ROBERT KAHN, Ph.D.
ORAL HISTORY

COMPUTERWORLD HONORS PROGRAM
INTERNATIONAL ARCHIVES

Transcript of a Video History Interview with Robert Kahn, Ph.D.
Founder & President, Corporation for National Research Initiatives

Recipient, 1996 SAIC Leadership Award for Global Integration

Interviewer:	David Allison (DA) Curator, National Museum of American History, Smithsonian Institution
Date:	April 20, 1995

DA: Let's start with your upbringing. What were the things that influenced you while you were growing up?

RK: I think the primary influences on me as a child were my parents. My Dad was a product of the Depression. He was a high school teacher, and later a principal at a school in New York. My Mother was a lady, a really a lovely person who unfortunately suffered from many heart attacks as a youth due to rheumatic fever. So she was largely bedridden while I was a youth, and that cast my sister and me almost in the role as housekeepers when most kids would be out playing in the streets. But I think even to this day the respect that I have for my parents and their sense of commitment, dignity and insight has really been the most significant factor by far for me.

DA: Did they approach you to choose science and technology as a career, or was that your own choice?

RK: No, I think they pretty much left my sister and myself free to make our own decisions about our own careers. And I think they would have been happy with virtually anything we did that we were happy with, and was productive from our point of view.

DA: Were you interested in science and technology type things in high school, or was it not until college you got interested in them?

RK: Well I was quite interested in issues over science and technology as a child. Not so much through school because for me as a youngster, school was almost too easy. There was not enough challenge in it for me. Things that we were taught in school, I thought could have been taught in a more compressed fashion. Homework for me was a rote exercise, doing what I knew I could do, and so usually I would put that off until the very last minute and then do it in lightning speed simply because the challenge wasn't there for me to get too involved.

I was much more interested in other things, like chemistry. I had a chemistry set that I used to do experiments at home. I had an erector set with which I used to build things. So I was quite interested in technological things, but generally not through school until I got into college. College really awakened me in a much more active way. So I became part of the process rather than being just a recipient in educational training.

DA: And you were a student at City College?

RK: I had actually spent the first two years at Queens College. In fact I started out as a Chemistry major. The thought was, not to actually go into industrial engineering, but I chose chemical engineering for whatever the reason - I am trying to pinpoint it - but I think maybe I didn't like the lab work so much. So I became more interested in electrical engineering, and I switched somewhere there in my first year. Then the program took me to City College at the end of my second year. So I spent most of the time at City College.

DA: Was there a professor that particularly influenced you at the time?

RK: Not in making the decision to switch. I later became very close with some of the faculty at City College, but of the faculty there, the only one I spent quite a bit of time with was a professor named Brenner, who was particularly influential in getting me to plan things out for myself and stimulating me to go on to graduate school. But I would have done that anyway.

DA: What did you think you were going to pursue as a career?

RK: I don't know that my career was part of a master plan. At every stage of the game, it seemed that the next thing to do was pretty well out there and just more or less self evident to me that that was the right next thing to do. So there was never a question in my mind whether I would or would not go to graduate school. I was going to go into graduate school and go into the PhD program, the only question was the timing and where.

DA: So you wanted to be sure to get into an advanced degree in whatever your field was going to be. Is that the way you remember?

RK: I think the whole notion of education was instilled in the family. We always had interesting discussions around the dinner table. Learning was a part of the culture in which we grew up. I never really questioned the fact that going on as far as the educational process went was what was going to be the case. Both my sister and I got advanced degrees, and I don't think it was ever a question whether we would or not.

DA: And you chose to go to City College. What was the thing that made you decide on that institution?

RK: I mentioned earlier that my mother had a series of heart attacks when she was a young woman. The net effect of that was to keep the family more closely bound. So I ended up going to school in New York simply to be closer to the family, although I had clearly explored going to college elsewhere.

When it came time to go to graduate school, my thoughts were still to try and be closer so I could assist my dad if he needed help, to be there for my mother if she needed the support.

Princeton was a fine school, still is a fine school. Its electrical engineering department was really being grown and born in many ways at that point in time. It just seemed like a good time to be there. The opportunity was right. I won a National Science Foundation Fellowship to go to the school of my choice and I selected Princeton for those reasons. And I was very happy with the time that I spent there. It was a very growing experience for me, in most dimensions.

DA: And this was now in the early to mid 1960's?

RK: I got to Princeton, I think in September of 1960 and I stayed there through roughly the summer of 1964 when I left with my degree and went up to MIT to teach.

DA: How would you describe the educational environment and the group of students at Princeton, in terms of how it shaped your later career?

RK: It was an extremely diverse group, a very bright group. I would say everybody in that graduate school program - certainly the people that came through the program - were brilliant in their particular fields. It was the first time I was put together with a group of people who, almost uniformly, were extremely bright and motivated and creative. The environment was just a very enriching one because I had the opportunity to deal with people in many different fields. And also rubbing elbows with people in the aeronautics field or mathematics or chemistry, history or English, was very good, it was what happened at that time. It was almost part of the planned milieu of graduate college, that you dealt with all these people, that you rubbed elbows with them, you were friends with them, you had dinner with them. You learned what they were thinking about, what issues or problems they were working on, and it was a very cross-cultural kind of activity.

DA: And did you work at Bell Laboratories while you were at Princeton?

RK: Well, when I left City College I was in one of these programs that didn't exactly end on an academic year boundary. It was a four-and-a-half year program because of the City College, Queens College split. There was a bigger chunk of humanities in the earlier part of the program at Queens College, so it was a four-and-a-half year program. I graduated in January of 1960 or December 1959, I'm not quite sure, but it was at that boundary. The official graduation was in June of 1960.

When I left I took a job at Bell Laboratories working in a mathematically oriented group run by Roger Wilkinson. It was a kind of operations research group worrying about global planning for the tele-traffic kinds of things in the Bell System. I went there with the idea that it would be a little bit of experience before going to graduate school, and the folks at Bell Labs knew that when I came in. Subsequently, when I won the NSF Fellowship in that following September, I went to Princeton full time. But I kept going back to Bell Labs in the summers, thinking that I might make that a career if I didn't go into teaching. I really wasn't sure at that time. So I went to Bell Labs in the summer of '61, and the summer of '62. The summer of 1963 I spent just getting my thesis thoughts in order. Then in 1964, I finished my thesis and went up to teach at MIT, which is what I finally decided to do.

DA: Well you had already had a good mix of seeing it from the corporate perspective as well as the academic perspective before you went to teach.

RK: I didn't have any direct experience teaching before I had gone up to MIT. And I liked Bell Laboratories quite a bit. I worked for Roger Wilkinson in my first job there. Roger was a grand old man of the Bell system, sort of a Walter Cronkite type of person. He used to sail his boats up in Westchester during the weekends, but he just had almost a grandfatherly view of how to manage the group, and he was well respected by everyone there. He could very well have been the last person to have a complete view of what the Bell system was about, from a point of this whole traffic engineering of it.

