Oral History of Robert Dennard

Interviewed by:
Gardner Hendrie

Recorded: July 20, 2009
Yorktown Heights, New York

CHM Reference number: X5377.2009

START OF TAPE 1

Gardner Hendrie: Well, we have here today Robert Dennard, who has graciously agreed to do an oral history for the Computer History Museum. Thank you very much.

Robert Dennard: You're quite welcome.

Hendrie: I think maybe the first thing we ought to try to cover is a little bit of your very early history, your family background, where you were born and brought up, what your parents did. A little bit of very early background.

Dennard: Okay. Well, I was born in Terrell, Texas, in 1932. That's an hour's drive, I guess, east of Dallas. My father was a dairy farmer at that time, which you recognize was in the middle of the Great Depression. We moved further east in Texas after four or five years. I really started my education in a one-room schoolhouse at the crossroads of a couple of country roads near Carthage, Texas. In about 1941, we moved to Irving, which is near Dallas, when it was a small town. It's presently a very large part of the metropolitan Dallas-Fort Worth area. It was a small town then.

Hendrie: Now, was your father still doing dairy farming?

Dennard: No, he'd discovered a job with a fertilizer company, actually. That's agriculturally related too, I guess, in Dallas. I continued my schooling in Irving, and graduated from the high school there.

Hendrie: Could I ask, did your mother work? Or was she a homemaker?

Dennard: My mother was a homemaker most of the time, but in Irving, she worked in the school cafeteria as a very nice dessert cook. I know I can attest to the fact she cooked some pretty good pies.

Hendrie: She was a pretty good cook.

Dennard: At home, as well as probably in school.

Hendrie: Did you have any siblings?

Dennard: I had three older, one brother and two sisters that were quite a bit older than me. I was a late addition to the family. It was pretty much like being an only child. They were away from home.

Hendrie: Were they a lot older?

Dennard: Actually, my sister and brother were both away in World War II.

Hendrie: Oh.

Dennard: At that time we were in Irving.

Hendrie: So they were 10 years or so older than you were.

Dennard: Yes.

Hendrie: Tell me what your interests were when you were young in elementary school, or going into high school? What sorts of things interested you particularly?

Dennard: Well, in elementary school, in just getting started there, we really lived pretty much way out in the country. My interests were a BB gun and just I had a lot of spare time on my hands to really admire the scenery and I rode horses occasionally. We had a horse on the farm. I helped my father sometimes. A little bit, not very much, with some of the agricultural things. I had chores to do. Chopping kindling for the fireplace and things like that.

Hendrie: Okay.

Dennard: I don't recall reading a lot at that time. Later on, when I was maybe in the sixth grade and beyond, I started reading quite a lot, quite a lot of science fiction. And my sister had left behind a really thick book of anthologies, a lot of H.G. Wells stories. And she left behind a complete selection of Ogden Nash poems, which I just loved. I know quite a few of those. And she left behind an album of Sigmund Romberg operettas, which I played many, many times. I learned most of the words to those. I really love that. She left me behind some really good things to start off some kind of intellectual career and I got some appreciation for the arts from that too.

Hendrie: What did your siblings do when they grew up?

Dennard: Well, my oldest sister was a nurse. She was a nurse in the Army. And eventually, she married a doctor. She was a school nurse. My brother worked, after he got out the Army, he got married and raised a family. I think he owned a filling station, and eventually went to work in the post office and retired from that at a pretty advanced age, I guess. He had a long career in the post office. My other sister married a farmer. She eventually did some work in a bank, a local bank. Most of my family stayed pretty much where they were and did whatever they could do in that same area.

Hendrie: Which of your sisters is the one who left you these treasures?

Dennard: The sister who went off in the Army as a nurse.

Hendrie: What's your earliest memory that you have of what you wanted to be when you grew up?

Dennard: You know, I don't think I formed any particular attitudes about that. When I was getting out of high school, my guidance counselor suggested engineering would be a good career because of my math aptitudes. She said there were good opportunities in engineering. I didn't know much about it. But that's what I signed up to do.

Hendrie: So you really got sort of pointed in that direction.

Dennard: Yes. But I was already interested in physics and things like that.

Hendrie: You enjoyed your science classes.

Dennard: What little science classes we had at that time in this class of about 90 graduates. But I had a very good math background. And very, very sound English training, which of course was actually very important to a career like this.

Hendrie: Were there any particular teachers that sort of inspired you at all in high school, before you got to college?

Dennard: As I said**,** I had some really good English teachers. And some of those, certainly, I found inspiring. I did a lot of reading. The math teacher was very, very thorough, and I appreciated what he was doing. And I thought most of the teachers were really quite good.

Hendrie: Were you encouraged by your family to go to college? Or did that sort of just seem like the next thing to do?

Dennard: I don't recall. It was not like these days, where people worry a lot about what kind of college they're going to get into and all of this stuff. I guess it just was a natural thing. But they encouraged me in the sense that I lived with them and they supported me by providing a place for me to live. I was looking for very local colleges. I had planned to go to a local junior college, which was not far away, and commute. Most of my classmates were going there, the ones that were interested in engineering.

Hendrie: What was that college?

Dennard: That was Arlington State.

Hendrie: And you were living in Irving?

Dennard: Yes, and Arlington StateJunior College was very close. I did not go there, as it turned out, because something happened that changed my career quite a lot. I had been playing in the high school band — they just started the high school band during my time. I was in the inaugural edition of this band. One day right after I had graduated, the director from the Southern Methodist University Band whose name was Dr. Oakley Pittman, came to visit me in Irving. I'm sure my band director had had something to do with organizing that. Dr. Pittman offered me a band scholarship to come to SMU.

Hendrie: What instrument did you play? I have to hear this.

Dennard: Well, I played an E-flat bass horn, actually, which is a little bit different. It's not quite as large as a double B-flat bass horn. It's a little bit smaller. He suggested SMU would be a very good school for me. He suggested that when you have an opportunity like that, why not take the best opportunity you have available? I got some general advice along with the scholarship, which I thought was very valuable. And I've used that general advice later on in other choices that I've had to make. That was my opportunity.

Hendrie: And you took it?

Dennard: And I took it.

Hendrie: Now, where was SMU located, relative to you?

Dennard: It was in Dallas.

Hendrie: Okay.

Dennard: It was a little bit longer commute, but I was still able to commute there. And having a scholarship did enable me to get through the first year. And after that, they had a co-op program. So I was able to work at Dallas Power and Light Company half-time, I guess it was eight weeks work and eight weeks of schooling, alternating. It took an extra year to get through my bachelor's degree.

Hendrie: That enabled you to continue and get your degree.

Dennard: Yes. And get a little experience with power, which is an interesting current challenge. Like how to put solar cells onto power lines, or to transmit the power through power lines. So I have a little bit of knowledge of what power technology is like.

Hendrie: While you were studying at SMU, what were the subjects in the engineering curriculum that you enjoyed the most?

Dennard: I guess the first subject was not quite in the engineering curriculum. The physics course, which is in the general arts and sciences, was what impressed me the most, by its level of difficulty.

Hendrie: Did you choose that, or was that a requirement?

