Case 11

Intel Corporation: 1968–2013

Charles W.L. Hill
School of Business, University of Washington
Seattle, WA 981095, June 2013

INTRODUCTION

In 2012 Intel was the leading manufacturer of microprocessors for personal computers in the world, a position that it had held onto for more than two decades. Over 80% of all personal computers sold in 2012 used Intel microprocessors. The company reported revenues of $53 billion and net profits of $11 billion. Meanwhile, Intel's only viable competitor, AMD, which in the early 2000s had been gaining share from Intel, lost $1.2 billion on sales of $5.4 billion.

Despite its historic dominance, the future looked uncertain for Intel. The rise of mobile devices had led to a strong substitution effect, with sales of PCs falling as consumers switched to smart phones and tablets for many of their computing needs. In the first quarter of 2013, global PC sales fell 14% on a year over year basis according to the research firm IDC. This was the worst yearly decline since IDC started tracking PC sales in 1994, and the fifth quarter in a row that PC sales had fallen. At the same time, sales of smart phones and tablets were booming. IDC predicted that sales of tablets would grow almost 60% in 2013, and that tablet shipments would exceed those of portable PCs.[1]

The crux of the problem for Intel is that most tablets and smart phones used microprocessors that are based on technology licensed from ARM Holdings PLC, a British company whose chip designs are valued for their low power consumption, which extends battery life. While Intel has a line of chips aimed at mobile devices—the Atom chips—microprocessors incorporating ARM's technology were found on 95% of smart phones in 2012 and over 30% of all mobile computing devices, a category that includes tablets and PC notebooks.[2] Moreover, in 2012 Microsoft issued a version of its Windows 8 operating system that ran on ARM chips, rather than Intel chips, creating a potential threat to Intel's core PC business.

THE FOUNDATION OF INTEL

Two executives from Fairchild Semiconductor, Robert Noyce and Gordon Moore, founded Intel in 1968. Fairchild Semiconductor was one of the leading semiconductor companies in the world and a key enterprise in an area south of San Francisco that would come to be known as Silicon Valley. Noyce and Moore were no ordinary executives. They had been among the eight founders of Fairchild Semiconductor. Noyce was general manager at the company, while Moore was head of R&D. Three years previously, Moore had articulated what came to be known as *Moore's Law*. He had observed that since 1958, due to process improvements the industry had doubled the number of transistors that could be put on a chip every year (in 1975 he altered this to doubling every two years).

Fairchild Semiconductor had been established in 1957 with funding from Sherman Fairchild, who had backed the founders on the understanding that Fairchild Semiconductor would be a subsidiary of his Fairchild Camera and Instrument Corporation on New York. By 1968 Noyce and Moore were chaffing at the bit under management practices imposed from New York, and both decided it was time to strike out on their own. Such were the reputations of Noyce and Moore that they were able to raise $2.3 million to fund the new venture "in an afternoon on the basis of a couple of sheets of paper containing one of the sketchiest business plans ever financed".[3]

When business reporters got wind of the new venture, they asked Noyce and Moore what they were intending to do, only to be greeted by vague replies. The two executives, however, knew exactly what they were going to do—manufacture silicon memory chips—they just didn't want potential competitors to know that. At the time, sales of mainframe computers were expanding. While these machines used integrated circuits to perform logic calculations, programs and data were stored on magnetic devices. Although inexpensive to produce, it was relatively slow to access information on a magnetic device. Noyce and Moore knew that if they could build a silicon based integrated circuit that could function as a memory device, they could speed up computers, making them more powerful, which would expand their applications and allow them to shrink in size.

These memory chips were knows as *dynamic random access memories* (DRAMs). While much of the theoretical work required to design an integrated circuit that could function as a memory device had already been done, manufacturing DRAMs cost efficiently had so far proved impossible. At the same time, some key research on manufacturing was being done at Fairchild. This research included a technique known as *metal oxide on silicon*, or MOS. Noyce and Moore wanted to mass-produce DRAMs, and after looking at other possible alternatives, they concluded that commercializing the MOS research was the way to do it. This prompted some cynics to note that Intel was established to steal the MOS process from Fairchild.

ANDY GROVE

To help them, Noyce and Moore hired a number of researchers away from Fairchild, including, most notably, a young Hungarian Jewish émigré called Andy Grove. At Fairchild, Grove had reported directly to Moore. At Intel he became the director of operations with responsibility for getting products designed on time and built on cost. Through the force of his own personality, Grove would transmute this position into control over just about everything Intel did, making him effectively the equal of Noyce and Moore, long before he was elevated to the CEO position in 1987.

Grove was an interesting character. Born in 1936, he went into hiding when the Germans invaded Hungary during World War II and managed to escape the Holocaust. After WWII, the tyranny of the Germans was replaced by the tyranny of the Soviets as Hungary became a satellite state of the Soviet Union. In 1956, after the failure of an uprising against the Soviet puppet government, Grove escaped across the border to Austria, and made his way to the United States. He put himself through college in New York by waiting on tables, and then went to UC Berkley for graduate work, where he received a Ph.D. in chemical engineering in 1963. His next stop was Fairchild, where he worked until Moore recruited him away in 1968.

Over the next three decades, Grove would stamp his personality and management style on Intel. Regarded by many as one of the most effective managers of the late twentieth century, Grove was a very demanding and according to some, autocratic leader who set high expectations for everyone, including himself. He was detail orientated, pushed hard to measure everything, and was constantly looking for ways to drive down costs and speed up development processes. He was known for a confrontational "in your face" management style, and would frequently intimidate employees, shouting at those who failed to meet his expectations. Grove himself, who seemed to enjoy a good fight, characterized this behavior as "constructive confrontation". He would push people to their limits to get things done. As he once noted, "there is a growth rate at which everybody fails, and the whole situation results in chaos. I feel it is my most important function. . . . to identify the maximum growth rate at which this wholesale failure begins".[4]

Grove demanded discipline, insisting for example, that everybody be at their desks at 8 a.m., even if they had worked long into the night. He instituted a "late list", requiring that people who arrived after 8 a.m. sign in. If people arrived late for meetings, he would not let them attend. Every year he sent around a memo to employees reminding them that Christmas Eve was not a holiday, and that they were expected to work a full day. Known as the "Scrooge memo", many would be returned with nasty comments scrawled over them. *May you eat yellow snow*, said one. A very neat man, if people's desks were messy, Grove would publically criticize them. According to one observer, "Andy Grove had an approach to discipline and control that made you wonder how much he had been unwittingly influenced by the totalitarian regime he had been so keen to escape".[5]

Grove controlled managers through a regular budgeting process that required them to make detailed revenue and cost projections. He also insisted that all managers establish medium term objectives, and a set of key results by which success or failure would be measured.