Also, for one period of time there I worked with a mathematician named Steve Rice, who is well known historically for having developed the mathematical theory of noise, which turned out to be critical in understanding how the performance of certain communication systems behaved. Steve was a master of applied mathematics. I mean he knew more about that field than anyone I ever ran into. I learned quite a bit about the application of mathematics to real problems from Steve while I was there.

I was really trained as a mathematician at Princeton, despite the fact that I was in the electrical engineering program. We took a lot of courses in electrical engineering, but mainly the program at Princeton was a self-teaching program. I spent more time in the library reading books on topics I was interested in, and was taking courses in physics and mathematics, as well as core competency stuff in engineering. But when I left I felt like I was really an applied mathematician, more than I was an electrical engineer. So when I went up to MIT, it was actually a very good fit because MIT was quite interested in the mathematical side of electrical engineering in that department.

DA: And your job was to join the Electrical Engineering faculty and teach at what level?

RK: Mainly it was teaching graduate level courses. I taught courses in communication theory, information theory. And what I noticed when I was up there, particularly through the interactions with some of the key faculty, was that they had a sense for engineering kinds of problems that I somehow didn't know where they got. For example, a mathematical course that looked incredibly interesting to me, and sometimes the reaction I would get would be, "You know, that's kind of interesting." And sometimes the reaction I would get is, "That's really interesting." And I wouldn't understand what the difference between interesting and really interesting was.

It turned out that the faculty up there had a lot more experience in the practical application of what was going on than I did, and that was something I badly needed. It was recommended to me that I take a year or two - this was in my second year - to take a leave of absence and get some practical experience. The advisor I was talking to at the time turned out later to be a very good friend of mine, Jack Rosencraft, and he gave me a list of people that he thought were mentors that might work well with me. He actually talked to them about my going to work with them for a year or two, and I ended up following that strategy.

I took a leave of absence with the thought that I would come back in a year or two. I went to Bolt, Beraneck and Neumann, which at the time was largely an architectural acoustic firm, a rather small firm in the Cambridge area. I don't know what their revenues were but this would have been 1966, and they were probably 15 or 20 million dollar total revenue at that time, I would guess. It was a fairly small company in those days.

They subsequently grew to be quite a bit larger, it must be three or four hundred million dollars now, if not larger. They are growing very rapidly and I think to the large extent, the fact that they are in the computer and communications business is due to the activities I was involved in back in the 1960's when I got there.

DA: I wanted to ask you about that and the growth of the Internet.

RK: You know it's very interesting, in the swell of history it's very hard for people to sort all these things out – the individuals, the collaboration. This was a case where we worked so closely together for so long that any attempt to tease this apart ends up being a disservice to another person so I don't even try. I really don't try. I mean I know what Vint (Cerf) brought to the party, and I know what I brought to the party. So it's sort of like saying we held a big spaghetti bash and he brought the eggs and I brought the flour and we all had the spices.

DA: When I talked with him, he felt exactly the same way. He was very much saying it was a collaboration.

RK: Yes it was. It was.

DA: Let's drop back now to the time you went into BBN. You talked about the size of the firm but as you came in, tell me about the group you worked with and what you were doing.

RK: Well, I came to work for a fellow by the name of Jerry Elkind, who very shortly after that, got to run to run about half of Bolt Beraneck and Neumann. BB&N at the time I joined it, was still largely an architectural acoustics firm. I think they were founded in 1948, or thereabouts, by Dick Bolt, Leo Beraneck, Bob Neumann, and a number of other people. The essence of that firm was from and about architectural acoustics; acoustics backgrounds.

The nature of the firm changed somewhat in the early 1960's as they got into computing in a small way. By the time I joined them in 1966, they had had some support in some activities in computing. They were working on artificial intelligence. They were working on time-sharing systems. But communications was not really part of the mix at the time that I joined them, at least in any substantive way. I guess people used dial-up modems to borrow time-sharing systems to run time-sharing applications. And they had been big in the application area, both in offering commercial time-sharing services and in providing medical information services.

But I got well interested in the networking area. That was the area that I charted out to work on. And in the period that I was there in 1966 and 1967, I really thought through all the issues that seemed appropriate to me, and thinking through how one might actually create a network to create computers. This was not a very fashionable topic at the time. You have to remember that in the 1960's, batch processing was still the main way that people did computing. You still submitted either bank tapes or punch card decks. The early time-share systems were pretty limited. The major companies had not really committed to them as main parts of their businesses. Digital was one of the first to produce both PDP 1 and PDP 6 to 10 series, and a few other companies like Scientific Data Systems, that produced machines that ran time-sharing systems. But, there was a lot of homegrown activity in those days; special purpose systems that were built, so it was really only in the research community that it made sense to try to move these machines. It did not really make a lot of sense to the operators of computer centers to leave their batch machines, because you could sort of get everything you wanted in one center.

So it was an area that I was working on. Not so much because I had expected that it would become a national network, but because I thought it was an interesting applied idea, and that I would gain some understanding in how to build some real things. I didn't actually expect to build it. I also did not realize that at the time that I was starting this work, that ARPA was interested in creating a network. This was an effort that really was an outgrowth of some early ideas that JCR Lickleider who founded the ARPA office, had. It was an activity that was pushed very strongly by Bob Taylor, who was at that time the office director, and by Larry Roberts, who he subsequently hired to go run the ARPANET project.

At the time I started working on it, I did not know that there was such a plan, or that ARPA had any interest in this, it was just interesting to me. Somewhere along the line, Jerry Elkind had come up to me and said, "You know I think the folks at ARPA would be particularly interested in what you are doing." I had been researching my thoughts on dealing with routing and buffering or float control, or error control, what have you, and I remember writing a letter to Larry Roberts, sometime in 1967 or thereabouts, basically introducing myself to him. You know, "I'm a researcher, been working on nets and Jerry Elkind suggested I write you this letter and let you know that I have some interesting ideas in case you want to hear about them." And I got a call back from the ARPA office, it could very well have been Bob Taylor that called me and basically invited me to come down. I met with Larry and with Bob back in the ARPA office and we got to know each other. ARPA subsequently issued and RFQ, essentially a request for proposals but they had spec'd out a kind of conceptual design for a network. It talked about using mini computers and linking many computers with high speed telephone lines, and the ability of computers to somehow inject messages into these nodes, have them broken up into packets and have the packets independently routed through and delivered at the other end. Very much the very same kind of idea that I was working on, but this was now documented as a real request, for a real network that someone would actually build.

DA: Can I get you to talk about some of the discussions…there was talk about ARPA's goal was packet switching for security reasons, but your interest in packets had lot of interesting facets.

RK: I just thought it was an efficient way to build a network to link computers; that is, the ability to send information in discreet chunks. It was not really motivated by security interests at all. In fact let me come back to the ARPA part of it, because exactly what the ARPA part of it was, is either a matter of personal opinion or record, depending on your perspective. At the time that I was doing this early work, I was not aware of ARPA's interests until I got really introduced to them and found at that they were serious about it.

