suggested was BB&N and…

Cerf: This is Bolt, Beranek and Newman.

Kahn: Bolt, Beranek and Newman was a small architectural acoustics firm at the time in Cambridge, up in the fresh pond section, just north of Harvard. He suggested I talk to Jordan Baruch, who I'm still very friendly with. Jordan later became Assistant Secretary of Commerce for productivity technology innovation and ran a big chunk of the Commerce Department. We had very good discussion back then. It turns out at the time that I was looking, I was looking at two things. The other possibility was the University of California at San Diego, which was just starting a EE department out there. It turned out Irwin Jacobs was exactly thinking of leaving at the same time and taking a leave of absence.

Cerf: Wow.

Kahn: We were looking at the exact two same jobs. If I remember, I was out in Hawaii that year. It must have been the summer of 1966. I was visiting Wes Peterson, who was one of the faculty members out there. He asked me, he said, "Did you hear what Irwin decided to do?" Because Irwin was about two or three months ahead of me in the decision process. I said, "No, I didn't. What did he…?" Because I knew whatever he decided -- it was the Pauli Exclusion Principal -- I was going to do the other thing. I was going to have to do the other. He said, "Irwin decided to go to BB&N." So I said, "Well, I guess I'm going to move to California." On the way back I stopped at Stanford and I was visiting Tom Kailath, who was then in the EE department. I think Tom may still be there if he hasn't retired yet. Tom said, "Did you hear what Irwin did?", because Tom always knew everything that was happening before anybody else knew it. He said, "Did you hear what Irwin decided to do?" I said, "Yeah, I understand he's going to BB&N." He said, "No, he's going to San Diego." So between the one day and the next, somehow Irwin had apparently changed his mind. I later was able to verify with him that after they decided to go to BB&N it was somehow a torrential downpour that night. I think Joan said, "Why don't we just go find someplace where the weather is a little better", and they just changed their mind overnight. So I ended up going to BB&N as a result.

Cerf: Did Irwin actually wind up joining the University of California at San Diego or did he go and start Linkabit right away.

Kahn: He became a faculty member in the department.

Cerf: Okay.

Kahn: As far as I know. I went out to visit him. I think it was the winter, it was like in December or January of either '68 or '69. I'm not sure which. I think it was '68 and I was en route to some vacation out in that area and I stopped by San Diego to visit him. But en route, I stopped in Los Angeles and I had dinner with Kung Yao who was a former student of John Thomas and was then on the faculty at UCLA. And Kung said, "Did you hear that Irwin, Andy Viterbi and Len Kleinrock just started a little consulting firm called Linkabit?" It turns out the company was an answering machine in Len's closet at the time. There wasn't much going on. But when I went down to visit Irwin he sort of said, "Yeah that's true, but how did you know, because we just decided it a couple of days ago?" So I gave him the story on that.

Cerf: That’s fascinating. Now we’re going to come back to BBN.

Cerf: Bob, you had mentioned, I think, an important life lesson when you had mentioned Wozencraft saying that he was asking this question, "Well, why are you doing that? What's the motivation?", and so on. Are there other kinds of life lessons that come to mind over this incredible time period of your career that you think are notable and worth sharing?

Kahn: Well, on a personal level, I think clearly the experiences I had with my mother's health as a young child told me how fragile things were. I sort of knew that right from the get go. I don't think I had much appreciation of financial stuff when I was growing up, I think for two reasons. Number one, we didn't have any really serious financial resources back then, and therefore there was nothing much to worry about because my dad had a fairly steady job and it was just sort of that's what life was like. But also, he was an accountant by background. There's something about me that always has taken the position that if somebody else is doing something, I'm going to do something different. I want to be always doing the new thing rather than adding on to some old thing. That's part of the "didn't want to be a little fish in a big pond" because that's somebody else's big pond. I think the fact that he was into that field sort of shied me away from it for a long time. I later learned actually how important and interesting it was. That may sound strange, and I think the same is true about law. My wife, Patrice, is a lawyer and I think she has shown me some of the more interesting aspects of law as intellectual domains. Particularly the copyright field, which she specializes in and how it's like this intricate tapestry. It's almost like in this book "Gödel, Escher and Bach" showed how music could have certain mathematical patterns and how art could have an Escher style. You can see the same things in both accounting and law and the like if you look at it that right way. So some of those were, I wouldn't say they were lessons as much as sort of awarenesses. One of the most difficult things for me to deal with was, as a young kid, I had always believed that if you had a really good idea from a research perspective that you ought to be able to get it out to the community. I remember when we were doing the original ARPANET design, it had been mandated that we use Bell 303 modems. These were 50-kilobit modems that you could get from the Bell System. They hadn't really been propagated very much. It was brand new technology for the most part. I know we're going to get into this topic later. But when I went to Bell Labs to find out how they worked - I needed to know all the details about this network that we were building - I discovered that they had a working Bell 304 modem at Bell Laboratories. And the Bell 304 modem [was] unlike the 303, which was basically a resistive-capacitive thing -- you pulse it and the circuit would die out and you do it again and you could tell which way the pulses were going. It's a trivial modem. The 304 modem was more complex. It used the same channel bandwidth, which was, I think 48 kilobits, one of the nominal structured bands that Bell used to break it down into voice channels. And they built this modem with different signaling convention that would get 108 kilobits a second on the same channel. I said, "Since we're paying for these channels, we should be using the Bell 304 modems because it would give us more than double the capacity". I couldn't convince Bell Labs or the Bell System to let us use it because it wasn't a product on the market. I said, "Well, but this is just a research experiment." They didn't want to inject it out because their business decision processes said that putting out something that would allow you to get more bandwidth on something, given the regulated environment in which they lived, was not as good a business strategy as requiring more inefficient modems to be used. So they didn't actually release it. It was, I don't think, ever released. I struggled with that. I actually wrote a memo saying, "Well, what's the purpose of research if it doesn't get out and you can't use at least for experimental…" It was very troubling to me. I ran into almost the same problem many years later at DARPA when we were doing the packet satellite program. At the time, there were no domestic satellites. This was, maybe, 1973 or something like that. But the only satellites we could actually access at the time were the military satellites, in particular, the DSCS [Defense Satellite Communications System] satellites or the Intelsat satellites. Those were the two options available to me at the time. I remember dealing with the Comsat folks to try and get them to work with us on a protocol that would essentially make a single satellite channel work like an Ethernet in the sky, so that you could send a packet up and it would broadcast it and the relevant ground station would pick it up based on identifiers in the signal and relay it to the appropriate parties. They wouldn't go along with it, because this amounted to a more efficient use of a resource that they had. It turns out their money was based-- I mean I'm now inferring this, but their profitability was based on a fixed percentage of assets. So as far as they were concerned, putting up another satellite would give them a bigger asset base. They could bring in more money. So they didn't want to do it. They never said that in any letter. I never heard it in so many words, but I'm inferring that that was the motivation. But I was now armed with this because I had already learned this lesson once before. So I went to my friend Fred Bond who then ran the Defense Satellite Office, which was part of D.C.A., I think, at the time. He was in charge of the DSCS satellite. He was in charge of all the satellite programs. Fred said "Fine. I'll be willing to give you some capacity on the DSCS system for experimental purposes." So I wrote back to George Lawler, who was then the marketing manager of Comsat who was handling this and I said, "We found another approach. We don't need the Comsat satellites." Now suddenly George was faced with the choice of not selling one channel vs. N, because he was thinking if we had like five satellite sites, that's ten different bidirectional circuits.

Cerf: Yes, exactly: $N^2/2$.

Kahn: $N^2/2$ bi-directional circuits versus one. So he'd rather sell 20 than one. But now, he had the choice of one versus zero. What was worse was, and I think this was the beginning of my learning how to invoke some political skills here. I simply pointed out to him that I had discovered that the authorizing legislation for Intelsat basically allowed the DoD to go directly to Intelsat if they wanted. They didn't actually have to go through Comsat in order to get channels, if it was a matter of national security. I said that if push came to shove we would invoke that instead of having to work through Comsat. So I thanked him for his help, appreciated all of it, understood where they were coming from and we were going to do a different way. One day later, there's a call from George. They've reevaluated the whole situation. They're now willing to work with us. That was a lesson that I learned and then I saw it directly applied. I've seen several. I guess the other major one that I would point to is that when you are running a program that for the first time gets a lot of money in it, discretionary money, that the sharks will come out of the water from every direction and try and take a bite out of it. We saw that in strategic computing in spades. Now, I sort of knew that, but I actually saw it directly happen. I think I also learned one other lesson when I started CNRI, and I know I'm getting ahead of the story here. But when I started CNRI, I really thought that there was a need for this country to invest in infrastructure -- information infrastructure in particular. I thought the idea was so compelling that industry would just jump in. We should have no trouble getting 100 of the big industries to put in $1 million each. That ought to be enough to get a program going DARPA-style. Didn't happen. It didn't happen for many reasons and I think it was the understanding of all of the reasons why it didn't happen. It wasn't that the idea wasn't good. Hardly anybody said this was a bad idea, just that the government should do it rather than us. Or, they don't have that much money for discretionary funding outside. Or, they were willing to fund it inside, but not out. I think that really gave me a much better appreciation of what the pragmatic possibilities were for actually raising funds for, let's say… This wasn't totally altruistic because I knew this would help the business base, just like the Internet wasn't altruistic either, it was a major contribution. But this was something that didn't fit anybody's immediate direct business need and therefore was very difficult to make the case that they ought to put in money external to the company when they could spend it internally on things that had more direct meaning and value to them. So that was another lesson that I learned.

Cerf: This sounds like the "What's in it for me?" problem, where in order to convince somebody to do something you have to have an argument that is persuasive from their point of view, as opposed to the one that persuaded you, which might be for different reasons.

Kahn: It's that famous story they tell about the salesman who walks out of somebody's house and somebody says, "Well, did you make the sale?" He says, "No, I didn't make the sale because they didn't understand." Somebody says, "The problem is you didn't understand how to make the sale."

Cerf: Exactly. We're still in the sort of the pro forma section of this interview, Bob. So the next question is: What do you think your proudest moment was or maybe your proudest moments, plural?

Kahn: That's so difficult a question. I think it has many different tentacles to it. The question is what are you feeling the pride about? Are you feeling pride about something that you did, or feeling proud about what somebody else did?

Cerf: It's a loaded question in many respects. Well, it could be, maybe -- I don't mean to distort the question necessarily -- but "things that were most satisfying" would be another way of asking this.

Kahn: Let me just first say something about the external part. I don't think that's really where we want to linger, but I remember when my mother went in for heart surgery the third time. She didn't come out it and she sort of had a hunch she might not because I don't think she… But she went through it because her family really wanted her to give a shot. Maybe there was a chance she could have gotten through it. But the alternative was she probably would have only lived for another month or two anyway. When I saw her decide to go through the surgery and make that decision, even though she was really unsure whether this was the right thing, there was a feeling of pride that swelled through me, but it was an external one.

Cerf: An impressive act.

Kahn: Yes, a very personal one too. If you're talking about pride in things that I've done myself or that have happened to me directly, I would say there's probably some-- It has to be associated with a set of recognition that I've gotten. You've gotten some of it as well. I don't know whether you feel that way. Things like the Presidential Awards that I've gotten - the National Medal of Technology and the Presidential Medal of Freedom. Those were very proud moments for me. The Turing Award, that was a proud moment. Some of the honorary degrees, especially the one from Princeton, which you're an Alma Mater. I don't tend to characterize things in my life by proudest moments and non-proudest moments. But I certainly think those would have had to be some of the ones.

Cerf: Any great golf games that you remember?

Kahn: I remember the hole-in-one I got. The one and only one I got, at Wintergreen once. The pro there took the golf ball, said they would mount it in a trophy and they never gave it back. But that was a really proud moment.

Cerf: You and Don Heath have played golf from time-to-time together. It seems to me I remember at least one game where you did very well and I don't know what they call them - the doubles thing - or that's tennis.

Kahn: Yes, Don and I played a number of times. In fact, I think we won a number of tournaments together. I think the one you're talking about… We won one at Avenel. I remember we won one at Hidden Creek where he's a member, or was a member. This particular one was a member guest tournament and I remember very clearly we were playing what was called "Ham and Egg" golf.

Cerf: I have no idea what that is.

Kahn: "Ham and Egging" describes an event where when one guy does really badly, the other one does really well. They basically count the way you do in the function of who did the best on the whole.

Cerf: Oh, that's a nice arrangement.

Kahn: So if was a hole where Don got an eight and I got a birdie, they count the three.

Cerf: Okay.

Kahn: Or, if I got an eight and Don a birdie, they'd count the three. Not that Don got many eights, but it was just one of those days where things were working. We didn't shoot that much better than our normal score. I think he shot an 81 and I shot an 88 that day. But when you looked at the combination, the "ham and egg" effect, we ended up with a 61.

Cerf: Wow.

Kahn: The funny part of it was, we probably missed seven or eight shots that were within two or three feet, little puts that we had missed along the way. We thought it was going to take something like a 54 to win the tournament. So when we finished at 61, we thought, "We're not even close." We were kicking ourselves, "How could you miss that little putt? I mean, it's right there." The bottom line was that the nearest competitor was like 68 or 69. We actually--

Cerf: That's great.

Kahn: But I wouldn't call that a proud moment. I think it was just an effective moment. We were good.

Cerf: This next question has to do with turning points, and it seems to me that you've already hinted at some or have been pretty clear about some. When Wozencraft told you that you should go off and get some experience, it seems to me that was a really important turning point for you. Can you think of others that you would lump into that category of important changes in direction that, looking back on it, had a big impact on the way your career has unfolded?

Kahn: Well, I think in recent years, the one thing I would say that probably made more difference than anything was getting my wife Patrice in the act with me, because she brought a perspective on things that I don't think I would have ever had myself. It was one of the really important things to understand how the perspective that she had, particularly the copyright and like, would make a big difference. I think it's really profoundly changed my view on things.

Cerf: We will come back to some of the more recent work that you've done. But I wonder whether your interest in digital objects and the like and the intellectual property that could adhere to them is partly driven exactly by this exposure to the legal side of things -- the concepts that you don't normally think of when you're building computer networks, for example.

Kahn: It certainly had a big impact. You were talking before about proud moments. I think in terms of the most important moments in my life, it was meeting Patrice and getting involved with her. It's clear that was just one of the most important things that ever happened and it probably did change the direction in which I went. When I started working on digital libraries, I thought the problem was technical. I began to see, mainly with the inputs from her, that the problem wasn't so much technical as it was social, legal, and political. People who were making the decisions and controlling things were not the technologists so much as the legal folks. Take any company that's basically in the IP business, whether it's a movie company or a book company, lawyers are pretty critical. You can come up with the best technical approach in the world, but it's got to make sense to the people who make those decisions.

Cerf: Just for the record, "IP" in this case refers to intellectual property and not Internet protocol.

Kahn: Right. So addressing the issues of the technology from the point of view of what the barriers are to getting it out was one of the fundamental insights that I had, and I would have to credit Patrice with that. It just stuck. So when you look at the digital object architecture, what you realize is it has the potential for allowing a lot of the dialogue between parties that might not normally happen, or would be hard to happen, to be possible to embed within the system. Not that technology is going to figure out what to embed, but that you have the mechanism to embed it suddenly makes it possible to envision technology negotiating things that are there on behalf of the decision makers in those companies. So it was a really fundamental insight that has really stuck to this day. While I love the technology - I love to work on the technological problems - just like you see in the network business today, not everything about networking is about how to get the bits to move faster.

Cerf: Absolutely true.

Kahn: It has to do with decisions that companies want to make about their investments, and what can happen, and so forth. The digital library field was particularly difficult because there's such a long history in librarianship. In my career, every time I've tried to talk to people who have that as their business - it doesn't happen so much now, but in the early days -- people felt the need to take an hour or two or three to explain that librarians really have been there before and that technologies can't march in and just suddenly understand things.

Cerf: This actually leads to a related kind of question. Looking back now, although we have not covered all aspects of your career and we are going to do that, as much as we can anyway. If you look back on all the various things that you've done, what sort of impacts do you see as a consequence of the work that you've done? How would you characterize the results of the research, or the results of the projects, or the results of the ideas, if you were to look back now over this 30 to 40 year period? What kind of impact do you see has happened?

Kahn: Well, clearly the thing that's gotten the biggest play in the outside world has been the Internet, which I guess was something that nobody could have predicted would have evolved exactly as it did. You know as well as anybody that this is not a thing that a single person did, or that occurred without the help of a lot of people over a lot of time. In fact, it really wasn't until the infusion of a lot of the corporate funds into pushing it out that it really began a global and worldwide thing. Plus the backing of governments around the world, because it wouldn't have happened if both of those hadn't happened. It might have been an interesting research project. I still think we're a long away from learning what the final impact of the Internet is. Networks are changing. New networks are coming up. People are constantly trying to reinvent networks. They're constantly trying to say, "Well, the old ones don't work. We need new ones." And yet, there are some things that are pretty constant in all of this as things go forward - the need for connectivity, the need for bandwidth, the need for reliability. I was struck, in a movie that I had put together for the unveiling of the ARPANET, to listen to all of the computer folks talk about what that was about. They all talked about it in terms of the technology of the time - the need for the powerful graphics simulator to connect to a remote time-sharing system. It was all in terms of what was there at the moment. Whereas when you listened to the communication folks, they would talk about the need for reliable communication. They would talk in things that didn't… They wouldn't say, "We need faster light emitting diodes to move it twice the bandwidth." I think we're now getting to the point where people are actually talking more in terms of the generics. So: the Internet is about connectivity. The Internet is about access to information. The Internet is about reliability. Those will probably be true ten years from now, 100 years from now, to the extent that that kind of capability exists. Now if we ever go beyond computers and beyond communications as we now know it to something else, maybe it will be different. But I can't imagine how ESP and gravity waves are going to take over very quickly, or what's going to supplant computers in the short term.

Cerf: If you were to think back for a bit from both your early years and maybe more recent ones, have you had people that you thought of as role models? You've mentioned several people that had a big influence on your life. But do you categorize any of them that way, or characterize any of them that way, or is that not part of the driving gestalt for you?

Kahn: Well, there are people who clearly had a big influence on my life. But if what you're asking is have I patterned my behavior after anyone else, I think the answer is probably no, at least not consciously. The individuals who probably had the most impact on how I am are probably my parents. I don't know to what extent that's just genetic and it's therefore in the genes and I behave that way, or because I saw the way they were and I try to act that way. It's not a very close fit, but it's probably as close a fit as there is.

Cerf: I would have guessed not. I mean I don't disagree with the point that you're making about your parents at all, but I'm thinking that your natural spirit of independence, I think, probably makes it unlikely that you would pick anybody to want to pattern your life after because you don't want it to be patterned after anybody in particular. You want it to be Bob Kahn's pattern.

Kahn: Well, if I were a professional golfer, I might give you a different answer, but I'm not.

Cerf: Do you have any advice that you would want to leave with either the current or the future generations about anything? Whether it's how you pattern your life or how you approach problems or what things they should be interested in or motivated by?

Kahn: Well, let's see.

Cerf: "Neither a borrower nor a lender be."

Kahn: I would say the first message, which is more of a professional one, is that if you really believe in an idea strongly and you think you have the ability to act on it, even if you don't know exactly how to make it happen, trust in yourself. Listen carefully to what other people are saying,, but if after listening to all of the advice, even if it's negative, you still believe that it's a good idea, go act on it. Trust yourself. Have faith in yourself. This was a lesson that I learned because when I got into the networking field in the first place, or I chose to get into the networking field, there were many people who thought this was a dead end. There was no "there" there. What did we have in the early days? Large batch processing machines. Maybe a handful of companies might have operating systems and time-sharing machines. But when AT&T chose not to get into that business early on, it made probably very good business sense for them because there just wasn't a business out there. There may be a research opportunity, but they have to decide what they want to invest in in R&D, just like is true with DARPA today. The research community might think that the following problems should be dealt with, but they have their own agendas. People need to decide for themselves what's important. So I think that would be message number one: believe in yourself. Trust your own instincts going forward. I guess the second lesson would be that if you get married, marry a good woman or the opposite because that can make a big difference in your ability to get things done. I've been very fortunate. I'm fortunate I have the added benefit of a collaborator in parts. If you're not into a good relationship, it's very hard to find the space and the time to do what you're going to do without being encumbered. So a good relationship is really important to being successful, in my view. Trust your instincts on that too. Let's see. Another thing that I might say is: leave a little room for serendipity in your life. That is, don't plan everything out because you just don't know what's going to happen that can profoundly change you. It could be something like an idea that just shows up out of the blue that completely changes things, or an off-hand comment that somebody you trust makes to you, or it could be a random thing that happened - a letter from somebody that makes a suggestion, or something that just happened by virtue of nature that changes things. Everything might not be good and some of that good stuff might be something that you just have to learn to deal with. But don't assume that the way things are now is the way it's going to be. You just have to roll with the punches as they come.

