MY LIFE IN SILICON VALLEY:
BRINGING NEW PRODUCTS INTO THE WORLD

Federico Faggin

Growing Up

Ever since I can remember I have been fascinated by machines. When I was five years old and growing up in Italy, I remember the day a huge steamroller-actually running on steam—came to town. I was so excited that I wanted to leave with the man driving the machine. The sight of a car, motorcycle, steam locomotive or airplane, and especially airplanes, would never fail to arouse my curiosity. I wanted to be a pilot. When I was eleven I saw for the first time a rubber-powered model plane fly. I was stunned, I couldn’t believe that one could build flying toys. I determined that I was going to build myself a model plane. Now I could fly, even if vicariously, without having to wait to grow up.

I started building my first model plane based only on my observation of that flying model, using generic material found around the house. Of course that plane couldn’t possibly fly as it violated all aerodynamic laws known to man. I was lucky, though, because while I was trying to make my plane fly in the same field where I saw the original model, its owner happened to come by in his bike. He was probably a young man in his twenties, but he looked like an old man to me. Curious, he stopped to take a good look at my plane. When I showed him what I had done he started laughing. I’m sure he had never seen such an unusual plane before. However, he sympathized with me and taught me what I should do, where to buy the proper materials and so on. I bought a book on how to build model planes, the first book I purchased with my own money, and I studied it avidly. My next plane made a straight flight of thirty meters before crashing—I mercifully called that my maiden flight. My third plane, however, flew very well. I had conquered the air. I was twelve.

That was the beginning of a passion that was an important part of my growing years until I graduated from high school. I didn’t realize until later how formative that hobby was. I learned the complete cycle of creation, starting with an original idea of a plane, visualizing its critical features, making a construction plan, procuring the materials, building it, testing it and enjoying the result. In a sense, designing and building model planes was more than a hobby for me; it was all that interested me, back then, and all my spare time was dedicated to it. When I was twelve, I severely injured my left eye. Now I could no longer become a pilot; undaunted, I decided that I would design and build real planes instead.

I decided to go to a technical high school, the Institute Rossi in Vicenza. I wanted a degree in aeronautical engineering. Unfortunately, the curriculum in aeronautical engineering was phased out one year after I started at the Institute, and I had to choose a different specialty. Radio engineering, for which I had genuine interest, was my next best choice. I must confess, though, that my desire to learn how to build a radio controlled model plane added considerable weight to my choice.

I did very well at school, graduating five years later the best in the Institute by a wide margin. During the high-school years I became also interested in computers and in physics and I enjoyed reading many books on those subjects.

My Early Work Experiences

I had not yet turned nineteen when I was hired by Olivetti to work as an assistant engineer at their Electronic Laboratories in Borgolombardo, near Milan. In 1960, Olivetti was the only company in Italy designing and building electronic computers. This was my first job. I was very excited to learn new things, earn some money and be independent. After a training class of a couple of months, I was assigned to a small group that was designing and building a small electronic computer. The arithmetic unit had already been designed and partly built, but the rest of it—the instruction set, the instruction control, the magnetic core memory and the input-output subsystem—still remained to be done.

Within six months I became the leader of the project and I was assigned four assistants to help me with the construction of the computer. I was very proud because all the people working for me were older than me—I was nineteen. I completed the design of that small computer and I had it running by the end of 1961. That computer used almost exclusively NOR gates made with diodes and germanium transistors. Faster silicon transistors were still too expensive, and integrated circuits were not commercially available yet. With 4,000 words of magnetic core memory, this computer was the size of a cabinet, 1.80 meter tall by 0.5 meter wide, and included about one hundred printed circuit boards (PCB).

Little did I know then that only nine years later, married and in a foreign country, I would be in charge of developing the world's first microprocessor. With it, a computer of roughly the same capability—in fact, twice as fast as my earlier machine—could be built, using only a single PCB, an improvement of more than a factor of one hundred!

During the summer of 1961, the idea of going back to school started to appeal to me. I became interested in Solid State Physics—I wanted to really understand how transistors worked—and I wanted a stronger theoretical knowledge in mathematics and physics. I left Olivetti at the end of 1961 and became a physics student at the University of Padua, close to my hometown of Vicenza.