Dennard: That was a required course. It was very interesting. Although it seemed to be that when I got half the problems right, I was really doing quite well, compared to a lot of the people in that course. It was a little startling after a very comfortable high school education. Which was not very challenging.

Hendrie: Yes, okay.

Dennard: And the math courses were quite interesting. I would say calculus was a very interesting course to me. And then the engineering courses, actually, elementary electrical engineering courses. I guess basic strength of materials, some of the basic civil engineering courses that were kind of applied things that gave me a chance to really understand some of the physics that I had had a little trouble with the first time around. I thought those were quite interesting. The early ones are the ones I remember the most.

Hendrie: As you progressed, did you specialize? You know, you were working for the Power and Light Company?

Dennard: No, I didn't anticipate I would be doing that. I thought I would probably be getting into electronics. But electronics was quite different. In those days, it was vacuum tubes. We had a lot of good control theory, which was pretty well advanced. That part was interesting, although it was pretty analytical and not too much applied to anything. I really enjoyed my courses in transmission lines, signaling on transmission lines. The labs in that were very interesting, and I think I got quite a good understanding of travelling waves on transmission lines. Although I guess I'm not sure what the applications of transmission lines were at that time. That was pretty much before a lot of high-speed networks.

Hendrie: But that's really valuable to understand.

Dennard: And I still use that a lot in my everyday work.

Hendrie: Yes, all right.

Dennard: In terms of specializing, no. I didn't know where eventually I would be. I was just trying to learn the basic stuff.

Hendrie: Did you continue all through your college career working, doing your co-op work at the same place, or did you go to different companies on the co-op?

Dennard: No, it was just that one company. I must have worked there on and off for four years.

Hendrie: Did you get any valuable experience in your co-op work? What kinds of things were you working on?

Dennard: Well, in fact, I started out helping survey crews put up stakes where lines were going to be constructed and poles were going to be put. And then I went out myself and did some measurements at places that had been set up to do measurements on the system. I got into a little bit of understanding about some of the somewhat more difficult or challenging things with regulation of some of the higher voltage distribution lines. It was not a great amount of experience.

Hendrie: How did you do in college, relative to your peers?

Dennard: I think I did very well. I was probably the first in my class. As I had been in high school, I was the first in my class. I think I was probably the best electrical engineering student in terms of grades and stuff like that. I eventually went on to take a master's degree. And I guess I did well in some of those courses. Particularly a physics professor there in solid state physics was very kind to me and helped me plan where I would go to graduate school. He suggested two or three universities and gave me good letters of recommendation.

Hendrie: So you did decide at some point during your undergraduate that you probably would get a master's degree?

Dennard: Yeah. I'm not sure if there was a lot of thinking at that time about continuing education. It might have been an alternative to being drafted into some of those wars that were going on at that time. The Korean War was happening when I was an undergrad.

Hendrie: About what time would you have been graduating?

Dennard: I graduated in 1956, with my master's degree. '54 I got my bachelor's.

Hendrie: Yes, I do think there was some draft going on, since I also graduated in '54, and I remember.

Dennard: I remember taking the Selective Service qualification test to stay in college, to show that I had sufficient aptitude to stay in college.

Hendrie: All right. You'd now gotten some introduction into solid state physics in your undergraduate program, toward the end.

Dennard: Yes. And the graduate program as well.

Hendrie: Yeah. So where did you go for your graduate work?

Dennard: I did my MS at Southern Methodist, but I got an offer of a fellowship at Carnegie Tech in Pittsburgh. And that was one of the top schools on my list, so that's where I went.

Hendrie: What did you get your master's in?

Dennard: Electrical engineering.

Hendrie: Electrical engineering. You had not switched to physics.

Dennard: No.

Hendrie: Or anything like this? Okay. Did you have to do a thesis project to get your master's at SMU?

Dennard: As part of the master's degree. Yes I did.

Hendrie: So what did you do it in?

Dennard: Well, it was an analog computer that I built to solve some differential equations.

Hendrie: So you built your own analog computer to do that?

Dennard: Yeah, out of vacuum tubes. Probably a pretty rudimentary one, but I learned how to assemble a thing like that and how to get it working, and how to hook it up to solve some basic differential equations.

Hendrie: You built this analog set up to solve some differential equations. Why did you need to solve these differential equations?

Dennard: To get my master's degree.

Hendrie: But what was the problem that they described?

Dennard: I don't think there was any particular special problem.

Hendrie: You just needed to do something, so you said, "This looks like an interesting project," went and did it. Did you look at any other graduate schools besides Carnegie Tech?

Dennard: I applied to MIT and Cal Tech and Carnegie Tech. Who else? I'm not sure. I think those were the ones. I think I was accepted at all those, but got a fellowship at Carnegie Tech, which enabled me really to do it.

Hendrie: So you went to Carnegie Tech. You were in a Ph.D. program?

Dennard: Yes.

Hendrie: What were you focused on in terms of a Ph.D. program there?

Dennard: Well, this was another quite different experience and another level of challenge. The classes there had a lot of good students. I started my research getting hooked up with a professor to do the graduate work. I was fortunate, I think, to be assigned to Professor Leo Finzi. His principal area which his students were working in was magnetics, actually. And so I had a project, had a particular magnetic resonant circuit involving a magnetic core and a capacitor and some questions about subharmonic oscillations in a circuit like that. That's kind of a specialized field, but something where there was still some lack of understanding. I was able to find some new things and get some things happening with experiments, and produce some theory to back up the experimental results.

Hendrie: At that period, people were working on magnetic circuits for digital computers and magnetic amplifiers.

Dennard: Magnetic amplifiers, yes, you're correct.

Hendrie: I just was curious whether that was the direction of the overall research?

Dennard: I'm not sure that my professor envisioned this becoming practical. But he was an authority in magnetic amplifiers. And they were being used, or proposed to be used, to control elevators and things like that. And Westinghouse was a local company that was interested in those kinds of circuits and materials.

Hendrie: But more for control circuits than building computers. I know there were some computers actually built out of magnetic circuits, I think there were a couple of kinds of magnetic-- Remington Rand actually sold a computer built out of magnetic circuits commercially in that era..

Dennard: Well, you know more about it than I do.

Hendrie: When did you get your Ph.D.?

Dennard: 1958.

Hendrie: 1958. That's relatively quickly. Yes.

Dennard: Let me just give a little more information, since you've asked. I enjoyed a lot of the coursework there. There was a particular kind of program called engineering analysis. There was a textbook called Engineering Analysis. And there were professors VerPlank and Teare, who had developed this course material. It was material aimed at training us to really learn how to solve problems independently, how to pose a problem, set up an approach, and proceed in a very systematic way. A lot of the problems involved solving differential equations. I think we ended up being able to solve most any differential equation that was known. They applied to solutions of thermal problems, and electrical problems. They were all the same differential equations pretty much. I thought that was the most thorough part and the part that's helped me the most in my career. So I wanted to say something about that.

Hendrie: I'm glad you interjected that. I wouldn't have gotten there, I don't think. When you graduated, what were the options you thought of as to what to do next?

Dennard: Well, by that time, I knew I wanted to do research. I interviewed Bell Labs, General Electric, and IBM, who were really the people who were doing something. IBM was just getting started. This research building where we're filming this today was just being built at that time.

