# Oral History of Federico Faggin

Interviewed by:
Gardner Hendrie

Recorded: Sept. 22, 2004, Dec. 13, 2004, March 3, 2005
Los Altos, California

CHM Reference number: X2941.2005



[Note: edits made in square brackets were made by Federico Faggin following the oral history]

**Session 1: September 22, 2004**

**Gardner Hendrie:** We have today Federico Faggin, who has graciously agreed to do an oral history for the Computer History Museum's oral history project. We really appreciate that, Federico.

**Federico Faggin:** My pleasure.

**Hendrie**: I think maybe you could tell us a little bit about where you were born and where you grew up. A little bit about the background that may have influenced, in some way, who you are.

**Faggin:** Yeah. O.k.

**Hendrie**: Or who you turned out to be.

**Faggin:** All right. I was born in Italy in Vicenza, which is a town near Venice in northeast Italy, and I was born in 1941 during the early phases of the Second World War. We lived in downtown Vicenza at that time. And when '43 came about, with the fall of the fascism in Italy, and the Americans coming up from Sicily all the way toward the north of Italy, the family moved to the countryside, because it was dangerous. In fact, our home eventually was destroyed by Allied bombing. I lived in the countryside until I was about eight years old, and then we went back to Vicenza and I pretty much grew up in Vicenza until I came to this country, except for a small parenthesis of work in Milan, or around Milan, depending on the case. So I was a boy that was very interested in mechanisms and making things with Meccano and stuff like that. I had a very important experience when I was 11 years old: I was playing with the kids and there was this "old" man, he was probably 20 years old, with a model plane in the field where we were playing. He was winding up the propeller, and then he let this thing go… And it was flying. I was absolutely amazed that you could actually build a toy that flies; I mean, only airplanes fly, not toys. So I was just running after this thing -- I was fascinated -- and I decided that I was going to build model planes. So, that was a passion that I still have to this day, where I build and fly [model planes]. Of course, now they are radio controlled; they've got three-phase motors -- electric motors -- and all that kind of stuff, with lithium-polymer batteries. But that was a very important part of my experience, because I didn't have any money, so I had to visualize a model plane the way I wanted it, and draw it; then make the plans; then build it, buying the basic material -- some sheets of balsa -- but cutting all the strips and doing all the profile of the wing and everything else. And building the plane, and then trying it out, see if it worked. And then enjoy it, right? So it was the complete process, from imagining a product to designing it, building it, and flying it.

**Hendrie**: Did you do that when you first started? Do you remember what your first one? Did you start with a kit and sort of make the rudiments?

**Faggin:** No. No.

**Hendrie**: Or did you just start?

**Faggin:** No. I just started with sticks and just made-up stuff, with the paper my mom used to cut models for dresses. And I made a contraption and, of course, it couldn't possibly -- it violated all the aerodynamic laws in the universe -- so it couldn't possibly work. But I went to this field to make it fly, and of course it didn't fly. And while I was there, the guy who had the model plane was coming by in his bike. So he stopped by and he started laughing at my thing. I was 11 years old. I mean, I was 11 and a half, right. He started laughing at me, and he told me: "you got to buy this stuff from the store; and you got to really do it this way… You don't do it this way, because it can never work." So he explained it to me and that was my introduction. So the second model still didn't fly, but it was better. The third model actually flew and then… I bought a book about how to do it, also -- which actually was the first book that I bought myself. In other words…

**Hendrie**: With your own money.

**Faggin:** With my own money. And it was my book as opposed to being something that my dad would gift me. So it was really something. It was actually the first book that I bought with my money and that I read cover to cover, of course. And one of the things that fascinated me in there was the fact that you could actually build radio controlled model planes. It was the beginning of radio control -- I'm talking about the early 50s now -- using ultra miniature vacuum tubes and, of course, they were with very rudimentary escape mechanisms, and so on. And I was fascinated. So that's how actually I got interested in electronics and radio technology, and so on. So I went into a technical high school, because I wanted to learn how to do that stuff. Also I wanted to learn how to build real planes. There was a specialization in aeronautical engineering…

**Hendrie**: Yeah, what kind?

**Faggin:** Aeronautical engineering. And also there were specializations in radio technology and...

**Hendrie**: So this was a really a pretty good technical high school back then.

**Faggin:** Oh, yeah, absolutely. I mean, a lot of the designers in Italy actually come out of this school, and they were not even engineers. In fact, I designed and built a computer when I was 19, after going to this school, working for Olivetti. So that was my first work experience in an actual project. So it was a school… I mean, math, for example, we did differential equation. I mean, calculus, differential equations. In some cases, I did the…

**Hendrie**: In high school?

**Faggin:** In high school. I did the cable equations for transmission lines, which is a partial differential equation. So, I mean, of course, it was only a special solution and all that. And it was 44 hours a week of lectures and laboratories. So it was very intensive. It was...

**Hendrie**: Sounds like a wonderful school.

**Faggin:** It was a very good school.

**Hendrie**: What did your mother and father do? Just to sort of… I'm trying to understand the environment at home a bit.

  

**Faggin:** Yeah. My father was teaching history -- History of Philosophy and History -- at a classical high school, and also Aesthetics at the University of Padua. And he was an accomplished writer. He wrote about 40-50 books in all subjects from history of art to mysticism -- about Meister Eckhart -- and all that. And he also had a translation of the Enneads of Plotinus from Greek to Italian. A very, very well received translation, and so on and so forth. So he was a very, very accomplished…

**Hendrie**: Very educated.

**Faggin:** Very educated, yeah, yeah. But most of all in the arts and the humanities, and not in science or technology. In fact, I was considered a bit the "mechanic" of the family.

**Hendrie**: And your mother?

**Faggin:** My mother was an elementary school teacher.

**Hendrie**: So they were both teachers in

**Faggin:** Yeah. In different levels, yeah.

**Hendrie**: Now did you have any brothers or sisters?

**Faggin:** I had an older brother, who eventually graduated in humanities and followed the path of my father. He became an expert in Flemish art. And also he wrote about [Flemish art and philology]… and he teaches now at the University of Padua. So I'm the second. And a sister, who became a fashion designer. She's very, very artistic. And my youngest brother, who is ten years younger, that… he's a lawyer. There ought to be a lawyer in the family, so he's the one.

**Hendrie**: Very good. Well, that's very interesting. What are your earliest… I think I know the answer, but I'll ask the question anyway. What were your earliest thoughts as to what you wanted to be when you grow up. Do you remember what were?

**Faggin:** Oh, yeah, well, the earliest, probably, I wanted to be an ice-cream man or something, you know, when I was four. But the earliest [thought] that really felt like what I really wanted to do, was after this experience with the model planes. So within a year or two I wanted to be an aeronautical engineer. So that's why I chose that high school and I wanted to build real planes and I wanted to do it in a hurry, too. I didn't want to have to go to university or all that, right? Yeah. I ended up going to university later on. But at that time I didn't want to. I wanted to just graduate -- it was a five-year program --and then go design planes.

**Hendrie**: So you completed that.

**Faggin:** Yeah, yeah.

**Hendrie**: So what did you do next?

  

**Faggin:** So then, after completing… now, I did not get to become an aeronautical engineering graduate of this school -- the title that you get in the school is called, "perito," which would be probably equivalent in terms of preparation at least to an AA degree here. But anyway, so I got that degree when I was 18, but the degree was in radio technology and electronics, because they had stopped the program in aeronautical engineering. So I went to this school only to find out a year later that I could not do aeronautical engineering, and I had to choose among other specialties. And I picked radio electronics and electronics.

**Hendrie**: Well, it certainly helped you with your model planes.

**Faggin:** That's exactly right. You've got it right. And the reason was that, you know, then… I'm going to build radio-controlled model planes. I was a one-track mind. And of course, as I learned more about it, I got fascinated by that field. And particularly toward the end of the high school years, I became fascinated with computers. So I started learning about computers on my own. And then I ended up working for Olivetti, which had a R&D laboratory in electronics for computers near Milan. And I took a job there as an assistant engineer.

**Hendrie**: Was that your first job?

**Faggin:** My first job, yes, my first full time job. Yeah. I did other things before, but this was my first full time job. And I did a training of about a couple of months with other engineers and other assistant engineers like I was, learning about transistors and logic circuits and the stuff that was part of the trade. And then I got a project, by luck, because the guy who was in charge of the project had to go to the military service. And so he left the project, just started, and he was developing a computer and I took over. And under the guidance of an engineer, I started working on that. Then this engineer got into a car accident and the project was mine. And so I actually completed the design of the computer and built the computer, on my own, with five technicians that were working for me. And I was 19 years old.

**Hendrie**: My goodness. Now what was this computer supposed to do.

**Faggin:** The computer…

**Hendrie**: What was the reason they were building a computer?

**Faggin:** The computer was a study to see how simple a machine you could build. And so of course it was the size of a… It was basically a rack.

**Hendrie**: Now what year is this?

**Faggin:** That's '61.

**Hendrie**: '61.

**Faggin:** A little bit of '60 and '61. All of '61. So it was a rack full of PC [printed circuit] boards, right. Each PC board was a rather small format and had a couple of flip flops or a few NOR gates, made with basically DTL, discrete DTL. I mean, ’60, ’61, it was germanium transistors; silicon transistors were just beginning.

**Hendrie**: So it was germanium transistors and diodes.

**Faggin:** Germanium transistors…

**Hendrie**: And diodes.

**Faggin:** And diodes, and resistors, and capacitors, right? So that was it. But I designed some of the circuits, and I designed all the control logic and the rest of it. I had a 4K by 12 memory -- core memory -- in that computer. So it was a magnetic core memory.

**Hendrie**: Yeah. You didn’t have to do that.

**Faggin:** No, no. The core memory I got, but all the control stuff I did. Yeah.

**Hendrie**: I was going to say, because this core memory was a total art form in those days. That’s really hard.

**Faggin:** That’s right. Yeah. Yeah.

**Hendrie**: O.k. Great.

**Faggin:** And the output was Nixie tubes and a teletype -- the input teletype – and I got that thing to work.

**Hendrie**: That’s wonderful.

**Faggin:** Yeah. And that served as a model for what then became, eventually, the Olivetti Programma 101, which was one of the first desktop electronic computers that they came out with. Of course, the technology, the memory technology, was different; it used a magneto-strictive thing [wire], because the core was too expensive. And each gate was put into a little block of discrete components that big [showing size with hands]. And it was obviously engineered differently, and also the architecture was changed. But it was the feasibility that you can actually build a useful computer with so many gates.

**Hendrie**: Just so many parts.

**Faggin:** Yeah, just so many parts. So I accomplished that. Then I decided to go back to university and I decided to do physics.

**Hendrie**: Now why did you pick physics?

**Faggin:** Because I wanted to understand more of the basics, like solid state. I loved transistors and I wanted to understand really how they worked. Not just that you get collector current is beta times base current, and so on… -- the simple models. I wanted to understand how they really worked, what was going on. And, of course, I was reading books and all, but quantum mechanics was not particularly obvious to me. So my background in physics was o.k. but it was not at the level that would allow me to understand deeply what was going on. So, physics was really my desire to deepen my basic understanding, because, when it came to engineering, I knew how to get around. But I needed a little more basic understanding.

**Hendrie**: A little more background theory. Understand more of the fundamentals of what's going on.

**Faggin:** That's right. That's right. Yeah, yeah. So I took physics at the University of Padua and I completed it… I got summa cum laude. And I did it in four years when typically it takes six years. And I was the only one to do it in four years. I distinguished myself, coming from a background that was not typical for an AA kind of person to actually continue with university. In fact, up until some time before, it was impossible, because people that did that type of school were… I have got to go. Sorry. We've got 4:00 o'clock, yeah.

**Hendrie**: We will take a pause now. We will stop and come back in a minute. All right. Let's continue. I have a question about just a little bit of detail about the computer that you ended up having to design, build and get working. You said it was designed to explore how small and inexpensive one might be able to make it. Now was this a twelve bit machine, or an eight bit machine?

**Faggin:** It was an unusual architecture. The basic idea of the architecture came from an article in an American publication. And fundamentally -- I actually forgot who the author was and so on -- but the basic idea was to make a decimal machine, where the accumulator was made of ring counters that then would carry over, and so you would count. Additions and subtractions were done by presetting the number and then counting pulses and letting the carry ripple. O.k. And the accumulator then was driving directly a set of Nixie tubes. And so the accumulator was a decimal… It was designed for decimal arithmetic. And, of course, we also could do, in the control system, we could do binary operations, and we had a full instruction set and so on -- which I actually ended up designing based on some of the ideas in that article and some books and so on. And so it was a hodge-podge of things. But you could do very interesting things with it. Yeah.

**Hendrie**: It was a real computer and it really worked.

**Faggin:** Yeah. It worked. It was absolutely… Yeah, yeah.

**Hendrie**: That's great. O.k. So that's interesting. Very, very. So you completed that at Olivetti and then what happens next?

**Faggin:** Yeah, well, as I mentioned earlier, I went to the university. And I did…

**Hendrie**: That's right. Yes. Sorry.

**Faggin:** Yeah. That's right.

**Hendrie**: And you studied physics.

**Faggin:** That's right, yeah. And so after I graduated I became – I don't know the equivalent title over here. It would be like… I was basically teaching for one semester and doing research in the university, teaching Electronics Laboratory to the students of physics for the third year. So it was like assistant professor, or a step below that. I don't really know. But I was paid, which was unusual for just graduates at that time. And I was very envied because there were people that would have to do two or three years of basically voluntary work to be able to proceed in the university.

**Hendrie**: Work their way into the system.

**Faggin:** Yeah. That's right. But I didn't like it. I didn't like the environment. It was too slow for me. And so I decided to…

**Hendrie**: Didn't you have any interest in getting an advanced degreed?

**Faggin:** No. That's the degree. It's a doctorate in physics. I mean…

**Hendrie**: Oh, you got a doctorate in physics.

**Faggin:** Yeah. I mean, there was no other degree. The thesis work was on flying-spot scanners. I designed the core of a flying-spot scanner, the opto-feedback loop, to have high accuracy for the reading of bubble chamber photographs. And I designed the whole system, built it, made it work and did the theory of it, all in one year, which was the minimum allowed time to be a thesis. So that's it.

**Hendrie**: So you used all your electronics background to do a practical...

**Faggin:** Of course. Of course.

**Hendrie**: Thesis.

**Faggin:** Absolutely, absolutely. Yeah, yeah. That's right. That's why I said I was an exception to doing it in four years, when typically it would take six years to get that degree. Although four is the absolute minimum that you can do it. They won't let you do it in less than four years, even if you are a genius. So there I was, and I was ready to do whatever I wanted, to some extent, obviously. And after six months, the full academic year, I started actually teaching before I even got my degree, but it was for a few months, then I got my degree in December and then I kept on teaching until June, the end of the academic year. Then I left and went in the industry. I ended up working for my first boss at Olivetti, who had started a startup company in Milan, which was kind of unusual, in thin film circuits. In those days there was still some kind of residual interest in that direction even though the main path was integrated circuits. You probably remember in those days people were playing with thin film circuits. I was there for less than a year, actually, but it was an important experience for me, because this little company – there were only seven or eight of us, maybe ten -- had taken the rep position for GMe, General Micro electronics, which was the first pure play MOS company in the world. It was a spin-off of Fairchild, here in Santa Clara, and so this is in 1966 -- I graduated in '65, December of '65, so it's in 1966. So because of this, my first assignment was to come here [in California] and learn about MOS circuits at GMe; take a course, understand their product line, so that I could be helping customers designing-in [such products]. Among other things that I was supposed to do, I was getting this company started in this business.

  

**Hendrie**: So, now, this was a startup, but it was acting as a representative, and designing thin film circuits…

**Faggin:** Yeah. They were doing things…

**Hendrie**: I understand. One thing to pay the bills, the other thing to…

**Faggin:** Yeah. That was the idea. But it was really a company where the guy running it was a physicist and really, frankly, in business he didn't know what he was doing. But, of course, in those days I didn't have any better judgment of business myself, and so… And, of course, I was attracted by the idea to actually come to the States and understand MOS devices, which I had read a little bit about before, but I didn't really know much about it. In those days they were still a curiosity and there were very few people who knew about MOS devices. Anyway, I did this course, and for me to come to Silicon Valley at that time was really a revelation. You see the environment and how people would work and having a sense of an environment that was really much, much more akin to my ideal place, my environment. So even if it was a short stay -- I was here, I was in the States, for about two weeks -- somehow a part of me said I would be back. And in fact, I was back a few years later, and I've been here ever since. But anyway, this company, however, could not make MOS devices, because in those days nobody knew how to make them, to make them reliable, because metal gate technology, which was the technology of the day, was not a good technology: It was slow and unreliable, and not living up to the theoretical promise of MOS devices. So then, while I was at this company -- its name was CERES -- GMe was sold to Philco-Ford, and then because Philco-Ford already had some reps in Italy, then they canceled the rep [CERES]. So then, all that activity went away, and I basically decided to leave the company and go to work for SGS-Fairchild, which was the only semiconductor company in Italy at that time, in those days. They were making, basically, transistors and power transistors, and they had a license to make planar devices from Fairchild, hence, SGS-Fairchild.

**Hendrie**: But they had gotten started independently as SGS.

**Faggin:** They started as SGS, yes, to make germanium transistor. In fact, those were the transistors that I used in my first computer. O.k. So, there you go, it's a small world, right? So after… While I was graduating they took the license of Fairchild, and Fairchild took a position in the company -- they owned 30% of the company -- and they were licensees for Europe -- manufacturing licensee for Europe -- of Fairchild silicon transistors and integrated circuits, which were just beginning, in those days, to come out. We're talking about '66; it was only a few years under their belt. So I took a position in R&D there and strong of my long experience with MOS devices [ironic] -- of course, those two weeks of course, that I took -- I was assigned the task of designing, of making, developing the process technology, and developing circuits with… MOS circuit. I had to be able to do it, right?

**Hendrie**: You had a PhD in physics, right?

**Faggin:** And I did.

**Hendrie**: Were you in Milan?

**Faggin:** I was near Milan, Agrate, Agrate Brianza. It's a lovely place. It's about, probably, 30 miles east of Milan. So actually, within six months that I joined, I had an MOS process technology and two integrated circuits designed, that then went into production. So I did that by myself with the help of a technician.

**Hendrie**: So how did you figure out how to…

**Faggin:** I got a lot of papers, some from the Fairchild guys that were at the Palo Alto R&D; they had also their internal reports on the theory of all the oxide, on how to make… The key thing was how to make the oxide interface clean.

**Hendrie**: Oh, o.k. That was the difficult process step.

**Faggin:** That was a difficult process step. The rest of it was relatively easy to do. I mean, it was a vanilla kind of process steps that you would use already in making integrated circuits. So that was not very much of a challenge, although for me it was new stuff, so I had to figure it out.

**Hendrie**: You still had to learn about it

**Faggin:** But it was not where the difficulty is. The difficulty…

**Hendrie**: You didn't have to invent anything.

**Faggin:** That's right.

**Hendrie**: You just had to learn.

**Faggin:** To learn all of it, and then adapt it to the situation that was there. But to really have clean oxide, it took a little bit of work. And eventually we succeeded. We figured out the type of annealing to do and all that kind of stuff, and how to grow it, and all that. So I joined -- it must have been May or so, of 1966 – SGS, and then at the end of that year, they asked me if I wanted to go to the States for a six-months period as an exchange with an engineer from Palo Alto Research Laboratory, at Fairchild, and myself. So we were exchanged and we came here.

**Hendrie**: So you're working in the research laboratory.

**Faggin:** I worked in their laboratory, yeah.

**Hendrie**: Yeah. Not in production, the production line.

**Faggin:** No, no, the research lab; yeah, yeah. And I also married that September.

**Hendrie**: I was going to say. I was getting to it. And when did you meet your wife?

**Faggin:** We actually met in our university years. In fact, we met at the train station, of all places, because we were both from Vicenza and we were commuting by train to Padua. So we met. And, of course, I had seen her before in Vicenza, so I introduced myself and that was the beginning of the relationship.

**Hendrie**: So now you've gotten… You were married at this time.

**Faggin:** So then, in September, I married and when they asked me to go to the States it was after I got married, when I came back after my honeymoon, probably a month later, or two. It was funny because the head of the [SGS-Fairchild] lab, who actually was a Canadian [Paul Beneteau], started talking to me in English to see if I could manage English, because obviously English is not my mother language. So I managed enough to convince him that I could handle the stay here in the job. That was a sort of an important test, because the rest of it, academically and in terms of what I could do, there was no question about that.

**Hendrie**: So he's perfectly happy to have you go to the States to do that.

**Faggin:** That's right, yeah. So he asked me if I wanted to go. Of course, I jumped up and down [of joy] and said, yes, I want to go.

**Hendrie**: How did your wife feel about this?

**Faggin:** My new bride was quite happy to get out of the fog of Agrate because winter in Agrate is not a pleasure. So she was quite happy to do that, as well. And that it was [going to be] a great experience -- six months, newly wed, it's an adventure! So she was happy to do it. So we agreed and the date was set for early February. And I haven't been back since. Except for, obviously, of course, for visits, but, yeah, I never left the States after that.

**Hendrie**: Just a little bit of background: Why did SGS and Fairchild want to do this? What was it? Just a technology transfer kind of thing?

**Faggin:** No. It was trying to create a closer connection between the two R&D labs, so that there would be people that knew people; that kind of thing. And then, as it turned out, while I was in the States, Fairchild decided to sell their interests in SGS-Fairchild, just about when my time was due to go back. And I decided to stay, because at that point they asked me to say. And I said, yes, of course. So at that point it became a permanent affair, as opposed to just a temporary thing.

**Hendrie**: Well, when you got to the labs, do you remember what the first things they… What happened? What were the first things they wanted you to do?

**Faggin:** The first thing that I did was to stop in Philadelphia, at the International Circuit Conference.

**Hendrie**: Oh, the Solid State Circuit Conference that was every year in New Jersey.

**Faggin:** Actually, it was in Philadelphia, University of Philadelphia, that year.

**Hendrie**: O.k.

  

**Faggin:** And so I went to that for a few days -- which was my first big conference, international conference. So it was good. And then I took the long road from Philadelphia to California by visiting Washington, D.C. and New Orleans. I might as well get a little bit of sight-seeing, do an extension of the honeymoon, right? So we basically traveled for about a week, ten days. And I got to California, it was the second half of February -- I don't remember the exact date, but it probably was the 20, 21, 22, of February -- and I had left, I remember exactly, the 9th of February, 1968. For me to see California at that time of the year, full of flowers and so on -- and for my wife -- it was like we had landed in Paradise. Because we left Milan. Milan at that time is cold and foggy and miserable. And here it was Spring. And so, I was already in love, but she fell in love, too, of this weather and place. Of course, in those days, there were orchards all over the place; it was not as built up as it is now, and it was much cleaner, although it is still reasonably clean. It was absolutely fantastic. That first six months that we spent here, were perhaps the best of our life. We just married, the weekend we would go explore places around: Carmel, Yosemite, whatever you like. So, it was a wonderful time, and we became very attached to this place. So my first… When I joined, I had the choice of working on [one of] two projects, of which there was some thinking done, but really no real work.

**Hendrie**: It wasn't something that had been underway for six months and just joining the team. This is a new project.

**Faggin:** No, no, no, no. I was assigned a project. It was my project. And I could choose between two. One was to develop, to design an MOS shift register, it was a big shift register, as a product. And the other one was the development of a self-aligned gate process technology using silicon as the [gate] electrode, instead of metal. And there was done some preliminary work, but really they were still… No one had figured out how to do it. Basically only some experiments showing that… In fact, Tom Klein, who was the guy who had done the experiment, had figured out that, because of the work function difference between P-doped poly-oxide-silicon, of that whole structure, you actually would have a reduction of threshold voltage of about 1.1 volt, which was an important element, too, of the technology. In those days, the threshold voltages of MOS devices, in practical processes, were around 5 to 7 volts. So you had to apply 24 volts [of supply voltage]. And so they were…

**Hendrie**: Oh, really, they were like 5 volts?

**Faggin:** Absolutely. Five to seven volts.

**Hendrie**: Five to seven, o.k.

**Faggin:** Absolutely. Yeah.

**Hendrie**: I didn't realize that.

**Faggin:** Yeah, yeah. And so then you had all kinds of problems with parasitic MOS devices, because you needed a very thick oxide in the field not to create parasitic devices, and therefore stray paths for current. And so it was very tricky; it was very tricky. You were walking a very narrow window of possibility, all using 1-1-1 material. In those days, there was no ion implant. It was before ion implant became a practical way to do things, and that was several years later. Although there were experiments going on in those days at Hughes, it was not available. This could have solved the problem. So, being able to lower the threshold voltage and use 1-0-0 material, instead of 1-1-1 material, would allow us to build integrated circuits with much lower [supply] voltage and so on.  So that was a good observation, but nobody knew how to do it, to actually build the structure, how to build integrated circuits.  And there was some work done at Bell Labs where they proved that, again, you could do some self-aligned transistors, but the process they had used -- which nobody told me about it; I found out later -- that process was not suitable to making integrated circuits; it was only suitable to make transistors, because basically the gate was a ring inside a hole.  And so you could not do it. Basically, in order to do this thing, I had to make two major contributions.  One is to figure out the architecture of the process.  How do you do it?

**Hendrie**: How do you build up the layers.

**Faggin:**  How do you build up the layers to make it happen, which nobody knew how to do, and so I invented that, and, also, to figure out how to etch reliably the poly-silicon.  Nobody knew how to etch poly-silicon reliably.  So I develop a special etch which eventually ended up being called, Freddy's Etch.

**Hendrie**:  Very good.

**Faggin:**  After experimenting and burning a hole in my shoe, because it's…

**Hendrie**:  All right, we're started.  I think we ran out of tape when you just mentioned, you know, that you developed this etch.  Maybe you could just carry on a little bit more with a little more data. That was very fascinating detail.

**Faggin**:  So, last time you asked me what kind of method I used to develop the etch.  And, basically, my first step was to find the chemist-in-residence, and ask him if he knew how to etch silicon, and he didn't know because etching silicon was not a typical step that you would do.  You would etch oxide or aluminum or nitride, but not silicon.  And so I looked, you know… I asked a bunch of other people around, and I don't remember who, but someone told me that he had an article on how to etch bulk silicon.  And so I got a copy of that article, and I found out in that article that a mixture of nitric acid and hydrofluoric acid would etch silicon. Of course they also were talking about the rate at which it would etch: it was way higher than what I needed.  But, you know, clearly, at that point, you have a starting point. You know that buffering a solution of those two acids would slow down the etch rate because it's normal, you know, I did my chemistry.  Then I started experimenting with different ratios of the three elements, and quickly I narrowed it down to reasonable proportions that seemed to work.  I had an etch rate that was at the proper rate, like it would take a few minutes to etch the poly-silicon, instead of an hour or two seconds, so that it could be controlled; and also the proper differential etching between the silicon and the oxide.  I also found that I had to grow an oxide layer on the silicon to improve the adhesion of the photo-resist on the poly-silicon otherwise there was a tendency for the etching to go under the photo-resist too fast.  And so, you find all those tricks, and within a short period of time I was able to have a reliable way to do etching of poly-silicon, and then I could proceed with the process.

**Hendrie**:  Very good, excellent.  And so how long did all of this take?

**Faggin**:  Two or three weeks; It actually took two or three weeks for both the figuring out the architecture of the process [and developing the etch].  In fact I proposed three different architectures, and settled on one of the three as being the simplest and the most likely to work.

  

**Hendrie**: Oh, yeah. Could you maybe just talk a little bit about what the three were? I mean, that might be interesting for historians.

**Faggin**: For historians, yeah. I called it process A, B, and C. Process A was what ended up being the process that was used by Fairchild, Intel and so on. And process B was a variation of process A, but it was less desirable because it had… and now I don't remember exactly, but basically, it was not a big change over the process A, but it had an undesirable characteristic: it had some deposited oxide that typically is dirty, where we could actually have grown thermal oxide. And so it was discarded. And process C, which was like process A, but then that process had direct contact between poly-silicon and silicon. It required one more masking step, but it would allow much higher density of interconnections. It was my preferred process. And then my boss, who was [Les] Vadasz, he was really a putative boss because I was really reporting into SGS-Fairchild, but he was my local supervisor…

**Hendrie**: He was supervising your work, he wasn't paying you.

**Faggin**: He actually didn't want me even to try. He said, "It doesn't work, it will not work," and he didn't want me to even try it, which, by the way, I didn't do. I actually tried it. As a good experimenter, I found a way to put it into the mask anyway. So, later on, that process was key to the microprocessor, exactly that process.