I enjoyed my university years immensely; I liked to learn new things and because of it I was an outstanding student. In many subjects, electronics in particular, I did much more work than what was strictly required to pass the exam. I have fond memories of afternoons spent in the library studying; and of the joy I felt when I could understand a new concept in physics, or figure out how to solve a difficult problem.

By the end of 1965, with a doctorate in physics, summa cum laude, I was ready again to go back to work. I was asked to stay at the University of Padua to teach and continue the research I started with my thesis, but after one semester I decided to leave because the pace of the university environment was too slow for me. I was hired by one of my old bosses at Olivetti who had just started a small company near Milano, called CERES. CERES was developing thin-film circuits and was also the Italian representative of GMe (General Micro-electronics), a spin-off of Fairchild Semiconductor, and the first company to make commercial MOS integrated circuits. I was hired to head up the business relationship with GMe and I was to be sent to Santa Clara, California for a one-week course in MOS integrated circuits.

Meeting Silicon Valley

It was my first transatlantic trip, and California loomed as the Mecca not only of the movie industry but also of electronics. Hewlett-Packard and Fairchild Semiconductor, for example, were already famous in 1966, and I was thrilled to be there. The man giving the course was Jim Imai, a GMe engineer, who three years later—how full of coincidences is life—as a consultant to Busicom, would recommend Intel as the best company to build their calculator chip set—the very project that led to the 4004.

I learned about MOS technology, a new subject in those days, and most importantly I had a peek at how people worked in the industry around the San Francisco Bay Area. I was immediately captivated: here the most important people were engineers, scientists and entrepreneurs, not politicians or pedantic professors of liberal arts like in my country. This was the place for me! I didn't know how I would come back, but I knew I would.

Back in Italy, we ordered a small number of GMe's 100-bit MOS dynamic shift registers—the most sophisticated product built in those days—but we never got delivery of them. Clearly GMe couldn't build them. Therefore my job as a Representative of this company didn't make sense, and I was reassigned to other projects. I left CERES nine months later to work for SGS-Fairchild (now SGS-Thomson) in their new R&D Laboratories. Because of my brief exposure to MOS technology, I was chosen to lead the development of the process technology for the fabrication of MOS IC's. Within six months, starting from scratch and with two assistants, I developed SGS's first MOS fabrication process and I designed two commercial circuits. I was promoted Group Leader for MOS.

One day in December, 1967, I was summoned by the R&D Manager. Speaking to me in English to see if I understood, he asked me if I was willing to go to Fairchild Semiconductor in Palo Alto, California, for six months, part of an engineering exchange program. I immediately said yes! Even if I had recently married Elvia and we had just moved into a new apartment. The prospect of working in the famous R&D Laboratory of Fairchild, the most advanced semiconductor company in the world, was truly exciting. My chance to go back to California had come.

I arrived in California in February, 1968, leaving the fog and cold weather of Northern Italy behind, to find sunny and warm weather and blooming fruit trees all over the Valley. Elvia and I were in heaven. My first project at Fairchild was to develop a new process technology to make self-aligned MOS transistors using gates made of poly-crystalline silicon instead of aluminum, as was currently done. The use of silicon as an MOS gate electrode had been demonstrated earlier at Bell Laboratories, but the method was not suitable to build integrated circuits. In only five months I developed the new process technology-called Silicon Gate Technology(SGT) —and designed and built an integrated circuit using it. This circuit was an eight-channel analog multiplexer with decoding logic that Fairchild could not reliably build with aluminum gate technology. By the end of 1968 this circuit, called 3708, became the world's first commercial IC to use gates made with polysilicon.

With SGT it became possible to make IC's up to five times faster than with aluminum gate, using half the silicon area, with much lower leakage current and with superior long-term reliability—the lack of which was the Achilles heel of MOS technology. It was this technology that made possible, a few years later, high-performance semiconductor memories, electrically programmable, non-volatile memory, and the world's first microprocessor—the 4004. SGT became the core technology for large scale integration (LSI), and today more than 90% of all IC's sold worldwide are built using silicon gates, processes that evolved from that original work.