At one point along the way there was another fellow who worked for Jerry Elkind by the name of Frank Hart. Frank ran a systems group in a different part of the building, and Frank showed up in my office, basically introduced himself, and said, "I understand you're interested in networks. We should talk." Frank came from an experience of real time systems building. He had come from Lincoln Labs, and he had a number of people that worked with him that had been involved in real time systems building. And it was very clear that he was interested in building this network. Whereas, at that time, I was mainly interested in how it could work. It was a complex systems engineering problem, and what was eventually listed in the RFQ, was not sufficient to build to because one had to really architect it. It's the difference between saying "I want a building that is 7 stories tall and has got panoramic glass windows everywhere," and then actually figuring out how actually you were going to build the thing. There were parts of the design that were described in the RFP when it first came out, but basically it was an intellectual challenge to take that description and figure out how you would actually build it.

I became very involved and agreed to work with Frank on actually writing down how it would actually work. And then the next step, after we had won it, based on the proposal that we put in, was to take that description and actually implement it, which had a whole other series of issues affiliated with it, because descriptions of how it would work in detail, had to be coded in software. It had to be coded to fit in a very small amount of memory that existed at that time. We were dealing with initially machines that never had more than 16,000 words of memory, and 16K of memory is tiny by today's notions. And I think we started by thinking it had to fit in 4K of memory because we needed room for buffers and we weren't sure how much would be required there. So the eventual task of taking the design once it was described in more detail and actually coding it in software was another task.

DA: Talking about the request from ARPA. Were there other serious competitors, or were you pretty much out there alone?

RK: I understood later than there were a number of other proposals submitted, and that it got narrowed down by the ARPA office to a handful. What I heard was that the choice of BB&N was based on the technical quality of what was in the proposal, and also the track record of the people. We had a very broad mix. We had people who had the ability of building something and making it work. We also had the capability of understanding the complexities of what it would take. I think it was based on that complex mix that BBN won the award. And I think it was proven out in the final analysis because we had a very tight time frame for getting something out there that would work, nine months as a matter of fact, and it was delivered, as I recall, a day or two early, for the first time.

DA: And you mentioned that the memory was a significant problem in actually realizing what you did. What were some of the other tough technical issues?

RK: From my perspective, the most difficult issues had to do with understanding the way in which the network would work. I was very concerned about issues of congestion control, efficient routing mechanisms, the possibility that the network could actually deadlock, where in fact things would be blocking other things from going and they both had to get through but they would be prevented by virtue of how the protocols would actually be implemented. So I was very concerned that we get it right because I thought that conceptually it was rather difficult to just build and know it would work.

I think there were other issues that Frank was more concerned about than I was. One of which was the whole issue of reliability. I mean, here we are going to put nodes in many universities around the country, and he was very concerned that students might go in and flip switches and pull plugs and otherwise mess with these machines. And also, they had to be flown around the country in a time when these machines were still pretty fragile as a technology. So he took extraordinary steps to make sure that this thing was going to work when it got where it got. Several years later we realized that perhaps such heroic measures were not necessary, but if it hadn't worked up front, I think we would have all regretted it. So I think those were some of the issues we had to deal with up front that were important.

The other one of course was the dealing with a very large community. We had people out there that needed to know what this was about, and had to figure out how to build both hardware and software to interface with it. One of the key elements in doing that was letting them know what the essence of the design was in time for them to be ready while it was still evolving. I think one of the key elements along the way was a document I wrote at BB&N. It was called BB&N report 1822, which spec'd out the interface between the host machines and the Internet; both hardware and software. And the interface became known as the 1822 interface. It was probably the one thing that the research community had to deal with that let them know how to relate to this network in the early days.

DA: You mentioned Vint Cerf, and others; please talk about your relationship with him and how that developed, and the whole community of people that were involved in this early development.

RK: I first met Vint at UCLA in some of the early testing work on the ARPANET. I had gone out there on several occasions to do testing on the network, both on its initial configuration, and later as it expanded and had more capability. Part of that activity involved connecting the UCLA host machine to the network. They were running something called the Network Management Center…Network Measurement Center perhaps, and they were taking network measurements.

The ARPANET had been designed, and was instrumented to allow measurements of the network to be taken and sent to any place that had requested it and that was allowed to do that. Vint was involved with the team of people there who were doing this measurement work. He was my principle interface at UCLA on this, and we worked together very effectively.

Dave Walton from BB&N had come out with me on many occasions, in fact on every occasion I was out there, and Dave and Vint and myself actually did most of the work on the early testing of the ARPANET. That's where I first met him. At the time he was a graduate student at UCLA. He subsequently got his Ph.D., and went to teach at Stanford. And while he was at Stanford he and I began to collaborate on what became the Internet, basically on TCP/IP protocols; but that it a whole other story.

DA: So you got involved with the ARPANET and then later that was a different story. Do you want to carry on the story of the development of the ARPANET from your perspective?

RK: Let me see if I can set the stage for what happened in the Internet development. The work that we were doing on ARPANET essentially built a packet switch that could be used by other people to communicate. It is what I think demonstrated the viability of packet switching as a communications concept. But in order to make it work, you had to plug the machines in, and machines had to figure out how they were going to communicate to each other. So in parallel with having to build the communications, there was an effort by a set of people led by Steve Crocker, another graduate student at UCLA, but the effort also involved Vint in a very central way, and others around the community, to figure out what kinds of mechanisms these machines should use to communicate with themselves. And they came up with one of the main contributions to the field, which is the notion of layered protocols; which is, if you could build protocols such that more functionality could be put on top of them, then they could essentially could be built on work that had been done previously. So one of the protocols that was developed was something called the Network Control Program or Network Control Protocol, it was known as NCP. This was the original protocol for host computers on the ARPANET.

This was not something I was spending any of my time on at all. In fact I was full-time busy trying to make the communications piece of the network work, the subnet portion. But when I first saw the design for the NCP, two things were clear to me. Number one; that it would work; but, number two, it wasn't the right interim choice for a number of reasons. One of the reasons it was clear that it wasn't, is that it depended on the ARPANET being perfect. That is, if the ARPANET ever lost anything, then protocols would need to be restarted and what you wanted was something that was a little more robust. Sort of like when you deal with a printer, and if the printer stops working, then you go fix it.

So the idea was dealing with problems if the ARPANET ever lost packets when it wasn't supposed to. The ARPANET was designed so that it didn't lose packets, but you could always have a case where a pair of machines went down due to lightning strikes or power outages – rare events but that could happen. The other thing was, it depended on a mechanism whereby you had to send something from one side of the net to another. And you had to have an acknowledgment for that thing before you could send the next thing. There was no way of essentially streaming data through the network like oil through a pipeline. Therefore, the longer the network delay, or the more the hops through the network, the less effective trip you could get, because the amount of data was calculated from round trip delay time. Later we discovered a number of other limitations that weren't apparent to me upfront, but I put that all beside because what they had obviously was workable in that context. In fact, it not only was workable, it worked for more than a decade. It was the mainstay for the ARPANET network for its first 13 or 14 years.

The thing about the Internet development was that it didn't really come from a grand plan for the Internet as we now know it. When I got to the ARPANET office in 1972, there had already been some early work on other kinds of network technology based on packet switching.