Cerf: There are two aphorisms that come to mind. One of them is "Man plans and God laughs." The other one is "Life is what happens to you when you were planning something else." Probably a lot of our careers would follow that latter pattern. Well let's try switching gears a little bit, Bob. Institutionally speaking, we haven't talked very much about your time at Bolt Beranek and Newman, and then at DARPA, and then at CNRI. So from the institutional point of view, I'd like to pursue that a bit. Of course within that context, you had many projects that you introduced or initiated. As we explore those different time periods, I'd like to expose some of the ideas and problems that you were grappling with. You took Jack Wozencraft's advice and you ended up at Bolt Beranek and Newman. What were you doing there? What was the plan?

Kahn: I don't think there was a plan at all. BBN was an interesting place in that time frame. It's sort of gone through a lot of transformations along the way. But back then, there were a lot of just smart people from both MIT and Harvard and elsewhere who were congregated there. They were pioneering a lot of the advances in computing technology and communications during that period. I think they were just happy to bring smart people into the mix and have them interact with other people. I met many people there who I got to know and have maintained long friendships with ever since. The environment there was unlike any other one I had ever been at. The gentleman who ran the part of the company that I was involved in was a fellow named Jerry Elkind, who was sort of a control theorist, experimental psychologist, computer scientist. He was sort of across the board. Jerry later ended up trying to do some things at BB&N that were pretty unusual, like he wanted to take the research guys and put them in running an operational part of the business. Jerry ended up being hired by Licklider over to MIT for a while before he ended up going after Xerox PARC where, I guess, Bob Taylor hired him in as his boss. I don't know what happened out there because I never really tracked the internal workings of PARC. I'm sure you could find people from the Xerox days that could explain that. I found Jerry to be a very interesting fellow to work with. He had another gentleman working for him named Frank Heart. So in some sense, Frank and I were parallel in terms of the reporting structure, but Jerry wore multiple hats. Frank had just come in from Lincoln Labs and he was trying to figure out what to do. In that context, they gave me freedom to go think about anything I wanted.

Cerf: Roughly, when are we talking about? Is this 1967, or something like that?

Kahn: I think I joined them in October of 1966, and I stayed there until October of 1972, so I was there just about six years. When I joined them my thought was to go back to MIT, but we got involved in the ARPANET project, which I can explain. I had to make a determination. Do I go back or do I stay with the ARPANET project? I opted to stay, figuring if they didn't want me in five years, then going back now was probably not the right thing, and I wanted to see this project through. I started working on networking when I was at BB&N because I found it an interesting topic. Most people in my field did not think this was an interesting thing to work on, for the reasons I said before. But I found it an interesting problem and I thought I was following Jack Wozencraft's advice. I knew why I was trying to solve these problems because I thought [getting] computers to work with computers was a good enough answer. Now, you might say today, "Well, computers working with computers -- for what purpose?" I hadn't thought it through to that level. I just had this idea that if computers could interact with each other, all kinds of good things could happen. That was a good enough answer to his admonition.

Cerf: Was Bolt Beranek and Newman's interest in networking -- did that precede the ARPANET project somehow? In other words, you were thrashing around in that area before ARPA released its request for quotation?

Kahn: That's right.

Cerf: That's interesting. I didn't realize that.

Kahn: The RFQ come out of DARPA some time in the summer of 1968. I had actually been working on most of the problems in networking -- buffer control, error control, remote control -- all of those issues.

Cerf: Communications related things primarily.

Kahn: Right. But to me, it was more with a leg in the computer side of the camp than pure communications, because I was not worrying about the traditional communications stuff. I wasn't worrying about phase lock loops and bit sync and BCH codes and the like. I was worrying about more the computer things. In fact, many of the things I was thinking about would have been almost antithetical to the communications view of the world, which was more dictated by Claude Shannon's inveiglements to optimize: find out how to use the maximum efficiency and bandwidth. That came out of all the theories of informational theory. But I was looking from the other point of view. How do you just make something work, even if it's less efficient? I was not trying to deal with optimization in the mathematical sense. I was trying to deal with: make it work in the practical sense. But to try and do a good job of making it as efficient as I could do by thinking it through. Or where I couldn't, to enable mechanisms to be invented in the future. A good example of that, which turned up in the ARPANET project, was when we had to specify what the routing algorithm was. We actually specified one. It was the "shortest path" algorithm. But the thing that we proposed to build was a mechanism whereby the entries into routing tables could be done by an algorithm. So you could pull this algorithm out and plug another one in. The internal structure wouldn't change, but you could maybe optimize the algorithm in the future. So I've been generating these memorandums and I remember Jerry Elkind coming to me one day and saying, "Why don't you share them with the folks at DARPA? They might be interested." I said, "Okay." I didn't really know much about DARPA at the time. I had learned about them when I got to BB&N. I thought my support was internal support for BB&N. I suspect they were probably charging to one of the DARPA contracts, but I don't know that for a fact. They sent some of those memos into DARPA. The next thing I know… I think I actually sent it in with a letter. I think I addressed it to Larry Roberts basically saying, "Jerry told me that you were interested in networking. So here's some stuff I have been doing." I probably have that letter in my files.

Cerf: In your archives somewhere.

Kahn: I didn't know Larry. I didn't know what he was thinking about. I certainly didn't know about the ARPANET project. I knew none of that at the time. Back comes a phone call from Bob Taylor who says…

Cerf: By the way, Bob's role at this point is running the Information Processing Techniques Office [of DARPA].

Kahn: Bob is running the Office, and Larry's role relative to him is somewhat confusing to me. I had thought Larry was a program manager in the office. But I've talked to Larry about that and he describes his role as a special assistant to the director, a place [??] in the office. So the exact arrangement there is unclear to me. I know Larry had… There were some issues about him joining and how he got to join. But I believe he and Bob Taylor were working collaboratively on getting an RFQ out for an ARPANET at the time, but I didn't know it.

Cerf: Yes. My recollection, just briefly, is that Taylor wanted Larry Roberts to come down from Lincoln Laboratory to run this program, and Larry didn't want to come. If I remember Taylor's story correctly -- perhaps you heard it too -- he managed to persuade the director of Lincoln Laboratories that he should send Larry down or he would have problems with funding.

Kahn: Yes, it's anecdotal, but I don't know that for a fact. In any event, Bob called me back and said, "We got your letter. We'd like to invite you to come on down and just sort of chat with us." I didn't know who he was at the time, but he explained. So I went down and met with him. I had a very nice discussion. They explained where they were headed and they were going to try and get this network built. If I could help them with it, that would be great. I think it was more of a "Who are you?" kind of meeting, get to know you. I don't remember when that took place, but it probably had to be sometime in the 1967 time frame. Sometimes in the spring/summer of 1968, the RFQ shows up on the street, and it lands on my desk. I'm reading through it and all kinds of stuff that are in my memos are in this RFQ. The error control mechanism I had described in my memos? The actual picture was in the RFQ. That's pretty impressive. It hadn't dawned on me at that time that I would be involved in anything to do with it because I figured… At that time, I was getting ready to go back to MIT. I hadn't quite been out the two years and the bottom line was: good, somebody's going to build this net. To me, the work that I had done was about as practical as I thought I needed to get to go back and be a mathematician. So Frank Heart shows up in my office one day and he says, "I understand you got the DARPA RFQ. Can I take a look at it?" "Who are you?" was my question. He said, "Oh, I'd like to…" Introduces Frank Hart and he runs a group. They like to build this sort of stuff. I said, "Fine." We started to chat and one thing led to another and Frank decided he wanted to put in a bid for the ARPANET. Except I ended up writing a proposal, which help us put it together. So I actually put together the technical part of the proposal. Severo Ornstein, who was one of Frank's people, was a really good hardware designer. A lot of what this proposal was about was describing what we would actually build. The software part of it you could sort of describe in terms of what had to happen; you didn't need to give them the lines of code. But if you were to create hardware, you needed to actually state what you would build. So Severo and I, I remember, closeted ourselves over at his house one day and he said, "Well what does this had to do?" I said, "Well, it has to take the DLE thing, detect it, look for a sync character" and he's designing the hardware right on the spot while I'm describing what it has to do, which was kind of an on-the-spot tutorial. It's like a Ph.D. program in hardware design right on the spot because I'm seeing ideas in my head suddenly turning into hardware right in front of my eyes. I'm thinking, "This guy is really terrific."

Cerf: This is without benefit of a computer-aided design either, I would assume. We were just at the point where you and Frank Heart basically are talking about the RFQ from DARPA about the network. You said you ended up writing the proposal to respond.

Kahn: Yeah, and I had a lot of help from Severo on the hardware. In fact, Severo was commenting on everything else I was writing so he was sort of the surrogate for Frank to make sure that what I was writing made sense because Servero could understand that sort of stuff. It really worked out very well. That's what became the technical proposal that went into DARPA. It was memorialized in the Report 1783 or 63, some number like that.

Cerf: Could I just ask one little technical detail? Did the BB&N proposal make reference to the Honeywell DDP equipment? Servero undoubtedly must have had something to say about that as the hardware guy. This was a combination of things that BB&N would actually implement in hardware themselves plus whatever else they had to do to adapt the DDP machines.

Kahn: What Severo was basically doing was designing the interfaces for the Honeywell machines so they could take stuff off the Bell modems, and take stuff in from the host computers, and doing it so it would be in real time. The actual choice for the Honeywell machine was probably made by some combination of Frank Heart and Hawley Rising, who worked for them. We knew we needed a main computer and our job was to go out and figure one. Frank's big issue there was he wanted to be sure it wouldn't break if you dropped it from a truck, so we ended up getting these ruggedized units that were big as refrigerators and heavier than a tank.

Cerf: They looked like half inch thick steel with screw eyes on the top where you could lower the thing down to the ground from a helicopter. The first one that was delivered to UCLA had all of the accouterments, and then people would ask "What are they going to do with this thing?", and the answer was "We don't know but they drop it it'll survive."

Kahn: The effort to actually create the ARPANET design was actually a pretty intense intellectual activity. In many ways the packet radio program, which we'll get into later, was a much more complicated development because it had so many other attributes to it. But because this [the ARPANET] was the first effort -- like the first missile that you send up or the first rocket into space, you don't know if it's going to work or survive. It may be very primitive by comparison, but because it's the first one intellectually it had more risk to it because we didn't know whether it would really work as planned or not. The specs that were out there for the net really were propagated by DARPA, put out by RFQ, but it was basically asking for something that… It was like when the space program came up they could put out RFQ asking for a rocket that would handle a certain payload and be able to come back to the earth, or whatever the properties were, but it wasn't a detailed design for how to build a rocket by any means. They required a rocket expert to actually do that. The spec that came out from DARPA for this network talked about half a second end-to-end delays in the net, it talked about the need to take in packets and deliver them reliably to the other end. But it was not a spec on how to build it. That's the contribution I really made, was how would you take this basic idea, this specification, and make it into something real, figure out what the strategy was. It was kind of like what Vint and I did together on the whole internet architecture. We had sort of started out with a basic idea of what it was, and then together we kind of worked out the details at that level. Now you still had to turn it into code. The guy who wrote the code was Bill Crowther in the case of the ARPANET. In the case of internet, the guys who wrote the code were Vince Grupa [sp?] at Stanford, and people at UCL in London, and BB&N.

Cerf: We'll come back to that one or otherwise we'll get our history all tangled up here.

Kahn: History is always all tangled up, as you well know.

Cerf: That's theorem number 205. Coming back to the ARPANET then, BB&N submits a proposal and it waits for some period of time and then it gets awarded. How much time does BB&N have, effectively, to actually implement the software and test it and get the hardware all put together and deliver the first thing?

Kahn: The history of the timing as I recall it was: the proposal went in sometime in late 1968, perhaps September or October or some time like that. The announcement of it was in December of '68 as I recall, and the start date was probably early '69, January, right after the New Year, that's my recollection. And it had a nine month delivery date. I don't know where they got the nine months from but it's very anthropomorphic.

Cerf: It's a common figure in the species, but anyway.

Kahn: The goal was to deliver the first of four nodes, which is what that project was about, to UCLA in nine months. In fact it got delivered supposedly a day or two early the story goes, and I think that's right. I went out and actually helped with the testing of that, both that first node and the first four nodes. In fact, that's how I first really got to meet Vint Cerf, who was then the point person on that for UCLA, working in Kleinrock's lab I believe. Vint was very helpful with a lot of the testing of that from the Sigma 7. Now when the first IMP showed up the Sigma 7 was not connected, as I recall.

Cerf: That's correct.

Kahn: That didn't happen until later. So the initial tests that we did were just tests on that machine itself.

Cerf: The IMP itself.

Kahn: And various testing things, because it had a teletype connected to and the like. I actually have a book somewhere about all the details of the testing that we did. Later on I went back with Dave Walden, who was one of the key software guys in that group, and we actually did the first four node testing of the network. In fact, I think we may have actually done some intermediate testing and I can't recall it right now. I have to go back and check the records, because it may not be that we waited until all four nodes were back before applying that- the intermediate test. But I do know we did the four node tests, I do know we did some of the one node testing right after it arrived, and there may have been one in between them. I have to check records.

Cerf: I honestly don't recall anything intermediate. I think my recollections most strongly are of the arrival of the machine at UCLA and all the trouble we had physically getting it into place. Mike Wingfield trying to work [through] a BB&N 1822 interface hardware design for the Sigma 7 computer that would drive it, or use it, and then you and Walden coming back either late in '69 or early in '70. My recollection is it was early in '70. By that time we did have the Sigma 7 installed because we went through a whole raft of tests to drive traffic into the IMP to see what it would do. Now my recollection is that when you came out you had in mind some problems that you thought would be exhibited under certain conditions, and your colleagues didn't think that they would happen, or that the probability was so low that they could be ignored.

Kahn: Are you really trying to get into all of that?

Cerf: Well, I think at least a little bit of that would be appropriate. I recognize we don't want to dive so deep into ARPANET that we never get to anything else, but I thought these ideas that you had that your colleagues kind of dismissed were pretty important suspicions.

Kahn: Well, I thought they were too. Katie Hafner is a wonderful writer. She's written for Newsweek magazine, she occasionally writes for the New York Times, and she wrote a book that was a really great read called "Where Wizards Stay Up Late" Katie sent me an early version of that book. She had gotten a recount of what happened here from a number of people. I don't know exactly who. And I read it and I said, "Katie, that's just not what happened." What she described was my going out to UCLA and running these tests and discovering that the network froze, locked up, sort of like the packets got all clogged up in some area. She said in this write up -- and it was just totally wrong, but she said in the write up -- that I then went back to BB&N and reported what I had discovered, and that I wanted to write a simulation program to go study it, and that Frank Heart basically said no because we have a real network, and he instructed people to go fix the problem. And then… A long story short, Katie ended up taking that part out of it. I wanted her to write it and do it right but I guess she didn't have enough time, so she literally deleted that whole part from the book. But the reality was I sent her a paper that I had published at a Princeton University conference describing the simulation that we had already built a year earlier, before we even got the contract on the network, to look at the problem of deadlocks which I knew were a possibility, and we had actually built the simulation. Warren Teitelman did the implementation and it was really a… Well, Warren was the BB&N LISP man. He was an InterLISP guy for a while, and I think he may now been at Google for all I know. I believe he's working…

Cerf: Well I'll have to look into that. That's interesting.

Kahn: I got a note from him recently. He has been doing some work for Google. Anyway, Warren wrote that part of… that program. It was a really interesting one because most people were working on model 33 teletypes, and this produced an interactive graphics display off the timesharing system, when people just weren't doing that. You could actually design a network, put the nodes, you could design what kind of traces in each of the nodes. You could watch the packets flow through the net, you could collect all the statistics and you could actually see the lockups, you could evaluate routing algorithms and the like. We built that and Jerry Elkind, who was then running the group, had just learned that the contract was awarded, and he said if you want to keep working with that simulation you ought to charge that to Frank Heart's contract because Frank Heart was the PI on the contract. I thought once we had won the contract I was going to bail out and end up going back to teach. Except that Frank refused to let me charge it to the contract and Jerry was saying you can't… This was Politics 101 right in front of me. Jerry said we can't fund it here, it's got to come out of that. Frank was saying I don't want you working simulations because we already have a real network that we're building, use that, and oh, by the way, if you want to charge it to the contract I want you over here working for me. After debating it I picked my office up and I literally moved over to the part of BB&N where Frank were. People didn't do that. This was the equivalent of moving from the math department to the chemistry department. It was a major move. People were, I think, very surprised when I had made the move, but to me it was survival. It was survival in two ways, number one, because I couldn't continue to do what I was doing without it and number two, they didn't understand that there was a real problem. In my opinion they didn't understand that. So we went over there, and what we ended up doing was not working on the simulation -- because I wasn't permitted to -- but trying to figure out how to make sure that the implementation didn't have these problems. But because I didn't control the coding -- Bill Crowther was writing it -- when I went out to do the field testing -- that was my job, both designing it and then testing it to see that it worked -- the first thing I went out to do was to show that it would lock up, because I had proposed, and wrote for the proposal, all kinds of mechanisms to deal with this issue. And it was overruled because the feeling was the likelihood of packets clogging up the net, because there were going to be so many of them going all over the place, was about as likely as all the oxygen molecules in this room suddenly showing up in that corner and therefore us all suffocating, which is not going to happen and if it does we'll deal with it later. That was the point of view. That's what I did: I went out and locked it up. So when Katie went to report it, she said I went out there and I discovered it from my experimentation, which wasn't…

Cerf: As opposed to deliberately inducing the effect, right?

Kahn: I induced the thing I knew about all along. We ended up fixing the problem eventually. It took a while. We came back and reported it. Dave Walden actually did the report, and nobody believed it, even though we had the physical data. It took a while before… Bill Crowther's conscience must have gotten to him, and Bill Crowther ended up writing his own simulation in the back room to sort of verify it. He came back and reported it to Frank, and said there's a problem here, the thing that Bob was talking about can really happen. Frank came to me and said how come I didn't know about this? So I said you did. I reported it, but he either didn't remember, or it hadn't gotten through, or he needed to hear it from one or two of his critical people. We then spent the next year redoing the whole implementation to deal with that problem. It was a major redo of the software. Finally fixed it, and when it finally got fixed my idea was okay, I'm off the hook now, I know this thing is going to work. I picked myself up and moved back across the bridge to the group that Jerry Elkind had been running. Jerry by that time had gone to MIT to work for Licklider, and Licklider had hired him over there. I think Danny Bobrow was running the group at the time, and Bert Sutherland was working for him. These were two people who later went to Xerox PARC.

Cerf: These are all legends in the industry of course.

Kahn: Along with Warren Teitelman of course. I moved back and shortly thereafter ended up planning the first public demonstration of the ARPANET at the Washington Hilton, and right after that ended up going to DARPA. Larry Roberts had actually offered me the job at DARPA something like May of ’71. I wasn’t prepared to do that at that point because we hadn’t finished what we needed to do. Eventually I decided that that was in fact the thing I would do, and I decided to go to DARPA literally to run a project in automated manufacturing. I was going to sort of say okay, that’s it for networking, I’m finished, I’m going to do something different.

Cerf: It’s a permanent infection that you can’t get rid of.

Kahn: You’re stuck with your own history.

Cerf: Something like that. Actually getting to DARPA in 1972 was a little more complicated than simply picking up and going there. My recollection is that Larry had asked you to organize the demonstration at the International Conference on Computer Communications in the Washington Hilton Hotel in October of that year, and you drew on a lot of the ARPANET community to do that. I remember being there, and many of my contemporaries too. But then it seemed to me that somehow or other there was a funny swap between you and Larry Roberts, because BB&N had gotten interested in the possibility of commercial networking and that manifested itself in the form of a company called Telenet. Can you tell us a little bit about your involvement in the formation of Telenet, and Larry Roberts’ involvement subsequently as you leave BB&N and join DARPA?

Kahn: My recollection is that BB&N had… There had been discussions about setting up a commercial operation out of BB&N, but the company had not decided to do it. I think Frank might have liked to do it. They had hired a gentleman whose name I don’t recall from Washington to come up there, and he left to start what became the very first packet company in the United States called Packet Communications, Inc. They filed with the FCC to get permission because they thought they needed to at the time. Probably did. But that company didn’t get anywhere. I don’t know the details of it, but it didn’t actually survive. But when he left, BB&N almost instantaneously made the decision to go down that path, almost as a antibody reaction I think. Steve Levy approached me. Steve was at that point sort of the entrepreneur in the company. I think he might have been the CFO or Treasurer. This was Steve’s bailiwick to do these sort of things. He and I went together and planned out the creation of this company, Telenet. Steve was the first president of it, and he wanted me to stay on with Telenet in some chief technology role. I had already agreed to go to DARPA a few months before. This would have been maybe May or June of 1972 that I told Larry that I would come to DARPA. Aside from the issues of integrity of telling somebody you were not going to do something you had decided to do, I really wanted to work on this automated manufacturing project. I thought it was really interesting in so many dimensions.