During June 1968, Fairchild Semiconductor decided to sell their equity interest in SGS-Fairchild and I was asked to stay at Fairchild, which I gladly accepted. My first official day of work, as an employee of Fairchild, was July 1st, 1968. On that day the rumor that Bob Noyce and Gordon Moore had left Fairchild to start a new company, was spreading like wild fire in the Laboratory. Soon came the official confirmation that, indeed, the two executives had left. Few days later Andy Grove joined the new start-up, and one week later Les Vadasz, my boss, also left. Within a couple of months two dozen more people left Fairchild to join the new company. This is how Intel started.

When my boss left, I suspected immediately that Intel would use the SGT. Suspicion became certainty when Intel hired the technician specialized in making the deposition of polysilicon for me. I knew then that I was in a race and I wanted at least to have the first product in the market with SGT. This happened during the next six months as I completed the project, and the 3708 became commercially available. Unfortunately, with the departure of so many good people, Fairchild was rapidly loosing momentum and I soon became disenchanted. It was time for me to leave as well.

The 4004

In early 1970 I had a meeting with my old boss, Les Vadasz, to explore job opportunities at Intel. I wanted to move away from process development into IC design, and I wanted a challenging project. Les asked me if I understood computers and logic design, I told him about my computer work at Olivetti, and I was soon hired to head up a new project. Vadasz was very mysterious about this project, only telling me that it consisted in designing four chips for a Japanese customer, and that some of the chips involved random logic design.

In April, 1970 I started my new job at Intel and I found out that my project was called internally the “Busicom Project.” Busicom was a Japanese manufacturer of calculators that had contracted Intel to develop a set of four chips to make a family of high-end calculators. Stan Mazor, an application engineer working for Ted Hoff, the Application Engineering Manager, told me the story: Busicom had visited Intel in the middle of 1969 asking Intel to turn into silicon the logic design of their calculators, which were built using small and medium scale integration (SSI and MSI) TTL components.

Busicom engineers had partitioned their design into seven complicated chips and when Hoff was asked by the marketing department to give his opinion about the project, he marveled at the complexity of the design. The Busicom design was essentially a ROM-based, macro-instruction, programmable, decimal computer based on shift-register memory, a design typical for the advanced calculators of the time. Hoff felt that there had to be an easier way to get the job done. By using Intel’s RAM memory technology, he argued, it might be possible to design a small computer, far more flexible and general, and perhaps simpler, than Busicom’s design.

Busicom’s engineers initially appeared unreceptive to Hoff’s ideas, but Hoff with encouragement by Bob Noyce, then Intel’s president, continued to explore options. Eventually a family of four chips was defined and Busicom management agreed, in October 1969, to proceed with Intel’s proposal and abandon their design. Mazor showed me the specifications for the chip set and the bits and pieces of circuit design that were utilized to make the feasibility study. He also told me that the next day Masatoshi Shima, a Busicom engineer who was part of the team that visited Intel back in 1969, would come from Japan to check on the progress of the project.

The next day Stan Mazor and I went to the San Francisco airport to greet Shima and meet each other. Shima said he had come to check the work and to please give him all the work to check. Back at the company I gave him the material I received from Mazor the day before, but Shima, after quickly scanning the material, became very angry at me because he had already seen it all back in October. He was furious that no more progress had been done since he had left Intel six months earlier; the schedule was irreparably compromised.

He kept on saying: “I came here to check and there is nothing to check. This is just idea.” He held me responsible for the delay saying: “How could you have made no progress in six months?” I kept on telling him that I had just joined Intel the day before, therefore I could not have made much progress in one day, but Shima was undaunted. It took about one week for Shima to understand that he and I were in the same boat; I wanted to have those chips done as much as he did, and if he helped me we could get the job done sooner. Just one day after starting my new job, I was already six months behind schedule! That was a personal record.

I worked furiously twelve to sixteen hours a day to make up as much time as was humanly possible. First, I resolved the remaining architectural issues of the chip set, and then I laid down the foundation of the design style that I would use for it, discarding much of the earlier work. Finally, I started the logic and circuit design and then the layout of the four chips. I had to develop a new methodology for random-logic design with SGT; it had never been done before.

I called the chip set “the 4000 family.” It consisted of the 4001, a 2-kb ROM with a 4-bit programmable input-output (I/O) port; the 4002, a 320-bit RAM with a 4-bit output port; the 4003, a 10-bit serial-input, parallel-output, static shift register to be used as an I/O expander; and the 4004, a 4-bit CPU.