**Hendrie**: Adding that step.

**Faggin**: Adding the step, what was called "buried contact." The buried contact between the poly [and silicon]. Direct contact, without having metal to connect the two.

**Hendrie**: Yes, not having metal between the two.

**Faggin**: So you could have metal running on top because, in those days, you basically had to do everything in one metal layer. Two metal layers were way too expensive, metal would break, and it was really…

**Hendrie**: They didn't understand how to solve any of the problems of dual-layer metal.

**Faggin**: That's right. So having that ability to use poly-silicon as an interconnect layer was very precious, and allowed to…

**Hendrie**: So you could take the gate and it could directly connect to… Electrically…

**Faggin**: That's right, to the drain of the prior stage so that you could have metal running on top of it. So that was a key. And of course, I did it because I recognized the value of that [step] because I was also designing circuits, I wasn't just the process development guy. I designed…

**Hendrie**: You were unusual in that sense. You were a circuit engineer who also was working on the process and a logic designer too.

  

**Faggin**: That's right, a logic designer too. I had designed and built a computer too, right. So I knew the whole spectrum. And so I wasn't limited in my view by the purview of a classical process development guy that would simply go down a certain path. So those were the three processes. And as I said, we settled for process A, but then in the succeeding weeks, I developed a test pattern to test the process. And in that test pattern, I put also test patterns for the process C: both a diode, which was created by a poly-to-silicon diode, and a resistor with poly-to-silicon connections, so that I also could -- and then an [identical] resistor without that -- so I could also find out the contact resistance of that poly-to-silicon contact. And those were enough test patterns to characterize the fundamentals of it. Okay. And so I got that [test pattern made] and as those masks were made, I designed the process flow. In other words, all the schedules for temperature, and the POCl-3 and all this stuff that you've got to do to build up the process. So I developed all the process, what is called the run sheet, right. And then I supervised the operators in the fab that we had, to go through the process steps and see to it that it was done according to my schedule.

**Hendrie**: Is this an R&D fab, a separate…

**Faggin**: R&D fab, yeah. In those days, each department had his own little fab, believe it or not. You know, because real men have fabs, right? So, the digital guys had to have a fab, the analog guys had to have a fab. I think there were three or four fabs in the research and development lab in Palo Alto.

**Hendrie**: Just in the R&D Lab.

**Faggin**: Yeah. They were small fabs, but, you know..

**Hendrie**: I understand, but..

**Faggin**: And, of course, in those days equipment was not -- it was not million-dollar equipment, right? But still, they were expensive. You had to still have a clean room, and you had to have aligners and everything, and furnaces, and all this stuff. So because I had developed the metal gate MOS technology at SGS-Fairchild a year earlier, I really was very familiar with all the equipment, how they worked, and all that. So, a few months into it, we got the first test pattern out, and we got -- if not at the first time or after making a few tweaks, soon after -- we got transistors working properly and the whole thing. And so, then, we decided at that point -- of course we had to characterize and do all kind of stuff in parallel -- to also design an integrated circuit using that technology to really test it out and see how it would actually compare -- that integrated circuit with another integrated circuit with metal gates -- and do an A-B comparison and see the differences. So we chose, Vadasz actually recommended it based on recommendation of the manufacturing guys in Mountain View, to use a device called 3705. The 3705 was an eight-bit analog multiplexer with decoding logic. So, you had these huge transistors, MOS transistors, that were basically like resistors: they were used in the linear region to multiplex analog signals. And so they had to be very low, very low on-resistance, and very low leakage, right? So that you could multiplex analog signals. And they had real trouble to making these devices in manufacturing, with low enough on-resistance and low enough leakage. They were real products that were used mostly by the military because they also needed very low power and all that. And there was decoding logic for the…

**Hendrie**: That was trivial decoding logic, the hard part were… What were essentially the switches.

  

**Faggin**: The switches, okay. They were switches, but they were used, as I said, in the linear region so that the ideal would have been 20, 30 Ohm, or something like that, but they were more like 150, 200 Ohm resistors, so I laid out the device. Vadasz didn't even want me to make it smaller. I could make it half the size, but he didn't want to. He was, in the past, pretty much of a task master in those days, actually, since I've known him, but after I left Intel, I never had anything to do with him again. But he was really a pretty hard-ass kind of guy, right. And he didn't want me to even make a smaller device. He said: "No, no! Do it just like the other one, of course in silicon gate, but like the other one." Okay, so I did it. And basically I built it a few months later, after we got to mask making and so on. I think I got the first devices to go out around June, June-July timeframe…

**Hendrie** What year now?

**Faggin**: Same year, '68. So, I got in toward the end of February and I got devices working around June. In June we got the first devices out, they worked, and it was very clear that they were very superior to the metal gate counterpart. And never mind the speed, the speed was much higher; the on-resistance was one third; leakage was 100 times less. Just...

**Hendrie**: Just a much better product.

**Faggin**: Yeah, much better product, right. So…

**Hendrie**: What about the ease of, what about the yield or the ease of manufacturing?

**Faggin**: Well it was very difficult to test, at that time, the ease of manufacturing because…

**Hendrie**: You're on this little prototype line.

**Faggin**: yes, number one. And number two, you've got to run some volume before you can test it. So you can not tell, right. Anyway, we got those devices to work, and then the big event… The big event is that, shortly after, [Bob] Noyce and [Gordon] Moore left to start Intel. And, in fact, they started the very day that was my first day as a full-time employee of Fairchild. So July 1st, 1968, that was the announcement date, was my first day [as a Fairchild employee] because I had converted from a guest to an employee. In June, or a month and a half before, I made that decision: Vadasz asked me -- I think it must have been probably in the May timeframe -- he told me that Fairchild was selling their interest in SGS-Fairchild, and asked if I was interested in staying. And so I jumped up and down and said, yeah, I'm interested in staying. And so, I soon got an offer, and then I told my employer that I was not going to come back.

**Hendrie**: That's very simple, I am just not coming back, California is too nice. The work is great and the climate is wonderful.

**Faggin**: <Laughing>. So, my first day, when I was switching from one employer to the other was July 1st, and that was the day that the whole lab was abuzz with the news that Noyce and Moore had left to start a company. And then, a week later, Andy Grove left. Andy was…

**Hendrie**: And what was Andy doing at the time?

  

**Faggin**: Andy, at the time, was the like a special assistant to Gordon Moore [the head of the lab]. He was a sort of right-hand man, at the lab. Before that, he was heading up the physics department of the lab, of the Palo Alto lab, yeah. And…

**Hendrie**: And Gordon was a physicist too, right?

**Faggin**: Gordon was a chemist.

**Hendrie**: Oh, he was a chemist.

**Faggin**: Yeah, Gordon was a chemist. And then, I think it was another week, Vadasz left. Then I started to smell a rat, right. And then it must have been a month and the rat at that point smelled: They hired the guy that was doing the poly evaporation. There was only one poly evaporation machine in the whole lab, so they hired "the" guy, the technician that did the poly evaporation for me, for the Silicon Gate Technology. So it was very clear what they were going to do.

**Hendrie**: Yes.

**Faggin**: Okay. And so at that time, I told the boss of my boss [Bob Seeds], because my boss could be the next one to go, and that was Tom Klein at that time -- after Vadasz left, Tom Klein took over – that...

**Hendrie**: Who was the boss of, who was Tom's boss, Tom Klein?

**Faggin**: Vadasz was Tom's boss before he left. When he left to go to Intel, to join his buddies at Intel, then Tom Klein was put in charge, only for a short period of time, and then somebody else was. But at any rate, so I told the boss of my boss, who was Bob… I forgot his last name, it might come back later [Bob Seeds]. I told him, "hey, I think Intel is going to do Silicon Gate Technology because I know that Vadasz was very keen on it, and so was Grove, and, you know, it's a much better technology, it's the right thing to do. And they just hired the [silicon deposition] technician, and so, watch out!" And he said, "Oh, if they do, we sue them." "Okay," I said, "but, we better hurry up because they could come out ahead of us." Because the whole transfer of technology from R&D to manufacturing was always a big problem in the Fairchild culture, in those days. So, anyway, that was the first inkling of what Intel was going to do.

**Hendrie**: And Intel did, eventually..

**Faggin**: Oh absolutely, they…

**Hendrie**: Did Intel make any metal gate MOS parts?

**Faggin**: Never.

**Hendrie**: Never.

**Faggin**: Never. No, they…

  

**Hendrie**: And any MOS they did. They did do some bipolar, did they?

**Faggin**: Yes. They did some bipolar. Part of that is because they were having some difficulty, early on, with the silicon gate technology, and they wanted to be in the market fast enough. But in those days you could design a memory in a few months: It was a relatively simple thing if you were doing something that was repetitive. The difficulty was to do random logic; it was much more complex because there were no tools, right? So that's how I started! And that takes me back to the July timeframe. And within a few months, since I was running more of this device -- it was called, by the way, 3708 -- and a few months later, toward the end of the year, it became actually a production product. It was sold -- and it was the world's first self-aligned gate MOS device ever to market -- and they sold it to a number of customers and it was built in the lab fab. So, I found out that the poly silicon -- evaporated poly silicon -- was actually cracking at the edges of the oxide well. So, basically, the process technology, the way that the architecture worked, you make a hole where a source, drain and gate were supposed to be, and then you grow the thin oxide and then you deposit poly silicon all over the place, right. And then you etch the gate and the silicon connections, and then you keep on going, right. So where the silicon goes up the thick oxide it would crack, just like metal would, evaporated metal, right, yeah. So, I started seeing that problem, and of course that was obviously not good, right. So then we found out -- actually Tom Klein did find out -- that there was one of the lab people who were experimenting with deposition of silicon using… I think it was silane. I don't remember the details, but basically a vapor-phase deposition of silicon. And of course, vapor-phase deposition allowed you to cover the steps. It was well known that you could cover the steps much better, using the same technique that was used to deposit oxide, the Vapox, it was called, which is vapor-phase deposition oxide. And so there was a reactor that we could use and try, and we tried to see if this deposited silicon would actually withstand… could cover the steps, and of course it was doing very well. And so we switched to that, and that took care of that problem. Then, further down the road, we decided to do another major piece of work that was going on in bipolar devices: It was using phosphorous gettering. And phosphorous gettering allows to… It basically soaks up impurities in the silicon and in the oxide, and it cleans up the system. Basically, you put a layer of vapor-deposited oxide, which is highly, highly doped with phosphorous -- and also [you dope with phosphorous] the back [of the wafer] -- and then you have a heat treatment, and the phosphorous becomes a gettering agent, and the impurities diffuse and they get trapped in this thing, then you strip this stuff out, and you…

**Hendrie:** And you just take it off it and throw it away.

**Faggin**: Take it off and move on with the process. And that's done towards the end of the process so you can clean it up. And of course, that was key to having a very low leakage, and really improving the reliability of the devices. That, together with the rest of it, was the key to making possible reliable MOS devices. [Phosphorous gettering could not be done with metal gate technology since, to be effective, it had to be done *after* the deposition of the aluminum gate, however the aluminum-silicon system could not withstand the high temperature required for the gettering to work.] Of course, the self-aligned gate would make them very fast, three to five times faster, and also make them smaller -- with silicon gate -- because then you didn't have the metal limitation of having to have a contact in the source, and then the gate, and then the contact in the drain -- only metal -- that would create a much longer device. And, of course, if you we're using my process C, my buried contact, then you could in fact have an even much higher density. Later on I will give you an example, but basically it was a factor of two higher density for random logic – substantial! In other words, half the cost; half the cost! And that was key to making the micro processor possible at that time -- I mean 1970-1971 -- otherwise it would have been prohibitively expensive to do.

  

**Hendrie**:  Too much silicon area.

**Faggin**:  Too much silicon area.  Yeah.  It would have happened later, anyway, but somebody else would have done it, perhaps.

**Hendrie**:  Okay.

**Faggin**:  So, anyway, with that last refinement, using phosphorous gettering, which also would help smooth out the oxide and allow metal to be much [less fragile]… It would break less because you also had rounding of the oxide steps when you did this.

**Hendrie**:  Oh, I see.

**Faggin**:  Okay, yeah, yeah.  So with that technique, although we didn't use that to the extent that later on Intel did -- Intel did a lot of that, but we did enough that we didn't have a metal problem – we were able to have reliable devices, and those are the devices that were life-tested, characterized to the seventh degree, and then allowed to be sold. And that was toward the end of 1968.  So, and basically at that point, we had the characteristics that I cited earlier on: 100 times less leakage. We were really in the sub nanoamp versus 100 nanoamp, that kind of stuff.  And stable devices, the whole thing. And very fast and very low on-resistance, and all the good things.  Okay, and so with that, my project ended, and then the issue was transferring this process to production now.  Now it's a demonstrated process: it works, it has the right characteristics…

**Hendrie**:  You had given it a number?  Somebody's spying them, you're selling them, R&D isn't going to make them forever.

**Faggin**:  That's right.  So a transfer engineer was assigned from Mountain View to transfer the process [to production], and…

**Hendrie**:  That was how it was done…

**Faggin**:  How it was done in those days, yeah.

**Hendrie**:  Okay.  So an engineer who worked in the production operation, would come to R&D, and...

**Faggin**:  Would come in -- that's right -- understand your process flow, your run sheet – the way it was called, if I remember correctly, in those days.  Understand it, and then adapt it, *adapt it* -- which, of course, meant starting all over again in many ways -- to the production, to the production floor, which was a source of problems because the new guy coming in had to put his own signature on it and begin to make changes, right.  So that was the beginning of a never-ending process.  On top of that, the design guys, the engineers, the people that were doing chip design, in production, for production parts -- there was no design guy to speak of in R&D, all the designers were in the Mountain View facility -- they were basically saying that they couldn't make… that my claims that you can make smaller devices were completely unfounded; it was all bullshit in that, actually, it was taking *more* room. And also, since they had to deposit silicon, it was costing too much, and so they were not receptive to the silicon gate technology despite... To me it was like: "How can you not see it? It's so  obvious."  And of course they were not even telling me directly those things, they were telling it to my boss, and so it was kind of underhanded. And then, when finally I found out what was going on, I said, “Well, give me an example, show me a circuit that you have tried with the silicon gate and substantiate your claim that it’s bigger than with metal gate.” So they come back, and they showed me a circuit. And basically what they had done, using silicon gate, they had tried to do it the same way they were doing it with metal gate, which of course you can not do, because it’s a different way to envision the topology. You have to do it differently. And so it shows that it was a classical thing, where people don’t understand the soul of the new thing, and they want to do it the way that they used to do it with the old thing -- with the new thing.

**Hendrie**: Right, yeah.

**Faggin**: And, of course it was coming out bigger than with metal gate, because it was a contraption! So I showed an example of how to do it, with silicon gate, that was in fact much smaller than in metal gate. I couldn’t convince them. “Oh, yeah but it’s a special case, in general it’s not [true].” And so it was a very difficult…

**Hendrie**: They wouldn’t learn.

**Faggin**: They wouldn’t listen, they wouldn’t think.

**Hendrie**: They wouldn’t think about it.

**Faggin**: They didn’t want to change, basically. To me that was actually the beginning of a lesson in sociology that was astounding because, to me, it was so obvious that the silicon gate technology was the way to go, that I was blindsided. I didn’t even [conceive people couldn’t understand]… How is that possible?

**Hendrie**: How is it possible that these people cannot see.

**Faggin**: Oh, then, there was this other thing: “With silicon gate technology you can not make bootstrap loads.” See, one of the key circuit tricks that people were using in those days were bootstrap loads, and bootstrap loads -- I don’t know if you know what a bootstrap load is?

**Hendrie**: Yes.

**Faggin:** Okay. So basically, it allows you to have the load device of a gate -- if you put a bootstrap load -- to maintain a constant gate-to-source voltage. As you swing the output, it bootstraps the gate up, and it maintains relatively constant voltage in between [gate and source]. And so you have, almost a constant current source for the load, which allows you to go much faster, and also to take the output swing all the way to Vdd.

**Hendrie**: Yeah, exactly. Which is, of course, you want a “one” that is at Vdd, and the ideal is a “zero” that is down at ground, or whatever the…

**Faggin**: Instead of Vdd minus Vt, and of course it is a Vt with body effect, which means that, for example -- remember, manufacturing was using 111 technology metal gates, so five to seven volts threshold; 24

  

volts [Vdd], okay -- the output swing without bootstrap load would be, maybe, 12 volts of something, okay. And, of course, after seven to eight volts it would slow very much, so the overdrive of the next stage was very little because, if you have seven volts of threshold on the other guy [transistor], the overdrive is 12 -- 7, you know..

**Hendrie**: And you are going toward 12 volts, and you're coming...

**Faggin**: Going to twelve works very slowly, right. But if you go through seven volts very fast…

**Hendrie**: Like to are heading to 24 volt

**Faggin**: …on the way to 24, you overdrive the device. And so the speed difference is enormous. The second thing that was critical to have a bootstrap load, was to have push-pull drivers. When you drive large capacitive loads, if you create a push-pull stage where both the transistor to Vdd and the transistor to ground have big fat devices, with a lot of current-carrying capability, if you drive the top device with a bootstrap load, you can get a lot of drive; the swing is not a lot, but because it is a big fat device -- and by being a push-pull you have only one or the other device [on] -- it was basically a poor man's CMOS technology, right.

**Hendrie**: I was just going to say, it's the concept that is fundamental to CMOS, yes.

**Faggin**: Yeah, yeah. But then you need to swing it all the way to 24 volts, on the fat device, the device to Vdd, in order to drive the capacitive load effectively. Of course, particularly for random logic, you have buses and stuff, and you've got to have that otherwise you can not do it; it's impossible to do it, okay. So, common wisdom was, that you can not do bootstrap load with silicon gate technology because underneath the silicon there is no junction; where, if you have a metal gate, you can put the other side of the capacitor as a P-junction -- in those days we had P-channel devices -- and then you put the metal gate, and then you have a capacitor, right. So, "You can not do bootstrap load, and you can not do bootstrap load," and that was it, right. And so, for a while, I believed this line too, but then a fair number of months later, I figured out that it was not true, that in fact I could make a bootstrap load, even with silicon gate because, in fact, the conditions that existed in the bootstrap capacitor were such that it was always biased to inversion, and so, although I did not have a real junction on the other side, I had a virtual junction; I had an inversion layer created by the fact there was an inversion layer [under all operating conditions]. And so I could, in fact, have a bootstrap load, and I actually did some test device, to make sure that they'd work, and I proved that they were working just before joining Intel. At Intel, of course they were not using bootstrap load because -- of course -- you couldn't do bootstrap loads with silicon gate technology.

**Hendrie**: Nobody had thought of this idea.

**Faggin**: That's right. And that was another key ingredient that was indispensable to make the microprocessor. Without that, there was no microprocessor. They would have been, probably, one fifth the speed, believe it or not. That's the difference that that simple trick did. So, the Mountain View guys were right to complain that they couldn't do bootstrap load, and that it was a fundamental limitation to do integrated circuit with silicon gate. Or else you would have to put an additional junction there – costing more money, so they didn't want to do that. And so they were reticent to using it. So, anyways, I sort of jumped back and forth a little bit, but the logical flow is lost if I stick to time only. So, going back, after I finished the 3708, we had production and so on…

**Hendrie**: So you got it into…

**Faggin**: Into production. No, no. Production meaning: shipping parts, but really done in the R&D fab.

**Hendrie**: Now you… Do you have any responsibility to carry it over into Mountain View?

**Faggin**: No. No.

**Hendrie**: That isn't the way that Fairchild works.

**Faggin**: No, in fact, after a few months of working with this transfer engineer, who was really not making much progress, and also he was not able to overcome the criticism on the other side, I actually went to the same boss and I told him -- the boss of my boss [Bob Seeds] -- that Intel was going to go with silicon gate. I said, "Look, Intel is going to do the silicon gate. I'm convinced, okay. I don't know if I approve, but I'm convinced that they are going to do it. Send me over to Mountain View to. I don't like to do it, but send me over to Mountain View to transfer the process because I don't like to be beaten by Intel. I want to be the first one to have complex integrated circuits using this technology because I believe in this technology, blah, blah, blah," right. And basically, he sat on it, he didn't do anything.

**Hendrie**: Didn't want to do it.

**Faggin**: I don't know why, I don't know what was going on, probably he didn't care. Basically, nothing happened of it. And that was something that I was going to do against my desire. [My] deepest desire was to continue to do research and develop new stuff. And so, next, I developed N-channel technology -- silicon gate technology. I developed N-channel devices, and then I developed CMOS devices using silicon gate. And then I worked on combining CMOS and bipolar in what I think, probably 15 years later, became BiCMOS. And I got that technology specked out, the process architected, when then I left Fairchild to go to Intel. So that was what I did. I also invented thin film [MOS] devices using poly silicon, which later on were used -- invented independently by other people -- they're now used in the TFT, in the flat-screen business.

**Hendrie**: Is that right? O.k.

**Faggin**: Yeah, seriously. I did the first devices back in '69, yeah. Etcetera, etcetera… I made a number of contributions, and it was a good period but it was getting frustrating, toward the end, because Fairchild, having lost all the people. Intel took about 30-35 people in a year and a half, out of Fairchild, and, of course, they were not the worst people…

**Hendrie**: And Fairchild had no defense against that? They didn't know about employment contracts and all those other legal…

**Faggin**: In those days it was not… I mean, there wasn't nearly the attention there is today to those matters. And of course, they basically took the technology that I developed, and they made it theirs. And they actually made me go… Before they left, Gordon Moore and Andy [Grove] and Vadasz, in agreement, they decided to make me go to -- in October of that year -- to ISSCC [International Solid State Circuit Conference] in, I think it was Washington, to describe the Silicon Gate Technology. And so, it became public domain. And I didn't realize [the implications of] what I was doing, but…

**Hendrie**: And they talked you into doing that? After they left?

**Faggin**: Well, not really, before they left, like in June. And they left in July.

**Hendrie**: Ah, but you had submitted the papers in June, so you're already committed, and your management wasn't going to tell you not to do it, because they had no ability to think through this strategy at all.

**Faggin**: Yeah, that's right, that's correct. Yeah, yeah. So it's interesting how, when you think back… I know I was a naïve kid, I was 27 years-old or something. So, well, it was October '68, so I was not even 27; I was 26, on my way to 27.

**Hendrie:** All right. We need to change tapes now.

**Faggin:** Yeah, that's it. I also need to call it for today.

**Hendrie:** All right, very good, excellent.

## Session 2: December 13, 2004

**Hendrie:** Today again, Federico Faggin, who is going to continue his oral history for the Computer History Museum. Thank you very much for being willing to do this, Federico.

**Faggin:** Thank you.

**Hendrie:** I think where we left off in the last session was, you were getting more and more frustrated with Fairchild and them not really seeming to do what they ought to be doing, or taking any advantage of some of the developments that you'd done, and your thinking about-- or you're about to leave for Intel. Could you tell us how, you know, who approached you? How did that happen? You know, what made you decide to go to Intel? Who invited you? Did they have to romance you? Or were you just ready as soon as they asked?

**Faggin:** Well, actually, I asked.

  

**Hendrie:** Oh!

**Faggin:** I called Les Vadasz, and I told him that I was frustrated, and that I actually wanted to move from process development to design. That I really wanted to design chips, using Silicon Gate Technology, as opposed to continue with what I had done prior, which was the development of processes. Although I had also designed a chip, but it was not a very complex chip [the Fairchild 3708].

**Hendrie:** Yeah, and it was really part of the test, to test the process.

**Faggin:** It was really part of the testing of the process, that's right. Although it became a product, but still it was part of a test, yeah. So.

**Hendrie:** Had you always wanted to design, or did this sort of come about you, "I've done enough process, it's time to do something different?" Where did the design bug-- has it always been there? Or why did you decide that's what you wanted to do?

**Faggin:** Well, I think that I always considered my activity in process development a detour from what I wanted to do at the outset, which was design.

**Hendrie:** Okay.

**Faggin:** But at SGS they didn't have an MOS process technology, so I had to develop that first, and then I also designed a couple of simple circuits. And at Fairchild I saw the opportunity to do the Silicon Gate Technology, which was once-in-a-lifetime opportunity, but after that I wanted to design, I wanted to go back to what I wanted to do, which was design.

**Hendrie:** And so did they say, "Yes," right away, "Come on over?"

**Faggin:** Vadasz, I don't remember if he said "yes" in the same phone call, in the same breath, probably he did. But definitely, if he didn't, then he called me back soon after to say, "Come over, and let's talk." So the whole interview process was actually relatively brief, because Vadasz knew me, so there was not an unknown factor.

**Hendrie:** Exactly.

  

**Faggin:** Of course, the fact that I wanted to do design was not obvious to Vadasz, and so he asked me what I had done in design, and so on. So at that point, my experience in designing and building a small computer, which I did at Olivetti, was a very, very important card that I played, and of course, that is what convinced Les Vadasz that I could do the job that he had in mind, which was developing the MCS-4 family, which included the 4004 -- those four chips, what at that time was called, the Busicom project. And also, kind of an interesting note is that Vadasz, when he interviewed me, did not tell me anything about the project. He said, "Well, we have a project that is pretty challenging, but I cannot tell you what it is." So, he left it open-ended, and I eventually took the job even without knowing what it was. But his word that it was challenging was good enough for me.

**Hendrie:** And he probably knew that if it wasn't challenging, you wouldn't be very happy.

**Faggin:** <laughs>

**Hendrie:** So he already knew what he was going to have you work on. Tell me about where the project was when you got there, and what had been done.

**Faggin:** Yeah, yeah. Actually Vadasz, in his description of the project, made me believe that it actually was in a much more advanced state of development than what I found when I joined. Of course, when I joined, the very morning that I joined, Stan Mazor, who was an engineer working in the application research group, that was headed by Ted Hoff-- Ted Hoff was absent, was traveling, the day I joined-- gave me the documentation about what had been done before, and told me that the next day Mr. Shima from Busicom was going to arrive at the airport, and he was going to come and check on the progress that was made in the project. So I had basically one day to figure out what I was supposed to do and get up to speed.

**Hendrie:** And figure out what had been done, and what needed to be done.

**Faggin:** Yeah, and what needed to be done, that's right. And so I went through the documentation, basically there was a formal documentation that was the description of a block architecture, without much detail, and then the instruction set of the CPU, and the basic organization of the family -- there were four chips. One was the CPU; one was a ROM with the input/output ports; one was a RAM with output ports; and one was a shift register, serial in/serial out and parallel out, static shift register; it was a MSI, medium scale integration, chip that was intended to be a port expansion for the I/O, so that was a simple chip. The other three were pretty much state-of-the-art chips. And then there were a smattering of trial circuit designs that later on I found to be totally useless, because they would not have worked, and that was it. And so...

**Hendrie:** There was no logic? Nobody had started a logic diagram?

  

**Faggin:** No logic diagram. There was none of that

**Hendrie:** Using MOS transistors.

**Faggin:** There was more useful work done in the area of the ROM, the 4001, the ROM chip and the RAM chip, simply because that was an area where Intel already had done some work. And so there were cells. It was a type of design that I had already done -- not the whole design, just the cell itself…

**Hendrie:** Yes.

**Faggin:** …a three-transistor cell for the dynamic RAM and a kind of unique serial type of ROM that was good only if you had a very slow ROM, for the 4001. And those cells I ended up using in those two chips. It was not the actual layout, but the basic circuit design.

**Hendrie:** Now had Intel already started work on the 1102 and 1103 at this point?

**Faggin:** Yeah, Intel had already announced, and was in production on the 1101. They were working on the 1102 and the 1103. As you know the 1103 was the one that ended up being in production; the 1102 was eventually scrapped. They were two competing, one-kilobit RAM, dynamic RAM, designs.

**Hendrie:** So they had experience, they'd been working on dynamic RAMs.

**Faggin:** They'd been working on dynamic RAMs, but they were still doing the layout, they did not have working chips yet of the dynamic RAM. In fact, I probably had working three-transistor RAM before the 1102 or the 1103 actually worked. But, in any event, it's immaterial, because my memory [the 4002] was not as challenging, in terms of speed and power, as the [1103] memory chip.

**Hendrie:** What you're doing at _____________.

**Faggin:** Right. I had 320 bits, but I had a bunch of other logic and stuff around it. So in terms of chip size and complexity, actually, my chip was more complex than the 1102 or the 1103, but it was less demanding in terms of speed and power. At any rate, the next day, after I spent the day looking over what was done and having a sense of what was the whole project, Mazor and I went to the airport to meet Masatoshi Shima, who was arriving from Tokyo. So we met, and Shima said, "I'm here to check. I'm here to check." "Okay, great." So we went to the office, and Shima immediately asked to check. So I asked Mazor if what I had was all that we had, or if there was something else, because Shima insisted that he wanted to check the logic, that was apparently promised that it was done. And Mazor said, "No, that's all we have." They had not done any more work, and he sort of tried to disconnect himself from the embarrassment. So I went to Shima and I said, "Here's what I was given, here's what I have."

