Reed Electronics Group

MICROPROCESSOR REPORT

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

GORDON MOORE AND CARVER MEAD

Two Pioneers Discuss Moore's Law and the Birth of an Industry

By Tom R. Halfhill {10/17/05-01}

To celebrate the 40th anniversary of Moore's law, the Computer History Museum in Silicon Valley invited Dr. Gordon Moore and Dr. Carver Mead to talk about the law, reminisce about Moore's distinguished career in the semiconductor industry, and discuss other topics. On the evening of September 29, the museum's auditorium filled to capacity with an eager crowd of museum members and guests. *Microprocessor Report* recorded and transcribed this special event.

Moore first made the observation known as Moore's law in an *Electronics* magazine article in 1965. But it was Mead, a professor at the California Institute of Technology (Caltech), who made Moore's law famous. Moore's original article didn't explicitly state the law in layperson's language and never referred to it as a law. It was Mead who boiled down Moore's technical observations about circuit integration and coined the term "Moore's law." Since then, Moore's law has captured the popular imagination to the extent that it's often misquoted and misapplied to other technologies. (See *MPR 12/13/04-02*, "Viewpoint: The Mythology of Moore's Law.")

Mead and Moore go back a long way. Mead received his Ph.D. in electrical engineering at Caltech in 1959 and taught there for more than 40 years. Mead first met Moore in the late 1950s, and they have been friends ever since. In 1980, Mead and Lynn Conway wrote *Introduction to VLSI Systems*, one of the most important textbooks in the field of IC design. Mead is also the founder and chairman of Foveon.

Moore began his long career in the semiconductor industry in the 1950s at Shockley Semiconductor Laboratory, headed by the mercurial William Shockley, who won a Nobel Prize for coinventing the transistor. Irritated by Shockley's abrasive management style, Moore and seven other engineers (dubbed the Traitorous Eight) left to found Fairchild Semiconductor in 1957. Later, Moore joined with other engineers to found Intel in 1968.

The anniversary event at the Computer History Museum was hosted by Dave House, a museum trustee. House, now retired, spent 22 years at Intel. Among other things, he managed the team that developed the successful "Intel Inside" marketing campaign for the Pentium processor.

Moore's Explosive Childhood

CARVER MEAD: Well, Gordon, it comes as somewhat of a surprise to many people who are interested in computation that the man who's probably done more to influence the face of modern computation than anyone else is a chemist. It's interesting. How did you get started being interested in chemistry?

GORDON MOORE: Well, being interested in chemistry happened when my next-door neighbor got a chemistry set when I was about 11 years old. In those days, you'd get some really neat stuff in them. *[Audience laughter]*

MEAD: Yeah, they were better than they are now.

MOORE: Oh, you can't get the stuff any more. And, frankly, I got interested in the interesting things you could do—the bangs, the smokes, and one thing or another—and decided very early, at 11 or 12, that I wanted to be a chemist, not knowing quite what one was, but at least you got to play with those fun materials. And I followed that route, you know. I toyed briefly with becoming a math major when I saw the simplicity of elementary calculus, but except for that, I struck pretty much to the chemistry end of things.

MEAD: So, early on you were making flash powders and gunpowder and things like that. Anything more in those early days—any memorable experiences?

MOORE: Well, you know, I developed along the route of finally culminating and turning out small production quantities of nitroglycerin, which I made into dynamite. A couple ounces of dynamite makes an absolutely fantastic firecracker. You can take a limb that big around and put one on the side of it, and the limb disappears. [Moore holds up his hands and wiggles his fingers.] I got all ten of 'em still. *[Audience laughter]*

So I maintained an interest in chemistry, had my home lab where I did turn out nitroglycerin. I'm using your hearing aids tonight because I blew my ears out as a kid. [Mead was a founder of Sonic Innovations, which developed a digital hearing aid.] You put a drop of nitroglycerin and a piece of filter paper on an anvil and hit it with a hammer, it makes the sharpest crack, and my ears would ring like heck for an hour after doing that. And I recognize now that I just blew them out then. So thanks for helping me be able to hear my wife again.

MEAD: Terrific. So, you ever make rockets or anything like that?

MOORE: Oh, sure.

MEAD: Any interesting incidents?

MOORE: Well, once a piece of our rocket fell on the neighbor's roof while it was still smoking. And we grabbed the hose and went up there and doused it, but the neighbor didn't take too kindly to that and called the Redwood City police. I grew up in Redwood City. And my father was a deputy sheriff. So the policeman—who my father knew very well, and I had a nodding acquaintance with—came out and gave me a lecture, and one thing or another, but that was about my worst experience with rockets. Making successful rockets was difficult.

MEAD: They're a lot harder than making the thing that just goes bang.

MOORE: Yeah.

MEAD: Did you ever end up with experiences on the other side of the law? I mean, you had an expertise that might have been helpful.