Hendrie: Oh, my goodness.

Dennard: They were hiring a lot of people. And a lot of my friends were going to IBM. Eventually, after the other interviews, I thought this was the best place also. I thought it was kind of exciting to come to New York and come to work for IBM.

Hendrie: All right. Was there a specific research group at IBM that particularly interested you?

Dennard: I don't recall exactly how I got into a particular research group. When I arrived, we were housed in temporary facilities while this building was being constructed. There was a laboratory that had been outfitted down in Ossining out of an old hospital. They had made laboratories out of it and offices. That's where all of the applied engineering work was done. There was another facility for mathematics and other things. In all it was not, I guess a very large activity, maybe a couple of hundred people or something like that in the early days.

Hendrie: And it was a group that did applied engineering?

Dennard: Well, it was research in engineering possible solutions for computing.

Hendrie: Let's take a break. We're going to change the tape.

Dennard: Great.

END OF TAPE 1

START OF TAPE 2

Hendrie: Tell me a little bit about what you were asked to do when you first arrived at IBM.

Dennard: Well, I was asked to work on a project that was somewhat related to the work I'd done in college. It was another kind of resonant circuit involving an inductor but it was, in this case, a linear inductor and a non-linear capacitor rather than a non-linear inductor and a linear capacitor.

Hendrie: All right.

Dennard: Not too much different from what I had studied. It was capable of being set to oscillate at some sub-harmonic frequency of an applied signal. There would be an applied power source, basically a sine wave at some high frequency, and this circuit would oscillate, say, at a third sub-harmonic of that frequency. Therefore, it had three different phases possible with respect to each other, and could represent ternary logic. That was the first thing I worked on. That was a lot of fun. We had someone doing the logical designs in ternary and some experimental work. That was actually a very interesting project.

Hendrie: Very. At the same time, there was a fellow by the name of Bob Winder at RCA Labs who was also working on ternary logic, just an incidental thing. It was a field of interest at that time.

Dennard: Yes, well, this was based on the Von Neumann patent, I believe, in this case. Von Neuman had asserted that epsilon, 2.718, was the optimum numeric base for digital computers. Of course, we did binary, you know, and today all computers still use base 2. But he suggested that 3 would be a better base. I'm not sure better in what way. Probably it would minimize the number of components. In those days people were concerned about how many components were used to do the logic.

Hendrie: And so that was why we were pursuing devices that behave in that ternary way?

Dennard: Yes.

Hendrie: How long did you work on it?

Dennard: I must have worked on that for a couple of years.

Hendrie: Really?

Dennard: We ended up writing a report. It seems like a couple of years. <laughs> We also did some work on magnetostrictive memory. That's essentially sending an acoustical wave on a metal wire driven by probably the magnetostrictive circuit. There was a transducer that developed the electrical signal and produced the acoustical signal.

Hendrie: All right.

Dennard: After that, we eventually moved into another location, the Mohansic Lab, which was constructed in Yorktown, north of here, a larger facility and a little better equipped. I got involved in a data communication project, which was quite interesting. We worked on developing a modem for communicating on telephone lines and we advanced the state of the art from 2400 bits per second to 4800 bits per second. It's something like 56,000 bits per second now and I never would have thought it was possible, based on the kind of telephone lines we had in those days. Phase jitter and things like that were really quite hard to cope with.

Hendrie: Now, was this the same group that you had started working in here that moved or did you switch groups?

Dennard: This was, I believe, another group, another manager. Which was common. Well, some of the same people were there. There was a fellow, Dale Critchlow, that joined IBM Research the same day I did, June 16th 1958. He had been working not far from what I was doing throughout a large part of my early career so we were working jointly on the project at Mohansic.

Hendrie: All right. Good. So the state of the art of modems at that period when you were working on it was that Bell Labs could supply commercially was 2400 bits per second, or was that not even commercially available at that point?

Dennard: I'm not sure they were being used. Personal Computers and things like that were not yet being used so I'm not sure what modems were even being used for. There were terminals, certainly, intelligent terminals, which were connected to mainframes. Possibly they could be remote and hooked up through —a modem. I'm not sure what the application was.

Hendrie: Well, the very early modems were used in the Sage system to send the information from the radars, digitize it at the radar and send it digitally to the Sage computers at the Sage Centers. So there were a few applications such as that military one.

Dennard: That kind of data rate? My.

Hendrie: I think they were very slow. I don't remember the rate.

Dennard: I think there were a lot of FSK phase shift modems that were even, you know, quite lower frequency than that.

Hendrie: Yeah. Okay. Good, so that was an interesting...

Dennard: It was interesting work. You mentioned Bell Labs. I know Bob Lucky did some early equalization work so there was a concept of time domain equalization which was just coming in, which was very good for correcting amplitude and phase errors for different lines which had different characteristics.

Hendrie: Bob Lucky is a past trustee of the museum.

Dennard: Great.

Hendrie: Were the active elements vacuum tubes still at this time, or were you using transistors?

Dennard: I think they were transistors, yes. Now, my memory is a little-- I remember the inductors and the capacitors...

Hendrie: Okay. It was probably in the transistor era.

Dennard: Probably discrete transistors.

Hendrie: Is there anything else that comes to mind that you think would be worth saying about this period when you're working on, in this case, the data communications modem project?

Dennard: That period of work was actually very exciting in the sense that, when we had developed this 4800 bit per second modem, my friend, Dale Critchlow and I were assigned the very pleasant job of going off to London to set up a transmitting site there with this modem. One of us stayed in London and the other, Dale, went off to Rome to set up a receiving modem there. We sent data from London to Rome on an international communication link and showed that this thing worked very well. But, among other things, we also had a really great time visiting a friend of mine in London and eventually ending up in Rome together at the other end of this project, and doing some great tourist activities. That was quite an opportunity. The person we worked for at that time had a great imagination for flair and drama and had figured out this would be something to do to get publicity.

Hendrie: Who was this that you were working for?

Dennard: His name was Emil Hopner .At some other point in time, he arranged for us to send, I think, several megabits per second, or megabaud, data transmission on a coax line from our Mohansic lab to somewhere in Chicago and back, I think.

Hendrie: Oh, wow.

Dennard: That was...

Hendrie: That's a great opportunity, especially when you're young, to go and do something like that. Had you ever been to Europe before?

Dennard: No, that was my first trip out of the country.

Hendrie: That's wonderful. Now, when you develop something like this and get it so it works and are able to demonstrate it. Did they, at that time move from research into the IBM product world?

Dennard: At that time…

Hendrie: And did this one ever move?

Dennard: Well, actually, my associate, Dale Critchlow, very shortly after, went off to La Gaude, France where there was a communication products division, and he worked there for a couple of years kind of transferring that work into the development laboratory and training people there in the fundamentals in this kind of communication work.

Hendrie: Okay. All right. Very good.

Dennard: Which was, I'm sure, an experience that he enjoyed a lot. <laughs>

Hendrie: Yes. Live in France for a couple years. So what happens next in your career?

