Silicon Valley, What Next?

Craig R. Barrett

The following text is based on the plenary address given at the 1993 MRS Spring Meeting in San Francisco, April 12.

Let me begin by putting tonight's topic—"Silicon Valley, What Next?"—in the proper context. We can talk about Silicon Valley, Silicon Gulch, Silicon Desert, Silicon Glen, or Silicon Island: the topic is all the same. What I really want to talk about is the future of the semiconductor industry, an industry which only a few years ago was declared almost legally dead in the United States, with no future. Today, that industry is alive and well and I hope to demonstrate that it serves as the foundation for the most vibrant manufacturing industry in the world today—namely, the electronics industry.

To accomplish this, I will draw upon a few of the materials principles I learned some thirty years ago. I'll talk about some of the major industry trends, some of the technology challenges, and my projections for what will happen over the next half decade or so.

In the worldwide market place, the electronics industry is the largest manufacturing industry in the world, and by far the fastest growing. It is currently estimated to be in the range of a $900 billion industry, much larger than automotive or steel. The semiconductor portion of the electronics industry is in the $60–$70 billion range, but it really forms the basis of electronics. The argument these days is that you can't have a viable electronics industry—the computer industry, home electronics, or any sort of electronics industry—without owning the base semiconductor industry that feeds it. So, most governments subscribe to the philosophy that you need to have a vibrant semiconductor industry to succeed, especially if you are interested in export marketplaces. Electronics is the biggest export market in the world.

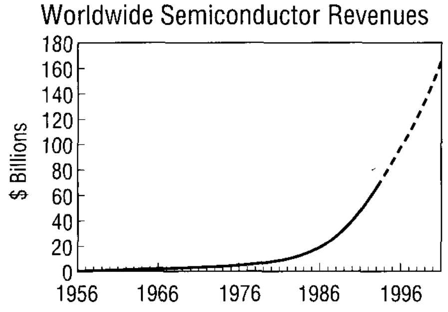

Figure 1. Worldwide semiconductor revenue growth. Source: WSTS, Dataquest.

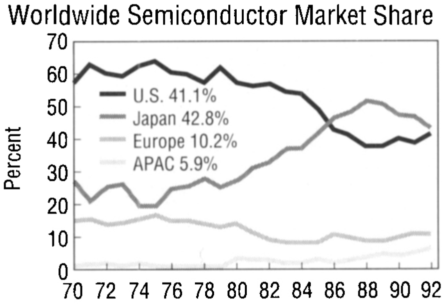

Figure 2. Worldwide semiconductor market share for Japan, the United States, Europe, and APAC. Source: Dataquest.

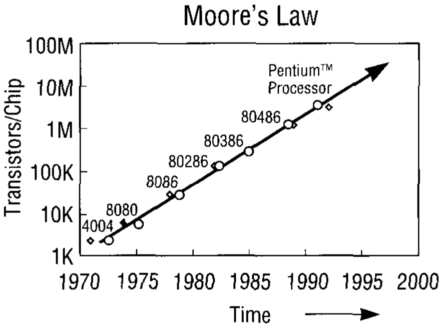

Figure 3. The "integration trend" given by Moore's law, which shows doubling of the number of transistors in an integrated circuit every 18 months for the past 30 years. Source: Intel Corp.

Semiconductor Industry

If you look at the industry in terms of its revenue growth over the years (Figure 1)—keeping in mind that exponential growth hardly continues forever—you will see that this industry has an enviable growth rate. Currently in the $60 to $70 billion range, it is projected by the year 2000 to be in the range of $150 billion a year. The automotive industry—or any industry—would love to experience growth like this.

With this kind of curve, if you just hold your market share, you have wonderful growth opportunities in front of you. Unfortunately, the semiconductor industry has been prone to price wars, something akin to what is going on in the airline industry in the United States today. In periods of excess capacity it has been a profitless industry for many of the participants. But it is, on average, growing rapidly. The United States was a world leader and 20 years ago held some 65 percent or so of the worldwide semiconductor market share (see Figure 2). That market lead steadily eroded, with the Japanese market share steadily increasing during the '80s. Western European market share went down slightly.

Over the last three or four years, the U.S. market share has rebounded. As I mentioned before, around 1988 many economic analysts and academic theorists projected that the U.S. semiconductor industry would die a natural death. There was no way it could compete in the world marketplace; the only hope for the U.S. electronics industry was in design alone, not in manufacturing. Fortunately, that scenario proved to be false. We began to regain market share in the late 1980s, and the United States and Japan are now locked in a tight struggle to determine who will be number one in the marketplace in 1993. I think both countries have viable competitors, and competition is going to be strong as we move forward.

The major growth curve in Figure 2 is really from the Asia-Pacific (APAC) countries. In countries such as Korea or Taiwan, it is a priority for them to get into the semiconductor business to fuel their electronic businesses. APAC has gone from virtually no market share to a 6 or 7 percent market share. They are doing to the Japanese semiconductor companies precisely what these manufacturers did to the U.S. semiconductor manufacturers a decade ago. That is, they are coming into the low end of the business, the dynamic random access memory (DRAM) area, buying market share, and making it uncomfortable for their competitors. Fortunately, for companies like Intel, we left the DRAM marketplace many years ago.

Many countries see the semiconductor industry as the lifeblood of the electronics industry, as represented by any one of their respective critical technology lists. Key technologies in each major industrial block are virutally the same and they all include electronics, photoelectronics, or electronic materials. The fact that these technologies are virtually the same has led me to an amusing conclusion about industrial policy and picking winners and losers. Regardless of whose list of winners you pick, you are essentially picking the same 20 or so industries or technologies, so it doesn't make any difference which list you pick. Just pick a list, and if you back those technologies, you probably will be backing the correct technologies for the next decade. Regardless of the source, the lists all cite microelectronics and electronic materials as key technologies for the next decade, so we can take some comfort in that.