MOORE: Oh, Carver must have been reading up on some background. My father was a deputy sheriff, thank you; he was chief deputy sheriff in San Mateo County for many years. One time they found a cache of tools, you know, pullers and drills and one thing or another, and it had a little bottle of yellow liquid in it. It looked like the kind of stuff you used for pulling safes apart, and they were concerned about what that yellow liquid was. So he called me down and asked if I could help him. I said sure. So I went down there and put the drop on the filter paper and hit it with a hammer—*bang!* And then they were really concerned about what it was. They didn't know what to do with it. "Oh, I'll take care of it!" And I went out to destroy the evidence. *[Audience laughter]*

MEAD: So that got you into chemistry then.

MOORE: That did. Yeah, it kept my interest. In fact, when I was getting my Ph.D. at Caltech, one thing you had to do for your final exam was to come up with ten propositions of things that were potential research topics and defend why they were good, and you usually did nine of these seriously and one a little less so. So I did a survey of the grad students in Caltech's chemistry department to find out how many of them got their early interest in chemistry through explosives, and over 80% did. So my proposition was that we should encourage that in order to help alleviate the shortage of scientists.

Trying Times at Shockley Labs

MEAD: So after you got your Ph.D., what happened then?

MOORE: Oh, I had to go to work. Betty put me through school, and got her Ph.T. signed by Mrs. DuBridge—

MEAD: A lot of people don't know that [Ph.T.] stands for "Putting Hubby Through."

MOORE: Interestingly enough, at that time, it was hard to find a decent technical job on the West Coast, so I had to go east for my first job. I went to the applied physics laboratory at Johns Hopkins for a couple of years, always anxious to get back to California if I could find a suitable job. And [William] Shockley was setting up out here, and I got a call one evening. He had gotten my name and thought he needed a chemist for his new operation, and fortunately, I recognized who it was. I had seen him give a lecture back there a month or two before. I had interviewed out here, actually at Lawrence Livermore, and they made me an offer, but I really wasn't interested in doing what they had in mind, and he got my name from them. And that's how I got into the semiconductor business. I knew very little about it.

MEAD: So what was your first assignment when you got to Shockley Labs?

MOORE: I'm not sure I specifically had an assignment. Shockley encouraged me to stop a couple of places on the way out. I went to Bell Labs for a couple of days. I stopped at the University of Illinois and visited, I think, Phil Handler. He wanted me to learn about surfaces, because surfaces were known to be a problem with semiconductors. I learned all I could swallow in a few days. I came out here and, you know, got involved in building equipment. The diffusion area was one that fell rather naturally in the kind of things I was used to doing. I've done a lot of technical glass blowing through my schooling and one thing or another, so I picked that up and starting building what we called glass jungles that went in front of all these furnaces and such. But there were a lot of things to do, and, of course, the most important thing was trying to learn something about silicon. I knew essentially nothing about it. I think [Intel cofounder] Bob Noyce was probably about the only one in the group that had any experience with semiconductors before we got to Shockley.

MEAD: So what was the experience like, working in that group with Shockley and those other smart, young people?

MOORE: You know, it was in the beginning, at least, a kind of standard startup experience. Except most of us weren't very prepared; we had a lot of learning to do. You have to recognize in those days hardly anybody knew anything about semiconductors, in particular silicon. Transistors around were germanium. Texas Instruments had some grown-junction transistors out, but Shockley focused on diffusion and silicon, and that's where we were trying to get enough knowledge to be able to make something. Initially he was going to make a transistor. Then he got concerned, I believe, because TI was already out there with a silicon transistor. [Shockley] decided he ought to make a device that didn't have any direct competition. He had invented something called a four-layered diode that he had high hopes for, so he kind of refocused the company's interest toward the four-layered diode. Some of us were still trying to make transistors, and any of you who have heard anything about the Shockley history recognize that we got into some management problems.

MPR

Dr. Gordon Moore

Shockley had unique management techniques. *[Audience laughter]* And a group of us went around him to Arnold Beckman, with Beckman Instruments, which was where the financing was coming from, and suggested that Shockley should be put in a position where he could be a consultant. Somebody else ought to be brought in to run the organization. Well, we were making progress, but we learned a bunch of young kids have a tough time pushing a new Nobel Prize winner aside. And Beckman became convinced that this would ruin Shockley's career, so essentially he ended up telling us Shockley's the boss, take it or leave it. We burned our bridges so badly by then, we felt we had to leave it. That was the origin of what turned out to be Fairchild Semiconductor.

Moore's Early Days at Fairchild

MEAD: So Fairchild then became the leading developer of all of these things. We used to call the little lab up there Fairchild University. There were a lot of graduates from there, Gordon, that populated Silicon Valley.

MOORE: It was a very fertile technology, and sort of every new idea that came along spawned a few companies. I heard somebody the other night say you could trace something like four hundred companies' origins to Fairchild.

MEAD: Wouldn't be at all surprised.

MOORE: That's a lot.

MEAD: So, you and I met around 1960.

MOORE: It was earlier than that.

MEAD: '59 probably.

MOORE: I heard you give a paper at a conference. It was your auto electrons through oxide kind of a device, and Carver was a professor at Caltech, and I was going down there. I walked into his office, and I think I asked him if he could use any transistors, wasn't it?