**Hendrie:** And I've been here one day!

**Faggin:** And I've been here one day!

**Hendrie:** So it's not on my watch.

**Faggin:** Yeah, and Shima looked this thing over and a few minutes later said, "This is not what I came to do!" He was really mad! He was saying: "This is idea. This is not design. I want… I came here to check, and there is nothing to check!" He was really, really upset! And I said, "Well, you know, I just joined the company yesterday." "Yeah, you bad! You bad!" <laughs> "You were supposed to do!" You know, he probably meant Intel, right?

**Hendrie:** Meant Intel, yes.

**Faggin:** But you know, his English was pretty poor, so fortunately I understood that; I didn't take offense. But it took me a while to calm him down, and eventually there were some agitated phone calls between him and headquarters, where basically the whole project was in question because they had a schedule, and it was clear that this schedule that they had been given was not going to be met, because no more work had been done since he left in October, I believe. Like five [or six months before.]

**Hendrie:** And you were arriving in when?

**Faggin:** I went there in February [correction: it was April 3, 1970]. So there were about five/six months of no more work. Basically, after Busicom and Intel agreed to proceed with the project, Intel was supposed to proceed with the design of the chips and so on, and having the layout almost beginning to start, but no work had been done.

**Hendrie:** Do you have any understanding looking back why no work was done?

**Faggin:** I think that they couldn't find anybody to hire, perhaps. I really don't know. Certainly, their preoccupation in those days was with memories, and so this was like a filler project for Intel. But I think it was a combination of negligence and not being able to find somebody to start the project.

  

**Hendrie:** How many designers did they have at that time?

**Faggin:** Well, very few. There were four or five at the most.

**Hendrie:** Really?

**Faggin:** Yeah, yeah. It was a very small team. The whole company had about 130-140 people, maybe 150, but most of the people were actually in production. You know, there was a fab, and there was an assembly line; they were all in the same building in Mountain View. It was not a big building, it was probably 30-40,000 square feet at the most. And so it was a relatively small company. The whole engineering department, which included both MOS and bipolar -- the numbers that I gave you were for the MOS…

**Hendrie:** Yes, for the MOS.

**Faggin:** ...designers [only] – was probably 30-35 people. So that was the total [number of designers and support people] for both bipolar and MOS.

**Hendrie:** Now were Karp and Regitz there by this time?

**Faggin:** Karp was there. Regitz was not there. Abbot, Bob Abbot was there. Let me see… Frohman, Dov Frohman, who was there. He had joined just one or two months before, and he was developing the floating gate...

**Hendrie:** Working on the EPROM, yes.

**Faggin:** The non-volatile ROM, and electrically programmable. And, let me see…

**Hendrie:** Hal Feeney hadn't join yet.

**Faggin:** Hal Feeney had joined, but he was out. I think he was getting married, but he had joined just a few weeks before I joined. That's it!

**Hendrie:** Okay.

  

**Faggin:** So the company was very few…

**Hendrie:** Very few.

**Faggin:** …very few people.

**Hendrie:** Yeah, yeah.

**Faggin:** There were three people in Application Research, with a few technicians on top of it. Three engineers, and I include...

**Hendrie:** You include Ted in that group.

**Faggin:** I include Ted on that. So, it was a relatively small team.

**Hendrie:** Okay. After you calmed him down, what'd you do next?

**Faggin:** Yeah, after I calmed down Shima, I told Shima that, basically, he could continue to bitch, or he could help me. And if he helped me, we could get there faster, but that, frankly, the predicament that we were in was what it was. And so only by working hard and working smart we would be able to solve this problem. Eventually he was given permission -- probably about a week after I joined, no more than a week -- he was given permission to stick around and help with the project, given that that's all there was. And so I started working furiously. By the way, I had no draftsmen. I was by myself, that's it! There were no people to help me. That was the end of it. The schedules that had been prepared were completely, completely wrong. I mean, the layout of the CPU was scheduled to take one month, the same amount of time that it would take to do a ROM or a RAM chip, which was what they knew how to do. Just to give you a sense of how incomprehensible that whole project was to Vadasz and to the other people that were there; they just didn't know what they were getting into.

**Hendrie:** They had no understanding of...

**Faggin:** No understanding of what they were getting into.

**Hendrie:** …how complicated logic is.

**Faggin:** Yeah, and what you need to do.

**Hendrie:** Yes.

**Faggin:** So, anyway, I decided to start work on the 4001 first, which was the simplest one, and then in parallel… I basically staggered the four chips one after another, so that while the 4001 layout was done, I could design the 4003, and while the layouts of the two chips were done, I could design the 4002; and as the layout was still continuing, then the 4004 could really be designed as well. It was like… What do they call it when they build a house? Fast-tract building. Fast-tract building, which is basically building a building as you are designing it as well: you don't design it first, then build it; you sort of design and build at the same time.

**Hendrie:** Yes, exactly. And everything sort of staggers...

**Faggin:** That's right, and so that's what I did. But the first thing was to actually develop the basic rules, the basic design methodology that I was going to use, because there was no methodology. And of course, Silicon Gate Technology requires a different style of design that you do with metal gate. And there was no random logic design with Silicon Gate; it was never done before. So I had to figure out how to actually design random logic in a power-efficient, area-efficient, and time-efficient manner.

**Hendrie:** Okay.

**Faggin:** And key to that was some of the work that I had done at Fairchild just prior to coming to Intel, specifically the work on the bootstrap load, which was a key component in metal gate technology. And people, engineers, at that time, thought that it could not be done with Silicon Gate unless we were using an additional masking step, which would've been too expensive to do.

**Hendrie:** I see, which would add the capacitor in with it.

**Faggin:** That's right, that's exactly right. So that...

**Hendrie:** How did you figure that out? I mean...

**Faggin:** Well, actually, it's physics. Frankly, when I was at Fairchild, that was a major objection that I got from the designers at Fairchild: They were resisting the adoption of the Silicon Gate Technology because we couldn't make bootstrap loads. In those days, the most effective way to design random logic with metal gate MOS, was to use four-phase design, because it was the most cost-effective way to do it. You could make small devices, relatively small devices; they would use very low power, but it was a dynamic design, so you could not do static chips. But, when you have a clock...

  

**Hendrie:** That's not necessarily a problem in a CPU.

**Faggin:** That's right. In a CPU or most synchronous designs, it doesn't matter -- as long as you don't need to store [data] for more than a certain number of clock cycles, that's fine. In fact, if you remember Four-Phase Systems, it was a spin-off of Fairchild Semiconductor. Lee Boysell, who was a key designer of Fairchild, left to form Four Phase Systems to make, basically, computer system using a number of chips with four-phase design. But with four-phase and with metal gate, you couldn't do a single-chip CPU, though you could do a CPU with a few chips. Just to tell you that the idea of a CPU made with MOS technology was more than in the air. The other way to do it was -- for static designs -- was to use bootstrap loads. And bootstrap loads typically would use two-phase designs which, if you wanted to be static, you would have to have all the latches, all the flip-flops had to be static flip-flops. But with a two-phase [clock] and a bootstrap load, and using pass gates, you could in fact have cost-effective designs. But you needed the bootstrap load.

**Hendrie:** You needed the boot-strap basically so that you didn't have the very slow charge up to the… through a resistor up the…?

**Faggin:** Well, because otherwise you would have to use so high [supply] voltage to be able to have… If you have a pass gate, and you don't have a boot-strap load -- the voltage out of any gate is one threshold below the supply -- then as you go through a pass gate, which in its gate has a signal which is one threshold below supply, then on the other side of the pass gate, you immediately have two thresholds below. Now the two thresholds, you've got to remember, they have body effect: so if you have 12 volt supply, you come down to where you have only five volts or so after two threshold voltages, and five volts is just not enough to drive another stage. So what we had to do, we had to have, in the pass gates, a signal that was at Vdd, with the same value of Vdd, so that we had a one threshold drop also in the gates driven by that pass gate. I think it's a little difficult to talk about this without making a picture or something, but at any rate, without using dynamic circuits -- two-phase dynamic circuits with pass gates -- the amount of transistors required would have been prohibitively large, and then the 4004 could not have been feasible. And so for example, many of the… Well, not many, but some of the circuits that were drawn as trial circuits, had this technique, and they would not have worked [without a bootstrap load]. So basically the question was, "How did I figure it out?" And what I figured out is basically physics: If you create a capacitor that, in the case of a boot-strap load is always negatively biased, you actually have an inversion layer under the gate, the silicon gate, and therefore you have a capacitor there, which is almost 100 percent of what you would have if you actually had a physical junction. And, as the boot-strap raises the voltage, both voltages [at both sides of the capacitor], up, that inversion layer stays all the time, and as such, you end up with a working boot-strap load. I actually made a test chip at Fairchild just before coming and I verified that it worked. So when I told Vadasz that I was going… That I needed to use a boot-strap load, he said, "What are you kidding? Don't you know that it cannot be done?!" And so I explained how it worked, I explained the principles of operation. And Vadasz mumbled something, and then a couple days later, I described that… either that or… Oh, yes, Vadasz asked me to describe that to Dov Frohman. Dov Frohman was the device-physics guy, and so his verdict would have been final. And so I explained to Dov how the boot-strap load works, and again, initially there was a sigh of consternation that, you know, “You cannot do it.” And then I explained. And I got a, “Hmmm.” And then, I never heard of it again. So not hearing anything of it means that it works. <laughs>

**Hendrie:** It means he agreed. You understood the physics well enough. Well, that you might not have understood it if you hadn’t come from a process background.

**Faggin:** Yeah, that’s right, that’s good. That’s exactly right.

**Hendrie:** Just the circuit design.

**Faggin:** That’s exactly right. So then, in fact, a few days later, I was walking over the drafting area, and there was Bob Abbott -- Bob Abbott was the designer of the 1103 -- and he was busily changing some of the X-drivers of the 1103, putting, what looked like to me boot-strap capacitors, under the direction of Vadasz.

**Hendrie:** Yes.

**Faggin:** And I said, “What? Are you making capacitors there?” “Yeah!” “But what are they?” “Well, Vadasz told me to do it.” And I said, “Do you know where they came from?” “No, he didn’t tell me.”… But you know, that was the mark that it was accepted as an idea at Intel… And, all of a sudden, it was utilized without ever acknowledging who actually did it -- which was typical in those days of Intel. And so I assumed that a “not no,” was a “yes.” <laughs>

**Hendrie:** Okay.

**Faggin:** So anyway, so that was a key element of the design methodology. Then of course there were many other pieces to it that were much more relevant when we got into the design of the 4004. Because the first two chips they were really basically what was known before, a little bit of extrapolation of that, and the boot-strap loads were enough to do the 4001 and the 4002. And, of course the 4003, which was the shift register, was a fully static design anyway, so there was not even a need for boot-strap load for that chip. Though I actually ended up using it because it was very useful for the output buffers, otherwise, it would’ve had very wimpy output buffers. And, as you know, we had to drive a lot of current, because the whole system had about 32 chips. And so the external, the four-bit bus, that was used for both address, data and instructions had a lot of capacitance, like 150 picofarad, and it was a challenging job to develop buffers that were fast enough to drive all this [capacitance] for a maximum system.

  

**Hendrie:** Very good, all right! So you figured out a methodology, and started working on the first chip.

**Faggin:** Yeah, the 4001. Yeah.

**Hendrie:** 4001.

**Faggin:** Yeah.

**Hendrie:** And how long did that take?

**Faggin:** I went directly from the function to the circuit design, because the logic design was minimal. And it's easy enough to design the circuits to reflect the logic that you want, anyway.

**Hendrie:** When you say you're going to circuit designs, is this literally a layout of the cells, or just...

**Faggin:** No, no.

**Hendrie:** It's just all the transistors.

**Faggin:** All the transistors, how they're hooked up, and their sizes -- of the transistors.

**Hendrie:** Okay, gate sizes.

**Faggin:** Yep, gate sizes and all of that.

**Hendrie:** Okay.

**Faggin:** That's what I call circuit design. And then if you have some challenging timing problem, then you would do a simulation. I was under strict orders to use simulation very sparingly because of the cost. We didn't have a computer; Intel was using a data service with teletypes.

**Hendrie:** Time sharing services.

  

**Faggin:** Time sharing services, yeah. And so I did not use any simulation at all for the 4001/2 or 3. We used a little bit, very little, on the 4004, the CPU. So, what I did to avoid simulation, again, using some of the work that I had done at Fairchild plus some work that I did at Intel, I basically developed a methodology of using graphic design, where you use normalized characteristics of transistors -- both characteristics of transistors, static and dynamic -- to find the proper sizes. In that way, I could do the design very effectively. I had developed a bunch of rules of how to do this that made possible to quickly size up devices, and be within, probably, 20-30 percent of the actual result that you would get, because the graphics were taken from actual transistors; they were not taken from equations. They were measured characteristics of transistors, and they were normalized to Z /L = 1. And so I had switching characteristics, I had the drain characteristics, the source characteristics: when you connect gate and drain together, and you look into the source; etcetera, etcetera. So, a bunch of curves like that. And with that I was able to make -- and [using the] slide rule -- I was able to make quick calculations, and be able to do the job very effectively and very fast. I think it took me no more than a week to design the 4001. I mean, the circuit design, so that I could then start the layout -- what was called in those days, “composite drawing.” Okay.

**Hendrie:** Now were you doing the composite drawing yourself? You did not have a…

**Faggin:** Yes. Well, what happened is that, in the meantime, it took a while to hire a draftsman, and I couldn’t find any draftsmen that had experience in chip design, so I ended up hiring Rob Sayre from Lockheed. He was a mechanical design guy, so he knew how to design mechanical things, but he had never designed a chip, he didn’t know what an MOS was, or a transistor was, for that matter. And so I had to teach him everything, which was an additional burden to me, because I couldn’t just give him a circuit and say, “Make a layout of this.” Right? On the other hand, nobody had done complex silicon gate layouts anyway. So there was a need of an engineer in those early stages to just establish again, the methodology of layout, and how to do it right. So, I ended up, basically, drawing everything -- all the cells, all the circuits -- I would do the layout myself, and then give it to Rob, and say, “Rob, put it in.” And do, of course, all the planning of the chip, because that’s the most critical thing: where do you put the various pieces. And you had to have techniques to assess roughly, how big they are going to be, what kind of form factor they’re going to have.

**Hendrie:** You have to guess correctly.

**Faggin:** You had to guess correctly, because if not, you...

**Hendrie:** It won’t fit.

**Faggin:** It won’t fit.

  

**Hendrie:** It'll bulge out.

**Faggin:** It'll bulge out, and you have to start all over again.

**Hendrie:** Yes.

**Faggin:** You have to erase and start all over again, which, of course, I didn't have time to do. So I also developed techniques, which then were refined later on for the 4004 of how to assess the sizes of things. And also to actually draw the circuit schematic in a way that would reflect the actual layout, so that there would be less of a translation issue. So if you do the right planning of the layout, after you do the circuit design, then you draw the circuit design in a way that reflects where the transistors are, trying to guess correctly how the lines, the metal lines are going to be, and whether you carry the signals in silicon or in metal. And you mark it properly. So that in one drawing...

**Hendrie:** It's actually all on the circuit.

**Faggin:** It's all on the circuit design. That's everything that you need to know to do a very effective layout.

**Hendrie:** Okay.

**Faggin:** And that's another innovation that I brought into this new methodology that I developed. So the layout started with the 4001. As I said, I had to do everything. All that I did was to draw freehand. Instead of drawing with [a ruler, I used] quadrille paper, so it's easier.

**Hendrie:** Yeah, on quadrille paper. Yeah.

**Faggin:** So it's easy to do and to scale right. And then you give it to the draftsman, and the draftsman would integrate into the total chip. The layout probably took about a month for the 4001, because it was relatively simple, and one draftsman was enough. Just to calibrate you, the 4004 took about three, three-and-a-half months with two, and sometimes three, draftsmen. So that's how much more complex it was.

**Hendrie:** Yes.

**Faggin:** While the layout was going on, as the draftsman was doing the layout, then I would design the 4003. In the 4003, I used a unique, very effective, static flip-flop that I had co-invented with my boss at

SGS-Fairchild [Fabio Capocaccia]. That was the flip-flop that we used for the static shift register, because it had to be a master-slave flip-flop. And then that went into layout, or soon after, and then [that layout] also Rod Sayre did. In those days, after the layout, it would go into a rubylith cutting.

**Hendrie:** Uhm, uhm, yes.

**Faggin:** So basically there was a cutting table, a large cutting table, and you would put the composite -- it was done in Mylar -- the composite layout in the table, and then on top of it, you would put this rubylith: It was basically Mylar with a very, very thin skin-- red skin -- that you could cut and peel off. And then, of course, you would make one rubylith per layer of the mask, where the composite layout had all the layers superimposed. And then each of those layers would be photographed with a huge camera, reduced to about 10x from -- typically in those days we were using 400 or 500x -- from 400 or 500x down to 10x, which was about this size [showing with hand]. And from 10x then to 1x in the step-and-repeat camera, [used] to make the masters, that were then used to make the working plates that were used in the [manufacturing] process. So it was a long process, and prone to errors, because every time you would make a translation, you could introduce new errors, because they were not done by computer, they were done by human beings. So, the 4001 was in ruby cutting, the 4003 was in layout, and that was a very fast layout; it probably took two weeks, because it was a small chip. But there again, even a repeated cell, you had to draw over and over again, because there was no method of repeating a cell. So, for example, the shift register, had ten nearly identical cells, and you had to draw every single one, because there was no way to repeat them [by machine].

**Hendrie:** Yes. ’Cause there was no automation at all.

**Faggin:** That’s right. And of course, then each one had to be checked very carefully because when a human repeats something, he is not like a computer, he might introduce errors….

**Faggin:** Yeah, so about that time was also when the 4002 was designed, and then, again, the layout of the 4002 started right after the 4003 layout was finished. So, then, we had both the 4001 and 4003 in ruby cutting, and the 4002 was being laid out. Then, soon after that, the design of the 4004 started.

**Hendrie:** So about what time is this?

**Federico Faggin:** So probably by now we are around -- let me see-- around May timeframe.

**Hendrie:** May of...

**Faggin:** Of 1970.

**Hendrie:** 1970.

**Faggin:** Yeah, yeah. Because I joined in February, it was February 8 or February 9 [correction: April 3, 1970]. And so, in fact, I think that September was the time when we received the first 4001's, they were fabbed. Because the process would take about a good three/four weeks of ruby cutting, and then the mask making would take another few weeks. And then the process, the actual wafer fab would take two to three weeks, more like three weeks with a hot run, what is called just a speeded up run. So, the whole process is around three months, easily. From the moment you have done the composite and the moment you see the chips it's about three months. So, you know that we got 4001's around September timeframe. All right, so the 4002 was actually started with a new draftsman, Julie Hendrix, who joined from Fairchild, so she had actually experience in MOS design. In fact, she was the draftsman who did the layout of the 3708

**Hendrie:** Ah, yes, for you.

**Faggin:** For me, at Fairchild, a couple years before. So now I had two draftsmen, and of course, Rod was initially helping on the 4002, trying to speed up the process, and by that time he was beginning to get the hang of it. But anyway, so I started the design, the logic design and circuit design [of the 4004] which, as I mentioned earlier, were done essentially together -- there was never a separate logic diagram drawn for the 4004 -- and where I developed the basic methodology in addition to what was done before. So taking all of what was done before, that I described earlier, adding some additional elements that were essential given the nature of the design -- that was essentially all random logic. And that's where I perfected the techniques of drawing these circuit design drawings in such a way that they would reflect as close as possible the way you would layout the chip. And also I decided to have a way to do the control -- that was not even attempted [before], and it was the most difficult part of the chip -- to do the control by basically creating some structured logic. And the way it was done, it was done in three levels where you would combine, first, the slow signals, then you would go to intermediate [speed] where you would have intermediate temporization, and finally, at the very end, you would have the clock that would drive in the signal [to the CPU's functional blocks]. And this was done so that it would be like if the whole information would flow properly through, and you would have the best way to have the shortest possible time [with the lowest power dissipation] between the decoding of the instructions and the actual control signals that would go and actuate the various functional blocks, like the memory, the ALU, etcetera. And then I gave that to Shima, and Shima essentially did the rest of the logic design by himself. Because at that point -- Shima was good at logic design, but I taught him the basic method: how to do it – he was filling in for all the various signals and so on, and he did a very good job. So he completed the logic design while I was busy as hell in the layout phase, because that was the most challenging aspect left to do.

  

**Hendrie:** Especially with all of the random logic.

**Faggin:** With all of this random logic. That's right.

**Hendrie:** And all of the internal busses and the problems with that.

**Faggin:** That's right. And I had never done a chip like that, so I needed to concentrate there. So Shima completed the logic, and then did for me some simulation -- again, all under my direction -- the simulation of the path from memory to internal bus. We had an internal four-bit bus that would then be reflected into, sometimes into an output, or also simply a memory. The internal bus was also used as a memory, because it was a dynamic circuit; it was a memory of the state of the internal system.

**Hendrie:** Ah, all right.

**Faggin:** But it was very tricky and so we wanted to make sure that we had the proper timing from memory registers to the internal bus and then out, to the outside, to the external bus -- the four-bit external bus.

**Hendrie:** Okay. Very good.

**Faggin:** And so that's it. Okay, that was the only computer simulation that we did, where we wanted to make sure that everything was fine. And then the layout that took about three-and-a-half months, where we also added one -- not for the whole period, but for a short period of time -- we added one more draftsman. So there actually were --most of the time, actually, probably two out of the three-and-a-half months -- there were three draftsmen. And at the beginning there were only two.

**Hendrie:** Now did you have any specific die size limitations that you had to squeeze through?

**Faggin:** Oh, absolutely.

**Hendrie:** A certain size die?

**Faggin:** Absolutely.

**Hendrie:** Where did those come from? And what were they? Do you remember?

  

**Faggin:** Well, they came from basically yield curves that manufacturing had developed. So, between those yield curves and the cost target, if the chip size was greater than, say, 140-150 maximum, but more like 140 mils…

**Hendrie:** Yeah, they were using mils in those days.

**Faggin:** Mils in those days, yeah. Then it would've been cost-prohibitive.

**Hendrie:** Okay, so it wasn't a cavity on a package or something like that.

**Faggin:** No, no, no. It was the actual chip size. There was nothing…

**Hendrie:** It was limited by yield curves and what the cost they had committed to.

**Faggin:** Yeah, yeah, yeah. The cavity limitation was quite a bit larger than that. But eventually it would've been also a limitation, because we wanted to have those devices packaged in sixteen pins. And so the cavity was relatively small…

**Hendrie:** Okay.

**Faggin:** …But it could stand larger devices than what was possible from cost reasons. One thing that I didn't say is that, another critical element to make the 4004 work, was the use of the [direct] poly to junction contact, what we called the buried contact. Remember? I think I mentioned to you that I invented it at Fairchild, and that technique really was giving me effectively two layers of interconnections, because I could do a lot of interconnection with the silicon layer by having this direct contact between the poly-silicon and the junction. And then metal could come over that, and so I had effectively two layers of interconnection, which in those days was rare, and the people that had it in bipolar, had big cost problems, because it was not yielding very much. So again, without the boot-strap load which allowed dynamic, two-phase dynamic circuits to be effectively done with a reasonable power dissipation, and the buried contact, I could never have met the goals of chip size that were realistic in those days. By the way, the 1103 also used the buried contact [and boot-strap capacitors]. In fact, that was the main difference between the 1102 and the 1103: The 1103 had buried contacts, the 1102 didn't. Of course, the 1103 was much, much smaller than the 1102 [and thus faster], and that's why ultimately it won, it won the race.

**Hendrie:** Yes.

  

**Faggin:** And of course, that was my idea. And I believe I mentioned earlier that I found out a couple of years later that Vadasz patented that idea as his. And that was part of the reason why I left Intel. But, at any rate, I had to put my foot down on that [the requirement to use buried contacts], because Vadasz didn't want me to use that, 'cause he didn't feel "safe" yet. But I told him that I had already checked those at Fairchild, and I had built devices like that, and they worked. And Intel had not done that yet.

**Hendrie:** At this time Intel had not.

**Faggin:** At this time, had not done that, yeah.

**Hendrie:** Okay.

**Faggin:** And basically I told him that without that [buried contact] I couldn't do it; it was impossible to do; nobody could do it. So he reluctantly let me do that, once again. That was the other piece of the critical elements that were required to make the 4004 possible. So, then, we come to September or so, when the layout of the 4004 was still going on in earnest. And the 4001 came out, and it worked first time, and that was a big relief for me, because it was my first complicated chip that I did, and it proved that the methodology that I had used for everything else was going to work. And it was going to be...

**Hendrie:** And the circuits that you…

**Faggin:** And the circuits that I developed worked, and the whole thing was properly working. So the chip came out, it was faster than was required, and it had all the characteristics we wanted, and that gave me, and also Busicom, a boost in confidence that things were progressing well. And then, a few weeks later, the 4003 came out, and that also worked perfectly. And again, one more notch…

**Hendrie:** Yes.

**Faggin:** And by -- let me see, September/October; early October probably -- we had finished the layout of the 4004, and then we had the ruby cutting of it. I know that Shima left around October, and it was before we were finished checking the rubyliths, but close to the completion of the ruby checking of the 4004. He also helped in all the checking phases of the project, because he was very good and very meticulous, and very attentive. And basically, he went back -- yeah, it was sometime in October -- he went back to Japan.

**Hendrie:** Now was he married?

  

**Faggin:** No, no, he was not married.

**Hendrie:** Ah! So this was not a hardship to come to work in the United States.

**Faggin:** No, no, it was not a hardship. No, no, he went back to Japan, because at that point, also, the job where he could help was finished, and he had to go back and begin to do the software development for the calculator, because that was his real job. His real job was to do basically what today we would call the firmware, develop the program for the calculator.

**Hendrie:** Now, I seem to remember reading somewhere that he might have done some of the original logic when he came to Intel, and they thought there were going to be seven different parts.

**Faggin:** Yeah, he did the logic design on those chips.

**Hendrie:** Of all of those chips, yeah.

**Faggin:** Of all those chips, yeah.

**Hendrie:** So he had experience doing logic design.

**Faggin:** He had experience in logic design, absolutely, yeah, he had. That's why I gave him the job to finish it up.

**Hendrie:** Yes.

**Faggin:** The logic of the 4004.

**Hendrie:** It was a logical thing to do.

**Faggin:** It was the right thing to do.

**Hendrie:** Yeah.

**Faggin:** He knew what to do. The key thing was to create the methodology right, because he didn't know how to do that...

**Hendrie:** The style of designing...

**Faggin:** The style of design.

**Hendrie:** With your-- with silicon gate and bootstrap loads.

**Faggin:** And by that time he had learned enough, he had seen enough of how the process was done so that he could follow that method; like understanding how the layout would proceed, so he could do the drawing the way that was very helpful to the layout process. So when...

**Faggin:** So, then, when Shima left, we had seen the 4001 and 4003 working. And then, in November, we received the 4002, and the 4002 also was working completely, except for one minor mistake, which however allowed to check everything, so that was it; they were still usable parts. And then finally, at the end of December, the first run of 4004 came out. By that time, I was close to the schedule that I had drafted very early in the project, and it was absolutely the best that could possibly be done. I mean, remember, from February [actually, April] to December having done four chips of that complexity…

**Hendrie:** Right.

**Faggin:** And at the same time I had to develop the testers -- I didn't talk about it -- I had to develop the test programs, of which Shima also helped, and begin to worry about the production testers. I mean, the whole thing. And so I used to work… my wife fortunately left that summer, went to Italy with our daughter, who was born in April -- in March, I'm sorry, of 1970 -- so she was a newborn. And she left for three months. That was a godsend for me so I could work 12 to 15 hours a day. And, you know...

**Hendrie:** And every day.

**Faggin:** And every day, and just get a lot of work done.

**Hendrie:** Wow.

  

**Faggin:** So, end of December, the 4004 comes out, and the wafers were given to me at the end of the day, probably 6:30 or so, and people were beginning to go home. I was actually happy because nobody would be around as I was, you know, very nervous.

**Hendrie:** Now these were just wafers. You were going to probe them.

**Faggin:** That's right, you probed them.

**Hendrie:** They hadn't been diced yet.

**Faggin:** And I had developed an actual tester, which was using some obsolete equipment that we had in the lab, and then adding a bunch of stuff to create an engineering tester for these chips as well. I had a technician to help me with that. So, I loaded the first wafer, lowered the probes, and… Nothing, not a wiggle. "Okay, well, bad die."