He instituted regular one-on-one meetings where performance was reviewed against objectives, holding managers accountable for shortfalls. He also required monthly management reviews where managers from different parts of the company would meet to hear a presentation of its current strengths, weaknesses, opportunities and threats. The goal was to get managers to step back and look at the bigger picture, and to encourage them to help each other solve problems.

Grove would also practice management by walking around, inspecting facilities and offices, demanding that they be clean, something that earned him the nickname "Mr. Clean". He pushed the human resource department to institute a standard system of ranking and rating that had four performance categories; "superior", "exceeds expectations", "meets expectations", or "does not meet expectations". People were compared against others of their rank. Pay raises and later, stock option awards were based on these rankings.

Despite his autocratic style, Grove was grudgingly admired within the company. He was a brilliant problem solver, a man with tremendous control of facts and details, someone who was determined to master the challenging technical projects that Intel was working on. Moreover, while he drove everyone hard, he drove himself harder still, thereby earning the respect of many employees.

THE MEMORY CHIP COMPANY

Making a DRAM using MOS methods proved to be extremely challenging. One major problem—small partials of dust would contaminate the circuits during manufacturing, making them useless. So Intel had to develop "clean rooms" for keeping dust out of the process. Another was how to etch circuit lines on silicon wafers, without having the etched lines fracture and break as the wafer was heated and cooled repeatedly during the manufacturing process. The solution to this problem, identified by Moore, was to "dope" the metal oxide with impurities, making it less brittle. Intel subsequently went to some lengths to keep this aspect of the manufacturing process secret from competitors for as long as possible.

Intel, of course, was not alone in the race to develop a commercial process for manufacturing DRAMs. Among the potential competitors was another semiconductor company started in 1969 by Jerry Sanders, a former marketing director at Fairchild. Sanders started his company with the help several other Fairchild employees who had not been recruited by Intel. Called Advanced Micro devices, or AMD, the company found it tough to raise capital until it received an investment from non other than Robert Noyce, who saw something he liked in the flamboyant Sanders.

Driven by constant pressure from Andy Grove, whose "in your face" management style was bearing fruit, albeit at some human cost, by October 1970 Intel succeeded in producing a DRAM chip, named the 1103, in relatively high yields (which implied that relatively few chips had to be discarded). The 1103 could store 1,024 bits of information (zeros or ones), which was 4 times as much as the highest capacity semiconductor memory device currently available. Since the fixed costs required to establish a manufacturing facility were very high, the key to making money on the 1103 was high yields and high volume. If Intel could achieve both, unit costs would fall enabling Intel to make a lot of profit at low price points. In turn, low prices implied that DRAMs would start to gain wide adoption among computer manufacturers.

The 1103 put Intel firmly on the map. The chip soon became the memory technology of choice for computer makers, and by the end of 1971, 14 out of the world's 18 leading mainframe computer makers were using the 1103. However, Intel did not have the market entirely to itself. Computer makers did not want to become dependent upon a single source of supply for critical components. To avoid this, most computer makers mandated that components had to be at least duel sourced, and for Intel, this meant that if it wanted business, it had to license its technology to other companies. Intel first licensed the rights to produce the 1103 to a Canadian firm, MIL, in exchange for an upfront payment and per unit royalty fee. Before long, MIL was competing against Intel in the market for the 1103, but MIL made a critical mistake in their manufacturing processes, and it wasn't long before a stream of former MIL customers were knocking on Intel's door.

Along the way, Intel received an inquiry from two disgruntled engineers at Honeywell, asking if Intel was interested in building memory systems. The idea was to mount thousands of 1103 chips on a circuit board that could then be plugged into a mainframe computer to increase its memory capability. Impressed by the idea, Intel promptly hired the two engineers and set up a division to do this. Before long, the new division was selling circuit boards to customers running IBM mainframes. This was something of a coup: IBM would not even consider buying the 1103, and had started making its own memory chips. Now Intel had access to a formerly closed market that accounted for 70% of all memory sales.

Around the same time, an accidental discovery at Intel led to a second product line—erasable programmable read only memory (EPROM). Read only memory chips (ROM) were finding wide applications in computing. ROM had desired data, a program for example, permanently burnt into its circuits. ROM was used to store programs, such as a machine operating system, or part of that system. The troubling thing about ROM is that if an engineer made a mistake in programming the chip, he would have to burn another chip, which was a painstaking and time consuming process. While exploring the reason for failure of 1103 chips in the manufacturing process, Dov Froham, another ex Fairchild researcher at Intel, found that the cause was that some of the "gates" inside the chips had become disconnected; they were floating. Froham realized that this flaw in the 1103 had a potential use; it might enable an engineer to design a ROM chip that could be programmed with ease in a few minutes. Moreover, he found that the data on such chips could be erased and rewritten by shinning an ultra violet light on it and the EPROM was born.

Engineers loved the EPROM chip, and once Intel solved the manufacturing problem and started to produce EMROM chips in large quantities, demand surged. Better still, for two years Intel had a virtual monopoly on the product. While other companies tried to produce similar chips, they were unable to solve the manufacturing problems, enabling Intel to charge a relatively high price for a product whose cost was falling every day with advances in cumulative volume.

THE BIRTH OF THE MICROPROCESSOR

By 1971 Intel had already created two revolutionary innovations in the semiconductor industry, the DRAM and the EPROM chips. A third, the microprocessor, was also created that year. The microprocessor was born out of an inquiry from a Japanese company. The company asked Intel if it could build a set of eight logic chips to perform arithmetic functions in a calculator it was planning to produce. Intel took on the project. Ted Hoff, one of the inventors of the DRAM, wondered if it might not make more sense to build a miniaturized general purpose computer, which could then be programmed to do the arithmetic for the company's calculator.