MEAD: I was teaching an undergraduate course in transistor electronics at the time. It had a lab with it. Of course, you're trying to teach people how use these things without any. The transistors you could get at the stockroom were just horrible things. So, that was a no-brainer.

MOORE: Yeah, I had a couple of big envelopes that we called cosmetic rejects. They didn't look good, but they still worked. And, I guess I populated your laboratory for a few years with those.

MEAD: We had a lot of very good projects come out of the lab because of those transistors. That was great. So then we started working together in various ways. I remember coming up essentially every week and visiting you at Fairchild. Lots of good discussions. There was one particular discussion I remember vividly about when you asked me how small a transistor could get. Do you remember that?

MOORE: I remember it rather vaguely, but I know you went away and figured it out. *[Audience laughter]*

MEAD: Spent the next 20 years of my life following that thread. That was an interesting time, because not everyone thought you could keep miniaturizing transistors.

MOORE: It still amazes me how far we've been able to go. To have a technology that hasn't really run across the problem that fundamentally stopped it, something that's changed as fast as this has is really amazing. Several times along the way, I think I've felt there were barriers that weren't out there very far. The closer we got to them, the further away the barriers went, or they just dissolved away. We make devices now that are really small. You know, it's nanotechnology by coming from the top down rather than the bottom up.

The Origin of Moore's Law

MEAD: I'd like to get back to the chemistry of all this, because Moore's law, in a way, is more about the production of devices than it is about the computers the devices go into. Do you want to talk just a little about the production of devices and how that really works? I'm not sure everyone here is very familiar with that.

MOORE: Well, let me take it slightly differently. There's a paper that [former Intel colleague and emcee] Dave [House] read from…I was being asked to write an article for their 35th anniversary edition [of *Electronics* magazine], and I was concerned that people thought integrated circuits were very expensive. They were, up until that time. You know, the military was the only one that could afford them. There were all kinds of arguments about why this was an expensive technology, but I could see in the laboratory things were really starting to change. I wanted to get the idea across that this was going be the way to make inexpensive electronics, not just tiny low-power stuff that the military was interested in. So that was kind of the motivation behind the article. And, what I could see was that our ability to pack stuff on a chip was going to continue to grow—partially because we were making things smaller, partially because we were decreasing our defect density. You could really see that there was a complexity. The number of transistors—the number of components in a chip that minimized the cost—was going to keep going more and more [toward] complex circuits.

So I just looked at the data we had for the first few years, you know, starting with one transistor, the kind in ’59 and ’61, the first integrated circuits and so forth, and saw that they had been about doubling every year. The next one coming had about 60 components on it. So I just took that doubling every year and extrapolated from 60 to 60,000 components for the next ten years, and said that’s what’s going happen—it’s going to make cheap electronics. I never had any idea it was going to be at all precise. You know, we were just trying to get the message across that by putting a lot more stuff on a chip, we were going to make much cheaper electronics.

MEAD: So during those early years, what was the mix of talents that were in your organization that were enabling that to happen?

MOORE: Well, not too different from the mix we have now. There were the people working on the processes, who in those days tended to be chemists and physicists rather than engineers. It’s just the way it started in the beginning. But we had a significant group that was doing circuit design, looking at decomposing systems in the logical kinds of blocks and the like, so we had a pretty good-size electronics group also. Hopefully, worked pretty closely with the process people, but not always.

MEAD: There was a time I remember you were plotting lots of things besides the number of transistors. You plotted die size, you plotted wafer size, you plotted design time. There were a number of things you kept track of during that period. Want to talk about those trends?

MOORE: Some of them were just breaking down the increases in complexity to see why it was happening. I think I decided there were three major contributors in those days. One, we were making things smaller, making the features smaller so you could pack more stuff in a given area. The second one, because you were lowering the defect density, you could make a bigger chip and get a reasonable yield. And the third one was, you were squeezing waste space out by being clever at the way things were designed and packed. Those were all making major contributions. And, as Dave said, in 1975 I changed my [Moore’s law] doubling every year to doubling every two years. The reason was, we had squeezed all the waste space out. The most complicated devices then were CCD memories, and this was a really active area right up against one another. There was no waste space left, so I figured we were going to lose that factor, we only have the other two going forward. It turned out CCD memories weren’t successful.

MEAD: Well, they turned into imagers.

MOORE: Yeah, for the same reason that they work fine in video cameras, they don’t work very well any place where they might get, say, a particle on them. [Former Intel sales manager] Ed Gelbach, sitting over here on the right, or my right, your left, remembers that well. We had memories that were making soft errors. They were making errors that wouldn‘t reproduce. And our customers were very concerned about that. Fortunately, we had CCD devices to study the phenomenon. It turned out it was the residual radioactivity of the packaging material. That just killed CCD devices as memories, but it did give us a good experimental gadget with which to solve the real problem.

MEAD: So down through those years, I remember you were also concerned about the design effort that went into these increasingly complex chips. I remember you had a little plot at one time that had the design effort going up sort of exponentially with a little moon at the top. So, talk a little to what happened there.

