An Interview with

PAUL BARAN

OH 182

Conducted by Judy O'Neill

on

5 March 1990

Menlo Park, CA

Charles Babbage Institute
Center for the History of Information Processing
University of Minnesota, Minneapolis

Paul Baran Interview
5 March 1990

Abstract

After a brief review of his education and work experience at the Eckert-Mauchly company, Raymond Rosen Engineering, and Hughes Aircraft, Baran describes his working environment at RAND, as well as his initial interest in survivable communications. He then goes on to describe the evolution of his plan for distributed networks, the objections he received, the writing and distribution of his eleven-volume work, "On Distributed Communications," and his decision against implementation of the network in 1966. Baran also touches on his interaction with the later group at ARPA who were responsible for the development of the ARPANET, and the cumulative nature of the inventive process. Baran refers to seven supporting documents during the interview. These documents are not included with the interview transcript, but photocopies are available from CBI. This interview was recorded as part of a research project on the influence of the Defense Advanced Research Projects Agency (DARPA) on the development of computer science in the United States.

PAUL BARAN INTERVIEW

DATE: 5 March 1990 INTERVIEWER: Judy O'Neill

LOCATION: Menlo Park, CA

O'NEILL: Why don't we start out with your background before starting on the RAND reports. Your education, work experience, that sort of thing.

BARAN: I attended Drexel Institute of Technology in Philadelphia and left there in 1949 with a BS in Electrical Engineering.

O'NEILL: Were you involved in any computer projects at the time?

BARAN: There were no computers at Drexel at that time. Or any real ones anywhere. My first job was with the Eckert-Mauchly Computer Company (which later became, Univac, from the name of the first commercial computer that they were building there).

O'NEILL: How did you start working at Eckert-Mauchly?

BARAN: I saw an ad in the paper. It was a job, and I needed a job.

O'NEILL: So you didn't have any particular interest in the company?

BARAN: No. Jobs were scarce at that time. As this was my first job out of college, I felt that I couldn't be too picky when there was a surplus of recently graduated engineers. I was a technician performing life tests on racks of components.

O'NEILL: What sort of components?

BARAN: The components were radio tubes, germanium diodes, resistors, etc. The germanium diodes were most interesting. No two were the same. I put together a device to measure each diode, so that they could be sorted into piles. They would use the diodes with low back resistance only in certain parts of the decoding circuits. We punched an extra blank hole in each aluminum chassis strip to accept a 6H6 vacuum tube just in case semiconductor diodes didn't work out.

I really didn't understand how computers worked and heard about a course being given at the University of Pennsylvania. I was a week late and missed the first lesson. Nothing much is usually covered in the first session anyway, so I felt it okay to show up for the *second* lecture in Boolean algebra. Big mistake. The instructor went up to the blackboard and wrote 1+1=0. I looked around the room waiting for someone to correct his atrocious arithmetic. No one did. So I figured out that I may be missing something here, and didn't go back.

O'NEILL: What convinced you that there was no future in computers?

BARAN: There were a number of factors. The main backer of the company had died in an airplane accident so the company was for sale. There was a sense of desperation. I recall once having to stop what I was doing, put on coveralls, and go into the machine shop to punch holes in chassis. Visitors would see people working cutting sheet metal rather than taking measurements on racks of components. External cosmetics don't breed internal confidence.

This lack of confidence was accentuated by the mean-time-to-failure estimates. Batches of components in the testing racks were failing too frequently for comfort. What was disconcerting, aside from the company running out of money and bringing prospective buyers through the place to see if they would buy the company, was that the mean-time-to-failure estimates were dismal. Sample components in the testing racks were failing too frequently. Even worse, there was an unexplained malady occurring called "sleeping sickness." If a vacuum tube was in the "0" state for a long time, it would on occasion apparently miss the first "1" that came by. After the glitch, it would go back to working fine.

Given the failure rates that I saw, I agreed with the common wisdom that these computers should be expected to have only a short useful period between failures. It wasn't exactly a confidence-building exercise.

What triggered me to leave there was a brown bag lunch with some colleagues. I was quietly munching on a (salami with cucumber on pumpernickel) sandwich taking it all in as the older and wiser engineers discussed the company and its future. I think it was Jerry Smollier "the power supply guy" who thought that the company had orders for three machines. "Okay" someone said, "let's assume that we get orders for another three - six computers. Okay, let's double it again." "What sort of business is it if your market is only a dozen machines?"

The logic seemed impeccable. If it took a room of people to program a simple problem for the computer, how many organizations could possibly need a machine to solve the exact same problem over and over again? Hmm.[1]

O'NEILL: Where did you go after you left Eckert-Mauchly?

BARAN: Having concluded that there was no future in computers, I joined the Raymond Rosen Engineering Products Co. in 1950. Raymond Rosen was a major RCA TV distributor. TV was growing rapidly at that time. Along with their TV franchise, RREP serviced RCA police and taxicab radios. A year earlier RREP had modified a radio and sold it to

[1]In retrospect a number of major "breakthroughs" occurred, the next several of which made nonsense out of my dire conclusions about the viability of computers. Example: marginal testing, where the voltage to each side of a flip-flop is varied to detect parts failing before they failed. Another was the discovery that the sleeping sickness was being caused by poisoning of the cathode coating and could be corrected. The idea of programming languages came along.

This experience taught me never again to ever underestimate what technology could do -- with enough work and dedication.

the White Sands Proving Ground. It worked. Now the company was in a small but rapidly growing telemetering business conducted in the corner of a warehouse building full of TV sets.

O'NEILL: What did you do there?

BARAN: I designed circuitry to correct tape speed errors when recording FM telemetering signals onto magnetic tape. I found that the corrector's performance was limited by unexplainable FM tape noise. Some batches of a tape were worse than others. I guessed that the tape was acting like a bow across a violin string. Three companies were making magnetic tape at that time, but only 3M was agreeable in working with us to help solve the mystery. The problem was tracked down to conical clumps of magnetic material. 3M then scraped these off with a knife blade in the manufacturing process to create what we called "deteatified tape." The term was used until marketeers at 3M realized where the name came from. They renamed it "Instrumentation Tape." Instrumentation tape later became a major market for them as data users record on tape only once saving the tape forever, unlike audio tape users who reused tape, so the market was much larger than anyone realized. Another lesson learned.

RREP soon became the leader in this new field of radio telemetering. RREP received a request for a telemetering system and a remote recording system. But the new system had to be up and working within two months to test the first Matador (long distance air breathing) missile. It was a worthy challenge. To meet the schedule, we had to gather up our laboratory breadboards and take them down to a beach house at the brand new Cape Canaveral installation. We worked like hell. The launch date came, and all our gear worked perfectly. The government people were delighted at seeing what radio telemetering could do using our equipment. I came back with a big pile of potential orders in my pocket and assumed that my boss would be happy.

I didn't appreciate the complexities of the business world. The older Mr. Raymond Rosen, the sole proprietor of the business, listened to my story quietly.

"So boychuck, so where are we going to get all the people we need to do the contracts?"

"We have to hire them."

"So, why will they work here?"

"We will have to pay more than they are now making."

"So, the guys fixing police and taxicab radios down the hall see this and come running to me crying for more money. If I give it to them I won't be able to compete."

"Gee, but the opportunity is huge. Telemetering and remote control are going to become big businesses one day."

"Telemetering, schemetering -- a flash in the pan! It will go away tomorrow. Now taxicabs and police radios I know. They are going to be around a long, long time after all this government electronics stuff disappears..."

I thought that the company was blowing away the opportunity of a lifetime to become a leading company in what would be a large industry. But, then again, he was rich, so he must know what he was doing. Clearly I didn't understand business like the old man did. Nevertheless I thought he was wrong, and I decided not to stick around and try to do a job with that level of management's commitment.