Dennard: Well, next in my career, I finally got the opportunity to come here to this building. Watson Research Center had been completed and was running. I got a chance to come and work for an old associate of mine, George Feth, from Carnegie Tech, where we had worked together. We were both students of Professor Finzii and knew each other well. He proposed I come and work on some stuff that they were doing. I think I was involved in one of the first projects in how to use tunnel diodes. We should call them Esaki Diodes because Dr. Esaki was here at this building at that time and had been hired by IBM after he received the Nobel Prize for his creation of what I'll call tunnel diodes, the more familiar term. I was involved in a project where I was pretty much the sole investigator to build fast memory out of tunnel diodes. One had been kind of designed by someone else and I was kind of perfecting it and putting it together. It eventually was kind of impressive. It could develop access times of, like, one nanosecond, which was very fast using some really high speed transistors. I remember, you know, the transistors in those days were put into-- some of the first integrated circuit boards were being put together so this was a discreet array of individual tunnel diodes but packaged densely together with drivers and receivers, but still discreet components.

Hendrie: And what year would this maybe have been?

Dennard: About 1964 - '63 or '64.

Hendrie: That's very interesting. Impressive. You took this project and built a working prototype.

Dennard: Yes. Which was built and tested and then pretty much put on the shelf because it really wasn't going anywhere. Let me characterize my whole early career as somehow a lot of research being done into various ways to do computing. Another one involved trying to figure out how to use microwave circuits, microwave amplifiers, which were really quite large traveling wave amplifiers and trying to put together things to do computing, thinking that somehow some high speed would come out of this. But keep in mind, in all that time, the transistor was still out there. I think, at the time that I actually joined IBM in 1958 was about the time that the first concepts of integrated circuits were being developed by Kilby and Noyce. So the future was there and it was kind of clear. But there was still a lot of exploration and other alternatives going on. I guess it was maybe understood that already transistors were the basis of the early computers that were being built at that time, largely with discrete transistors. Integrated circuits were just really beginning to happen in the middle '60s. They were really getting into computers.

Hendrie: So the mainstream may have been-- was sort of on going?

Dennard: The mainstream was going on and, as it turned out, the mainstream and evolution, the evolutionary part of the mainstream would really become so powerful that it would make any of these other approaches seem inadequate. It left them in the dust.

Hendrie: Left them in the dust, yeah. Well, it's still pretty impressive, though, one nanosecond access time, tunnel diode memory.

Dennard: Actually, it was, yes, but it was a very small array. I've forgotten how large.

Hendrie: It was very hard to actually generate data to put into it and look at the data that's coming out. Circuits were not that fast at that period.

Dennard: We had a lot of very big racks of high speed pulse equipment and pretty good oscilloscopes but it was all lab equipment. The laboratories were really quite well equipped to do measurements and things.

Hendrie: Okay. But to put that instrumentation into a computer might have been the hardest part. The tunnel diodes might not have ended up being the problem.

Dennard: Yes. <laughter>

Hendrie: In the sense of being able to make use of that speed.

Dennard: Well, about the same time, some people in our Fishkill labs had developed some small memory chips using bipolar transistors to store maybe 16 bits of information at nanosecond kinds of speeds, which was the much better approach.

Hendrie: Yes. Well, after you worked on the tunnel diode and got that working or demonstrated it, as you said, you didn't continue to work on it. That wasn't what you did.

Dennard: No, in fact, I feel very fortunate that I got an opportunity after that to join the early microelectronics project here at Yorktown, which was exploring how to use MOS transistors in integrated circuits. And with an ambitious goal, in fact: they used the words "large scale integrated circuits". They already had some idea that the longer term goal would be to integrate thousands of transistors into integrated circuits. That was considered large scale integration. Thousands of logic functions in a chip.

Dennard: They used terms like monolithic circuits at that time. All of the circuits on a single stone. <laughter> That project was started. My old friend, Dale Critchlow, came into the project and became my manager at this time. We were responsible for designing the transistors, the MOS transistors, and characterizing them and understanding how to build digital circuits from them. Quite interesting work.

Hendrie: Very good. How many people might have been involved in this project, in the microelectronics effort at this point?

Dennard: That particular project, which was aimed at N-channel MOS transistors, probably was 30 or 40 people.

Hendrie: And who would some of these be?

Dennard: Possibly some other people working on some memory design, static RAM designs outside that.

Hendrie: So that was a significant effort, then?

Dennard: Well, for that time, considering it was a fairly simple process that only involved four masking steps and each basic processing step took about a day to complete. We could get pretty fast results compared to today, where things are so complex and it takes so long to go through and the masks are so expensive. It was a good number of people and sufficient to do some pretty good work.

Hendrie: So there were people involved in processing and there were people in all sorts of aspects, not just designing of the circuits?

Dennard: Right. There were solid state physicists and materials specialists and...

Hendrie: And you had a lab that you could actually fabricate things and all of that?

Dennard: Yes.

Hendrie: Well, that could really be interesting.

Dennard: It was very interesting. We ended up building and characterizing a lot of devices, coming up with a process that was pretty stable, learning how to design some of the basic logic circuits and building some small integrated circuit chips. A 55 circuit chip was one of the milestones that we put together. Then we also wanted to get into the memory thing because IBM had pretty well decided that magnetic memory was kind of near the end of the road and that semiconductor memory would be a good thing to replace it.

Hendrie: Really? Interesting. Now, what timeframe is this, do you remember? What year are we in that you sort of shifted into this or that the whole microelectronics program got started? Did you move into that group when it got started or had they been doing some work on this N-MOS process for awhile?

Dennard: I got involved pretty much near the beginning, certainly before the process had been stabilized. <laughter>

Hendrie: You couldn't make anything that worked initially.

Dennard: No, not ever two things that worked the same way for quite awhile. And there were a lot of stability problems.

Hendrie: What year would you say this was?

Dennard: I got involved maybe in 1964 or '65. It was in 1966 that we really began to seriously look at memory and how to do memory in this technology.

Hendrie: Very interesting.

Dennard: My boss gave me the responsibility to start designing some memory prototypes and to try to figure out the best way to do memory. There was an approach using a six transistor memory cell that had been proposed by Lew Terman. They had done a lot of circuit work to show it could be made very fast using small sense signals on the bit lines. Those sense lines could be wired out to bipolar support circuits, which could sense very small signal levels very fast. That combination would give a very fast access time.

Hendrie: Was this sort of an MOS implementation of a six transistor, bipolar memory cell, that could be done with ordinary transistors, a six transistor flip flop, you know? Or was it a unique circuit...

Dennard: The circuits are actually quite different. The bipolar transistors were on the preferred stream, I think, and it was assumed they would be the ones that would win. The early versions, the 16-bit chip, had been done with bipolar, not MOS There was a special bipolar circuit called the Farber-Schlig Circuit, which was developed here by fellow researchers. You know, the bipolar transistor and the MOS transistor are quite different when you start putting together circuits.

Hendrie: Yes, I'm aware of that.

Dennard: There were certain attributes of those bipolar circuits and how they were sensed that were quite different than sensing MOS.

Hendrie: MOS was different.

Hendrie: Now, other people were working on gates and logic elements but you were trying to think about memory, am I projecting something that wasn't true or...?

Dennard: I think I'll just retrace my story: we learned how to do the logic elements first and we pretty much understood that and had built a small functional 55 circuit logic chip. We turned our attention towards this memory project to see what we could do about the memory problem.

