

INTRO 1: It is now my great pleasure to introduce Dr. Gordon Moore, our first plenary speaker for the 50th anniversary of ISSCC. Dr. Moore graduated from the University of California at Berkeley with a Bachelor of Science degree in chemistry in 1950. He received his PhD in chemistry and physics from the California Institute of Technology in 1954. Dr. Moore joined the technical staff of the Applied Physics Laboratory at Johns Hopkins in 1953 where he did basic research in chemical physics. He joined Shockley Semiconductor Laboratory shortly after its founding in 1956 in Palo Alto, California, where he worked on semiconductor process technology with William Shockley. Dr. Moore co founded Fairchild Semiconductor Corporation in Mountain View, California in 1957 serving as Manager of the Engineering Department until 1959 when he become the Director of Research and Development. Fairchild produced the first commercial integrated circuit during this period.

In April 1965 while at Fairchild, Dr. Moore wrote his famous predictive paper on the benefits of integration. This prophecy later became known as Moore's law. In July 1968 Dr. Moore co-founded Intel Corporation to develop large-scale integrated products beginning with semiconductor memories. Intel soon produced a number of products based upon LSI technology including the world's first microprocessor. At Intel, Dr. Moore has served as Executive Vice President and Chief Executive Officer and the Chairman of the Board. In 1999 he became Chairman Emeritus. Dr. Moore is a director of the Giliad Scientists and a board member of the California Institute of Technology. He is also Chairman of the Executive Committee of Conservation International, a member of the National Academy of Engineering, and a Life Fellow of the IEEE. Dr. Moore has received numerous awards for his contributions and revolutionary vision in silicon semiconductor industry. In 1990, Dr. Moore was awarded the National Medal of Technology by President George Bush. In 2002, Dr. Moore was awarded the presidential Medal of Freedom, America's highest civilian honor, by president George W. Bush.

At the 1979 ISSCC Dr. Moore presented a plenary talk entitled: Are we really ready for VLSI square, in which he explored the implications of sub micron lithography and million device chips on integrated circuit architecture and design.

We are greatly honored to have Dr. Moore back to present the first plenary talks of the 50th anniversary of the conference. Please join me in welcoming Dr. Gordon Moore.

Slide 2:

We are greatly honored to have Dr. Moore back to present the first plenary talks of the 50th anniversary of the conference. Please join me in welcoming Dr. Gordon Moore.

Applause

Well, thank you, it's a pleasure to be here. This conference has grown a bit over the years. I didn't make the first few. In fact, I didn't get into the semiconductor industry as you just heard until 1956. And even then I didn't go to Solid-States Circuits Conferences. Chemists didn't know an awful lot about circuits. But when chemists starting making circuits in the early 60s, I did attend several of the conferences. At the time they were all in Philadelphia. In Philadelphia in February it was all that easy to convince your spouse that is was really hard work, that you weren't going on a boon doggle. I'm not quite so sure that San Francisco is that easy to sell these days? But many of the parameters related to the semiconductor and solid-state circuits industry have shown exponential dependences over the years. But no physical quantity can continue to change exponentially forever. There's always some kind of a catastrophe if you project it far enough into the future. So what I want to do today is look at some of these exponentials and maybe give some idea of where they might go and talk a little bit about how we're going to deal with the looming catastrophe that people seem to be projecting, as they look further down the road.

SLIDE 3:

The first one I want to look at is the growth of revenues in the industry. This is a phenomenal growth industry. It has grown 80 fold over this 35-year period I've depicted here. A compound annual growth of some 14% even with the flattening of the last few years. The dips and bumps on this exponential curve may not look so severe, this one perhaps is a little more severe than the rest, but I look back here and realize that in 1974 and again in 1984-85 Intel had to get rid of a third of the work force. Exponentials tend to kind of distort things that are more nearly experienced linearly when you're actually there. But while this is phenomenal growth, if you want to see the real underlying growth of the industry, look at the output.