O'NEILL: What did you do next?

BARAN: I spent a year doing some independent consulting. And, then while living in the East got married to a woman ("girl" in those days) who was from California. Naturally, we had to settle in California. I sent off some resumes to California companies that were recruiting in the East. If the job worked out, great. But, if not, the consolation prize was free transportation.

I interviewed with a group from Hughes Aircraft, with what was then the Ground Systems Department. I thought that it was an odd name. The primary business of the company was Air Force airplane electronics. The Department had about 22 people at the time. It was called "Ground Systems" because of the department head, a guy named Barlow. His idea was that as a new activity he would automatically receive all requests for proposals that came through the front door that were not specifically for airborne systems. His scheme worked beautifully. Ground Systems has since become a major part of the entire Hughes Aircraft operation.

Initially I worked in the Systems Group on radar data processing systems; later, in Studies and Analysis. As a system's engineer I was specifying transistor subsystems. But I hadn't the foggiest notion how and why a transistor worked. I had been away from school six years and the technology had just whizzed by. I still really don't know how a transistor works. -- Well, really I do know; it's just that I find myself unable to believe those electrons and holes can be so smart, while people are so dumb. I thought I had better find out what was going on, so I started taking extension classes at UCLA at night.

O'NEILL: Were you working on your degree at UCLA, your master's degree in computing, while working at Hughes?

BARAN: I didn't set out to get a degree. It was a case of wanting to know how transistors worked and to be able to use them in circuits. Taking one course led to another. I wanted to take a computer course that would explain to me why 1+1=0. With a course here and one there I had much of the course requirement for an M.S. degree out of the way. I was lucky and encountered several super teachers in this UCLA night school while working at Hughes. For example, Montgomery Phister teaching logic design, Willis Ware -- a super communicator able to explain the underlying concepts of digital computers and, later, Jerry Estrin.

When I decided to formally go for the M.S. degree I had the remarkable good luck in drawing Jerry Estrin as my advisor, and later my thesis advisor. Jerry had recently arrived at UCLA from The Institute of Advanced Studies at Princeton. Being new he wasn't yet inundated with students clamoring for his time so I was able to work with him on a one-to-one basis on my thesis -- which was on adaptive character reading. He would read the same literature and chide me if he found something I hadn't read. He kept me continually challenged. He has a wonderful way of finding out what you knew and what you didn't. He would gently, but firmly, focus you into your weakest areas. At this time, I was working at RAND. Jerry Estrin convinced me that since I had so many courses out of the way, why not go on for a Ph.D. RAND allowed me to take time off during the day, but business travel was increasing at the time causing me to have to miss more and more lectures. But, the final decision was made one day when I drove in to UCLA from RAND and couldn't find a single parking spot in all of UCLA and the entire adjacent town of Westwood. At that instant I concluded that it was God's will that I should discontinue school. Why else would He have found it necessary to fill up all the parking lots at that exact instant?

In retrospect, the only way that not obtaining a Ph.D. degree has affected my life is that when I'm at home and the telephone rings and someone asks for Dr. Baran, I have to hand the phone to my wife and say "It's for you dear". (Evelyn has a Ph.D. in Economics from Harvard, so she is the better educated one in the family. One of her endearing qualities is that she doesn't remind me about it too often, particularly when she was teaching money and banking at UCLA while I was a student.)

O'NEILL: You started at RAND in...

THE RAND CORPORATION

BARAN: I started at RAND in 1959, after Hughes which was 1955.

O'NEILL: When you started at RAND, were you in the Computer Sciences Department?

BARAN: Yes, the Computer Sciences Department of the Mathematics Division.

O'NEILL: Were your initial interests in communication systems?

BARAN: Not directly. RAND was a most unusual institution in those days. If you were able to earn a level of credibility from your colleagues you could, to a major degree, work on what you wanted to. The subject areas chosen were those to which you thought you could make the greatest contribution.

This period was the height of the cold war. Both the US and USSR were building hair trigger nuclear ballistic missile systems. The early missile control systems were not physically robust. Thus, there was a dangerous temptation for either party to misunderstand the actions of the other and fire first. If the strategic weapons command and control systems could be more survivable, then the country's retaliatory capability could better allow it to withstand an attack and still function; a more stable position. But, this was not a wholly feasible concept because long distance communications networks at that time were extremely vulnerable and not able to survive attack. That was the issue. Here a most dangerous situation was created by the lack of a survivable communication system. That, in brief, was my interest in the challenge of building more survivable networks.

O'NEILL: How did RAND determine which problems it would work on?

BARAN: RAND at that time worked almost exclusively for the Air Force. The Air Force would fund RAND once a year. RAND had what today would be considered a remarkable amount of freedom on how it would spend its money. However, there were requests made by the Air Force to study specific subjects that RAND felt should be performed to be responsible.

Every week RAND management circulated the letters from government agencies requesting help on various projects. If a staff member was interested and could make time he or she could sign up to work on the project. Otherwise, if no one volunteered, RAND would respond with a form letter saying something like "Thank you for your interest but after review the subject was determined to be inappropriate to RAND".

RAND also conducted a number of formal projects, generally of a support nature for the Air Force. Here the gridlines were narrow and the staff was directed in usual project fashion. These support projects were regarded by RAND management as being necessary to preserve the Air Force's goodwill to allow other parts of the organization greater freedom. RAND also had an internal support function, which of course was directed. What made it all work so effectively was a very intelligent management that understood the research process.

O'NEILL: Did your survivable networks project start with a letter from the Air Force?

BARAN: Not directly. It was a well known problem. I had either already convinced myself that it was a solvable problem before I came to RAND, or I came to that position shortly after arriving at RAND in late 1959.

RAND had a very effective internal, open communication system. Staff members had ready access to knowledge of critical defense issues. RAND was a multi-disciplinary organization that permitted broad gauged insight and understanding of national strategic issues and their relative importance. By way of history, RAND was initially set up in 1946 to preserve the nation's operations research capability developed in World War II.

O'NEILL: Did you have a lot of interaction with the people in the Communications Division at RAND?

BARAN: No. This developed over time. I must confess that initially I got off to a ragged start in my relations with the Communications Department. The Communications Department had very competent radar and analog transmission people and great experts in satellite communications. But it lacked computer technology based skills. The stuff I was writing was pretty far out and they would tend to get a little huffy about it, at least at first. The opposition went away with time.

"What in hell is this new guy in the Computer Science Department mucking about in communications?" Generally, RAND didn't suffer from turf problems, but here was an exception. The Computer Science Department was primarily a programming support arm at that time. The electronics people in other departments tended to have a radar or an analog communication background. They did not understand, nor fully appreciate, what was happening in the world of digital computer technology and what this technology could do in the future. For example, their mental image of a computer would be that of a huge installation, when I spoke of file cabinet size units. Thus, many of the things I thought possible would tend to sound like utter nonsense, or impractical -- depending on their generosity of spirit for those brought up in an earlier world.

Computers and communications were at that time two totally different fields. It was difficult to talk about error free transmission to experts in analog transmission who had no appreciation of what digital technology might be able to do. It wasn't just the RAND Communication Department. Initially, when I went outside of RAND I found the mental set even worse.

These were two completely different worlds, communications which were analog and computers which of course were digital. The underlying concepts were different. When talking about what digital technology could do for survivability I found the major challenge to be analog communications engineers unable to grasp new concepts.

O'NEILL: So you had an interest in survivability, or second response to a first attack?

BARAN: Yes, primarily the survivability of strategic command and control communications.

O'NEILL: That was from early on at your time at RAND, so approximately 1959.

BARAN: Yes. It may be helpful to understand that time period.