The 4001 was the first chip designed and laid out. The first fabrication of the 4001 (called a *run*) came out in October 1970 and the circuit worked perfectly, considerably lowering my anxiety level. In November, the 4003 came out and worked as expected. In the same month, the 4002 also came out. It had a minor error which was soon diagnosed and corrected. Finally, the 4004 arrived a few days before the end of 1970. My hands were trembling as I loaded the two-inch wafer in the probe station; the moment of truth had arrived. Without the 4004 there would be no Busicom calculator, and the 4004 was a much more complex circuit than the other three.

As I probed the first circuit, I found no life. I went to the next circuit: same, no life. I started sweating. I was glad no one was in the lab at that time to see my defeat. “How could I have possibly screwed up so badly?” I was thinking. Finally I examined the wafer under a microscope and I soon found out that one of the masking layers had been erroneously omitted in the wafer processing. I started breathing again, my reputation was intact, I had done no wrong. It took half hour to resolve the mystery; the longest half hour of my life.

Three weeks later, in January 1971, a new run came. I received it in the evening hours and I set myself to spend the night probing it. I was praying for it to work well enough that I could find all the bugs so the next run could yield shippable devices. My excitement grew as I progressively found the various areas of the circuit working. I was alone in the lab, having only the hum of the instruments to keep me company, and I was intensely happy. By approximately 3:00 a.m. most areas of the chip had been successfully verified and I went home in a state of mixed exhaustion and elation. Elvia woke up and asked me: "Is it working?" "It works!" I said. That was the night that the world's first microprocessor was born.

By March, 1971 all the chips were ready, including four ROM chips containing the software for the calculator—we call it firmware nowadays—that Shima had developed back in Japan. Soon the good news came from Japan: "Everything works!" Now we could go into high-volume production. It took eleven months to go from idea to a fully working product, a major feat!

Now that the microprocessor was a reality, I thought the chip could be used for many other applications. However, Intel's management disagreed, believing the 4000 family was good only for calculators. Furthermore the chips had been designed under an exclusive contract; they could be used only by Busicom.

The opportunity to prove that the 4000 family was good for other applications came when the need for a production tester arose. The tester was clearly not a calculator application, so I decided to use the 4004 as the tester's main controller. Using the new electrically programmable ROM (EPROM) chips that had just been invented at Intel by Dov Frohman-Bentchkovsky, I designed the hardware and the firmware for that project, avoiding the mask-programmable 4001's, not appropriate for a one-of-a-kind project.

This work gave me much insight about what could and couldn't be done with the 4000 family. At the successful completion of the tester I had additional ammunition to convincingly lobby for the 4000 family's introduction. I suggested to Bob Noyce that perhaps Intel could trade some price concessions for nonexclusivity. (I had heard from Shima that Busicom was hurting in the market place and needed lower prices to effectively compete.)

During the summer months of 1971, Intel's management obtained nonexclusivity and decided to broadly market the 4000 family. In November 1971, the 4000 family, now known as MCS-4 (for Microcomputer System 4-bit), was officially introduced with an advertisement in *Electronic News* magazine. The main caption read, "Announcing a new era in integrated electronics, a microprogrammable computer on a chip." This is perhaps the only ad I know where hyperbole was justified.

More Products

My next project at Intel was the 1201, the world's first 8-bit microprocessor, which was introduced in April 1972. The history of the 1201 is quite interesting: In December 1969, an officer of Computer Terminal Corporation, a shift-register customer of Intel, based in San Antonio, Texas, visited Intel. CTC wanted Intel to modify one of its existing products—the i3101, a bipolar 64-bit RAM—to make a 4x16 stack memory to be used in their TTL microprocessor, intended for their new generation intelligent terminal, the Datapoint 2200.

Hoff and Mazor studied the request and determined that the CTC processor, of which the stack memory was part, didn't appear much more complicated than the proposed 4004. They concluded that it would be feasible to make a single-chip, 8-bit microprocessor. They drew up a target specification, and CTC contracted with Intel for the development of such a device, a custom product embodying CTC 8-bit microprocessor architecture.