MOORE: Well, that was a concern. A variety of things happened. You know one of the most important things for the industry overall was actually a book that Carver and Conway—

MEAD: Lynn Conway.

MOORE: I couldn’t think of her first name—wrote on systematic design of CMOS circuits. [*Introduction to VLSI Systems*, 1980.] Now it trained a whole generation, at least a generation of engineers, who came out of school finally knowing how to design integrated circuits. But it turns out we never used it. *[Laughs]* We had our own way of doing these things, and it was hard to convert over. The way we did it at Intel resulted in somewhat smaller chips with the same amount of functionality, but with significantly larger effort. But Carver’s design technology certainly ended up, I guess, spawning the custom integrated-circuit business and continuing to nurture it. But the systematic design with a variety of different approaches has got us off that curve that was growing to the point where we wouldn’t be able to design the products we make today at all.

Designing Intel’s First Microprocessor

MEAD: I remember in the early days, Intel thought of itself more as a memory company. Now it’s more of a computer-chip company. How did that whole thing play out?

MOORE: Well, my view is slightly different than that. We thought of ourselves as a large-scale integration company. We set up to make complex chips. The problem the industry was having at the time Intel was formed is that you could build a much more complex chip than made sense. They tended to become unique when they got big, used only once in a computer or something like that. There were only

10,000 computers a year being built, and you couldn't get in every one of them, so the design cost was absolutely dominating. The fact that you could make the chip didn't make much difference. The net result was that successful semiconductor companies were the ones with large assembly plants in Southeast Asia. It got to the point where the piece of silicon cost less than the packaging material.

When we set up Intel, we wanted to go the other way. We wanted to find something where you could get the cost back into the silicon, if you wish, where a new company could compete. So we looked for complex circuits you could make in large volume. Semiconductor memory was the outstanding one we could see, because all digital systems needed memory of some sort. It looked like there were good opportunities there. So that's where we got started. But then we started looking for other examples of complex circuitry to build in large volume. And, it was just at the time when the electronic calculator was coming along.

Unfortunately, or maybe fortunately in retrospect, all of the established calculator companies had made a deal with the established semiconductor companies. So we were a little guy, a little bit late to the party, and we probably wouldn't have been chosen anyhow. But we got introduced to a startup [Busicom] in Japan—which is almost an oxymoron—that wanted to build a family of scientific and business calculators. And they had done all the electronic design. They had thirteen very complex circuits laid out. They were looking for someone to build them. We sat down and looked at them, and we really wanted to make these complex circuits as a second example of something that we could build in large volume. We had a little engineering group that was struggling to make the few memory chips we were doing, and we could no more take on making thirteen custom logic circuits than fly.

Fortunately, one of our engineers—Ted Hoff, who we hired from Stanford—we wanted somebody with significant computer experience when we brought him into company—he looked at those and said, "Gee, I think we could do all of these calculators with a general-purpose computer architecture, and I don't think it would be much more complex than the memory chips we're making now." And we thought, gee, that's fine. We'll only have one chip to do and a couple of memories rather than thirteen chips. We might be able to tackle that one. So we had to sell the Japanese on abandoning all the engineering they'd done and looking at a general-purpose computer architecture. I thought that was going to be a really tough sell. So they came over and visited, and I remember going in and starting the presentation, and the chief technical guy says, "Fine, we'll do it your way." *[Audience laughter]* So we got the order.

Then we still had to put together an engineering team to actually do it. And they came back several months later to see how much progress we had made, and it turns out that only the week before we had hired the engineer who was going to run the project, Frederico Fagin. So Frederico had a week to look at what he was going to do before he had to give his first progress report to the customer. Anyhow, we succeeded in making that, which became the first microprocessor, the 4004. And Busicom did start selling their calculators, but they were thinly financed and came around looking for a lower price. And we said, "Gee, the only way we can give you a lower price is if we have higher volume. Can we sell it to other people?" So we actually got the rights to sell this custom chip for noncalculator applications more broadly in the market, and that gave us an opportunity to really start trying to sell these chips.

MEAD: I remember when you donated some chips to me down at Caltech. I had to sign a little thing that said I wouldn't build calculators. *[Audience laughter]*

MOORE: Anyhow, a few months later, their financial problems worsened, and they were even willing to sell us complete rights to the chip for returning the money they had paid for the development, which was something like $65,000. So Intel got full control of the microprocessor back.

MEAD: What were some of the other early applications?

MOORE: The first applications of that were really peculiar. I remember one of our board members asking in one of the meetings, "When are you going to get a customer I've heard of?" *[Audience laughter]* We had someone in the Central Valley [California] automating a chicken house. I still haven't quite figured out what it did there. We had a blood analyzer and just a variety of miscellaneous things. It was just a general logic replacement. That was the advantage of a microprocessor. It really was a general-purpose logic function. You could program it to do about anything by putting a program in a read-only memory.

Moore's $15 Million Wristwatch

MEAD: Since that time, Intel has become very influential in the larger view of computing. You've made mainframes, you've made watches. Why don't you tell us the watch story?