1. The US and the USSR distrusted one another.

2. Each superpower regarded the other as a potential enemy having the capability of mounting a surprise attack.[2]

3. The physical vulnerability of the strategic weapons systems at that time created a doctrine in which the counter attack would be launched immediately on detection of an attack.[3]

4. Humans and defense systems are highly prone to errors.

5. It's too easy to accidentally misread the signals in an environment of mutual distrust and paranoia.

6. If we could wait until after attack rather than having to respond too quickly under pressure then the world would be a more stable place. This meant that enough strategic forces would have to survive the first attack to return the unfriendly act.

7. This meant as a minimum hardening missile sites and dispersing aircraft, both of which can be done.

8. To coordinate the response, to surrender[4] or to accept a surrender to stop a bloody war required a survivable network. But we didn't know how to build survivable networks.

9. Our hypothesis was that somehow it should be possible to build a survivable system. Learning how this objective could be accomplished was my personal motivating interest. Almost all my work in this field was completed in the

---

[2]Throughout history successful generals make their plans based upon enemy capabilities and not intent.

[3]The other day I heard a revisionistic view of history of that period that claimed that the US was considering attacking the USSR in a first strike. This is nonsense. I never heard anyone in strategic planning ever propose a first strike against the USSR. Rather the temptation, or the necessity to go first, was recognized by the fear brought on by the differential effectiveness of the USSR's ability to keep secrets that might tempt them to pull off a successful first strike.

The focus of all those that I knew were concerned about the nation's defense was on avoidance of war. I never encountered anyone who deserved the Dr. Strangelove war monger image so often unfairly ascribed to the military by fiction writers in the late 1960s.

[4]The issue of surrender was never formally discussed at RAND. RAND was precluded from doing so by an act of Congress. There was a clear but not formally stated understanding that a survivable communications network is needed to stop, as well as to help avoid, a war. But this rationale was generally stated only implicitly because of this quaint Congressional edict.

1960-62 time period. It was not done out of intellectual curiosity or a desire to write papers. It was not done in response to a work statement. It was done in response to a most dangerous situation that existed. After 1962 it took another year or two more to respond to the many objections raised, first orally and then in writing.

The major defense objective for the last 40 years has been to prevent the world's superpowers from stumbling into World War III. It was an extremely expensive standoff. But, in retrospect, it worked. The threat of an attack by the USSR today is no longer regarded as a realistic concern. But, in 1960 it was an entirely different environment. The great communications need of the time was a survivable communications capacity that could broadcast a single teletypewriter channel. The term used to describe this need was "minimal essential communications," a euphemism for the President to be able to say "You are authorized to fire your weapons". Or "hold your fire". These are very short messages. The initial strategic concept at that time was if you can build a communications system that could survive and transmit such short messages, that is all that is needed.

My first scheme was a proposed communications network based on an idea initially proposed many years earlier by Frank Collbohm, President of RAND. Collbohm proposed using existing AM broadcast stations to relay messages from one broadcast station to the other. Why broadcast stations? HF ("short-wave") communications would be destroyed by high altitude nuclear bursts affecting the ionosphere, rendering short-wave communications impossible. On the other hand the lower frequency broadcast stations during the daytime depend upon "the ground wave" and are not affected by ionospheric changes. Thus, the short range broadcast stations would continue to be operative. But, given the short ranges involved, you would have to repeat the message from station to station. To do so quickly and accurately using voice would be impractical. I showed how with little digital logic at each of the nodes you could get a message cross country and showed how survivable the system would be against enemy attack. Since it only had to be a simple message, you could flood the network with that simple message to avoid routing considerations.

The way RAND works is that if you have an idea that you think is ready for the outside world, it will be carefully reviewed internally. Then, and then only, will you be allowed to formally present the work outside. It's all a matter of quality control. When you pass the quality control filter, you are encouraged to go out with a set of briefing charts to convey the message to all places in industry, academia, and government that would be in a position to comment intelligently. So that is what I did. The major initial objection to the scheme was its limited bandwidth. The generals would say, "Yes, that would be okay for the President. But I gotta do this, and so and so gotta do this, and that command gotta do that. We need more communication than a single teletypewriter channel."

After receiving this message back consistently, I said, "Okay, back to the drawing board. But this time I'm going to give them so damn much communication capacity they won't know what in hell to do with it all." So that became my next objective. Then I went from there to try to design a survivable network with so much more capacity and capability that this common objection to bandwidth limitation would be overcome.

O'NEILL: Were you able to convince the Air Force to actually fund a study to produce the reports?

BARAN: Oh, no. I didn't have to. RAND was already funded. I never had to go through the proposal begging stage coupled with a delayed approval process famous in government funded activities. RAND received its money once a year and it was allowed pretty much to do what it wanted to do. There are a few projects that they would be informally obligated to do. These are readily manned because while everybody likes to think that they are able to come up with their own projects in their own field of interests, as a practical matter most people in research send the message, "Gee, what am I supposed to be doing?" Those who made the mistake of asking got the project work to do.

Very quickly at RAND I received what would now be an amazing amount of freedom. I could do whatever I wanted to do. The only thing that RAND management did require was that my underlying assumptions be realistic and the logic consistent. I received strong support and encouragement from both Computer Science Department and top corporate management, including the freedom to travel without excessive justification.

I still was taking flak from other parts of RAND, but this diminished over time as their objections raised were being answered one by one. RAND was by far the most effective research organization I have ever encountered in my life, in part because of that freedom. It trusted the people, and the trust was honored. I might mention that this degree of freedom was not a result of management laziness. Rather it was the result of management wisdom, dedication to intellectual honesty, trust in individuals, and a true understanding of the research process.

I never had to waste time begging for money or writing proposals. Management ran interference so others could carry the ball. The Computer Sciences Department was managed by giants in the field of computer research management -- John Williams, Willis Ware, Paul Armer, Keith Uncapher -- all now so well known to the Charles Babbage Institute and to ARPA. If we had more time I could tell you of the major contributions they made to the work on this project. They were great men to work for, and I was always able to count on their continuing active support whenever I had to play an adversary role to accomplish my objections. Whatever success resulted was in a major way due to their intelligence and research management skills.

O'NEILL: How did the eleven volume report as such start to develop?

BARAN: Where do they come from? Why produce so damn much paper? Good question. Why would anybody in their right mind voluntarily grind out so much paper? I did it only with great reluctance, and in piecemeal fashion. Here is how it happened. I had this set of briefing charts and would present the concepts throughout the relevant military and R&D community. The responses were mixed. Some thought it great. Many others said something like, "Since it hasn't been done, it probably won't work." More useful were comments like, "It probably won't work because ..." and then would give me a reason. Most times I could answer the question with confidence. And sometimes I was less sure of myself. I would have to go away and think about whether the objection was in fact valid. And, if so, how one would circumvent the problem noted. This made for a lot of detail paper. It is easy to propose a global concept. It is far more difficult to provide enough details to overcome the hurdles raised by those that say "It ain't gonna work."

The basic network configuration was simple. Avoid any central node. Build a distributed network of nodes, each connected to its neighbor. How much redundancy of connections are needed for survivability?

The first interesting thing I found out very early in the game in early 1960 was that it would only take about three times as many links as the minimum needed to connect all the nodes to produce an extremely robust structure. That is, any node that survived the physical attack would almost always be able to communicate to the largest group of surviving nodes. That was a most fortunate finding because it meant that we would not need to buy a huge amount of redundancy to build survivable networks -- just a redundancy level of maybe three or four would permit almost as robust a network as the theoretical limit. A pleasant unexpected result. In order for a distributed network of the type proposed it would require that signals would go through many different nodes. This would mean all-digital transmission, smart switching, but offer potentially very high capacities. Simple concept; now for the realities.