Industry Trends

Let's look at some of the trends that drive the industry. They are the integration trend, the anti-inflationary trend, and the capital-intensive trend. Figure 3 shows the integration trend. Moore's Law was first formulated by Gordon Moore, founder of both Fairchild Semiconductor and then Intel. Moore, in the late 1960s and early 1970s, projected that every 18 months or so the number of transistors in an integrated circuit would double. He kept plotting this trend on his semi-logarithmic plot for years and he kept expecting the curve to bend over, but it hasn't bent over for the past 30 years. It still continues on its trend. If you carry Moore's projection out to the year 2000, it predicts a gigabit RAM—a microprocessor with about 100 million transistors, and a computational speed giving about two billion instructions per second. That will be your standard desktop computer.

Figure 3 shows how Intel microprocessors follow Moore's law, but the same curve applies to static RAMs, DRAMs, Motorola microprocessors, etc. They are all roughly parallel. There are no physical limitations that would prevent Moore's law from continuing. As I will point out later, to maintain Moore's law, what is needed is a few billion dollars of research and development, and much of it is materials oriented.

Figure 4 shows the anti-inflationary trend. From the production of the first commercially available DRAM (the 1103) in 1971 or 1972, to today's four-megabit or 16-megabit DRAMs, the cost per bit has gone down by over a factor of 1,000 (Figure 4a). I consider that to be relatively anti-inflationary compared to most things today, including my tax rates.

Figure 4b shows the number of DRAM bits that have been produced. This year there will be something like 1,016 memory bits produced worldwide. If you divide that by the number of people—every man, woman, and child on the face of the earth—you'll come up with a startling statistic: This year, the "average" person will use about 10 million bits of dynamic memory.

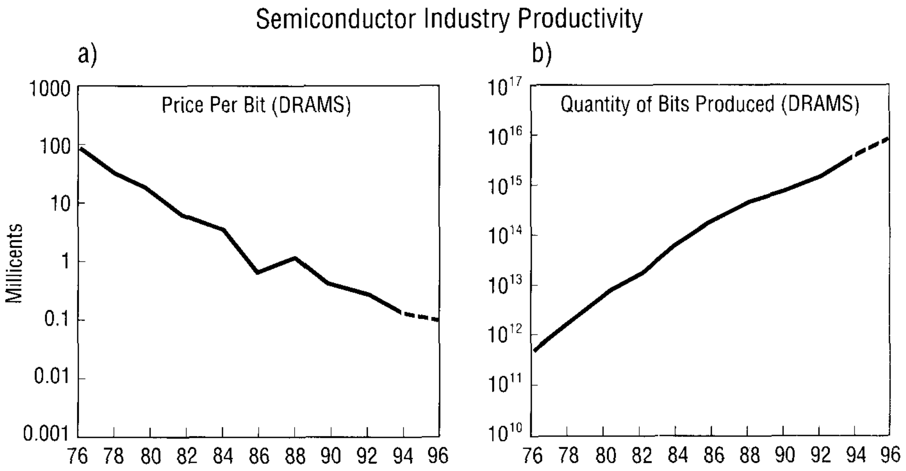

Figure 4. Anti-inflationary trend showing semiconductor industry productivity relative to (a) price per DRAM bit and (b) quantity of DRAM bits produced. Source: Dataquest.

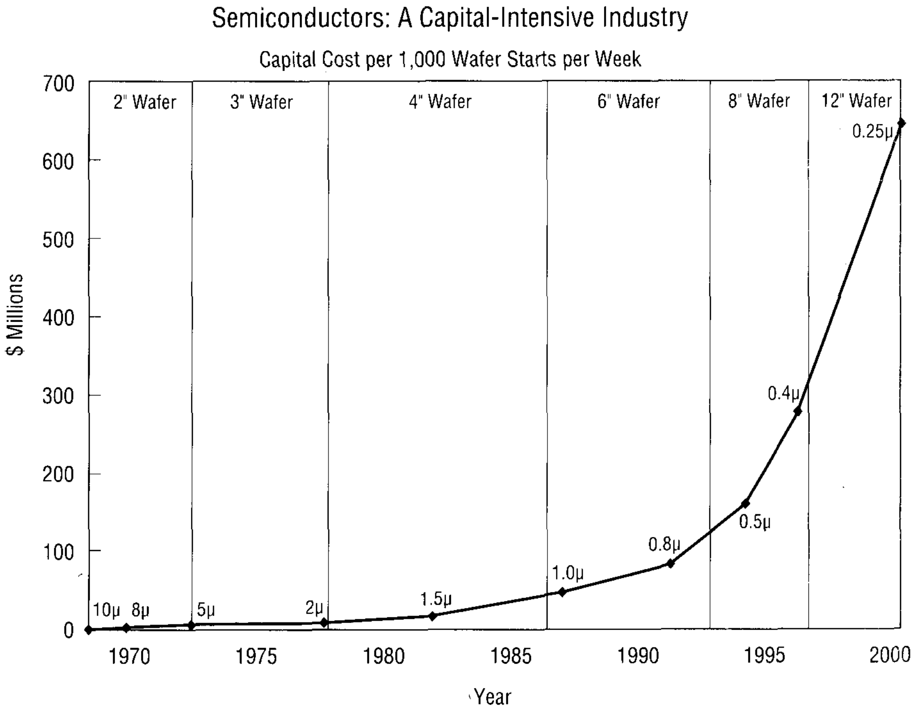

Figure 5. Growth of capital cost per 1,000 wafer starts per week in the semiconductor industry. Source: Intel Corp.

The capital-intensive cost factor of our industry is shown in Figure 5. We judge capital intensity by the cost of building a manufacturing plant. Typically, we build manufacturing plants that produce about 5,000 wafers a week. Today, it is 5,000 eight-inch silicon wafers a week. By the end of this decade, it will be 5,000 12-inch-diameter silicon wafers a week. Today, one of those manufacturing plants costs roughly $1 billion. Recently, we announced the building of such a plant in Albuquerque. A 5,000-wafer-a-week plant, building chips on eight-inch wafers with a 0.5- or 0.4-micron line-width dimension, has a capital expense of $1 billion. I like to compare that to the first wafer fab facility that Intel built in California. For instance, I look at the line item which shows me how much the cafeteria costs; today cafeterias associated with these manufacturing plants cost us $3–4 million. Then I go back to the 1970s when our first wafer fab facility cost less than $1 million, fully equipped.