One of which was radio based packet switching – in fact there had been there had been an effort out of the University of Hawaii, led by Norm Abrahamson and Frank Quo, to demonstrate how you could send packets over radio waves. They built a system called the Aloha Net. Basically it was a one-hop system, from a user's terminal to a central station, but it was meant to circumvent what was then, the very noisy telephone lines in Honolulu.

We developed a packet radio system that was an extension of that, subsequently. It was like the ARPANET except that all its nodes were small mobile computer controlled nodes. This could very well have been the first defense application of microcomputers. We used the national M-16, the first 16- bit micro net that was on the network as I am aware. It also used spread spectrum - prior to then, people did not think that that technology could be made in small radio units. So this was a system that could be all mobile. So you could use it to have an ARPANET on ships at sea, or on an airplane, or with a field army that was moving, or whatever. There was also some work that had been started to apply packet switching to satellite channels, satellite that had a footprint that was broad enough to be seen in many geographic areas could have the effect of similar transmissions being received by multiple sites, or selected sites, or one site, so it's like an Ethernet in the sky.

When I got to the ARPA office, I started the packet radio network and picked up on some of the work on packet satellite that had been going on. At the time, the plan for packet satellite work was to treat the satellite as if it were an imbedded part of the ARPANET. That is to say, if you took an ARPANET node, called an IMP, the satellite code would be in the upper part of memory and the ARPANET code would be in the lower half, and the communication between them would be internal memory transfers in the machine.

The concern that I had was that this was not expandable into an open architecture type of environment. So one of the very first things we did was to modify that approach; the satellite net would be a separate net. So we created the packet radio net as a separate net and we certainly had the ARPANET as a separate net. We knew another net that might be on the horizon. In fact BB&N was just setting up Telenet, which I worked with Steve Levy on at the time. So the idea of the program at the time, was to find a way to take multiple independent networks and to both link the networks together through the use of intermediate translating boxes, which were called "gateways;" and to find a way to get the machines on opposite ends of these networks to communicate with each other through the hybrid of networks.

Now the issues were not all apparent when we started, and how we would deal with them. We had problems where one machine might take packets that were this size, but the next system might take ones that were only half the size. So you might have to chop things up along the way. They might have to take different paths. You might get duplicates, and you would have to put them together, and so forth. So I had a pretty good idea of what the network part of this might look like, and how the end-to-end protocol should work to improve upon the performance of the NCP protocols, but I didn't know what all the issues were of imbedding these protocols in the very large number of different operating systems that were involved at that time.

That is why Vint and I became such close collaborators as we did. That was a large part of the issue of developing of the NCP protocols. You could describe the protocol in it's functional state, but it's very hard to describe all the issues associated with imbedding it into every one of the computers and making applications run on top of them. So that was the particular expertise that I was looking to Vint for, and when the two of us got together and started talking about it, the issues just started blending and merged and what came out of it was like taking a beater and putting it into a bowl with ingredients and out came something that was a real blend. So the net effect of the TCP/IP protocol was really a joint effort that is really hard to tease apart.

DA: You say when you got together, was this at professional meetings? Did you have conferences? What was the work style that you had?

RK: We would get together on occasion and then we would usually work furiously for a day or, half a day. It would either be a trip that I would make to California, or a trip that he would make to Washington. I recall that we met several times on weekends where he would take a place and we would work at the airport, or at the ARPA office, and we would just furiously hash around on issues. Finally we set aside one period of like two or three days during the summer of 1973, out in California. I remember we went into one of the rooms at what was then the Cabana Hyatt, a hotel right across from Ricky's Hyatt House on El Camino, and we just sat down to write the document that would eventually become the paper that we published on TCP/IP.

At the time we were actually calling it TCP because we had sort of bundled it all together. We hadn't yet conceptually decided to split it up into two pieces. This was something that came about a number of years later, I would guess around 1977. But the functionality was all in this original design that we had had. So the idea was that there would be a protocol header that could be read by the gateways all along the way. That was essentially the IP piece of the protocol, an end-to-end piece that would let the machines at the two ends figure out what to make of all the stuff that was coming in - how to format it going out, how to check some of it, and the like. And that became the TCP protocol.

But it was not a continuous activity. We would work together feverishly and get back together again a few weeks later or a month later, and eventually there was one continuous spurt of two or three days to just sit down and write the whole thing out. I recall that very vividly. Eventually we both went over to Sussex, England, and presented the results at a meeting. I remember we barely had time to get it typed up before we had to go. So the conclusion was that I would go first and get a secretary to do the typing of it, and then he showed up a little late for the meeting with a box of the papers. That was the first time we actually presented it publicly to anybody. That was in September of 1973. It was later published by the IEEE. That was the first time it was published in their *Transactions on Communication* in May of 1974.

DA: The picture I get is of two guys sitting in a hotel room, working on some kind of board or papers together, ordering in pizzas, is that how it was?

RK: Yes, it was kind of like that. I can't remember exactly what was in the room. We had a big table that we were working on and we certainly had a lot of pencil and paper, and I'm sure there was a board there because Vint liked to get up and draw these spider drawings, as he would call them. He liked to think from picture on board, and often times we would have a conversation back and forth and he would say, "Let me draw a picture of that."

For me it was okay to keep it in my head, but for him, I think he liked to actually have it on a board, and it served as a useful tool in communicating because I could see what his idea was and we could work off that. But it was basically the two of us writing it in that room. We literally wrote it to the letter, no personal computers, no workstations, just pencil and paper, well, maybe we used pen.

DA: And through this whole development it was really you two, I know there were people on the fringes, but you did the critical ideas. Did you farm any of it out, or have consultants?

RK: About the only thing I would point out was that at this particular time, this was largely as an outgrowth of the work on packet radio, and Hawaii work before.

The folks at Xerox Park had been working fairly quietly on the development of something that later became known as the Ethernet. They had just started and would later complete a set of protocols of their own. I think they called them "Park Universal Packets" or something.

Those protocols were a competitor, and I know that if it weren't for the fact that it was not done in quite the open fashion that all of the ARPANET was done. It was sort of a timely development by Bob Metcalfe and David Boggs and a number of others.

DA: What factors would you say were critical to the subsequent development of the Internet, from the early protocol development to the kind of things we see now?

RK: Well that's actually a very deep question, in many ways. Let me just give you the context in which we were operating. In 1973, there were no Local Area Nets around. People did not have them. There were no personal computers or workstations around. What you had essentially was a lot of large mainframes and a few mini-computers that were connected, but they were not the workhorses. So in one sense the architecture that we came up with in the early days adapted itself to this new world of personal computers and work stations and Local Area Nets as we now have them, was probably more a result of the nature of the way it was done, because we were not designing it for that in the early days. We did it to deal with independent, autonomous nets and when we had Local Area Nets show up, it just tended to apply to it directly. So we were fortunate in the way we had done that.

I guess I would say there were a few things along the way that were crucial. One was the sustained support by ARPA in the early days - when I was running the programs - and when Vint was running it after me within the office. The second thing was the adoption by the DOD of the TCP/IP protocol as a standard for interconnecting networks. That gave it an imprimatur that I think it badly needed at a time when most everybody else was trying to ignore it because they were assuming that the ISO protocols would somehow take over.