The project was given to Federico Faggin, an Italian engineer who made some of the basic breakthroughs on MOS technology while working at Fairchild. Although the Japanese company subsequently decided not to build the calculator, Intel pushed ahead with the project. Faggin, who worked 12 to 14 hour days for weeks on end, produced several prototypes in short order. (A source of irritation for Faggin was that despite the long hours, his boss, following Grove's lead, constantly complained that Faggin was late for work!)

Due to Faggin's efforts, by November 1971 Intel had its third product, the 4004 microprocessor. In an article in Electronic News that accompanied its introduction, and which described the 4004 as a computer on a chip, Gordon Moore heralded the 4004 as "one of the most revolutionary products in the history of mankind". No one paid much attention. People in the computer industry viewed the 4004 as a fascinating novelty. Although small and cheap, it could only process 4 bits on information at a time, which made it slow and thus unsuitable for use in the computers of the time. The 4004 was followed by the 8008 microprocessor, which could process eight bits of information at a time. Although faster, it too was a product in search of a market. In an attempt to speed adoption, Intel started to sell development tools that made it easier and faster for outside engineers to develop and test programs for new microprocessors. Slowly the microprocessor began to make inroads into the computer industry, primarily in peripherals such as printers and tape drives.

THE PERSONAL COMPUTER REVOLUTION

By the mid 1970s and embryonic new industry was appearing, the personal computer industry. A company called MITS based in Albuquerque, New Mexico produced the first true personal computer. The MITS Altair used an Intel 8080 microprocessor, which was priced at \$360. The first program offered for sale with the Altair was a version of the BASIC programming language, written by Bill Gates and Paul Allen, and designed to run on the 8080. The two had moved to Albuquerque to be near to MITS, and they had established a company of their own, Microsoft. The Altair was sold primarily to hobbyists who wanted to write computer code at home (for which Microsoft Basic came in handy).

In short order, a number of companies sprung up making personal computers. The most successful of the early companies was Apple Computer, which introduced its revolutionary Apple II in 1977. By this time, a number of other companies were also producing microprocessors, including Motorola, whose processor Apple used in the Apple II. The Apple II was a big commercial success, in no small part because it was easy to use for, and because one of the most successful early programs, a spreadsheet called VisiCalc, was written to run on the Apple II.

The commercial success of the Apple II got the world's largest computer company, IBM, to take the nascent personal computer seriously. IBM started to develop its own personal computer in 1979 in a top-secret project. To speed the product to market, IBM took a monumental strategic decision—it decided to use "off the shelf components" to build the PC rather than develop everything itself, which had been the norm at IBM. Originally the company planned to use a microprocessor from Motorola and an operating system called CP/M from a company called Digital Research. However, Motorola was late developing its product, and Digital Research's CEO, Gary Kildall, proved to be difficult to work with. Casting around for alternatives, IBM contacted Intel, offering to purchase it's latest microprocessor, the 8088, which was a derivative of Intel's 8086 chip. However, IBM did not tell Intel what the microprocessor was to be used for (originally Intel was told that it was to go in a printer). As part of the deal, IBM insisted on alternative sources for the 8088. Reluctantly Intel allowed AMD and a number of other companies to produce the 8088 under license. A 1982 cross licensing agreement with AMD, which gave AMD the right to produce the 8088 chip, would come to haunt Intel for years to come.

For the operating system of its first PC, IBM decided to use MS-DOS, a Microsoft operating system. Originally developed by Seattle Computer, and called Q-DOS (which stood for quick and dirty operating system), Q-DOS was purchased by Microsoft for $50,000 when Bill Gates heard that IBM was looking for an operating system. Gates renamed the product, and quickly turned around and licensed MS-DOS to IBM. In what was to be a stroke of genius that had enormous implications for the future of all parties involved, Gates, sensing that IBM executives were desperate to get their hands on an operating system in order to get the IBM PC to market on time, negotiated a nonexclusive license with IBM.

Executives at Intel, who by now had realized that IBM was developing a personal computer, were profoundly unimpressed with the choice of MS-DOS and Microsoft. After a visit to Microsoft, one Intel executive noted: "These people are flakes. They're not original, they don't really understand what they are doing, their ambitions are very low, and it's not really clear that they have succeeded even at that."[6] For its part, Microsoft had to produce a version of MS-DOS that would run on the Intel microprocessor. From now on, like it or not, Microsoft and Intel would be joined at the hip.

Introduced in 1981, the IBM PC was an instant success. To stoke sales, IBM offered a number of applications for the IBM PC that were sold separately, including a version of VisiCalc, a word processor called EasyWriter, and well-known series of business programs from Peachtree Software. Over the next two years, IBM would sell more than 500,000 PCs, seizing market leadership from Apple. IBM had what Apple lacked, an ability to sell into corporate America.

As sales of the IBM PC mounted, two things happened. First, independent software developers started to write program to run on the IBM PC. These included two applications that drove adoptions of the IBM PC: word processing programs (Word Perfect) and a spread sheet (Lotus 1-2-3). Second, the success of IBM gave birth to clone manufacturers who made "IBM compatible" PCs that also utilized an Intel microprocessor and Microsoft's MS-DOS operating system. The first and most successful of the clone makers was Compaq, which in 1983 introduced its first personal computer, a 28-pound "portable" PC. In its first year, Compaq booked $111 million in sales, which at the time was a record for first year sales of a company. Before long, a profusion of IBM clone makers entered the market, including Tandy, Zenith, Leading Edge, and Dell Computer. This entry led to market share fragmentation in the PC industry.