SLIDE 4:

This is the number of transistor ship per year as near as I can estimate over the same period. And this has grown $8\frac{1}{2}$ orders of magnitude over that period. 300 million fold, now that's a growth industry. Maintaining an average growth of about 80 % per year over this whole same time period and having a significant period in here where it actually doubled every year. That is we had more electronics built in the year than existed in the beginning of that year when we started out. Truly phenomenal growth. I tried to illustrate this number you're getting up near 10 to the 18th. I've used raindrops falling on California. Neal Wilson, the Harvard biology expert on ants, estimated that the number of ants in the world is 10 to the 16th to 10 to the 17th. So for years I used that. Now each ant has to carry 10 to a hundred transistors if it's going to take care of its load. (Audience laughter).

Perhaps another way of looking at the size of this number, I estimate that the number of printed characters produced every year is between 10 to the 17th and 10 to the 18th. That's all the newspapers, books, magazines, Xerox copies, computer print outs that you throw away, everything. That's the same order of magnitude as the number of transistors that the industry sells. (Of course we print a lot of them we don't sell). Those are where those little red dots are when we get through sorting them. This is phenomenal and we sell them for about the same price as a printed character in the Sunday New York Times.

SLIDE 5:

But, the power of the industry is when you divide one of these plots by the other and get the average cost per transistor over this period of time dropping from about a dollar to 2/10 of a micro buck. And that's the average transistor cost during this time period. If you look at D-RAMs, it's a order of magnitude below this. You get 50 million transistors for a dollar these days. And not only the transistors you get all that circuit design and inter-connections thrown in free. It really is a spectacular industry completely unparallel by anything else I can see.

It's hard to realize while looking at that, in about 1960, I don't know the exact date, the industrial engineers at Western Electric, looking at the transistor estimated that eventually it might cost about 60 cents to build one, looking at time and material studies and the way things were happening at that time. This multiple million-fold reduction in cost requires some really special things in the industry. First of all, it requires a technology that has phenomenal capability. I think there's a unique technology underlying the

industry. And a fantastically elastic market. It has to be able to consume ten to the 18 transistors a year and growing. And also in order to make it happen it has to have the contribution of a lot of people doing circuit design, clever extensions, developing the capability to continue it and technologists that keep the technology moving rapidly.

SLIDE 6: Now go back a bit in time and this is what a transistor looked like about 1959. I didn't have one from 58, which was the first year that transistors were manufactured on a whole wafer essentially, rather than one by one. But this was a one-inch wafer back in 1959. A real breakthrough at the time because one of my first contributions of the industry was proving when it went above $\frac{3}{4}$ an inch in diameter, the yields went to zero, the material was so bad. This shows one of my technical contributions, that flat part of the bottom of the wafer that lets you align it in 2 dimensions. (Audience laughter). And there are about 200 possible transistors on a wafer like this and if we were lucky something like 10% of them were good. (Audience laughter). Now of course we can get a lot more microprocessor chips on a wafer, each with millions of transistors. Very, very much better yield.

SLIDE 7:

If we take the next step from these early planar transistors to the early planar integrated circuits this was one of the very first micro logic circuits made. The first integrated circuit was kludged together by Jack Kilby at T.I. Fairchild had the planar technology in place, so we had the technology to make the things practical and Bob Noyce knew how to extend that technology to get to something that would be worthwhile. This is one of the first ones we built. This was a flip-flop consisting of 4 transistors and 6 resistors. It was one of the micro-logic family. The die is round so we could bind it right to the header with little dobs conductive epoxy since the yield on making 6 bonds to a chip was going to be so low, we thought we needed something that didn't detract from the yield from the assembly point of view. It's a terrible picture. I'm sorry that this is what we have left to show you of those early days. Maybe it shows a bit about the problems we had.