Dennard: And so this was really the atmosphere at the time that I developed this new concept of how to do memory. It was much more simple, using what is today called the DRAM cell, a single transistor and a capacitor to store the information. That occurred to me one day quite out of the blue, after I had attended a research conference where the guys who were developing the next generation magnetic memory described their latest plan...

Hendrie: This was an internal IBM conference?

Dennard: Yes, internal here in Yorktown Research. They described a plan to build some kind of laminated memory board where they would use etched copper lines in a laminated memory board with deposited and etched thin magnetic films at the intersection of the copper lines.

Hendrie: I remember people working on thin film magnetic memories.

Dennard: Yes. They described their thing, which had a very large array with lots of bits. I forget how many but what impressed me about it was that it seemed that it was a very simple, basic technique compared to what we were doing, it was just very basic. I went home that night wondering whether we couldn't do something as simple and basic in our technology. I thought about using capacitors as memory elements, which was very basic. And, when I thought through how to store the charge and read it out later, I got the idea to put a voltage into a data line and steer it through a transistor into this capacitor and turn the transistor off and leave the charge on the capacitor. The writing was very simple. The reading was a little bit more challenging. The first idea was that the capacitor would just be another transistor and then I could read out the conductivity of this transistor.

Hendrie: So the capacitor would just be the base of another...

Dennard: The gate.

Hendrie: Yeah, the gate. Yes.

Dennard: The gate of an MOS transistor.

Hendrie: Right, it would be the gate.

Dennard: And I could connect that to a sense line. But I thought that was kind of complicated. Also, just using those two transistors, it was very hard to read because, if it was indeed turned on by the signal, then it would be a very low impedance and, if I had a lot of low impedance transistors on a common bit line, it would load down the line. I was trying to figure out how to actually do that. It ended up I would either have to add another transistor to isolate this particular cell from the other cells and have another select line for reading, which means what I'm describing here is a three transistor cell..

Dennard: But this was getting more complicated and I wanted it to be simple. So I started looking at various ways, could I really do it with two transistors? It was very difficult. And then finally I found sometime one night, (I was working on it pretty much at odd hours) I discovered that I could read the charge out of this capacitor back through the same transistor that had written it and make a small signal on the bit line, the data line that had been used for writing the information. I could get at least a small signal there. I was familiar with how to sense small signals from my static work on the six transistor memory cell so I thought that was very logical and very possible. I got very excited. I really had simplified this thing to a single transistor and a capacitor and I felt that was probably as simple as it's going to get.

Hendrie: Very good.

Dennard: So we started studying it. I did a paper, an internal paper at that time, about comparing this approach to the other more complicated approaches and it came out very good. It used very low power and achieved quite decent speeds and, of course, potentially many more bits per chip.

Hendrie: That's a great story. I love it.

END OF TAPE 2

START OF TAPE 3

Hendrie: You were just discussing what you did after you came up with the idea, sort of made a little comparison and presented it to your colleagues.

Dennard: Yes. I think this comparison was really done because there was interest by some government defense agencies that supported research in this area. I think we presented this information to them, which was kind of like releasing it, in a way, to the public. It's kind of classified and all that, but they certainly were moving around all of the other people working in microelectronics, and discussing ideas like this with them.

Hendrie: Was this NSA that you were doing work with?

Dennard: I think NSA, yeah.

Hendrie: They were a huge supporter of computer research.

Dennard: They were stimulating work in this area. I gave this talk at an internal IBM conference, where it fell on somewhat deaf ears, I think, because I didn't get much feedback at that time on it. Maybe the idea of a memory cell that you had to refresh after a small fraction of a second was not initially appealing. A lot of my associates were kind of skeptical about it and thought that static RAM would eventually (now that we have the term dynamic RAM, the six-transistor cell became known as static RAM.)be a much easier thing to do. They're right, but they also weren't right. I went to one of the upper levels of management here and suggested that this was a very important idea that we should try to rush it into a patent application. That was accepted and we did that, filed a patent pretty early, within five or six months of the invention. But then most of the processing technology at that time really wasn't ready for this-- at least the technology we were building. The N-FET had a lot of leakage challenges, so it really wasn't ready to build MOS DRAM. The six-transistor cell was a much safer thing, so we chose to design a test site using the six-transistor MOS RAM cell to be the prototype for our entry into memory with MOS technology. I actually designed a 512-bit cross-section static RAM chip and came up with some of the architecture using the balanced bitlines, the two-bitlines, using a way to form an array and select various bitlines, and to decode that array and to read out the data. We built that test site, and it worked great. It had something like 50 nanosecond access time. That pretty much completed our work on the MOS technology and this prototype of what the principal product was going to be.

Hendrie: So you actually built a 512-bit chip and fabbed it and got it running? You said a test site. Could you clarify what you did?

Dennard: A test site is a cross-section of a-- so we had a fully populated set of bitlines and a fully populated set of word lines.

Hendrie: But not all the other quadrants, yes.

Dennard: Not the whole thing right.

Hendrie: Not enough to demonstrate the...

Dennard: Enough to demonstrate all of the functions working the way they would work at real speed.

Hendrie: And issues of parasitics and all of the stray capacitance. All of the things that could affect you..?

Dennard: Yes. We put these chips into modules, and I worked out ways to shield the low-voltage lines going out to the sense amplifiers, to keep those free from influence by the large signals that were being used to drive the array—those were like 12-volt circuits in those days.

Hendrie: That's right. N-channel, yes, was 12 volts, or 15 volts. Yeah, it was high voltage.

Dennard: Very high-voltage circuits. So we had some 12-volt signals on this module and then several hundred millivolt signals going out to be sensed by the bipolar sense amplifiers. So a kind of system was put together, at least, with a lot of modules and a lot of support circuits, and tested. The technology was pretty well qualified, so that technology was turned over to our development labs for further development.

Hendrie: So would that be Fishkill?

Dennard: Fishkill and also in Burlington.

Hendrie: Oh, up in Burlington, too. You had MOS processes at that point in both of those facilities?

Dennard: Burlington became the principal MOS... Well, actually, I shouldn't say that. It actually was a more international thing. The Boeblingen Lab in Germany actually—or maybe the factory there was actually the principal factory where the MOS technology was first manufactured.

Hendrie: Really?

Dennard: Yes. We eventually manufactured one-kilobit chips and two-kilobit chips at that location. And some associated logic programs were started.

Hendrie: With, fundamentally, the six-transistor MOS static RAM scheme.

Dennard: Yes.

Hendrie: Was any work ever done on trying to integrate the sense amps, to build an MOS sense amp onto the chip so that you could get rid of the...

Dennard: Oh, I think those products at that time probably used—they used on-chip sensing. I think it was under that first prototype...

Hendrie: How far did you get in terms of actually doing any fabrication work and trying the single-cell?

Dennard: That's going to be another story. At the time of the invention, I simply did a prototype using discrete components and tested a simple circuit.

Hendrie: You were just mentioning the microelectronics program, or at least the program you were involved in at that point, was shutting down, after you had delivered the prototype, or the demonstration, of the MOS six-transistor SRAM.