A third thing that helped along the way was the actual split of the ARPANET into two pieces, which caused the initial genesis of the Internet, for real use to take place. And the problem here was, with the ARPANET growing and more and more nodes going on it with the military getting increasingly interested, the feeling was they needed a separate net because they couldn't have something that they were learning to depend upon be operated and maintained by universities. They really wanted more control over it, and they wanted to use it for more military things. So there was a decision to split the two, which meant that we had to make TCP/IP the protocol of choice on the ARPANET to start. And the transition from the NCP protocols to the TCP/IP protocols was actually done on January 1st of 1983.

I had been expecting that Vint would manage that transition, but he left to go to MCI roughly in September, or early October of 1982. So I picked up the Internet program at that time, which Vint had been running for the previous several years. I managed it through the TCP/IP transition and I brought a fellow in to run it, who became very effective on the process side of things, Barry Lyner. We'll get back to Barry in a moment, but I think the split of the ARPANET into a residual ARPANET and then a military portion which was called MILNET, actually created a real user base that depended on these protocols.

Another thing that was crucial in this grand scheme of things was the transition of responsibility for the networking for CSNET to NSF. This occurred shortly after I left. I left ARPA in 1985, September of 1985, and Barry Lyner left around the same time. ARPA had made a conscious decision at that time to stop playing a lead role in networking and transition it to NSF. Eric Bloch had just taken over as the director of the National Science Foundation and was quite interested in taking on that responsibility and propagating the technology to a larger part of the science community.

I think that was a very important step, because NSF's approach to running the network was very different than ARPA's was. You were already in an era now when Local Area Nets could now be procured commercially. They started out in the early 1980's with workstations and LANS, so it made sense to get other parts of the community involved. They could go out and buy these things turnkey. They did not have to be networking wizards to figure out how to do their work. And the approach taken by NSF under the aegis of Steve Wolf, who was a fellow Princetonian graduate when I was back there, was to have a very open approach to networking. His view was, let's spawn lots of nets.

He created the notion of a higher speed backbone, which he called the NSFNet to run at one and a half megabits per second. This was thirty times faster than the ARPANET which was running at 50 kilobits. And even though the ARPANET seemed like very high speed when we started, everybody was saying that was too slow by far. There was a lesson to be learned there. So Steve opened it up, created this thing called the NSFNet. Little regional nets flourished. These were not Bell regionals. These were NSF regional network mainly run by either academic groups or those who had those interests at heart. That was a very critical development to it getting the kind of populism that we now see.

But I wouldn't be fair to this topic if I didn't say that perhaps <u>the</u> key issue of all in the development - that enabled this to happen, was the breakup of the Bell system in 1984. I don't think AT&T saw it as in their corporate interests to have the network's state of communications be based on an open architecture notion that anybody can plug their own private networks into, or commercial networks into, and would all work in this democratic, egalitarian mode that the Internet has evolved to. But with the breakup of the Bell system, and then with alternative long distance carriers coming into play, it was actually possible to explore these issues and other carriers were very interested. It wasn't very surprising that when the NSFNet contract was awarded, that MCI was one of the key players in that, and was one of the essentially new guys on the block at the time. So I think that was quite an important development.

Another important development was the commercialization of the Internet. Very few people realize it, but CNRI was actually responsible for the very first approved commercial use of the Internet. In 1989, we had worked with MCI on a research problem to explore the interaction of incompatible mail systems, that is mail systems that were not designed to work with one another, heterogeneous is perhaps a better term for that. And we had figured out how to make the two work together.

In fact, that was a project that Vint ran because he had been very involved in the development of the MCI Mail system. He designed the MCI Mail system when he went back to MCI. The federal government approved our using that as a commercial application. Well, very shortly thereafter all the other carriers who had similar things were doing it.

The Internet was starting to be more visible at that time; this was late 1989. A Congressman from Virginia named Rick Boucher, introduced a bill, which enabled the National Science Foundation to take a dual use role in its NSFNet for this. Not only would they use it for R&D purposes, but they would also be able to make this network useable for commercial activities as well. They could spawn a commercial activity if they wanted. You would have to get the exact wording of that bill, but basically it freed up NSF to let commercial kinds of things happen.

I guess the final thing that really blew the thing wide open was the introduction of point and click interfaces. There were two point and click interfaces that I was aware of. We actually built one in the late 1980's to work with the National Library of Medicines databases but it was very restricted to that one. With funding from ARPANET and NSF, we funded the NCSA, which is part of the supercomputer center at the University of Illinois, to build a point and click interface to their simulations. The people who were building it didn't actually see the benefit, although they thought they did, of directly connecting it to the web. But one of the fellows down the corridor simply took those ideas and created the Mosaic interface, which was the interface between the point and click world and the web world. That just caused the fury of interest in the Internet like we've never seen since its inception. It was growing at a faster rate than it had ever grown before, the number of official networks, probably close to 100,000 now, and there were domain names close to a quarter of a million and maybe more, by now.

Another thing I should mention about all this activity, because it often comes up and you mentioned it earlier, is the extent to which security considerations plays a role. I'm inclined to say that they played more and less of a role than any of us would have liked. In the early days - I think you have to actually talk to the people at ARPA who were making the decisions about the network - but it was very clear that not only everybody there understood that this technology has powerful potential for the Department of Defense. It had capabilities for survivability and for performance that had no counterpart in circuit switched networks.

On the other hand, I think what was really driving them technologically, was its utility in the community for promoting innovation, for efficiency in the management of their operations. So in fact, when you actually look at the motivations for it, they were many. And clearly the issue was its use in Command and Control and help with the military kinds of things, they were on the list, but that was not the driving force for the research community. In fact it was only many years later that the network was used by the military, when they felt comfortable with it.

Now security was something that was dealt with early on. It wasn't dealt with in a way that was really compatible with the consumer markets. I think public key encryption was one of the important attributes that was later picked up on. Of course this was a development that occurred in the late 1970's in part at Stanford and in part at MIT, and some of that was ARPA funded, and probably some of that was not. But I think public key encryption as a notion was important contribution – the idea that you could use two different keys – one for encryption, and one for decryption - and that you couldn't deduce the private one from the public one. That led to the commercialization that you now see growing on the Internet in recent years.

DA: During this time you dealt with the numerous issues that you just touched on, you personally transitioned from working at ARPA to coming to CNRI. Do you want to say a little bit about CNRI is and why you founded it?

RK: CNRI is a non-profit, R&D organization. It's main focus at the moment is on information infrastructure, the national information infrastructure in particular, but of course it has got tentacles into the global – what Senator and later Vice President (Al) Gore called the "Information Superhighway." When I was running the ARPA office, which supported most of the computer research, it was pretty clear to me that we needed a major effort on infrastructure development in the country. I likened it to the state of the automobile industry at the turn of the century where you could build the best automobile technologically, but it was ultimately constrained by the road system that you had to work with. You needed road maps. You needed stop signs. You needed the protocols. You needed the things that would enable it to really grow. Computing at that early time was not only pretty powerful and expanding, but the actual usage at the time was a small fraction of what the possibilities of use were. With a greater infrastructure in place I thought we would lower the barriers to use of this technology so it really could become an engine for economic growth in the 21st Century.