By 1982, Intel had a replacement chip ready for the IBM PC, the 80286 microprocessor. The 80286 was desperately needed since the 8088 was painfully slow running some of the newer applications. IBM introduced a new PC, the AT, to use the 80286 chip, and priced it at a premium. Demand was so strong that IBM put the AT on allocation, which opened the door to clone makers, particularly Compaq. By now, 70% of the microprocessors sold to PC manufacturers were made by Intel, with AMD accounting for a significant portion of the remainder. For the 80286, Intel had cut the number of licenses down to 4. It also ran an intensive marketing and sales campaign, called Checkmate, which was successful in getting many Original Equipment Manufacturers (OEMs) to use Intel's version of the 80286 in their machines.

THE DRAM DEBACLE

In 1984 Intel booked revenues of $1.6 and made almost $200 million net profit, up from $134 million in revenues and $20 million in net profit a decade earlier. The growth had been dramatic. However, Intel's share of the DRAM market had been sliding for years. New entrants, particularly from Japan, had been grabbing ever more DRAM sales. They had done this by undertaking large scale investment to build efficient fabrication facilities (fabs) and paying meticulous attention to quality and costs, doing everything possible to drive up yields. One source suggested that while peak yields and U.S. DRAM plants, such as Intel's, were around 50%, in Japan they were closer to 80%. This translated into a huge cost advantage for the Japanese producers.

The American manufacturers, Intel included, had made the crucial mistake of underestimating the Japanese threat. Demands from computer companies for second sources had helped to facilitate diffusion of the underlying product technology and commoditized DRAMs. In such a market, advantage went to the most efficient, and this was the Japanese. Moreover, Japanese companies seized the lead in developing more powerful DRAM chips. While Intel had created the market for DRAMs, and dominated the market for 1K chips, in each subsequent generation it fell further and further behind. By 1983 when fifth generation 256K DRAMs started to appear, Intel was a year behind in the development cycle and as a consequence, was at a distinct cost disadvantage when it introduced its product.

Somehow, despite Grove's aggressive leadership, Intel's share had fallen to only 1% of the total DRAM market. To regain market share, management understood that Intel would have to build a new fabrication facility, at a cost of $600 million, and throw company R&D resources behind an effort to bring a next generation 1 megabyte DRAM chip to the market. To make matters worse, the DRAM market was in a big slump, bought on by overcapacity as a result of aggressive investments by Asian producers, and Intel was losing money in the DRAM business.

Faced with this bleak prospect, Intel's senior management had to decide whether to continue to compete in the DRAM business, the market they had created, or to focus resources on the more profitable microprocessor market. It was not an easy decision. Irrespective of the economics, there was enormous emotional attachment within the company to the DRAM business. Many at Intel wanted to build a 1 M DRAM. There were also valid arguments for staying in the DRAM business. Some thought that DRAMs were the technology driver in semiconductor manufacturing, and without the knowledge gained from making DRAMs, Intel's microprocessor business would suffer. In addition, there was the argument that customers would prefer to buy from a company that offered a full product range, and if it exited the DRAM business Intel would not be able to do that.

As Andy Grove describes it, a crucial point arrived when he and Gordon Moore were discussing what Intel's strategy should be. Grove asked Moore, "If we got kicked out, and the board bought in a new CEO, what would he do?" Moore's reply, "he would get us out of memories". Grove then said, "why don't we just walk out of the door, and come back and do it ourselves." It was one thing to make the decision, another to implement it. Grove removed the head of the DRAM division, recognizing that he was not the man to wield the ax, and replaced him with another manager, who promptly "went native" and started to argue for going ahead with the 1 megabyte DRAM chip. He too was replaced, and a year after the decision was made, Intel finally exited the DRAM business.

THE MICROPROCESSOR BUSINESS

In 1987 Gordon Moore stepped down as CEO of Intel, passing the torch on to Andy Grove, although Moore remained as Chairman. Grove, who held the CEO position through until 1998, and was then chairman until 2005, had no intention of letting Intel's dominance in microprocessors go the same way as its DRAM business.

Chip Design

By now, it was well understood at Intel that the market had an unquenchable thirst for more powerful microprocessors. Software was advancing rapidly, with new applications becoming available all the time. Running these applications quickly required more computing power, and users were willing to pay a premium for this. Intel knew that consumers would only be too happy to replace their old PCs with better, faster machines. It thus became critical to develop and introduce newer microprocessors. At the same time, the market demanded backward compatibility. The new machines had to run older software, and this implied that each new generation of chip should be able to run older programs. This requirement implied that too a degree, Intel was locked into the microprocessor architecture that had started with the 8086 (from which the 8088 was derived), and continued with the 80286. The next microprocessor in what was now known as the x86 architecture was the 80386, or i386 for short.

First introduced in October 1985, i386 was a 32-bit microprocessor that was much faster than the i286. Intel had been trying for over a year to get IBM to introduce a machine based on the i386, but IBM seemed to be dragging its feet. The problem for IBM was that an i386 PC would be very close in power to minicomputers that IBM was making a lot of money on. Fearing that i386 machines would cannibalize its product line, IBM seemed to want to keep the i386 of the market as long as possible. At the same time, Apple computer had introduced a new machine, the first Macintosh, which used a Motorola microprocessor. The Apple Mac was the first computer with a graphical user interface and a mouse. As it started to gain market share, Grove feared that the market might switch to the Apple standard, making it more critical than ever to get i386 based machines on the market.

Intel had an ally in Compaq Computer. In 1986, Compaq took advantage of IBM's sloth to be the first to introduce a PC built around the i386. Compaq seized the lead from IBM, other computer makers quickly followed, and from then on, IBM started to lose influence and share in the PC business. As the high margin i386 chip gained traction, Intel's sales exploded, hitting $2.9 billion in 1988, while profits surged to $450 million.

Over the next two decades Intel continued to drive the industry forward with regular advances in its x86 architecture. These included the i486 (introduced in 1989), the first Pentium chip (1993), The Pentium Pro (1995), various derivatives of the Pentium Pro architecture, and more recently, its 64-bit Core 2 Duo and Quad processor line, first introduced in 2006. The latest Intel processors have pushed the limits of performance by building two or four processors into a chip. Intel prices new chips at a premium then drops prices as manufacturing yields improve. It is not unusual to see prices drop by 30–50% in one year.