Cerf: We need to come back to that because we don’t know about that, but let’s—

Kahn: It was a big project. I agreed to work with Steve over the next several months to plan Telenet. It was just Steve and myself doing it. None of the people that were in Frank’s group were involved in this. In fact, it was a big surprise to him, I think, when he found out that this was happening. They set up Telenet, I think it was in August of ’72 formally. It was housed up in the Boston area at BB&N. When I joined DARPA, which was in October, Larry Roberts left at about the same time to become the second president of Telenet. Most people think he started it, but he was actually the second president. He persuaded Steve Levy to let them move the operation down from the Boston area to the Washington area, and they set up shop somewhere around 16th and K or 17th and K, in one of those corner buildings. So that's how Telenet got its operation. Larry had left and he became involved in that for quite a while, and I ended up at DARPA trying to run this manufacturing program.

Cerf: Did Larry also bring Barry Wessler with him to the Telenet activity?

Kahn: Yes, Barry had been a program manager at DARPA. I think he went with him. Interestingly enough, Barry was in one of my classes up at MIT when I was there so I actually knew him from way back.

Cerf: I'm not surprised. This is a small community and there is an amazing amount of interaction and coincidence I think.

Kahn: I believe Barry and Larry worked together on that for quite a while.

Cerf: When you talked to Larry about coming to DARPA, what was it that was proposed that you do there?

Kahn: They had proposed, in their congressional filings, to set up a program in automated manufacturing. It was purported to be a fairly large program, hundreds of millions of dollars kinds of numbers, and back in the '70s that was real dollars.

Cerf: That was real money.

Kahn: I was interested in it for a number of reasons, because the whole idea of how you make things in factories by specifying designs and causing machinery to get put in place that would cause it to be built. I could see network activities taking place on the factory floor, I could see database systems, I could see artificial intelligence, I could see computer aided design. I could see real utility, which was starting to more and more influence what I did. I found it kind of a fresh, brand new area. Nobody had really done this before and that appealed to me. Leaving the networking stuff behind was difficult, but I didn't see any way that I could be very much more productive in the networking field at that point, because to do it would have required that I stay at BB&N to do it. I was now back in a different division. We had all of those old political schisms, and funding schisms, and I was looking for the next thing for me. BB&N had the view that the ARPANET was going to be the network for the world in general. It didn't quite turn out that way, but it was a big deal for the DOD, particularly in terms of the defense data net that it became part of later on. But this seemed like a really good move for me. So that's what I went to DARPA to do. About a month or two after getting there, Larry broke the news to me that -- it must have been very quickly he broke the news to me, I don't know when he learned, he might have even learned before I came but I've never asked him -- that the Congress had killed the program.

Cerf: Actually what this says is that you and Larry didn't switch positions at the same time. You came to DARPA while Larry was still running the office--

Kahn: We overlapped for—

Cerf: --and then you left subsequent to that.

Kahn: I came in in the end of October of ’72. He left in September of ’73.

Cerf: Oh, so there was almost a year of overlap.

Kahn: There was close to a year of overlap.

Cerf: I had misunderstood that this was almost simultaneous, but it was not.

Kahn: No, and it wasn’t a swapping of positions because Larry replaced Steve Levy running Telenet. They hired him in as the CEO, whereas I had no formal role with Telenet, ever. I was on the BB&N staff working with Steve in setting it up, but I declined to get involved with it directly.

Cerf: You weren’t an officer of Telenet, you were just helping facilitate that.

Kahn: I had no role with Telenet -- by my choice -- but I worked very closely with Steve in setting it up.

Cerf: Sure. When you came to DARPA and this automated manufacturing program gets canceled, what happens then?

Kahn: See, I don’t know that it got canceled. I don’t know that it ever got started.

Cerf: Fair enough.

Kahn: I was sort of hired to go do that on the assumption that it was in the budget and would happen, but it never actually happened. At that point I seriously thought about just… that I came there to do that thing, and I was packing up and maybe go back up to Boston. Probably not back to BB&N but probably back to MIT.

Cerf: Is it fair to say that you could not, under government rules, have then taken over responsibility for the ARPANET project because you’d been working on it at BB&N? Or was there any issue there?

Kahn: The bottom line is I don’t know, but I chose not to get involved in any event, partly because of the fact that that’s where I had been. There were enough issues that I was part of at the time, personality conflicts, technical judgment issues. It just made no sense to me to go there because all it was going to be was arm wrestling one way or another. So I never got involved with the ARPANET project in all the time I was at DARPA.

Cerf: Larry continued to run it while he was still the Office director?

Kahn: Well, Larry masterminded the whole thing when he was at DARPA. It was his specification. He pushed it. The ARPANET would never have happened if it weren't for what Larry had done. Of course, he got support from Bob Taylor, he got support from Charlie Hertzfeld who was the director when it started, he got support from Ed Rechtin, he got support from Steve Lukasik. He had continuing support along the way, but Larry was the key.

Cerf: When he left a year later, was Steve Walker the one who ended up with the responsibility for ARPANET at that point?

Kahn: Steve Walker was responsible for a period of time. I think he was responsible during that initial period, and then Licklider had come in in 1974 when Larry left, and he hired Craig Fields to join the office and Craig I think took over the ARPANET responsibility and ran that for a while. But I had… They… People would talk to me, they'd ask me what about this, what about that, because I was a kind of font of knowledge, having been instrumental in initiating it. But I just wanted nothing to do with the project myself.

Cerf: At least not in any official capacity.

Kahn: I didn't want to have to deal with the BB&N folks on that particular project.

Cerf: After the manufacturing project doesn't happen, what did you wind up occupying yourself with?

Kahn: Well, I was getting prepared to leave. Larry basically sat me down and said look, you can't do that -- you shouldn't do that -- you know as much about networking as anybody in the country at this point. I think that probably included him. He said we got too many options and opportunities in the future, more networks we can create, and you need to make that all happen. It was persuasive enough for me, having just moved down, having just bought a house. To just undo all of that in the span of no time just didn't really appeal to me. So I said okay, I'll give it a try. I got involved in setting up a number of programs. I think the one that was most immediate was the packet radio program: to apply packet switching technology to mobile radios. People had been thinking about it. The idea of such radio nets did not originate with me. Norm Augustine and Frank Kuo had gotten funding years before to build something called the ALOHAnet on Hawaii to link radio terminals—

Cerf: I'm sorry to interrupt. I think you meant Norm Abramson.

Kahn: Yes. What did I say?

Cerf: Norm Augustine who was over at… It's Norm Abramson…

Kahn: Norm Abramson, right.

Cerf: …at University of Hawaii.

Kahn: Sorry. I know Norm too. Norm Abramson and Frank Kuo, University of Hawaii, had gotten funding for something called the ALOHAnet,to link by radio terminals to the central computer at the university. So the idea was about, but nobody had ever tried to actually build a whole network of radio nodes that could do relaying.

Cerf: Do you recall roughly when the ALOHAnet work was going on? The reason I'm asking is that there is some correlation with work on Ethernet and Bob Metcalfe. There is a confluence in here of another person whose name hasn't come up and that's Steve Crocker, who I went to high school with and later was at UCLA with as a graduate student. But the reason I bring it up is that Metcalfe was exposed to the ALOHAnet at Steve Crocker's house here in Washington when Steve was still at DARPA. Could you say a little bit about when the ALOHAnet was happening, and when did Steve join DARPA or ARPA at the time? Was it before or after your arrival? Do you remember?

Kahn: Well, he was there before me. I think he showed up there sometime in the 1970 or '71 time frame. He left shortly after I got there. We did not overlap very long. My recollection was that the original funding for the ALOHAnet actually came out of either ONR or AFOSR. It was a small thing that DARPA then picked up and helped to grow into a larger effort, and that it started roughly at the same time the ARPANET did. So I would have dated it -- and this is without checking my historical facts -- I would have dated it roughly to 1969. What happened before then, preplanning, I just don't know.

Cerf: This sounds pretty consistent, because my recollection is that they named the device that was handling the radio packets the Menehune, which was supposed to be an IMP or a little—

Kahn: Hawaiian IMP.

Cerf: --a little leprechaun, and that probably was a side effect of having seen what the ARPANET packet switch was called as well.

Kahn: I think Dick Bender, who later worked for CNRI, actually at one point told me that that was the name he chose because he was working on the project as one of the people. But I don't know that for a fact, and I may be recollecting it wrong.

Cerf: That sounds like about the right time frame, because my recollection is Metcalfe goes out to University of Hawaii just before he comes back to Xerox PARC and invents the Ethernet with Dave Boggs, and that's all like 1973. So what happened? You've been exposed to this packet radio idea. Has the project actually started at that point or are you—

Kahn: No, there was no project in place. Larry had been thinking about it, I guess. As with many things, had Larry stayed at Lincoln Labs I'm sure he could have built the ARPANET. Had Larry stayed on at DARPA I'm sure he could have made the packet radio program happen. It wasn't like every one of these were inspirational things. The hard part was, I think, in conceptualizing them, and architecting them, and causing them to be created. When I first started working on the ARPANET I thought it was a new idea with me. I was not familiar with the stuff that Len had done, or Paul Baran had done, or Donald Davies. I was just exploring what I thought was an interesting idea. When I brought it to the attention of the folks at DARPA I had no idea that they were even thinking about networking or knew any of that stuff. But the packet radio work was a clear extension of the ARPANET. Instead of the nodes being fixed nodes with a lot of wire connections, they'd be mobile nodes with a lot of radio connections. So the idea was evident, as to what this was. But how to do it was a matter of a program, and that's what I set up to formalize.

Cerf: What was the motivation behind doing the ARPANET in the first place, and what continued that motivation in packet radio and packet satellite, which we've sort of bumped into a little bit.

Kahn: Vint, I'm surprised you don't know by now.

Cerf: I'm sorry? [laughs]

Kahn: That's a tongue in cheek comment. I know you know well. So just for the record, I would say that the motivations were multiple. Motivation number one was to explore the area of computer network resource sharing, so different computers… We're dealing with an era back then where the machines were all different, we didn't really have a lot of portability, and so if you wanted to use a resource you had to go there or get there through a network. Programs on ILLIAC weren't going to run anywhere else, and there were a lot of specialized resources as well. Trying to figure out how to share resources. If you had a simulation program running on one machine and somebody wrote one on another, how could they work together. It was an opportunity for the future, to break new ground. That was one motivation. A second motivation was specifically to allow the research community access to specific resources. Accessing a very large database storage system wasn't so much an interaction thing -- although maybe you want to bring it back to another machine -- but maybe you just want to see a data point on a screen. I kind of lump those two together in the computer resource sharing activity. But the second question was: how could you build a network that would facilitate that in an economic and efficient way? The use of the regular telephone network just didn't seem appropriate. Number one, the bandwidth was too small. Maybe you could get 2.4 kilobits, but people were thinking we need at least an order of magnitude more than that. How could you possibly have a network that only had 2.4 kilobits of bandwidth in the backbone? And especially if you're multiplexing a lot of users. Number two, how could you do it reliably, because phone lines made errors and computers needed perfect communication. The bottom line was it really needed a rethinking of how to do the networking in a cost effective way. You could have bought lines and hooked everything up in a fully connected way. But this idea of packet switching had been floating around, and so DARPA was interested in showing that packet switching in fact was an efficient way of dealing with computer communications. They were trying to demonstrate a new mode of communications. Then finally, another motivation for it -- although it wasn't so much funded in the research community as through interactions with some military sites -- was to show its utility for command and control on a more global basis. But that was not driving it. That was one of the stated applications of it that they would explore. I think the motivation was more of a research motivation at the time.

Cerf: The command and control argument surely must have resonated with the people who funded DARPA, recognizing that computers could be used in command and control to manage resources better than they would otherwise. I seem to remember from that time period phrases like 'force multipliers', that would allow a smaller force to overcome a bigger one because you could manage the resources better with the help of computers.

Kahn: It is certainly true everybody understood that the potential was there. At some level you could argue that it was supported because of that potential, certainly at the level of the director or down below. But if you look at the people in the trenches who were building the technology and doing it, they thought they were solving a technical problem. This was not a military problem that had some urgency and they needed to decide how to get a solution by two weeks from Friday…

Cerf: In fact, the art of ARPA, I always thought, was transforming problems that were known to be faced by the military into problems that the research community could work on without having very much awareness or exposure to the specific military application. But since you mentioned packet radio and mobile computing, it's pretty clear that that's tactical. And there is the ships-at-sea problem which must have led to the use of satellites for that communication. So somewhere in this picture the packet satellite pops up too. Is that part of the ensemble of things that you were working on?

Kahn: I would take a lot of credit for the ARPANET work, certainly designing that. And I would take much of the credit for the packet radio stuff -- architecting -- and I was sort of the chief architect for that and orchestrated the building of it too. The packet satellite work had a very different history. I would argue that as a network concept Larry Roberts really started that one, but he was never able to follow through. He had done some early analysis of slotted ALOHA on satellites as opposed to plain ALOHA. ALOHA was a technique, a name of a technique that the ALOHAnet had created, where if you had a radio net, just send whenever you feel like sending it, just hope it doesn't collide with something. It would be like a traffic management system that basically said: go anywhere you want any time, and if you collide we'll retransmit you. Not very good for cars, but it did work for packets. If the traffic was light, most of the time you wouldn't collide with anything. Well, it turned out there were some theoretical limits. Larry showed that by time-slotting the channel you could avoid some of the collisions that would occur when things were very close to each other, by just making sure that they missed a little bit, and then you could double the capacity. So people were thinking about that, and he had even funded BB&N to build a variant of an ARPANET IMP that had more memory, to control a satellite channel. Unfortunately, at the time he was not thinking about the satellite channel as being anything more than a modem connected to the IMP. So it was just like any other connection except it went through a satellite, and there were no satellites to use. So even though BB&N had been funded to write the code, there was no deployment planned that was possible at that time, because nobody had worked on it. And the biggest issue with that whole approach, as I discovered when I got to DARPA, was there was no way to separate that into two separate nets. That is the satellite net, as it finally became known, actually became a separate net because we split it apart and we defined interfaces between the satellite piece and the ARPANET piece that allowed the standard ARPANET interface to be used to the ARPANET. We defined a separate interface to the satellite net, and that allowed for a device to be put in place which we called a gateway, that eventually became critical to the original evolution of the internet. Well, as long as both pieces of software were in one machine, then all the interactions between the two were in core memory transfers controlled by one group at BB&N, and to me that was a nonstarter. The internet never would have happened if that was the model that we applied in every single case; that is every single other network was like a modem connection to the ARPANET. The ARPANET would then be the central thing that kept propagating. We'd never get different administrative controls, different administrative participants, different political… The open architecture that led to the internet never could have happened that way.

Cerf: We were just talking a little bit about packet radio and packet satellite, and you mentioned the gateway notion. This is a point where we come to Yogi Berra's fork in the road -- which we are going to try to take, right?

Kahn: Well, let's take it.

Cerf: There are two ways to do this. You have a long history at DARPA. You were there for 13 years and you did a lot of things in addition to the things we talked about so far. What I'd like to suggest is that we stay on this track, continue talking about the open networking ideas which I think you actually were working on before you got to DARPA, and then talk a little bit about how all that evolved into Internet. And then come back to look at some of the other programs that you were involved in while you were at DARPA, so as to try to keep this thread more coherent. You mentioned just briefly this open architecture idea. Was that something you brought to DARPA or is it something that evolved after you got there?

Kahn: Well, it must have been in my head but I don't recall consciously working on that before I got to DARPA. What I had been working on at BB&N was the notion that we really needed a different kind of protocol in those computers than the original NCP one that you and Steve Crocker and others worked on. That protocol, for a variety of reasons, ended up being more like what I would call protocol that could deal with a device like a line printer, where if stuff didn't get properly conveyed to that device, you'd press a reset button and redo it. You wouldn't expect radio signals to jam the printer. You wouldn't expect lots of packets to show up on that interface, and so forth. The model of the ARPANET was a very reliable communications device that would just convey things, and so the device at the other end was just like a line printer. And I knew that in the world of the future we needed to look for something else. We might have radio links. We might have all kinds of other things. I was not thinking of multiple networks, but that we needed a protocol in the machines that would behave more like a communications protocol. I actually wrote a paper on that subject, and it was called something like "Communication Principles for Operating Systems" that Bob Metcalfe cited in his thesis. It was a start down that path. That's really what I brought to DARPA, along with a lot of expertise in building networks. When I began to work on creating these other networks at DARPA, that's when the notion that we needed a way to make them all work together showed up. That's where the open architecture notions came from. I think it's fair to say that I had a pretty good notion architecturally of what they should look like, and a fairly good level of specificity, without working out the details of course. You and I, probably in a way neither of us could have done all by themselves, really worked that out in a lot more detail at the specification level. The idea that you would have gateways between these nets that could handle the protocol translations -- that could deal with all the impedance matching of error messages and whatever -- that's all part of what I was looking at at the time. The question was how do we go from a very general idea like that to something specific that could deal with this? I knew we had to address the devices in a more general way, because the ARPANET lets you address wires on it, so if you had something on a net three away how are you going to specify what that was? I figured things would have unique identifiers but I didn't know what they would be at the time. I started working with Vint Cerf, party of the first part, and together we ended up detailing that. Vint brought detailed knowledge of protocols to the table. He actually came out of the computer science side of the world. That's a world that I hadn't come out of. I came out of the communications side. So, we brought different things to the table but it was a collaboration like no other one that I ever had in my career, because it just clicked. We often just saw eye-to-eye on everything. Every now and then we wouldn't, and once that would happen we'd have this heated debate and we'd pretty soon come to the conclusion that the problem was just we had different assumptions about something or other. In later years when we were working together when that would occur, one or the other of us would sort of say, "Okay, wait, let's stop. Let's go back to our assumptions and figure out what it is, because if we're disagreeing on this it's probably in an assumption, not the reality." I was saying "Here's what we need to do" and I'm thinking five years from now, and Vint is thinking "I got to do it in the next two weeks. How’s this going to work?” Between the two of us we sorted that out. We could try and see through any arguments that were of fundamental disagreement on the technical level, but I'm not able to stipulate any right now that didn't have any genesis in different assumptions that were brought to the table. So, together we were able to work out the nature of what this architecture looked like in fairly detailed form. We wrote it over… We had a lot of meetings on the subject, both at DARPA and at Stanford on the west coast. At one point, we then decided we got to get this down on paper. We got a conference room in the old Cabana Hyatt, which was on El Camino right opposite the old Ricky's Hyatt.

Cerf: South of Stanford University, yeah.

Kahn: Which I guess neither of them are there, although the Cabana building may be there with a different name. It was one of the conference rooms on the first floor, with a name like the Lanai Room, or something like that. I can't remember exactly. And we sat down to just write it. We wrote it over the course of I think it was a weekend. And I remember that so vividly because we pretty much had in our heads what this was. It was just a matter of getting it on paper. I remember saying to Vint at the time "Do you want to start or shall I?" He said, "No, I'll be happy to begin." He took a pencil and a pad, because that's all we had at the time, no laptops. And he sat there for all of about four or five minutes and nothing was happening. Not a word was going down on the pad, nothing. I finally said "How you doing?" He said, "Well" he said, "I don't know where to start." And the space was so large. So, I said, "Fine, why don't you give me the pencil?" So I started to write it out. And I wrote. Not that he didn't have it in his head, he just wasn't sure where to start, I guess. So I wrote out the first eight or nine pages. Then he said, "Okay, give it back to me." And so he took it and he wrote the next 20 or 30 pages. Then we edited, and reiterated, and that's what produced the very first draft. Which I wished he had kept a physical copy of it, because it's probably worth its weight in--

Cerf: We had my secretary at the time, Caroline Tajnai, type the thing up. I thought after she typed it there was no reason to keep the written version, so we threw it away. I threw it away.

Kahn: So, she threw it away.

Cerf: There's a life lesson hiding in here, of sorts anyway, at least a practical one. Because Bob's style of writing things is to sit down and just start writing, and it doesn't matter in what order, the point is getting the content down on paper.

Kahn: Getting it out of your head, yeah.