The analog AT&T long lines telephone system at the time would never switch through more than five switching nodes. One was preferable, five was the absolute worst case. If you tried to connect too many analog transmission spans in tandem, the voice levels would be wrong and the noise and distortion prevent you from hearing very well. This meant that we had to use digital signals to allow regeneration of the waveshape. The idea of traversing so many tandem nodes proved to be a mental block for many analog communications engineers. But I would tell them that we would use digital transmission then being developed at Bell Labs by John Mayo and colleagues. That technology allows one to be able to reconstruct the signal before it is irrevocably changed. Well, here is a proven concept, yet it was hard to take, given the mental set of the older telephone people. Next we had to chop everything up into 1024 bit packets[5] so each packet could find its own way through the network. Now that concept was even more difficult for an old telephone guy to swallow.

[5] I called them "message blocks." The name "packet switching" came from Don Davies independently.

And then you had to tell them that each packet will find its own route on a statistical basis to get where it wants to go. After I heard the melodic refrain of "bullshit" often enough I was motivated to go away and write papers to show that algorithms were possible that did in fact allow a short message to contain all the information it needs to know where to go. To do so we had to simulate the network's performance. We then were able to show that it did not take very long for the self-adaptive behavior to occur efficiently so the network would be able to learn quickly where each node was even though each node had zero information at the start. A byproduct of this phenomenon was that network users' "names" never had to be tied to a physical location. The network could learn where its users were and be able to route traffic efficiently. Well, these sorts of statements in the old day of relay switches were eyebrows raising. You needed a lot of proof. That meant more simulation. Computer simulation of a highly parallel network chewed up a lot of computer time. But we were able to get the clock cycles we needed.

As this discourse was going on, the next level of questions was "Okay, theoretically it might work, but it will be prohibitively expensive. You need a huge computer at each of these nodes." I said, "No, I don't think so," but the burden of proof was implicitly placed on me that then required that I paper design the hardware computer needed at each switching node. The preliminary design showed that it really didn't take all that much equipment, so the cost could be reasonable.

The next level of objection, really a request for definition, was "You are going to have a lot of users on the network so you must define another layer of the switching system to connect large numbers of users to the net- work. And you have to show that that is reliable and cost effective." Okay, back to the drawing board again. This question required me to write still another report, this time on how you connect the users to the system. The basic idea that once you went digital, signals from data, teletypewriters, facsimile, and voice would all be digitized. Since everything was in bits and all bits are the same it was easy to build a universal interface. This proved to be another hard concept for older telecommunications experts to swallow without audibly choking.

O'NEILL: So, at that time, you were trying to include voice along with these other things and make allowances for that all the way through?

BARAN: Yes, why not? I mentioned that at that time there was work going on in Bell Labs, on what is now T-1 digital voice transmission. This proved that high digital data rates were feasible over existing copper pairs of the telephone system. The most outspoken of the "it ain't gonna work" school were the most senior AT&T technical management people. AT&T at that time held a total monopoly of all long distance communications within the US. All military telephone traffic in the US went out on AT&T lines. AT&T headquarters people insisted that the concepts proposed were not feasible. Bell Labs was comprised of two parts. The analog people who behaved just like their analog counterparts elsewhere and thought that the proposed system couldn't work. The second category was those digital people in Bell Labs highly competent in the digital art. These people had no trouble at all understanding what was being said. But the Bell Labs digital experts tended to be regarded at the time as being "not practical" by the people at AT&T. So when I say AT&T, I mean the headquarters group, then at 195 Broadway in New York City.

One helpful factor during this period was that AT&T was composed of gentlemen. Talk politely to them and they would invariably talk politely back to you. They may not agree. But generally you could expect a formal polite response -- unless you pressed them too hard. (A few exasperated souls at AT&T were not always able to restrain their temptation to send a few zingers via the back channel -- "Don't listen to that guy Baran, he is full of crap and doesn't know what in hell he is talking about...") But face-to-face the discussions generally proceeded with polite listening, generally followed by the same phrased party-line policy positions expressed in generalities. AT&T was in an awkward position here as the workability of such a network constituted a multifaceted threat. AT&T had long denied that there were any vulnerability problems with their network. They even blocked the military from the data needed for proper analyses. Their claims of invincibility were based upon distortions of fact, concealed weaknesses, and statements phrased in the common public relations style of the 1950s and '60s. Deny that any problem of any sort exists. Keep it all positive. (Corporate communications in general evolved to a far more honest level in the 1970s. But, back in the early '60s denial and concealment of problems at all costs were considered as proper corporate policy. Truth wasn't discovered to be a viable corporate communications policy until many years later, when investigative reporting made it far safer to immediately confess the existence of problems since it became nearly impossible for corporations to conceal really embarrassing facts. With the old worker loyalty ethic gone secrets became harder to keep.) Beyond the fear of the admission of concealed weakness there were, in retrospect, several basic corporate strategic constraints that prevented AT&T from seriously considering building an all digital packet switched network.

1. AT&T was a monolithic system. Everything added to the system had to work with all the old equipments of the past. Evolution, not revolution, was the prevailing concept. At that time there was simply no way to bring radical new technology into the plant and have it gracefully coexist with the old.

2. AT&T was a totally integrated system. While it did have digital transmission under examination, it was in the context of fitting directly into the plant by replacing existing units on a one-for-one basis. A digital repeater unit would replace an analog loading coil. A digital multiplexer would replace an analog channel bank: always a one-for-one conceptual replacement; never a drastic change of basic architecture.

3. The idea of changing the basic system architecture built over the years that was then producing the best telephone service in the world was pure heresy.

4. And lastly, those in the top technical decision making levels of the AT&T corporate structure simply did not understand nor did they appreciate what digital technology from the computer art would mean for telecommunications in the coming decades.

I think that AT&T's views were most honestly summarized by AT&T's Jack Osterman after an exasperating session with me. "First, it can't possibly work, and if it did, damned if we are going to allow creation of a competitor to ourselves."

In the early days, anything that AT&T didn't make couldn't be connected to the telephone system. A company that made a plastic mouthpiece cover called HushaPhone, to allow voice privacy in large offices, was sued in each state of the Union in the 1950s. An undertaker who gave out free plastic phone book covers was sued by the telephone company.

AT&T's rationale was that the plastic phone book cover with the undertaker's advertising took revenue away from AT&T and thus lowered revenues, reducing the quality of telephone service to all. Utter nonsense. But when you are sued by an AT&T subsidiary (the telephone company sued separately in each state jurisdiction) you quickly learned not to take on AT&T, irrespective of the merits of the argument. It was this deadly environment of heavy handed opposition that blocked technical innovation in communication outside of AT&T's domain. Trying to stop technical innovation is like trying to stop a river. You can dam it up, but as the water gets higher and higher, watch out when the dam breaks. It was this attitude of arrogance, I believe, that lead to the unfortunate later fracturing of AT&T. It might not have occurred if AT&T had allowed a modicum of freedom to others to use their network as is now commonplace.

This was the environment of the times. AT&T was the proverbial 800 pound gorilla you had to deal with. One does not take on the gorilla unless there is no alternative. And at the time there was none.

Our objective at that time was to have AT&T build the network for the Air Force. The choice was simple, the Air Force had the dollars; AT&T had the monopoly on long distance communications. Everything dealing with long distance communications fell into the domain of the AT&T monopoly. Everything dealing with the telephone was theirs. Nothing got done unless they wanted it to be done. Nothing could connect to the telephone system without their approval. (The Carterphone decision did not occur until 1968.)