Now, you might think integrated circuits were obvious solutions or immediately accepted. This was not the case at all. This was a tough sell. Our customer in those days, the technical interface with our customer, was typically the circuit designer. We were making transistors and things like that. And to go to a circuit designer and tell him that "Hey, we're gonna do

your circuit design for you,” wasn’t something that was sold very easily. The reliability people looked at these and they were used to taking transistors and measuring their parameters over a period of time, looking for drift. They said, “Gee, we can’t measure the transistors here. You can’t get a hold of them. How do we know if it’s reliable or not?” I remember going to one aero-space company said that we use 16 different flip-flops in the systems they built. They could never use a standard flip-flop; they had an expert for each of those. The really had to be specially designed.

Then Bob Noyce made another one of his major contributions to the industry. He said okay, “We’ll sell you the circuit for less than you can buy the transistors and resistors to build it yourself.” And that was a major break through. The fact that it was also considerably less than it cost us to build the integrated circuit at the time, was of little consequence. We had to develop a market for them and of course the solution that the semiconductor industry developed was that whenever you have a problem, you lower the price. That’s the way you solve all of these things. Let the elasticity of the market bail you out. Well, during that time, in the early circuit days, I was the director of R&D at Fairchild and had a little bit more visibility than most people did in where integrated circuits were taking us. And I was asked to write a paper for the 35th Anniversary issue of Electronics magazine where I predicted what was going to happen over the next ten years in the component market.

SLIDE 8:

And that’s when I plotted this curve that eventually became known as Moore’s law. Extrapolating from about 60 components to 60 thousand over ten years was pretty aggressive. I never expected it to be precise. I was trying to get the message across that this was going to be the cheap way to make electronics putting a lot of it on a chip rather than building it up from individual components soldered together. But, in fact, it turned out to be more precise that I ever could have imagined if I put the data on it here.

SLIDE 9:

This is what happened to a few of the points over that time period with the most complex circuits available. They fit amazingly well on it and my friend Carver Mead, a professor at Cal-Tech, dubbed this Moore’s Law.

SLIDE 10:

And what happened since then, well in 1975 I updated it – added the purple line. And I argued why the slope was going to change from doubling every year to doubling every 2 years. My data said it ought to change right then but the stuff I saw in the laboratory, which was generally CCD memories, suggested that there were a few generations that were going to continue to double. So I said double again for another five years and then go to the slope of every 2 years. And CCD memory has never quite worked out that way. It was a problem called soft errors that came around. CCD's are much better imaging devices than they are memories.

SLIDE 11:

So instead of waiting 5 years to break the actual data broke at about that time. I got the slope pretty close, but that 5-year hiatus really made it crucial. It didn't fit as well as it might have.

SLIDE 12:

So, this shows what a wafer in early production days, at three-quarters of an inch, the size of a nickel, looks like compared to a modern wafer. One of the other things that I projected in 1975 a bit tongue in cheek, was what was going to happen to wafer size again using my semi log paper to extrapolate.

SLIDE 13:

One of my colleagues showed what actually happened. I had remembered that I had only predicted 56 inches but he went back and got the original data. And not only have things like this occurred but the structure has gotten complex.

SLIDE 14:

If we take a cross section through a modern process, a modern seven layer metal process with things like tungsten plugs down her at the bottom always the interface is in one thing or the other – This is amazing to me. I remember when the principle argument to go from bipolar circuits to MOS circuits was process simplicity, you only needed 5 masks. We succeeded in driving that up to 25 and I don't know when it's going to stop. This actually surprises me. One technology that made this possible was the idea chemical mechanically polishing after each layer so you maintained a flat surface if you build up insulators and metals and developed a structure. Without that the topography became so complex that you couldn't get more than 2 or 3 layers of the metal before it was unmanageable. The IBM invention of the chem-mechanical processing really allowed this to continue. And I suppose

if I plotted the number of metal layers versus time, it would also be close to an exponential, although I haven't plotted it in that direction.