Dennard: We may have some overlap here. I'm not sure how that's going to work out. So why don't you ask a question, something like—I think we probably finished that one. So you ask a question.

Hendrie: How did you finish up this project, before you wound it down?

Dennard: Well, as might be expected, there was a lot of documentation. Actually, I don't think I expected it, but my boss said we're going to write a very thorough treatise on what this program was and what we learned. So we started actually a very intensive effort to do that. First of all, we took a lot of data on the MOS transistors that we were, by that time, able to fabricate with very good characteristics. We made a lot of measurements, and we put measurements and models and design work and everything that we leaned in the program together in a really large book that was maybe an inch thick. It became basically later on the bible of MOS technology in IBM. That was really a substantial amount of work that we transferred, and the people in the development divisions were able to use those, and they were quite successful. Now, at the end of this period, we really got transferred-- we totally shut down the line that was manufacturing or producing these transistor chips, and a lot of us got reassigned to another data transmission or communications kind of a program. We were working with some of the first kind of token ring architectures and things like that, and we were certainly figuring out how to use our MOS technology in these kinds of things. We had a lot of collaborative work with our groups in Raleigh, North Carolina, who were doing a lot of communications products.

Now, the next phase of activity came in the early 1970s when a lot of us were organized into a new research program that was aimed at a very aggressive goal. Someone came up with a target for reducing memory cost to 1 millicent per bit, which was several hundred times less expensive than what was the going rate for memory bits then. It required a great deal of effort. It was based on the idea that we had some really advanced electron beam lithography work and very good optical groups, and we thought we would be able to do integrated circuits at much, much smaller dimensions. To get to that low cost, we also figured that we would use the simplest form of memory, which would be this one-transistor DRAM cell, which up to that time had not been used in any products. Although some early papers were available. The first paper really about the design of a one-transistor DRAM chip-- came out at about 1971, I think. We were trying to really take advantage of this one-transistor idea. I was manager of the transistor design group that was given the job of designing transistors, and to do this application, miniaturize transistors. When we did some sizings and figured out what it would require to get 1 millicent per bit, it turned out that we would have to greatly reduce dimensions. The prevailing dimension at that time was about 5 microns-- the most aggressive stuff we did-- We would have to reduce them down to 1 micron, a factor of five reduction in size, which would give 25 times more transistors per unit area.

Hendrie: When you say 1 micron, are you talking about the gate width, or the minimum dimension?

Dennard: The typical integrated circuit dimension, the minimum dimension, would be 1 micron for a line space or a separation between two lines, between conductors, for example. Or, in the case of the MOS device, it was the channel length, the separation between source and drain. Some of the contact hole dimensions would also be minimum dimension. Those were the basic steps in the MOS technology. You call these layout dimensions of the integrated circuit-- the wires and the transistor elements themselves. Our group had the job to figure out how to design these reduced size devices. We were trying to take a pretty big step, and we knew that when we made these devices, if we took our existing devices and just pushed them down to these dimensions, they wouldn't turn on and off properly at all. We knew that as we just reduced the dimensions just a little bit, the thresholds would drop, and the depletion layers around the source and drain tended to start merging with each other, and pretty soon punch-through occurred-- the transistor wouldn't cut off at all. So we started thinking about how to design this.

Hendrie: How you are going to get from here to there?

Dennard: We knew the depletion layers were going to be part of the problem, and we thought about that, and we finally figured out that to reduce the depletion layers, one way would be to reduce the voltages. Increasing the doping in the substrate would be another way. When we really looked at it a little bit, we realized that there was something going on here, that we could produce some kind of scaled result, that if we just reduced all the dimensions of the device and we reduced the voltage in the same proportion, and we increased the doping in the same proportion, that the depletion layers would scale down just perfectly with the other dimensions, and the device should have the same electric field properties and it should operate basically the same way-- identical kinds of characteristics. We could express the results in a kind of non-dimensional way, in terms of dimensional analysis. We could plot normalized characteristics for this device that would look exactly the same as characteristics of the big devices, except in a normalized fashion they would lay right on top of each other. Actually, we were pretty excited by that prospect, and we had an associate, Hwa Yu, who actually built some small one-micron devices. When we measured them-- lo and behold-- when we normalized them, the characteristics were perfectly the same. Now, by normalizing, I mean that we measure the characteristics on a voltage scale that's reduced by the same amount as a scaling Factor. And there's also some normalization on the current scale, which I can't remember right now, but that's beside the point. The exciting thing is when we really examined the circuit behavior, we found that that would scale also, that circuits would behave in a very identical fashion, except that when plotted on the time scale the results would be much faster. In other words, the time scale was being compressed by the smaller dimensions. Of course at the lower voltage we were operating these devices, we also got much lower power consumption per device. In fact, we computed that the power density of it would remain the same for this constant electric field scaling that we were doing, so that a one-watt chip would still be a one-watt chip even though you could cram 25 times more transistors into it. And they were operating five times faster. That was a very impressive result. So we published a paper on this, along with experimental results of these scaled devices, in 1972. In 1972, we actually gave a paper at the IEDM about this. I kind of recall—I sometimes think I must have dreamed this—but I recall that when I gave this talk and talked about scaling the gate insulators of these MOS transistors from 1000 angstroms to 200 angstroms, I really recall that a lot of people in the audience laughed. Because the idea of building such ridiculously thin gate insulators was unthinkable. They all knew that 1000 angstroms—you couldn't go below that and get any reliability. But of course we were scaling the voltages down, too.

Hendrie: Exactly. And they weren't thinking about that.

Dennard: So the scaling caught on pretty fast though. In spite of those early laughs, people started doing this, and I started hearing a lot of references. We eventually wrote up the work. The IEDM only published abstracts of papers at that time. I coauthored a paper with Alec Broers in 1973 at an ECS symposium in Chicago, which disseminated all of the results. Then we wrote a paper in 1974 and published in the *Solid-State Circuits* journal, which became by far the most referenced paper in the history of that journal, three-times more referenced than anything else that's been published there. That is how we got into the business of scaling. Interestingly enough, we didn't just scale devices. When I look back at the history, I found that after we published in December 1972—the scaling paper—at the ISCC that came a few months later in early '73, we showed some results of one-transistor DRAM cells built using these small device dimensions. When I read the literature for that year's Solid-State Circuits conference, I found that there was an early one-transistor manufacturing chip being planned by, I believe, Mostek, that appeared in that publication. It was still at the design level, I believe. There was some fabrication, but I think it was very much in the laboratory. Our transistors were about 20 times smaller in dimension than-- but we only had a small array of transistors-- about 20 times smaller in dimension than what was planned to be manufactured by others at that time.

Hendrie: Was your ability to make the very small dimensions primarily due to your electron beam ability that other people didn't have? I don't know whether they had it at the time.