By continually increasing the performance of its chips, Intel was able to vanquish several potential competitors, including a series of fast chips from AMD in the early 2000s, and several chips based on an architecture known as reduce instruction set computing, or RISC, that during the 1990s seemed to threaten Intel's market dominance. One notable RISC chip arose out of an attempt by Apple, Motorola and IBM to seize momentum away from Intel with a RISC processor called the PowerPC. However, few companies outside of Apple adopted the processor. The limited volume meant high costs, which were further compounded by manufacturing problems at Motorola, and the PowerPC never gained wide acceptance. In 2006, Apple effectively killed the PowerPC when it announced that it would henceforth use Intel microprocessors in its machines.

Following Moore's law, successive generations of Intel chips have used ever-smaller micron geometries to cram ever more transistors on a chip. Intel's 8088 chip, introduced in 1979, had 29,000 transistors, the i486 chip, introduced in 1989, had 1.2 million transistors, and by 2012, its most powerful PC chips contained 1.48 billion transistors. By 2012 Intel was working with such small sub micro geometries that more than 100 million transistors could fit onto the head of a pin! Compared to its original 4004 chip introduced in 2012, the chips Intel was producing in 2012 ran 4,000 times as fast and each transistor used 5,000 times less energy, while the price per transistor had dropped by a factor of 50,000. Driving forward chip design and production requires very heavy R&D spending. By 2012, Intel was spending over $10 billion a year on R&D, or 19% of sales. This was split between spending on chip design, and spending on improving manufacturing processes.

Manufacturing Processes

Designing and manufacturing these devices requires constantly pushing against the limits of physics and technology. Microprocessors are built in layers on a silicon wafer through various processes using chemicals, gas and light. It is an extremely demanding process involving more than 300 steps and, on modern chips, 20 layers are connected with micro circuitry to form a complex three-dimensional structure. Intel is pushing the frontiers of sub micron geometry. The company is currently is producing transistors that measure just 22 nanometers, whereas most other semiconductor manufacturers are still making 45 nm or 32 nm chips (a nanometer is *one billionth* of a meter). Intel newest factory in Arizona, designed to come on line in 2014, will push this frontier still further making chips that have just 14 nm geometry. To carve features this small on a silicon chip, Intel uses a technique known as extreme ultra violet lithography. This is a way of printing circuit patterns onto silicon chips that goes beyond lasers and lenses, and utilizes xenon gas and microscopic reflectors. If it sounds incredibly complex and esoteric, this is because it is at the leading edge of what is scientifically possible. Indeed, each new generation of Intel chips relies upon pushing processes beyond what was attainable just a few years earlier.

So complex is the manufacturing process, that the high tech fabrication plants, or *foundries*, required to make microprocessors cost up to $5 billion each. By 2012 Intel had 16 of these plants around the world. Too equip its plants, Intel works very closely with equipment vendors. Due to its scale, Intel enjoys considerable leverage over equipment suppliers. In some cases, Intel will design a new machine itself, and then have equipment vendors manufacture it. In others, Intel works closely with the vendors on the design of a piece of equipment. As a result, Intel itself holds hundreds of patents relating to the processes for manufacturing semiconductors. Whenever equipment is developed specifically for Intel's requirements, vendors are generally prohibited from selling that equipment to other companies, such as AMD, for a given period.

When installing new equipment, the goal is to gain manufacturing efficiencies through increased yields, or other process improvements. For example, in the 2000s Intel switched from using 200 mm to 300 mm wafers in its manufacturing processes. The larger wafers allowed Intel to put more microprocessors on each, increasing throughput and significantly lowering costs. Intel is currently working to develop the commercialization of 450 mm wafers and is forecasting that it will start to make microprocessors on 450mm wafers by 2016/2017. If it can achieve this, it will be the first in the world to do so. This may give Intel an advantage in manufacturing efficiencies that will be very hard for other chipmakers to match.

To boost yields, raising the percentage of processors that come of the line operating perfectly, Intel uses sophisticated statistical process control procedures. Since even a microscopic piece of dust can contaminate a chip, the specifications that Intel works to are extremely demanding and tight. Over time, Intel has turned yield improvement into a precise science. With each succeeding generation of microprocessor geometry, the company seems able to achieve a steeper learning curve. By constantly pushing out the envelop with regard to manufacturing technology, product design, and yields, Intel has reportedly been able to reduce its unit manufacturing costs for a processor by as much as 25–30% a year.

Typically, Intel will refine new manufacturing processes in one factory, perfecting yields and reducing costs, and then transfer those processes to other facilities. To do this, it relies upon a methodology known as "Copy Exactly!" Under this methodology, engineers spend up to four years perfecting a new manufacturing technique in one of Intel's development factories in Hillsboro Oregon. Once they are satisfied with the results, they work to meticulously import every last detail to other factories around the world. Engineers strive to duplicate even the subtlest of manufacturing variables, from the color of a worker's gloves to the type of fluorescent lights in the building. Employees from around the world spend more than a year at the development factory, learning their small piece of the new recipe so they can bring it back to their home factory. The idea is to capture the infinite number of intangibles that have allowed a process to succeed in plants that have already brought it online. According to one Intel manager: "It's not just there's a specification or a recipe or a program you put into a machine. It also is what the human being does and how they interact with the machine."[7]

The extremes to which Intel engineers go to control the precise conditions in its dozen or so factories has become legendary. A few years ago Intel engineers were trying to figure out why one plant in Arizona wasn't hitting the benchmarks achieved at another in Oregon, where the processes were first developed. Then it hit them: Arizona's desert air was so much drier than the air in Portland, and the engineers in Arizona were skipping several steps taken in Oregon to dehumidify. Intel scientists theorized that the dehumidifying, besides removing water, also eliminated impurities such as ammonia. So engineers began adding water vapor to the air in the Arizona foundry, essentially making Portland air, and then subjected it to the same dehumidifiers used in Oregon. It worked! According to one engineer, this "shows the level of things you've got to worry about when you try to make something as complex as the chips we make."[8]

Intellectual Property

From the i386 chip onwards, Grove was determined to ensure that Intel was the only supplier in the world of its architecture. AMD, however, believed that under the terms of the 1982 technology sharing agreement between the two companies, it had rights to Intel's designs. Intel simply refused to hand over technical specifications for the i386 to AMD, sparking off a lengthy court battle between the two that persisted until 1995. In the end, the two chipmakers agreed to drop all pending lawsuits against each other, settled existing lawsuits, and signed a cross-licensing agreement. Irrespective of the final settlement, AMD had spent $40 million a year on legal fees alone. Senior management attention had been diverted by the ongoing legal battle. AMD had been slow to develop its own version of the i386, waiting instead to get specifications from Intel, which Intel only shared after ordered to in a 1990 ruling.