My challenge was to convince AT&T to cooperate. It wasn't easy. I recall walking into a room of AT&T engineers and started to describe how the network would work. One of the older analog transmission guys said, "Wait a minute son, let's try that again. You mean you open the switch here before the traffic has emerged from the end of the cross country circuit." I would say, "Yes." He raised his eyebrows, looked at the others shaking their heads and said, "Son, this is how a telephone works." It was pretty patronizing from time to time, until I learned to use Western Electric part numbers. This greatly improved the interaction. At least I didn't start out with an image of complete stupidity. Nevertheless, I don't think we were ever really taken seriously by AT&T during that time.

O'NEILL: When did this take place?

BARAN: This was ... let's see ... I started work in the simpler systems shortly after I arrived at RAND in 1959. By 1960-61 I was fairly deep into the stuff. At the time I was starting to write the stack of reports, I had already convinced myself that everything should work. I just had to keep coming back to answer questions that demanded more detail.

A secondary motivation, but one that provided me with the background that I needed, was serving on a Department of Defense Ad Hoc Committee at that time.

O'NEILL: This committee was meeting in the early 1960s?

BARAN: Yes. In the early 1960s.

Shortly after I joined RAND, the Department of Defense received a request from each of the services to build a big record communication switching system. The Army wanted to build one, the Navy was going to build one, the Air Force was going to build one, as would NSA. The DoD figured it would take a good hunk of the defense budget, so they created a committee to look at each of these systems and determine which they should buy and flush the remainder. Willis Ware was invited to serve on the committee, but he had too much on his plate so he recommended that I serve in his stead. This Committee was headed by Bob Scherer of DDR&E (Office of the Deputy Director of Research and Engineering).[6] The committee met every other week, usually in Washington. I would spend one week in California where I lived, and then one week in Washington; back and forth every other week for about a year or so. A result of this committee work was the opportunity to learn how these huge communications switches worked. Most importantly I was

[6] If I recall correctly, other members of the committee included Morris Rubinoff (Univ. of Pennsylvania), Gerry Dineen (Lincoln Labs), Elmer Shapiro (SRI), Marlin Kroger (IDA).

able to ask some really dumb questions. Dumb questions only are allowed without giving offense if you are a child or a member of a distinguished evaluation committee. "Why are these things so big?" "Why do these switching systems require rooms and rooms of stuff?" Answer: much of the capacity is journalling and storing traffic that passed through the node. "Well, why do you want to store all these transactions?" "Well, we always have in the past." It seems that the real reason that every communications center office was built with that burdensome capability was to be able to prove that lost traffic was someone else's fault. "I can prove it left here okay." I learned that these and similar antiquated requirements could be thrown out if we were able to operate in real time. The big switches were implicitly designed for an era where bandwidth was very scarce.

Because of the heavy travel involved I had lots of uninterrupted hours on airplanes, flying there and flying back. The airplane seat is a wonderful place to spread out and work uninterrupted. In those days, RAND allowed its staff to travel first class if the trip was over two hours. This gave me the equivalent of one day per week to work in a comfortable setting. Almost all my writing was done either in airport waiting rooms or in airplanes. Today the difference in price between coach and first class fare is so great that few organizations can justify first class travel any longer. In the old days the price differential was small. It's a pity because present day steerage type, crammed together air travel poses an invisible impediment to economic productivity by the business traveller.

O'NEILL: Was the AUTODIN system one of these systems that was being reviewed?

BARAN: An early version of the AUTODIN system was one of the systems. We looked at such systems and we would wince. "God, why all this junk?" After listening to that, I said, "Well, gee we don't need this and we don't need that if you move data real fast, you don't have to worry about keeping records, and that would simplify everything down to a very much smaller system." Again, these were hard concepts to swallow if you have done something a certain way all your life. Where were we?

O'NEILL: Well, you were talking about actually producing reports and the reasons for producing them.

BARAN: Yes, back to the reports. The key content was completed by mid 1962 and the writing cleanup in 1963. One of the very few areas that I found fault with at RAND was that they are so damn careful about the review process that it could take a year to get a report through the referees, make changes, and rereview the documents. Some department editors were utter nitpickers when it comes to reports. Since the profession staff had great freedom of expression, those doing the editing sometimes felt it equally necessary to display their own creativity, regardless of their understanding of the content. The editorial staff had a propensity to rearrange commas. It took a long time to go through the review process. I must share the blame for the long delays because I did not have time to spend looking at the iterative minor changes. I let the reports stack up in my in-box for months per iteration. We decided to hold up all the reports and release them at one time to simplify the cross referencing effort. This meant that the reports did not get out of RAND until about August 1964. But those that were working in the community had early access as the work was developing.

I wrote a short summary of the work, for the *IEEE Transactions on Communications Systems* March 1964. This paper is essentially the same as RAND paper 2626 that was published in September 1962. I looked to the IEEE paper as a pointer to the detailed reports.[7] Anyone reading the *Transactions* could get the full set from the worldwide RAND depository libraries. There are several hundred of these depository libraries today, and probably over 100 at that time. In those earlier days copies of all reports were mailed out free of charge. There were other RAND reports written about the subject prior to the August 1964 publication date. RAND reports tend to be distributed widely in their community of interest. Each author provides an initial distribution list of all the people he or she think should be on the list. RAND then adds their own list for government agencies, depository libraries, etc.

TAPE 1/SIDE 2

---

[7] Also see the full page abstract in Scanning the issues: Distributed Communications IEEE Spectrum August 1964, p.114.

O'NEILL: So you felt that the reports were well distributed?

BARAN: Yes, that was the intent. At the time I felt that they were widely distributed. But, sometimes reading some comments about these papers at a later date, I am less sure about it. Of course, it is one thing to write a report and another thing to read it. I think I saw a number somewhere -- one of those numbers you don't believe until you think about it for a while -- the average technical paper is read by about six people. At first you laugh. Then you say wait a minute, how could that possibly be so? And, if you think what reading a paper really entails; not skimming it, but sitting down and reading it from beginning to end, including the references. Using this definition little of the stuff that crosses people's desks really gets read. I recall seeing the pile of the eleven RAND reports on one person's desk who later denied ever seeing the reports. I reminded him that I saw the reports sitting in a pile on his desk at the time I visited and that we chatted about them. He took on a blank expression and then changed the subject.

O'NEILL: Were the reports classified? There is conflicting information in the literature about the status of the reports.

BARAN: Eleven reports of the series were not classified. Two were classified. So, instead of eleven reports there really were thirteen reports. Two of them came out later and were classified. One was called "Weak spots and Patches." I asked all the readers of the eleven volumes to try to poke holes in the system. Two weak spots were found. Both were patched and described in a subsequent volume that was classified. The second classified volume dealt with cryptography -- not major to the concept at all.

O'NEILL: So all of the eleven reports were published and not classified?

BARAN: That is correct. All eleven volumes were unclassified. The two classified ones issued later and did not add much to the discussion.

O'NEILL: Were there discussions with either RAND management or the military at that point about classifying any of the other eleven reports? Or was it clear to everyone that you did not want to classify them?

BARAN: We chose not to classify this work and also chose not to patent the work. We felt that it properly belonged in the public domain. Not only would the US be safer with a survivable command and control system, the US would be even safer if the USSR also had a survivable command and control system as well! There was never any desire for classification of this work. Classification tended to be restricted to analysis of vulnerabilities; how many sites would be damaged, what our assumptions were of the other guys' capabilities, versus our capabilities. They are about the only things that ever got classified.

O'NEILL: There are a few co-authors on some of the reports. Were they your staff or were they also independent researchers interested in this area? How do they fit into this?

BARAN: Well, there are two other people's names found in two of the series: Joe Smith and Sharla Boehm. Joe Smith and Sharla Boehm programmed two different network models. The results were critical assumptions so we conducted two independent approaches and then compared their results. Joe Smith was the author of one of the reports examining the effects of a "hog" node studying network behavior. Barry Boehm wrote a report (not part of this series) in which he suggested an improvement in routing algorithms.