The interesting part of the capital intensive trend is that our industry has been driven on the anti-inflationary curve. This, in turn, has been driven on more functionality per dollar, and functionality has been driven on increasing yields and increasing productivity. Capital intensity is increasing at such a rate that capital costs, or the depreciation costs of our wafers, is projected to be well over 50 percent of the wafer costs as we move into the 1990s. That may put a crimp on some of the anti-inflationary nature of our business. It also provides some insight into why the location of these manufacturing facilities is not very dependent on the labor rate these days; the labor content is less than 5 percent of the total wafer cost. So the wage rate of the local workforce makes little difference. What really makes a difference is the tax rate on the facility, the depreciation rates, and so on.

A Look at Chips

Let me give a couple of examples of typical integrated circuits. Figure 6 shows the Pentium™ microprocessor. This is a follow-on to the Intel486™ CPU. It is a good example of a state-of-the-art microprocessor. It has over three million transistors in it. It is about 700 mils on a side unmagnified. It runs at about 100 million instructions per second. It is built up of about 16 or 18 mask layers. Each mask layer is somewhat akin to the complexity of a street map of metropolitan Los Angeles. And all you need to do is superimpose those 16 or 18 layers on top of each other with enough registration that they all line up and the part works. The three-dimensionality of our integrated circuits is an increasingly important factor in their complexity, particularly in their materials complexity. In a three- or four-layer metal interconnect, the aluminum silicon metallization is connected by tungsten plugs between the layers. The last two layers are relatively flat because they are mechanically ground flat. For those of us who joined this industry 20 years ago, and have since then been trying to keep dirt out of our manufacturing facilities, the concept of taking these expensive wafers and polishing the active surface flat with an abrasive compound didn't make any sense when it was first suggested. But it has been engineered into a highly manufacturable process.

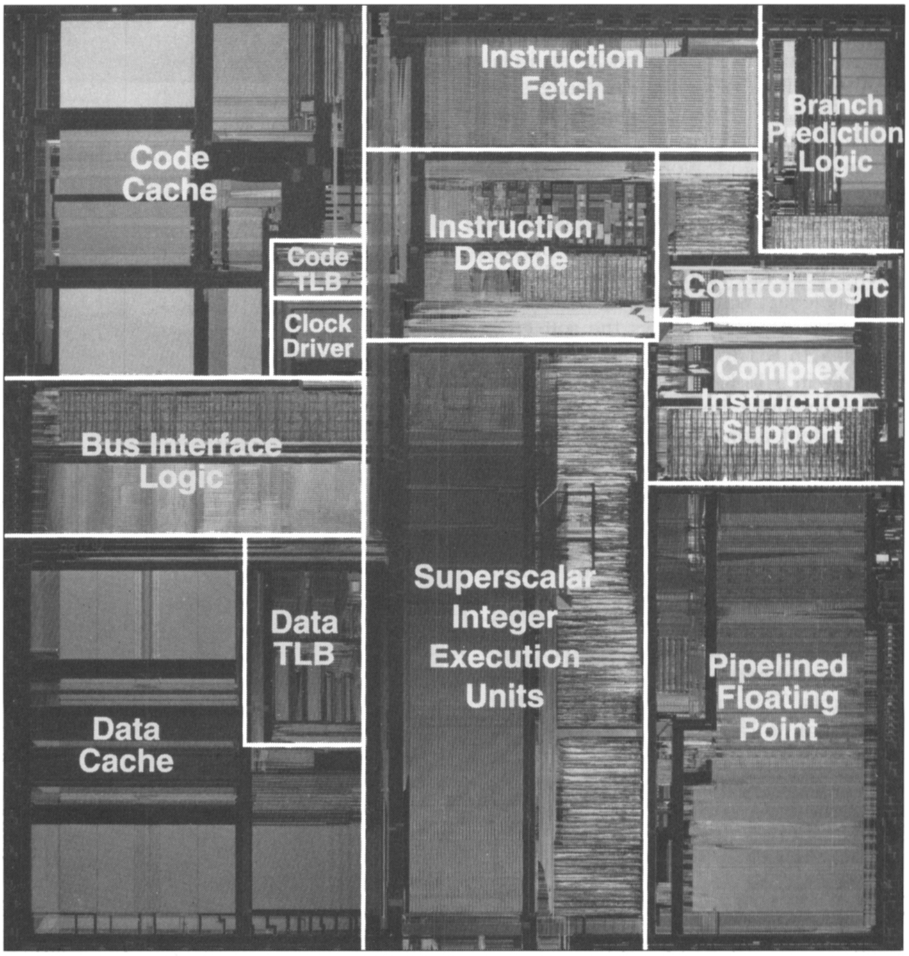

Figure 6. The Pentium™ microprocessor. Source: Intel Corp.

On the other side of the semiconductor business, the memory side, there is the eight-megabit flash chip. In the future, this technology, which is a nonvolatile read-write memory, will be used increasingly to replace DRAMs. The floating gate in a memory cell of a flash memory chip sits on top of a little bit of gate oxide. The gate oxide is roughly 100 angstroms plus or minus one angstrom thick. I have often asked our technologists how you control something that is thinner than a molecule of the material that you are trying to grow. I have never received a reasonable answer to that question, but they do tell me that it is 100 plus or minus one.

The interesting technological challenge from a materials standpoint concerns the memory cell, which is programmed by injecting some hot electrons through the gate oxide onto the floating gate. The floating gate is a little bit of polysilicon with other materials, surrounded by an insulator. You put about 50,000 electrons under that floating gate, then you assume that it will hold that charge for the next ten years. The programming margin involved allows the loss of a maximum of two or three electrons per day from that floating gate. That may not be difficult for one gate, but there are eight million of these floating gates in this circuit, and the next generation will have 16 million, and the generation after that, 64 million; and as the number of electrons you program onto each gate goes down, you soon need an escape rate of less than one electron a day as the basic requirement for long-term reliability. This presents some interesting materials challenges.