Cerf: Getting it out of your head. And I was trying at the time to figure out how the whole thing was going to look, and then start writing it at the beginning, the proper beginning. So I was stuck because everything depended on everything else and I couldn't figure out where to start. So I learned a vital lesson, although I don't exercise it as well as Bob does. He says, "When you have to write a paper, just start writing. It doesn't matter because you can go back and edit, especially thanks to computers." So, thank you.

Kahn: Well, I learned a corollary of that from Vint Cerf, which he probably doesn't remember right now. But when you're trying to make something happen, the first step is just do something. And Vint at one time had a big sign over his door that said "Do something." Just do something. It's sort of the implementation equivalent of…

Cerf: Actually, the origin of the "Do something"… We're jumping forward in time here, because the person that told me that was Josh Lederberg, who is on the board of CNRI, and we were talking about digital libraries. I was describing all of the things that could be done, and Josh stopped me part way through all this, and he looked at me and he said, "Do something." I thought that was pretty good advice coming from a Nobel Prize winner. We will come back to CNRI later on. I want to try to figure out what to do here. We've got this paper. By this time it's September of 1973 and there is an opportunity to present the ideas before what was called the International Networking Group.

Kahn: Which was meeting in Sussex, England sometime in September of that year. I go over there. Vint arrives a few days late with the box of the papers that his secretary has now managed to type up, minus the originals. We present the material at this working group meeting. The reaction of people there was, I think, muted relative to what it should have been at the time. People often ask, "What was the press reaction to the first installation of the IMP at UCLA?" And we said, "Well it was a non-event." It was on the shipping dock and the engineer brought it in and it wasn't written up in the press. It was not like Lindberg's first flight to Paris. Nothing happened. Well, nothing much happened at this meeting either. We gave the presentation and people said, "Yeah, very nice" and that was sort of that. We went back later, revised the paper in cosmetic ways, mainly so it would be suitable for publication, and IEEE Transactions on Communication published it in May of 1974. We were later told, to our great surprise, that it's one of the more valued historical papers. That was very interesting, but you have to have the original journal, which I do.

Cerf: Oh, boy, I don't anymore.

Kahn: Well, you weren't a member of IEEE at the time.

Cerf: It's now worth at least $7,500, according to one of the antiquarian bookstores in New York, who is offering one of theirs for sale for that amount of money.

Kahn: Except mine has my name on it with the mailing label.

Cerf: Holy smoke.

Kahn: So that's like having a mint.

Cerf: Wow, you're right, that would make it even more valuable, yes. Well, if your retirement plan doesn't work out that might be the beginning of something new.

Kahn: No, I think I'll just keep it.

Cerf: So, we're at the point now where we've got this paper published, but the work continues at Stanford.

Kahn: And UCL and BB&N.

Cerf: Well, no, actually not. Because my recollection -- maybe you have a different one -- is that the specification, detailed specification, before we could write code, took most of 1974. And I thought that we didn't engage BB&N and the University College of London until we had the full TCP spec, which didn't come until December. That would have been my rendering of the milestones.

Kahn: One would have to go back to the record, and I don't know exactly where that would be located. It may be in the archives. But BB&N was actually working on two separate things on the Internet side at that point in time. One was the TCP/IP spec, which was more dependent on the work that Vint and company were doing. But they were also building parts of the packet radio system. The packet radio system was controlled by a mini computer that had a station in it and that station actually had embedded within it an interface to the ARPANET. So in some sense they were going down the path of building gateways before we had the full TCP/IP spec. There never was a full spec for a gateway written. Nobody ever wrote a document that said in detail "This is how you implement a gateway." All we did was specify what the functionality of the gateway had to be. And some of the work on the IP addresses really didn't come until later either.

Cerf: Actually, I think either you discounted a particular document or you didn't remember it. But Virginia Strazisar at BB&N wrote a document, IEN 30, called “How to Build a Gateway”. This was BB&N's prescription for what you had to put into the gateway in order to make it work.

Kahn: No, I know. But that was what BB&N was saying about how they would do it, and that's during the whole period while they were already working on the gateway stuff. I'm just saying that when I said that BBN was working on it, they were working on that in parallel with the work that was going on, on the TCP/IP side of the world.

Cerf: That document came out-- this one we could check -- I'm pretty sure that document came out after the TCP specs had been done, because it was all about encapsulation. The station which had the interface to the ARPANET in, it eventually became the gateway between ARPANET and packet Radio Net. But a lot of those things didn't get tested until '77.

Kahn: ‘77, '76, that's right. No, when I said they were working on it I don't mean they were necessarily implementing it in the form that you were talking about. But as part of their work that was one of the things they were supposed to be doing.

Cerf: They had to, because the only way they could… The packet radio system was being implemented at SRI right?

Kahn: They were doing the integration and field tests.

Cerf: Right.

Kahn: The radios were being built by Collins Radio in the Dallas area, and the control part of it was the station, which was a PDP-11 minicomputer, and BB&N was building that one.

Cerf: And it was Jerry Burchfiel’s group that was doing the station?

Kahn: Jerry, and Ginny was part of that, a number of others.

Cerf: We did some tests, if I remember right, of packets flowing between the packet radio net at SRI and ARPANET, somewhere in the '75 or '76 timeframe, before we had implementations of the full TCP/IP.

Kahn: Uh huh.

Cerf: Anyway, my recollection is that we didn't start seeing code coming out of BB&N for TCP until sometime in '75.

Kahn: Oh, I'm sure that's true. But was there some part of what I said that you were having trouble with? Because I don't think what I said is challenged by this. I think that's exactly in agreement with what I said before.

Cerf: No, the question was when did UCL and BB&N start working on the Internet? My recollection is that they didn't start until the '74 TCP spec was done.

Kahn: The TCP stuff. But what I was talking about, the BBN stuff, it was on the gateway stuff and if I threw…

Cerf: That was on the packet radio program.

Kahn: …if I threw UCL in the mix I probably should not have. They were not, at that time.

Cerf: So we've got this time period now from '74, '75, through '78 when this protocol starts getting more and more refined and everything else. In that same period of time I leave Stanford and come to ARPA.

Kahn: Sometime in 1976, as I recall.

Cerf: Right, yes, in August or something like that. What I'm curious about is how we characterize your involvement in Internet as this all unfolds. Because for a while you get to focus on a lot of other programs at DARPA because I'm at ARPA helping, but then I left in 1982.

Kahn: Between 1972 when I came, and 1976 when you arrived, I was basically running every one of those programs, with the exception of the ARPANET program which was run by other people. We had efforts to build end-to-end security devices, that I started. We had programs to work on sending packetized speech, the precursor to VOIP technology today. Remember, we were sending packetized speech through a 50-kilobit network, and one conversation in principle could have been, I suppose,e 64 kilobits per second. It's hard to get all those bits through. We actually were demonstrating how to do it. Later we were using compressed speech techniques using LPC, CVSD, and other techniques. We were actually able to send full blown conversations through the old ARPANET. It required changes in certain protocols, because we couldn't acknowledge every packet because you needed to maintain continuity, but we did that. So there were a lot of programs. I was running the packet radio program, the satellite program, the security program, the command and control program, all of that stuff. When you showed up in '76, by that time I think I was Chief Scientist and probably Deputy Director.

Cerf: You were Deputy Director at that point.

Kahn: I was more than delighted to see you show up, because we suddenly had somebody else who was capable of providing vision and leadership for that. We worked together on that pretty closely for the next six years. I would say that the most significant thing that we did during that timeframe was to set up the initial things that became the Internet standards process going forward. That's a very interesting issue because Vint and I had many discussions about what would happen if one or the other of us were hit by a bus, or the equivalent. I was more concerned about him, because I knew what I would do, and would have to take it all over again and I really couldn't. I didn't have the time and couldn't afford to. So I implored him to set up a kitchen cabinet of some sort to help with the management of the program, mainly so people in the community would know about what's going on. He kind of resisted that notion as a notion, but then he came up with the equivalent with a different context. He said, "Why don't we set up a group to help us configure the Internet? People are going to have to know how to deal with that." And he called it the Internet Configuration Control Board, which got populated with roughly a dozen implementers from around the country, mainly people he knew well and that he had been working with. I think we got Dave Clark from MIT to chair it. That became the sounding board for what we did. It didn't have any real responsibilities at that point, but it did have a lot of knowledgeable people who were in very close contact with us. We were still making the decisions at DARPA, but it was a start of something of significant proportions later on. In 1982, when Vint left to go to MCI for the first time, I took over the program again and started to run it. During that period, we had brought in another gentleman who is now deceased, Barry Leiner, and Barry took over part of the packet radio program. He had been introduced to a lot of these programs during that period of time, but he really didn't know much about the Internet at that point. So I took it over and ran it for another year and tried to get Barry to become more involved in the Internet so he could pick it up. His approach to things was very different from Vint's. Where Vint could literally pick up anything and engineer it on a moment's notice, and do a good job, because he was an engineer and he was an implementer as well as a pretty good manager of others, but he was good at doing it himself as well. Barry was not of that ilk. Barry was a good process guy. Barry had a lot of technical expertise. But he wasn't about to make decisions about this, because it was a big system that by 1982 had already become pretty complex. He was more comfortable relying on another group to come back with the recommendations. Or, I suppose he would have taken them from me, although I didn't interfere too much at that point. I introduced him to the ICCB. Now Vint had gone. He became uncomfortable with that whole mechanism. The reason was, when Vint had started it, it was like 12 people and a few what I call strap hangers, people who just wanted to sit around and listen. But by 1982, because that was where almost all the really interesting discussion was happening that was happening anywhere outside of DARPA or in the community in a collective way, people just wanted to sit in and listen. I don't know the exact number but there were several hundred people that typically we had to invite into these meetings, so you couldn't have a meeting of 12 by getting a small conference room. You had to have an amphitheater -- theater in the round kind of style -- in order to have a meeting of just the set of 12 people. Barry looked at it and said, "No." He and Dave Clark got together and they decided that they would like to recommend a restructuring of it. Namely: we get rid of the ICCB, create another body which they called the Internet Activities Board, which essentially had the same people plus or minus a few, and to take all of the discussions that were going on in that group that people wanted to listen to and put them into a series of task forces underneath. There were ten of them initially. One dealt with autonomous systems, end-to-end routing, privacy, security. It was one of those task forces whose job it was to maintain the punch list, because Vint was gone by now, and the Internet was sort of being born from the introduction of the TCP/IP protocols. What we needed to do was get all the hosts on the net sort of aligned. This other group had the responsibility for the so-called punch list. If you ever bought a house you know that before you take possession you want to make sure the water works on this, and the switch is okay there, and the plug is installed here, and the leak is fixed. They are the details. That's what's the punch list. This group, which was at that point chaired by Ed Cain from DCA, was in charge of the punch list. The other ten working groups could do whatever they want wherever they wanted, and it didn't have to all be in one place, and they could meet asynchronously in different places. That was the formation of what's become the current standards model for the Internet. This was an early form of it, but it had the right shape and form. And it grew over time. From ten initial groups a few years later it was up to 50. Suddenly the ability of the IAB -- this 12 member group -- to manage all of them, because they were chartering the groups and they were overseeing them, just became too large. The responsibility was then given to this other task force, the Internet Engineering Task Force, which had pretty much prosecuted its punch list, to manage the other working groups. There was another one set up to deal with research, but that was the beginning of the structure of the modern standards process that's been used to manage the Internet. That took hold and started to grow. In 1986, after I had left DARPA, Vint had also left DARPA a few years earlier, and Barry was leaving at the same time, DARPA made the decision to get out of networking. Normally they go into an area they've funded for a while and then they move on to the next thing. Well in this particular case, they decided it's time to let it go. NSF was very interested in taking it over. Erich Bloch had recently joined as the head of NSF, and he wanted NSF to get into networking to broaden it out, and to the science and education community more broadly, and they did. I think one of the seminal decisions that NSF made -- and I attribute it to Steve Wolf but I'm not sure who else was involved -- was to build upon what DARPA had done with the TCP/IP protocols.

Cerf: Actually, Dennis Jennings takes the credit for having made the original decision to move to TCP/IP, and he wasn't in the networking group I don't think.

Kahn: He was in the supercomputer side.

Cerf: Supercomputing thing. Dennis was on leave from the University of Ireland and claims that that was his responsibility.

Kahn: Let me roll back a bit then.

Cerf: Okay.

Kahn: If you really want to go through that change.

Cerf: We also need, by the way, to roll back to 1983 because we skipped an important milestone that we should talk about.

Kahn: Yeah, I'll get back there in a minute. But I want to point out that the National Science Foundation had gotten really interested in networking by the late '70s. They had a few aborted attempts to do something but they were really serious about wanting to do something. They had a meeting at the University of Wisconsin that I attended. Out of that came the basis of a project called the CSNET project, which was really an effort that allowed the research community to get on a commercial net, which in this case was Telenet. We worked the interconnection between Telenet, the ARPANET, and worked out all the arrangements with NSF to allow that. Vint was the chief technical liaison to NSF on that, helped with the engineering of it. I handled the political administrative agreements to enable that to happen. By 1983 or so, NSF had made a decision that they wanted to work more closely with ARPA on supercomputers and we actually had an agreement with them to support that. It didn't actually go as planned because they wanted something to happen instantly and this was in a period where DOD was growing a defense version of that. The net result was that while they made a decision to go with TCP/IP, it was because they were going to build on top of the ARPANET using these new protocols. The eventual decision of NSF was to build their own network called the NSFNET, and that's where Steve Wolf essentially decided to build upon stuff. So NSF had some incremental decisions along the way to do that, and there were many other structures and bodies, networking councils, and the like that had gotten set up along the way. But that was sort of what happened on the NSF side. I think they did great credit to the science and engineering community, the research community in general, but I think their decisions made it possible for the work that we did to become widespread and well known. If we now roll back the clock a bit, 1983 was the time that we actually converted the protocols on the ARPANET from NCP to TCP/IP. The plan for doing this had been worked out a number of years in advance. Vint, I think, had been working with Jon Postel, and Jon had actually written down the plan for what the transition would look like. We announced it to the community, DARPA did, probably at least a year or two before the transition date.

Cerf: Like mid-'81 or thereabouts we told everybody that by January 1 of '83 we were going to convert over. Incidentally, Dan Lynch had a role to play in this too, also at ISI.

Kahn: By the time the transition came about, Vint of course had now gone; he was at MCI. So it fell on me to manage the transition from the one protocol to the other, which was actually instructive in many ways, because it didn't go at all like I expected. Here we had a plan for what was supposed to happen, but it didn't quite happen that way. We kept citing dates and asked people are they ready, but even a week before the transition we were getting requests from many people to delay it. "Are you really serious about doing this transition in a week? We're not ready yet?" "Well, you've had a year or two to get ready. Why not?" "Well, we didn't think you were serious." We then had to figure out what's the strategy for not breaking the net, because not everybody is going to be able to make the transition. We actually ran both protocols in parallel for a period of at least six months, until finally we were able to just sort of shut down the old. The ARPANET then ran with the TCP/IP protocols until it was decommissioned sometime in 1990. I think it was actually in 1990 that the last node was decommissioned. By that time the NSFNET had taken over as the backbone, along with a lot of other networks. And commercial nets that were now beginning to become part of the picture, although they weren't officially allowed to connect to the NSF or any of the government nets. That didn't occur for another three years, until a bill passed that was sponsored by Rick Baucher, a House member from Virginia, was actually passed in early 1993 that enable the NSF to open up the NSFNET to commercial traffic.

Cerf: Actually, this may be a place where we don't have the same recollection, because now we're getting all tangled up in CNRI stuff. Maybe we'll have to come back to this, but I thought that by 1989 when we brought up the MCI Mail Internet connection, that same year there were three commercial networks, and I thought they were allowed to interconnect with NSFNET.

Kahn: It was an agreement to allow electronic mail interchange. In fact, it was pioneered with the MCI mail, and the discussions that you had with the FNC [Federal Networking Council] was what enabled that. Or maybe it was still the FRICC [Federal Research Internet Coordinating Committee] at that time, the predecessor to the Federal Networking Council. Soon thereafter, very soon thereafter, AT&T and Sprint were allowed to do email. There were federated commercial nets through something called the Commercial Internet Exchange, the CIX, that had been set up, and there were a number of NSF-supported nets that were becoming murky participants. For example, there was a network that was set up under NSF aegis in New York, with support from New York State and Nynex, called the NYSERnet. One of the participants in that activity was a fellow named Bill Schrader. Bill had essentially left -- I think he was president of NYSERnet -- he had left to set up a separate organization called Performance Systems International, which later became PSInet. He took the people and he took the technology, so there was really no issue about whether they would continue with them or not. I remember having breakfast with Cas Skrzypczak, who was the chief technology guy for Nynex at the time. He hadn't heard about it. We were at a governor's meeting in New York at the time. They ended up deciding to support it because they weren't trying to ruffle feathers, but I don't think they were very happy about what happened. But the bottom line was that the PSInet activity, which had a lot of the research community on it, was now instantly a commercial net. Suddenly you had the interface to the NSFNET, and you had both allowable and unallowable traffic and no way to sort it out. So even though there was a policy that defined what the use was, the reality was there was no visibility into what was actually happening. So you're right about what was going on. It was a rather gray and murky area. But the bill that made it all legal was the Baucher bill. I think it was put forth in late 1992.

Cerf: That's right.

Kahn: I think it was formally passed in early 1993, and it gave NSF the congressional imprimatur to open up the NSFNET to commercial use.

Cerf: They had an “appropriate use” policy at NSF that would have officially restricted the commercial traffic to flow only on the commercial systems. But as Bob points out, things were so interconnected that you couldn't tell. And there wasn't any good way to mark the packets as “this is a commercial packet” and “this isn't”.

Kahn: Right.

Cerf: Well, unless you want to continue on the Internet path, I was going to suggest to go back now to talk about what happens, particularly as Barry Leiner comes in, in 1983, to help you.

Kahn: Barry Leiner comes into the program. He actually joined DARPA in 1980.

Cerf: In 1980?

Kahn: I believe, yeah. Or maybe '81.

Cerf: You're saying Barry and I overlapped more than I thought.

Kahn: You guys overlapped at least a year.

Cerf: Yeah, that would be '81 to '82, or something.

Kahn: It could have been '81 he came.

Cerf: I remember fairly low overlap. He was working, and we knew about him because of the packet radio program, because he was working on detectability...

Kahn: That's right.

Cerf: …of the signals in the packet radio net. But in any case the point--

Kahn: We can get the exact numbers if you like, but you guys overlapped for a while.

Cerf: Yes that's absolutely true, because there had to be some kind of a handoff. The thing I wanted to come back to, though, was the other activities that occupied your time at DARPA. You've already mentioned the packet radio, and satellite, packet speech, and security, and the general command and control program. But it seems to me that one of the most impactful programs was the VLSI program, because it had such an influence on the number of people who knew how to program or how to design VLSI chips. And also the strategic computing program. So, I thought maybe you should tell us a little bit more about those.

Kahn: The 1970s were a very turbulent time at DARPA. We were fighting the Vietnam War. There was a presidential resignation. There was a lot of alienation in the research community. DARPA saw the loss of many of its key researchers, some were for commercial opportunities, some went to Xerox PARC, some went elsewhere. But the overall budget that was available for research and information technology, which is really what we had responsibility for, had dropped precipitously between the period of roughly 1973-1974 and maybe 1978-1979. During that period I would bet more than two-thirds of the funding, maybe three-quarters, actually disappeared as research funding. The budget didn't go down by that much but a lot of it was shifted off into command and control experiments and other things that the research community wouldn't have viewed as research. They weren't necessarily inappropriate things to invest in, but the research program itself was being very badly gutted at the time. My concern was… In fact, one of the interesting things is people have often asked me why I spent so much time at DARPA. My view was that this was the best hope for the research community at that time. If we walked away from it and let it go, that in the final analysis everybody would regret it. My view was to try and build it back up again, do what I could. Fortunately, I had the support of the director. At that time when I became the deputy director, the head of ARPA was a fellow named George Heilmeier, who was a graduate student along with me at Princeton, so I knew him well. I would say he actually gave me a tremendous amount of support while I was there, and enabled me to do a lot of the things that I did, with his backing. He got replaced by a gentleman named Bob Fossim, who I also got to know very well and had great respect for, and he also gave me tremendous support. Of course, the program then we managed to grow it significantly under Bob Cooper when he was there. He had a much more defined mission and focus which I could really appreciate and Bob and I actually ended up working together to grow that program significantly. But the first thing that I did was in figuring out what to do to at least repair the damage, get back to where we were before, was to create a program in VLSI design inside of DARPA. What I did was to build on some work that Carver Mead had been doing for many years in trying to characterize VLSI foundries. He had come up with a standard way of describing designs that could run on multiple foundries. It allowed you to break the barrier of having know proprietary design rules which made it impossible to teach in the classrooms. Even if you could teach it in the classrooms, there's only one foundry that could implement it, so you weren't able to really do justice to the field. In the middle 1970s, it was also clear that computer science needed to deal with computer architecture issues as part of their life blood, and that without some ability to do VLSI design that field would suffer pretty dramatically. We were also moving into a range where it became clear that sub-micron VLSI design capabilities were about to show up. Carver Mead wrote a wonderful paper, I think with Ivan Sutherland and Tom Everhart, that had a big impact on me in terms of just laying out the semiconductor part of the field. I knew that designing devices that had sub-micron dimensions would be very difficult to do by hand. It's one thing to design a chip that might have 100 elements on it. It's another thing to design a chip that might have a billion elements on it. You really need computing help. Therefore, it was an appropriate thing for the information technology part of the world to deal with, rather than just the material sciences part. So I managed to get support for a new program in VLSI design, and we built upon what Carver had done. Carver and Lynn Conway got together and wrote a textbook on the subject, called "An Introduction to VLSI Communications" -- or "Design", I forget the title -- that was used in a lot of classrooms. We provided support for a lot of the universities to make available money for research in that area and we provided a fabrication design service to do that. Xerox PARC actually prototyped it, with Lynn's help. But Xerox I do not believe wanted to continue to support that. It wasn't part of their long-term business interest or strategy. So they actually approached me -- some of the key people there -- and asked if we could help. Lynn was, I think, part of that process. I ended up moving it to the University of Southern California under Keith Uncapher's aegis. Danny Cohen became the initial program manager for that and they named it MOSIS, which stood for "MOS Implementation System". My one contribution to that was to introduce the "I". Danny wanted to call it M-O-S-E-S and I thought that sounded just a little too--

Cerf: Too biblical.