O'NEILL: So they were assisting you by doing the programming?

BARAN: It was more than that. Their work represented contributions in creating insights into this new field of understanding distributed network behavior.

O'NEILL: Was their background in programming?

BARAN: Sharla Boehm was teaching high school until she got tired of snotty Santa Monica high school kids, so she walked across the street and became a programmer at RAND. She was remarkably effective.

O'NEILL: I was curious about the general nature of interaction with the other people.

BARAN: As time went on the support received from the other RAND departments increased and the work became less controversial. The way RAND works is that you build informed alliances with people in other departments who share interests in the same problem.[8] The eventual level of cooperation was excellent. But I did get off to a cool start. At RAND, whatever is published is always under the author's own name, rather than that of the corporation, or department. RAND is one of the few places you can read a report and know who really wrote it.

O'NEILL: There are references to other work on networks going on at the same time. For instance, the work at SRI. They published some reports in December of 1961.

BARAN: I don't recall this one, but I will read that paper.

I had my assistant, Wyn Wilks, go over to the Stanford Library and dig up a copy of this paper: "Link Error Control and Network Route Selection", *IRE Transactions on Communications Systems*, p.328+, 1961, by R.C. Amara, H. Lindgren and M. Pollack. We were both flabbergasted, Roy Amara the principal author is a neighbor and has been one of my closest personal friends for the last 22 years! You can imagine how stupid I feel in not immediately recalling his paper.

Clearly I must have read it at the time, as it is relevant to problems that I was working on at the time. I didn't get to

[8] A list of colleagues who worked together can be found in the Preface of Vol. 1 of the series.

know Roy well until the mid 1960s, and that was in a different context. I read Roy's paper with great fascination, trying to trigger recall by imagining myself back in time. I must confess that there was no clear image.

My next step was to try to take the clues we have today to see if they allow us to infer a reasonable scenario. We know that the IRE's paper's publication date was December 1961. A footnote mentions an earlier presentation at the IRE meeting in Feb. 1961, and that the work was sponsored by the Army. I would guess that the work was probably done about 12 to 18 months prior to the publication date. This would be in the middle of the period when I was defining the entire system. In the community interested in communications survivability, there were no significant barriers to information flow. I recall visiting SRI on numerous occasions, giving briefings of work in progress, and swapping information and ideas.

From the clues, it is very likely that we would have met, that it is likely we would have had discussions, and that it is certain that I would have seen this paper. I just wish my memory were better. But this was almost 30 years ago.

O'NEILL: Here is an *IRE Transactions on Communications Systems* in December 1962 by Prosser. There was not an explicit mention of him in Volume 5. Yet I have seen a suggestion that it might have had an influence.

BARAN: The Prosser paper is another one that I reread after you pointed it out. This is a little different as I recall hearing the name Prosser mentioned about five years ago, and reading the paper at that time. I didn't draw a recall then either. But it is clear that I screwed up in not including this paper in Vol. 5, History. What happened was that I completed Vol. 5 in 1962 (the last reference was dated 1961), but I was so busy working on other activities that I didn't make as much activity to tidy up these references. This paper had a December 1962 publication date so his work occurred earlier. His work was sponsored by the Air Force. I did visit Lincoln Laboratories many times, giving briefings and talking to people during this time period. Thus there was likely information flows in one or both directions. (Information flows tend to be bilateral.) But then again I am unclear on the details of the time. There were also informal RAND papers written that described the early work[9] so the basic ideas were in place and widely discussed in the 1960-62 time frame.

It should be understood that I was never working in a vacuum. Everybody working in this field sought to encourage others by swapping ideas and insights. It was not at all like today's commercial biotechnology developments where there is a premium on secrecy for commercial reasons.

O'NEILL: There were a few references to earlier papers that the reports superseded. Are those the ones you are referring to?

BARAN: Yes.

O'NEILL: You mentioned before about trying to get wide distribution. We talked about the classification issue. How did you decide who should review these? You came up with the review list?

BARAN: You raise two questions. The first is on the classification review. The second is on the selection of the reviewers of the draft for technical quality control. If I recall correctly, the actual security classification is done by the Air Force acting upon suggestions by the author, and by the RAND management. The reviewers were suggested by the author in this case. The reviewers were selected for their knowledge of the field; the background to be able to understand the report. And the volunteers had to be willing to read the volume of paper.

O'NEILL: You list in Volume 1 a list of actual reviewers who reviewed one or more of the volumes. Did you get a lot of reaction and feedback? Was there a lot of interaction going on as they were being reviewed?

[9] For example, RAND paper P-1995, May 29 1960, "Reliable Digital Communications Systems Using Unreliable Network Repeater Nodes" that provides a good statement of where I was heading and why.

BARAN: Yes, I think so. By then most of the feedback was old hat. I don't recall receiving any violent reviews, if that is what you mean. Remember, when these reports finally emerged from RAND they had lost their initial controversial flavor. The work had become highly respectable, so the reviewers merely provided RAND with the comfort of knowing that the reports met professional standards. By this time, the only major outpost of frontal opposition was AT&T. Everyone else's objections had muted.

O'NEILL: I just want to be sure I get the dates right. Is that all part of this long review process you were talking about, once the reports got written?

BARAN: Yes. The drafts were essentially compiled in 1962. Almost everything was done. It was a matter of finishing the loose details for publication.

O'NEILL: Were there reviewers as the drafts were being completed? Was that an interactive process with any of the people listed as reviewers?

BARAN: It has been 25 to 30 years ago so you must forgive me if I don't precisely recall how the review process was conducted. I believe that we had early informal reviews for most of the work in 1961-62. And, as the final reports ground through the mill, there was a second formal review process. I believe that RAND required three reviews per report. The final review occurred late in the game when I was already off working on other projects.

There are two levels of internal review within RAND. Your colleagues are providing constant feedback. And before any briefing all department heads would form a review committee. These people were great critics. These older department heads were researchers in their early days but were out of date. They made superb reviewers, being expert in knowing how to go for the weak spots in presentations logic.

Each outside briefing conducted also constituted an informal review process. We got back questions, comments, opinions, etc.

O'NEILL: Would people request information from you and you would then give a presentation on it?

BARAN: Any time any group who had the slightest resemblance of a need to know, or would have an academic interest in the subject, would receive a briefing. I don't recall turning down any request. If it were to be to only a person or two, such as the press media, they would be invited to come to RAND to visit. RAND received lots of visitors.

There is an expression in RAND that it takes about 30 to 60 briefings to sell an idea. This set of ideas was no different. It was expected to take a lot of work to get the idea sold. I believe that RAND archives will probably contain the names of groups that were briefed, and distribution lists for its reports. I have some old calendars (1965+) that have some meeting dates that I may be able to find, and I suggest that you contact RAND for their files.

O'NEILL: Well, I will see what I can get.

BARAN: When looking at report distribution lists there is a caveat to be kept in mind. Just because a person receives a report doesn't mean they read it. You may not always be able to infer who knew what, and when they knew it, just by the date of the report, or even if they were on the distribution list.

O'NEILL: Exactly. But it would be nice to clarify the issue.

BARAN: There are a few organizations that are automatically placed on the distribution lists. These would include the Air Force, DoD, ARPA and people and organizations known to be working on Air Force or ARPA contracts in related fields. These people would be placed on the distribution lists when reports came out. It was assumed that each research agency would provide copies to their contractors, in the event they were overlooked from the distribution lists. The objective was to insure that everyone working in the field received copies, including and especially AT&T. However, initially it was a field with relatively few players, so there is a chance someone could be behind the door when the reports were passed around. In this community the key people knew what was going on by the informal channels. Only those far down the organizational structure got their information via the report channel.