I really would have preferred to do it in the government, at that time. That's where I was. It was a very natural vehicle. But that was during the Reagan administration and I did not get a lot of support for that as an idea for a government activity. Although I did get a lot support for it as something the private sector wanted to do. In fact many of the words we are now hearing from the Clinton/Gore administration, we were hearing back in those days as well.

So, after thinking about it for a while and talking to a number of people about the possibility, I realized that this idea was too hard to just convey in words. What we needed was some examples of it – some organization that could take a leap doing that. And I thought I was probably in as good a position as anybody to try and make that happen. I had good contacts with industry. I had good contacts with the academic community and of course, I had good contacts with the government, coming from there. So the idea with CNRI was to play a leadership role in helping to foster this idea of a national infrastructure through R&D types of projects, and by working with all of the parties, to sort of bring them together in a sort of coordinated fashion.

We got our initial support for this activity from industry. Some of the computer companies were first off the block to help support us; places like Digital Equipment and Xerox and IBM. But shortly thereafter some of the communication firms were fairly quick off the block to help us as well; places like MCI and Bell Atlantic.

We have gotten support over the years from many more organizations. In fact we currently work with something called the, "Cross-industry Working Team," which is a collection of some 50 or so high tech companies and publishers and banking industry, semiconductor industry, wireless – virtually all aspects of related things who work together in understanding language, the terminology, getting commonality on architecture issues, a common understandings of how to proceed in the future. That's been pretty effective so far. So we've been getting support from a lot of companies right now with this work.

Over the years CNRI has played a pretty important role, I think, in furthering the cause of infrastructure development in the country. One of the things that we have done over the course of the last decade is continue to maintain the process by which the Internet evolves. We started that during our days at APRA, myself and Vint and Barry Lyner – I want to come back to Barry because he played an important role in this whole process. But also subsequently after CNRI got started, by virtue of running the whole secreteria of the IETF, which is the body that is generating the protocol, or the grist for the protocol mill as it were. The IETF essentially adopts the standards and we also worked with a number of organizations to create another non-profit, which we housed internally in CNRI for a while, and then later formally incorporated and spun out, called the "Internet Society." That was set up really to be the home for the Internet process, in terms of it being the home for the protocols, and essentially be able to apply the "Good Housekeeping" seal of approval, and in general being the information arm for the uses of the Internet for some of the providers, as to what is going on in the Internet process.

When we got started with CNRI, we actually had three notions that we were going to pursue, one of which we really pushed very hard. That was the idea of a national digital library. The idea we had at the time was to use active agents in the library, we call them "knowbot programs." These would essentially be mobile pieces of code that would go through the network on behalf of the user. So unlike point and click where you have to see what you are wanting and you point and click on it, these, you would describe to it what you wanted and they would contain enough expertise inside these knowbot programs to go to various sites on the Internet and execute there appropriately. We had some prototype systems that were developed along the way. We continue to work on that, but the whole idea of digital libraries has really expanded in a significant fashion. We've gotten into some implementations. We built prototype systems for the Copyright Office of the United States, where they can take submissions across the Internet digitally and actually do registration of copyright, intellectual property that is submitted. The registration system doesn't have to keep the digital objects, what they keep is unique identifiers for those objects, and the fingerprints of those objects. It's actually quite a nice activity.

We're also working with the Library of Congress on broader digitization of their collections using an architecture that we had originally conceived of, but was further refined through interactions with the community, and I would say particularly through interaction I had over a three year period with a fellow named Bob Wilensky, who was the Chair of the Computer Science Department at the University of California at Berkeley. Many attributes of the Internet architecture that I developed back in the early 1970's for interfacing between heterogeneous library systems, it's a more complicated problem there because there is no one single objective for the library that's as best efforts simple as packet delivery, but it's been a very interesting development. We are currently implementing many of the component pieces of the infrastructure to make available to the community. We'll see how that proceeds.

DA: The digital library was one initiative; you mentioned there were others major initiatives at the time.

RK: The other ones we were interested in at the time didn't get too far. One was to try and work with the community in building codified knowledge about things. Working on the content side of things where we would capture in computerized form, knowledge about certain areas that were critical; like how to build a jet engine, or how to build a refrigerator, or laws of Physics, to whatever. But I think the technology was premature at that time, and the support was just not there. So we put it on the back burner. We hope to get back into it. Meanwhile ARPA has been supporting a lot of interesting activity in the technology part of that, at a low level, but it will probably incubate up pretty soon.

The third one is an area that has actually got quite a bit of interest in the community, and that is the use of the network for handling online electronic transactions. We're seeing quite a bit of that in the electronic commerce domain. We have some ideas that we wanted to pursue that are a little beyond what you are seeing now in the buying and selling aspect of things, but it's an area that we continue to be interested in on working on globally. We're collaborating with two other groups; Financial Services Technology Consortium, and CommerceNet, on electronic payments form. So we're slowly getting into those kinds of things, but I think mainly our efforts have been in the digital library area over the last decade. Plus of course, our major efforts in the networking area, one of which I should mention, which is support to the whole Internet process.

The other one was a major effort that we undertook over a number of years period to actually give a major portion of the field, a very high-speed network. We talked abut 50 kilobits as being a high speed, and then one and a half megabits as being high speed, but here we're talking about leveraging the possibility of fiber optics running at speeds of anywhere from 622 megabits per second, all the way up to point-four gigabits per second. And we got support from ARPA and NSF to do that.

We set up five test nets in the United States with support from a lot of people in the community. We have like 40 different organizations involved with us; major industrial players, major academic players, and an enormous amount of contribution from the carriers who donated the fiber optics, which even today we couldn't afford to buy if we had to.

Not only did we actually get those test nets built in five different parts of the country, but the spans ranged from 30 miles to 1-thousand miles, just to give you the geographic spans. The networks were used both to understand architecture issues - these networks helped the technology get built and pushed forward - and also to understand the kinds of applications that might be appropriate, at least within the scientific community. We never began to tap what the industrial sector might do, just beyond R&D, and that's a whole area I hope can get back into one of these days.

The demand for high-speed networking has grown significantly since we did that project. And although the test beds will not all survive in their original form, they were all different and one of a kind, my guess is that the demand of high speed networking will lead to gigabit nets on a national scale. Initially you will see three, which is not quite gigabit, but eventually you will see twelve, you'll see forty-eight and higher, and I think that's really right around the corner at this point. I think we can take a lot of credit for pushing the technology. We did not invent ATM, the community did all that, but we helped to accelerate that by working with them, and we hope to do that in the future.

DA: A lot of people see the Internet from the outside as being chaotic in how it's growing and developing. You were involved in creating the original protocols. Does it look that way to you? Has it become what you expected?

RK: It's interesting how some of the most important things in the world are not centrally managed. Like government and the world economy, and yet somehow they seem to manage. People might look at the world economy and say its chaotic, yet the economy goes on and it doesn't collapse because there is no central manager.