Kahn: -- biblical, yeah. So that's how MOSIS got created at ISI. They've continued to run it to this day, in support of the entire research community. That was one of the things we did. The research community got to use it, and it made it possible for people to teach VLSI design in classrooms, to get students to be able to implement it and to do really quick turnaround, low cost designs, and it turned out to be a big win. There were a lot of issues associated with, it as to whether that was a good thing to do or not. Maybe the physicists should do the complicated designs and you should use microprocessors for the non-complicated ones. I talked to Bob Noyce who was then at Intel about this at great length, because he was opposing this VISIC program that the Pentagon was funding because he said the really good guys in the commercial world…

Cerf: I was going to suggest that there’s an interesting analogy between what was done in the VLSI program and what was done in the Internet program, in the conceptual sense. In the VLSI program, as I remember it, you were trying to have people design circuits that were fabrication independent. The primary parameters had to do with line width and didn’t have to do with specific physics of particular fab lines. The idea was to be able to design circuits that could be implemented in multiple fabs. It occurred to me that there’s an interesting analogy to the Internet, because we intended that the packets be carried over arbitrary communication systems, that we were insensitive to which kind of communication or transmission system was in use. So there’s an odd parallel there that hadn’t occurred to me until you started talking about a VLSI program.

Kahn: It’s a valid parallel. There are some things that are obviously different and some things that are the same. The big problem was, at that time, the fab sites all had proprietary design rules. It wasn’t so much that we were trying to do fabless ones. I would have been perfectly happy if every one of the fab sites had said you could make the proprietary design rules publicly available. Then you would end up with something at least you could teach in the classroom and they could publish papers on it and say what they did. But it had the added advantage and if you could have ones that would run across a multiple set of process lines, you would have much more powerful design methodology. Carver was the one who came up with the original design methodology. I think Lynn helped with the exposition of it. I don’t know exactly what role they played, because they worked together pretty closely. She’s a very good writer, very good explainer, and it really made that textbook come alive. But I can’t believe he checked every single [fab] line every place in the whole world. So it’s very likely that whatever designs he had, had a good chance of working, maybe in fact did work in all. But we didn’t know that, and he didn’t know that, because he could only say, “Of the ten that I know about in detail, I think this will work on all ten of them.” So those are what we used. We had a small set of fab sites that we were dealing with, and these designs worked across all of them.

Cerf: It seems to me that this is an example of the power of abstraction, in some respects. By getting things at the right level of abstraction, you can generalize a lot more readily than you could otherwise.

Kahn: Given that you’re raising that issue, I want to say that one of the difficulties with deciding what to do, and what level of abstraction to use, had to do with saying with specificity what it was he wanted fabricated. You can imagine at one level the designer could have said, “I want a chip that solves the DNA sequencing problem.” Then it’s left to the designer on the other side who knows about silicon and the physics to figure out how to do that. He’s got to solve the whole problem himself. You’re down one level and the designer could say, “Well, I want ta chip, but here’s the block diagram of it, It’s got a modulator here and a detector here and a rectifier there.” The designer has a clue as to what you want, but he’s still going to do the work. As you go to lower and lower levels, you could have somebody provide a circuit design to implement the circuit. It doesn’t say where on the chip the resistors and transistors and stuff go, but he’s spelling out in some detail exactly what kind of thing he wants implemented. A level down you could imagine a gate diagram. Here’s where I want the gates to go. Not the Bill Gates, but the transistor gates.

Cerf: The T-gates.

Kahn: And one level down from that you can imagine a stick diagram, which literally outlines where the base and the collectors and the individual transistors might go. Some higher level functional spec of a lower level design. What we ended up doing… If you view this as kind of a parabola, at the very bottom of the parabola there’s a point where the two sort of merge, because every one of those other ones leaves enough room for ambiguity that something could happen that makes the design not exactly what you wanted. What I opted for in this program was to pick the bottom of the parabola, the very bottom. If you were to draw a horizontal line and figure out where it hits tangentially, it was a point. That point was where the designer specified the artwork to go on the masks to do the fabrication. So there could be no ambiguity in the communication. Therefore we knew it would work and it would be what the designer wanted and the fab lines could built it.

Cerf: Fascinating.

Kahn: It was a choice with specificity that’s very low level. We have the same problem in the networking world where some people might complain TCP/IP is pretty rigid. It’s got a lot of protocol parameters, but it would be much better if you just had these two computers talk to each other in English. They can figure out what they want to do. Maybe that wouldn’t work quite as well. What would they say about timing issues and fragmentation and how would they figure out addresses? Could that be done? Perhaps. Maybe the future will tell.

Cerf: Of course, the big question is will all implementations inter-work? That’s the important part, at least for Internet. There’s another major program that you started while you will still at ARPA.

Kahn: One other thing, Vint, just to close on this.

Cerf: Oh, sorry. Yeah.

Kahn: I was mentioning about Bob Noyce and the conversations he-- the reason he liked the DARPA VLSI program, because it was focusing on training individuals. That is, it was producing the human capital for the future of their industry and it made a really big difference. There was controversy about it. Was it right to get university people trying to figure out how to do VLSI design or not? But if you focused on where the expertise was, not every one of them would be a world class computer architect. Not every one of them would be a world class physicist. But every one of them who came out of it would have the potential to learn at a later point in time. So he saw that as the big value. I think that's been one of the main contributions. They've done some interesting things. The geometry engine that Jim Clark designed at Stanford became the basis for Silicon Graphics. There are probably a half a dozen other major things that you could point to that made a big difference. When all is put together, the biggest contribution was training people in that field and invigorating the computer science field in the process.

Cerf: I was about to shift into another program that you started sort of late in your time at DARPA, and that's the Strategic Computing programs. Maybe you could tell us a little bit about the origins of that and what motivated it and what the outcomes were.

Kahn: By 1992, we had pretty much gotten the funding level in IPTO back to about where it would have been…

Cerf: You don't mean '92. You probably mean--

Kahn: '72.

Cerf: '72, right.

Kahn: Thank you.

Cerf: Twenty years later.

Kahn: 1972 we had-- no, I mean 1982.

Cerf: '82.

Kahn: In 1982, we had pretty much gotten the funding levels in IPTO back to roughly where they would have been had the impact of the Vietnam War on research [not?] taken place. Not quite as high as it was, but within about 10%. Much of the anguish had now gone and the question really was, where do we go from here? At any point in time you're always looking for a rationale for the next thing. If you trace the history of DARPA's investment in computer science, the motivation for starting that office in the first place was based on a vision that, I guess, was largely attributed to Licklider. Many other people may have had pieces of it, but it was a vision that talked about moving from batch computing to interactive computing, moving from numerical computing to symbolic computing, moving from a localized environment to a networked environment. Much of that by the early '80s had come to pass. We no longer were buying UNIVAC machines and the big mainframes. IBM still sold some. Many of the companies still sold them, but for the most part, a whole new way of doing computing had arisen, both through minicomputers and now with the microcomputer coming of age. What was vision for the next 10 or 20 years? That's what we were struggling with at the time. Also, we knew if there was a vision that the computer science community could benefit from it. We were looking for ways to get more money into R&D, because the agenda of things that needed to be worked on was hardly fulfilled at that point in time. Also, the Reagan defense buildup was happening. So we knew that there was opportunity. The program that I put together with a lot of help from Bob Cooper--in fact, in many ways the decision to push it was Bob's, but I ended up structuring the program--was Strategic Computing. The whole goal there was three or fourfold--big items. Item number one, let's figure out a way to have all this computing technology that we've been working on for the last 20 years at DARPA--and it's now starting to take hold--be made available to industry, so that the DOD can take advantage of it. So goal number one was getting it to DOD. The strategy was you need industry help, because that's what the DOD has to rely on. Furthermore, with all of that buildup, we also recognized that we'd have to get some of the other offices at DARPA involved. Bob Cooper was very interested in funding some early applications. To be quite honest, I was somewhat skeptical of that from a technical point of view, because to build the applications before the technology that it was depending on had been started on, much less created, didn't seem to make good technical sense. But I think Bob's motivation was different. It was more political and strategic. What he was worried about was if you didn't take the other offices and get them thinking about what you're doing, [then] at the time you're ready to really do something, they won't be ready for it. Number two, the politics of taking all the money and building it up in one office versus sharing a little bit of the wealth with the other offices politically probably didn't work. I understand that much better now than I did at the time, because I was arguing technically at that time that it didn't make technical sense. But it wasn't a technical issue. It was totally a political and strategic issue. So we got the program actually launched and going. It was the first really large-scale program in computing that had ever been funded anywhere in the DOD. Which isn't to say that lots of money hadn't gone into computing-related stuff. I don't know when the B-1 bomber program started, but I'm sure there was a lot of computer stuff on that that required a lot of coding by different contractors. But it was mainly contract software development for specific applications, as opposed to opening up a whole brand new area and field and research. It was like a billion dollar program that we got started. I was very pleased at that. Bob and I went around to the research community, at his instigation, to sort of give them a heads-up on it. He kept telling them, what do you think about the program? Because we had framed the program in a way that essentially covered all the research interests that I thought the research community ought to care about. They actually ratified that. There was no project that somebody might want to work on that couldn't be fit into some piece of this overall program. It was into architecture, it was into applications, it was into AI, it was into VLSI. It just sort of covered the waterfront. He wasn't sure, but I think he got a lot of reinforcement from the field that said, "Yeah, the structure of the program is fine." He also pointed out this was going to make a big change. Are you guys ready for it? I think that might have been misinterpreted, because I think they were thinking a lot more money for the universities and therefore, it's going to change the nature of the way the universities are. But they were all sort of saying, "Yeah, we're ready." But what he really meant is that industry is now going to start to play a big role. I think that wasn't exactly what the community was thinking about and expecting, but it turned out to be a pretty good synergy. More money went into research. A lot of that money, I would say the preponderance of it, went into funding industrial groups to build up capabilities in multiprocessor architectures and speech understanding systems and image processing work and the like -- real state-of-the-art applications. But what it did more importantly was it took the IPTO budget and it grew it by leaps and bounds. It went from mainly the $40 million-ish level to maybe the $90 million level with the addition of VLSI and other growth that got built into the program. That took it up from $90 million to roughly the $200 million level, if not a little higher. It was a big push. By the time I left, we were certainly one of the biggest offices, if not the biggest office, at DARPA. Mainly the Strategic Defense Initiative had taken the laser stuff away. But when I started it, we were, I think, probably the smallest office or one of the smallest. It was a big change at DARPA at that point. It also changed DARPA in a fundamental way, because suddenly now you had all the other offices getting interested. Today, if you went into DARPA and you were to ask them, "What money are you putting into computing?" what you will hear back is "Every office is doing some." But the mission offices are focused on applying it to mission purposes and the research offices have carved out generally some very specific goals. Like the information processing office has been working feverishly on large supercomputers and on learning technology -- very specific missions. That's typically what DARPA has done. It's formed its agenda and it's focused on it. Whereas, at the same time, the National Science Foundation budget, this is now circa 2006, has also grown significantly. But their purview has basically been the research community, so they go after the research community to find out what are the issues that are appropriate to focus on. NSF doesn't ask the defense department what to work on and DOD doesn't feel like its job is to go ask the universities what should it work on. They all have their own mission and their own constituency, their own set of problems. That's where we are today.

Cerf: Just out of sheer curiosity, Bob, did the Strategic Computing initiative have any impetus coming from what the Japanese call the Fifth Generation Computing program? As I recall, there was a lot of visibility associated with that, but I don't know if it motivated any of the strategic initiative.

Kahn: That's a great question. When I took over the office, one of the things I was trying to push was more money into informational infrastructure kinds of things. I wasn't getting much traction. I wanted to build a knowledge-base for science and technology and I wanted to work on a variety of things that really didn't happen. We were fortunate in getting the VLSI program going, but I was trying to do more back then. It wasn't really going very effectively in getting more money. In fact, it was a big enough achievement to get where we were to get back in the norm. Somewhere along the line, the Japanese started their Fifth Generation program with a stated goal of overtaking the US and moving ahead. They put all their marbles on a language called Prolog, which you may remember from back then hasn't really made a mark on the external world to any great extent. They were dealing with pattern information processing systems and fuzzy logic and stuff like that. There was a lot of concern that we needed to do something to not be overtaken by the Japanese. The head of defense research was a gentleman named Dick Delouer [ph?], who I think believed very strongly that we should be combating that. That's why I believe I did not have the discussions with him. When Bob Cooper and I talked about this, he said, "I don't want to worry about the Japanese. I want to worry about what's right for defense. What we do should make a big difference for defense." All of the stated goals inside of DARPA, no matter what you thought might or might not be the matter of reality, was let's do something that makes a big difference for defense. That was the internal DARPA view and position. Like many other dual-use technologies, you could have asked about the Internet. Was the Internet to create this worldwide global goal? Maybe some people had that in the back of their heads, but if you asked DARPA, they would have given you a different point of view. In fact, if you really go back to the early days of DARPA and looked at the requirements for getting DARPA programs going, it was almost always to identify a specific defense need, to find a person there who would be a partner with you, maybe get them to fund part of your work, work out a transition plan. If you think about the Internet as an example of a project, at the time we started it, the rest of the military didn't have any timesharing computers. They had a few batch machines that were used for various missions like accounting where interactions were not de rigueur. As a result of having no interactive machines, they had no interactive computer networks. Imagine trying to sell somebody a program to link together the networks they don't have for the computers that they don't have, and a transition plan, and funding. It was basically not an argument that had any traction. It didn't fit into the mold. Fortunately, we were able to pursue that in the guise of the packet radio program, because it was a very specific need. So we were able to develop the early TCP/IP protocol, spec that way and get the initial implementations and demonstrate it. At that point, it got a life of its own. It's a tough sell.

<break in tape>

Cerf: You were about to say something else about the Strategic Computing initiative.

Kahn: The one other thing about the Strategic Computing initiative which is probably not understood that well by others was the infrastructural component of that program. We had argued that in order to enable the research community to do a lot of their tests, they needed high performance work stations. We managed to work, both through a combination of different schedules and purchasing agreements, the ability to get a lot of workstations to the research community. At that stage, there were very few-- the IBM PC had just come out. It wasn't powerful enough to do much of what we're talking about. Apple wasn't on the horizon yet. Getting the hundreds of machines out into the research community and start to see the community that could make use of this was an important thing that Strategic Computing enabled. Furthermore, it enabled the initial Internet connections to happen at most of these places. We gave the institutions money to buy the local area net connections, get the machines, get hookups. This was a time when some of the commercial stuff was if not fully commercial, a glint in the eye. I'm not sure when, for example, the early nets got formed, but it was an early California network that interestingly was named Cerfnet. Not, I believe after you, the California Educational Research Federation.

Cerf: Foundation, right.

Kahn: Foundation. There was a CALnet of some sort. There were a variety of research nets being formed all around and people had to get connected to it. They generally had fees associated with them. People had to pay for that. The program enabled all that. A lot of the infrastructure for the early Internet happened through Strategic Computing, which caused it to get funded and in place. Very few people realize that, but without that program, we would not have seen the uptake of the Internet after January 1st of '83. There wouldn't have been enough machines to use it. There wouldn't have been enough networks to connect to it. It would not have been properly seeded. So I think that's a little known aspect of it. The other thing was: an adjunct of the VLSI program was the decision to fund the University of California at Berkeley to build a virtual memory version of the Unix system. The motivations of that were twofold. One, AT&T had obviously developed Unix and was making it available through licensing arrangements through, I think maybe Western Electric. The university could get licenses to use the existing Unix software, but they couldn't get virtual memory systems out of AT&T. So we funded Bill Joy, who was then an employee at Berkeley and who later became one of the founders along with Andy Bechtolsheim of Sun Microsystems, to actually build a virtual memory version of Unix. The motivation was number one, for image processing. We had a very large image understanding program that was ongoing. This was before Strategic Computing. These image databases were large. The existing software just couldn't deal with it properly. But more importantly, the VLSI program also needed it. A lot of these VLSI designs were really huge and went beyond the bounds of what the typical system could deal with. We had also funded BBN to build some of the early protocols, the TCP/IP protocols for the Unix platform. The decision was made to have Berkeley put the BBN Unix software package on the Berkeley distribution tape. We just essentially told them, "Please do that," because we were funding both and they needed this protocol. Bill Joy's reaction to that posed a problem for us, because he was not comfortable putting somebody else's software on that. But his system didn't do it. What he ended up doing was sort of, I presume, just looking through the code and re-implementing it into the guts of the Unix system and parts in the kernel. I don't know how he did the actual implementation. He made it work officially in the local area network environment that he was working with. Previously, many people had looked at this. Dave Clark was one of the first to pioneer putting TCP/IP on local workstations and getting it to work in local area nets and making local nets work. But this was a very specific project. Bill Joy actually did that code again. It worked fine in the Berkeley environment. But he made some efficiency optimizations at the time. He took the checksums, and instead of putting them at the back he put them in the front... I don't know all the details. But the bottom line is his implementation, which worked fine in that environment, wasn't compatible with either the specs or whatever other people were doing. It took a while to sort that out. My recollection was at least a year and maybe more. Jon Postel was involved and perhaps lots of others. That development turned out to be extremely important, because it's one thing to go around to people and say, "Please download the TCP/IP code and oh, by the way, please integrate it into your operating system and oh, by the way, please integrate it into all your applications," which at the minimum was a job for a wizard and hard work and you have to then do it for every other application. Whereas, by getting it through a single package from a commercial vendor, you could go buy their workstation or their operating system and it came packaged for free. So it became the carrier of that part of the technology base. A big difference. Now the gateway stuff, you still had to configure your own: get a piece code from BBN, get an interface from here, buy a LSI-11 from Digital [Equipment Corporation] or however you're going to do it. But shortly thereafter, probably in 1985 or 6, we saw a number of companies emerge that became powers in the whole network field, including certainly Cisco, which was a spin-off of Stanford. There were some companies up in the Boston area that worked closely with MIT.

Cerf: But it would have been Proteon who was the one who first came out, as I recall. And 3Com and Bridge Communications had some role to play here, although they were often operating below the level of IP. But eventually at least 3Com moved into gateway and router technology, too.

Hendrie: And then there was Wellfleet in the Boston area.

Kahn: Wellfleet. Well, I don't know whether they were doing the same thing. Bob Metcalfe related to me how unhappy he was, because Vint had managed to convince Bob to go build a TCP/IP package and sell it, and Bob had done that. Meanwhile, unbeknownst to Vint, Duane Adams, who was working for me, had gotten the Berkeley guys involved in it. It was the net effect of this whole effort. Suddenly, Bob saw the investment that he made being undercut by this free package coming out of Berkeley. He thought Vint was responsible for it.

Cerf: So yeah, he was very angry with me about that and used to periodically poke me in the ribs and say, "You cost me a bunch of money." Of course, later he admitted that he'd actually made a million bucks off of his own implementation, which was called Unit or something like that.

Kahn: You cost him all the additional money he didn't make.

Cerf: I'm sorry?