Dennard: Maybe you could say it was a laboratory thing, where we took advantage of this electron beam lithography. Now, not much later, we decided on a much more ambitious thing, and we started building an 8000-bit array of memory cells using this electron beam technique. We completed that and published the results in 1975 at an ECS conference. Dr. Yu was the main author of that paper, and he did some wonderful work. In the course of that, he really used reactive ion etching for the first time in integrated circuits. It had to be the first time, because he put the name-- reactive ion etching-- in that paper, which is still used today for that process. He showed how to etch the polysilicon gates with totally straight lines with no bias, no change in line width. In those days, we used some fancy techniques called metal liftoff. Aluminum was the metallurgy, and we had a liftoff process so the aluminum was deposited into the place where the photo-resist had been removed with the electron beam lithography. That also was a non-biased-- it was a process that could make very small lines, basically square lines. It was amazing, the alignment in those, with the electron beam system was just perfect. You could not tell that anything was out of line. They were just perfectly matched, each mask level to the preceding levels.

Hendrie: Really? Oh, that's amazing.

Dennard: When I look back on it, I think, "Well, so what happened?" So about 1975, at the time we had done this, was the time that DRAM-- now called DRAM, by the way, for the first time, with an initial "D" rather than "dynamic RAM." The initial "D" in front of RAM was becoming used for the one-transistor cell at the time the one-transistor cells came out. The history of one-transistor manufacturing was that the first dynamic RAMS were some dynamic RAMs using even eight-transistor memory cells, I believe. There was a company called Dynamic Memory, I believe.

Hendrie: Really? Okay. I have not heard of that.

Dennard: They were not very successful with eight transistors per bit, because Intel, as you know, came out with a three-transistor dynamic RAM memory cell in the early ‘70s, which was really quite successful and quite a good product. I should explain that IBM in this same period, in the early ‘70s, began shipping tremendous quantities of static RAMs to populate all of our large systems. The static RAM chips, which had 1K-- well, first they were the bipolar static RAMs, which were, I think, 512 bits per chip. Later, in IBM, we went to the MOS technology for 1K and 2K-bit chips.

Hendrie: But still static.

Dennard: Still static RAMs.

Hendrie: Were they fundamentally using the same six-transistor array, or cell?

Dennard: Yes, they were. There was another phase-- I don’t know at the 4K level or 8K level-- where we used a four-transistor DRAM cell, which is not very well known to the rest of the world, I think. It was an internal IBM product that was invented by Lew Terman and Dominic Spampinato. This was after the one-transistor DRAM. It’s not that we progressed from six to four to three to one, you understand.

Hendrie: I understand. Right. You went six, one, and then...

Dennard: Six, one-- well, conceptually, at least. So the maybe 8K dynamic RAMs using the four-transistor memory cell, which was, if you think about it, really quite a good topology, because it had balanced bitlines; it had balanced sensing lines. You had differential signals, and you had very, very high-speed operation with the four-transistor DRAM cell. It was kind of a forerunner in topology to the folded bitline one-transistor DRAM architecture, which came later but has the same advantages of balanced sensing.

Hendrie: Balanced. Then you can use the differential amplifier.

Dennard: During the time we were doing this miniaturization work, we also were working with our development labs on trying to help them make the transition from the static RAMs into dynamic RAMs. This 8K bit chip that we scaled to 1 micron was in fact one of the prototypes that was being developed in our product division. It was the same chip; we just took it and scaled it, just reduced all the dimensions.

Hendrie: Took the same circuit design, including the drivers, and the integrated amplifiers and everything.

Dennard: Because everything had become digitized by that time. The masks were being produced by digital techniques. So we just took the descriptions for the artwork and changed the scale and programmed the electron beam to run this thing. It was very easy to do.

Hendrie: That's very good.

Dennard: But as I was saying, we were planning how to get into the one-transistor DRAM business. Eventually there was a proposal to do a 64-kilobit chip, which really is what was the first IBM entry into the business. There was a lot of work to get the sense amplifier schemes for that working well. Because now we're getting to lots of sense amplifiers on a 64K-bit chip. Up to then, they hadn't been so concerned about power, and those sense amplifiers tended to have a lot of active current in them.

Hendrie: They were linear circuits.

Dennard: Sometimes they were balanced in the active state with both inverters turned on, just shoving the outputs together, which was kind of power inefficient. So we had to come up with some real dynamic sense amplifier techniques, where they were operated strictly with no DC paths. That was a very interesting development. One of the persons working with me came up with something called a charge transfer technique that was used in some of the first sense amplifiers, certainly in the 64K-bit product.

Hendrie: So you worked on a number of the difficult problems that needed to be solved to get dynamic RAMs from a research project to a product.

Dennard: Yes. We also came up with the idea of bitline and wordline redundancy. We knew that the chips with all of this complexity-- one defect would cause the chip to be bad and not be shipped. So we worked on the concepts of putting in spare bitlines and wordlines, and having some fusible links to reconnect, reconfigure the chip after test to eliminate bad bitlines and bad wordlines. When the yield is poor, this can improve the yield really remarkably.

Hendrie: Do you know whether that technique ever went into production?

Dennard: That technique was announced at the 1979 ISCC in the IBM 64K-bit chip. In the same session, there was a Bell Labs 64K-bit chip that also used that technique. Bell Labs used laser cutting of polysilicon links, and IBM used high-current, blowing-out fuses electrically, to personalize these chips. Those two papers came totally independently with no knowledge of either group that the other group was working on it.

Hendrie: That's pretty interesting.

Dennard: Yes. Redundancy is used, and has been used since then in I believe all DRAM chips, and SRAM chips as well.

Hendrie: So it became adopted.

Dennard: Even integrated SRAM arrays built into microprocessors and other highly integrated chips use these redundancy techniques today.

Hendrie: Because it's just such a sensible thing to do.

Dennard: I want to mention the name of my associate, Stan Schuster, who worked out and wrote a paper describing quantitatively how much yield improvement you get from redundant circuits based on the Poisson distributions of defect densities.

Hendrie: Wow, all right. Very good. We should maybe do a quick pause. We have to change a tape.

END OF TAPE 3

START OF TAPE 4

Hendrie: Do you remember what the first IBM system to use the single-transistor cell RAM—which I believe you had mentioned was the 64k—was? The first IBM single-transistor system product?

Dennard: Yes, that was announced, you know, I said ’79; I think it could have been the ’77 ISCC. I’m having a little trouble with the dates there. But the first product, as far as I have been told, was the system 8800. I’m not really familiar with that system, it’s a small system I gather and probably maybe a slower speed one that took advantage of the somewhat slower speed of our one transistor DRAMs. They were not aimed at high speed at all. The Static RAMs had access times something like 50 to 60 nanoseconds, and although the DRAM products could have been capable of that, they deliberately produced access times more than 10 times slower than that because they were sufficient for main memory. By that time, we had developed the cache technologies to the point that the DRAM was some several cache levels removed from the main processing. So, they really didn’t require very fast speeds as long as they were very, very much faster than the hard drives, obviously, and made a very nice match to the system speeds. I think it was an 8800 system. I don’t believe I ever saw one, but I believe I remember those numbers. That would be the first. Of course by now it’s used in all the systems and probably got that way pretty quickly after the first few generations of DRAM.

Hendrie: Are there any other interesting stories that you might recall during this era of the early single transistor DRAMS?