In the case of the Internet, what I would focus on is the process by which it evolves. The architectural framework, which we developed in the early days, was sufficient technically to let you understand how you put things together and how things could communicate. But what was really more fundamental to this process in the long term was how you could ever cause those things to change and improve and adapt and evolve over time, particularly if you got to the point where you had thousands, tens of thousands, even millions of users, hosts and networks. We're there right now. There many millions of hosts, hundred thousands of networks formally registered on the Internet.

Key to this development early on was some work that Barry Lyner had done. I mentioned him earlier as the person who took over for me in the late 1983, roughly. Barry took a look at the process which we had been using at that point, and decided that it was really not workable for the long term.

I had become very concerned in the late 1970's when Vint had taken over the Internet program and it was still continuing to grow, that if anything happened to him, I would have to come in and be there full time. And I had too many other things on my plate at the time to attend to. So he and I talked about what it would take to get more people from the community involved in the decision-making that was going on. Vint set up something called the Internet Configuration Control Board, better known as the ICCB, which consisted of about twelve key implementers from the community. They would meet and he would work with them convey to them what he was thinking and doing, what they were up to. And over time the discussions in those groups, which were often in conjunction with some of the packet satellite meetings or the radio meetings, whatever – the discussions got to be so interesting that many other people asked if they could sit in.

So at the end of 1982, when I had taken it back again, a meeting of the ICCB would often have a very large number of people sitting around listening to the debate. I would guess it could have been a hundred or more. It was very difficult to call a meeting if you had to coordinate it with hundreds. When Barry Lyner took over he said that just didn't work for him. So he reconstituted this core group, I think it might have been slightly augmented at the time, but basically about the same twelve people, and he called it the Internet Activities Board. He created underneath it twelve task forces, each one with a Chair to look at certain considerations of the time, and protocols or autonomous systems, or security. There was even a group in there called, "Internet Engineering," which was intended to focus on some of the nearer term requirements of the Internet.

So this was all set up roughly in the late 1983 time frame when Barry took over. And over time, this turned out to be a marvelous managerial stroke, because over time, as more and more people wanted to get involved, they didn't have to get involved in the IAB, which could act as a little deliberative body, but they could create more task forces and the task forces could meet anywhere and not everybody had to come to all of them. There was a certain stroke of autonomy that got injected here.

Over time the IAB took on - and this was Barry's preference as how to operate – more responsibility for essentially approving the standards. The task forces would generate the grist, and the mill was the IAB, and the IAB would essentially approve the standards. We would go through a set of tracks; something would be identified as just an informational item, and it would be sent out through the same mechanism that the ARPANET originally used, this request for comments series that was used to propagate information out of the ARPANET, was now used to propagate information about the Internet, which really overlaid itself on top of the ARPANET.

The task forces were able to meet independently, but their output went to the IAB for standardization. In fact, it was a grass roots standards process. Things really didn't get on a standards track unless they were in widespread use, and once they were on the standards track, they were in various stages of approval until eventually they would become official protocols. But by the time they were approved, they were in widespread use.

So it was the exact inverse of the process by which, let's say the ISO, where they would declare the protocol as a standard, and then hope people would implement this.

DA: So this was really bottoms-up.

RK: Bottoms-up instead of top down.

DA: That leads to a question about the role of the United States. So much of this early Internet development happened in the United States, and was driven by a relatively small number of individuals and organizations. Will this American leadership, or domination by the United States, will that continue?

RK: Well, I would argue that it isn't really dominated by the United States, and it hasn't been for a long time. The interesting thing about that process is that it empowered people in the community to basically take charge of what goes on. Anybody can participate in the work groups of the IETF. In fact, what I didn't say is that when the number of working groups started to grow, the feeling was that the IAB was no longer the right body to control them. So one of the work groups, called the Internet Engineering Task Force, was designated as the body to coordinate all the work of the taskforces. So inside the Internet Engineering Task forces was a task force of task forces.

The IETF actually does all the work. In the early days it provided grist for the mill of the IAB. But after the Internet Society was formed, and after a lot of discussion in the community, the IAB proposed that the Internet Society be created to accept the IAB as a part of it. This has cause a split between the IAB and the IETF, which I guess had not been consulted in the process. In order to sort of assuage the relationship between the two, the IAB gave its powers of standards-making to the leadership of the IETF.

So today the IETF in fact, is the body that not only does the work of developing the options, but it also approves the standards. The Internet Society is intended to be the housing body for those standards as they evolve.

When you look at the role of the US government in this whole thing, originally it was deep and central with DARPA, and later gets propagated widely by the NSF taking a clear leading role. Today the role of the US government in the Internet is barely discernible. I mean, it occurs in the form of a little bit of support to CNRI and other bodies to help them with some of the process related things. For example, CNRI stills runs the IETF secretariat, with a little support from the federal government, although a lot of that is offset by attendance fees paid by people who come to the meetings. The US government has never weighed in to the workings of the IETF in the time that we have been involved. In fact their instructions have been to facilitate and support it.

When CNRI started it, of course the chair of the IETF worked here. At one point when Vint became chair of the IAB, he was here also. So we actually have leadership of it all working here as well as supporting it. So there was always the question of what hat are you wearing at any one point in time. CNRI was set up as a place that could make this happen, could facilitate this to happen in the community at large. So when the head of the IETF went to work for ANS later on, we could facilitate that and support that. We simply made it possible for them to continue and to elect whom they wanted to follow on in their footsteps.

That's a role that I think we're happy to play as long as it's necessary. If there comes a point which may very well occur over the next few years, when it's no longer necessary for US government support to go into the hopper, then the US government won't have any of those tentacles involved. They have supported John Postel as, what we call the Internet Assigned Numbers Authority, and that was by delegation. We had actually delegated that role to John back at the time we were at ARPA, for assigning numbers in the ARPANET context. We had quite a different space of participants then than we now do, but that's also grown over time. They (the US) support a group called Network Solutions, which is now part of SAIC, to handle the domain name registration for a certain number of names. But even then it's a decentralized function when you realize that the registrations in other countries around the world are often handled just uniquely in those countries.

DA: So this is a global phenomenon.

RK: It really is a global phenomenon, and where the US government has been involved in the Internet in a direct way, it's been early on, very tight, and in recent years much higher level. And ever since the NSFNet was decommissioned, they were really not in the commissioning of network services anymore, although they still do and will probably continue to provide advanced research networking capabilities or one sort or another, that go beyond the commercial Internet as we now know it is providing.

DA: Have you been surprised at which the pace the Internet has been used for electronic commerce? That seems to be the drumbeat which people keep hearing, and yet it seems to be developing relatively slowly, at least compared to some of the hype that was generated.

RK: Some people might say that it is growing faster than anything else they've seen in history. On the other hand, I think it's again, the infrastructure issue. It is still inhibited by the kind of infrastructure that you have in place. Therefore it probably is - you're probably right - a tiny fraction of what is ultimately going to be possible. You heard that old expression, "You ain't seen nothing yet." I suspect that is true, and what even exists in two to five years, may be radically different from what we now see.