What Lies Ahead

Let me propose to you what the industry holds for the future—what the industry's characteristics will be by the year 2000. Again, I don't think that there are any technological challenges in the way, only some hard work in terms of the technology creation and the manufacturing facilities. You just need a few billion dollars. We will have 100 million to one billion transistors per chip, depending on whether it is a microprocessor or memory chip; a DRAM takes one transistor per memory cell, so that is a gigabyte memory. We will be talking about 50-angstrom oxides, and storage oxides with a billion storage gates per chip. Minimum dimensions will be in the range of 0.15 to 0.2 microns. Chips will be roughly an inch, or two centimeters, on a side. Chips will run at 500 megahertz. Five to six layers of metal interconnect will be stacked on top of the active silicon. And the facility to build these will cost probably in the range of $3 to $4 billion, if you want to build them in commercial quantities.

There are a couple of materials problems associated with the chips of the year 2000. To put down a 50-angstrom oxide, you would like the silicon beneath it to be relatively flat and smooth. Figure 7 shows an atomic force microscope (AFM) scan of a silicon surface that was put in deionized water for a period of time. The big bumps on the silicon get replicated in the oxide grown on top of it. This may be a relatively simple problem to solve, but when you get down to 50-angstrom resolution, it can be a difficult problem to overcome.

Another problem concerns putting isolation trenches into these circuits to electrically isolate adjacent memory cells. Figure 8 shows an anisotropic etch trench filled with polysilicon and oxidized to make an isolation trench. Because of the sharp edges resulting from oxidation, dislocations are created in the nearby active circuitry. Having done my PhD work worrying about dislocations in materials at high temperature, I look upon this problem affectionately.

An interesting problem that hit our industry about ten years ago and still confronts us today, is the concept of soft errors in DRAMs. Figure 9 shows the soft error rate for a 16 K DRAM. It really shows the error rate in a 16 K memory dependent on the incident alpha flux. In a memory storage cell, incident alpha particles dissipate in silicon by creating electron hole pairs. The electrons that the alpha particle creates can be collected in memory storage capacitors, where they disrupt a zero state to a one state or vice versa—depending on your nomenclature—and create a soft error. The arrows in the figure indicate the alpha particle flux from naturally occurring compounds like the plastics or ceramics used to encapsulate the circuits. This one observation by itself was enough to cause the entire DRAM industry to recalibrate itself. The industry had been on a relentless march to decrease the size of storage capacitors in their memory circuits until they ran into this fundamental limitation. This has caused the storage capacity to stay in the range of several hundred thousand electrons as the minimum storage size in any DRAM circuit.

Now, let's consider the five layers of metallization that our circuits are going to have in the future. In a typical cross section of two layers of metal, with a tungsten plug in the via between them, you might find a little bit of undercut. What if you have millions of these on a circuit? The undercut increases the local current density and causes electromigration failures at that voided region underneath, which is totally

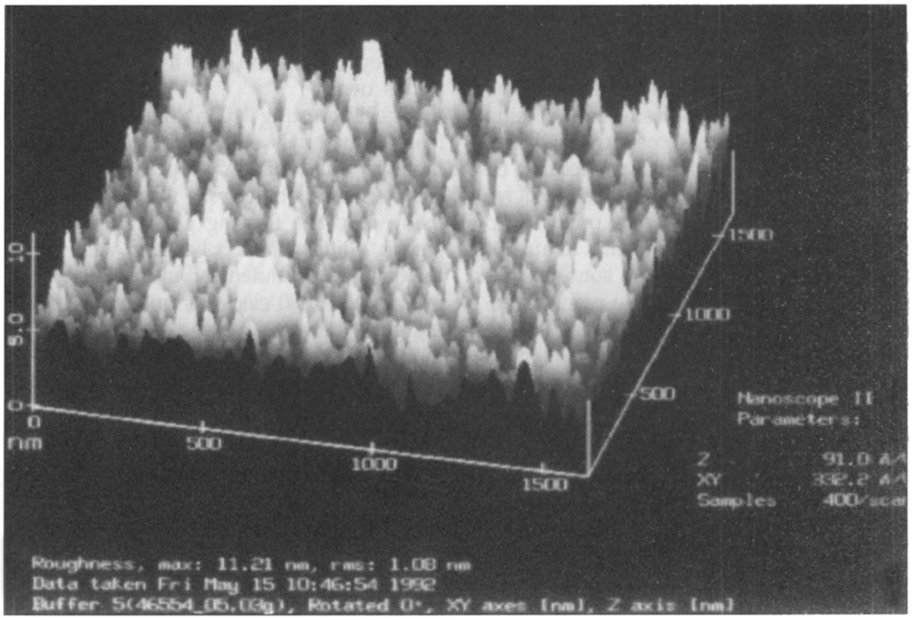

Figure 7. Atomic roughness shown by an atomic force microscopy scan of a silicon surface after exposure to deionized water. Source: Intel Corp.