**Hendrie:** Bad die. Next die.

**Faggin:** Next die: same thing. "Well, I'll try another one." Next die: same thing. "Uh, uh… Well… Maybe… a bad wafer." So I picked up another wafer, put the wafer down, tried the first die: same thing. And now I know I'm in deep, deep trouble. But at the same time, if this was the first chip that I had done, I would probably have freaked out, but still I had sweat coming down…

**Hendrie:** Of course!

**Faggin:** But something had to be wrong. It was not possible, you would have some signal, something.

**Hendrie:** Yeah, something.

**Faggin:** Something would show, something; it would be wrong, but it would be something…

**Hendrie:** Right.

**Faggin:** So, finally, I gave up and I took the wafer under the microscope to take a look, and sure enough, one masking step was left out: it was the buried contact layer.

  

**Hendrie:** Oh!

**Faggin:** So, maybe 60 percent of the gates were not connected, right? So there was nothing working, obviously. But that [incident] made me lose about three weeks of a schedule that was already very tight; just a simple screw-up in manufacturing. So, I was really pissed, because up until that point I would've been very close to the schedule that I had promised. But there you go…

**Hendrie:** Now was everything… There was no separate R&D line?

**Faggin:** No, it was all done in production.

**Hendrie:** And Intel had this philosophy of only one fab, so you don't have to move things.

**Faggin:** That's right, yeah, only one fab. But of course, I was using a standard process. The process was not developed explicitly for the 4004. So, it was even not an issue for me. But, anyway, I was really very upset that evening. But hey, c'est la vie! And so the next wafers came out toward the second half of January, I don't remember the date, but about three weeks later, and, again, at the end of the day. I was very nervous. And this time, I put down the probes, and… Eh! The signals are there; well, now I got something to start, right?

**Hendrie:** Yeah, now let's see whether it'll compute anything.

**Faggin:** I had already pre-decided what I needed to do to check the various building blocks and I had some specific test programs that I had developed, and I tried them all, and they all worked. And then I started going beyond that. And basically, by 3:00, three-and-a-half o'clock in the morning, I had checked all the basic things, and everything worked. So then, that really was one of the most elating moments of my life; when you see the culmination of, at that point, ten month worth of work – 10-11 months worth of work -- and see that things are really running the way they're supposed to be.

**Hendrie:** Exactly! That's really gratifying.

**Faggin:** That was it! That was a major, major moment. And I remember going home in a state of excitement, and my wife was waiting for me. Not actually waiting, she woke up, she was in bed, and she woke up and said, "Hi, how's is it doing?" And I said, "It works!" <laughs>

**Hendrie:** Yes!

  

**Faggin:** That was great! That's the essence of the development phase of the 4004. Of course, then, there was still a bunch of work to do to bring the product to production, to do the characterization, to develop the production test programs. We had bought a tester – a production tester for the [4000] family -- that was a so-so kind of a machine, but it would do the job. And so we programmed that.

**Hendrie:** Can you talk a little bit more about the testing, and how you did that initial testing? And then what the approach to testing and manufacturing was?

**Faggin:** Well, the approach was to develop a test sequence. First of all, you would do what is called the parametric test. In other words, you measure opens and shorts and leakage, and break down voltage and so on. At the pins, right, because we're talking about a tester that would probe in all the accessible points, which are the pins of the device: the pads of the die and then the package pins. And so, you do that first phase, and if something doesn't go, you basically reject the die and you go to the next; in a packaged device you reject the packaged device and you load, with automatic loaders, the next device. And then, of course, you go through the functional test, which basically has to check the internal logic, and if the device performs properly. And the philosophy here is to develop a test program where each transistor in the chip is checked, at least once, in each one of the two states: open or closed. Okay? On or off.

**Hendrie:** Yes.

**Faggin:** And although it's always difficult to guarantee that you do that effectively for those devices, those are the rules. So, in developing a test program, we went through all kinds of convolutions to make sure that we had properly tested that [on-off condition]. I also took care of [the requirements of] testing in the design itself, by sending out from the chip, during idle times, internal information, like flags, for example or the content of other internal flip-flops, to be able to test without having to go through very complicated convolutions.

**Hendrie:** Without indirect inference...

**Faggin:** Inference, exactly.

**Hendrie:** ...as to what the flags were.

**Faggin:** Particularly flags, because the flags are set at every arithmetic operation, and if you want to find out if a flag is correct, then you've got to test it with a conditional jump, and send out some signal in some of the pins, and that would take way too long. So, I came up with that idea during the design, [and the 4004] was designed so that it could be tested effectively.

**Hendrie:** Okay, and it required a little bit more logic, but you thought it was probably...

**Faggin:** Yeah, but very little.

**Hendrie:** Very little, but worth it.

**Faggin:** Yeah, it was worth doing because we had a lot of dead cycles in the internal bus, where I could use those internal cycles. There was nothing useful…

**Hendrie:** Yeah, I understand.

**Faggin:** Nothing useful going on, then I would use that time to send things, like internal registers content, out or internal flip-flops, out. So then, obviously, the development of that test pattern was tedious, because it was like designing, like writing a program in machine language, not even assembly language, machine language, right?

**Hendrie:** Where literally, at every cycle of the tester, you had to define what's going to happen on every pin of the tester.

**Faggin:** Absolutely, yeah.

**Hendrie:** Every pin of the tester.

**Faggin:** Absolutely.

**Hendrie:** How many pins on this tester?

**Faggin:** Well, I don't remember, we had 16 pins [on the devices to be tested] -- well, probably 40/42/44, of that order.

**Hendrie:** Yeah, yeah.

**Faggin:** But of course, we only had 16 pins so that was definitely not a challenge. The problem was to develop the test pattern so that you would see all the transistors, internally, in one of those two states. And also, for the internal memory, you wanted to see if there was any pattern sensitivity in the internal

  

RAM, and things of that sort. So, there was some test that would allow to test that pattern sensitivity, etcetera, etcetera.

**Hendrie:** All right.

**Faggin:** So it was sophisticated work. At that time I had...

**Hendrie:** What was the equipment that you were using? That you bought?

**Faggin:** Pacific Western, Model 10 or 30. I forgot...

**Hendrie:** All right.

**Faggin:** But a commercial tester that was okay, but it was not so great. It was all that we could find that was cost effective, because the big testers, like the Sentry tester of Fairchild, were way too expensive, and we didn't even have one of them...

**Hendrie:** Okay.

**Faggin:** Then later on, as a matter of fact, because this tester was too expensive to use for [wafer] sort, I designed and developed a tester myself, in my group, to test the 4004 -- actually the whole family -- but in particular the 4004, which was the most sophisticated one. Basically, a box with logic where I used the 4004 itself to create the pattern. And so it was a comparison test.

**Hendrie:** Ah!

**Faggin:** Where all I needed to do was to write a...

**Hendrie:** You wrote a program.

**Faggin:** A program, an assembly language program, and then the device under test was compared with the 4004 results; that way it was much more effective. In fact, that job was what gave me the conviction that the 4004 was good for other application than calculators. We'll get there, but basically, at that time, people thought that the 4004 -- the whole 4004 family, which later on was called MCS-4 -- was only good for calculator or calculator-like products, and not for other applications. And my position was that it was not true; that it actually was very good for control applications as well, and I used the opportunity to build that tester to show Intel's management that you could do other things. The tester was controlled by the 4004. So there were two 4004's. There was a 4004 that was used for control of the tester itself.

**Hendrie:** And then there was the...

**Faggin:** Plus the 4004 that was producing the signals to be compared to the 4004 under test.

**Hendrie:** And was this used still at wafer probe, or...

**Faggin:** At wafer probe.

**Hendrie:** This is for the functional test.

**Faggin:** And then eventually it was used also for finished devices, because it was good enough for the functional test of finished devices, it was actually better than what we could do with the PW tester, the Pacific Western tester.

**Hendrie:** Okay. Wow. Okay.

**Faggin:** And it was a box about this size [shows size with the hands] with a number of printed circuit boards. Actually they were not printed circuit boards, but they were wired…

**Hendrie:** Just wired. A wire-wrap.

**Faggin:** because, you know, it was one off [one of a kind].

**Hendrie:** Yes, effective.

**Faggin:** It would not have been cost-effective. They were… there was solder, it was one of those… Anyway, it was not wire-wrap.

**Hendrie:** Okay.

**Faggin:** Wire-wrap became popular a few years later, yeah.

  

**Hendrie:** All right.

**Faggin:** At least at Intel. <laughs>

**Faggin:** I should say that in the 4004 there were a couple mistakes that were found after I finished checking it, but they were very simple to diagnose and correct. And by March timeframe we had everything working. And in fact we had received already the codes to put into the ROMs. There were four ROMs that were used in the first product, the first Busicom calculator. So we received those around January/February, and by March we had -- no it was May, sorry, it was May when we had everything [working]. By March we had working 4004s.

**Hendrie:** Was it just one more spin on 4004.

**Faggin:** One more spin on the 4004. We sent the parts to Japan and, one week later, Shima told me that everything was working. He actually had built a simulator in Japan; a simulator of the 4004, with random logic, so that he could develop the software, otherwise he would not have been able to check his software well enough to be able to then commit to ROMs.

**Hendrie:** To ROM, yes.

**Faggin:** That's right, that's right.

**Hendrie:** He needed it to actually test it.

**Faggin:** He needed it to do a little better work. That's right. And so he was able to actually do that with the 4004. Not only had he tested that with the simulator, but also when I sent him [the 4004] in March...

**Hendrie:** He could replace the simulator with a 4004.

**Faggin:** Replaced the simulator with the 4004, and tested the whole system correctly, except he was using a RAM to hold the code, simulating the 4001 ROMs.

**Hendrie:** Oh, yeah, he was going to burn it into ROM. Put it in RAM.

**Faggin:** Of course, yeah. So, he tested the whole thing, and then they sent him the code [the 4001s] at the end of March, or early April. And I know that in May the calculator was completed. It worked with all the chips of the family.  It was an engineering prototype, and it is the very engineering prototype which is in the Computer History Museum.

**Hendrie:** Very good!

**Faggin:**  Which I donated to the Museum.

**Hendrie:**  Thank you!

**Faggin:**  Because the president of Busicom, after the completion of the project, gave it personally to me as a gift, in recognition for my work in making this possible.

**Hendrie:** Oh, very good.

**Faggin:**  So.

**Hendrie:**  That's wonderful.

**Faggin:**  Yeah, so that concludes the first part of the story of the 4000 family.  But in the meantime, as you know, there had been another project going on at Intel -- what was called the 1201, that was the Datapoint CPU -- the first 8-bit CPU. In fact, when I joined Intel, the project for the 1201, which later was called 8008...

**Hendrie:** The marketing cover was 8008.

**Faggin:**  The marketing name was 8008.  The project for the 1201 had started.  Hal Feeney was in charge of the project, and Ted Hoff and Stan Mazor were helping along, but there was not a lot of help there necessary; it was mostly design, because the whole architecture was the Datapoint 2200, intelligent terminal architecture.

**Hendrie:**  Yes.

**Faggin:**  So the architecture was already done; it was Datapoint's. In those days the company was called CTC, Computer Terminal Corporation, which had produced a product called Datapoint 2200. Eventually, the name of the company was changed to Datapoint.  So it became known later on as Datapoint, but in those days it was CTC.  And Vic Poor, as you know, was the VP of Engineering of the company at that point, and he actually, he and his people, were the designers, the architects of that microcomputer that was supposed to be built using TTL. So it was a TTL design, and Datapoint wanted to have it converted into an MOS chip. The actual story is that around… It must have been around October, maybe even September-October ’69, before I joined Intel, so I heard this secondhand, Vic Poor went to Intel and they wanted to have a special [bipolar] chip, a custom chip, to do the stack pointer of the machine in a single chip, because otherwise, they would’ve had to use a lot of TTL gates to do that. And it was a 64-bit memory, essentially, but organized as a first-in/last-out stack pointer. And he went, as all the customers would go with special requirements, to the application group, which was Ted Hoff’s group. Ted and Stan wanted to take a look at the design, and when they saw the design that they were trying to do, they -- with a lot of guts, I must say -- they said, “Well, we could do everything in one chip.”

**Hendrie:** Yeah, not just the stack...

**Faggin:** The whole thing! Yeah, “We could do it all in one chip,” you know? With the strength of already having taken a look, having developed the architecture of the 4004.

**Hendrie:** Yeah, but 4004 was anything but done.

**Faggin:** That’s right, yeah, that’s right. This looked like a little more complicated than the 4004, but not that much more complicated. So they took the gutsy move of saying, “Well, yeah, we can do that.” So out of that came an engagement between the two companies, and out of that came the hiring of Hal Feeney to lead the project. And Hal had developed some random logic circuits at a company before -- I forgot where he came from, probably General Instrument, yeah, they had MOS.

**Hendrie:** Yes, I believe that’s correct.

**Faggin:** I think that’s where he came from.

**Hendrie:** Because he had done MOS.

**Faggin:** He had done MOS before, but not silicon gate, because of course.

**Hendrie:** Nobody had.

**Faggin:** It was only Intel and Fairchild in those days. So he had this project. And in fact, when I came to Intel and I found out what was going on, I was disappointed because it was clear to me that the first microprocessor would have been the 1201. Because I had four chips to do, and the 4004 would’ve been the last one, and he had a head start and the architecture was already done. And so I was disappointed, but hey, c'est la vie….

**Hendrie:** So you sort of wished you, maybe…

**Faggin:** Oh yes, I wished I had that project, as opposed to what I had but, of course, that was not the situation.

**Hendrie:** It's not your choice.

**Faggin:** That's right, it was not my choice. And after awhile I forgot about it. I was so busy doing my own thing that, clearly, I didn't care anymore. But then, as it turned out, the 1201 had some difficulty getting started. Clearly, the challenge was enormous for Hal [Feeney] to figure out how to do it, and the bottom line was that Hal was given another project, a memory project, that came along in the meantime: it was a 512-bit, static RAM.

**Hendrie:** MOS static.

**Faggin:** MOS static, for a customer. It was twice the size of the 1101, in terms of [complexity]. But that project didn't go anywhere because the customer disappeared. So, basically the [1201] project was put on ice.

**Hendrie:** They did the same thing to that project they did to the Busicom project.

**Faggin:** Because, first of all, Ted Hoff and Stan Mazor could not give proper guidance, nor could Vadasz give guidance to Hal; it probably was way too complex. And basically, that was it, so it languished, and then it was stopped. Also, sometime later, after it was put on ice, CTC decided to proceed with the TTL implementation anyway. So there was less of an urgency, because the window of being done with MOS was…

**Hendrie:** They weren't calling up every week, as to where it is.

**Faggin:** That's right. Yes. So the whole thing kind of cooled off, and that was the end of the project for awhile.

**Hendrie:** Now when was it put on ice, do you remember? You weren't working on it?

  

**Faggin:** Probably April [correction: July] timeframe. So it was put on ice just around the time that I was doing the logic design of the 4004, around that time. So we had not seen the whole design on the 4004 yet. But, of course, I was very clear to Vadasz, in my weekly reports, that that was not a trivial task, to do it right. So then in January [1971], after the 4004 worked, now there was an interest in something like the 8008, by a Japanese company, Seiko Precision, which in those days they wanted to -- it was called Seikosha -- they wanted to make a desktop scientific calculator, similar to the HP… Remember the HP, was it 9100?

**Hendrie:** Yes.

**Faggin:** Yes. And they wanted to use the 8008 as the basic engine of this machine, because at that point Datapoint was no longer interested in the [chip]. By the way, Datapoint had also commissioned Texas Instrument to do a single chip design of the 1201/8008, which eventually was done, and it was announced, I think, between March and May. I would have to go take a look at some of my notes, but probably it was between March and May of 1971. So it was actually announced as a CPU-on-a-chip, before there was any announcement of the 4004 -- public announcement. Because the 4004 was a custom design, there was no public announcement whatsoever. Although it was after we presented a paper on an IEEE conference, down South, in L.A., and definitely after the 4004 worked, because the 4004 worked well in January -- in March it was fully working. But, it was really a PR stunt: that chip never worked, and it was never used, never produced, and that was the end of it. But there was a patent done, and they [TI] basically patented the architecture of Datapoint… Anyway, this is TI, and there's been a big legal stuff, ever since, around that whole thing. At any rate, it's interesting how this story drags on. But comes January, with the interest of Seiko, the project is resumed. At that point, Hal Feeney was working for me; he helped me, toward the tail-end of the MCS-4, the 4000 family, with the test patterns for the production. So he was helping me doing that. And then, when the go-ahead came again, for the 8008, then, under my direction, the 8008 was restarted.

**Hendrie:** Okay. And this would be approximately when?

**Faggin:** January, 1971.

**Hendrie:** Ah, okay. So there was a whole six months. Now what were you doing?

**Faggin:** No, it was all much more than six months, because it was stopped around April, May, a year before [correction: July 1970]. But there was not much work done beyond, basically, the logic equations for the control, things like that, which Hal Feeney did properly. But no circuit was ever done, no circuit design was ever done.

**Hendrie:** Okay. There was no circuit diagram.

  

**Faggin:** No, no circuit diagram, no. Basically it was a product definition, with a little more of the basic control logic -- basically the control equation, the control logic equation -- of the control [section] for the 8008.

**Hendrie:** Now what were you doing between when the whole chipset appears to work, in May, and all the way around, till 1971. What were you doing in the rest of 1970?

**Faggin:** All the 4001, 2 and 3 and 4. Right?

**Hendrie:** Okay.

**Faggin:** That was the year where the 4000 family was designed. 1970, right?

**Hendrie:** Oh, yes. Okay. 1970 was when….

**Faggin:** From February [correction: April] 1970 to March 1971, that was the time when I was almost 100% occupied by the 4000 family

**Hendrie:** Okay. But, when Hal Feeney started back up on this project in early ’71, in January or February.

**Faggin:** Yes, he started in January, yes.

**Hendrie:** In January, you become….

**Faggin:** He started in January ’70, with the project, [correction: March 1970] but then it was stopped.

**Hendrie:** And then it began again in…

**Faggin:** And then it began again in January, a year later, in January ’71, under my direction.

**Hendrie:** Under your direction, even though you were spending a lot of time, obviously, on the 4000.

**Faggin:** Yes. Yes, but there were….

  

**Hendrie:** You could do that.

**Faggin:** I could do that. And basically, Hal did not take a lot of direction after I explained the methodology that I used, how to do it, and all that. After that, he did it by himself. The key thing was to figure how the methodology, how do you do this stuff? Okay?

**Hendrie:** Yes. Exactly. What are the circuit styles you are going to use….

**Faggin:** They key circuits you're going to use, the key techniques that you're going to use. How do you organize it? So I helped him for the first few months, where we worked closely together and I showed him how to do it, and so on and so forth. But then after that, there was regular supervision, of course: we would have regular weekly meetings and go through the issues and problems, or if there was a problem, he would come to my office and discuss it; that kind of stuff, like that. But it was his project and he brought it to completion, by himself. And then there was only one major problem, after he went to marketing, where there was a problem in the memory of the 1201, and it was my job to actually find the problem, debug it, and solve it and so on. But that's fine.

**Hendrie:** What was that problem?

**Faggin:** Basically, the memory was losing its memory. And it was similar to…

**Hendrie:** This had been designed as a dynamic 3 transistor chip?

**Faggin:** Yes. That's right. Yes. And this is the same problem, actually, that I found on the 4002, and the 4004 and that I corrected, between the time that I found it -- which was as I was characterizing the parts -- and the time that they went into production. So that did not have any impact on the schedule because people were just getting ready for production and so on, but I had to change a bunch of things. That was a problem that was found also in the 1103, that required to back-bias the substrate, and it forced Intel to go from a 16-pin 1103 to an 18-pin 1103, which was like the sky had dropped from heaven. I never seen so [many] long faces at Intel, over this issue, because it was a religion in Intel; everything had to be 16 pins, in those days. Everything had to be 16 pins. Okay?

**Hendrie:** Okay.

**Faggin:** And even the original 1201 was 16 pins, with incredible sacrifices of function, because we were losing a lot of the advantages of silicon gate by having to multiplex [many signals on a few pins], because of this requirement.  It was a completely silly requirements to have 16 pins.  It was God-given 16 pins.  I don't know what….

**Hendrie:** Where there no 18-pin packages?

**Faggin:** But people were using 40 pins and 48 pins already!  And all the calculator chips that were used in high volume production were 40 pins.  So it was a religion.  And, in fact, one of my major comments, early on, even on the 4004, [was that] we were throwing away a lot of performance in the architecture because of the insistence of having 16 pins.  It was crazy.  And of course, for the 8008, it was even worse because you're forced to have a lot of interface devices between the CPU and the memory, because you had to de-multiplex, to put latches, to do all this kind of stuff.  So the actual 8008 architecture was not a very good one -- not the CPU architecture -- but the external architecture: the bus structure and so on was a contraption.  And, in fact, it contributed to its demise; the 8008 never went that far, and it was replaced by the 8080, which was a much better design: my idea and what I wanted to do right the first time around, but I couldn't do it.  Anyway, back to the beginning of the 1201. The design, as I said, started in January.  The actual circuit design was done in a few months, and then the layout started, probably around April, or so, that year.

**Hendrie:** Well now, the 1201, originally, the way CTC was designing it, their original concept, it was a serial machine.  Were you aware of that?  Well, I found that out from Victor Poor.

**Faggin:** From Vic Poor?

**Hendrie:** That it was originally designed to work with an MOS…

**Faggin:** With shift registers?

**Hendrie:** …serial, with shift register memory.

**Faggin:** Yes.

**Hendrie:** Was what it was going to… In fact, I think it was bought from… Intel supplied the serial memory.

**Faggin:** Yes.

  

**Hendrie:** And it was going to be a bit serial machine.

**Faggin:** Yes.

**Hendrie:** Originally.

**Faggin:** Yes.

**Hendrie:** And then I think Stan Mazor may have convinced them to go…

**Faggin:** Yes. To use dynamic RAM.

**Hendrie:** …to make it a parallel machine, fundamentally.

**Faggin:** Yes. No, I wasn't aware of that. Because the technology, in those days, for terminals, was, really, serial.

**Hendrie:** Yes, of course.

**Faggin:** They were shift registers. And, in fact, the only high volume products, in 1970, when I joined Intel, were shift register; it was not the 1101. The 1101 was available in the marketplace, but people were not using it, in any quantity. And so Intel quickly designed a family of shift registers that were pin compatible with a National Semiconductor line of shift registers. And because of Silicon Gate Technology, they were better shift registers: they had a better retention time because with silicon gate, leakage is much less than with metal gate, because you can do great gettering, phosphorus gettering, reduce the impurities, after the whole structure is sealed. And so, they basically took the market away from National Semiconductor. So, the real hope for RAMs was pinned on the 1103, because the 1101 never really made much hay, and the density requirement and speed requirement were not met -- also the power requirement -- by that chip. And only the 1103 was beginning to promise to catch up with core memories, and also with shift registers. Shift registers had 6 transistors per cell, just like static RAMs, even in dynamic shift registers, where dynamic memory had 3. And, of course, in those days, people were talking about one-transistor and one-capacitor dynamic cell, which eventually…

**Hendrie:** Eventually they did figure out how to do it.

  

**Faggin:** …ruled the world, from the 4K on. So anyway, parenthesis closed on that, it's not surprising that the original architecture would have been using shift registers -- for the [Datapoint] 2200 -- although I wasn't aware that Intel convinced Vic Poor to actually use a.…

**Hendrie:** A parallel.

**Faggin:** Yes, RAM, as opposed to serial.…

**Hendrie:** A parallel approach.

**Faggin:** Yes. Which really meant, from a CPU architecture, very little.

**Hendrie:** Yes, exactly.

**Faggin:** But, certainly, from a system point of view, it was a big difference.

**Hendrie:** Oh yeah, big deal.

**Faggin:** But, see, that was the same thing again: It was re-using what was learned from the 4004, because that was the same thing that Busicom had. See, Busicom also had a CPU [design]. So it wasn't that nobody was thinking CPU's in those days.

**Hendrie:** Yes, exactly.

**Faggin:** They had a CPU, but it was designed around serial memory. And it also was using more than one chip, because it was unnecessarily complicated. Calculators, in those days, also used serial memory, because, obviously, it makes sense. You do bit serial operations as the data shifts out, and it makes sense. And, in fact, the key suggestion of Intel to Busicom was to use dynamic RAM, because they only had 3 transistors, instead of 6 transistors for the shift register.

**Hendrie:** Exactly.

**Faggin:** Therefore, they were even more practical, more cost-effective. [Dynamic RAMs] They're not only better architecturally, but also more cost-effective than shift registers. And that was a major contribution that Ted Hoff had at that time, to basically articulate that [idea] to Busicom and get Busicom to buy it.

  

**Hendrie:** Okay.  So back to the… I think we got sidetracked a little bit on the 8008.

**Faggin:**  That's all right.

**Hendrie:** It's alright to weave around.

**Faggin:**  To weave around, right.  So, basically we get to….

**Hendrie:**  Now you're supervising Hal, and he's working on the 8008.

**Faggin:**  So we got to layout around the late spring, and summer, and we actually saw first silicon out, of the 8008, toward the end of '71.  And there were a few minor mistakes even there.  And then Feeney went to marketing, in early 1972, and so we ended up, in my group, finishing the 1201 -- by that time called 8008 -- which was then introduced in the market -- I think it was April of 1972 -- and just [made] available [commercially].

**Hendrie:** Now, did you hire another engineer to finish it up, or somebody you already had worked with?

**Faggin:**  No, I already had people, at that point, I had a number of people.  I had a… Young, is his last name.  He's a Chinese engineer.  I forgot his first name. [He was Young Feng]. And Young came in toward the very end of the 4004 [project], again helping with bringing it to production; test patterns again, because there was a lot of [work]. I explained how we were making those test patterns: it was very laborious.  And, of course, it took an engineer to do it, a junior engineer, but still it took an engineer, because it was not something that you could have people who don't understand….

**Hendrie:** They have to be able to go through the logic and think about…

**Faggin:**  And understand it, that's right.

**Hendrie:**  …if I do this, this transistor will wiggle.

**Faggin:**  This will happen.  Exactly right.  So, I need to go back because, at that time, I was frustrated because all this work, it was for a custom chip, a chipset, for a Japanese customer. That was it, right?

**Hendrie:** Right.

  

**Faggin:** And I really believed that the way to go was to have microprocessors out there. It was a major innovation in the industry. So did Feeney, so did Ted, and Mazor, for that matter. But, in detail, we were differing. I was convinced that the 4000 family was actually a worthy product to be in the market -- not just the 8008. There was like a camp, where I was the only one to support the 4000 family as a product worthy of being out in the market, where internally everybody thought that it was too limited of a micro controller -- microprocessor -- and that the 8008, on the other hand, would be appropriate to be out in the market. So I took it upon myself to convince management that they should go [out] and announce the product, and market the 4000 family. And there were two problems: One was the belief of its limitations, that I needed to overcome. The second was that it was an exclusive custom project for Busicom. So Busicom would have to give permission to do this. The first objection I overcame by designing the tester that I mentioned earlier, where I used the 4004 [as a controller]. I used Frohman's 1702's -- they were the first EPROM devices, so that they would emulate the 4001-- so that I could build even a one-off system, and I could develop the system with that, and also show that you could do that so that the customer could do the same thing. That was the basis for then figuring out how you market this stuff. I built the tester; it worked; I showed people how to do it. I showed people that this is an application where you'd use the 4004 as a controller, and it is very effective in doing that, and I convinced them. And then, as far as Busicom, through my direct connection with Shima, I found out that Busicom was in trouble – financially -- that they were not that competitive with their calculator, because Intel was charging too much. So I came up with the thought that if I convinced Noyce to lower the price, in exchange for the exclusivity, then we could do that. And that's what I suggested to Bob Noyce. Bob Noyce obviously had to agree, because that's what he did a few weeks later. I had found out he was going to Japan and I met him on the hallway and I said, "Bob, I heard from Shima that they are having trouble. But if you want to lower the price, I'm sure that they might rescind the exclusivity, and then we can announce the product." And by that time, since I had already solved the problem of usability, all the pieces were there. And then, a few weeks later -- actually, probably a month later -- after the trip of Noyce, there was an agreement that they were going to do that. And so, by summer timeframe, the go-ahead of going public with the 4000 family came. And that's when a [new] marketing group was formed, within Intel. And Hank Smith was the manager of marketing for microprocessors, and both Hal and myself, and also Ted and Mazor, worked very closely with the marketing group, to really help out, map out how to go and market the stuff, and so on and so forth. So, that was what happened. And by November, 1971, that's when the formal announcement, the now famous ad: "A new era in microelectronics has arrived," or something like that, [announcing] a computer-on-a-chip, blah, blah, blah. So that was the ad that they….