The design work on the 1201 had already started a few weeks before I joined Intel. Hal Feeney was the engineer in charge of the project. I remember my disappointment when I found out, soon after joining Intel, that the 4004 was one of two CPU-on-a-chip in development at Intel. Since I had four chips to design, and the 4004 would be my last one, Feeney—having only one chip to design—surely would be the first to create a microprocessor. But things turned out differently.

For a variety of reasons, the 1201 design was put on hold soon after it started, and resumed in January 1971, under my direction. Feeney was assigned to work for me and together we designed the 1201. Feeney did the detailed design following the same methodology I had created for the 4004.

Architecturally, the 1201 was very similar to the 4004, and its logic and circuit design, after the experience gained with the 4004, proved to be not too difficult.

However, being an 8-bit CPU, it required more transistors than the 4004, about 40% more, therefore the chip was larger and more expensive to manufacture.

The design and layout of the 1201 progressed throughout 1971, and in early 1972 we received our first wafers. There were only a few minor errors, promptly corrected in a new revision of the masks. Feeney in the meantime had been transferred to a new marketing department, under the leadership of Hank Smith, created to properly commercialize the first microprocessors.

The only difficult time with the 1201 came toward the end of the project, in April 1972. I was characterizing the chip when I noticed some intermittent errors at high temperature. I was under a lot of pressure because the 1201, marketed under the name "8008," had just been introduced in the market. Customers wanted it, but Intel couldn't ship it because it didn't work reliably yet. For a few days I couldn't figure out what was happening, but after a feverish week of work I understood the problem and found a solution. It was a tough problem at the crossroads of device physics, circuit design and layout: the charge stored in the gates of the transistors in the register file was leaking away due to substrate injection. I finally found a way to modify the circuit and the layout, permanently fixing the problem. The 8008 could now be delivered to the impatient customers.

Late in the summer of 1971, I went to Europe with Hank Smith to give a series of technical seminars on the 4000 family and the 8008, in anticipation of their announcement. It was my first visit to customers, and it was an important experience. I received much criticism—some of it valid—about the architecture and performance of those early microprocessors. The more computer - oriented the company I visited was, the nastier people's comments were.

When I returned home, I had an idea of how to make a better 8-bit microprocessor than the 8008, incorporating many of the features customers wanted: most important, speed, ease of interfacing, a better interrupt structure and some other new features. I drew up the architecture of a new chip intended to use the new and faster n-channel MOS process Intel had just developed to fabricate the first 4-kb dynamic memory. By early 1972 I started lobbying for the new chip. However, Intel management wanted to see how the market would respond to the MCS-4 and the 8008, before committing more resources. I was frustrated because we were wasting time, Intel was not the only company to see the potential of microprocessors and other companies were in hot pursuit, Texas Instrument, Rockwell and Motorola in particular.

Eventually, in late summer 1972, I was given permission to start the project. I hired Shima from Japan to lead the project, under my supervision. By the time Shima arrived, in November 1972, I had completed the basic architecture and the design feasibility so that detailed work could start. The new microprocessor, called the 8080, was introduced in March 1974. It was an instant success. By early 1974, the market had finally accepted the new direction and microprocessors were utilized in hundreds of different products by Intel’s customers. However, many potential applications could not use microprocessors because the 4004 and the 8008 were not powerful enough. It was the 8080 that blew wide open the microprocessor market; what the 4004 and the 8008 promised but couldn’t deliver, could now be done with the 8080. The microprocessor had come of age; and its value could no longer be debated.

Due to the success of the 4004, 8008 and other projects I led, and with the rapid growth of Intel, my level of responsibility steadily increased. By early 1974 I was promoted to the position of department manager in charge of all MOS IC design, except for dynamic memories. I had 80 people in my department and more than a dozen projects underway simultaneously. Among my accomplishments I remember in particular a custom circuit I designed for Mars Money, Inc. in 1972. It was the first circuit to combine random logic with on-chip EPROM to create the function of an adaptive coin reader. In 1972 I also designed a chip that allowed the 4004 to interface with standard memories, in particular EPROMs, useful for low-volume applications, where the expense and time required for ROM tooling was prohibitive. In 1973 I designed the 4040, a much improved successor to the 4004. And in 1974, with Dick Pashley, I created the first fast, static, n-channel, 5-volt RAM, called the 2102A, by combining depletion load transistors with low-voltage, n-channel MOS technology. We achieved an access times of 80 nsec where before 300 nsec was hard to get.