MOORE: That's another example of a complex ricochet, thinking large volume. At least that was the idea when we got into it. I guess my view of that was if you're going to put a lot more functionality in a watch, not knowing quite what I expected, it was kind of like a PDA that I envisioned. We pushed that and found out we were in a different kind of business. When we got out of the business, the silicon content of the watch cost less than the push pins on the side of the case for setting it. It didn't go in the direction of a lot more functionality at all. People wanted it to tell time, maybe the date, and the month, and that was about it. So we can thank TI for getting us out of there fairly cheaply. We were selling $150 watches, or trying to, trying to make this some kind of prestige item. It told time more precisely than anything else. TI, one year, got into the business and said $19.95, and the next year they said $9.95. They've made it the cheapest way to tell time rather than the best way to tell time. So we got out of there. For years I wore what I called my fifteen-million-dollar watch to remind me to be careful about consumer products in the future. *[Audience laughter]*

MEAD: How about the other involvement in computation generally? Intel's had lots of ways of looking at that.

MOORE: We've done a variety of things along the way. One business that the microprocessor spawned is what we called development systems. They were actually special-purpose computers designed to help our customers design and debug their hardware and software. And for several years, essentially through most of the '70s, we did a significantly larger business in development systems than we did in microprocessors. But we were always assured by our sales department that every one of those development systems out there was a silent salesman that was going to result in at least $50,000 worth of components someday. It was a good business. Essentially, we were building an engineering workstation and didn't recognize it. We talked about pushing it harder in that direction, but we never got the energy together to do it.

Another Break for Intel: the IBM PC

MEAD: So then, of course, the PC happened, right?

MOORE: That's right.

MEAD: Talk us through that one.

MOORE: Well, this happened at a time when we were competing principally with Motorola to get our 16-bit processor established in the market. We had a major program called "Crush" internally. I'm not sure we can say that these days. That was designed to get as many design wins as we could for our 16-bit processor. I think we had a goal of 2,000 in '79 probably. We actually succeeded in getting closer to 3,000 design wins. We were very good. Our sales effort was great. Buried among those was a design at IBM where we didn't quite know what the application was, but design wins at IBM were always important. It was a big enough company that, you know, their secondary products were pretty important. And it turned out to be the PC. So they actually used the 8-bit bus version [8088] of our 16-bit processor [8086], which was cheaper to make than the 16-bit one. Since then, we've been riding the PC as it has developed. It's been a marvelous business.

MEAD: There were some important decisions made along the way. I remember you talking about backward compatibility one time.

MOORE: Well, we had a lot to learn, and I think we all did. I remember Dave House and I talking with [Philip "Don"] Estridge. Estridge [the IBM vice president who supervised the PC division] ran the IBM thing. We were talking about software. And Dave, being a computer guy, said, "Oh yeah, you recompile." Estridge put us pretty darn straight on that. You take it off the shelf, undo the shrink wrap, and it goes into the computer and it works. No such thing as recompiling. So we learned a few lessons along the way.

MEAD: I remember one time you telling me that because of the increase in complexity, the part that had to be backward compatible was always the small part of the new thing.

MOORE: That's absolutely true, yeah. It was very important that this software continue to work, so you had to make sure that you carried all the works and one thing and another from one generation to the next. You didn't have the opportunity of throwing away what you had done so far and starting all over. We tried that once on a parallel project, though. Fairly early in the 8080 days; it would have put it in the early '70s. We put together a team and told them to be unfettered by compatibility or anything; go out and do it right. We have one more chance. You got involved in the beginning of that one.

MEAD: That was Justin Ratner and those people, yeah.

MOORE: Bill Latin, the group, the product, or the project was the 432. It was sort of a 32-bit processor. It had every concept known to modern computer science buried in it. But it was the wrong approach. It turned out to be kind of a minicomputer put on a piece of silicon. It didn't take advantage of all the things that uniquely had made the PC successful. It resulted in being the subject of several computer-science classes around the country, but, as a product, it was a real failure. Some of the things that came out of it got incorporated in devices later down the road, but that one missed the boat completely.


MPR

Dr. Gordon Moore and Dr. Carver Mead

More Transistors Than Ants

MEAD: Along the way this has been a fantastic evolution process. You and I have talked down through the years about the process, and you've said some things I've quoted many, many times as Gordon Moore's other laws. Not everyone here knows Gordon Moore's other laws.

MOORE: Including me. *[Audience laughter]*

MEAD: I remember one time we were talking about how far you can see ahead. Do you remember what you said about that?

MOORE: Well, what I usually say is two or three generations is about as far as I've ever been able to see, and it keeps receding as you get closer.

MEAD: Right.

MOORE: There always seems to be a barrier out there about three generations away.

MEAD: But it always seems to be that far away.

MOORE: Yeah.

MEAD: Yeah, yeah. You also talked once about our ability to predict in terms of how long it's going to take to do things, and how far we're going to evolve over any given period of time. Do you remember that one?

MOORE: No, I don't.

MEAD: Well, I will quote you and then you can have rebuttal time.

MOORE: OK.