Dennard: Well, not too far along into the DRAM evolution, a very interesting thing happened. Intel suddenly announced to the world that they had found out that something called “soft errors” had been happening in DRAM products. I believe it could have been in the 64k-bit chips or maybe the next generation after that. They had done some measurements and found some bits being dropped. They figured out that it was due to ionizing radiation, and due in particular to alpha particles being emitted from the package in which the chip was sitting. Radioactive materials were contaminating some of the ceramics. That was a very exciting time for me, because I felt some personal responsibility for DRAMs. So we got involved very quickly to try to understand what was going on. A guy in my group, George Si-Halasz developed some analytical techniques for understanding in a kind of quantitative way how this was happening; how the collection of these ionized excess carriers took place and what we could try to do at the device level to minimize the problem. And, of course, Intel also apparently did a lot of work themselves and should be commended for finding and explaining this phenomenon very well. Actually, after a flurry of activity, it became, for a while at least, kind of a non-issue because it was relatively simple to put in some kind of coating layers to protect the chip from the particles. The particles didn’t travel very far and they would be absorbed by a few hundred microns of material so the chip could be coated and protected. It became later on a problem now, that’s reoccurred in CMOS technology. It turns out that some of the packaging techniques we used which were what you call C4—I think it’s Controlled Collapse of solder joints—also had the problem. Solder, being made of lead, which is a very heavy element and so it has some isotopes of radioactive material inherently involved in most lead. And so to the degree that we've continued to use lead, we've continued to have to keep investigating soft errors. As we scale to smaller and smaller dimensions and lower and lower voltages, soft errors have become almost, well impossible, to prevent. Nowadays it's more like finding ways to correct them by having some rad-hard circuit techniques, or error correction techniques to take care of occasional loss of information. That first discovery was an very exciting time.

Hendrie: Did you stay working in MOS technology and/or memory?

Dennard: I was very fortunate to keep working in MOS technology, you know, the rest of my career and pretty much to be working on scaling. And a lot of the time it was pretty easy, but each new generation – brings new challenges. What generally happens is that some things that were negligible become not-so negligible as you continue to scale, and cause some new problem. Resistivity is an example of that. As we scaled further, we had to eventually use more metallic kinds of process steps. We had to put in silicides, for example. It had more metallic conducting and better conductive properties than the basic semi-conductor that they're coating or interacting with. So we have those new kinds of process steps that had to come in. I worked on scaling for every generation from 1 micron, 0.5 micron, 0.25 micron…—We eventually came to understand, due to Moore's Law that told how fast scaling was expected to proceed, what all the future generations would be and have a roadmap at least for what the dimensions would be. I worked on all of those generations up to a point. Of course, some of the younger guys have taken over and we've developed at IBM the SOI technology, which was kind of unique development.

Hendrie: What does SOI stand for again?

Dennard: Silicon on Insulator. It's a relatively thin silicon layer on an insulator on a silicon substrate. It's a dielectrically isolated device. That was a development that has made a big difference and generally has been applied to our IBM products for quite a long time. I want to just say that during this work, I've been always mostly looking ahead at some number of future generations in trying to anticipate what problems will arise in getting the new lower operating voltages. Getting people to really get comfortable with the idea of scaling of voltages has been one of the most difficult things.

Hendrie: What's your analysis of why people have a problem with that?

Dennard: Well it was because when PCs became a product there was kind of a standard interface that developed, which was 5 volts. People designed a lot of peripheral equipment to come together and plug into the PC, and the interfaces were all defined. Everybody thought it was going to be that way, like forever, it appeared. There was another interesting phase when CMOS had scaled down, and really was threatening to take over the mainframe business, and 5-volt bipolars were about to disappear. There was also the idea of combining CMOS and bipolars into this technology called BiCMOS that was an interesting perturbation. That was mainly an attempt to maintain this high voltage interface at 5 volts or maybe 3.3 volts which is compatible with the lower level of the TTL interface back then. So, breaking through that and finally showing that to scale CMOS down into quarter micron channel lengths, you absolutely needed to use these lower voltages to maintain reasonable power levels. That was a key step in the sense that once we got to 2.5 volts, then going below to 1.8, 1.5 and 1 volt, which is pretty much where we are today, happened rather earlier --

Hendrie: It didn't matter so much because you destroyed the interface to TTL.

Dennard: We've broken the barrier.

Hendrie: You've broken the barrier, yes.

Dennard: So that was a very interesting period there. The interesting aspect of scaling, that was one of the ones where I spent a lot of effort and time investigating voltage scaling issues, was whether or not noise would be a problem? A lot of people thought, "Oh, noise would be a problem." No, the noise is generated by the voltage. And as you lower the voltage, the noise gets lower too. So it really is not a problem.

Hendrie: Any other comments or things you'd like to talk about during your career that you think our audience might find interesting?

Dennard: Well, I'm not sure. Nothing pretty much occurs to me– Should we wrap it up now?

Hendrie: Do you have any advice for young people who are interested in science and technology, who seem to have a bent in that direction?

Dennard: Yes, yes I do. I'd like to say that I've had a long career, and one thing that's really interesting from my present perspective is to look back at what things were like, 40 to 45 years ago and think about how the things that we've done in that time period in my career and in the field in which I'm involved, in microelectronics, have really completely changed the world—a lot of really astounding things have happened. We have capability that we take for granted today that was just really amazing when you think about it. And so the question for these young people is, if you're just starting in this field now, well just think about what things might be like 40 years from now and what you would like to see the world like 40 years down the road. I think that things will change, but the interesting thing is that these young people have a chance to really participate in the change and really be the enabling element for this change. It means that there is opportunity there. I mean, these things don't happen by themselves. It takes real people, you know, making these breakthroughs. I think the key for the people if they want to be involved in something like this is that they need to have the attitude that they really can contribute to something totally new. And they can. It's not like the invention and new kinds of ideas are reserved for any special class of people or something that's reserved for a few. I mean, anybody can participate in this. I think my very humble upbringing, which was fully described earlier here certainly establishes the fact that I'm a very ordinary person, with a very ordinary background and upbringing. Somehow now you have to get the training to do this stuff. It's not enough to just think creatively. Once you've posed the question, you've got to answer the question.

Hendrie: You have to demonstrate how to do it.

Dennard: You've got to come up with something. Sometimes those solutions are somewhat technical, they're not necessarily complicated, but they're technical and they involve technical content. So, you have to have the background and you have to do the preliminary work. Hopefully, it's the kind of work that is done in certain institutions that tend to have the mission and give the support that someone like me at IBM has had. Anyways, that's some advice for the young people. Attitude is the primary thing, I think, because I can see that a lot of people that don't succeed don't really think in terms of succeeding. They're kind of content to do their daily job, the assigned job. But sometimes, you know, it's the job that's not really quite assigned that leads to the new things.

Hendrie: Well, thank you very much. Is there anything else you'd like to say?

Dennard: Well, no. It's been a great career for me. I'm very thankful to IBM, where I've been here for 51 years, for providing this environment where I could work and be productive. There's a lot of teamwork involved in things like this, although I might have talked about individual contributions that I've made. I think as I described it also there are lots of people involved. I've learned from a lot of people and now I'm busy teaching a lot of other people, which is my principal thing today. So it's been a great career and I'm still enjoying it very much, every day.

Hendrie: That's wonderful. Well thank you very much for doing this interview.

Dennard: Thank you.

END OF TAPE 4

END OF INTERVIEW