Kahn: You cost him all the additional money he didn't make. But meanwhile, I don't think he really appreciated the fact that you had nothing to do with the Berkeley decision to go forth. It was all handled by the grant.

Cerf: By the time that was happening, I was at MCI. In any case, the workstations, the Berkeley release, and local area networking, Ethernet especially, contributed enormously to the rapid spread of Internet technology, at least in the research community, because they could buy this stuff off the shelf.

Kahn: And eventually everywhere else.

Cerf: Yeah. I'd like to move to the next major phase in your career, Bob. You left DARPA at the beginning of '85? Or end?

Kahn: September 1985.

Cerf: So the end of '85 or near the end, and started a company called the Corporation for National Research Initiatives. What were you thinking about doing? What was motivating that particular path?

Kahn: I left in September. I actually stayed on as a consultant to DARPA. For a while there, I was trying to help Saul Amarel take on the job. So I actually was sort of full-time in DARPA for a while after that.

Cerf: Saul would become the director of the Information Processing Techniques Office?

Kahn: Right. Saul had been a professor at Rutgers. He came in to replace me running the office. I actually was helping him. Just to pin the dates down, I filed the papers to incorporate CNRI, as it's now known, on January 29, 1986. The reason I remember that date is because, just like I remember the date of my mother's first attack because it was Roosevelt's-- the day that he had passed away. This was the date that the Challenger exploded.

Cerf: Oh geez.

Kahn: As we were going through the papers, we heard the announcement of it. So I'll never forget that.

Cerf: So you've just had the 20th anniversary of CNRI this January, is that right?

Kahn: Well, it depends on what you think is the start date.

Cerf: The incorporation date is at least one milestone.

Kahn: That would have been there, and maybe we should still do something, although we haven't yet. The first operations didn't occur until May of '06 [sic]. At that point, it had one employee. Not doing much, I might add. In June of 1986, we got employee number two, namely you, Vint Cerf. I think we got employees three and four a couple of months later. During that period, CNRI was born with the goal of focusing on fostering R&D for a national information infrastructure. I need to say a little bit about my motivation for why I did that. I would have preferred, quite frankly, to do that out of DARPA, because we had a large budget, we had the imprimatur of the government. It's much harder to make infrastructure things happen in the country when you are one of many in the private sector. It's one thing to be respected. It's one thing to have a little bit of capital. It's one thing to believe you can do that, but to really do it effectively, it needs a sustained push over a long time. Government is able to do that; That's how many of the infrastructures that we know were put in place. But the Reagan administration's view was that this was industrial policy, that if we're going to create infrastructure, the private sector should do it. The arguments made perfect sense, except that it was very hard to do, because the private sector didn't see it as their obligation. I had discussions with many people, many CEOs about that. They all thought it was a great idea. How many of them really understood it, I don't know. I know that a number of people told me ten years later, without naming names, that when I first told them what I wanted to work on, which was a national information infrastructure, they thought this was the equivalent of sort of mining the clouds for precious materials. It didn't seem like there was a there there. But, of course, people weren't thinking about Internet. They didn't know about what infrastructure really was. Even today, if you said to somebody, "Give me a good definition of infrastructure," they would be hard pressed. Even the notion of what a distributed system is, is somewhat elusive. I remember having discussions about even that topic with various people who were very good technically. I said, in DARPA days, I wanted to set up a program on distributed systems. They said, "Well, we already have them." I said, "Can you give me one example?" They said, "Yeah, well we have these sensor systems. We have sensors all over the ocean. We bring them back to this machine and they compute here. We have these seismic arrays and we bring all the signals back and we compute there." I said, "Yeah, but that's not a distributed system." "Yes it is. These sensors are all distributed." So then you get into the issue of what does it mean to be distributed and what does it mean for them to have to work together? That's only the starting point for getting to infrastructure. There are many people who say, "Well, aren't memory chips infrastructure?" Well, I would say they're infrastuff. They're not structure, they're--

Cerf: A new term in our vocabulary.

Kahn: It's like oil is infrastuff for the road system, the pipeline. But it's very hard to describe that. It was especially hard back then, because there were no good examples that they could typically relate to. The ARPANET was a research project. There probably wouldn't be more of them. So it was a hard message to get across. Most people didn't understand it and therefore, there wasn't a lot of support that I was getting from anyplace for this as an idea. But my thought was this was sufficiently important. I guess hadn't internalized how hard the sell was going to be. I might have been still in the mode of saying, "They're just not getting it. I should try harder." But I hadn't yet figured out how to make the pitch so that people would latch on. I went to a variety of organizations trying to convince them. They all said it was a great idea. They really did. I had no feedback that said, "This is a really bad idea. Don't do it." But they weren't necessarily willing to pony up any funds. We raised enough initial funding to get the place going. People that were particularly helpful were IBM, Ralph Gomery committed to help. He said, "You're going to have trouble with other companies." Digital Equipment was very helpful. Sam Fuller, and I guess Gordon Bell, must have been involved. Xerox, in terms of Bill Spencer and maybe some of the senior people were very helpful. MCI was one of the early funders. I would single out Dick Liebhaber there, who decided to make an investment in us. I think it was more an investment, in that case, in what you were trying to make happen that you believed in. But they were a strong supporter for a long period of time. Bellcore, mainly through a number of people who helped to do that. Dave Sincosky [ph?] for one. I'm sure George Heilmeyer and Bob Lucky were helpful. Bill Newport, who was then a senior VP at Bell Atlantic, and Bell Atlantic itself, which is now of course Verizon. Those are some of the initial folks that helped us get going with enough support that we could really launch the CNRI in at least an organizational way that was effective. But the problem I ran into was it was too hard to raise the kind of money to run a really big national program. Could I have done it at DARPA? I think, given the political context, the answer was no, because it was sort of ruled out. But we were also facing a series of other things that were emblematic of the times. First of all, the Reagan defense buildup had just about played out. I sort of left at the top of the build-out, so there was really no more new money going into DOD at that point in time. I knew that we would be in defensive mode for the programs that we had created. Since the other offices at DARPA had not gotten a big part of it, they would be inclined to go after a fixed pot for more share for their new ideas, too. This would be totally digging in your heels and defending things, but we had some more issues on the table. Gramm-Rudman was on the horizon and it eventually came back and bit DARPA pretty heavily. Number three, there were an increasing number of “black programs” that were being started. Things that I didn’t know about, couldn’t talk about if I did.

Cerf: By this you mean classified programs?

Kahn: Classified programs of one sort or another that I don’t even know what they were, because I was never privy to them or read into them. But there were many of them for which funds needed to be found, if they weren’t separately appropriated. Finally, we were starting at a path that we hadn’t seen in our field before, where there were earmarks for things. You colloquially call that “pork”. We were faced with some of that as well coming down the pike. I just knew that the difficulties of maintaining this were going to be totally in defense mode, no new things and that just wasn’t me. Not that I couldn’t have done it, but I had all these ideas percolating through my head of things I wanted to do. That was another motivation for leaving to start CNRI. I couldn’t do what I wanted to do and I couldn’t do anything new, so [it was] time for new leadership at DARPA, I felt. Besides which, I had been there 13 years at that time.

Cerf: That’s a long time.

Kahn: Which, by today’s standard, is six sigmas past the normal lifetime. It was the right time. I wasn’t drained. I wasn’t out of ideas. I was motivated to do it, but the timing was right for me to leave. So I set up CNRI and I think we’ve made major contributions over the course of the 20 years that we’ve been in existence. Some of them you were very helpful in making happen. I’d be happy to recite them, because we’re talking about the organization. Some of which I was directly involved with. I think, on balance, it was a major assistance to the country on a number of different areas. Today we have major efforts in microsystems and nanotechnology support to the whole country, sort of like a Moses for micromechanical design, which is a very interesting problem in its own right. Unlike VLSI, where you can announce a CMOS process and a lot of people could design for it, every design in MEMS and nanotechnology sort of requires its own process. It’s like programming your foundry like you program a computer, putting that in place and then managing the designs through it. We’ve been doing that now for about seven years and it’s growing fairly rapidly. We’re trying to make it self-sufficient and maybe turn it into its own little stand-alone activity. It’s been very successful and I think it’s tremendously helpful to DARPA, but [also to] a lot of other people who don’t have government support. As you well know, we provided the major support for the Internet standards process. This was an area that you and I talked about back in the early 1980s. NSF needed the help. They came to us and asked us. A few times we declined, because we thought it was too administrative and not research enough. Finally we acceded and agreed to do it. I think we did a very good job. For a long time we did it under a cooperative agreement with the National Science Foundation, through 1997. You were the PI [Principal Investigator] on that for a number of years. We had I think several other people followed it. In 1997 we put that activity in a separate subsidiary organization that I asked Al Vezza from MIT to come down and run when he left being the associate head of the Lab for Computer Science. That continued to provide the administrative support until literally last December [2005], when we sold that subsidiary to a local company called NeuStar [Secretariat Services, LLC] so that they could continue on. So we’re no longer in the process of administering that. The locus of support for that is now really coming from the entire community. We took all the intellectual property that we held and we put it in a trust. It was an idea that some people didn’t describe it in quite that way, but Patrice [Lyons] actually took that notion, formed it in a way that was workable, and made that happen. So we put all the intellectual property that we owned in the trust. We got the Internet Society to put intellectual property that they owned in the trust, and the IETF community as a whole now has the leadership for it and we're completely out of it at this point. In some sense, we kind of went soup to nuts on that and the thing continues with a life of its own at this point. We'll have to see how that plays out.

Cerf: On the IETF stuff: if I remember right, Phil Gross was the first Executive Director that we hired to run the Internet Engineering Task Force effort. Then when he departed to be vice-president of ANS, Advanced Networking Systems, we hired a guy named Steve Coya. So CNRI has really had a very long history in the IETF area.

Kahn: Well, let me just dot some "i"s and cross some "t"s there. The IETF, as I said earlier, really started as one of ten task forces to sort of maintain the punch list, the original Internet Activities Board. As that grew in time, it took on a larger set of responsibilities. Phil was not the first head of that as a titular group. I think, in fact, maybe Mike Corrigan came in...

Cerf: Mike Corrigan was the first IETF chair.

Kahn: ...for awhile and then Phil sort of took over that activity sometime, I would guess, late '85, early '86.

Cerf: He actually supported… He was at MITRE and I think Corrigan had had him under contract and so Phil was supporting Mike when Mike was formally the chairman of the IETF.

Kahn: That's right. And Phil then was recruited to… He was sort of doing it without any direct support, because all the people that were participating in that were either funded by corporate support to the extent there were corporations involved back then, but many of them were just supported by DARPA on legacy funding and contracts that were about to expire. So this did not have any long life continuity. Of course, NSF decided to come up and pick up the support for that, and they ended up funding CNRI to do that. So that kept the activity alive and productive going forward. So by the time Phil had taken over, he was still at MITRE and he was brought in to CNRI to essentially run that day to day under this cooperative agreement, but Phil was the chair of the IETF. Phil was, therefore, handling, under your oversight at that time, both the technical and the administrative responsibilities for the organization. It still had a lot of technical folks involved so he wasn't doing everything, and in fact he wasn't doing everything administrative either, because we had support there. But the locus of activity was at CNRI for both of them. When Phil left, he was replaced by Steve Coya, who had come from MCI also. I think you knew him before well, and ended up recruiting him. Steve became the first executive director.

Cerf: That's right.

Kahn: There was no such title before, because Phil being the chair of the IETF, what was Steve's role? He wasn't going to be the chair so, as an Executive Director, he managed the administrative role. So at that point, which was roughly 1991, CNRI basically allowed the technical aspects of that to go with Phil, and I recall the discussions that you and I had as to whether that was a good thing to do, whether we should try and retain it. Because our agreement with NSF, basically, you know, empowered CNRI to create a secretariat for the IETF and the IAB and, therefore we could have argued, relative to that contract, that we should retain it at CNRI. And I think you were pretty persuasive in saying, "Let it go." I'm not sure whether I would have made that decision on my own, but I probably would have come out there in the final analysis, because there was no good alternative. Phil was liked, he was doing a good job, he wanted to continue it. Al Weiss, who was the head of ANS, was willing to let him do that. It wasn't actually clear what a better solution would be, so I probably would have come out in the same place. But, anyway, that's what happened. So Steve came in as Executive Director and he ran that operation as the Executive Director until he left, basically, which was in, I think, maybe early 2005 or late 2004. He was there for quite awhile.

Cerf: If we were to go back to the first year of CNRI, if I remember right, one of the things that was a highlight from my point of view was an interview, or a hearing actually -- it was a senate hearing -- that then Senator Al Gore held, and you were one of the witnesses at that hearing. My recollection is that that was the first instance of the term "information infrastructure" showing up at least in the recorded vocabulary of the annals of the congress. Do you remember that particular hearing?

Kahn: Yeah, I do. It had to do with the formation of the High Performance Computing program, which I used to think of as a follow on to the Strategic Computing but it was sort of broader in the government, so it wasn't just the DARPA program, it had many, many participants, including DARPA. As they proposed it, it was an intended response to a charge that Al Gore had made to the administration to worry about linking the NSF supercomputer centers with fiber optic links. The administration had chosen to respond and they put together a plan. In fact, I remember the very first plan that they put together really started out under the auspices of FCCSET, I think the Federal Coordinating Council on Science, Engineering and Technology. But there was a plan that was put forth, and instead of just focusing on what Gore had asked for it came up with a much broader agenda. It talked about not only building hardware, it talked about software systems. It talked about even human factors and some of those issues. So he held the discussion to open it up: I asked for this and you came back with this and what's the rationale and motivation? I believe that was the session that I participated in, and that was where the topic came up. I remember raising the subject and I remember at the time he asked me if I could give him two or three examples, and I think I did at the time. Probably one of them was digital libraries, one of them was a national medical exchange, so you could find medical information. I don't remember all the others but when the thing ended, he said, "Well, just for the record, could you give me your..." I think I said to him I developed a whole list for the National Academy back then. He asked me if I could give him the whole list and I said, "I'd be happy to" and I later did. The thing he was referring to in the National Academy was actually interesting, because the kind of things that we do had been misconstrued by so many people over a long period of time who just didn't know what they were, or didn't know how to bin them, what bucket or what box to put them in. So one of the things that the Academy did was they reconfigured, restarted something called the Computer Science and, I think it was called, Technology Board.

Cerf: Are you sure it wasn't “Telecommunications”?

Kahn: It's now “Telecommunications”...

Cerf: Oh, okay.

Kahn: ...but I think it was originally ["Technology"] The same acronym. Like IAB continues, CSTB continues but the words mean something different. I was one of the original members of the reconstituted board chaired by Joe Traub, who has been until recently, maybe he still is, chairing the current CSTB, although there were other chairs in between. One of the first things that we did, and it was a fairly august group of people, was say: what are the things that this board should be worried about? This was 1986, I believe, and I had already started CNRI, so we were now known to people there as focusing on the National Information Infrastructure. And we were the only ones who were doing it, because it was sort of nonprofit in the national interest. The topic came up and was put on the board and it was rejected; not necessarily by the people in the room but the view was that the National Academy should not be focusing on something that was the proprietary interest of any one company. I said, "Well, National Information Infrastructure is like working on clean air, you know, it's something that's going to be a benefit to everybody, it's not a product or something we're selling, or unique to us." Well, a number of years later this same topic came up again and it got on the agenda unanimously, and nobody could ever remember why it didn't get on originally because they were worried about why isn't it on it at that point in time. So sometimes that happens. I'll give you one other example. It sounds a little embarrassing in the retrospect of hindsight but it was the reality that we were dealing with then. Nobody knew what infrastructure was. So we had this infrastructure thing on the board and we were trying to figure out, well, what should we focus on doing? Well, we needed to get support for this effort from the board under which the CSTB was located, I think it was the physical sciences board. And so, Irving Wladawsky-Berger, who was then a member of it, and I were charged with going and giving the briefing. And the first question that came out was, "Well, what do you mean by infrastructure exactly? What is it?" There weren't, to my recollection, any computer science folks on the board, or if there were they were sort of people that worked in multiple areas of which computing was one. We kind of looked at each other, how do we want to try and explain this in concrete terms? So we said, why don't we try email? We said, email is a good example. People knew about it, people were using email, and there are national issues and policy issues having to do with getting email out into the country and interoperable systems and dovetailing everything. The reaction at that time was bewildering to both of us, because they said, "Well, how does this help the scientific community?" We said, "Well, the physicists can you use email, and the chemists can use email" and they said, "Fine. If you want to work on email for physicists, that's fine. If you want to work on email for chemists, that's fine. But you shouldn't work on email in general." And we said, "It's the same problem. Physicists are real people in their other lives, physicists often talk to non-physicists, everything is interoperable and that's what makes infrastructure good. You can talk about the highway system for doctors, but it's the same highway system that the taxi cabs and the truckers and then everybody else uses so you can talk about that-- it makes no sense to do that." Well, we didn't make any traction and it disappeared. But a few years later the same question was put forth and at this point there were no objections because it seemed logical. So I think it was part of the education process of people, to figure out what we were talking about, what it means. We've seen that oftentimes. It's one of my bigger frustrations. I've learned to deal with it, that sometimes a vision that you can see is hard to convey to somebody else because I lack the vocabulary to map it to what's in their head. It's like trying to teach quantum mechanics to somebody who is observing the physical world, or maybe even is talented technically. It takes some getting used to, because some of the notions [are] not visible, counter culture, and after a while you start to think in those terms.

Cerf: They're even counterintuitive. One thing I would like to achieve in this interview, Bob, is an appreciation for the scope of work that was done at CNRI and is still going on. The things that come to my mind very quickly are things like digital libraries, the gigabit networking program, collaboratories, the interconnection of MCI mail and the internet. You mentioned earlier the microelectronic mechanical systems, the knowledge robot ideas that came out of the digital library work, the digital repositories, digital object identifiers and the systems that go along with that. I'm not sure how to try to capture a lot of that, but the first question is: what did I leave out? And the second, of all those various things, are there some that you particularly would like to make sure we capture in this interview?

Kahn: Well, you'll have to have some but I don't know how important it is to be exhaustive here. Let me just assert that we've covered our role in helping the internet to become real, and especially our role in the internet standards process. I'm not going to go back to the microelectronics-- the MEMs, microelectronic mechanical work and the nanotechnology work because we could spend hours on that subject alone: how that works and how we can help people in that area, sort of like in a Moses system. Let me start with the gigabit test bit project because that was an early one that we got started. As part of that program, which was funded by the National Science Foundation, who we funded… In fact, when I first wrote the proposal for doing this, it was in a very abbreviated form and it was intended for NSF to do the work on their own because I was trying to assist Erich Bloch in understanding how NSF itself could help the community. If they're going to get into networking, make a big step, don't make a little one. I suggested that they go down that path because I thought the technology was able to be harvested, it wasn't quite out there yet and, rather than focus on something more near term, take a big DARPA-like step. I think the sense at NSF was that they don't know how to do that exactly, that a lot of the technology is in the research labs. You can't buy it. They don't have the staff to manage it, like DARPA might manage a project. That what they could do is procure a network. So they made the decision, during this whole process, that they would, in fact, put in place a one and a half megabit net, which became the NSFnet, and get behind that, which they could procure and they could manage the procurement. Somehow they left it that they would have a separate research program to deal with gigabit nets, but it wouldn't be an operational build of any sort. I recall that Dave Farber was actually the first one that approached me on that topic. I mean, I had been interacting with Erich in any event, because I knew...

Cerf: This is Erich Bloch?

Kahn: Erich Bloch. I knew him long before.

Cerf: Erich was, at the time, the head of NSF.

Kahn: He was the director of NSF.

Cerf: Right.