MEAD: This was a discussion with Gordon, and he said, "You know, engineers, we're always way too optimistic in the short run. But in the long run, it will evolve much further than we can see now." Do you remember saying that?

MOORE: OK, I've said something like that, yeah. I don't even think it's original.

MEAD: I've quoted that many times to many people. I thought that was one of the most wonderful—it's given all of us a belief in the future, in a country that hasn't always had such a deep belief in its future. I think, for me, one of the wonderful things about Moore's law is that it's a tangible thing about belief in the future. You've given all of us a great deal of belief in the future, and that's a tangible thing, not just a rosy glow over burning embers. So what do you think from here? You said something about how many transistors for every ant on the planet, Gordon?

MOORE: Well, I've used several of those analogies, but Neil Wilson, Harvard, I guess he's an entomologist, conservationist, ant expert, estimates there's between 10^{16} and 10^{17} ants on Earth, and we're beyond 10^{18} transistors per year produced now, so that's 100 transistors for every ant. *[Audience laughter]*

MEAD: That's some sort of calibration.

MOORE: The one I've used recently, I estimate there's something like 10^{18} printed characters a year. That's all the newspapers, books, magazines, and Xerox copies, and papers you throw away from your computer printer, and everything. So we actually make more transistors than printed characters, as near as I can estimate these days. We sell them for about the same price. In fact, transistors are cheaper than printed characters in the Sunday *New York Times*.

MEAD: So printing on silicon a functional transistor is cheaper than printing on paper a readable character.

MOORE: Yeah.

MEAD: So that's something about economics.

MOORE: I was trying to get the idea across in the 1965 paper [defining Moore's law] that this was going to be the cheap way to make electronics.

MEAD: And it certainly has come to pass.

MOORE: Yeah.

Audience Question-and-Answer Session

Q: I have two curve-fitting questions. Looking back at Moore's law in history, what has it really turned out to be in terms of doubling every so many months? And is there an equivalent [law] for the cost of the fab over time? Has that doubled every so many months? *[Audience laughter]*

MOORE: Well, on the first one, the first ten years where I predicted ten doublings, we only got nine, but it was pretty close. Beyond that, where I predicted every two years, it's been closer to every 21 months. It's been a little bit faster than every two years. And recently, if anything, it's accelerated. It's like the expansion of the universe. All of a sudden when you don't expect it, it's expanding faster instead of slower. Part of the reason is that the participants in the industry now recognize that unless they stay at least on that curve, they fall behind. The economics of the business are peculiar in that the next generation of technology always makes things cheaper than the current generation. So if you are a generation behind, you have a cost disadvantage and a performance disadvantage, so everybody really sees the need to stay right up here to the front if they are going to make devices that are really leading edge.

On the other thing, I used to plot the cost of fabs versus technology. Took me a long time before I could say "billion" when talking about fabs. Now it's $3 billion. But we flatten that out amazingly well. It's taken a long time to go from $2 billion fabs to $3 billion fabs. In the process, we've gotten the 12-inch wafers, so we get a lot more out of a $3 billion fab now than we used to get out of a $2 billion one. So on the cost per unit area of silicon, I think we've actually gone down.

Q: Fewer and fewer American students are getting into engineering. How do you think the U.S. will do in the scientific world in the future?

MOORE: I think that's a major problem. There are fewer of them, and they're a lot more expensive than engineers are in other parts of the world who are, if not as competent, almost as competent. Engineering increasingly is an international operation. I don't know quite where that settles down, but it's a problem I think, as a nation, we really have to face. Do you have a comment on that Carver?

DAVE HOUSE: Carver might have an idea on that. He's in the education business.

MEAD: Yes, I think it's a big problem. It's mixed up with our immigration policy. We used to get the brightest people from overseas coming here, getting advanced degrees in the U.S., and staying here. Now we force them to go back when they get their advanced degree. The smartest thing we could do as a country is to make sure that with every advanced degree, we have a green card issued. That would be extremely important. *[Audience applause]*

Q: What was the single most significant impetus that sparked your exodus from Shockley Labs and propelled you to found Fairchild? I don't know if it was personalities, or—

MOORE: Well, there were a variety of problems. I don't know if I could pick the one that was singly the most important. Things like deciding he was going to give the entire staff lie detector tests to find out what the source of some little problem was…I think if you talk to the eight people that left [Shockley Labs] to set up Fairchild, you would probably get eight different answers to that question.

Q: It's been suggested by some futurists that your observation, Moore's law, applies to a larger landscape, the green technologies: genetics, robotics, integrated circuits, and nanotechnology. What do you think about that?

MOORE: Well, I am amazed at in how many contexts Moore's law is used, and I am happy to take credit for all of them. *[Audience laughter]*

Q: But do you think it was a valid observation?

MOORE: I can't comment. I haven't made the plots myself.

Q: When the eight of you founded Fairchild, I believe you hired a CEO from another semiconductor company. What were some of the competitors out there, and how did they compare to Fairchild? And what was your first fab like?