Intel Inside

For years, Intel had viewed its customers as original equipment manufacturers, focusing its marketing efforts on engineers within those companies. But the nature of the end market was changing. By the early 1990s increasingly sophisticated customers were making their own purchasing decisions, often in computer superstores, or buying direct from companies like Dell and Gateway. Consumers now had influence on the process, and could exercise choice over not just the machine, but also the components that went into it, including the microprocessor.

In 1991, Intel started to market directly to consumers with its *Intel Inside* campaign, effectively telling them that a computer with an Intel chip inside would guarantee advanced technology and compatibility with prior software. Supported by slick advertisements, the campaign was a stunning success. Within a year, Intel was listed as the third most valuable brand name on the planet. In 1993 Grove was able to claim that the number of consumers who preferred a PC with an Intel microprocessor had risen from 60 to 80%. By 1994, some 1,200 computer companies had signed on to the campaign, adhering "Intel Inside" logos on their machines, or including the logo on their product ads.

Complicating matters, one aspect of the long running legal battle between Intel and AMD was a trademark dispute. Intel had claimed that "386" referred to its trademark, and competitors like AMD could not use it. However, in 1991 a court had ruled that the name "386" was so widely used that it had become generic. The ruling infuriated Grove, who believed that clone makers would now be able to piggyback on Intel's marketing campaigns for the 386 and 486. He then made the suggestion that the next chip, which was to have been known as the i586, be given another name that could be trademarked, and the Pentium was born.

Forward Vertical Integration and Customers

Intel vertically integrated forward into the production of PCs in the mid 1980s, selling "boxes" without a screen, keyboard, or brand logo to well known computer companies who put there own brand on them and resold them. The move led to complaints from several of Intel's customers, who felt that Intel was indirectly competing against them in the end market and lowering barriers to entry into the PC industry. After push back, in the early 1990s Intel exited this business. However, the company continued to make motherboards, which are large printed circuit boards that hold the microprocessors, other critical chips, slots for connecting memory and graphics cards, and so on.

Intel's move into motherboards assured more rapid diffusion of each new generation of chips by making it much easier for PC companies to incorporate those chips into their machines. The move infuriated PCs manufacturers such as Compaq and IBM who generally made their own motherboards. Compaq had been able to gain a competitive advantage by bring PCs containing the latest generation Intel chips to market early. Compaq responded by trying to reduce their dependence on Intel. They used for chips from AMD and initially refused to participate in the Intel inside branding scheme. However, by the mid 1990s Intel's position was so strong that this had only marginal impact on the company.

Intel continued to make motherboards through the 2000s, even though profit margins were lower than on sales of stand-alone microprocessors. By 2007 some 24% of Intel's revenues came from the sale of motherboards. At this point, large branded OEMs with a global reach (HP, Dell, Lenovo, Acer, Toshiba and Apple), accounted for about 50–53% of global PC sales, with the remainder being captured by a long tail of smaller local brands. As of 2012, some 18% of Intel's total sales (stand alone chips and motherboards) went to Hewlett

Packard. Dell Computer accounted for another 14% and Lenovo for 14%.

The Microsoft Connection

Throughout the 1980s and much of the 1990s, the relationship between Intel and Microsoft, was an uneasy one. When Microsoft introduced Windows 3.0 in 1990, its first operating system with a graphical user interface, it boosted demand for new PCs to run graphics heavy programs. The same happened when Windows 95 was introduced five years later. In both cases, Intel was a beneficiary of the resulting upgrade cycle. Intel clearly needed Microsoft, but that did not mean that they respected the company. Intel was frustrated that Microsoft did not seem particularly interested in optimizing their software to run on Intel's chips. Microsoft's engineers seemed more concerned with adding features to their products, than in streamlining code so that it took advantage of the full capabilities of Intel's microprocessors.

Microsoft, one the other hand, was interested in making its Windows operating system as ubiquitous as possible, and that logically implied making a version of Windows that would run on other microprocessors, such as the new generation of RISC chips. During the 1990s Microsoft was eyeing users of powerful computer workstations, many of which used RISC chips. This was a potential nightmare for Intel, and it became all to real when Microsoft announced the development of Windows NT, a high end version of Windows that would run on both Intel and RISC microprocessors, including the PowerPC. What stopped the nightmare from occurring was the development of the Pentium Pro, which was so fast and efficient that it effectively eclipsed rivals who used RISC architecture.

Reflecting these underlying tensions, relationships between Andy Grove and Microsoft's Bill Gates were often rocky, and there were reports of meetings dissolving into shouting matches. This started to change in the mid 1990s. It may have been that after the failure of the RISC challenge to Intel, the two companies, and their respective leaders recognized their interdependence and decided that cooperation was better than conflict. Beginning in 1996, quarterly meetings were held between Grove and Gates, aimed at coordinating strategy and resolving differences.

In 2012 new cracks began to appear in the symbiotic relationship between Microsoft and Intel when Microsoft introduced a version of its Windows 8 operating system that would run on ARM processors. For Microsoft, this was a logical move given its strategy of having Windows 8 run on all devices, including tablets and smartphones where the low power consumption offered by ARM processors was highly valued. Microsoft reportedly made the decision to produce an ARM version of Windows 8 because Intel's atom processor consumed too much power to make it a compelling choice in tablets. The move opened the door for PC manufacturers to start building machines that ran on none Intel chips.