Kahn: He had just come in, in 1984, to head NSF and I had been interacting with him when I was at DARPA to just try and help him. Dave came and said, "You know, they're going to go build a one and a half megabit net. I think that CNRI would be a good place to help them make a gigabit capability happen as an R&D thing." I remember looking it and saying, "Well, you know, if you're going to go do an NSF, then maybe you want to build the one and a half megabit one. That might be interesting." As I recall, you, at the same time, were probably having discussions with DARPA about doing a one and a half megabit thing for DARPA, which decided that "been there, done that", they're going to move on to the next thing. So they didn't bite. So NSF funded Merit to create the NSFnet and I put together the proposal, with help from many people in the community because it was representing different points of view as to how to do this, that became the proposal. The actual way this worked was quite interesting because I wanted to do a solicitation. I didn't know what all the good ideas were. This was not intended to be a “Bob Kahn figure it all out and go implement it” project. I wanted to put out a call for white papers and find out what the ideas were, but NSF didn't have a mechanism to allow that directly. They said, "Can't do that. You've got to propose something specifically, if you want, and we'll consider it." I said, "Fine". I wrote a proposal and I got significant input from Dave Farber in the process and I got letters from people who said they were willing to participate with us. I got it from AT&T, and I got it from IBM, and I got it-- I don't know who they all were at this point any more, but a lot of the big players said they were willing to work with us. They'd supply people and technology and the like, and I got something like eight or ten universities to say the same thing. People like John Turner from Washington University and I think even somebody from Penn, although I don't know if it was Dave himself, Dave Clark up at MIT, there were a number of the relevant players in the field at the time or many of them. And I was willing to proceed on that path or I wouldn't have sent in the proposal, but I kept thinking, maybe there are better ideas out there that I don't know about. This is just a selection that I made based on what I now know, because I was forced to make a selection. So I sent the proposal in to NSF and it came back basically saying, "Well, we like the idea, but we don't understand why you picked that set of participants." <laughter> This was couched… I mean, they knew what was going on here and they said, "We'll give you some preliminary funding to go find out what the best ideas are in the country." I think they funded a half a million of what eventually was a $20 million grant, and we put out a call for white papers. We ended up being enabled to do what I wanted to do in the future, but because they directed it rather than I proposed it, it was apparently allowed. Although I understand there was some debate within NSF as to whether this was really an appropriate path to follow at all, and that maybe NSF should really be doing all of this and the like. The way I did that was kind of interesting because I inverted the process. Normally the government will have an advisory panel consisting of people from the universities and research labs that know about the area, give them the technical evaluations and the government makes the decisions. I flipped it around. I had an advisory panel made up entirely of government people, some from NSF, some from DARPA, some from DOE, some from NASA, some from NIH, right across the board I think one or two private sector folks were involved, by choice of the government They actually did the vetting of all the proposals. They broke them into three categories: the A category: must fund; the B category: okay to fund if you need them and have extra money; and the C ones, which were: these are really bad, don't even consider them further. We put the program together mostly from the A and some of the Bs. But I then had to go out and solicit carrier support and I remember having a discussion with Erich Bloch, who called me, and he said, "Bob, we're going to fund this". A $20 million grant from NSF was a large grant. I think it was $16 million initially, including research, the actual proposal was more. They took out the research piece and just left the testbed part in. They basically augmented it later with some more money, so that's how it ended up at $20 million. But that was probably worth four times as much if you went back 20 years. Erich and I had a discussion on the phone. He said, "Bob, we've decided we're going to fund this," and he said, "It's unusual because, normally, we ask for all the commitments from all the parties ahead of time and you don't have them." And I said, "I can't get them if I don't get the government money" and we proposed five test beds and I said to him, "Quite frankly, I don't know that I can create all five of them, but I think I have a good chance to get at least three, and I'd be very disappointed if I don't get two." He said, "I just want you to agree that you won't spend all the money and accomplish nothing." <laughter> So he was kind of worried about how it would look over there, and I said, "You have my agreement." In fact, I was very miserly on how we spent the money until we essentially locked in all the agreements. We had discussions with lots of players, and in the final analysis, there were more than 40 organizations that were involved, a lot of carriers, some computer companies, some universities, labs. I think we funded about 10 or 12 of those, and the rest were voluntary contributors. We built five test beds and they all happened. We got 100% and, out of that program some very interesting things happened. One was the first deployment of an all-optical link between Wisconsin and Illinois. AT&T built that one with the erbium-doped amplifiers as the first deployed operational link anywhere, I think, in the country. No electronic repeating, I mean. We had the first deployed ATM switch in a carrier's central office. It may have been the first OC-12 switch; I'm not sure whether there was an OC-3 one before. This is the switch that worked in 622 megabits. That was deployed in one of the carrier offices in North Carolina as part of a testbed there. We funded Illinois to build a point-and-click browser. They were building a database of their simulations and they kind of wanted to work on digital libraries as part of that. You may recall that CNRI actually built a point-and-click browser for the internet in the late 1980s, not very well known, but the particular browser that we built was part of an effort with the National Library of Medicine.

Cerf: I remember that.

Kahn: When we demonstrated it, the reaction that we got was that doctors would not rely on point-and-click and using mouses and we needed to have a forms-based system and so we ended up basically rebuilding that into a forms-based system for access to Medline. Tim Berners-Lee says he built an early point-and-click browser for the web on a NeXT machine out at CERN. I have a feeling we might have been earlier than him, but neither of those actually made it out because the one we built was sort of unique to the National Library of Medicine and it involved generating KnowBots behind the scenes and automatically figuring out how to translate what ended up being forms-based commands into database commands of different database systems. But we funded Illinois to build a point/click interface to their simulation stuff that was stored elsewhere and I have Charlie Catlett, who was then the manager of that project, at NCSA basically saying that the one thing he regrets -- this was like 1993 or maybe 1994, at a major public conference, I still have it on videotape -- that they did not decide to hook that browser up to the web protocols. He says they debated it and they made a conscious decision not to, because they would have had to reprogram all their simulations in HTML and they...

Cerf: Oh, my.

Kahn: ...thought that was a waste of time.

Cerf: Oh!

Kahn: So they didn't do it, and these guys down the street picked up on that and I don't know the details of whether they used their code for experimentation or just did a clean slate implementation, but they built what became the Mosaic browser. So I think you could argue that that program was the genesis of the Mosaic browser at some point in time. Of course NSF was involved in funding both so NSF can take credit either way. But the Mosaic browser really changed the landscape and it's what turned the web from something that was just out there and not really making much progress into something that your grandmother could use, because you didn't have to learn procedures and you didn't have to wire things up. You could now punch a button and get something. It was a major social contribution to make accessible something that you really needed to be a computer scientist before, or have the same mindset, to be able to use. So I think that was an interesting thing that happened out of that. We had a number of other things. We had 12 -- some large number -- of supercomputers in one of the testbeds that linked JPL, Cal Tech, the San Diego Supercomputer Center and Los Alamos all hooked up doing really interesting things like evaluating hydrogen fluoride laser dynamics, like how does the lasing occur in high-powered lasers? Weather simulations. Really a major effort. Of course, one of the testbeds was totally focusing on networking research. We were trying fast switches of different kinds involving IBM and Bellcore and MIT and Penn. All in all, I think that was a major project. I think it put high speed networking on the map, and I think it wouldn't have happened without some major commitments by carriers to help out. MCI was a major contributor to two of the testbeds, the one up in the northeast corridor called, I think we called that one AURORA.

Cerf: AURORA. Yeah, I actually remember these.

Kahn: And the one out...

Cerf: That included the Bellcore people...

Kahn: Bellcore and IBM, MIT and Penn. The one out in the southeast, which we called CASA.

Cerf: You had another one called BLANCA, as I recall.

Kahn: Yeah, that was some combination of Bell Labs, NCSA, University of California; it just sort of worked out. And CASA was from, like, Cal Tech and San Diego, so the names just sort of showed up. It looks like they were designed to be Spanish in origin, but it didn't quite work that way. I think that really did put high speed networking on the map. Unfortunately, having now built those nets in the mid-1990s, there was no program to fund research on high speed nets. It didn't materialize and the bottom line was, nobody took the initiative to find that funding at that time. The carriers said, "Well, why should we keep these nets up if there's nobody funded to make use of them?" We ended up dismantling the whole thing.

Cerf: Yeah.

Kahn: A few years later, they set up something called the Next Generation Internet Program, or something like that, and they suddenly put up money for building it, but now we had lost the commitment of the carriers to actually make that technology available. So we are here today, in 2006, with less capacity, at least in some places, than we had more than a decade ago within the research community. But people now appreciate that high speed nets have some value, including the industry, so maybe we lost a decade but my guess is we'll pick up and continue on our path and progress. So the gigabit network project, I think, was a really important one. One of the things that I don't really take a lot of credit for in terms of vision at CNRI but was true of an individual was we helped take a language that had been incubated at CWI in the Netherlands and grow it into a major prototyping language called Python. It was really the work of a single individual and his band of merry men, Guido van Rossum, who I believe is now a Google employee.

Cerf: I think that's right.

Kahn: He generated a marvelous language. When people list the languages that they'd like to make things work with, they almost always list C+ or C++, they almost always list Java, they almost always list Python as three of the critical. Sometimes they'll list other things Perl and the like, but I think that was a major contribution. The hardest part of that contribution, interestingly enough, was something that Patrice solved, with a little help from Guido, which is, for all the while he was working on this, I was not willing to license the software. The reason I wasn't because, every time I went to Patrice and said, "We need..." that's my wife, Patrice, "...to rent a license," she says, "Well, do you own it?" <laughter> "Where did the pieces come from?" And so I'd go to Guido and he'd say, "I have no idea." It came over the transom, my friend Joe here… And then the answer was, "Well, did they give you the rights to use it?" "Well, no, I mean, they're just my friends." So we had all this stuff and I just was not able to… It's not that I was refusing him in some abstract sense, I just didn't have the ability to do it because I didn't know whether we owned something to license it. When Guido finally decided to leave, which was in March of 2000, this was just before the dot.com boom, he was going to...

Cerf: Yeah, the “dot bust”.

Kahn: Well, yeah, just before the bust. Before the boom busted, he had gone out to take a job at some Silicon Valley firm, which six months later was bankrupt.. He didn't make his millions or billions out there, and eventually he ended up finding a job at Google after three or four intermediate places. But the bottom line was, he eventually understood the need for some regularity. It was not an easy thing for him to decide, but he went back to all his friends. He got, they were kind of blankets: “to the extent that I have any rights in anything that you might have of mine relating to Python, I hereby…” They couldn't even say what it was any more, it had been so iterated on, but we got to the point where at least we felt we had confidence in what we had. And since some of that was derived back from CWI days, Patrice went back to the Netherlands and she had to find out what were the laws applying to employees back then. How do you find the 1984 Dutch laws? It turns out they got it from different places. But we worked it all out and finally were able to license it. Then the most interesting issue was they said it really needs to be compatible with the GPL, that license that Richard Stallman and the folks at the Free Software Foundation had been generating.

Cerf: Why is that?

Kahn: Because a lot of the people who used Python also used GPL code.

Cerf: Oh. And so the results of their work get sort of conflated. Oh, my.

Kahn: That's right. And the GPL regs, which, in fact, I think were part of the reason why the whole open source movement started, said, "If you use any of our software in this collection," then they later modified their rules somewhat but, back then, I think it was, "If you use any of our software, then anything you do with our software has to be put back in the pool."

Cerf: Yes.

Kahn: So it wasn't very friendly to companies trying to build new products -- systems where they wanted to own their stuff or maybe even sell it. The open source thing was intended to get away from that. They wanted sort of transparency in a lot of different dimensions like: at least let us see the code, and don't make the licenses un-understandable, and things like that. I don't think they explicitly prohibited you from owning it yourself and selling it, and things like that. It actually was a very interesting fork in the road that was taken by those guys back then. I think that's something we can be proud of. It's, again, a thing we did for awhile that's got some lasting value out there. I think that was good. The thing that I am most involved with and, in many ways, most interested in, and I think I'm pretty happy with the way it's gone, is the work that we've been doing on the digital object architecture. It has its genesis in work that we had started on digital libraries. In fact, you were intimately involved in that back then. We were starting down the path of thinking that the way to deal with libraries in the network environment was by essentially sending out these entities -- mobile programs -- that could look at different places, pull back information and bring it back. We actually built some early prototype systems that could go out and do similar things, like name kind of "whois" like databases, bring it back. They all came back in different forms, put it in standard forms. You wouldn't know you were dealing in multiple things because these programs would be going out and doing different things on your behalf. We ended up, in the early 1990s, getting the first serious funding to work on that. It was part of a effort called the Computer Science Technical Reports Project that DARPA funded. This was before the web had emerged because that really didn't become visible, publicly, in any significant way, until sometime in the 1993 timeframe, after the work at Illinois had become very public. As part of that project we were asked by DARPA to work with a set of five universities who all wanted to digitize their collections; I mean, the historical stuff. DARPA really didn't have the time or the energy to get into it. The program manager at that time was a gentleman named Brian Boesch, who was working in distributed systems, and he asked if we were interested in that and I said, "Yeah, very." So I actually wrote a proposal to work with that community and I called it something like “linking electronic libraries”. The model that I had at that time was that, if we can enable people to build different electronic libraries, then it was an interesting challenge that was very much like the original internet problem that we focused on, where you didn't want to have to deal with somebody in this collection of networks by specifying what net they were on and what path to get to get there and what protocols to use to talk to them. You just wanted to send email to them and have the system figure it out or whatever. So the whole notion here was that, if you had these different systems, and, abstractly, I was thinking of them as different informational systems, although we ended up focusing on digital libraries, which was a more manageable subset, sort of like countable numbers rather than uncountable, still big. The whole idea was, if you had different information systems and you wanted to pose a query for which the answer was in the collection of systems, how could you do that in a way that would not have to have you know about all the details in these different systems? But, by virtue of having the digital library example in front of me, it gave it a lot more specificity and we could start asking… Instead of asking an abstract question like, "What was the nature of the Greek city state?" or, "How many ways can I count such and such?" or "What's the nature of life?" Things that there may or may not be answers to and then the question domain could be so open-ended. With the digital library world, you were much more focused, like, "Get me today's New York Times" or "Get me the memo I wrote yesterday on such and such". Now, it could be more detailed.. You could ask a more research-y question like, "What was Thomas Jefferson's favorite reading material?" But my view was, it was a very concrete focused problem one could work on. Let's say that I knew exactly what I wanted. Now somebody had done the research and knew what Thomas Jefferson's favorite book was. How could I get that book if it were on the net in digital form? The thing that we created, and, in effect, was sort of part of the thinking all along, was that that these entities on the net would be digital objects, and these digital objects would all have identifiers, and those identifiers could then be used to figure out where they were on the net because you'd have a general purpose resolution system. If you send an identifier into the resolution system, it would give you enough information to have the machinery find it. Like here are the ten places that have it, they might have IP addresses, they might have DNS addresses, they could have handles. They could have any way of identifying the locations, but you could then take the one identifier and find what you want wherever it is on the net. Or you could do other things we thought about. You could authenticate the material when it got back because you could put authentication information. You could have security of all sorts embedded in there as well. That was the basic idea. So we took on ourselves to work on this problem of the generalized architecture for linking these different libraries. I saw it intrinsically as the internet problem reposed at the application level. Instead of having packets moving around, you have digital objects. Instead of worrying about resolving IP addresses or DNS names or using routers, you'd have a resolution system. We actually built such a resolution system. We called these things handles and the resolution system was called the handle system, but what it really is, is a general purpose resolution system that will resolve anything. In fact, we even put a big chunk, if not all, of the DNS on there in getting ready for the Y2K transition, with help from Network Solutions. We showed that the system could deal with that as easily as anything else, except that doing it to do DNS resolutions was like using a supercomputer to add two and two. You had more power in the handle system than you needed for just that task. On the other hand, running the handle system costs four orders of magnitude less than running the DNS system because, there's just not that much cost in running it currently But it's managed in a different way and all the management is built into the system. So that was one of the pieces of the architecture. We knew that identifiers would be important. But what I did with the university community was I funded them not only to do the digitization, but I asked them what research you would do on that. So they all built their own libraries, all different, and they all proposed research programs to use their collections once digitized and so some very interesting things happened. Carnegie did one that involved video and audio stuff, and Berkeley worked on certain environmental stuff in the state of California. Everybody had… Stanford had their own set of proprietary stuff, student theses and the like. We put that program in place and it provided the support to go work on it. I was very interested in how… well, there was a lot of disagreement up front on how the system should be structured, and particularly what the role of semantics should play in it. There was a lot of commonality on other grounds and it brought to mind more clearly than ever the importance of intellectual property and rights clearances and all that.

Cerf: Yeah. Okay. So we had just gotten into the digital objects world, and I'm wondering whether you would like to say anything about the-- you were talking about the concepts but these notions are now actually in use.

Kahn: Yeah, so let me explain what's happened. The whole idea of the digital object architecture is that you have these objects, they are structured, they're parseable by machine, they can therefore be moved from machine to machine and the elements of it understood. The elements are not designed by standards. The structure of it is designed by a meta-level standard, so everything is type-value pairs and you can define new types and resolve them through the resolution system. Once one of these objects shows up, you can figure out what the object means because it's in this meta structure. The identifiers are critical to these objects, just like file names are to files, and there's a resolution system that resolves it. Well, we built the resolution system. We call it the Handle System. It's been on the internet now since 1994. For a long time, ISI ran a part of it. There was an agreement I had with John because he actually wanted to move over toward handles, but he said the politics of getting this through the system is so impossible and, of course, he was in the process of trying to disengage, even. But we've had it deployed. It's been seven by 24 for about a decade now. It's widely used by the publishing industry, but what's even more important is not only do they use the system but they use the identifier means for all of their electronic journals or at least the participating sites, which is most of the big publishers. They set up some organizational structures, there's something called the International DOI Foundation, which is kind of like an ICANN for the publishing industry. They like to think of it as broader but, mainly, they're dealing with publishers at the moment. They do a few things in that organization, including qualify registration agents, worry about policies and procedures, licensing trademarks, marketing and the like. They've got some number of registration agents around the world who they give pretty much carte blanche to work with other parties who want to identify things. Any large publisher is not one of them but they could work with one of those parties to identify all their materials. The IEEE, for example, puts out electronic journals. Their goal was -- and this was part of the architecture originally -- to make it last over long periods of time. If it's a short period of time, you could do a web-like thing, just cite where it is, and as long as nobody cared, that website will probably be good for six months, especially if the people there are keeping it up. But the IEEE would like to put an electronic journal on the shelf and have somebody in 500 or 1,000 years take it off the shelf, be able to click on what's on the back and still find it just like if this were a library and they kept all the collections on the shelf, you could go to the right part of the library, pull that collection off and find the reference. They don't want things to somehow disappear. Mobility is probably treatable in lots of different ways But when you're dealing with objects that are intrinsically dynamic and linked to lots of other objects, all of which can move, building a system that allows this to survive not only movement but technology changes over millennia, is a real challenge. But I believe the use of identifiers to do this enables that to happen. You've got to maintain the system below it, but the things that are embedded in the content don't have to change in order to allow this to happen. That was part of the design goal for this system and I think that's what we have out there. That's the real power of it. It won't be felt until people are beyond the period where the old systems don't work any more, but that was one of the goals. So that's been out there. It's widely used and it has these very nice properties. We also put out the software for it under and Open Source license, mainly because of the interest of the folks in the global grid community who wanted to be able to use it and give it to companies to do software development. Any company that wants to can take the software and make use of it. The only thing that we collect as part of this activity is a small charge, as part of the service agreement, to help support the global registry, which is really what you need to know where to go to do the resolution. Because unless everything is resolved in one place, you need to go somewhere else to figure out where, and that somewhere else could be a big distributed system rather than one location. Also, because many of the people who are using this system insist on using legacy technology, and I call the web legacy technology because it's whatever it was before the system was made available, we built another mechanism whereby you can get ultra reliable access to the system through an intermediate level that doesn't require any changes to the current internet protocols. It's a really interesting contribution. It's really only relevant to the web stuff in its current form, but there's another incarnation that's much more broadly of interest but I don't have time to go into it right now. Another thing that we've done with some industrial support is to look at the application of identifiers and the digital object architecture for the networking part of the problem. We decided not to get into that at all up front because it was so political, there were so many parties, the internet hadn't really taken shape, I didn't want to disrupt all that. But the internet's now gotten to a point of stability, it looks to me like there are real strong interest in trying to reinvent things going forward. And so what we did was, with some industrial funding, we looked at the application on this, the network management, network provisioning and network operations. It turns out you can not only identify all the information flowing in networks as digital objects -- which means that you can store stuff and you can access it, you can put it in archives and it's accessible you can do post-mortems after the fact -- but literally you can identify every element of the network as a digital object, even though it's physically some hardware. Sort of like your being an identifier and saying it happens to be housed here right now but maybe, in another life, we'll house it somewhere else. So if a router somewhere were to go down, you could literally pick that router up and move it somewhere else and you don't have to worry about changing all the parameters of everything internally. Because that router, in your network concept, was a logical element rather than a physical thing, and that physical thing was only its current place of instantiation and you can move it elsewhere. So think about the network as now totally fungible. It's an area we've been pursuing and I think it's very powerful and it's just one more application but at a lower level. So we worked on repositories to store digital objects and we've gone through three generations of them. One generation was part of kind of a monolithic digital library that we worked on with the DOD. Another one was part of a modular implementation of a digital library that we worked on to enable other services to be plugged in by other parties. And the third one was one we did with some folks in Homeland Security that was a realtime one to keep track of very high stream entries. If you're dealing with a library, the number of books that are going to enter into the system per second is probably a fraction. You might have a few a day, you might have tens or hundreds a day.