**Hendrie:** That they put together.

**Faggin:** Yes, they put together, and that set the world in a new direction.

**Hendrie:** Yes. That's great.

**Faggin:** Yes. And then a few months later, so that's November… April of the following year -- '72 -- the 8008 was announced. And so now, the world's first 4-bit and the world's first 8-bit microprocessor, were Intel's doing. In the meantime, during the summer -- late summer of '71, it must have been probably August, September timeframe -- I went to Europe with Hank Smith. So that was just the beginning of the…

**Hendrie:** The marketing push.

**Faggin:** The marketing… yes, well, it was not even a marketing push. It was really fact-finding, right?

**Hendrie:** Yes.

**Faggin:** Because, at that time, we had not announced anything. So we were going around, under non-disclosure agreement, to some key customers, to investigate their interest in these products. So it was done also to accumulate marketing intelligence to decide whether it was the right thing to do to announce this thing, was there enough of a market, and so on. And so, part of that fact-finding mission, I went to Europe with Hank, and we visited a bunch of customers, and we gave presentations on both the 4004 and the 8008 to a bunch of customers. And I remember vividly that the more a customer knew about computers, the more hostile they were on our products. Whereas people that had a problem to solve -- like we went to a company that wanted to make blood analyzers, and they wanted a controller to control the blood analyzer machines -- they were as happy as they could be because they finally had a solution to a problem. Where the people like Nixdorf -- I went to Nixdorf Computer -- they crucified me, because the 8008 was not fast enough, and it didn't have this, it didn't have that… You call this a computer?

**Hendrie:** It wasn't a real computer.

**Faggin:** It wasn't a real computer. Okay? And I remember that it was very frustrating, but also very useful, because I got a lot of input and a lot of information. And out of that trip came the idea of the 8080. Yes. And so, by having inputs directly from the customer, and knowing what could or could not be done, that gave me enough [information] to then propose the 8080.

**Hendrie:** Here are the things we ought to do.

**Faggin:** Yes. And I came up with the architecture of the 8080, and I proposed the 8080 soon after I went to do that trip. That was even before the 8008 was finished. But then, Vadasz didn't want to hear about it. It's a sad story because we almost lost the leadership in microprocessors because of the delay of Intel to respond to my urgings. And basically, from when I proposed the 8080 to when we agreed to go do it, it was at least 6 months or 7 months. So we basically lost 6 or 7 months of lead time. And when the 8080 came out, 6 months later, the 6800 of Motorola came out. So we really reduced our leadership that was a year, more than a year -- which was, in those days, the cycle time to make a chip -- we reduced that, we cut that in half, because the response, internally, was, "Oh, no, no, wait. Let's see how the market responds, and blah-blah-blah." And that was frustrating. Also, at that time, the key thing was to use a 40-pin package and really make a much more rational I/O and bus structure than was there before. Also using N-channel technology that, at that point, had become available at Intel, so the combination of N-channel technology, plus parallel operation, instead of a multiplexed operation, would give us a [speed] advantage of at least a factor of 5 over the 8008; plus much more usability, because instead of having to use about 30 TTL packages to interface the 8008 to memories, and therefore losing all the advantage of a single chip implementation, it would take very little logic -- external logic -- to actually interface.

**Hendrie:** Okay. So had N-channel been developed by Intel?

**Faggin:** Yes. N-channel had been developed for the 4 kilobit dynamic memory.

**Hendrie:** Oh. So other people had N-channel for MOS.

**Faggin:** Yes. Other people had N-channel. N-channel was the obvious next step. And Intel had just developed it. In fact, I don't think there was even a product out yet, when I came up with the 8080 architecture, but it was….

**Hendrie:** But you knew that it would be there.

**Faggin:** But it was coming. That's right. So it was the right thing to do. I was taking advantage of what was coming, and happening. And it was high voltage N-channel: it was 12 volt.

**Hendrie:** Oh. They hadn't done?

**Faggin:** No. They hadn't done the 5 volt yet. The 5 volt came in a bit later. And there is a story around that too, because the first implementation of 5 volt, N-channel was not good enough. Then I got involved, later on, fixing that, and that's when I recommended to use depletion load -- in the transistors -- that would improve dramatically the performance. And that's when the 2102A, which was a 1000-bit static RAM, was redesigned, using this new technology that I pushed to develop. It doesn't take a lot of work to do -- it's just a mindset, the right mindset. So, with that then, we went from 1 microsecond access time to 200 nanosecond access time. It was just….

**Hendrie:** That is dramatic.

  

**Faggin:** Night and day. And that, of course, was the new technology that I used for the Z-80. Because, it's like anything else, any step of the way you want to have new technology, ideally, and with smaller size transistors, and faster transistors. So, first it was P-channel, [then] N-channel, then C-MOS, then scale, scale, scale…

**Hendrie:** Let's see. Where were we?

**Faggin:** 8080.

**Hendrie:** 8080. Exactly.

**Faggin:** Yes, the genesis of the 8080.

**Hendrie:** And you were talking about your trip to Europe.

**Faggin:** The trip to Europe, yes.

**Hendrie:** Do you remember some of the specific things that you heard there, that sort of got you thinking?

**Faggin:** Yes, well. One of the major criticisms was the very wimpy interrupt structure of the 8008, for example; I remember clearly that. And, in fact, it was barely usable -- part of that also because of the limitations of the pins that we had. But, also, architecturally, it was not done right. And the speed was a big issue, and the fact that the interfacing was so clumsy with the 8008. Of course, that much I knew already, I didn't need them to tell me that. And there were some issues… I remember some issues with… yes, they had some problem with the fact that, for example, they didn't know when they had overwritten the stack, or things of that sort.

**Hendrie:** Some subtle but important things.

**Faggin:** Subtle but important things, because in the 8008, we only had a four-level stack. So that's when I went to 8 [levels], which would drastically reduce the issues, and because to actually check for that [overflow], would have been kind of difficult. With a much more sophisticated architecture, with multiple level interrupts, we could have done an internal trap; an interrupt on overflow. But that would have been costly. And so I did not want to do it that way, on the 8080. But, by doubling the stack, that took care of a lot of that [issue]. And things of that sort… Now, I don't… it's been a long time.

  

**Hendrie:** It's hard to remember those. I think you do very well, frankly, in terms of remembering some of these details.

**Faggin:** Yes. But anyway, a lot of the complaints, though, were really not real. They were complaints that were indicating their frustration with the fact that the semiconductor industry was getting into their turf. So, it was more like -- the innuendo was -- you have no business designing CPUs. You are a chip maker. Make chips! We'll buy from you, and don't screw with our business! That kind of….

**Hendrie:** Yes, coming from there, yes.

**Faggin:** Yes. That kind of emotion behind -- of course not said -- behind those words. And so there were a lot of criticisms. It was really bitching. But there was no substance in that.

**Hendrie:** Yes. But you'd find those more from people like Nixdorf.

**Faggin:** Of course.

**Hendrie:** Than an end user, who doesn't….

**Faggin:** Oh, absolutely. Yes. That's right.

**Hendrie:** Isn't feeling those feelings.

**Faggin:** Yes, Nixdorf and… what was that company, the English company?

**Hendrie:** ICL

**Faggin:** ICL, yes. Those were the guys. They were, oh, oh… So also Plessey… They were making computers in those days, or something… The Plessey Computer Group, or something. Those were the guys that really gave me a hard time. But, it was good.

**Hendrie:** Okay. Yes.

**Faggin:** At the end of the day, that was also used to make lemonade… A good lemonade. So, then, I developed the architecture of the 8080 toward the end of '71, and sent a memo to Vadasz, urging him to start the project.  And, as I said, nothing happened until April, May, a year later. So, we lost at least six months in the momentum that we had [originally] gained in the marketplace, because of that delay. And all the delay was because the [main] preoccupation of Intel was memories. The microprocessor was a sort of a, “well, let’s see if it works first, and let’s try it out.” There was no commitment, really, behind it, no real commitment behind it. They were saying, “Oh, let’s figure it out, see if it works.” And the bottom line that was there for a lot of time -- many years, even after I left -- was, “Oh, we make microprocessors so that we can sell more memories.” That was the punch line of their strategy. There was not really an understanding that microprocessors were actually changing the industry. And it wasn’t until the microprocessor was firmly entrenched -- so, in the early ‘80’s -- before it was clear that something fundamentally different had happened in the marketplace.

**Hendrie:** That this was really….

**Faggin:**  Really.

**Hendrie:**…and you could bet the company on this one.

**Faggin:**  That’s right. I saw it coming, and that’s why my leaving, in the end 1974 -- leaving Intel -- was because I wanted to make a company dedicated to microprocessors, and micro controllers, which I saw as a major business opportunity. So I saw that [market] as something that was happening, and I wanted to take advantage of it. So, back to Intel. Part of the job that I had was also to do some custom circuits. But, in that area, Intel was never really serious. They were sort of playing, without a real commitment. It was almost like a filler. So when it’s slow times, “let’s do some custom circuits.” Then, when the time gets tough and you don’t have enough capacity, “forget about custom.” That kind of thing. Of course, you cannot play the business that way. You either commit to do it, or you don’t. But you cannot go half way.

**Hendrie:** Especially changing your mind, while you’re doing somebody’s circuit.

**Faggin:**  Yes. But one of the reasons why the 8080 was not done in time, was also because I got busy with a custom circuit. It was a chip for Mars Money. And it was one of those things: Friends of Gordon [Moore], or friends of Bob [Noyce], would go in and say, “hey, how about doing something for us? And blah, blah, blah… ” Then they would convince Gordon, let’s say -- Gordon Moore, right -- to do something like this, and then Gordon would come down to Les [Vadasz], “let’s go do this.” And that was it. And then I was stuck doing this stuff.

**Hendrie:** What was this?  What was the company?

  

**Faggin:** Mars Money, is the company that does the M&M's. Mars Money was a subsidiary of Mars Company that does the M&M's. Sorry, Mars Money doesn't do the M&Ms, the Mars Company does the M&Ms. And Mars Money was a subsidiary that was developing coin recognizers, basically machines to recognize change in a vending machine. And they had contracted a consultant to develop an architecture, and a basic technology for that. And so there was a conceptual design, but it was never even tried in practice, other than the basic sensor technology, which was developed completely. The conceptual design was that, basically, you have a ramp, you throw a coin, and as the coin goes down this ramp, there are some electromagnetic sensors that, as the coin comes down, they change frequency. Basically, they are resonating circuits and the coin….

**Hendrie:** Yes. Isn't it a classic metal detector?

**Faggin:** And have enough… Yes.

**Hendrie:** In the view of, it changes the inductance or something, in the circuit, and there's a frequency change.

**Faggin:** That's right. They change enough of the inductance to affect the frequency, so you would have a shift in frequency that is characteristic of the coin that you send down. And, of course, coins are multi-metal. It is not a single metal. It's like laminated metals… I don't know what they are called, but there is a core of one metal and then there is a skirt of a different metal. And so there is a very unique signature, except the signature is dependent on changes from machine to machine. So you had to store values, that are taken at the time of the testing of the system, because they are different from machine to machine: They depend on the sensor, how the sensor is mounted, and so on. So it is something that you cannot predict ahead of time because the window is narrow enough that you cannot [otherwise discriminate] -- if you put too broad a window, then slugs can be used to pretend-money, pretend-coins -- but if you do it this way, then that frequency pull is characteristic of actual money.

**Hendrie:** So the idea was to, rather than adjusting the machine for a constant signature, they would just -- however it came out -- then they'd record what it was and then they'd use that. That required a computer.

**Faggin:** That's right. That required a computer. And they came to Intel because Intel was the only company that had non-volatile memory, and then you could store those values in non-volatile memory. And they [also] wanted to do a single chip. And so it was a very sophisticated problem. Of course, we told them that the right way to do it was to use a microprocessor, because really it was a simple enough problem that it could be done with a microprocessor, but then it would require an external, non-volatile memory, to store these values. And that would have been too expensive. So, basically, they convinced

  

Intel management to go ahead and do this. And, of course, then I was stuck with another project… (laughs)… which then consumed me because any project is a project.

**Hendrie:** You're going to make it work.

**Faggin:** I'm going to make it work. And, of course, nobody had ever put non-volatile memory, together with random logic, and the issues of testing and the issues of doing it right, and so on and so forth, [are not simple]. You have high voltages flying around. It's not a trivial thing to do because you've got to then program this device on the board itself. This was not what was done before; it was done in the programmer, not in the board itself.

**Hendrie:** Of course. Right. Exactly. Not in the operational board.

**Faggin:** That's right.

**Hendrie:** Lots of complexity.

**Faggin:** So, lots of complexity. And, again, that was actually a very difficult project to do because it also had to be done at very low cost. And I put a lot of hours, and a lot of work on that, and it eventually worked and it went into production. But it took our attention away from what we were supposed to do, which was [microprocessors]. At that time I also had come up with the idea of the 4040, which is an improved 4004 microprocessor, which eventually we did, but much too late. And so we actually were losing momentum, until April, May, when sufficient market response came that microprocessors were really [accepted]… They were announced in November, so, 5-6 months later there was enough response by the market that management felt that it was a good thing to do to proceed with the 8080. And so I got permission to start the 8080. But by that time we lost a good 6 months. It would have taken only a few tens of thousands of dollars to buy a six-months market window. Anyway, still, we did not lose the game, in the end, but we let the gap be narrowed a bit more. So, with the 8080, what I decided to do was to hire Shima, because of all the people that I had worked in the past, I felt he was the one that would require a minimum amount of supervision on my end, and he would deliver the goods. So it took a little while because, obviously, he is a Japanese citizen, and it was not easy to get him out. But, eventually, he came aboard. And I had already done a lot of work with him, in preparation of having the permission to proceed. I had already talked to him, I had already convinced him to come. I had already done enough preparatory work that when I was told to go ahead with the project, then it was only a matter of a few months to get him to come.

**Hendrie:** But now this is poaching an employee of a customer.

  

**Faggin:** He had already left Busicom.

**Hendrie:** Oh, alright.

**Faggin:** Okay. He had already left Busicom, he was working at Ricoh. But still, in Japan it's difficult. So, Bob Noyce had to ask permission to the CEO of Ricoh; it's all formal stuff anyway. Obviously, Shima wanted to come and that was the end of it.

**Hendrie:** It was going to happen.

**Faggin:** Yes, it was going to happen. But this way, it was done properly and politely, and so on. So, Shima came a few months later. In the meantime, I completed the architecture and did all the preparatory work, because it was to use N-channel. So there were some changes to be made, adjustments to be made on the methodology that we were using for P-channel. So, many circuits [needed] to be re-characterized for N-channel. We were still using the [same] methodology of [the 4004]

**Hendrie:** So you needed to get new graphs, new tables….

**Faggin:** So I needed new graphs, new stuff. So I prepared all that stuff before Shima came on board, so that Shima could hit the road running. And then I worked very closely with him, again for a few months, because he had never designed a circuit before. He never designed a chip. He'd helped me, but he was really… he was not….

**Hendrie:** He'd done logic, but that's different from doing transistors.

**Faggin:** That's exactly right. Yes, that's right. So I taught him. And he's a smart guy he quickly learned. And within a few months again -- like I did with Feeney -- he was completely self-sufficient. And he proceeded pretty much on his own, again with a little supervision from me. And the 8008 [I meant 8080] it was about a year and a half design cycle, because the chips were getting more and more complex, and the methodology of developing these chips had not changed: We were still cutting rubylith, we were still doing all this stuff; and the chips were getting more complex and the 8080 was already more than twice the transistors of a 4004.

**Hendrie:** Do you remember what the 8080?

**Faggin:** Transistor wise, it was probably 4500, 5000 transistors.

**Hendrie:** And the 8008?

**Faggin:** The 8008 was 3,000, a little more… just about that. And the 4004 was 2200-2300 transistors, around there. So, then the Z80 was about 10,000. So it was almost doubled every generation. The 8008 was really half a generation: It was P-channel, same design rules of the 4004, just 8-bits. In terms of internal complexity and instruction set, it was about like the complexity of a 4004, the same number of registers, just a little bit bigger word for the address.

**Hendrie:** Yes. So longer… Wider internal buses. That's about it.

**Faggin:** Yes. 8-bit bus, instead of 4. But, basically….

**Hendrie:** That's not a lot.

**Faggin:** Yes, not a lot. But generation to generation, we would double the number of transistors. So, 2200 for the 4004, about 4500 to 5000 for the 8080, and about 10,000 for the Z80, and then about 20,000 for the Z8000: it was the following generation. It was just like that. So we get to the middle of '72 [when the 8080 project started], and by the end of '73 we had first silicon out of the 8080. And, again, it worked pretty much completely, except for a few mistakes.

**Hendrie:** Now did you make any changes in the instruction set, besides the stack…?

**Faggin:** Oh, yes. The 8080 had many more instructions than the 8008. It was cleaned up. The interrupt structure was cleaned up, the bus structure was cleaned up, more resources, more registers, twice as big a stack pointer, etcetera, etcetera. So a lot of improvements.

**Hendrie:** Now, did you work out all of those?

**Faggin:** Yes, I worked out all of those, except for a few instructions -- a couple of instructions, one or two, that Stan Mazor suggested, and that's about it. But it was essentially my architecture.

**Hendrie:** Okay. And this wasn't done in terms of any sort of conjunction with CTC?

**Faggin:** No. This was done completely independently.

**Hendrie:** Because they all went off and made a second generation of things.

**Faggin:** But I recommended to use the 8008 instruction set to be machine code compatible for the 8080, because it made sense. It was the right thing to do, and to protect the software investment that our customers had made.

**Hendrie:** Exactly, of course.

**Faggin:** And that was the right thing to do.

**Hendrie:** And it proved to be absolutely the right thing.

**Faggin:** And it proved to be the right thing. That's right. And I repeated that on the Z80 at Zilog. Yes. So, actually, the Z80 architecture, which I also did, was an even bigger step up from the 8080 than the 8080 was from the 8008. So in the Z80 architecture we had more than twice the number of instructions, and all kinds of 16-bit modes, as well as 8-bit, and all kinds of bit operations, from bit-set on any bit of the word…

**Hendrie:** All sorts of things like that.

**Faggin:** All sorts of things, yes.

**Hendrie:** A whole bunch of sophisticated….

**Faggin:** Many more addressing modes and so on. So the Z80 was really a much better architecture than the 8080. But, of course, we learn by doing. And, in fact, the Z80 is still today in volume production -- 29 years later.

**Hendrie:** Wow. It was a great part.

**Faggin:** Yes. And a lot of people learned microprocessors by studying Z80's at school. At least a couple of generations of students did that.

**Hendrie:** Buy one…

**Federico Faggin:** Of course.

**Hendrie:** …for a buck 29, eventually.

  

**Faggin:** Eventually. Yes. Of course, now it's probably embedded, because they are mostly embedded. It's a few cents anyways.

**Hendrie:** Well let's go back though to the 8080 -- we're in the 8080 story.

**Faggin:** Yes. So the 8080 was completed at the end of '73 and it was announced in the market… It must have been March, April, of '74.

**Hendrie:** Now, did you get involved in any of the applications or trying to promote this, or was that more the marketing group, people that had moved into the marketing group?

**Faggin:** No. At that point marketing had become pretty self-sufficient.

**Hendrie:** Pretty effective.

**Faggin:** Yes. Occasionally I would take a customer….

**Hendrie:** Well, they had a chip designer in it, if nothing else, in Hal Feeney.

**Faggin:** …customer visits and so on. But, fundamentally, marketing was self-sufficient. So, after the initial kick-off -- that was toward the end of '71-- I got involved less and less with customer activity. But I still would take a number of trips a year to customers; but it was more to looking ahead than to support the existing products.

**Hendrie:** Now, when you hired Shima, was this a point at which you stopped doing designs pretty much yourself?

**Faggin:** Pretty close. Within a few years, my career of designer, where I would do the design, or I would be a de-facto project leader ended. And, by early '74 or late '73, there was a reorganization, and I took over most of R&D. And so I had all of the MOS design. Except for dynamic memories, all of those designs were mine. And I had all the engineering services: all the layout, all the drafting, draftsmen and so on, all that part was also a part of my department. So I had a department of about 80 people, with about six sections. So I was doing custom circuits, static memories, EPROM's, microprocessors, and all the peripheral circuitry in microprocessors. Did I say timing circuits? Yes, I probably did. And then services. And then there may even be something else… Oh, ROMs. So, basically, every category of MOS devices, which by that time was almost 100% of Intel.

**Hendrie:** There still was a bipolar.

**Faggin:** There was still a bipolar [activity], which was there for a few more years, and then it was killed. Then Intel exited bipolar.

**Hendrie:** I remember there was a bipolar chip set that Ted Hoff had.

**Faggin:** Yes. Ted Hoff did, which never made it in the marketplace… it was never successful in the marketplace.

**Hendrie:** Is that right?

**Faggin:** Yes.

**Hendrie:** Never got adopted by any major manufacturer….

**Faggin:** No. It never got adopted. No. It was a bit-slice [architecture], remember?

**Hendrie:** Yes, I remember.

**Faggin:** I think AMD was one of the pioneers in that bit-slice technology and technique. And AMD was quite successful. And then Intel was following on the footsteps of AMD there. Ted Hoff did the architecture and [Jean-Claude] Cornet, who's a French engineer, did the actual design. But it didn't do much in the marketplace. It was a flop. And they were doing some static, fast RAM; they were also doing some programmable ROM, with fusible links, in bipolar. Yes, fast ROMs in bipolar. And some other… I think they had a very fast content addressable memory, about 16 bits.

**Hendrie:** Oh yes, I think I remember that.

**Faggin:** 16 bit of content addressable memory -- something like that. And that was about it. So the bipolar family was relatively small. And then eventually -- probably in '77, '78 -- they basically exited the business.

**Hendrie:** And shut down the line.

  

**Faggin:** Shut down the whole thing, yes. Probably the line was shut down a little later. First they stopped new designs, and then they disbanded the design group. And then they kept production going for awhile.

**Hendrie:** When the customers… There weren't any customers…

**Faggin:** Well, they had customers. They had to support customers for a little longer. And then they had end-of-life buys, and that was the end of it.

**Hendrie:** So what do you do while Shima's designing the 8080?

**Faggin:** Well, when Shima was there [working] on the 8080, I had 6 or 7 [other] projects going on. There were some custom chips: there was a set that was done for a terminal manufacturer; it was a 3-chip set for an old terminal: there was a memory controller chip, a CPU chip, and another chip – I've forgot right now. Oh, yes, an I/O chip. And, again, something stupid, that we should never have undertaken but, again, from the top down… There was another friend of Gordon Moore…

**Hendrie:** What was the company? Do you remember?

**Faggin:** I think it was Ricoh.

**Hendrie:** Okay. You think it was Japanese company.

**Faggin:** Yes, it was definitely a Japanese company, but I'm not sure 100% that it was Ricoh, but almost sure. And so, again, [we were] taking a lot of time to go do something that was unique and not very useful. And, in fact, that [chip set] never went into production. And then we were doing a single-chip calculator, for desktop calculators. Again, a custom… Actually, no, that was not a custom design; that was our own design. But we had actually bought the logic design of that [chip] from the same engineer for which we were doing the chipset for Ricoh. So that same engineer -- he was an American engineer, working for a small R&D contingent, here in the Valley -- had done this design. And so we actually bought it, and we put a single-chip [calculator] together with that [design]. But, again, that was never pushed in the market place, it was another one of those things that was done halfheartedly. Then, at about that time, I took over the 2102A. And I put a lot of energy there. What happened was that in '72, Intel had….

**Hendrie:** For our listeners, let's just explain about the 2102.

  

**Faggin:** The 2102 was the [world's] first 5 volt, N-channel, static memory: 1000 bits. It was a very nice product, architecturally, but the process was slow. This memory was spec'd at 1.5 microsecond access time, and there was a selection at 500 nanoseconds. The market really wanted 500 nanosecond, 700 nanosecond, or better, and Intel could hardly make enough of these 500 nanosecond [chips]. In fact, they were in trouble. Burroughs was the big customer -- I believe the only customer in volume -- that was buying 500 nanosecond parts -- they were a selection of the distribution. The customer was screaming and yelling because they couldn't have enough: The yield was very low at 500 nanosecond, maybe 10 - 15% of the distribution was at that speed. So, Vadasz asked me -- in desperation -- to develop a faster 2102, which was later called 2102A, and he proposed to use, basically, thinner oxide. And I came back and I said, "look, yes, thinner oxide helps a little, but it's not the solution, we need depletion load because we are starving for voltage." See, the depletion load gives you full supply at the end, not 1 Vt below. So, you have 5 volt supply and 5 volts output in any gate, instead of say 2.5 volts. 2.5 volts minus the threshold gives you a very small overdrive, where 5 volts minus the threshold -- which is around 1-1.5 volts -- is a large overdrive. So it makes a big difference in speed. And also a depletion load is like a constant current source, so you really have much higher [and faster] positive swings….

**Hendrie:** Oh, okay. And it just discharges that capacitance, just like that.

**Faggin:** Yes. Oh, it's just perfect, right. Just perfect. And, at that time, Intel was in the timing circuits, which was also an activity that I picked up a bit later, so they were doing CMOS devices, so they had an implanter. Of course, you needed an implanter to be able to do depletion load. They had everything that they needed to actually do this. And Vadasz didn't want me to do it. I later on found out why: because he had sent a memo to Gordon Moore saying that depletion load was not the way to go, that thin-oxide was the way to go, and therefore he didn't want to look like a fool. But I didn't know it at that time and I didn't care if he was looking a fool. If he wanted me to do the job, that's what needed to be done. And so, I actually ended up having a fight with him. And I said, "look, Vadasz, if you want me to do the job, you've got to give me depletion load. If not, get somebody else to do it because it's not going to work, okay?" And so he left the meeting, slamming the door, and said, "okay, do what you want!" Okay, I'll do what I want…

**Hendrie:** That's just what I wanted to hear.

**Faggin:** That's right. And it's exactly what I did. And Dick Pashley, who was a PhD from Caltech, had just joined Intel. He was assigned to me, and it was his first chip. And then a processing guy was [assigned], to develop the process technology. That was quite simple to do, really, at that point. It was an extra masking layer and you had to decide the right dose -- the implant dose -- and that was it. So we had the process within a few months. And also the design was a relatively [straightforward]….

**Hendrie:** And it wasn't hard to characterize.

  

**Faggin:** It wasn't hard to characterize. No. So, basically, within six months from when I got the project, we got devices that came out. And right there, we measured some devices going at 80 nanosecond access time, 80 nanosecond access time!

**Hendrie:** When you're having yield problems.

**Faggin:** Well, you had yield problem at 500!

**Hendrie:** Now, was this the same design rules?

**Faggin:** Same design rules.

**Hendrie:** There was no shrink here?

**Faggin:** Same design rules. No shrink.

**Hendrie:** Or anything like that going on?

**Faggin:** No. No shrink. I just used depletion loads, and a little thinner oxide. That's it! Not even that much thinner, because we didn't want to take reliability hazards.

**Hendrie:** But you said, well, I'll use a little bit of that too.

**Faggin:** Yes, of course. Why not?

**Hendrie:** Why not?

**Faggin:** That definitely helps. But 90% of the speed contribution was the depletion load; 10% was the little thinner oxide. That's it. So then, within 3 months after that, we were in production, and we saved the day with Burroughs. And, at that point, when I saw those parts, I said, "Les [Vadasz], now we can go after bipolar memories." Because bipolar memories using isoplanar process -- in those days, Fairchild was a dominant player in that [field], they had static memory using isoplanar process and it was a much more dense process than Intel had -- were around 60, 70 nanosecond access time. And, I said, "with this technology, we can push it a little and we can get there, and we can beat bipolar." And that was not even thought possible, in those days.

  

**Hendrie:** And we'll have higher densities.

**Faggin:** Absolutely.

**Hendrie:** And obviously much lower power per bit.

**Faggin:** And that was the beginning, you remember. And then I left, six [months], probably about a year, after that [project]. And Dick Pashley took over, and that became one of the most profitable areas of Intel, for a number of years, where they were making these very fast static RAMS. They [the RAMs] were used in cache memories all over the place, and so on. So that direction started with….