THE BARRETT ERA

In 1998 Craig Barrett succeeded Andy Grove as CEO. A former Stanford engineering professor who had become chief operating officer of Intel in 1993, Barrett's tenure as CEO was market by an aggressive push into new markets. By the 1990s the Internet was starting to take center place in computing, and Barrett saw opportunities in extending Intel's reach into chips to drive computer networking gear and wireless handsets. Moreover, Barrett was concerned that without product diversification, Intel would not be able to maintain its growth rate given the maturation of the PC market in many developed nations. In his first three years as CEO Intel spent some $12 billion on acquisitions and internal new ventures designed to strength the company's position in these emerging areas.

Barrett's push into these areas failed to yield any quick returns. By 2004 Intel only had 6% of the market for chips used in networking gear, and 7% of the market for processing chips within wireless phones. Part of the problem; Intel ran into stiff competition from embedded competitors. In the market for wireless phone chips, for example, Intel was competing against the likes of Texas Instruments and Qualcomm, both of whom had a strong market and technological position.

Moreover, Barrett's tenure was marred by some embarrassing product delays, capacity constraints that drove some customers to AMD, and product recalls. To make matters worse, in the early 2000s AMD seized the lead in chip design for the first time, and for two years AMD could boast that it was technological leader in the industry until Intel recaptured the lead with newer chips. Complicating matters, the PC industry went through a sharp contraction in 2001 that led to slumping sales and profits for Intel. While the industry recovered in 2002, growth rates since 2002 have been lower than in the 1990.

Some observers have blamed the problems of the Barrett era on management issues at Intel. The company, they say, had become too large, too bureaucratic, and was no longer the egalitarian entity of its early years. The "constructive confrontation" of the Grove years, which had kept managers on their toes, had been replaced by an autocratic culture dominated by people who got promoted for managing upwards. A management vacuum following Grove's departure led to a lack of accountability and control. To quote one critic: "In the Grove era, each leader who spearheaded an unsuccessful attempt left the company after the project failed. However, throughout the Barrett era each figure head has remained at Intel after the project failed".[9]

PAUL OTELLINI'S PLATFORM STRATEGY

In 2005 Barrett became chairman. Paul Otellini replaced him as CEO. Another long time Intel employee, Otellini was the first Intel CEO to not have an engineering background (Otellini was an MBA with a career in finance and marketing). As head of company wide sales and marketing, Otellini gained prominence at Intel during the late 1990s by pushing the company to adopt a more aggressive approach to market segmentation. By the late 1990s prices for low end PCs were falling to under $1,000, and in this commodity market OEMs were casting around for cheaper microprocessor and motherboard options. Ontellini came up with the idea of reserving the Pentium brand for higher end chips, and creating a new brand, *Celeron*, for lower performance chips aimed at low cost PCs.

In the early 2000s, Otellini pushed for the creation of the Centrino chip platform for lap top computers. While Intel engineers were focused on designing faster more powerful processors, Otellini argued that lap top users cared more about heat generation, battery life, and wireless capabilities. The Centrino platform was designed for them. It combined an Intel microprocessor with a WiFi chip (for wireless networking), and associated software. Personal computer manufacturers were initially skeptical about the value of the Centrino platform. For a while they continued to buy an Intel microprocessor while purchasing WiFi chips from other companies. But when performance tests showed that the Centrino platform worked well, most manufacturers shifted to purchasing this platform for their laptops and Centrino quickly became a recognizable brand.

Introduced in 2003, the Centrino was a huge hit, and helped to pull Intel out of its sales slump. Indeed, by the late 2000s Intel was dominating the market for lap top chips with its chipset offerings. Upon succeeding Barrett, Otellini called for the Centrino strategy to be applied to other areas of the computer industry. He wanted Intel to design separate "platforms" for corporate computers, home computers and lap top computers. Each platform was to combine several chips, and focus on providing utility to a specific customer set. The platform for corporate computers was to package a microprocessor with chips and software that enhance the security of computers, keeping them virus free, and allow for the remote management and servicing of computers (which could bring large cost savings to corporations). The platform for home computers was to combine a microprocessor with chips and software for a wireless base station (for home networking), chips for showing digital movies, and chips for three dimensional graphics processing (for computer games).

The goal was to enable Intel to capture more of the value going into every computer sold and that should increase the company's profitability and profit growth. To implement this plan, Otellini announced a sweeping reorganization of Intel, creating separate market focused divisions for mobile computing (lap tops), corporate computing, home computing, and health care computing (which Intel regarded as a promising growth market with its own unique set of customer requirements). Each division has its own engineering, software and marketing personnel, and is charged with developing a platform for its target market.

To further the strategy of capturing more value going into every computer sold, Intel moved into the graphics chip business, integrating graphics capabilities into its chipsets. Although Intel gained some share at the low end, ATI and Nvidia currently dominate the high-end graphics chip business. The most important and demanding applications for graphics chips are computer games. In 2006, AMD purchased ATI for $5.4 billion, signaling its intention to bundle both microprocessors and graphics chips together.

In mid 2008 Intel introduced a new line of low power consumption chips called Atom that were aimed at mobile internet devices (MIDs)—which was then defined as devices between a smart phone and a conventional laptop and included net-books (very small laptops meant primarily for web surfing). At the time the Atom chip was introduced, Apple had yet to revolutionized the computer market with the introduction of the iPad, although the iPhone had been introduced a year earlier. Unfortunately for Intel, smart phone and tablet makers, including Apple, quickly gravitated to low power consumptions chips based upon technology pioneered by the British company ARM Holdings Plc. The main advantage of ARM technology was that it generated far more computing power per watt than alternative designs, which implied extended battery life, a key requirement from consumers. ARM does not manufacture chips itself. Rather, it licenses its technology to other companies, including Apple, Samsung, NVIDIA and Qualcomm, who incorporated it in their chip designs. They then get the chips made by contract manufacturers. By 2012, ARM chips had become the de facto standard for mobile devices such as smart phones and tablets, leaving Intel at the fringe of the market.

INTEL IN 2013

Paul Otellini retired in May 2013. His legacy was a mixed one. On the positive side, he had helped Intel to reassert itself against a resurgent AMD and cemented the company's dominance in the PC market. The company's revenues grew from $39 billion to $54 billion, earnings per share increased from $1.40 to $2.39, and Otellini left Intel with a commanding market share lead in its core business. Moreover, its manufacturing capabilities remained unmatched in the industry. On the other hand, Intel had largely missed the move towards mobile computing, despite the introduction of the Atom chip, and the company was struggling to gain share against ARM chips.