SLIDE 15:

Of course the big driver for the improvements over a period of time has been the ability to make smaller and smaller features. The original planar transistor didn't push that too far, but starting with the early integrated circuits, we've been on a very constant change for many years of cutting the dimensions in half about every 6 years. That's 2 steps generally at 0.7 in a step. So you double the density every 3 years. Every 6 years I went down a factor of 4. And of course I would expect that to start rolling off.

Actually the opposite has happened. When the last few generations are projected going forward, the time period between generations is a step down to 2 years rather than the 3 we've done historically. This is amazing to see a curve like this accelerate rather than start to round out. And you see some of the other comparisons you might hear. People always start with the human hair. A human hair is a few orders of magnitude away from where we are now. We're working at an eighth virus sizes and working down towards single molecules. I think that's a pretty big molecule here. But it's still quite a way to go before we get to that point.

Of course, I guess we're well into the realm now, that's called nanotechnology. We're doing nanotechnology from the top down rather than the bottom up. Doing it from the top down at least for the electronics, we can continue to connect everything. The challenge of the people building single electron transistors from the bottom up is how they put a billion of them together and make it into an interesting, useful function. It'll take a while to see how that's going to happen.

SLIDE 16:

From a linear view -- the exponentials tend to distort how we look at things from a human point of view. If you just look at the qualitative progress over a period of time taking in INTEL 1978 technology and modern technology, this is a 6 transistor static RAM cell compared with one contact opening in the 1978 technology. The continued evolution of these things over period of time really has a qualitative effect not just a quantitative one.

SLIDE 17:

I actually think we're breaking the laws of Physics in some of this. We're printing 50 nm lines with 193 nm light. We're printing lines at a quarter of a wavelength of the light. This is something that I would have thought was impossible. If I go back to the other curve, I remember thinking that 1 micron was probably as far as we were ever going to go, you know, a couple of wavelengths of visible light. We wouldn't have a way around that problem because imaging really requires... well, its hard to make images smaller than the wavelength of light. Then we moved to ultra violet light and I thought maybe a quarter of a micron, maybe we'll get there and that'll be as far as we're ever going to go. We blew through that. Now we're sub-tenth micron, 90 nanometer technology in production quantities and we're looking at printing lines a quarter of the wavelength of light. It really is amazing that we've been able to do that. Lasers are great for improving some of these capabilities. This requires that you can use a low contrast optical image, contrast drops as you do this, but you use a very high contrast photo resist with adequate controls and you can actually print these fine things.

SLIDE 18:

For those of you who haven't been to a lab recently this is a sort of cutaway picture of what a modern step-n-scan production tool looks like. It's a several million-dollar tool using an eximer laser and a variety of things. It's a far cry from the original systems as Bob Noyce and Jay Last put together at Fairchild to do our initial lithography. Well we've been working hard on these x and y dimensions; we haven't neglected the 3rd dimension, the vertical one either.

SLIDE 19:

The minimum insulator thickness has stayed off the line that goes on exponentially. This one probably surprises me more than the line width it. I remember doing a back of the envelope calculation about the time Intel was formed back here and convincing myself that statistically if you went to layers that were less than about a thousand angstroms, 100 nanometers thick, you'd probably get enough fluctuations that things wouldn't be very good. But I didn't realize, I should have I guess, the force was with us. Chemical forces really help. And you don't get a statistical layer of atoms coming down when you oxidize the silicon, you get nice chemical reactions that really do maintain the integrity of the layer down to the very, very narrow layers. Here we are a couple of nanometers thick physically, electrically they look a few nanometers thicker than that because I guess the electrons can't

get all the way to where we see the edges. And if you look with a transmission electron microscope that silicon substrate, the insulator layer and poly crystal and the silicon on top, you see these structures really are getting down to a few molecular layers thick. Well, you can't go much further than that, but you don't have to in this case.