O'NEILL: You mentioned how AT&T felt as the reports were being written. Why do you think the system that you proposed was never implemented?

BARAN: Interesting story. Let's go back in time.

By August 1964 the reports finally emerged. In 1965 RAND made a formal recommendation to the Air Force to proceed.[10] By then RAND was 100% behind the project. And, the Air Force was sold and also totally behind it. The Air Force then created an evaluation review committee. It was run by MITRE on behalf of the Air Force System Command. The MITRE evaluation group went through each of the parts of the proposal. I recall that there was a separate evaluation of costs; an evaluation of cryptography, and every other part of the system. Their conclusion was positive and they recommended proceeding. This was now about 1966.

Now let's go further back into history. In about 1949 the Department of Defense was created to unify the separate and competing military services. The consolidation into a single Department of Defense took place on paper quickly in 1949 but not in practice. In the 1950s and early '60s each military service acted as a Balkanized domain. Each had its own power base. For example, the Air Force had the bombers, therefore the Air Force was responsible for strategic defense. RAND worked for the Air Force, not the DoD. On paper DoD had responsibility, but the Air Force had the people. Power went to those with people. When McNamara became Secretary of Defense in the Kennedy era (1961+) he

[10] Copy attached.

immediately set out to implement the Defense Reorganization Act of 1949. These were not easy moves to make. You have all sorts of turf issues -- numbers of slots that are open for advancement, that sort of thing. The Department of Defense took more and more power away from the services. As part of this move the Defense Communications Agency was created. This new agency assumed responsibility of all long distance communications for all the services. To avoid jurisdictional disputes the DCA was run by an organization headed by one Air Force general, one Navy admiral, one general from the Army, etc. And, worse, this early DCA had near zero technical competence in digital technology. When you staff an organization of this type you sometimes acquire those not wanted by their parent organizations. Even in 1966, the DCA was extremely weak in technology. If you were to talk about digital operation they would probably think it had something to do with using your fingers to press buttons.

DCA has long since corrected their initial technical weaknesses and has since acquired some very good people and become competent in these technical areas. But at the time (1966) DCA primarily consisted of operational people, with little understanding or real interest in high technology.

Now let's return to the MITRE Committee proceeding under the Air Force Systems Command's direction. The authorization paper work to proceed took DoD approvals. At this time the General Counsel of the Department of Defense determined that as a long distance communications system, the task of building the new network would be assigned to the Defense Communications Agency.

A key person in DoD communication funding decisions was Frank Elldridge, Jr., Special Assistant, Command Control and Communications in the Office of the Assistant Secretary of Defense, Controller (Systems Analysis). Frank was a part of the team of analysts that Secretary of Defense McNamara assembled, and like several in that group had come from RAND. At RAND, Frank was a project leader in command and control survivability studies. And he was a very early and a very strong supporter of this work on distributed communications while at RAND. And he was a personal friend.

Frank and I agonized over this one. We agreed that DCA had the charter. The legal determination had been made. We also agreed that the then present DCA wasn't up to the task. I felt that they could be almost guaranteed to botch the job since they had no understanding for digital technology, nor for leading edge high technology development. Further, they lacked enthusiasm. Sometimes, if a manager doesn't have the staff but has the drive and smarts to assemble the right team, one could justify taking a chance. But lacking skills, competence, and motivation meant betting on a sure loser.

We found ourselves agreeing that DCA should *not* be given the funds to proceed, as the chance of their success would be too low to justify the risk. This risk was compounded because we both knew that if the project turned into a botch, it would be extremely difficult to get it going again. Detractors would have proof that it couldn't be done. We decided to wait until an organization with the requisite competence could be found that could take on the task within the DoD restrictions.

O'NEILL: So no implementation rather than have a bad implementation?

BARAN: It was important. Yup. That was a hard decision, but I think it was the right one. We could have wasted a lot of money.

O'NEILL: And this was 1966?

BARAN: Yes. I checked my old calendar, which I save as a name reminder. I found an entry on Wednesday, 15 December 1965 for 10 AM to 12 noon. I was meeting with Segerstrom of the Air Force Systems Command Electronic Systems Division at MITRE. I believe that this was a meeting of the MITRE Committee. This committee completed its work in about three months or so (that's a guess) so this period must have been in mid-1966. I'm sure that I have other records, but it would take a lot of digging.

O'NEILL: How would you characterize your interaction with the developments in networking after 1966? After the decision about DCA?

BARAN: By 1963-64 I was spending almost all my time on other activities. About the only thing that I was doing during this period was presenting results of this work in briefings when requested. In 1966 I was occupied in other activities.

O'NEILL: Okay.

BARAN: I did continue during this period to encourage others to study the characteristics of networks such as I had examined. For example, I was an ACM National Guest Lecturer at the time and spoke at a number of campuses on network survivability and on computer privacy.

I also taught the subject of this work in detail at a University of Texas Computer Sciences seminar course, and at a summer session (1966) at the University of Michigan. The objective was to get the word out.

What I did not appreciate at the time was the commonality of the alternative approaches to building such networks. I simply had approached building the network as an engineer. I assumed that there would be many different ways of building such networks, each with interesting, new and different properties. I viewed my work as an existence proof. I believed this was not the only way to build such networks. As it was only the first one I considered in detail, I felt that there must be totally different ways to build better networks. My view was not to say this is the only way to build such networks. Rather, "Here is one way, why don't you see if you come up with other better ways. In any event, we know we certainly can do better than we're doing with our highly centralized networks, so let's start to expand our thinking."

O'NEILL: Did you attend conferences with ARPA people after 1966?

BARAN: Yes.

O'NEILL: So you were around at these meetings?

BARAN: Oh, yes. I attended some ARPA meetings. The Computer Science Department at RAND received some money from ARPA. There were a number of highly innovative things being done at RAND under ARPA auspices at the time so I did get involved tangentially.

O'NEILL: Did you go to any of the ARPA principal investigator meetings?

BARAN: Yes, I have attended a couple of these. However, some may have been related to work that I was doing in the 1970s, but I did speak to the ARPA people on many occasions.

O'NEILL: There is an ARPA PI meeting listed at the University of Michigan in 1967 where they discussed the ARPANET.

BARAN: I checked my calendar, I did not go to one in 1967. But, as I mentioned I did lecture during the previous summer at the University of Michigan on the distributed network (packet switching) to spread the word.

O'NEILL: Were you personally involved, did people come to you with questions about what you had written, or ask you to review their work? How would you characterize your interaction with people like Larry Roberts and Bob Kahn?

BARAN: Well, I did not meet Bob Kahn until much later. I met Larry Roberts on a number of occasions, both in Washington and at RAND.

O'NEILL: When you say Roberts came out to RAND, you mean he visited? When would that be, do you remember?

BARAN: I really don't remember the dates, so I checked my old calendar (copy attached) and here is what I found: On Tuesday, 28 February 1967 I find a notation on my calendar for 12:00 noon Dr. L. Roberts. On Tuesday, 31 October 1967 I see a notation 9:30 AM to 2:00 PM for ARPA's (Elmer) Shapiro, (Barry) Boehm, (Len) Kleinrock, ARPA Network. On Monday, 13 November 1967 I see the following: Larry Roberts to abt (about?) lunch (time?). Art Bushkin = 1:00 PM. Here. Larry Roberts IMP Committee. On Thursday, 16 November 1967 I see 7 PM Kleinrock, UCLA - IMP Meeting.

I might say that my calendar tends to be incomplete and just shows dates not on formal itinerary sheets.