More worrying still, PC sales were now in decline as demand switched towards tablets. That being said, no one expects the PC to disappear. Indeed, there is a belief that sooner or later the need to replace aging PC inventory will lead to a robust replacement cycle. There was some hope that the introduction of Windows 8 in 2012 might stimulate replacement demand, but many consumers were put off by the new tile based interface Microsoft utilized on Windows 8, and replacement demand remains muted for the time being.

That being said, there is a silver lining in the rapid switch towards mobile computing: Increasingly, these devices are using high-speed wireless links to store data on "the cloud" and access applications that resided on "the cloud". At the heart of the cloud are very large server farms containing hundreds of thousands of PC servers that are networked together. Most of these servers, as it happens, are based on PC architecture and run on Intel microprocessors. Thus the growth of mobile devices that are connected to the Internet through the cloud could result in more server farms and more demand for Intel microprocessors going forward. Nevertheless, for the time being Intel is clearly fighting headwinds in its microprocessor business.

Otellini's successor as CEO is Brian Krzanich, the former COO. A long time Intel employee, Krzanich made his mark in the company as head of the manufacturing organization. His elevation to the CEO position probably speaks volumes about the importance Intel attaches to the manufacturing aspect of its business. A key task for Krzanich is to make sure that the company remains relevant in the post PC era.

Intel is not sitting back and letting ARM chips dominate the mobile device market. It is introducing a new generation of its Atom chips that appear to be far more competitive with ARM chips, and deliver similar performance per watt. These are 22 nm chips and will be manufactured using the latest technology. If the new generation of Atom chips are competitive, it is possible that Microsoft will again focus just on writing Windows to run on Intel architecture, since producing two versions of Windows is a costly exercise. This could provide upside for Intel, particularly if Windows 8 and its successors gain traction in the tablet and smart phone markets—although to date that has yet to happen. Even if the Atom chip is successful, however, the economic impact for Intel might well be muted by the lower average selling price of chips for mobile devices, as opposed to PCs.

Another aspect of Intel's current strategy is to defend the laptop market from encroachment by ARM chips. In 2013 Intel introduced its Haswell chips that can run PC software but have longer battery life. Reportedly, laptops running on Haswell chips have a battery life of up to 10 hours, which represents a 50% improvement over prior generation chips and comparable with the battery life for a tablet.

Although Krzanich seems to be following the script laid out by Otellini, it is clear that he faces significant challenges going forward. The task for Intel is to remain relevant in the post PC era, to hold the rise of ARM chips in check, to continue to dominate its base, to revitalize, if possible, its long-term symbiotic relationship with Microsoft (a company that is itself facing significant challenges), and to gain meaningful traction in the rapidly growing mobile device market where Intel so far has been little more than a bystander.

REFERENCES

1. Anonymous, "Intel's Right Hand Turn", *The Economist*, May 14th, 2005, page 67.
2. Anonymous, "The empire strikes back", *The Economist*, December 2nd, 2006, page 69.
3. Anonymous. "Battlechips". The Economist, June 7th, 2008, pp 75–76.
4. B. Colman and L. Shrine, *Losing Faith*, Colman and Shrine, 2006.
5. R.P. Colwell, *The Pentium Chronicles*, John Wiley, New Jersey, 2006.
6. C. Edwards, "Getting Intel back on the Inside Track," *Business Week*, November 29th, 2004, page 39.
7. C. Edwards, "Shaking up Intel's Insides", *Business Week*, January 31st, 2005, page 35.
8. P. Frieberger and M. Swaine, *Fire in the Valley*, McGraw Hill. New York, 2000.
9. A. Hesseldahl, "AMD vs Intel: The challengers new plan", *Business Week Online*, July 11th, 2008, page 1.
10. A. Hesseldahl, "AMD wins another round against Intel", Business Week Online, June 9th, 2008, page 12.
11. Intel Corporation, Form 10K 2007.
12. T. Jackson, *Inside Intel*, Penguin Books, New York, 1997.
13. Lashinsky, "Is this the right man for Intel?". *Fortune*, April 18th, 2005, pp 110–120.
14. R. Parloff, "Intel's worse nightmare", *Fortune*, August 21st, 2006, pp 60–70.
15. B. Snyder Bulik, "Intel's New Strategy Demands a New Partner", *Advertising Age*, March 14th, 2005, pp 4–5.
16. Intel Corp. 10K Statement, 2012
17. Intel Corp: Assessing Intel's Atom Tablet Opportunity, *Morgan Stanley*, May 15th, 2013.
18. Vivek Arya, "Haswell: Mobility of a tablet, power of a PC", *Bank of America Merrill Lynch*, May 30th, 2013.
19. Anonymous, "Chip of the old Block: Intel v ARM", *The Economist*, May 2nd, 2013.

NOTES

1. T. Samson, "IDC: PC shipments worst than predicted. tablet shipments get better to exceed PC shipments by 2015", *InfoWorld*, May 28th, 2013.
2. D. Traviosm, "ARM Holdings and Qualcomm: The Winners in Mobile", *Forbes*, February 28th, 2013.
3. Tim Jackson, Inside Intel, Penguin Books, New York, 1997, page 18.
4. R.S. Redlow, "The Education of Andy Grove", Fortune, December 12th, 2005, page 116.
5. Tim Jackson, Inside Intel, Penguin Books, New York, 1997, page 33.
6. Tim Jackson, Inside Intel, Penguin Books, New York, 1997, page 206.
7. Anonymous, "When Intel says 'Copy Exactly', it means it", *Chinadaily.com*, May 30th, 2006.
8. Anonymous, "When Intel says 'Copy Exactly', it means it", *Chinadaily.com*, May 30th, 2006.
9. B. Coleman and L. Shrine, Losing Faith: How the Grove Survivors led the Decline of Intel's Corporate Culture (Logan and Shrine, 2006), page 117.