SLIDE 20:

If we go to a material with a higher dielectric constant we can actually get higher fields in the silicon with a thicker dielectric. This is an experimental structure and the capacitance goes up something like 60%. But the great deal is the leakage current. In here you get down where you get a lot of tunneling. Leakage current decreases a hundred fold. So there are things that can be done as we approach some of these limits that preserves our rate of evolution of the technology without getting off the trends we've been on historically.

SLIDE 21:

These have all led to dramatic increases in the performance of the electronics that we build. Here I've plotted processor performance. I'll excuse myself for using only Intel data, but that's the easiest for me to get my hands on. And you can see that over this time period from the first microprocessor in 1971 to modern microprocessors we've had about 5 or 6 orders of magnitude improvement in processing on an IC. That's a compound growth rate of about 50% or a doubling in performance every 20 months or so. Those of you have heard Moore's law quoted as doubling every 18 months, notice I never said 18 months, I said one year and two years. One of my colleagues at Intel, Dave House translated that into processor performance and decided that went a little faster than a number of components. So he was the guy who said 18 months, not me. (audience laughter) But, it's pretty close to what's happened on the processor performance over the period. And it shows no sign of decelerating. In fact, if anything, over this part of the curve, it's accelerated.

SLIDE 22:

Now, we get some problems coming along here. One of them that I think is an important feature of this conference, and that is what is happening to the power? And here I'm looking at 2 contributors to the power. The active power of the processors which are getting up to power dissipations of a fairly bright light bulb and power densities the we use to strive to get in power transistors. This is getting to be a problem. I don't want a kilowatt in

my laptop. It would be very uncomfortable. So from practical reasons the power is going to have to roll off there.

The thing that is probably more disconcerting here is what is happening to the power contribution from leakage, not the active power of the device. That is a steep exponential. We've been fighting the power for quite a while and our best tool for fighting power has been the power supply voltage and you get the square dependence on it.

SLIDE 23:

I suppose we can call this an exponential also, its kind of a step wise exponential. We use to live with 12 volts and we went to 5 volts. We were going to stay at 3 for awhile. But we discovered that cutting the voltage was so nice for power, that we just kept right on cutting the voltage. Now every processor seems to pick an optimum voltage and we continue to lower this principally to get the lower power, but of course, if we have much thinner dielectrics, they like lower voltage also. Now, again, this can't go on forever. You need at least a few hundred mili-volts, I don't know exactly what, just to overcome some of the noise problems that will exist in digital devices. I suspect something around 1 volt is going to be a limit, but I sure have been wrong on a lot of these other things that I suspected were going to have limits, and you folks probably know that a lot better than I do. We chemists don't understand this kind of stuff so well. Especially that we don't look at it closely anymore.

SLIDE 24:

But there are new things we can add that help with a lot of these problems. For example, moving from silicon to silicide gates helps considerably. We can continue to improve the performance of the transistor by purposely straining the silicon in the region. You strain it one way for the n channel device and the other way for the p channel device and improve the performance of both transistors. As we look a little further down the road we see things like a high dielectric constant, gate dielectric, which gives us the higher electric fields while still keeping the leakage current down. And new transistor structures.

SLIDE 25:

One of which is the tri gate structure I show in here. This is an interesting device because it kind of changes at least the way I've always thought that transistors around completely. And we got to make very, very narrow line

webs down here in the 100 nanometer or less, range. All of a sudden it doesn't make sense to deal with thin cells in the other direction anymore. We just make a relatively thick film and make the transistor move in the other direction. Here the gate wraps completely around the silicon, which is sitting on a silicon insulator in this case. So you get a source, and a drain, and a gate that depletes the silicon from 3 sides. You can make a fully depleted transistor in this manner and cut the leakage current dramatically and get very high performance. It's the kind of thing that you don't think about till you get down to the dimensions where it suddenly becomes possible to build this kind of structure. This kind of a transistor might carry us significantly further than we can go with only the techniques we thought available in the past.