My nearly five years with Intel were the most productive time of my engineering career. I personally designed or supervised the design of at least two dozen commercial ICs－I don’t remember them all－growing immensely in both my technical and managerial competence. By the end of summer 1974, however, I grew restless. Intel had become a large company and the environment had become a bit stifling to me. Intel was started with the mission to be a memory company, and microprocessors were important only insofar as they helped sell more memories; I didn’t like that.

Moving On

My time at Intel, while extremely productive, had not been completely satisfying. Most things I did I had to push and argue to get done, and I felt unappreciated, since the attention of management was primarily on the memory business. The idea of starting my own company, a company dedicated exclusively to microprocessors, began to surface in my mind. Silicon Valley, even in 1974, was a hotbed of start-up companies. Why not do what Noyce and Moore had done six years before? I approached Ralph Ungermann, one of my managers, and asked him: “How about starting our own company to make microprocessors?”“Let’s do it!”he said immediately. Without any specific plan, having approached no one else, Ralph and I left Intel and founded Zilog.

We rented an office in downtown Los Altos in a shopping mall, and we set out to figure out what we were going to do. The first product that came to my mind was a single-chip microcomputer, I called it the 2001. It was a full computer on a chip, not just the CPU. In other words, I wanted to integrate an 8-bit microprocessor with the ROM, RAM and input-output electronics, to create a single-chip computer for simple, low-cost applications.

This device would have been the first of its type in the market—now these chips are called microcontrollers and represent the bulk of the consumption of devices containing microprocessors. But after further reflection, I decided against the 2001. I couldn’t see how we could possibly get the millions of dollars required to start our own wafer fabrication facility, and without it we could not compete with companies having their own manufacturing capability. We had to purchase wafers from wafer foundries at prices much higher than normal manufacturing cost; a major problem for a price sensitive device like the 2001. I reluctantly abandoned the 2001 idea and turned my attention to a less cost-sensitive product. I still remember a Saturday morning in December 1974, when the basic idea for the Z80 came to me. I said triumphantly to Ralph, “Yes! Let’s make a super 80!” By that I meant a product that was instruction-set compatible with the 8080, but much more powerful than the 8080.

During the following months I developed the basic architecture, the instruction set and the design feasibility of the Z80, using a manufacturing process with depletion loads, similar to the one I co-developed at Intel for fast memories. The Z80 was conceived from the start to be part of a family of peripheral chips working seamlessly together. The other members of the family were a programmable parallel input-output port, a programmable counter timer, a direct memory access controller and a serial input-output controller. By adding standard memory to these chips, a powerful computer could be built.

By April 1975, Ralph and I were able to secure a financing of $500,000 from Exxon Enterprises, the venture capital (VC) arm of Exxon Corporation, then the largest corporation in the world, and one of the few companies active in venture financing in 1975. 1975 was a recession year and most VC firms had been seriously affected by the excesses of the late sixties and early seventies. There was no money to go around. During 1975 the total VC investment in the USA was $10 million, the lowest it has ever been. We considered ourselves lucky to have found some money.

In April 1975, I hired Shima from Intel to lead the Z80 project and during the next few months, Ralph and I put together the team that would develop the Z80 microprocessor, the development system with its software, and would then proceed to design the remaining peripheral chips—eleven people in all. By March 1976, the Z80 was fully functional and the development system ran for the first time with an actual Z80 chip instead of its emulator hardware. We had done all that work in eleven months, spending a little more than $300,000, an incredible feat even in those days.

The Z80 was a announced in May, 1976 with a color ad in *Electronic News* whose caption said: “The battle of the 80’s” and compared the Z80 to my prior chip, the Intel 8080, showing its superior performance.

The Z80 became wildly successful, in fact, even today, twenty-one years after its introduction, it is still produced in high volume. More than one billion Z80 have been produced so far by a number of suppliers scattered throughout the world.

With the Z80 my engineering career ended. As CEO of Zilog, the task of leading a fast-growing company soon took all my energy. From that point on my career took a turn and I became an entrepreneur and CEO, starting and leading high-technology companies in Silicon Valley.