**Hendrie:** With that change.

**Faggin:** With that change.

**Hendrie:** Add the depletion load to a product [you were] already making.

**Faggin:** Yes, that's right. And, of course, I never got credit for it. Pashley got all the credit, as if he had started the whole thing. But it's an old [habit] for Intel. But that was really very important. And what it did, it gave me a process technology for the next generation of microprocessors, because that's what I used, that same technology, for the Z80. Same technology, because that was a great technology! Yes. That was the best way to do circuits. It actually had more enduring value than P-channel or N-channel, the conventional N-channel. The N-channel with depletion load stayed longer in production than any other of the prior generations. P-channel was only a few years.

**Hendrie:** I know.

**Faggin:** Then N-channel was also, a few years [later], replaced by depletion load; and depletion load stayed until CMOS took over, which was really toward the middle of the '80's. So it was, therefore, [used for] almost 10 years: It was the process "horse" of MOS for a long time and then it was superseded by CMOS.

**Hendrie:** Where are we? Okay. You had just gone through one of the projects you were doing, you know, while the…

**Faggin**: The 8080 was being designed, yeah.

**Hendrie:** While the 8080 was being done, and was this the…

  

**Faggin**: 2102A.

**Hendrie:** Yes, exactly. Any other things that were going on then?

**Faggin**: Then another thing that I started around that time, toward the end of the 8080 project, when I took over the timing circuit business -- you probably remember that Intel had purchased Microma. And Microma was an American company that was one of the first, if not the first, company to design an electronic watch. Remember the electronic watch debacle?

**Hendrie:** I think I have heard that. But tell us the story.

**Faggin**: So, Microma was a company that was purchased [by Intel] -- I think was purchased in ’72, around that time -- and they had developed a liquid crystal electronic watch that was quite expensive: $250 a pop. And, as I said, it was perhaps the first, arguably the first, liquid-crystal electronic watch in the market. At about the same time, or maybe a little later, was, remember, the Pulsar [watch] that had LEDs, and you had to push a button to light up the LEDs?

**Hendrie:** Yes, yes.

**Faggin**: In fact there was even a 007 movie where 007 had one of these watches… Of course in bed…

**Hendrie:** Yes, in bed, right.

**Faggin**: …with some gorgeous lady. So Intel, then, because of this acquisition, became thrust into the timing watch circuits design.

**Hendrie:** And that was intentional, to get in to it, yes.

**Faggin**: It was intentional, yeah, that’s right. And so they developed a CMOS technology. They developed technology to do sub-threshold [operation], because you had to have very low power, right? You had to have a one-year battery life. And you had to have very low voltage, and so it was a special process, and you had to operate devices at sub-threshold [below the normal threshold voltage of the MOS transistor]. And there was a very capable engineer that was running the group -- it was a small group. When I took over the group, my contribution there was to start down a path where the architecture of the time-keeping was basically designed around a micro-controller. And so, instead of each chip being different -- you had to have random logic, and then if you had to have a model change, you’d have to make a new chip and so on -- I basically argued for an architecture that was a specialized processor doing the handling of the time keeping and the watch functions, that was much more cost effective. And you could do a family of watches with one chip as opposed to one-watch, one-chip. So that [idea] was adopted and eventually became the way to go [for the entire industry]. But, by the time that was all done, or soon after, Intel got out of the whole business because it was not a good business. As you know, Texas Instruments got into it and they lost a lot of money. National Semiconductor got into it, also with a lot of losses. Basically, the semiconductor industry decided that they had to become watchmakers, but they didn't understand anything about consumer business and how to market the watches and so on. So it was a bloodbath, and all of them, eventually, got out of the business.

**Hendrie:** Because they didn't understand.

**Faggin**: Didn't understand it, yeah. So that was another thing that I was doing while the 8080 was being developed, and then, of course, I was worrying about a bunch of peripheral circuits that needed to be done also, in order to help out the building of micro computers. In other words, serial controllers, parallel ports -- programmable parallel ports -- and programmable peripherals. And so I began to worry about that. So, by early '74, we also had a group that was beginning to develop peripheral circuitries.

**Hendrie:** Okay, and that was sort of their charter to work on.

**Faggin**: And that was their charter. Yeah, yeah. And so that's when the serial controller was done, and the programmable parallel I/O controller was done and so on. And, basically, my last year at Intel, was mostly managing, managing and, obviously, giving technical direction, but I was no longer involved [directly] in any design.

**Hendrie:** Managing and leading.

**Faggin**: And leading, yeah. Actually, leading is probably more correct. In that year or so that I was in charge of this large group, we had close to 20 projects going on at the same time. So, it was a lot of projects.

**Hendrie:** So you had a lot of designers going on, and then of course all the other people.

**Faggin**: Of course, as I said, 80 people in the group. So I mean it's a...

**Hendrie:** So yeah. So maybe, what? 15 and 12, 15 designers?

  

**Faggin**: Oh, the designers? No, no. In some cases there were more than one designer per chip. If they were complex chips, there would be a senior designer and a junior designer learning and helping, that kind of stuff…

**Hendrie:** So it's starting to become that kind of company.

**Faggin**: That's right, yeah.

**Hendrie:** Now who were you working for in this time?

**Faggin**: For Vadasz.

**Hendrie:** Vadasz?

**Faggin**: I worked for Vadasz the whole time, the whole time.

**Hendrie:** The whole career was there working with [Vadasz].

**Faggin**: The whole career, yeah, yeah. When the last reorganization that I was involved with occurred [at the end of 1973, early '74], Vadasz became VP of R&D, and in his group were folded Application Research, so Ted Hoff worked for Vadasz from that point on, and also the Bipolar Design -- that was a separate group -- also ended up reporting to Vadasz. And so when that big reorganization occurred, Vadasz took charge of a bunch of other things, and I was given almost everything that Vadasz had [before]. It was my responsibility after the reorganization. The only thing that Vadasz kept without giving to me was the Dynamic RAMs. That was one of the core…

**Hendrie:** Was clearly a very important...

**Faggin**: …core, core business for Intel, and also I would not have wanted it because that was completely consuming. And it was really hard for Intel to keep pace. At that point, Micron had actually taken the lead. Not Micron, sorry…

**Hendrie:** You mean Mostek.

**Faggin**: Mostek, sorry.

  

**Hendrie:** Exactly, that's why I was going to say they were…

**Faggin**: Micron came out of Mostek. Yeah. Mostek actually took the lead. In the 16 kilobit design it was Mostek that set the pinout and set the standard. Intel tried, but they lost out to Mostek. So, so there was a lot of pressure in that group, when I was there, and I would not have wanted that group for nothing in the world.

**Hendrie:** Didn't Mostek also, you know, develop a successful, you know, single bit? Did they develop a successful one transistor per bit design before Intel, or…

**Faggin**: At about the same time.

**Hendrie:** About the same time.

**Faggin**: About the same time.

**Hendrie:** They were both obviously working on it, because that would be the holy grail of getting bigger.

**Faggin**: Yeah, yeah, at about the same time, yeah. Basically, Mostek adopted the Silicon Gate Technology. See, Mostek started with metal gate with ion implant, which was the only other way to make sensible MOS devices in those days -- with ion implant. Without ion implant, the technology was really very poor. And Mostek got down that path. And then…

**Hendrie:** Well, and they used the ion implant to do the depletion loads.

**Faggin**: And also to control the field threshold [voltage] because, if you know about [MOS] technology, the MOS devices, before the silicon gate, had a very narrow window of operation, because you were caught between breaking the metal because the oxide thickness was too thick.

**Hendrie:** Yes. And it couldn't go over the steps.

**Faggin**: Yeah. And metal couldn't go over the steps, on one hand; on the other hand, having all kinds of parasitic devices because the field threshold was lower than Vdd, the supply [voltage], and therefore you would have all kinds of problems. Plus dirty oxide that could not be cleaned up with gettering because you could not do the gettering after the metal was deposited. All of those problems went away with the Silicon Gate Technology. But Mostek made the metal gate last a little longer, stretched it out by using ion implant, which was a new technique in those days -- with aluminum gate – and got more mileage out, but still they had to capitulate and take the road of silicon gate, which they did with the 4K [dynamic RAM] memory.

**Hendrie:** With their 4K memory.

**Faggin**: And from that point on it became all silicon gate. And they were the first to use depletion load on 5 volts, also. Okay.

**Hendrie:** Ah, okay. To do 5 volt depletion load N-channel.

**Faggin**: Yeah, N-channel, yeah. So Intel was not the first to do depletion load N-channel, Mostek was. Yeah.

**Hendrie:** Wow, okay. So now you're…

**Faggin**: So now I'm...

**Hendrie:** In the last year you're managing and…

**Faggin**: I'm managing a bunch of people, and I'm getting restless.

**Hendrie:** Managing isn't your most fun thing.

**Faggin**: No, it is fun. I mean, that's what I wanted, right? But at the same time 1974 was a recessionary year, as you remember.

**Hendrie:** Oh yes, I forgot that.

**Faggin**: In the middle of '74, Intel had their first layoff. It was tough, and I was getting a little frustrated with the management style of Intel which was shifting more and more toward the management style of Andy. Andy was taking more and more power -- Andy Grove. And Intel had started the signup sheet where if you came after eight o'clock you had to sign your name, because you're late and you're supposed to be at work before eight o'clock. I used to work nights, so if I would come at eight-thirty, then I would have to sign up, then be reprimanded that I came in late. And I was really upset about that. I felt that it was very unfair, that it was going to level out everyone, that basically…

  

**Hendrie:** It was more important when you came to work than what you accomplished.

**Faggin**: Yeah, that's right. At what time did you go to work. Right. So, then I found out, in the early part of '74, that Vadasz had patented the buried contact. Dov Frohman told me. He said, "do you know that Vadasz patented your idea?" because Dov Frohman was at Fairchild, he knew about the buried contact idea of mine. And so he said, "Do you know that Vadasz patented your idea?" "Really?" I said, and I became very, very upset, and I went to Vadasz and said: "Is it true that you have patented my idea of the buried contact?" And he said, "Hmm… let's go talk to Grove!" And he took me to Grove; didn't say a word. And it was clear that he knew exactly what he was doing, and that Grove had something to do with it, also. And, in fact, I found out later that he [Vadasz] tried to get the processing engineer that developed the process [Tom Rowe] to patent the idea, but the guy said "No, no, that's not my idea. You told me what to do, I don't want to have my name on that." And so, he [Vadasz] basically had to put his name. So, anyway, it was one of those things… Intel was not an easy company for me to work with, because of this attitude. First of all, I never felt acknowledged for what I did. I always felt like I was never good enough: do more! And they didn't appreciate what I did because they didn't even understand it. It was almost like, "oh, give it to Faggin, he gets it done," so what's the big deal? Okay. And frankly, I showed that, in their own business, which was memories, just by taking over one project, I changed completely what could be done with the 2102A, ok?

**Hendrie:** And Noyce was not really connected that much.

**Faggin**: Noyce was already preoccupied with something else. Already in '74 he was not as active and, really, Andy was the power.

**Hendrie:** And Noyce was not running the company anymore. So he was setting the…

**Faggin**: Not really, not that much. Well, he was still CEO in those days, but he was not as active as he was in the very early days of the company. And Andy had really taken over the leadership role of the company, as a matter of fact, yeah.

**Hendrie:** Okay. All right. Yes.

**Faggin**: No question about it.

**Hendrie:** So then it's his style that starts to permeate through the company…

**Faggin**: It was his style, yeah, that's right. Exactly right.

  

**Hendrie:** As opposed to Noyce's.

**Faggin**: As opposed to Noyce's, and I could see that, and I didn't want to have much part of that. So the combination of all those elements, that were negative elements, plus the positive element which was, "hey, if I start up my own company, I'll learn a lot. I'm going to work as hard as at Intel, and if I'm successful I will make much more money that I can make at Intel." In all of that, the appeal of the positive [elements] and the sort of the pushing of the negative combined, basically, to make me decide in the late summer of '74 that I wanted to start my own company.

**Hendrie:** Now, had they been reasonably generous in stock to people, you know, to design engineers? People at your level at least had accomplished a lot or even… you were senior manager at that point.

**Faggin**: I had a little more stock than an average engineer, which was absolutely not right, not okay. And so I was not…

**Hendrie:** So they weren't saying thank you with stock either.

**Faggin**: That's right, or with salary either, okay? And so I was not treated that well compared to what I had accomplished for the company. But it wasn't so much that because if you start your own company, the difference between a founder, and even a department manager, as I was there, would still be a big difference. Right, so…

**Hendrie:** Oh it's huge, yes.

**Faggin**: Yeah. So, at that point I was past that. But also from a financial motivation, it was not very strong for me at Intel. So, by the end of the summer, I decided that I wanted to do it. And I wanted to do it before I would get involved with something else, and that was the right time to do it because otherwise it would have been difficult to get out. I mean, if you are in the middle of a new major project, the next generation…

**Hendrie:** Yeah, you've signed up for another project. How are you going to…

**Faggin**: And I didn't want to do that. So that was the right time, and one night I asked Ralph Ungermann, who was one of my managers, out for a drink, and...

**Hendrie:** Now, Ralph Ungerman was involved in Intel?

**Faggin**: Oh yeah, he was working for me from day one.

**Hendrie:** Oh, he was working for you. Oh, he was one of your managers that worked for you.

**Faggin**: One of my managers, he worked for me, yeah.

**Hendrie:** Oh. And then what was he doing?

**Faggin**: Well initially, he came on board, and he took over custom circuits. But, as I told you, the company was never particularly…

**Hendrie:** Very interested.

**Faggin**: …interested in that. So he was handling that for a while. And then he took over the microprocessor business. So in early '74, I gave him microprocessors, which included -- well the 8080 was done at that point, pretty much -- but included all the new peripheral circuits that we were going to do, and also the new processors.

**Hendrie:** And were they starting new processors yet?

**Faggin**: No, no, we never started a new processor because much of the attention in those remaining months before I left -- and he left -- was toward getting the peripheral controllers started.

**Hendrie:** Getting all the support chips. Getting a real family so that you could actually build something useful.

**Faggin**: All the support chips. Yes, that's right. Build something well, very effectively, yeah. And so we did that.

**Hendrie:** Okay.

**Faggin**: And that's why I wanted to leave before [starting a new microprocessor platform] because it was the beginning of, "what are we going to do next?" And so on. And I just wanted to get out…

**Hendrie:** 4040 had not been invented yet?

  

**Faggin**: Oh no, the 4040 was almost done at that point, or all done.

**Hendrie:** Oh, okay, so they decided to go do that.

**Faggin**: Yeah, they decided to do that, yeah. That was decided at least a year before, but it was kind of late now because it was still having to be compatible with the 4004, so it had to be P-channel, and so it was too late. It should have been done when I proposed it first, then it would have gotten a little more life out of it.

**Hendrie:** Oh, it had to be P channel? It was not an N-channel?

**Faggin**: No, no, because you had to be compatible with the 4001. You had to be able to use the other parts as well, like 1702's. We did not have any N channel EPROM yet. And N-channel EPROM came out right after I left.

**Hendrie:** So, there wasn't a…

**Faggin**: So there was only P-channel EPROM, and that was key to be able to properly…

**Hendrie:** To build micro controllers.

**Faggin**: To build micro controllers, yeah. And that was also one of the reasons there was a little bit of lull in the push toward N-channel processors because of the lack of EPROMs and so on, that were voltage compatible. Anyway, so…

**Hendrie:** You took Ralph out…

**Faggin**: I took Ralph out…

**Hendrie:** For a drink.

**Faggin**: For a drink. And I said, "Ralph, how about starting a company -- a micro processor company -- you and I?" And he said, "Okay!" He didn't even ask, "what do you have in mind, what do you…"

**Hendrie:** Really.

**Faggin**: Yeah. He said, "Okay!"

**Hendrie:** Now, had he been a design engineer?

**Faggin**: Yeah, he had been a design engineer, though he was never greatly accomplished in chip design. He was much more of a logic designer, a system level kind of designer himself.

**Hendrie:** Okay. And that made sense in the custom chip business.

**Faggin**: Yeah. That's right. Yeah. And he had designed a few chips, but he was not…

**Hendrie:** That was not his biggest skill set.

**Faggin**: Yeah, biggest skill, yeah. He was much more of a system designer. Which was good because then he was complementary to my skill set.

**Hendrie:** Exactly.

**Faggin**: Because again, we needed to do development systems, in my mind I wanted to do micro computer boards, OEM micro computers, that kind of stuff. And so it was important to have somebody with that background. So that was the beginning, and then he resigned, and I resigned after he did. And then I was asked to stay another couple of months to help out in the transition, which stupidly I did. I was such a naive guy because I didn't have to do it, right?

**Hendrie:** You didn't have a contract saying…

**Faggin**: I didn't have a contract; it was just out of trying to be nice. But in fact, I hurt myself by doing that. Anyway, I stayed for a couple more months. Nobody knew that I had decided to leave, and so it was also hard for me to kind of play the game knowing in my heart that I was leaving in a few months. And Vadasz didn't even take advantage of that time to find a replacement. So that, basically, I worked two more months without the benefit of having really achieved the opportunity [to replace me]…

**Hendrie:** Yeah, what your opportunity was.

**Faggin**: So I left, and my last day was -- actually, I remember because it was Halloween of 1974 -- October 31$^{st}$. And that was a big a big deal because my group was big, and I was respected by my people. A lot of my people actually got upset by my leaving, and upset even with me [personally] that I left. You know, how could I leave? Because I'd always been very supportive of Intel, very supportive of whatever we did. I was not sharing my frustrations with them.

**Hendrie:** Yes.

**Faggin**: And that was the beginning of Zilog.

**Hendrie:** Now had you found the funding or had you talked to anybody about raising money?

**Faggin**: No.

**Hendrie:** Just going to leave.

**Faggin**: Leave. I didn't want to think, even, about what was the first product I was going to do. I was just…

**Hendrie:** Just you get ..

**Faggin**: I didn't want to be contaminated with any of that stuff, okay? Most people do a business plan already, already talk to somebody. I didn't want to do that at all.

**Hendrie:** They started development in their cellar?

**Faggin**: Yeah, absolutely, in their cellar or something.

**Hendrie:** A lot of things like that happen.

**Faggin**: No, no. I did not do any of that. I just wanted to get out, and then I'd take care of it at that time. Yeah. So what we did then, we did consulting for a little while, while we were developing the business plan, we were developing the basic idea of what we were going to do. And in fact, the first idea that I had was the idea of the micro controller: A single chip computer that was very cost effective and I/O-oriented, because that was what you do, right? You basically "bang" those I/O's, right? And I called it the 2001, as a working chip [name], and I developed the architecture. That's what later became the Z8, the Zilog Z8.

  

**Hendrie:** All right. Okay, yes.

**Faggin**: And then -- this was toward the end of November when I did that -- I decided against it because the idea was that we would use a foundry. I think Zilog was the first fabless company in the world. I mean, really.

**Hendrie:** Yeah, really. Everybody thought you had to do it yourself, even people who…

**Faggin**: Of course. That's right. So we actually went into business as a fabless company, and then we had to -- because there was nobody that was reliable to fab for you -- we had to create our own [fab]. But we started with the fabless model, before it was known as a business model, using Synertek.

**Hendrie:** Oh, all right.

**Faggin**: And Synertek was a company that had started as a fabless [not fabless, but rather, a chip manufacturing] company for corporations that wanted their own custom circuits. So, not quite the fabless model. In other words, people that are captive users of chips would have – and some of them where their investors, like, I think, Honeywell was an investor, and other companies were investors -- their own products. So they were basically a product company, but with a fabless model to be a captive source for some of their investors and potentially others like them, but not people that would sell chips in the marketplace.

**Hendrie:** Yes, okay.

**Faggin**: People that would…

**Hendrie:** People that needed chips, needed, right…

**Faggin**: Needed chips for their own products, yeah. Systems companies, yeah.

**Hendrie:** And might have been tempted to build their own fab to do it, but recognized hey, there is a much more cost effective approach.

**Faggin**: That's correct. So like an intermediate step from…

**Hendrie:** I understand.

  

**Faggin**: From fab to fabless.

**Hendrie:** Yes.

**Faggin**: And so with that decision, it was clear to me that the 2001, this 8-bit micro controller, was not going to be cost effective because we would have to pay too much per wafer, right. And because I could see that the marketplace for this device was going to be very price sensitive, I couldn't see [how] to have enough [profit] margins. And that's when I came up with the Z80 [idea].

**Hendrie:** Okay.

**Faggin**: I still remember it was a Saturday. I was debating what, you know…

**Hendrie:** You and Ralph are…

**Faggin**: Myself. No, I actually did all that, I did all of this myself. Ralph was not involved much at all because he was doing more of the consulting… It's a long story, but he and his wife had started a company [about two years before] that was consulting in testing. And so when we started this [new] company [Zilog], as a shortcut, I bought 50 percent of this company [Ungermann Associates].

**Hendrie:** Oh. Of his…

**Faggin**: Of his, okay. So we were 50/50, okay, and I was the CEO and he was Executive VP, and his wife was CFO. And we were using this as a way to pay the bills, pay the office, while we were building the business plan and the…

**Hendrie:** For it, yes.

**Faggin**: For this company [Zilog]. So I was the one really putting together the business plan and putting together the basic product strategy and so on, and Ralph was doing a little more of the…

**Hendrie:** Testing consulting.

**Faggin**: Consulting work. Yeah. Because we had to put some food on the table too, right.

  

**Hendrie:** You do have to get some money, yes.

**Faggin**: So, it was some time in December, it was on a Saturday, and I was debating what was the right thing to do. And that's when it hit me, it just hit me all at once. It was just like it's own gestalt, and I said, "Super-80!" And Super-80 meant having envisioned what I wanted. It meant: a five-volt, depletion-load micro processor with a lot of improvements which I already could see reasonably clearly, like twice as many registers, more instructions, more addressing modes, more bit-level instructions, a better interrupt structure yet, and designed as a family from the outset. So, where the peripheral controllers are a patchwork done afterwards, this is done all together. You envision a system…

**Hendrie:** Envision the system, then partition. This is the CPU, this we'll put in this chip, this will go put in that chip.

**Faggin**: That's right, yeah. And so that whole thing made a lot of sense to me, and it also made a lot of sense in conjunction with the manufacturing strategy that we were adopting where we had to pay, obviously, twice as much per wafer than if we were making our own. So it had to be a high-level processor to make sense. And that was the decision. And then I told Ralph, and Ralph said, "Yeah, yeah, good. Let's go." Then, in the next few months, I designed the architecture and the instruction set, the whole thing, and then, of course, much of the effort was trying to get money.

**Hendrie:** Yes, I was going to say that.

**Faggin**: 1975 was the worst possible time to start a company in the history of the US high-tech industry. Believe it or not, in 1975, the total venture capital investment in this country was ten million dollars, in that year. Ten million dollars! All the VCs were gone [underground]. Nobody had money, nobody; lots of people had gone out of business. The ones that were in business were just in a wait and see [mode]. And basically, from the excesses of the late 60s, early 70s it was, Pooofff!

**Hendrie:** It was a wasteland.

**Faggin**: Wasteland, okay. And we were lucky to find…

**Hendrie:** So now who did you, tell me about some of the people you found.

**Faggin**: We found Exxon Enterprises.

**Hendrie:** You found Exxon.

  

**Faggin**: Actually, it would be more correct to say that Exxon Enterprises found us. And actually that's an interesting story because soon after Ralph and I left, we called Electronic News.

**Hendrie:** Called what?

**Faggin**: Electronic News.

**Hendrie:** Oh, Electronic News, yes, yes, sure.

**Faggin**: The, you know, the principal rag…

**Hendrie:** Of course, they were the principal rag at that time. Right.

**Faggin**: The principal rag, right, that's right. And we told them that we had left to form a microprocessor company. Okay. So they printed a nice article in Electronic News, "Faggin and Ungerman leave Intel to found microprocessor company," blah, blah, blah, blah, blah, right. And that got the attention of Exxon Enterprises in New York. Exxon Enterprises was a subsidiary or an affiliate, probably an affiliate, technically, of Exxon Corporation. And they had started dabbling in the information technology, and they were funding both internal efforts and outside companies. I got a call, soon after the article appeared, by one of the young partners there, asking me if we were looking for money. And I said, "Well, not right now actually, we're just putting together a plan of what we want to do, and we don't even know how much [money] we need yet. But hey, if you come around the area, we would love to talk to you," okay.

**Hendrie:** Right.

**Faggin**: Actually that was a good way to hook 'em. [Laughs.]

**Hendrie:** Right. Exactly.

**Faggin**: Because if you need those people, right…

**Hendrie:** If you are too anxious…

**Faggin**: My instincts were good. And so, that was the first conversation, and then a few days or a week later, the same person called and said, "We're going to be in the area, can we drop by?" And I said, "By all means. Yeah, let's come over." And so I had a week to put together a business plan, the beginning of a business plan, to present to Exxon so that we would make a reasonable impression, right. So two of them came over a week later, and I had put together a business plan around the 2001.

**Hendrie:** Okay.

**Faggin**: Because at that time I had… In fact, I think that the idea [of the 2001] came kind of instigated a bit by the fact that I had a meeting a week later. So it was like, "I better hurry up here." I'm not sure of that, but I think it was like that, that [the 2001] came after my first conversation with Exxon. Not the second, the first, recognizing that there was an interest there, so I'd better hurry up and [come up with something].

**Hendrie:** Yeah, you better hurry up and decide what we're doing.

**Faggin**: And decide what I want to do here, right? And so I put together a business plan. At that time my wife was the executive secretary that was helping us. So she typed the -- not a very thick -- business plan, talking about the 2001, the market that it would open up, and so on and all that. And then we presented that, and they were very impressed and they said that they wanted to continue to talk. So then, when in December I came up with the Z80 idea, which I called Super-80 to start, I was a bit concerned because, I said, "now I'm going to go back to Exxon Enterprises and tell them that I already changed my mind, than what I sold them already, I've decided that it's not good enough…" Right?

**Hendrie:** You have a better idea. Yeah, right.

**Faggin**: But I felt that that was the right thing to do because I didn't believe that it…

**Hendrie:** Yeah, you didn't believe in the other. You didn't want money to build something you didn't believe in.

**Faggin**: That's right, that's right. And so I said, "well, if they don't want to do it, it's still not a good idea for us to do that." And so I decided to go for it. And I said, "Look, this is the reason, and that one [the 2001] will be the next product. So, the first product is going to be the Super-80 -- the Z80, what we later called Z80 -- and the next product is going to be the Z8," -- I didn't call it Z8 at that time, but 2001 -- "and so it's still part of the plan, but it's much better if we do this first because you've got better margins, and better position in the market." So, they bought it. But then it was kind of hard to negotiate the deal with them to get the money. I decided that half a million dollars was what we needed to basically prove the concept and go to market and have the beginning reaction of the market, and also to build a development system and all the essential software to aid the customers to design in the parts. So they [eventually] agreed with that, and to make a long story short, by April of '75, we had a verbal commitment that they were going to do it. We still hadn't signed a deal, we didn't have the money, and so I said, "Okay, fine. If you have your real commitment, then we need to start. We cannot keep on waiting or spending time here, because otherwise we lose the market window, right." So they agreed to give us $20,000 a month until the deal was closed so that we could hire some people and begin work. Because at that point, by February or so, I was finished already with the architecture, and we could have started designing, right.

**Hendrie:** You wanted to start designing it, and building it.

**Faggin**: So from February on we were wasting time. In April they agreed -- we lost about a month and a half or so -- and so I hired Shima, who in the meantime was pressing me. He actually said, "I want to come with you." And I said, "Yeah, but you have to wait until, until…"

**Hendrie:** I have money.

**Faggin**: Yeah, I'm ready." And then just in March, I think, he called me up and he said, "You've got to hire me soon because otherwise I have to get involved in the next project, and I don't want to do that." So that's also part of the reason why we did it this way. And so, in April, Shima came. We also hired a friend of Ralph -- he was a system guy -- to do the development system together with Ralph. And…

**Hendrie:** Who was that?

**Faggin**: His name was… it escapes me.

**Hendrie:** That's all right.

**Faggin**: It's okay. I'm getting tired.

**Hendrie:** Yes, okay.

**Faggin**: Yeah. Doug Broyles was his name, Doug Broyles, Doug Broyles. And then we hired also an engineer for the software, Dean Brown…

  

**Session 3: March 3, 2005**

**Hendrie:** Well, we have here Federico Faggin, he is continuing his oral history for the Computer History Museum and I think where we left off at the last session, you had just hired Shima at Zilog and you were starting to really… to go to town and design the Super-80, not to be called that too long.