From time to time I would drop by to say "hello" and chat with the people at ARPA, as at that time (1966-67) I occasionally did a little consulting for DDR&E which was about a corridor or two away in the Pentagon.

I didn't meet Kahn until later. ARPANET went out for bids. BBN won the bid, and Kahn worked for BBN at the time.

O'NEILL: Do you know if RAND received a copy of the RFQ?

BARAN: Oh, probably a courtesy copy. RAND would never competitively bid. And certainly would never bid on a project of this sort. It would be totally inappropriate to RAND. (RAND is sometimes called Research And No Development).

O'NEILL: When did you leave RAND?

BARAN: I left RAND in mid-1968, but continued as a consultant for a bit.

O'NEILL: Did you see the RFQ -- did you review it at all?

BARAN: I do not recall, but in that community everybody knew what everybody else was doing.

O'NEILL: And were you still part of the community?

BARAN: Tangentially. I was off into four or five other different things at the time. I had become interested in long range planning (after doing a rotten job on the subject at Eckert-Mauchly). In 1968 I was in the process of helping to set up the Institute for the Future. We did a little work under a Ford Foundation Grant at RAND, considering the feasibility of the Institute. And I was doing a number of other things as well. Keith Uncapher is probably the best guy to describe that phase. Do you know who he is?

O'NEILL: Yes. He is part of the advisory committee for this project.

BARAN: Keith is probably the best person to say who did what, when. Willis Ware is another, and Paul Armer is a third. These are probably your best sources for unbiased inputs. They were very close to what was going on.

O'NEILL: What was your general reaction to the development of the ARPANET? Were there any surprises, for instance? Anything that you found peculiar or interesting?

BARAN: Well, yes. I recall chatting with Larry Roberts about the low data rates he was discussing. I was thinking in high data rates, to avoid long delays in the communications process. We chatted about that -- it was a cost and availability of circuits issue.

I might say that Larry Roberts is a man I admire greatly. His drive, intelligence, and single objective mindedness is what made the ARPANET happen.

TAPE 2/SIDE 1

BARAN: There was a meeting at UCLA about a year ago bringing together all the many people involved with the ARPANET. It was a delightful nostalgic time with many people who I have not seen for a long time. It was a great old home week. I enjoyed it very much.

Bob Kahn came up to me and said, "I didn't know anything about your work." I said, "That's fine, Bob. If you say so, I'm sure it was so." But, I sensed an unnecessary defensive air. Bob's great contributions over time are so well known that I didn't understand his defensiveness on the subject.

Larry Roberts described his primary motivation for building the ARPANET as resource sharing -- which it clearly was. During his early planning for the ARPANET Roberts described the idea somebody proposed was that the resource sharing to be done on a single centralized computer. Roberts described that he was against a centralized approach, and he mentioned the RAND reports either caused or were a factor in specifying a fully distributed approach.

What I think probably happened was that Kahn, who worked for BBN, got the job of writing the software in response to a specific request for proposal. In speaking to Larry, I think it fair to say that Larry had seen the reports, but probably Bob Kahn did not. If Bob said he didn't know about the reports, then I'm sure that he had not. Bob probably came to the project from a focus on the programming issues. Thus the reports may have held little relevance to him. And, hence, of lesser interest. Another factor is that he may have been so busy implementing that he had to minimize his reading time.

O'NEILL: Did you know Donald Davies over in England?

BARAN: I met him after the fact, many years later. I did not know him at the time. He had a nice way of putting it, he said, "Well, you may have got there first, but I got the name first."   He is correct. Packet switching is a far more graceful name than Distributed Adaptive Message Block Switching. Precise, economic, and very British. I take my hat off to him for coming up with such a wonderful name.

O'NEILL: So at the time he did not send you any of the reports he was writing?

BARAN: No, that came late in the game. I had completed all my early work by then.

O'NEILL: Oh, in terms of your work?

BARAN: My work was 90 percent over by 1962. Completing the drafts and getting them through the system was done by 1964. His work on short packets came later. He said he was thoroughly embarrassed when somebody sent him a copy of my report after he had done his work. He said he did not know about my work. I certainly believe him. But from time to time I wish he would correct the typo in his book that accidentally misdates my IEEE paper.

O'NEILL: Oh, I see it lists your IEEE paper as March 1969 instead of 1964.

BARAN: It's fun to see many people refer to that paper with the 1969 date year after year in footnotes and in bibliographies. It's obvious that they haven't read the paper, only the reference to it. This confirms my long-held suspicion that far more people are writing papers than are reading them. And, just because someone cites you as a reference, it doesn't mean that you can assume that they have read your paper. Otherwise how could a botched reference go undiscovered for so long? Interesting.

O'NEILL: In general, how do you feel about how the story shakes out right now? Do you pretty much agree with ... I know at the meeting at UCLA last year, the anniversary, you were given an award. I don't know any details about that, but I saw that in a report. Can you talk about what the award was for?

BARAN: Primarily, living long enough.

O'NEILL: That's what they cited?

BARAN: The citation says:

> The UCLA Advanced Computing Technologies 1989 ACT the Pioneer Award.
>
> In recognition of his early 1960s conception of an all-digital, computer controlled nationwide network using packet switching. Baran's ground breaking RAND reports "On Distributed Communications" have had a major impact on communications systems and information networks throughout the world.

The way someone characterized the UCLA meeting overview of the early history of the ARPANET was that there were two separate periods.

The first period, you should pardon the expression, was the "Baran Era... or pre-ARPANET." Everything beyond that date is the ARPANET Era. And, as I said at that meeting, "Hey fellas, I didn't do the ARPANET. That was Larry Roberts." (I am from time to time unfortunately given credit for things I haven't done, and conversely, I lived only in the period 1960 to about 1967 at the latest, and wrote a lot of papers that defined the interesting properties of packet switching. After that, it's someone else's era.)

O'NEILL: That really covers the questions I had prepared. Is there anything else you'd like to add? Any other comments you'd like to make?

BARAN: I think that covers all the points you were looking for.

My experience with innovations is that everything has a predecessor event or events. Generally when the next generation of ideas and effort comes along, what has gone before becomes irrelevant. Then the following generation comes along and there is the same shift of focus. The process of technological developments is like building a cathedral. Over the course of several hundred years: new people come along and each lays down a block on top of the old foundations, each saying, "I built a cathedral." Next month another block is placed atop the previous one. Then comes along an historian who asks, "Well, who built the cathedral?" Peter added some stones here, and Paul added a few more. If you are not careful you can con yourself into believing that you did *the* most important part. But the reality is that each contribution has to follow onto previous work. Everything is tied to everything else.

Too often history tends to be lazy and give credit to the planner and to the funder of the cathedral. Maybe we should take the care to avoid the simplifications and say, "Okay, this person did this or did that, and that person did so and so." No single person can do it all, or ever does it all. But we are lazy and tend to give all the credit to a single person most closely identified with an activity and forget all the others who really made it all possible.

O'NEILL: That is the reason for having a professional history done of this area. We try to weave in some of the past, to get more perspective. Thank you very much.

BARAN: Thank you for coming by.

END OF INTERVIEW

Attachments – photocopies are available from CBI

1. U.S. Air Force, Project RAND, Recommendation to the Air Staff, *Development of the Distributed Adaptive Message-Block Network*, August 30, 1965.
2. Letter from the Air Force, October 11, 1965
3. Chapter by Paul Baran on packet switching from *Fundamentals of Digital Switching*, 2nd ed. (1990).
4. Paul Baran, Mathematics Division, The Rand Corporation, P-1995, *Reliable Digital Communications Systems Using Unreliable Network Repeater Nodes*.
5. Calendar pages.
6. *IEEE Spectrum* August 1964, p. 114.
7. List of RAND depository libraries.