It would take too long to describe the history of Zilog, the many products we created, the good and the difficult times, and the vital role that Zilog played in the early years of the microprocessor. At Zilog I learned the difficult lessons of running a company in a highly competitive business. I did many things right and many things wrong, and, of course, I learned from my mistakes much more than I learned from my successes.

In the late 70s, Exxon Enterprises, which had controlling interest of Zilog, decided to consolidate all its investments in information technology into a group called Information Systems Group, publicly stating its intention to become a dominant force in the information business, and targeting IBM as the company to beat. This is not what I expected. Ralph and I wanted to create a public company, not a division of Exxon Information Systems! Exxon Enterprises had originally agreed with our vision, but now the plan they obviously had had all along began to take shape. This turn of events was a major blow to Zilog and one of the key reasons that induced IBM to choose Intel instead of Zilog microprocessors in their first personal computer. That choice sealed the fate of Zilog as an also ran. At the end of 1980, discouraged after unsuccessfully trying to convince Exxon Enterprises to divest themselves from Zilog, I left the company.

I learned the hard way that technical excellence is only one of the ingredients necessary to succeed in business, and that the best product doesn't necessarily win in the marketplace.

Entrepreneuring

The following seventeen years I co-founded and was the CEO of two other companies, Cygnet Technologies and Synaptics. At Cygnet we developed and introduced in the market in early 1984 an intelligent telephone that, connected to a personal computer, allowed the PC to become a voice and data workstation dedicated to managers, people concerned with communication much more than with computing.

At Synaptics, the company I started in 1986 with Professor Carver Mead of the California Institute of Technology, we set out to develop the technology for learning systems, based on artificial neural networks and analog VLSI. After six years of research, we came to the conclusion that while neural networks were superior in many ways to conventional technology for recognizing complex patterns, there was not enough architectural commonality to warrant the design of complex analog VLSI learning chips. We decided to use our expertise in neural network technology to pursue human to computer interface products, based on giving senses to the computer—touch, hearing and sight. In 1995 Synaptics introduced the touchpad, a new type of pointing device based on determining the position of a finger in a sensing surface—just like the skin of the computer. The touchpad has been very successful, capturing in three years 50% of all notebook computers sold worldwide. The company has grown rapidly and is quite successful.

Synaptics has also developed a system for computer entry of Chinese characters using a pen and a special touchpad. This handwriting recognizer is based on neural network technology and surprisingly allows faster data entry than writing by hand. This important result is achieved by recognizing the characters as they are written. Since, on average, only half the number of strokes forming a Chinese character need to be written before the computer can correctly recognize the character, the user can write twice as fast with the computer than writing by hand.

At the personal level, I have many interests. Understanding how the brain works is one of them, and many years ago I spent four years studying neurobiology to gain some insight into this marvelous creation of nature. Studying the brain I got interested in consciousness, this most mysterious property of our brain. Consciousness is our ability to know that we know, to form pictures and feelings in our body based on our sensory information and our thoughts. These images form in the screen of our mind and are intelligible to some virtual *entity* inside our head. While we seldom think about consciousness, we use it all the time—without it, we wouldn't even be aware of existing, we would behave just like sleep-walking automatons.

But what is consciousness, really? How can it possibly emerge from the operation of a complex machine? If it can, then it will be possible to make a computer which is conscious; maybe even a computer which is more intelligent than we are. If consciousness, on the other hand, exists outside the material world, how do we reconcile this possibility with physics? How does it connect with the material world? How does it exchange information with the objects described by quantum mechanics?

These are clearly questions at the boundary between physics and metaphysics, but then, what do you expect from the son of a philosophy professor?

My life has been good and productive and I am very grateful for it. I feel fortunate to have been able to live in one of the most intellectually stimulating places on Earth—Silicon Valley—and to have had the loving support of my wife Elvia. God willing, I intend to continue to work for the rest of my life to advance the state of knowledge in the area of artificial intelligence. I firmly believe that our search to learn how to make a machine intelligent will naturally lead us to deepen our understanding of our human nature, learning what are the unique properties that make us human, and what is our role in this wondrous universe. I wouldn't be surprised if in due time we will discover that we all possess a deep spiritual dimension connecting us to the purpose of the Cosmos.