**Faggin:** That's correct, yes.

**Hendrie:** All right. So maybe you could take the story from there.

**Faggin:** Okay. So we are probably about April of 1975; that's when Shima started. At that point I believe we had worked out some kind of a loan arrangement with Exxon that they were going to pay our costs until an investment would occur. And we started in earnest the actual design of the Z-80 by first completing the work on the characterization of the process technology that was supposed to be used. That's a task that I did, working both with Synertek, which was a sort of a wafer foundry in those days, and also working with Mostek. We were already talking to Mostek as a potential second source, at that time, and they allowed me to actually have access to some of their process information so that we could characterize the process, so I put together design rules that would allow the Z80 to be run in both factories. And, in fact, they were…

**Hendrie:** A compromise between the two processes.

**Faggin:** A compromise, yeah, that would allow easily porting the Z80 into at least three companies because I knew the design rules of Intel, so it would definitely work at Intel, it would work at Synertek and it would work at Mostek. And also designing all the same kind of methodology that I had used earlier for the Intel processors, doing quick graphics design to be able to quickly do what you needed to get done, because, although at that time simulation was more accessible than it was when I started the first microprocessor, it was still expensive. And it still would take time and so we were going to use it sparingly. So by June of '75 we got our money. We got half a million dollars from Exxon Enterprises and with that we could really start the business. So we started hiring some additional people at that time. When Shima came onboard, also Doug Broyles came aboard. Doug was a systems guy and we decided to hire him also early because it was essential to start the development system in parallel with the design of the chip. Otherwise we could not have early design-ins of the Z80 without the development environment that involved both the development system and a minimum amount of software like, an editor, an assembler, etc., etc., a debugger, and so on. So shortly after Shima came onboard we also hired Dean Brown, Dr. Dean Brown. He was a software expert and he came to head up the software activity. So by the time we got the money we already had five people in the company and then we hired the other people within a short period of time of having the money and we built the company to 11 people and with those 11 people we went all the way to, essentially, samples of the Z80.

**Hendrie:** So it took 11 people…

**Faggin:** Eleven people.

**Hendrie:** That was the team that, basically, did the whole development.

**Faggin:** They did the Z-80 chip, the development system, all the software, all the product literature, all the rest of it, finance, administration, the whole thing, the whole thing for the company, from that point on until March of '76. That's when we had fully working samples of the Z80.

**Hendrie:** So it only took you one year to--

**Faggin:** It took a little less than one year, yeah. Well, the architecture was done, as you remember, starting in December timeframe when I had the idea until Shima arrived and then it was refined between Shima and I, but much of that work was done earlier. Still, if you compressed that work, probably it was a couple months' work before April.

**Hendrie:** Now was this…

**Faggin:** Yeah.

**Hendrie:** I'm sorry.

**Faggin:** Yes, go ahead.

**Hendrie:** Was this the design… did you do the design in this style with, you know, basically, at the transistor level so there was only one translation…

**Faggin:** Oh, yes.

**Hendrie:** Only one translation to the actual rubies

**Faggin:** That style…

**Hendrie:** Yeah.

**Faggin:** That style remained with all the microprocessors that I was involved with, including the Z8000 later on, yeah. It was only by the late 70's that computer-aided design began to demand -- and also the complexity of the design -- began to demand a change in style. But up until that point, that methodology was quite effective and we continued to use that, yeah. So, the layout of the Z80 started at about June-July timeframe, and so at that time we hired two draftsmen, one senior and one more junior draftsman, but still experienced in MOS layout. And things were going quite slow and we were slipping. So I decided to take over the whole layout effort and I basically laid out with my own hands two thirds of the Z80, where the other two draftsmen did one third. I basically worked from that point, from June until October, 80 hours a week straight. I mean, it was really a grueling kind of work, but I had promised that by March $8^{th}$ or March $9^{th}$ we were going to have samples and, by the way, to the day, we had samples.

**Hendrie:** Oh, very good.

**Faggin:** Yeah, that was really something. I mean, the whole team worked very well. But had I not stepped in we would have been late by months. And also if we had put more draftsmen… Given that you don't have computer-aided-design, you cannot make mistakes. If you make some mistakes then you have to erase everything and start over again. There was no CAD in those days. So the ability to plan ahead and know how you would use the area properly, and so on, was absolutely essential to do the job well. And I think, frankly, the layout of the Z80 is a masterpiece. I think it was exceptionally well done. And, of course, Shima did a super job in the design and everything else. He handled all of that and I was basically helping him, you know, just…

**Hendrie:** Yes. You didn't need to step in at that at all.

**Faggin:** No, I didn't need to step in but if I didn't, you know, we would have been at least a couple of months late if not three months late.

**Hendrie:** What-- how big a chip did you end up that first--

**Faggin:** It was – with those first design rules-- it was about 200 by 200 square mils. So it would be, what? 100 mils is two and a half millimeters, so it's five millimeters by five millimeters.

  

**Hendrie:** And what was the basic gate… what was your sort of minimum gate width? Just so I understand the design rules you were using.

**Faggin:** I think it was still six microns in those days. Yeah, I think it was still six microns. Yeah, six microns. The actual pushing of the design rules and the ideas of scaling were just beginning in those days to be understood, but I did the process so that it could be scaled. So, there was the notion of scaling and, in fact, we scaled soon after. When we transferred the product to the Mostek fab, later on, it was scaled by about 20 percent, so the six microns became five microns and the chip size was reduced. Then we scaled it one year later when we moved it in our factory, we also scaled it again. So it was already predisposed for scaling. At the beginning we didn't want to take risks and we didn't know where we were going to be [doing the manufacturing], so the most important thing was to have it work.

**Hendrie:** Yes.

**Faggin:** So the first silicon was received -- I still remember -- it was received in January [1976], relatively early January and, basically, there were only a few mistakes. I mean, we could almost run a program, whatever programs were ready in those days when we were still working on the development system. By the way, the development system was done by developing it using a logic simulator of the Z80, and that was done not so much for checking of the logic of the CPU, but so that we could develop the software in parallel, otherwise we…

**Hendrie:** Yes, you needed something to run the software.

**Faggin:** We needed something to run the software, right.

**Hendrie:** Now was this a logic simulator, a hardware logic simulator or was this--

**Faggin:** Oh, no, no. Hardware. It had to be real time, right? It had to be real time. So, what we had [was] this big thing made of hundreds of little -- not hundreds, but tens of wired-wrap boards. Ralph Ungermann and Doug Broyles spent many nights wiring up this [hardware] simulator. Shima did actually the translation of his logic into the building blocks…

**Hendrie:** So he turned it into…

**Faggin:** He turned it into…

**Hendrie:** TTL or whatever.

  

**Faggin:** It was low-power CMOS.

**Hendrie:** In low-power CMOS. Yes, okay.

**Faggin:** Yeah. And so we had that thing done in… Or, low power TTL – I'm confused now. Anyway, it was the lowest possible power that we could [get] because, otherwise, it would have been hundreds of watts.

**Hendrie:** You'd then have another problem, heating problem.

**Faggin:** Another problem, right. So that's right. So we designed that -- yeah, yeah, it was lower power TTL. So they designed that [simulator] and then it would end up in a 40 pin package that would plug into the printed circuit board of the…

**Hendrie:** Ah, so you took all the inputs and outputs from it, made it into a cable, designed it into a…

**Faggin:** And the cable would go into [a 40-pin socket in] the board, ready for the Z80 to be plugged in. And so, in the process of doing that [simulator], there was basically only one logic error that was found. That helped the [logic checking] process. But, in fact, it would have been found at the first silicon anyway, so it was not a big deal. The real purpose of that [simulator] was the…

**Hendrie:** Was the software, debugging the software.

**Faggin:** But it also tells you how accurate was Shima in his design, right? Basically there was only one logic error in the whole Z80.

**Hendrie:** Oh, my goodness. Now how many transistors were… do you remember roughly?

**Faggin:** About ten thousand.

**Hendrie:** About ten thousand?

**Faggin:** Yeah.

**Hendrie:** That's pretty good.

  

**Faggin:** Yeah. So in January we got the first samples. Most of it worked. It took about two or three weeks to debug it and find out everything. And, you know, there was Shima -- he is like a steel-trap mind -- he would figure out all the patterns to check all the cases and he would be there with a patience that was incredible. And just stick to it, and he would go through…

**Hendrie:** Just systematically go through.

**Faggin:** Step by step, by step, by step, by step. And he'd say: "it works!" and then, "This works! This instruction works!" And then step, by step, by step, "That instruction works!" [Laughs.]

**Hendrie:** That's wonderful.

**Faggin:** That was something else! And, you know, at the end he said, "Checked it, all. Okay. Found these few mistakes and fixed it." We went through another spin and…

**Hendrie:** Did those… Did the other mistakes, the non… Were there other logic errors or was it circuit problems?

**Faggin:** No. No circuit problems.

**Hendrie:** No circuit problems.

**Faggin:** Only logic problems.

**Hendrie:** Only logic problems.

**Faggin:** Yeah. No, no, speed was fine, power was fine, everything else was fine.

**Hendrie:** Yeah, okay.

**Faggin:** No race conditions…

**Hendrie:** No voltage sensitivities or anything like that.

  

**Faggin:** No voltage. No, no, no, none of that, no.

**Hendrie:** Good.

**Faggin:** No, because in the design rules that I set up there was so much margin that, basically, if you just followed the rules, you get it right.

**Hendrie:** You get… Okay.

**Faggin:** And in the few areas where there was a critical design, like I/O's – the input/output -- where you had noise margins [to respect], we did simulation on those, so those were done. And again, there was enough margin that different processes would be accommodated. And so, then, after changing the masks and getting a new run going, we finally had the next revision of the silicon coming out, either the eighth or the ninth of March, I don't remember exactly but, anyway, early March. And what we did was wonderful because, at that point, the development system was at a point where it was already executing code, and so on, completely. And so, instead of testing the chip -- we didn't have a test program, a complete test program, anywhere at that point -- we were still…

**Hendrie:** Did you have a tester?

**Faggin:** What? Well, we had [yet] to develop the tester.

**Hendrie:** Yeah, so you developed your own. You didn't buy a Fairchild tester.

**Faggin:** No, no, no. They were way too expensive anyway.

**Hendrie:** Yeah, didn't have the money.

**Faggin:** Yeah, didn't have the money anyway. In fact, that was something that was the next project to do, the production tester, but for now we had it jerry-rigged enough that we could go through the full characterization and the full testing, but with pieces, as opposed to the whole thing. Then we needed to put the pieces together into [the full tester]. But the architecture of the tester was already understood because we had been doing testers at Intel, both Shima and I, for a number of years, so we knew exactly what we wanted. But the beautiful thing was that that evening, after we received the wafer, we set up the probe station and so from the wafer, with the probe station, we ended up with a 40-pin package that would plug into the Z80 socket in the computer board of the development system. And we typed, "Control C" and we got, "choonk… kachoonk… kachoonk…," we got the prompt back. First die!

**Hendrie:** First die?

**Faggin:** First time, first die that we…

**Hendrie:** First die you chose on the wafer?

**Faggin:** First die that we went down on the wafer!

**Hendrie:** Oh, my goodness.

**Faggin:** And that meant that 95%…

**Hendrie:** Oh, I'll bet you went out and… you quit and went out and celebrated then.

**Faggin:** That's exactly what we did… 95% of the chip, at that point, you can be guaranteed that it works because, just to do that simple thing, you had to execute…

**Hendrie:** Almost everything has to work.

**Faggin:** Yeah, that's right. So I went out and bought a bottle of champagne and then we celebrated that evening. Then, the next day, Shima started going through his stuff. But it took very little time, having gone through the first time, to really verify that everything worked. So at that point, it was clear that we were in business. In fact, out of that first run, we actually ended up selling samples of the Z80. One of the first Z80's, if not probably the first -- I think it was the first commercial Z80 that we sold -- we sold to Cromenco. And Cromenco was a company that then became a major computer manufacturer; they had a line of computers, computer boards and so on. They were quite successful, all based on the Z80, starting with that sample that they bought from us in State Street [downtown Los Altos, CA]. One of the other early parts that we sold was to NEC and, when I sold it, I kind of bit my lip, I couldn't say, "no," but I knew that they were going to copy it. And, in fact, a year and a half later, an NEC version of the Z80 came to market. And, of course, it took much longer than otherwise would have happened because we put a bunch of traps. I don't know if I did mention that?

**Hendrie:** No, you did not mention that. Maybe… Yeah, tell us about that.

**Faggin:** Oh, that was beautiful.

  

**Hendrie:** This is good.

**Faggin:** Between the first version of the Z80 and the final version, I told Shima: "Let's put some traps so people that try to copy it are going to copy also the 'mistakes.'" And they're "mistakes" that are invisible because, [since] we had depletion-load devices, we could make active devices depletion, which, of course, you would never do in your right mind, right? But because the process of implanting would leave no trace, you could create depletion transistors and nobody would know that they were depletion, expecting them to be normal transistors, just like normal transistors, right? So we created all kinds of fake logic that basically had the [active device replaced by a] depletion-load device that, whether you were driving it or not, was always a short to ground. So…

**Hendrie:** Yeah, they didn't do anything.

**Faggin:** They didn't do anything, right? So we did that and we put about nine traps scattered around. The one that was the most obnoxious trap was actually an error that Shima had made, a logic error that Shima had made, that was not even caught at the simulator. It was a logic error -- so that he actually made two logic errors [Laughs] -- a logic error that, basically, made a jump-on-condition not work if certain things had happened quite a few instructions before. It was a very obscure error that Shima was able to catch, instead of a customer catching. Typically, those kind of errors are caught by customers, because you don't even think of testing them.

**Hendrie:** To even try that, right.

**Faggin:** Try that [condition], right. So then, what we did was to basically fix the error by basically changing some of the transistors in the logic depletion -- the active transistor, so-called active transistor -- making them depletion mode so that the thing was corrected, the error was corrected, leaving the circuit [topology] the same as it used to be. And so there, anybody, anybody…

**Hendrie:** They would have…

**Faggin:** Would have thought it was right.

**Hendrie:** Yes. That they've copied it correctly and they'll put it into production and you have this little test loop and if you run it on it…

**Faggin:** And then they were done in a progressive way, in this sense: the first trap would inhibit all the communication from the internal bus so, basically, the chip would be deaf and mute and that would be it, right? So that you couldn't talk to it, you couldn't understand what was going on. So you had to go around it [with another silicon spin, losing about one month], and then they were progressively more subtle, okay? And then eventually there was this final crowning jewel of a thing, with this conditional jump issue. Anyway, it's funny because many years later, like in 1980 or so, I was taking an MBA class at Stanford -- a Summer Stanford MBA -- and there was an engineer there that used to work at Intel and for a number of years he was at National Semiconductor. National Semiconductor had copied the Z80 but moved it into CMOS. So they had made a CMOS version, for low-power application, of the Z80 and this guy was the project leader. And when he saw me, he said, "Damn you, Faggin…" I said, "What's your problem?" "Damn you, Faggin and your traps…" "Oh, you got caught in the traps, right? Yeah, I'm glad," I said. "How long did it take you to fix it? To find them all?" He says, "Nine months." "Wow," I said. "Nine months? We thought that a good guy could take six months to do it."

**Hendrie:** Very good.

**Faggin:** So, anyway, we know also that NEC got caught in those traps. So it bought us nine months to a year, basically. It was nine months to a year for the competition to come in and copy the device. To my knowledge, I don't know that anybody else has done something like this and we have used that [trick] in every other product that we've made.

**Hendrie:** Is that, right?

**Faggin:** Yeah, of course, yeah. So anyway, with this aside, with the working Z80 now we could take off so we, basically, by March, we had already identified a number of new employees that we wanted to hire in the company, but we were waiting for working silicon before pulling the trigger. So, soon after, the company went from 11 to probably 20-25 people.

**Hendrie:** You needed to hire, I assume, sales and marketing people? You hadn't done any of those.

**Faggin:** That's right, none of that. Sales and marketing; we needed to hire manufacturing people, we needed to hire all kinds of people who were product engineering people for the characterization and so on. Plus, as you remember, the Z80 was a family of components, so we wanted to start designing other members of the family. In fact, we had started already: we had a relatively junior engineer already doing the PIO -- the parallel input-output programmable chip -- so that it was already being designed during the end of the Z80 design [phase]. This guy helped initially with the Z80 and then he was freed up by the time the layout was finished, and so he had a few months of design already in by the time the Z80 was working. But then, of course, we needed to hire more people for the other chips, as well as to speed up the PIO [project]. That chip was called Z80-PIO. I still remember that in March we took advantage of an existing local show to show the development system and the Z80. So we had a sort of an informal announcement of the Z80 at this local show and then we hired a good PR and advertisement firm, and we started getting ready for the formal announcement of the chip, which happened in May of ’76. And there was a beautiful ad. It was: “The Battle of the 80’s,” and there were lead soldiers, some mounted in horses, fighting a sort of a battle scene in the bottom of the ad. And then there was a comparison between the Z80 and the 8080. So it was clearly an attack on Intel and I was told that the day that the ad hit the Electronic News [paper], Bob Noyce put that ad -- it was a double-page ad -- in [a wall of] his office and he was chastising his people for having lost the lead.

**Hendrie:** That’s very good.

**Faggin:** Anyway, around the time of the announcement, we also announced that Mostek would become our second source. We had signed an agreement. We had decided to have that because for a small company to get designed-in it would have been impossible. For microprocessors, you know, people had to make so much of a commitment in software development that they wouldn’t trust a fly-by-night start-up company. So it was necessary for us do that. So we did it and it was an immediate success. I mean, the Z80 caught the [wave]. Also it was a very special time, as you remember. It was a time where personal computers were beginning to emerge. The computer hobby clubs were all over the U.S. and it was a time of ferment. Lots of young people [were] getting enthusiastic about computers, so the Z80 became a darling of that time. And, over the next few years, it was designed-in in thousands of different products. So it was very, very successful and the company started growing quite rapidly. As a matter of fact, in February of ’76 we were 11 people and in late ’78 we were over a thousand people. That tells you the incredible growth that we had to go through to keep up with this. Originally, we were not thinking of having our own fab, although it was not excluded, but it would have been prohibitive for us to raise the money required to do that.

**Hendrie:** That wasn’t part of the original plan, at all.

**Faggin:** It wasn’t part of the original plan, but obviously it was clear that we needed to have a fab as soon as possible after the success of the Z80.

**Hendrie:** Because you were--

**Faggin:** Well, otherwise we would have been dependent on Mostek or Synertek, we would not have been able to be competitive and it would have been game over. So we were able to convince -- I was able to convince -- Exxon Enterprises to give us the funds to start a wafer fab. And that actually happened in May of that year, so that those months following the first working Z80 were months where many, many major decisions were made and a new course was set for the company. So in May of that year -- 1976 – I also hired Len Perham from AMD. He was a manager of process engineering at AMD, and I hired him as VP of manufacturing, and his first task was to design and build the factory. And so he hired a couple of people and we started in earnest to do that and, believe it or not, from scratch -- no building, no factory, nothing -- in January of the following year -- in eight months, seven months, whatever -- we had Z80's working, built in a [new] factory that was built up from nothing.

**Hendrie:** Wow. Oh, my goodness.

**Faggin:** In June we decided to start the building. We moved into the building in October of '76. Then all the equipment was [put] in and engineering lots were beginning in December and, as I said, by January, we had working Z-80's built in our own factory. And, of course, in those days, you could build a factory with the price of one [wafer] aliner of today. It's just unbelievable. Yeah. And then, by early '77, the building that we built was not large enough so we started moving down Bubb Road -- this was Bubb Road, in Cupertino -- and there were sufficient vacancies becoming ready for us that we just took over five buildings in Bubb Road. So we had a little makeshift campus of Zilog over the [next few] years because of the growth that was so enormous. Also in '76, we hired a person to start a factory in Manila and the Manila assembly line was done in June '77 and by…

**Hendrie:** Now this was… Was this a wafer fab or was this for bonding?

**Faggin:** No, no. Assembly line, for packaging.

**Hendrie:** Yes, for packaging.

**Faggin:** Yes.

**Hendrie:** Okay.

**Faggin:** So the Manila factory was done in '77 and by the end of '77, we had our own… [The Manila factory] was built up from scratch and the whole thing went up -- by the time we said, "Okay, let's go," -- in about six to nine months. By the end of '77, we were self-sufficient in manufacturing with both wafer fab and assembly. That's why in '78 we were over a thousand people because a lot of people were…

**Hendrie:** A lot of production.

**Faggin:** A lot of production people. Manila was probably 500 people. So it was the big chunk and then the fab in Cupertino was also quite a few people, probably a couple hundred people.

  

**Hendrie:** Can you give me some sense of how the revenue built in the first years and when you went profitable?

**Faggin:** Yeah. The revenues, It's very hard now to remember the numbers… Roughly…

**Hendrie:** Roughly, yeah, roughly.

**Faggin:** Roughly. So the first year, '76, we probably had a million dollar -- it was mostly development systems because people had to…

**Hendrie:** Yeah, people had to design the system and then go into production themselves.

**Faggin:** -- design the system and then go into production. Yes, so it was around that. And '77, I think, was already nine or fifteen, I mean, it was already…

**Hendrie:** Somewhere around there.

**Faggin:** And then… so maybe, let's say, nine. Then, the following year it was about 25.

**Hendrie:** Very nice.

**Faggin:** And then '79 was around 40, 45. Okay.

**Hendrie:** That's very nice.

**Faggin:** Yeah. And at that time, Exxon wanted us to invest. We actually never turned profitable in those years simply because we were spending 35% in R&D because Exxon were pushing us to accelerate things and so on…

**Hendrie:** To build market share.

**Faggin:** To build market share.

**Hendrie:** They weren't asking you to become profitable.

  

**Faggin:** That's right.

**Hendrie:** Okay.

**Faggin:** Yeah. And so… but, you know, we would have been profitable if we had a more realistic R&D [expense] of 10-15 percent. By the way, for the standards of the day, even 10 percent was high because, let's face it, the whole Z-80: [the chip], development system, software and so on, the whole thing, was done with $400,000. I spent $400,000 by March; from the beginning of the company until March, and so I had $100,000 left of the first half a million and then, of course, I had more money coming in, right?

**Hendrie:** Yes, but…

**Faggin:** But look at that, right?

**Hendrie:** That… Well, 11 people. Look at 11 people.

**Faggin:** That's right.

**Hendrie:** Look at doing product like this…

**Faggin:** That's right.

**Hendrie:** …with 10,000 transistors on it and getting it…

**Faggin:** Plus the system, plus the development system.

Q: …and plus the system and getting it -- and the software -- and getting it…

**Federico Faggin:** And getting it into production.

**Hendrie:** …into production with 11 people.

**Faggin:** But I was not in production yet, but it was, it was…

  

**Hendrie:** Well, it was ready.

**Faggin:** That's right.

**Hendrie:** It was done.

**Faggin:** Yeah.

**Hendrie:** We were talking about the revenue and that you didn't need to make a profit and some things along those lines. We'd sort of gotten off on that track as opposed to the development of the company, but I think we sort of came to a stopping place.

**Faggin:** Yeah, that's right. I'm sorry, I thought I was done with that question, yeah.

**Hendrie:** I think you are done with that. Maybe you could talk a little bit more about what you did…

**Faggin:** After the Z80?

**Hendrie:** Yes, after the Z80.

**Faggin:** Sure. Obviously, in the rat race that the microprocessor business is, if you have a successful product, you have to have a next one, and so on, right? So, soon after the Z80 [was completed], I hired Bernard Peuto who came aboard probably even as early as April of '76; and his task was to architect the next generation microprocessor, which was obvious that it was going to be a 16-bit; we were running out of room in the 8-bit world and…

**Hendrie:** Where did Bernard come from? Where did you find him?

**Faggin:** Bernard was a computer architect at Amdahl Corporation, he was a PhD, probably [from] Berkeley, but I'm not sure, I'm not positive. He was a computer scientist and he had a lot of experience in computer architecture, but particularly in measurements of efficiencies of instruction sets and so on. So his task was to create an architecture that was the next platform, with the best possible characteristics. That eventually became, the Z8000. Of course we knew that Intel was already, or would soon be, working on a 16-bit. In fact, what Intel did after we left, was to try to block us at different times by doing an upgrade of the 8080, which they called the 8085, and then a 16-bit microprocessor, which was the 8086, and then a 32-bit system that was the 432. I think it was basically an ill-fated architecture that was…

  

**Hendrie:** Yes. That's quite famous in the…

**Faggin:** Yeah, where they spent an enormous amount of money and nothing came of it.

**Hendrie:** Did they do it in Oregon or?…

**Faggin:** Possible. Yeah, yeah, possible. Justin Rattner, I think, was the architect and the guy in charge of the project. And also they did a single-chip micro controller. As you remember the Z8, which was [intended to be] the successor product to the Z80 in the low end, the Z8000 being in the high end, was the first product that I conceived, and at that time I called it 2001; it was a single-chip micro controller. Intel found out that we were planning to do that because I unwittingly told to a prospective engineer [from Intel] that I was trying to hire, that this was the plan that we had. So they basically started a project also to block us there, and that became the [Intel] 8048. Yeah. So, in the chessboard they had put gates all over, in the effort to stop us. At any rate, the 8085 was, however, no match for the Z80. It came out, I think, about six months after the Z80 and it was really, as I said, not a match at all. There was nothing there. It was just a clean-up of the 8080, basically a clean-up, meaning making it work at 5 volts, adding a few instructions, but there was nothing there. In fact, it was not particularly successful in the market. The Z80 really took over the 8-bit market, few exceptions made for a few high-visibility projects, like Apple -- they used the 6502 -- and some other high-volume applications, but very few, using the 6502. But the general market, the broad market, really went Z80. So then we started the Z8 and the Z8000 almost simultaneously, also ourselves, and we hired some [additional] people. We hired a project leader for the Z8 and Shima was going to be the project leader for the Z8000, with the architecture being done by Bernard. The architecture of the Z8 was based on my early work on the 2001 and then we assembled a team of engineers, it was like a committee, and it was completed by this committee of which I was part as well. Basically [out of it] came out the architecture of the Z8, which was quite successful. And Z8's and Z80's are still in high-volume production today, as you know.

**Hendrie:** That was the Z8…

**Faggin:** Yeah.

**Hendrie:** Instruction set compatible with the Z80?

**Faggin:** No.

**Hendrie:** It was a different instruction set.

  

**Faggin:** Different instruction set, yeah, because in the micro controller world it doesn't matter, really. You had very little [legacy software]. And there the whole point was [fast I/O control] -- and that was one of the contributions that I made with the architecture: I wanted to be able to operate directly on the I/O because normally, in a typical architecture, you take an input from an input port, you move it to the accumulator or to a register, and then you do some operation and then you move it out. With the Z8, on the other hand, I wanted to write a bit directly on a port or I wanted to be able to add to the port directly, so that the port became like an extension of memory, okay, an I/O port…

**Hendrie:** That's very similar to what Digital Equipment did in the PDP-11 with the Unibus where I/O locations were in fact memory locations so it was sort of homogeneous. Same idea.

**Faggin:** The same idea, yeah, and of course I didn't know this, but it made sense to me because that's what you -- in a micro-controller --that's what you do most of the time. You're basically banging bits on the I/O, right, particularly at the level at which things were done in those days. So to be able to do it cost-effectively, with the fewest amount of bytes in [the] program, as well as, as fast as possible, with the least amount of movements, that was the right thing to do. So that was the story of the Z8 and the Z8 came out… I think by the time it was finished it was '78. So it was in the architectural development phase for a few months, and then [the] design started soon after, and it took about a year and a half -- from the time the actual chip design started, after the architecture was done -- until we got samples. The Z8000 took quite a bit longer [than the Z8] for the architectural stage, and the Z8000 -- while it was a very good machine -- had one major architectural flaw: a decision that was made that was not good, and that was to use memory segmentation. So, memory segmentation was actually an Achilles heel of the whole architecture. We could have gotten rid of it, if the Z8000 had been successful, by evolving out of it, but we never got a chance because that [the Z8000] never quite made it against the competition of Intel. Intel, by the way, also had a segmented architecture, so it was not unique in that sense, but it was not the best choice. Motorola, with the 68000, chose a linear addressing and that was much more effective, because in those days the graphics applications were beginning to appear and they were beginning to be important, so we were basically not effective with our architecture in that space, and that was a major mistake. As it turned out, the Z8000 did not make it for other reasons, but certainly it could have survived as an architecture if it was not as limited.