Suffice it to say that I hope that the Internet doesn't become only a medium for commerce, in the sense of production, consumption, buying and selling, where people have something to offer and you click on it and send in your credit card number and take delivery of something, maybe even through the US mail. Of course you would have to figure out how to send a pizza through the Internet, but you can buy a pizza through the Internet.

The whole evolution of this as a medium for collaboration and interaction among people is one that has hardly been explored. You know you can use it for email interchange, and there are a variety of early examples of early applications that involve groups that participate and collaborate, but I think that's the real opportunity space. Because somewhere between the end of the spectrum that says the Internet is simply a distribution system, actually a dissemination system for information, and the Internet is a replacement for telephone calls, where in the telephone case, nobody cares about intellectual property that you want to retrieve, you are creating it on the fly as you do it. There's some big, wide intermediate space here where people are going to work with each other to create and maybe store, manipulate and modify and make available to others, the results of their collaboration.

DA: Tell me what in your view, what the use of the Internet looks like let's say, 10 years from now. What's different?

RK: Well, I don't spend most of my time buying things. I spend a lot of my time thinking about things, and if the Internet can help me do a better job of thinking and planning, that would be very helpful. I spend a lot of my time interacting with other people on planning and scheduling things, and if it could help with that, that would be really great. I spend a lot of my time going to meetings, and if the Internet could facilitate making available those things, or having me participate and make me feel like I'm really there without the need to get on a plane and travel or wait on the tarmac for ten hours, that would be fine. And I suspect we're going to see most of that start to come into fruition.

Just like you can say, "what's different about the telephone system now than it was 50 years ago," the answer is at some level, "not much." Yes, you can do direct distance dialing instead of using an operator. Yes, the quality may be a little better, and maybe the phones look a little different and you get more choices, maybe even higher or lower bills. The fact of the matter is that there has been an enormous change within the telephone system in 50 years. The technology now bears no resemblance to what it was then; the potential and capability bears no resemblance. And that's an application that is very one-dimensional; individuals talking to individuals. I think it's going to fan out.

I think the Internet dimensionality along those lines, is just going to keep pace with the proliferation of possibilities. You've got, instead of one organization trying to evolve something; namely AT&T, which did a marvelous job for a long time; you now have everybody around the world who has a conceivable interest able to contribute, both of the level of the technology, the equipment, the service, the applications, as well as the use.

And it's all driven by this marriage of the computer and communications technology, and my guess is, it will amaze us all. Any linear projection is bound to be wrong diverted by all the normal lineararities and innovations that come along that we never though about before.

DA: Do you have any concerns about things coming ahead?

RK: I guess my biggest concern is about society itself. The biggest danger to man is probably man himself or mankind I guess. I don't see the Internet as being a larger threat to mankind than people are. People still have the ability to do wondrous things for their fellow citizens, or not as the case may be. Often times we enter into blind alleys. We do things we just didn't know enough not to do.

I worry a lot that the concern today about limiting the ability of materials to be made available on the Internet, very much going to interplaying with notions of free speech in the United States. Today the notions of transmitting something over the Internet, really doesn't apply in most of the cases. You have a mechanism now, which makes it very easy now for someone to go in and reach out and grab something on the Internet. So the fact that you wrote something or did something that is determined to be obscene by somebody's judgment, doesn't mean you transmitted it because you happened to make it – someone else had to go in and get it with the appropriate devices to filter and make clear what you're getting it. It seems to me that is a manageable proposition. Whatever is on the Internet or cyberspace, or whatever you want to call the information superhighway regime of the future, you've got stuff in the real world today that isn't digitized, isn't electronic, and some that is, that ought to be equally troublesome to the people that are concerned about the Internet. Whatever strategy you adopt ought to be something that will apply in the other case as well. I can't see people banning certain kinds of activities in this country which they might not want to see in their communities. It's protected by free speech, and they ought to do the same thing on the Internet that they would do in their own community if they became offended. I think the Internet, much like I think of the airways for speech. It really is that kind of standard.

DA: From a technical standpoint, do you feel the technology is going to hold up to this growing demand and load on the system?

RK: Sure. The fact that it's commercial just guarantees it to me. If you have a case where there is some much demand for something that people can't provide it quickly enough, they're going to get set up to provide it more quickly. The problem is when people get overwhelmed and they don't realize that the demand is there. There is money to be had in service provision, somebody is going to be out there to provide it, that's the essence of the marketplace and I believe that.

DA: Any other thoughts you want to get down while we are here today?

RK: I think one of the questions I get asked all the time is, whether what happened with the Internet was anticipated. When this work began back in the early 1970's, to anticipate this growth, you would have had to anticipate the break up of the Bell system. You would have had to anticipate the transition to the NSF. You would have had to anticipate the creations of LANS and workstations and the like. You would have had to anticipate so many things that would have been too incredible. To say you could have anticipated all of that would have been too incredible.

On the other hand, one could have anticipated back then that you could have had a large number of networks of kinds yet to be determined, all interconnected with computers interchanging interesting stuff – of forms yet to be determined – plus some old stuff, like emails and files and the like. You could have imagined that you needed a process to manage it. You could have imagined that the carriers eventually would have gotten interested if there was an opportunity for them to do this in a way that would lead them to a business. Much of what actually happened you could actually foresee, but to talk about the specific events that happened, no. I think the biggest uncertainty was, would there be enough support to continue this on for its direction?

Vint and I remain the custodians for an award that was given out in 1993 to the Internet as "Product of the Year." It was probably the first time the product of the year was 20 years old. The fact of the matter is that that's when people woke up to it. That's when they were first aware of it, and I think that's something I could not have foreseen. That the continued support would have been there for that 20 year period to allow it to pop up above the surface.

The other thing that occurs to me is that this is a field that is still in its infancy. For all the people that think that the Internet has matured and that the information age is done, I keep thinking of the people that thought that everything about data communications was wrapped up in the 1960's, or that physics had done everything that it needed to do back in the 1920's and there was nothing left to discover. I think if you look back 100 or even 50 years from now, this will be seen as the really early pioneering days. The days when every television set is plugged into the Internet, and the broadcasts aren't coming over cable or over the air, that's probably only a few years away. In today's paper there was an announcement that Zenith is going to make televisions that are Internet compatible. Well, that's a matter of today and next year and the year after.

If you go out 10 years, who knows what's going to be in your hand held terminal? It's going to have the ability to have more computation in that terminal, and more ability to interact with other computational resources on the Net, that just the whole notion of how we relate to computing may change dramatically.

So we may not see any more work stations occupying space on your desk and you might not see the big operation systems that we now have, you may have other mechanisms. More on the lines of network-based functionality that really define what the future is like, and it may be all invisible. It may be in your glasses in your clothing and the memory in your shoes and who knows how that all works. I'm sure even those observations will look silly in 50 or 100 years because we've got other more majestic ways to think about it that are just not within our scope of comprehension at this point in time. So it's still pretty early.

DA: That's a great way to end our conversation. Thank you so much for your time and observations. What a pleasure!

RK: Thank you.