You're not going to have 10 million a day. People just aren't going to create that many that are readable or interesting. But if you're dealing with sensor transactions or other kinds of things, you might have very large numbers of things entering into it. So we built a real time version, and right now we're working on melding the real time version and the modular version to build a real time modular version. We've applied that in many different contexts. One of the things we pioneered was the creation of compact metadata registries. The organization, in the publishing industry that helps them manage their activities is an organization called a CrossRef. It's located up in the Boston area. CrossRef is a very interesting organization because it's got a commitment to help the publishers manage their identifiers in perpetuity. So you might ask, "Well, what happens if CrossRef goes out of business?", but that's an organization, that's really a membership organization of about a thousand publishers. So they all have to decide to go out of business before that would happen. So I think it'll have some longevitys. The way they do it is they will typically charge a one time fee for working with a publisher. I don't know the details of their business arrangements -- I never asked -- but let me give you an example. Let's say they charged a publisher a certain amount per identifier. if you gave them a book, for example, let's say. a simple model might be one identifier for the book, but that means you get the book, the whole book and nothing but the book. But suppose the publisher wanted to transact business on a chapter or a table or a figure or references or whatever, they might want to give out 10 or 100 identifiers to pieces of the book and, generally they'll identify what they want to transact business on. They're not going to sell you a letter of the alphabet. They're not going to sell you a word in the dictionary. They're not going to probably sell you a sentence, maybe not even a paragraph that might even be “fair use” in some contexts. But they might be willing to sell you a substantive part of it for use in a course or whatever, or a figure, or a cartoon maybe in a magazine. So what they would do is have a single fee -- well, maybe they do a block grant, will do so many books or journals per year or so many identifiers for a set fee. For that one fee, they will create the identifier, they'll get it into the resolution system, they will annotate all of the pieces in the back. Let's say there were references back there, they'll put in the right identifiers, they'll make it all work as a clickable and they will put the metadata into their metadata registry and they will maintain that for the life of the identifier. I call it “cemetery funding.” <laughs> People don't like that idea but it's sort of like one payment up front because you don't know who 50 or 100 or 1,000 years from now will be interested in paying for that identifier because who knows what it was from way back when? But if you can have one fee and they'll then continue it from then on, that kind of works. So that's the way they work. They pay a small amount of fee up the chain. It would be like a registrar in the ICANN context paying -- or maybe a registry paying -- a small fee to ICANN per year for what they do. Maybe ICANN could have a different model, but they pay, like, four cents and it's a declining scale, so when they have a lot of them, they pay less. It's a one-time fee for creating these identifiers, and when all is said and done, that's how the IDF gets its support. But let's say theu charge -- I'll pick a number; I have no idea what the real number is -- let's say they charge a dollar and a half per identifier. That does all those things: creates the metadata, puts it in the system, they maintain it forever for the buck and a half. They might put on a four cent charge, so they might charge you $1.54 instead of charging a buck 50. Four cents goes up to the central body, they keep the other buck and a half and they can set their own business schedule any way they like. That tends to really work pretty well and it's a model that does work. We were very involved in helping to create the first metadata registry. The DOD has gotten very interested in this area for managing training and equipment kinds of documentation. The program is called the Advanced Distributed Learning Program. It's run out of the Pentagon. We actually built the metadata registry for them. It's now deployed at the Defense Technical Information Center, DTech, They actually keep the materials there, too, if they're internal DOD, I believe, but they may reference things that are outside that are just useful for what they do. The DOD has mandated, with that registry, that there be a specific metadata standard used, it's called SCORM and they've mandated the use of the handle system, which has been part of what DTech has been doing for years. They sort of manage that for the DOD. It's a nice application of the technology. So the repository stuff, the metadata registry stuff, the resolution on the system and we're working on other attributes which would provide advanced services on top of that. But I think the thing I'm most interested in and have been very active lately is turning that into an archival system that can last for a very long time What we built is a reference model. It's on the internet right now. It runs at 600 megabits a second. DARPA is hand picking the people to make use of it and they're going to give it free to some large number for awhile. It's got some very interesting attributes of how things work. The most interesting thing is it allows you to move stuff from archive to archive, so it's got a inter-archive open architecture kind of nature to it that's really very powerful. It's something I'm very interested in right now because the ability to move things in this hyperlinked environment and have everything work is really interesting. Particularly since we disentangle attachments from email, and so attachments might go one place and the email might go another and you want it to all play together. I think if there's one message in all this, it's the importance of identifiers. Because the technology for how you store stuff and how you parse it and what the particular things mean and what the standards are will change over time. But the idea of having things linked in some abstract level through the use of identifiers I think is a very powerful construct and one that increasingly, over time, people will understand the importance of. Especially when you don't put semantics directly in materials, but use some other external system to figure out how to map those identifiers to the right level of resolution information to take whatever the next step is. It also lets you move things around with impunity because the identifiers don't change when you move, and the resolution system can sort of do the mappings. That's one other area. There are lots more that we're working on. We're pursuing, I will say with some concern on my part, because I always think some level of management is important and, in this world that I just described, the identifier system, at some level, I think needs to be managed. I don't think you can let it go freeform and not know who's where. That's one of the reasons that ICANN has been so important to the internet going forward for the DNS. But short of that, I'm trying to understand whether there is a model of the future in which everything can be sort of unmanaged. People complain about the internet today as having nobody in charge, or no one party in charge, but yet there are management structures in lots of different places. How much of that can you actually relax, and how much can you do in the way of peer to peer networking, such that you end up with a robust system in the final analysis? That's one of the other things we're looking at. So CNRI is into lots of things. I think one of the things that we've done that's really nice is an infrastructure series that we've published. We're looking at a number of other issues. One of the ones that we're particularly interested in is: can we do something to affect manufacturing at the level of maybe semiconductors or MEMS or nanotechnology to really make it cost effective to do small job lots? While part of that has to do with getting cost effective equipment, because almost all the equipment that's made available in industry today is basically in a semiconductor mindset where you want to make, you know, 10 million or 100 million things through the piece of equipment so it's designed for high throughput and they're high cost. But, when you divide it by the number of objects, they come out pretty cheap. If you only wanted to make 500 objects or 1,000 objects, you can't afford to buy a lot of multi-million dollar machines. It becomes the equivalent of, why is the proverbial hammer so expensive, or worse? It's because you've got a lot of overhead that you have to then apply, unless you can cover it some other way. But it turns out that, even if you could get the equipment costs down to more reasonable levels, they still won't be cheap enough. They've got to be reused. So the real task there is to figure out how to get the equivalent of what we're talking about in the network, commonality of abstractions so that they can work across multiple platforms in terms of manufacturing process steps. So that given pieces of equipment without change can do multiple kinds of processes, not just that one step. That you can package them up, take this equipment, this equipment, and, in different combinations, you know, piece A, B, and F, put it together to do this generic process or that generic process. The real key to making this work, I think, is going to be designing what the abstractions are one level above the equipment and that's another thing that we're currently working on.

Cerf: I'm sure that we could go on and on here but there's only a finite amount of tape. I'd like to suggest just a couple of other specific things. First of all, with regard to the longevity of property in a repository, one of the problems that has both a technical and an intellectual property element to it is being able to interpret the object after many, many years have gone by You can imagine one possibility is you're doing a search of the internet in the year 3000, assuming the internet is still around or whatever its successor is, and you encounter an object which is essentially interpretable by Microsoft PowerPoint 1997.

Kahn: Try VisiCalc.

Cerf: Well, that would be another good example.

Kahn: Do you have a copy of VisiCalc around?

Cerf: The answer is yes. I don't know that it will execute, however. I have a five and a quarter inch floppy disc, which I don't have a reader for, that if I had the right Apple IIe hardware...

Kahn: So all you do is take your Blackberry or whatever your device on which your email arrives and you plug it into this old floppy five and a quarter inch reader and it runs VisiCalc on your attachment and you're all done, right?

Cerf: Yeah. So the issue here is not just the physical medium, which I think you could argue could be converted and transmitted, as long as you can convert bits into different storage formats, but the problem is what do you do about the software that interprets the objects and maybe even the operating system that was needed to run the software that interprets the objects? And the reason it's an intellectual...

Kahn: What about, what about the hardware in which the operating system has to run?

Cerf: You know, the House that Jack Built, yes. So here's the question. Have you already tackled the problem of obtaining rights to use the software 1,000 years from now to interpret-- I'm asking this in part on the grounds that I don't think we're going to be able to always transform a particular object into the next version of something that can interpret it. Backwards compatibility is not always guaranteed. For example, Microsoft recently announced that it's not going to support Windows 98 any more, and so if you don't have any rights to Windows 98, if you don't have the ability to get the source code or do something, it may be that things that were only workable in Windows 98 will never work any more. Or after a period of 20 or 30 or 40 years. Is there anything in the repository ideas that would allow you to capture the necessary software in order to keep interpreting objects over really long periods of time?

Kahn: Well, I think this is not a technical problem. I mean, you and I could sit down and we could design a strategy that would do it all. The real issue is, do we have the right to do that and do we have the underlying mechanism to capture everything that we need?

Cerf: Yes, that's why I'm asking the question.

Kahn: If somebody built the general purpose emulation machine and then every architecture that were ever put out had the right emulation pattern to it and we could preserve that over time, we might be in pretty good shape. If everyone could agree that Java was the answer forever, then all you'd have to do is worry about getting a Java compliant environment underneath it, however you chose to do that, but that's not what we face. That's a technical challenge that could be solved with enough time, energy and money put out, I believe. The more interesting question is not even can you get the rights to the software because, sooner or later, it's going to be out of copyright. Sooner or later. If you're really talking about 1,000 years from now and you just kept your copy on your machine long enough, assuming your machine still ran long enough to produce it, you could put that in some kind of an archive and maintain it and there are people who are trying to build software constructs. You still need to figure out what are you going to run it on whenever, but let me assume that problem is solved, too. The really interesting problem to me, and the one that we are trying to clear rights to -- let me say we will be trying to; we've had some discussions about how to go about this, but, ultimately, it will probably be up to the lawyers to figure out how -- is how to get rights to the proprietary data structures that all these programs use. If you want to generate a good job of building a metadata registry, it's not enough to know that this object showed up on December 3rd of 2007 or that the author called it Phil. You probably need to know something about what's inside of it. Now, you could ask the author to generate all the metadata for you and give it to you separately, but most people aren't going to want to do that, and for a lot of stuff that you're collecting, it would make no sense to do that. Let's say this was an image that was coming in, and let's say you had 100 of them that you wanted to put in or maybe it was even a stream of them. You're not going to want to go through every single image, I believe, and figure out what it is to say about it from a metadata perspective. And yet when you go in to try and access it at a later time, you're going to need to know things that you didn't think about. I didn't think to mention that you were in that picture and, well, maybe you could be found because you're so distinctive that anybody can pick you out of any picture at a moment's notice, but most people can't, or can't be found that way. Unless you've got something that knows how to do that in these proprietary data structures, you're in trouble. So that's one thing we can do. Second thing that we can do is map those things, if we know what kind of thing they are, into some standard form. I'll give you an example. I think what we now know as spreadsheets will have some kind of longevity going forward. I don't know exactly what form they'll take, they may be 3D spreadsheets, they may be in hyperspace, who knows what they'll look like? But let's say, for the moment, there's a generic kind of an object called a matrix, that's what a spreadsheet is. It's got so many rows, so many columns, so many cells This is something we've actually been working on. Let's say we could describe that in some generic form. You give me a spreadsheet and I don't know what's in the proprietary structure but I know it's got m rows and n-- you tell me or I can figure out -- it's got m rows and n columns and I can figure out what's in every row and column. So my object says the ij-th element is such and such, and maybe I store the original and maybe I store the abstracted version. Then, at a later point in time, 1,000 years from now, if I pull this out, let's assume the semantics are such that the word matrix means something. It hasn't changed its nature. So they know this is something that's got rows and columns. They can pull it back out. They can then import it into any system of the future that's got all the bells and whistles. It knows about colorization in printers and screen formats and all the stuff that you need to manifest the data but you're able to get it and import it. Now, they might, you could argue that, maybe, 1,000 years from now, somebody will build the interfaces to important, Microsoft spreadsheet or the VisiCalc spreadsheet but I wouldn't bet on it. I don't even bet that they'll keep things in ten years from ten years ago, much less, you know, for a millennia. But I think you can get into that intermediate stage and actually make things work.

Cerf: This sounds to me like it's an area where lots and lots of research would be well worth the effort.

Kahn: And I think that's exactly what's going to happen. We may play a part in helping to organize some of that. But just to show you how difficult a problem this is, what do you think could be easier to characterize than an image? I mean, all it is, is a set of pixels.

Cerf: Yes. Although as soon as those images are compressed in any way, suddenly you have a whole algorithm to worry about, to undo that to generate the appropriate pixels.

Kahn: I'm going to say that compression is not allowed. We have so much memory you don't have to compress it. You're right. If you compress it, you got to understand how to uncompress it and the like. But let's just say you have an image, a certain number of pixels. You'd think that interchangeability between any two parties would be identical and yet I'm sure if you go to the camera industry today that they all have different ways of standardizing things. The reason is that no only do they have different ways of characterizing the elements of their pixels, but they probably all embed metadata and what kind of camera it came from and they have all this other stuff and it's all in some probably proprietary data format, unless they agree to share it all. So it's more than just the images. I often think about what's called knowledge attachments. If you've got a map of something, it's one thing to have something that will enable you to understand just the map but suppose somebody then annotates it to tell you what the street names are and the buildings, the legacy of the buildings, how much the building cost, where the sewage drains, You may not know how to process that image if you knew about processing maps, because you don't know how they coded it and what all the details are. What you think are images or pixels turn out to be something different. So without knowing how it was all coded together, you know, you're dead in the water. So I think keeping one's leg up on how that's going is going to be important. think the only way to do it, quite frankly, is to specify standards for people who write programs who know what their proprietary data structures are. Not to disclose them but to generate relevant metadata that might be of interesting according to some external standards, so that the programs can generate it without necessarily, disclosing proprietary data structures.

Cerf: In some ways, this is reminiscent of what happened with the web, because the original HTML was being generated by hand. Then eventually, that got a little tiresome and people wrote programs that would allow you to do direct manipulation of the content, images, and text and generated the HTML for you. Getting something to generate the appropriate metadata is probably part of the problem. I have to say, when I think about spreadsheets, that trying to retain their interpretability over long periods of time is really tough because, if you look at the content of any one of the matrix elements, you could have all kinds of references with identifiers in those references...

Kahn: And it should.

Cerf: ...that could be… Well, that's what gives spreadsheets their power but, when they start making references to other spreadsheets by file name or by URL or something...

Kahn: Dead in the water.

Cerf: ...you're-- well, or at least you have a really tough transformation problem if you're trying to retain that over a long period of time.

Kahn: Yeah, like 1,000 years from now.

Cerf: So, in some ways, there's a digital tower of Babel facing us as we start accumulating more and more digital information if we don't start tackling these problems, so I'm glad that we had time to bring them up. Well, I think we should try to round this off with a couple of future-looking questions, if you have the energy to do that. Since you've had such a key role in the development of the internet, maybe you'd care to speculate about what its future is likely to be like. There are any number of axes in which to answer that. One could be the technological one, another could be policy, another could be governance and another could be applications and you can probably think of others. But as you think a little bit about this thing that you were so central in creating, do you have any current sense of where you think it's likely to go and in what directions it's likely to evolve?

Kahn: Seems to me you've been even more vocal than I have on this subject in different quarters. I'm sure you've got your own views. I don't really know where to begin on that. I could tell you what's obvious now, but remember I said leave a little room for serendipity. Things that are going to shape the internet in the future are the unknowns that just change the dynamics of everything in unpredictable ways. I am sure the Founding Fathers must have debated what this country is going to look like in 1830 or in 1870. They would have had to throw up their hands, because how could they have predicted the depression of 1930? How could they have predicted the Vietnam War? How could they have predicted what the world looks like right now? We build an organic system and let it go. The internet is really a dynamic organic system. I could tell you that wireless will become a more endemic part of the system in all ways. We could even speculate about whether wires will go away or not. I don't think they will, by the way. Or maybe they will because there will be fiber somewhere or other. We could speculate about how much bandwidth will show up. We could speculate about how much information is relevant to keep. We could speculate about will people ever throw anything away any more. We could speculate about the liability of keeping all that information around. Maybe it's worth spending a lot of time to get rid of stuff, or maybe it becomes mandatory to figure out what you need to get rid of and budget it. I was really surprised to see, in the financial standards, how they require that you put on balance sheet the cost of rectifying problems that they know about that you might have to deal with in the future, like asbestos removal from a building. It's no longer possible to ignore that. You now have to figure out what it will cost you to remedy that, put it on the balance sheet...

Cerf: As a liability.

Kahn: ...as a liability right now, even if you don't do anything about it for a long time, because ultimately, when that building comes down, you're going to have to deal with it as a separate issue. You can't just demolish the building. So, who knows? But I certainly think that the future of the internet is going to be tied to information: information access, information management, collaboration among individuals to a much greater extent than it is now. I think we'll see the world wide web as a first step in information access. There will be many other systems that get evolved as we go down the pike. Just like the Arpanet was not the first and only network that took over the world, the web has got a lot of worldwide penetration and it's in its name as well but, eventually, there will be other good ideas that people have. We can't predict what they are and they will embed themselves. I think the ability to access information of all kinds at any time is going to have profound implications that I wouldn't even want to speculate on right now. Because if you can know whatever you want whenever you want to, it can have the appearance of making you an awful lot more savvy about things than maybe you really are, but who knows what the effect of that is? You might as well ask, "Would this be a better planet if everybody were smarter?" Or if everybody were dumber? Or if the span were bigger or smaller? And I don't think you can answer that because everybody adapts. I remember once visiting my sister, whose husband was a doctor in the public health service. They were in Kahli, Columbia. I'm used to living in a place where it might get down to minus 10 or 20 in the wintertime and it could be 100 or even 110 in the summer, so a fairly big range. In Kahli, Columbia, I think the low for the day was, like, 70 and a half degrees when we were there and the high for the day was maybe 72. There maybe was a one or two degree differential. And the people would come out in the morning with sweaters on <laughs> and they'd take them off in the afternoon because they were all sweaty and hot and then they'd put the sweaters back on again at night because they adapted to the dynamic range that they saw. I think it's very hard to predict exactly what kind of dynamics we'll have in different things going forward in the future. I think the one thing that's going to be with us from now on is the need for pervasive connectivity, and right now we see it mainly in an earthbound context. Some folks, yourself included, have been thinking about the planetary reaches. I'm not sure what an individual will need to do relative to, you know, planets like Jupiter or Pluto or whatever but who knows? Who knows what businesses might develop about information access to things you would think have no value today but that could develop as well. Information you think is really important today may turn out to be of no value in the future. Things that you didn't think you could even know about turn out to be the most critical things in the future and you can't live without them any more. We may start to measure things that today get written off as unimportant because minor changes are indicative of major calamities down the road and we know how to do those correlations. I'd love to be around to see how it all evolves, but I'm sure that what we will have is the need for this technical connectivity. We now know the power of the computing technology regime, but I don't know whether that's a passing fad and will be replaced by something else that's got those same kind of capabilities and more. Microfluidic computation systems, things that could be embedded, empowered by chemical reactions in the body, who knows what the future really will hold? But whatever it is, I think the ability to link these kind of things together will be with us for a long time. So, to that extent, I think the internet legacy, regardless of what the protocols are, whether it's TCP/IP or it's next generation or whether it's IP or it's next generation, I think what we were able to demonstrate was the ability of things to interconnect on a global basis, maybe beyond. And I think it will be as important to the world as identifying for the first time the importance of transportational systems, and it doesn't matter whether it's your feet or a car or a train or a plane or whoever knows what comes next. If somebody generates an ESP system to move you from Point A to Point B, you'd say, "Well, that's just another example", even though the technology may be pretty hard. So we don't know what it's going to look like. But I think there's something fundamental going on here, and I think it will be with us for a long time. I was delighted to have been part of it. I was especially delighted to have worked with you on it, and in many ways even more so with Patrice because we did very intimately, but I think the world will look back and see this as one of the major contributions in history.

Cerf: Thank you, Bob, for taking so much time. This is pretty special, I think, because it demonstrates, this long discussion demonstrates, the breadth and depth of contributions that you've made personally. I'm glad to see that the Computer History Museum is recognizing that. I'm sure they'll take some small bits and pieces out of this many hour discussion to share with the people who will celebrate with you at the Museum event that recognizes your fellowship. I wish I could be there with you and my calendar isn't going to permit but I'll be certainly there in thought and, again, appreciate very much your time today. Thank you.

Kahn: Thanks very much. Thanks very much for enabling all this.

END OF INTERVIEW