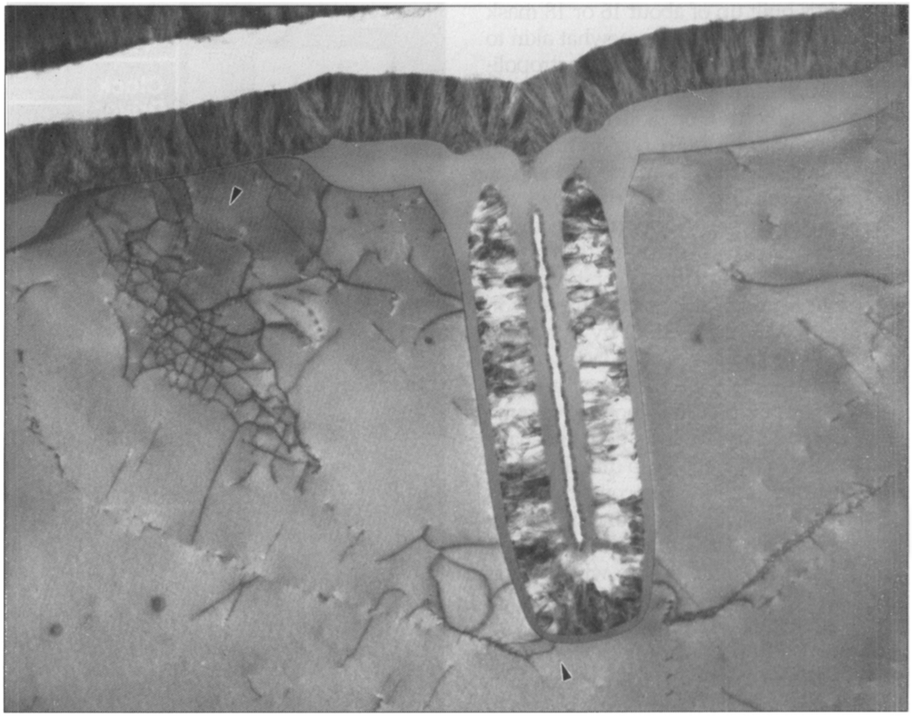

Figure 8. Dislocations generated during trench isolation. Source: Intel Corp.

invisible from above because it is covered by other metal layers, or covered by itself.

These are just a few of the problems facing the industry. Other areas of concern include everything from high-integrity dielectrics in the 50-angstrom thickness range; high and low dielectric constants, depending on whether you want storage capacitance or minimum interconnect resistance; planarizing materials when you have five layers of interconnects; encapsulating bare chips; plasma sources to do etching that won't destroy the oxides underneath the 50-angstrom floating gates, and so on. This wide range of topics generates a lot of research. Figure 10 shows just what the United States spends on semiconductor research. It is an ever-increasing curve. The one monotonically increasing function in our industry is how much we spend on research and development.

In 1992, commercial semiconductor houses (excluding those captive semiconductor facilities and companies such as IBM and Digital Equipment Corp.) spent in the range of $3.5 billion a year just on research and development for semiconductor materials. In 1992, about $4.5 billion was spent on research by the industry, including captive companies, and about $1 billion was spent by the U.S. government. The industry spends about $200 million in SEMATECH and about $25 million through the Semiconductor Research Corporation (SRC) supporting university research. Of this, I would conclude that about $2 billion is spent directly on materials-related process technology. I would also suggest that one of the reasons the U.S. semiconductor industry has had a resurgence recently is that these research dollars have continued to be spent year after year in increasing amounts to solve the basic problems that confront the industry.

Major Players

Let me now make a few projections about startups, the future of this industry, and who the future players will be. I'll start by looking at who the major players are now, how much market share they command, and what changes we have seen in the past few years. The top ten semiconductor companies in the world consist of three U.S.-based companies, six Japanese companies, and one Western European company. For the last decade, they have held about 55 percent market share, which means that they have been growing as fast as the industry as a whole. If you included the next 10 companies on top of this, you would find that they comprise most of the rest of the industry. In the future, semiconductors will be an industry for big companies; you should not expect to see many startups.

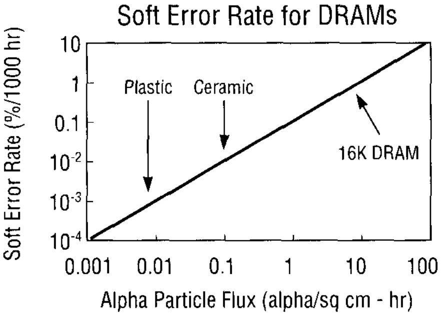

Figure 9. Soft error rates for DRAMs. The arrows indicate alpha particle flux from the plastic or ceramic materials encapsulating the circuits. Source: Intel Corp.

And any startups that you do see in this industry will not be manufacturing-related, they will be design-intensive only, and they will buy their manufacturing capacity from one of the larger companies.

Figure 11 shows the number of startup companies in Silicon Valley over the last six or seven years. In 1985, there were some 60 companies started in Silicon Valley in the semiconductor arena, all only performing design functions. Last year only 18 semiconductor companies were started. Again, they were all design-related. Now that may sound like a lot. It sounds like a lot to me because in the early 1970s there were only 10 or 15 semiconductor companies a year

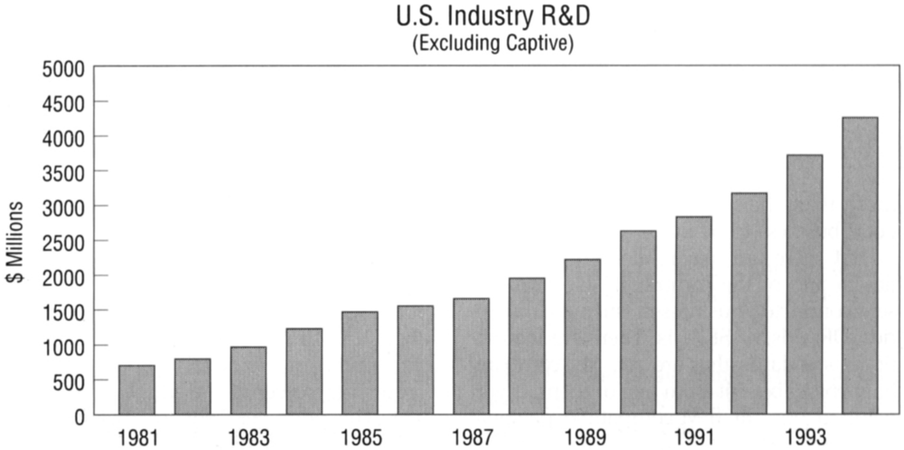

Figure 10. Growth of U.S. semiconductor industry research and development funds. Source: Annual Reports, Analyst Reports.