MOORE: Well, Fairchild was still pretty early in the game. Texas Instruments was the biggest semiconductor company in those days. The others were Philco, Sylvania I think, Transitron was still around—

MEAD: RCA did some stuff.

MOORE: RCA some, yeah. But TI was the one we certainly focused on most. They were the ones that were furthest along in silicon. Our first fab was down on Charleston Road, where Fairchild set up initially. It's California's one-thousandth historical landmark, actually, the building. It was, I guess you'd call it, an R&D building. Two stories in front, one story in back, a total of 14,000 square feet, as I recall. We split it down the middle and decided one half would be R&D, and the other half production, if we ever got that far. I remember laying out the fusion area. We had six tables. Each table had two furnaces on them, so you had twelve furnaces. [Fairchild cofounder] Jean Hoerni looked and said if you staggered them, you would get an extra table in here, so we said OK, we'll use that for expansion. *[Audience laughter]* Our nitrogen supply had a manifold with two cylinders on each side. You could switch from one to the other. It was expandable by another couple of cylinders, I think, if you ever needed more. And that was the beginning of it, where we developed the first processes, made the first devices. When people started buying them, we found out we had to make a lot more of them.

MEAD: What were the wafer sizes then?

MOORE: The wafer sizes varied by where they were cut on the crystal, you know. *[Audience laughter]* The crystal went kind of like this *[makes a wavy motion with hands]*. One of my first contributions was to show if they got bigger than three-quarters of an inch, the yields went to zero, so they were typically one-half to three-quarters of an inch, and you'd get a few transistors on that if you were lucky. Things have changed a lot since then. *[Audience laughter]*

Q: Recently a number of factors combined to create a panic about power dissipation as a fundamental limit, and the heat generated—you know, 150W microprocessor chips. Intel has responded with a major effort, and they have solved the leakage problem, and they have created multicore processors. So now people don't appear to be so concerned about heat as a fundamental limit. My question is, how much time did we buy before it again becomes a fundamental limit?

MOORE: I may not be as close to that as I should be, you know, I'm a little technically obsolete these days. But power is something you can trade off with other things. Intel finally got around to focusing on it in the last couple of years. Now we make much lower power processors and still getting the performance advantages. The things they are doing will last for at least a few generations. I don't know when power will come back to bite us again. I think it's an engineering trade-off you have to make. One of our guys was plotting the power density coming out of these processors, and he went by the power density of a nuclear reactor, and then in a few years we got to the power density of the surface of the sun if we didn't take a different tack. So he succeeded in convincing the people internally that power was something they really had to focus on.

Q: I'm curious about the early days at Intel. How many times—or did you have any experiences at Intel—where you thought it was the end of the road, that the company wasn't going to make it, or finances were running out?

MOORE: That's a question you'd get a completely different answer to if you asked me or if you asked [former Intel chairman and CEO] Andy Grove. *[Audience laughter]* My view is that it was a very smooth startup. The goal we set for ourselves, we figured we had to get to $25 million in revenue in five years so the established companies couldn't get us out of business. We actually got to $63 million, so we blew by that. We raised money ahead of when we needed it, so we never really got things too tight. I thought it was smooth as could be. Andy says it was the toughest time of his life; he thought we were going to go under about every week. I think it's our different attitude.

HOUSE: Only the paranoid survive. *[Audience laughter]*

Q: I have a question for Dr. Mead. You've alluded to the almost religious reverence that we hold Moore's law in, and we also keep joking around about the fact that it is probably the most misquoted law in history. Why do we keep misquoting it?

MEAD: Well, I can't speak for people that quote things, because I have personally been misquoted greatly, and I know Gordon has also. It's what I said before about Moore's law. I think it's extremely important that there's a living example of people's belief in the future bringing it to pass. Because that's really down deep what Moore's law is about. If you don't have that belief, you won't put in those extra hours, and you won't figure out the problems of the process, and you won't get the yield better, and it won't happen. And semiconductors are certainly not the only arena where learning curves of that sort work. And so I think it's not amiss to take that belief system into other parts of our industry.

HOUSE: I would comment that probably at Intel we contributed to that, because we noted that the doubling of transistors every two years was a nice factor, but it wasn't very useful in selling microprocessors. But that doubling every two years allowed the performance of microprocessors to actually double about every 18 months, and that was the number that we would tend to push from a marketing standpoint, because it was applicable in selling microprocessors, particularly during the RISC and CISC wars. So we would promote the doubling in performance every 18 months, and then people would come back and say transistors would double every 18 months.

MOORE: That's right.

HOUSE: And mix the two.

MOORE: Dave House deserves credit for changing it to 18 months. *[Audience laughter]*

HOUSE: Well, I always quoted it as a corollary to Moore's law.

Q: In the past several decades, the field of semiconductors has been overwhelmed by electrical engineering and physicists. In the early days, you were training as a chemist. How did that enable you at Fairchild to better see some of the problems in surface physics, silicon semiconductors, that you were able to solve problems because of your chemistry background and chemistry training? We don't have any chemists in the field very much now.