SLIDE 26:

Of course we have to be able to print finer and finer lines. To do that we need shorter wave lights, we can't continue to print at a smaller and smaller fraction of the light we're using. But the technology generation that we're seeing now, even by the regular scaling we've been doing where we've reduced the line width by a factor of about .7 with every technology node we go through, from today's volume production we can make conventional transistors down to the 30 nanometer range. It takes 2 or 3 years between generations. 3 years is the typical roadmap numbers the SIA puts out, the Semiconductor Industry Association. Two years seems to be what we've been doing over that last few years and expect to be doing for the next few generations. So, we're taking 10 plus or minus 2 years of conventional scaling. Below 20 nanometers it's not clear with the conventional devices will work, but something like that thin transistor I showed on the last slide looks like a very realistic possibility. To make these very narrow lines, we need a step in lithography.

SLIDE 27:

And this is the tough transition that we're going to have to go through, but we've been working on a technology using what is called an extreme ultra violet (EUV). This is a wavelength range that used to be called soft X-rays but X-rays are kind of a bad name in the industry so the name was changed. But with the experimental system that currently exists, we can print 50 nanometer lines and spaces, small dots and continue to be improved. It gets to be much more complicated optically. There are no materials that are transparent really in these ranges so you can't use lenses. You can't use masks to put the radiation through; everything has to become reflective

optics. Even air is an important absorber, so you have to work in a vacuum. Simple mirrors don't work and so we go with metallic mirrors. The mirrors we used here were something like 88 layers with alternative new materials that give you the high reflectivity activity that we need. So working with about 13-nanometer light that's something less than a tenth of the 193 nanometers that's used in production now. And the industry has been investing in this kind of a system for quite a while.

SLIDE 28:

This is the proto type that exists in Sandia in Livermore. An industry consortium has been working with Sandia and with the Livermore National Laboratory to develop this useful technology. We're working with them because they use the technology that generally came out of the old Star Wars program. And it seemed most productive to continue with the group that had been working with the technology rather than starting all over someplace else. But as you see, this is a pretty complicated looking machine and its every bit as complicated as it looks. The mirrors have to be figured significantly better than the Hubble telescope and then we have to be able to make them on a production basis, so it's a real challenge.

SLIDE 29:

But one thing it will do, it will keep us on the exponential of the cost of the lithography tool. (Audience laughter). Well as you see a lot of things are changing exponentially. There's reason to believe that each of these has a problem. But I remember thinking a million dollar machine was going to be prohibitively expensive for the industry. Though we've blown through 10 million dollars for a lithography machine and looking at things that are significantly beyond that. The feature is the productivity of the tools has increased at the same time that the cost has again allowing us to decrease the cost of the transistor to make cheaper and cheaper electronics.

SLIDE 30: But we ought to remember that no exponential is forever. Your job is delaying forever. In the 40 years since the first commercial integrated circuits and the 50 years since the first commercial transistor and the first ISSCC, I think we built a fantastic industry. It is the most complex processing industry that I can identify by a significant margin. We manufacture the greatest number of items when you count that item as a transistor rather than the product they we typically sell. I don't think you can count bits on disks because they're not individually manufactured. They're just an area that happens to get polarized in a particular way. We've

had million fold, actually 10 million fold cost reductions and we've passed it on to the consumer. Again something that no other industry I can identify has done. I discovered recently or was told by an economist that the semiconductor industry had become the largest manufacturing industry in the U.S. as measured by value added because we start with sand, and a lot of value is added in there. (audience laughter) And of course, the electronics industry taken worldwide is the largest manufacturing industry that's here.

But there's still a lot to do. And I think there's a lot of life left in the technology that we've been developing and a lot of clever ideas of how to extend it in some non-conventional ways. And its ability to infiltrate sort of everything society does is tremendous. It is really an ubiquitous technology and something that I think will continue to have a very important role for the foreseeable future and beyond what I can foresee certainly. I am certainly honored to have be part of it. Thank-you. (applause)