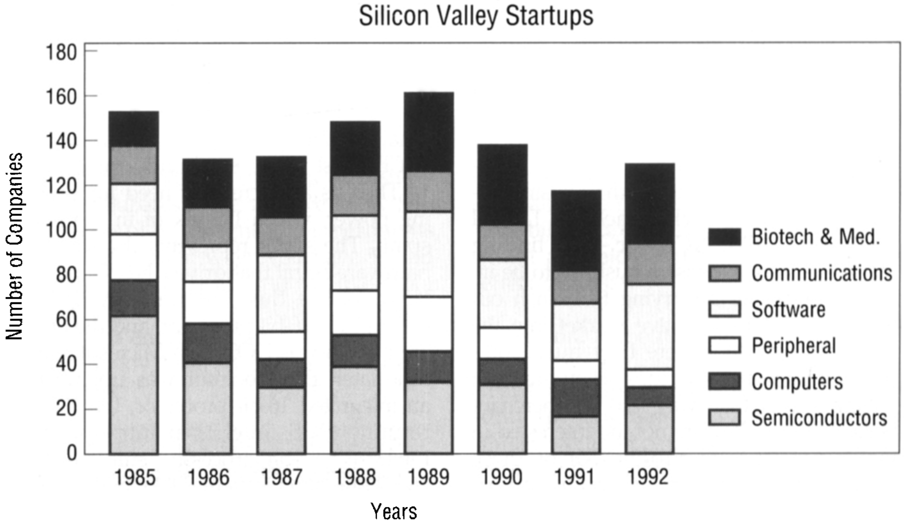

Figure 11. Silicon Valley startup companies. Sources: Ernst and Young, Dialog.

started in Silicon Valley. So if you consider bulk numbers alone, there are about as many companies starting up today in Silicon Valley as there were when Silicon Valley got its name. The difference today is that all of the startup companies are design houses; none of them are complete design and manufacturing facilities. In Silicon Valley and around the world, biotech and software companies are the ones that are starting up in greater numbers.

There must be some barriers to entry. One is the cost of doing business. It is difficult to go to your local venture capitalist and say, "I want a billion dollars to build a manufacturing facility." So there is a major necessity for capital. Intellectual property is an important consideration as well. It turns out that most of the basic patents in the semiconductor industry were filed many years ago. If you want to get into this business, it is hard to do so without violating those patents. Many companies make a major business of collecting money from their royalties. For example, a company like Texas Instruments collects about $300 million a year in patent royalties—somewhat more than their operating income. It is a good business to be in.

But there are major intellectual properties in any of the commodity electronic or semiconductor businesses, such as memories, DRAMs, or SRAMs. There are international standards that are set. My company happens to benefit from one of them. If you want to buy an IBM-compatible personal computer today, it is usually an "Intel-architecture-compatible personal computer" because the standard is an Intel microprocessor. It is difficult to break into the microprocessor business because semiconductor companies protect their intellectual property with patents. This situation also makes it difficult for new and emerging technologies—such as ferroelectrics or gallium arsenide, an emerging technology since I was a graduate student at Stanford—to break into the business.

In addition, many companies who for years thought they needed to have semiconductor technology to do their business—companies like Hewlett-Packard, Digital Equipment, IBM, NCR, etc.—are finding that it is a very expensive business to be in. These companies are trying to branch out from their internal captive market into the merchant market, where they must compete with any new startups. They are even willing to buy their way into the merchant business. Plus, they are not too interested in making short-term profits. All of these things make it difficult for startup companies to enter the semiconductor field, which favors the established companies, rather than the startups. That is one of my major conclusions. If you look at "Silicon Valley, What Next?," you'll find that the majors will continue to grow, and that new companies in this field will be design-intensive, not full manufacturing houses.

Now, assuming that my thesis is accurate, we end up with ever-increasing levels of integration, as well as ever-increasing microprocessor power and memory power. What in the world are you going to do with it? Why would you buy the next generation? The answer is that you *should* buy it.

Taking A Look Back

Figure 12 shows processing power relative to generations of microprocessors. I could have illustrated this for the 68000 series from Motorola, or just about any other microprocessor family. It is interesting to look back. The 8088, which was introduced by Intel in 1979, didn't find its way into the IBM personal computer until 1981. It's refreshing to recall that the IBM PC has been around for only 12 years. Some of our children think that PCs have been here forever. Some of us didn't have PCs when we went to graduate school. Some of us still don't know how to use our PCs. But if you look at the first IBM PC, which was built from a 16-bit processor with an 8-bit IO (because all the peripherals were 8-bit at that time), you'll see that it had a processing power of about 0.3 MIPs. A PC built using the current state-of-the-art Pentium™ processor, has a processing power of 100 MIPs, or 300 times that. The price per MIP for that first IBM PC—if you had bought one in that time frame, with a printer, etc.—was $8,000 per 0.3 MIPs. Today you can buy an Intel486™ CPU-based PC with 50, 60, or 70 MIPs for under $2,000; so the cost per MIP has gone down substantially.

Need for Processing Power

In the scenario in which cost per MIP has gone down and MIPs are going up, you may ask, "Why should I want this?" I'd like to show you why you should want it, and what is going to drive not only our industry, but the rest of the semiconductor industry.

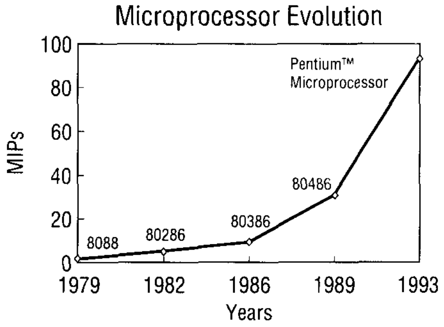

Figure 12. Processing power with generations of microprocessors. Source: Intel Corp.