MOORE: I was a physical chemist. In fact, my Ph.D. says chemistry and physics. You learn on the job. There wasn't an awful lot to learn about semiconductors at the time I started with Shockley. I learned enough as we went along. I was always more comfortable in the processing end. I never learned much about circuitry or things like that. But I think that any good technical background works pretty well in the technology industry. Noyce was a physicist, Andy Grove is a chemical engineer, Craig Barrett a material scientist—you know, all different but reasonably related fields.


MPR

Dr. Carver Mead

Q: When you wrote your original paper, how much of your thinking was that you were analyzing a trend versus throwing down the gauntlet for the industry to follow that trend? And now 40 years later, how much do you think the industry would have followed that trend on its own, versus your setting a goal which the industry marched toward?

MOORE: Well, like I said, I was really just trying to get the idea across that this was a way to make cheaper electronics. And I wasn't throwing down a gauntlet. I was trying to change the customer's mind. Up to that time, there were all sorts of arguments why this would never be cheap: you take an eight-transistor circuit, and we know you get 20% yields on transistors, so you would take 0.2^8, your yield would look terrible, things like that. Arguments that the reliability wasn't good because you couldn't measure the individual components. So I was just trying to shoot down some of that. A lot of these things would have happened anyhow. The natural rate at which technology changes—I'm sure if I hadn't plotted some of these curves, somebody else would have. I don't know if the industry would be very different now than it is.

HOUSE: So now we've learned that Gordon wrote Moore's law for marketing reasons. I guess he's the first marketing guy. *[Audience laughter]*

Q: I'm not going to ask you an engineering question, but I was wondering if at Fairchild and Intel, did you ever continue to kind of sneak out and blow stuff up?

MOORE: Chemistry sets don't have the good stuff any more. *[Audience laughter]*

Q: Looking to the future, how many years out do you see Moore's law continuing? What do you see as the next major challenge we have to overcome? What are our prospects for overcoming it?

MOORE: Well, this is a question I always answer by saying I can never see more than about three generations, and we can still see three or four generations of the technology. The generation now is a couple of years. It used to be three. So I can see six to eight years further into the future. Probably push beyond that, I don't know. The fact that the materials are made of atoms is a fundamental limitation, and it's not that far away. You can take an electron micrograph through some of these devices, and you can see the individual atoms in the layers. The gate insulator in the most advanced transistors is only about three molecular layers thick. We can change materials, it will fatten it up to ten again, I suppose. That cuts the leakage currents down. But we're pushing up against some fairly fundamental limits. So one of these days we're going to have to stop making things smaller.

Q: My question was originally about different substrates for integrated circuits, but since we are talking about atoms, what about different orderings for substrates, such as crystalline structures for printing circuits? Do you see a future in that?

MOORE: Different substrates than silicon? I'm not quite sure I understand your question, but silicon is just a marvelous material once we learn to work with it. I don't know anything else where you can get the huge single crystals and work with them as effectively as we do, so there's not much pressure to change substrates. There's not much motivation for it, really.

For More Information

To learn more about the Computer History Museum in Mountain View, California, visit *www.computerhistory.org*. A video recording of the Moore-Mead event is available on the museum's website at *www.computerhistory.org/events/index.php?id=1125619052*.

Q: I was thinking in terms of the heat problem, and maybe different substances have different properties.

HOUSE: Heat dissipation, and whether we would change substrates for that reason.

MOORE: Silicon is a good heat conductor. Diamond would be better, but technologically, that doesn't make sense yet. It's hard to find many things better than silicon that fit in with the rest of the technology.

Q: Seems like even a few years ago you could have eight young people go out, raise a few million dollars, and you'd have a semiconductor startup, and you could get pretty far with it. Right now it seems like, well, you need eight young people and you need a very rich uncle with $40 million–$50 million to get anywhere with that. Where do you see innovation going in the chip industry if everything is getting so expensive?

MOORE: Well, set up a software company. *[Audience laughter]* There's no question the silicon technology has gotten very expensive, very complex. It's hard to do on a small scale. Except for some of the specialized applications, it's not a good area for startups any more.

HOUSE: I guess Carver has done some work to make the silicon business a software business. With tools.

MEAD: It's one of those things that always happens when you have certain aspects of a business that are capital intensive—then they get shared. So that's what the silicon foundries have done for startup companies. There are quite reasonably funded fabless semiconductor companies which are still happening, and yeah, they are mostly software companies.

Q: You know back in the '70s and '80s when I worked at Intel, we thought that actually Moore's law was the iso defect curve [a defect-density curve that plots wafer yields normalized for die size; named after the contoured iso-lines that show elevations on a topographical map]. So I always wondered, how did you come up with the iso defect curve, and just some funny stories if you can remember.

HOUSE: That was the law wasn't it?

MOORE: Yeah, I have to do a little definition here. In order to characterize the process, we needed some measure that was independent of the size of the die. So I developed a curve where presumably you could normalize the die-size out. The same defect density was the iso defect idea. It was based on such hokey science that I never let anybody know the origin of it. It's been hidden in my idea. It amazes me that they still use it at Intel. They have much better models of yields now than we had then. ◇