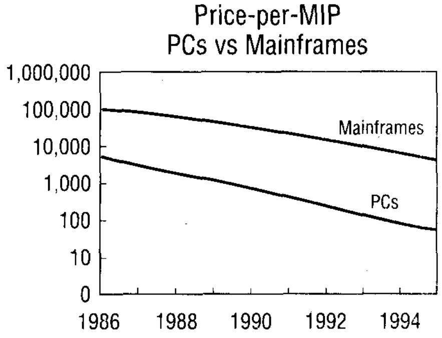

Figure 13. Price per MIP of PCs versus mainframe computers, as a function of time. Source: Dataquest.

There is an increasing need for processing power, which I think of in terms of a spiral. The spiral represents the software/hardware spiral that drives the PC industry. It works like this: The first IBM PC was driven by a 16-bit processor, and the operating system was DOS from Microsoft. Sometime later, the processor was upgraded to an integrated 16-bit processor, the 286, still running DOS. Intel then introduced the Intel386™ 32-bit processor, still running a 16-bit operating system, DOS, but it ran it a lot faster than the 286, so a lot of people bought it. Then Microsoft came out with the Windows™ operating system, still 16-bit, not taking advantage of the 32-bit processor. You couldn't run Windows™ on a 286, you had to have a 386 chip. So everybody had to have a 386 CPU. And Windows just barely ran on a 386 CPU, so you really needed an Intel486™ CPU for it. Every time new software comes out, it causes the hardware to lag. New hardware is needed to run the software, because that software hardly runs well on the processor on which it was introduced. We are just about to get our next introduction of operating systems from Microsoft and from a variety of other operating system vendors. They will run on an Intel486™ chip, but they won't run well, and you will need a Pentium™ processor, or the generation after that, Intel's P6, to run them.

We've got you, folks, because most of you don't like your keyboards; most of you like wonderful graphics; and most of you would like to have natural data. You would also like to have good, sharp, VCR-type images. You would like to have video. You would like to be able to talk to your PC. You would like it to do all sorts of things which require new applications, and new applications eat up processing memory, power, and DRAMs. This is the reason the DRAM and microprocessor business has been continuously expanding. The same goes for the hard-disk and disk-drive controller business. If you buy a hard-disk drive today with less than 200 megabytes, you will regret it because you know the next operating system is going to eat up your entire hard-disk space. Then what will you do?

So I want to conclude this part of the talk with a curve that shows the price per MIP as a function of time (Figure 13). This is a basic problem for the mainframes of many manufacturers because, although the price per MIP in a mainframe is going down substantially over time, the price per MIP in a PC is about 100 times cheaper. You can spend about $200 per MIP in a PC, or you can spend in the range of several thousand dollars per MIP in the mainframe. That is the reason people don't buy mainframes any more. And there is no apparent slowdown in this curve. In fact, as the MIPs and the microprocessor continue to move forward, they will continue to drive down the cost-effectiveness of these localized work stations.

Return of the Dinosaurs

Let me conclude as follows. I don't think there is any slowdown in the technology. All it takes is a few billion dollars a year of R&D. U.S. companies are willing to spend that, and so are many companies in other countries. But it is a case of "the return of the dinosaurs"; those companies who were thought to be too large, too lethargic, and not entrepreneurial enough, are the only ones with the capital to invest to make this happen.

- There is certainly no shortage of materials-related problems to make this industry move forward. It is built on materials science. All aspects of it—from growing 12-inch-diameter defect-free silicon wafers to diffusion in the solid state, thermodynamics and solid-state kinetics, growing and etching films, resist, photolithography, and the five or six layers of metal interconnects—are materials problems. In my opinion, these problems are not handled well by most of the electrical engineers who go into this field, as opposed to the materials scientists. So we need more materials scientists.
- I think there will be fewer and fewer startups in this industry that do design and also manufacture the product. There will still be design houses, but the people who do the manufacturing are going to be the folks who can afford to spend a billion dollars a year or so in capital investment.
- I think there is an increasing role for university research in this field. The university research system in the United States is probably the best basic research institution in the world. Today that research institution, through a decrease in government funding, is being starved for funds; but I think it is still the premier research system, and will in the future play an increasing role in solving some of our basic problems. That has to happen for two reasons: (1) the problems need to be solved, and (2) we need the graduate students who come and work in industry. There is no shortage of problems to be solved at either end.
- My last prediction is that several of the captive semiconductor manufacturers around the world, those companies who use their proprietary semiconductor output to feed their own operations, will first go into the merchant semiconductor business. Then, as they come under increasing cost pressure, these companies will decide that they don't want to become semiconductor manufacturers at all and will go out of the semiconductor business. This is not a particularly attractive projection but it is true in the United States, Western Europe, and Japan.

Academia vs. Industry

Let me conclude my presentation with a short story about why I left Stanford and went to work at Intel 20 years ago, and why it may be a good opportunity for some of you to do the same thing. I spent several years as a graduate student and about 10 years on the faculty of Stanford; I thought I was reasonably educated about the high-tech industry, when I got the itch to go off and do something besides technical papers. I believed that high-tech industry was really the semiconductor industry, which at that time was just getting started. So I visited Intel and began asking learned, academic questions about why this happens, and why turning the knob that way makes that happen, and I found out that there was an immense amount of black magic involved. Over the years, we have fixed a lot of that. We use statistics tremendously well. I wish that those of you who are here from academia would begin to teach statistics in your courses. We use statistics extensively, we use databases more; and we know how to do complex designs of experiments, which is important when you have 50 variables and you don't want to run five billion experiments to find out which variable is causing the effect you are looking at.

But frankly, I still see a lot of black magic in our industry, and we need a lot of bright, young minds coming in with new ideas and new concepts. I see just as much opportunity in the industry today as there was 20 years ago. It is a lot bigger than it was 20 years ago. The problems we are trying to solve are a lot tougher—or a lot smaller, depending on how you look at them in terms of going down to angstroms of thickness as opposed to tens or hundreds or thousands of model layers before.

Material Matters is a forum for expressing personal points of view on issues of interest to the materials community.

But I think the opportunities are there. I think that where the United States and companies like my own are concerned, we have proven the naysayers incorrect—those who said we couldn't support a manufacturing industry, a technology-intensive industry. We were very happy last year when we outran several foreign competitors and became the largest semiconductor company in the world. We intend to stay there. I have a running bet with my friends at Motorola, who say they are out to get us; but we are going to stay ahead of them. We can only do it, however, with the best and brightest minds from the university system joining us. So, to the students here who received the graduate student awards, send me your resume. Thank you very much.

Question: What is Intel doing to stem the tide of escalating capital costs in the industry?

Answer: One of the things we try to do is play our part in SEMATECH, which is the industry consortium that concerns itself with manufacturing cost effectiveness. We also try to be as competitive in our manufacturing as we can to dilute the impact of the capital cost by amortizing it over as much output as possible. But, frankly, while we are anti-inflationary, the equipment suppliers are the most inflationary industry I know. And it is getting to be a substantial problem. As I said, when you look at 70% or 80% of the cost of the wafer being capital depreciation, you don't have a lot of room left. But that, in my opinion, is why this industry is going to be dominated by big players. I have not seen a viable proposal for a small cost-effective wafer fab facility, and this billion-dollar level is about the minimum at which you can operate for cost effectiveness.

Question: About eight years ago, in an edition of *Solutions*, the CEO of Intel stated that the company could not afford to continue in long-term research, that certain things wouldn't pan out in the next 20 years or so, and that they were going to stay in more evolutionary, short-term work such as better oxides, smaller linewidths, and so forth. At the same time, Texas Instruments, the example you cited, has continued to invest in quantum effects and very long-term devices and, as you said, is now living off its patents. And in ten years, when these patents have run out, TI may still be living off its patents. What is Intel's current philosophy in research? Is it revolutionary or, rather, evolutionary, as it has been in the past?

Answer: Intel started as an outgrowth of Fairchild R&D, which was a basic R&D lab. The founders of the company started it as a product development operation, not a basic research operation. That was the basis of the comment, made by either Bob Noyes or Gordon Moore, that you read in *Solutions* magazine. Our current position is still that our research outlook is typically no more than five years or so, about two generations of our product cycle, or two generations of the process technology. We have a small basic research group that does two things. The group conducts a small amount of basic research in-house, and it works with universities around the world to monitor the key research elements that we think we should follow. We have not set up a large basic research lab, and we don't have any plans to do so in the near term, either. We are still very much an evolutionary research and development establishment, as opposed to a basic research establishment.

Question: I have two questions. First, my resume is ready—when can I give it to you? Second, you mentioned that the minimum lithographic dimension is targeted for about 0.18 microns. You also mentioned about five to six levels of metallization. Now each layer of metallization has got some variability and it gets compounded as it goes from one level to another. How can you achieve these goals with the existing technology?

Answer: Oh, you can't. My comment is that that will be the common technology in the year 2000. And, by the way, people do four and five levels of metal today, so five or six levels of metal is no big extension of the current technology. Certainly, the minimum lithographic dimension for volume production with 0.2 micron is a couple of generation steps forward in the wafer stepper business. But my prediction, within that time frame, is that the solution will not be x-ray lithography, but optical lithography.

Question: I have heard gripes from industrialists that university research is meaningless, and that universities are simply training grounds for scientists. What do you see as the role of U.S. industry, considering that the semiconductor field is becoming more capital intensive and that the gap between university research and industrial needs is going to widen? What do you see as industry's role in making university research more meaningful?

Answer: Well, I think the industry's role is severalfold. One role is to fund meaningful research in the university environment through agencies such as SEMATECH and SRC. Secondly, all of the major companies have direct relationships on their own with professors, graduate students, and the major universities, and they work on research projects or fund them directly. And thirdly, I think industry can provide a training ground or an industrial experience for university faculty members and graduate students, an opportunity for them to come and spend some time in the industrial research environment.

Question: In your slide on critical technologies, you pointed out that optolectronics is high on the list. Can you look into your crystal ball for the year 2000? What role do you see for optoelectronics in your product mix?

Answer: The real role for optical components in our product mix is for running at greater than 100 or 200 megahertz on a chip, for communicating on chips, and possibly for chip-to-chip light transmission through a local gallium arsenide diode or something of that sort, heteroepitaxially grown on the silicon.

Question: I think you may have opened somewhat of a Pandora's box. With the number of students out there looking for jobs right now, I don't envy your secretary in the next few weeks. And that brings me to a plea. Looking at some of your viewgraphs, the points are well made and also fairly clear, but you are sending an additional message to materials scientists, and electrical engineers as well. We see this in our admissions process. The students know where the jobs are, and they are not in materials, and they are not even in the circuit areas; they are in software. And you folks need to be hiring those students in larger quantities if you want to have the supply that I do believe you are going to need, say, in the next decade. So, please hire more materials scientists, as well as electrical engineers.

Answer: We do in fact hire a bunch of electrical engineers, and a bunch of computer scientists, and a bunch of software engineers, and a bunch of materials folks. You know, in reality it is probably a relatively tough period of time, looking at the electronics industry as a whole, with IBM downsizing, and DEC downsizing. Our stated role, although we don't have a quota, is to try to hire 500 new college graduates per year as a minimum base line, to keep that pipeline fed. This year I am anticipating that the number is going to be closer to 750 to 800. Several of us do have a materials background, and we like to run into folks that we can communicate with, although there is clearly a limited appetite that the industry has for people in this discipline.

***Craig R. Barrett** is executive vice president and chief operating officer of Intel Corporation, which recently was named the largest manufacturer of semiconductors in the world. Barrett received his BS, MS, and PhD degrees in materials science at Stanford University. Following completion of his PhD degree, he did a postdoctoral fellowship in England and then, in 1966, joined the faculty at Stanford, where he stayed for the next eight years. In 1969 he received the Hardy Gold Medal of the AIME. Also, with William Nix and the late Alan Telleman, he co-wrote a well-known textbook in materials science and engineering. Barrett joined Intel in 1974, became a general manager in 1980, and by 1984 was a vice president. In 1990 he was named executive vice president, and recently was appointed chief operating officer of the company.* □