**Hendrie:** Just added a level of complexity that fundamentally…

**Faggin:** Well, yeah, basically if you want to process an image, you have artificial boundaries that are created by these 64k…

**Hendrie:** Segments.

  

**Faggin:** 64k segments. They are way too small segments. They should be much larger segments. You begin to do a convolution operation over [an image], what do you do at the boundaries? It's a very messy problem and so it was really something that was not understood in those days, because graphics applications were just beginning to emerge…

**Hendrie:** And segmentation was the latest thing in computer architecture…

**Faggin:** Segmentation was a big thing in computer architecture. So we basically took more of a theoretical approach, as opposed to a pragmatic approach, to really understanding more of the application and the customer needs and so on. And I have to take my blame as well for that, although I was no longer in architecture at that time. I was running the company -- which at that time was more than I could handle, as a matter of fact. For a period of time we were also designing testers. We wanted to make a business out of testers which was -- with what I know today -- a silly thing to do because we were basically throwing away resources down a road that was not really that important to the company.

**Hendrie:** And you just couldn't afford to put your very brightest and very best people on the testers to compete with somebody that's just doing testers.

**Faggin:** Yeah. That's true, though they are very different people. We didn't put chip designers to make testers.

**Hendrie:** But still really great engineers still do wrong things.

**Faggin:** Yeah, still the wrong thing to do. It shows the lack of experience that I had and Ralph Ungermann had. Basically we were trying to do too many things. Part of that [problem was that] we were strongly encouraged to do more and more by Exxon Enterprises because they saw that we were so good at what we were doing, they were throwing too much money at us, in a way. And so, the combination of lack of experience on my part and lack of real leadership at the Board [of directors] level…

**Hendrie:** You didn't have any really experienced businessmen or venture capitalists.

**Faggin:** That's right.

**Hendrie:** To sort of ask questions…

**Faggin:** To ask questions and guide us. In fact, the opposite [was happening], they were misguiding us. And very often my first instinct was correct but then I was talked into things that I thought was expedient to do, when my instinct was against it. One of the big lessons, by the way, for me, that I learned during my tenure at Zilog, was to trust my own instincts. Instead of pushing them away, to actually allow my instinct to emerge and tell me what it needed to tell me.

**Hendrie:** Yes. Very good.

**Faggin:** At any rate, that [tester development] was part of the excessive R&D that we were spending as well, which was, in that case, unnecessary. But the company was doing well, despite the fact that it was very demanding to grow at the rate at which we were growing, particularly with the diversity of technologies and products that we were handling. Another thing that happened at the company, sometime later, was that Ralph Ungermann and I had a falling. This happened around ’78, ’79. It started around ’78, and, probably, by ’79 it was completed, and Ralph left the company, but basically…

**Hendrie:** What happened?

**Faggin:** What happened was that Ralph really wanted to be the CEO and I was the CEO, he was executive vice president, and his desire was to run a company, my desire was to run a company, so there you are: there was a conflict and it’s a natural conflict. Again we were not mature enough to manage it. I was not mature enough to manage it. I looked at Ralph as a friend and I was unable to confront him on this issue. So he was running behind my back and kind of rallying people [against me] and so on. Basically, I was unable to nip the thing in the bud…

**Hendrie:** And to sit down and say, “We need to come to resolution of this.”

**Faggin:** And to have sufficiently clear, unambiguous leadership to actually do it right. So, it clearly was something that I caused by my lack of experience and my lack of being in touch, again, with my own feelings. I was kind of pushing it down hoping that we could work it out and that kind of stuff, which of course I learned in subsequent years, that that’s not the way that you should do it. But it was also a difficult thing because, as I said, I was looking at Ralph as a friend so it was very difficult for me to confront him and…

**Hendrie:** He was your partner.

**Faggin:** He was my partner, yeah. But it got to a point where it was untenable and we had to part company, so he left. I asked the board to intervene because it was, basically, either I go or he goes, because it’s not a tenable situation in the company. That was crippling the company. When you have internal strife this is not good, it is not good, and that did not help the company. Then, on top of that, Exxon Enterprises decided that they were going to compete with IBM and -- of course, that’s what they intended to do all along, but I didn't know -- and become an information company. So they decided that they were going to combine the various companies in which they had investments, some of which they had started outright, some of which they had started, like ours, as a venture capital investments, into one organization. Of course I felt betrayed from that side as well because I was very, very clear with Exxon that we wanted to have an independent company, go public, go the classical route of a startup and they said, "of course, that's good, that's right, that's fine," but, in fact, we were betrayed. We were misled because that was their plan all along. [All] that was also weighting very heavily and now they had controlling interest -- by far controlling interest of the company -- so I felt straight-jacketed. Internally, having problem with my partner; externally, having problem with the only investor that we had at that time. Again I learned a lesson never to rely on one investor. Going back, I can actually see a time when I had an opportunity to change [the situation of having a single investor], although the consequences, if I had taken that road could have been the loss of the company, and I did not feel like risking the company. Basically, what happened was that, at the time of the second round of financing, when there was a fair amount of money that had to come in, because we also were setting up the manufacturing…

**Hendrie:** You had to raise a lot of capital to set up the manufacturing.

**Faggin:** Yeah. It was still modest [by today's standards]. I think we raised 6 to 8 million dollars, at the most. Nowadays you do it for just the first round of [investment of] a [modest company], but they were different times. So at that time there was some interest by venture capital because by early '76 the market conditions had changed. Seventy-five was impossible…

**Hendrie:** It was a terrible year.

**Faggin:** The total venture capital investment in the U.S. was $10 million, believe it or not -- in one year, total -- and we got half a million <laughs> of that 10. Everybody, all the venture capital had gone underground. Many had gone under and basically there was no money to go around. So at that time we were among the lucky ones to have found Exxon Enterprises to actually put money in our company. But then in '76 there was interest and Exxon did not want anybody [else] to come in. I insisted a little bit and they said, "No, we don't want it. Really, we have been good to you," and they basically played it that way. At that time my instinct, again, told me, "You should bring somebody else, just in case," although at that point I had not really understood very much about it because I was so focused on making the Z80 a product. In the first year of the company, I was just an engineering manager, like I was at Intel, so I had not…

**Hendrie:** The work you were doing was…

**Faggin:** That's right. Yeah. So, at that point had I said, "No, I do not want it. I want to bring somebody else," and if they had said, "Then we are not going to participate," I [could have] said, "Well, then forget it, we leave and we basically close down" or something. I had to be ready. In other words, I had to be [prepared] to go all the way.

**Hendrie:** Exactly, for them to say, "No, we're out."

**Faggin:** That's right. I had to be ready to go all the way. Again, these were the same symptoms, [the same weakness,] that later on with Ralph Ungermann, my partner, I did not go all the way. So it was the same problem, same issue. And of course, at that time I didn't realize it, I didn't recognize that, but that was part of my later growth.

**Hendrie:** But you didn't have a wise man on the board to tell you, "maybe you want to do this…"

**Faggin:** I'm not… I'm just…

**Hendrie:** Somebody who would already have been through this and the scar tissue.

**Faggin:** Of course. Of course. I'm saying it not to crucify myself but, at the same time, I'm acknowledging what happened. In fact, I have to say that I'm happy to be able to say it, the way I'm saying it now, because I have learned on my own, without a mentor in a way, out of the hard knocks of life, right, by having scars. <laughs>

**Hendrie:** You figured it out, yes. I call that a scar tissue. That's exactly what I call it.

**Faggin:** Yeah. Anyway, at that time, had I put my foot down and said, "No, I'm not going to go along with this, we've got to have somebody else," probably the course of Zilog would have been different because, in all fairness, the potential of Zilog at that time was to be what Intel became. I don't mean it with an attitude of neither sour grapes nor of blowing smoke. It's a fact because we actually captured the 8-bit market, and the 16-bit -- the Z8000 -- was substantially better than the 8086. It had the same [basic] architectural flaw that the 8086 had but it was a much better product, better architecture, a lot cleaner, cheaper, and so on. And the choice that was made by IBM, was predicated exactly by the very problem that I was talking about, which was Exxon Enterprises being the only investor. So we were basically Exxon, a competitor, and therefore there was a directive from the top not to use Zilog products at IBM. So that's why we were not even considered at IBM.

**Hendrie:** You weren't even in the race.

  

**Faggin:** We were not even in the race. That's why, had I changed course early on, and had Exxon only a minority investment at that point -- which was still possible, given that the big money was coming [in] with the second round… [we would have been in the race]. But I didn't see it and I didn't do anything about it. So that's again, other aspects of the story but important aspects of the story, because after all, we are talking about my experience and what I learned, not just technically, but also…

**Hendrie:** Exactly. Absolutely.

**Faggin:** By the school of hard knocks of life. <laughs> So, when Ralph Ungermann left, it must have been, probably early '79 or late '78…

**Hendrie:** Was the Z8000 out yet?

**Faggin:** The Z8000 was… If it wasn't out, it was about to be out.

**Hendrie:** Approximately. Okay. Fine.

**Faggin:** And then I replaced Ralph Ungermann with Manny Fernandez. He was VP of a division of Fairchild, and Exxon Enterprises had hired him, independent of Zilog, to run a company they had started. And when he went into this company, looked at the technology, looked at the potential, he basically said, "There is nothing here, we need to close it down." So, basically, he was hired to run it and he killed the whole thing and then, at that point, he was out of a job. Exxon actually pushed him a bit to take over and to take the job of Ralph Ungermann. So he joined the company and we continued on, and I stayed on until the end of 1980. I actually was not as effective in the last year as I was in the early years of the company, partly because Manny Fernandez was another one like Ralph Ungermann, who also wanted to run the show. And so, having gone from one situation to another identical situation, and I was getting tired of the Exxon game and what they wanted to do, I decided that I would move on. I was basically burned out also. Can you imagine? I had a perforated ulcer in '76 that just absolutely wiped me out for a little while. Fortunately, and this I owe it to my wife, [all ended well]. First of all, I didn't know that I had an ulcer, although I felt some hunger pains before eating, I didn't think much of it. I would eat and everything would be fine and then one night I was having dinner at home and, "Pooh…," like a stab, absolutely like a stab, and it was excruciating pain and I said, "Elvia, take me to the emergency." So we went to the emergency of Stanford [hospital] and they couldn't tell what was going on and they couldn't give me any painkillers. Everybody tells me that that's one of the worst pains [there] is to have a perforation of the ulcer, and not doing anything [about it], I literally felt like banging my head against the wall to put myself off. It was that excruciating, and because they didn't know what was going on they couldn't give me any painkiller, I was there for two hours. They took an X-ray and they couldn't see anything. So they were just waiting to see if it would work [itself] out, and then, after three hours, I started to develop fever. So it was clear that I had an internal infection and so they said, “We’ve got to operate,” and then at that point they suspect an ulcer, a perforated ulcer: the symptoms were all pointing there but it was not by no way…

**Hendrie:** Not for sure.

**Faggin:** Not for sure because they didn’t see anything in the x-rays. So they wanted to do exploratory surgery and my wife was very courageous and said… she basically confronted the doctor and said, “You’re not going to do a resection of the stomach,” because in those days, if you had a perforated ulcer, standard procedure was a resection of the stomach: they cut down part of your stomach to reduce the amount of acid that the stomach produces and it leaves you not with your full capacity, basically. She said, “You’re not going to do that!” The doctors took her seriously -- obviously they would have done it if it was essential, but it was actually [not so]. They do a vagotomy, typically, and they cut a piece of the stomach -- so, the doctor found out that it was an ulcer, they sew it up and they didn’t do anything [else]. So I owe it to my wife and to her confronting -- by the way, it was just the opposite of me, right, <laughs> at that time, <laughs> -- confronting style, to actually have… She basically made the doctor a little scared that he would have had trouble later on, if they…

**Hendrie:** Yes. Exactly, if he did something—

**Faggin:** And so, a few weeks later, I was perfectly fine and everything was done and, after a few months, I had no more consequences and that was the end of it.

**Hendrie:** They did go in and sew it up or they just…

**Faggin:** No. No. They went in and saw that it was in the pyloric orifice, that’s for some reason why they didn’t see it in the x-rays, and they just sew it up. That’s it. That’s it. So they didn’t cut anything. They didn’t…

**Hendrie:** They didn’t take anything out.

**Faggin:** They didn’t take anything out, yeah.

**Hendrie:** That’s the important thing—

**Faggin:** So I woke up just with a cut in me but no parts missing. <laughs> That was in ’76…

**Hendrie:** That was a sign that this was very, very stressful…

**Faggin:** Very stressful, yeah. Though, later on, I had recurrent bouts with ulcers but at that time I knew what to do, and also in those days there were new medications, like Tagamet and those things, that would take care of it. But eventually it was found to be of a bacterial origin. It's called a Helicobacteria or Helicobacter pylori, or something like that, and after they put me through a regimen of antibiotics, a very specific regimen, from that point on I never had anything. So it was [completely cured].

**Hendrie:** Well, wow. That's wonderful.

**Faggin:** Yeah. Anyway, this is a parenthesis but it shows [the stress I was under. In a sense, Zilog was for me an experience that was both exhilarating and excruciating at the same time because of the convergence into my psyche of incredible pressures, and incredible highs; being able to build a company from scratch: 1,000 people after three years, and at the same time all this stuff around. It was really living intensely <laughs> but perhaps a bit too intensely, and certainly without the level of guidance that would have been appropriate. So, by the end of '80, as I said, I left. I actually engineered my exit by showing the way to do it to Exxon and basically said, "Look, I want to get out but I don't want to hurt the company so why don't you…" In their new big organizational structure they had a group vice president position open for the computer systems group which eventually was the group under which Zilog would be one of the companies reporting into. And so I said, "Why don't you promote me there? Okay? I'll take that position. I'll try to map the course for Exxon for six months. I'll take a look if I like it. Chances are that I won't like it and then [in that case] you pay me, you buy my stock and I'm gone and everything will be quiet and will be happy. Right?" And they bought that [solution]. So I resigned from my CEO role, I took over that. Of course, [I stayed] in the board of the company [Zilog]. The company was reporting to me, in a sense, hierarchically, but basically the company was running on its own. I was a board [member and] my power was sort of behind the scenes, so to speak, but it was very clear after a month that my instinct was right, that…

**Hendrie:** This wasn't the kind of thing you wanted to do, in a company like Exxon.

**Faggin:** A company like Exxon. In fact, I had clear signals that that was the opposite of what I wanted to do. So, when my time was up, I got a big check and I became a millionaire before I was 40 and with my big check on tow I went home. <laughs>

**Hendrie:** At least they bought your stock, yes, but in some sense it was your only exit strategy, to get out and get some liquidity from your success.

**Faggin:** That's right. So, when this happened it was 1981, it was right around May or so of 1981.

  

**Hendrie:** How did they value the company at that time? Do you remember what they set as a market cap or how they figured it out because there was no public market for it.

**Faggin:** I know, I know. We, basically, haggled a little bit and I agreed to a price that seemed reasonable and that was it. Yeah. So there was not an independent evaluation. It was okay. They were not unfair to me at that point. It was okay. And it was the time when the marginal tax rate was 70% and…

**Hendrie:** Oh, yes. I don't really remember but yes, of course I remember.

**Faggin:** The time when interest rates were in the high teens and it was kind of funny because I remember the first thing that you do, of course, when you have something like that, you put the money in the bank, right? So I put it into basically what was the equivalent of the money market, or a little more than that in those days, and I was getting about $2,000 a day of interest. <laughs> So for a young kid like me that never had that much money…

**Hendrie:** That's pretty good.

**Faggin:** <laughs> So that's pretty good. I didn't know how I could possibly spend $2000 a day <laughs>

**Hendrie:** You didn't know what to do. Oh, wow. That's great.

**Faggin:** Yeah, and so I took off for a few months and went to Italy, just trying to decompress and trying to figure out what I wanted to do [next].

**Hendrie:** All right. Good.

**Faggin:** Now we are on vacation.

**Hendrie:** Exactly, ok. So yeah now what are you going to do with the rest of your life.

**Faggin:** Yes, what am I going to do with the rest of my life? That's exactly what was running through my head. My desire was to start another company, but at the same time I was telling myself, "Well, it's a lot of hard work. Is that what you want to do?" And so there was a period where I was taking it easy and I was vacationing. At that time I had a small child. No! Two small children, because my second [child, a] son was born in '79 in November. And my third child, a son, was born in November of 1980. So basically [in] early '81 I had a baby and a toddler. And of course a daughter who was born [in March 1970] just before I joined Intel [April 3, 1970], so she was eleven years old at that point;  yes, eleven years old in ‘81. So, it was good to spend some time with the kids and I enjoyed the boys immensely. That was good, it was a good time as well, but I was a bit restless. I didn’t know what I wanted to do and so that whole period was enjoyable but also somewhat concerning because I didn’t know what I was going to do when I was growing up.  Much of the summer was spent in Italy vacationing, seeing my parents and friends and so on. And then I started looking around toward the end of the year, and deciding what I wanted to do. To make a long story short, I decided to start a company in voice-data workstations. Basically in those days there was a lot of talk about creating an environment for managers because managers and consultants are basically communicators. They don’t do a lot of spreadsheets; some of their people do spreadsheets, but they communicate. So how to speed up and create an environment where communications are more effective -- so went the theory -- was a critical part of what would improve the productivity of that class of people. And of course, as you remember, in 1981 the IBM PC was introduced and it was almost an instant success. So toward the end of ‘81 I was trying to figure out if I could do something with it, somehow, something that was going to create a combination of a personal computer and a communicator for both voice and data. So that was the germ of the idea. And then I started talking to some people and by early ‘82 I assembled a team of people, got some financing and started Cygnet Technologies.

**Hendrie:** Now where did you get your financing from?

**Faggin:** The financing I got from Merrill Lynch Venture Capital. And the person, the senior partner in Merrill Lynch, was George Kokkinakis who used to work at Exxon Enterprises in the venture [capital] part of Exxon Enterprises, in the early days when Exxon was still much involved in the venture business. He had since left Exxon, also, because obviously they had stopped doing that. They wanted to create this monolith. And then I had another couple of investors coming in, but the lead investor was George Kokkinakis. The business idea was to create what we then called the Cosystem, basically an intelligent telephone that would bolt to a personal computer, an IBM PC, and together -- seamlessly communicating, talking to each other -- it would create an intelligent voice and data workstation. Actually, the idea was pretty good, except the timing was wrong in many ways, as we will see in a second. So, by the time we got started with funding and everything else, we were probably toward the middle of ‘82, and then it took about a year and a half to develop the product. In early ‘84 we introduced the product.

**Hendrie:** Ok, now who… tell me just a little bit more. Where did you get the team that was going to go do this?

**Faggin:** Ok, the…

**Hendrie:** You didn’t have a Shima that was sort of ready made right there.

  

**Faggin:** That is true. That is true. And I was moving away, I was tired of semi-conductor, obviously right? Because I had -- Andy Grove had asked me [in 1981] to come back to Intel, by the way -- a number of opportunities, obviously, in the semiconductor industry, but I was tired of the industry. I was ready for a change, for a shift, and I wanted to go more into the system business, understand more of that part. So, the marketing vice president came from one of the Exxon companies and I had met [him there]. He was the VP of marketing for Qwix, the company that was making the electronic typewriter, you may remember that, and so he joined me. The engineering guy had been the VP of Engineering for Two-pi Corporation. They had developed a computer that was -- I forgot now exactly -- it was a low cost version of an IBM computer, but I forgot now exactly [the model number].

**Hendrie:** I think I remember.

**Faggin:** Yes, Two-pi Corporation. And he was the VP of engineering, and basically those were the two key guys. Then I hired the finance guy, the administration and finance guy. He was the Zilog auditor.

**Hendrie:** Yeah, and they each are people you had met at various times.

**Faggin:** The classical thing, right? You basically hire people that you know, that you've been associated with and that you know they can do the job. So then, of course, we had to hire, to find some key guys because then, obviously, I had very little knowledge of telephony and communications. And so we found a guy that actually [became] the vice president of marketing -- what's his name? Klein. I always get confused with his first name because I know a number of Klein's… Jerry Klein, yes -- so Jerry Klein knew this engineer and we hired him from back East. Little by little we built a team. The team did not have to be that large but certainly more than eleven people. Because when you build a system you need software people, you need hardware people and so on. But anyway, the product was introduced at the PC Fair in San Francisco in early '84. And actually it won the prize for best product -- best peripheral product -- of the Fair. And it was a good product: It was well designed, well conceived and it was very useful. In fact, we were using it internally in the company and it was extremely useful. It had a form of electronic mail that was very clever. Basically, using commercial telephone lines, it would deliver the mail directly from co-system to co-system. A light to show you that you had mail [would turn on], so that you knew that you had mail. And you could make electronic mail with cc to other people, and the system would take care of sending the mail everywhere it was needed. If the line was busy it would try again, and so on, until it was delivered, giving you acknowledgement that it was received. It was a very good system. For example, with this system, you could -- in the middle of a phone call -- you could say, "Let me send you this information," and send to the other guy that you were on the phone [with], the screen that you had on the PC. So now you could talk about the data that you had in the PC and share the data. All of that in one phone line, with a conventional telephone. So it was a nice system. But two things happened. The first thing, just soon after we announced the product, there was the breakout of AT&T and the Baby Bells -- our system was really a system that had to be bought by the communication czar of a company, because you needed similar systems to create a network – and basically, the whole communications industry went into hibernation for a couple of years. Nobody wanted to do anything because there was confusion in the industry, because of the breakout of AT&T. That was definitely a very difficult environment to sell, all of a sudden. And, number two, [we had] a major recession. In ‘84, ‘85, as you remember, [there was a] major recession. So those two things combined with, frankly, a solution looking for a problem. Basically we had a great system, but the market wasn’t ready for it. People were still trying to learn how to use PC’s and so we ended up selling about 5,000 units the first year. So we sold some, right, but, really, we needed to sell 15,000 to begin to take off, that is to have sufficient funds to be able to then create the next generation of products, and really take off. So we basically never made it to that level of success, and we decided to close the company. Actually, we decided to sell the company. We sold it to… I forgot the name [of the buyer] now. It was a PC Company that bought the company. But, basically, it was not successful. It was sold in ‘86. We tried to survive by developing some software products, by having cut down considerably the number of people in the company, putting ourselves in survival mode, and seeing if the market would respond. But at that point, when you loose the initial momentum, it’s just very difficult… So, as I said, we tried for one more year past the point where it was almost certain that we were not going to make it. We tried to make it. We made a software product that would give us enough to survive, but basically there was no longer enough there. So we sold the company and I moved on and joined a start up company. As a matter of fact, I was a co-founder of that company while, at the same time, I was trying to wrap up Cygnet Technologies, selling the company and so on. I was consulting for this company and, as I said, a co-founder. This company was later called Synaptics.

**Hendrie:** Ah… ok. But it wasn’t at that time?

**Faggin:** It was not at that time, yeah.

**Hendrie:** What was it called, do you remember?

**Faggin:** It was called the Meno Corporation.

**Hendrie:** The what?

**Faggin:** Don’t ask me.

**Hendrie:** The Meno?

**Faggin:** Meno.

**Hendrie:** Meno?

  

**Faggin:** M-e-n-o. I think it's a Socratic something, it's one of those wet dreams or some VC.

**Hendrie:** Ok, very good.

**Faggin:** But, at any rate, this company started not by my instigation, although I participated, as I said, in the startup phase. And it started by a venture capitalist named Kevin Kinsella. They had a small, sort of like a boutique, startup fund. And his job was actually to put together the core team and help finance the first round or the seed round. So he was not even a seed-round venture capitalist, just a startup artist. He had read something about neural networks and his hunch was that there was something there. He had found a neurobiologist at UC Irvine who was, according to him, doing some seminal work in neural networks, and then he was looking for a CEO. His technical team was: professor Lynch of UC Irvine, one of his post doc and a, soon-to-be PhD, student – all were supposed to be the technical team, working of course all part time. And then he was looking for a CEO. And Lauren Yazolino, who was my VP of engineering [at Cygnet Technologies] -- who had since left the company because there was no longer a job for him in the company, as the company was winding down -- was interested in the position. So, Lauren, knowing that I was interested in that kind of area, intellectually interested and curious about it, asked me to go with him down to UC Irvine and help him evaluate whether there was anything there or not, by talking to Professor Lynch. And so I went with him. That was early '86, so quite a bit before I sold Cygnet Technologies. I went down and I listened to stuff that for me was very new. Obviously I never studied neurobiology and I knew very little, but it was kind of fascinating to hear about possible theories of learning with potentiation of synapses and so on and so forth. And so I got interested. Then Lauren asked me if I saw a potential. I said, "Well yeah, long term potential but I don't know enough about it. But I'm interested to kind of keep an eye on it, and at this point I could help you as a consultant, or something, in the meantime." Something of that [sort], I said. But then I got very intrigued. So, when I went home I made a computer program where I tried to make pattern recognition using the networks that Dr. Lynch described -- networks that are in the orthorhinal cortex -- by using the potentiation rules for synapses that he had found. I actually could make something work: I could actually recognize patterns and so I kind of got excited about it. I thought, wow that's an interesting thing, right? This [work] actually gave impetus to this whole company to start. Because before it was just neurobiology but there was no way to see how it could really do anything [useful]. And then with the work that I did -- I didn't know anything about the work that had been done already in the industry or in the research community -- it looked like there was some potential. So that was enough to start the company, and begin to put a business plan together and see if we could get funding for the company and so on. So I continued to work, sort of part time, to try to understand not only what it could do but also how could I implement that. So I came up with the idea of creating synapses by making synaptic weights by using floating gate structures, where you would store analog voltage in a floating gate structure and finding ways to make circuits that would allow this learning to occur. That was the initial work that I did, and that was enough to actually get some seed funding by TVI, Technology Venture Investors that, by the way, was a very successful [VC]. It has since been dissolved as a venture capital [firm], but that was the only venture capital that invested in Microsoft. They turned, I don't know [exactly], ten million dollars into a billion dollars -- something like that. And then, the Sprout Group, was the other venture capital. We raised a little bit of money, just to study the potential of this technology.

  

**Hendrie:** Who was the venture capitalist from TVI?

**Faggin:** He was Marquard, who was in fact…

**Hendrie:** Dave Marquard, the one that…

**Faggin:** Invested in Microsoft. And I still believe he is in the board of Microsoft.

**Hendrie:** He also invested in Sun, that was on the board of Sun and then he had a little conflict he had to get off one. And he chose the right one to get off of. He's still on the board of Microsoft.

**Faggin:** So Dave was the one, and then Keith Geslin was the one from the Sprout Group; he was the partner involved. Although Dave was the key guy to make the decision, to help make the decisions, he did not come to our board because he had too many commitments. We had Jim Bocknowski, who was also partner of TVI, and later on left TVI and started his own fund called Delphi Bioventures. So Jim Bocknowski and Keith Geslin and Kinsella were the people from the venture capital [side] in the board. And, for a while, I was consulting and helping trying to figure out what we could do. We were supposed to get a lot of help from the people down at UC Irvine, but there was very little coming from their way. It was a lot of talk but nothing there. So, of course, Lauren, being more of an engineer than a scientist… [was not into theoretical studies]. I'm a physicist, so I have a little bit of a scientist in me although I've done mostly engineering work. But I'm interested in more speculative things, and so I was interested in this kind of stuff. I started studying biology and, actually, it was a very fun period of my life because I got to study a bunch of stuff that I loved to know about.

**Hendrie:** You had a chance to learn a whole bunch of stuff, right?

**Faggin:** It was great. But at the same time, from a business point of view, I just couldn't see any way that the technology could achieve anything useful in any short [amount of] time. At the same time, Kevin Kinsella kept on pushing to get a product out the door and I said, "You've got to be out of your mind. We don't even know what we have first, and you want a product? What are you talking about?"

**Hendrie:** Yeah, what product is it that you want?

**Faggin:** So I started looking around to see if I could team up with somebody. And then I found out that Carver Mead at Caltech was working on neuromorphic structures, mostly, in his case, not for learning but for sensory inputs. So he was working on the silicon cochlea and silicon retina, you probably have heard about these things. And, since I've known Carver since my early days of Fairchild, I went and visited him and he showed me around [the lab] and he showed me what he was doing. And I saw the potential synergy there because, obviously, you need sensors which are neuromorphic to match with neural networks that are supposed to do the recognition of the patterns that are created by the sensors. And so I felt there was synergy. Obviously Carver knows technology and science. So he was the right person to team up and I asked him if he was going to join the company -- obviously as a consultant, doing a day a week -- see if he could do that. And in October he decided to do that and he came on board. So we started a collaboration, working together. Then, by early ’87, it was very clear that the company, the way it was started it, didn’t make any sense because the people at UC Irvine were basically not contributing anything useful to our technology. And at the same time Kevin Kinsella kept on pushing to get products when we didn’t even have a technology. And so I basically went to the board and I said, look either we restart the company on the right track, by going into a research mode for a number of years and figuring out if there is anything here and then, if there is something, then we can decide what product to make and go down that road, or else we have nothing and we should close it down; there is nothing here, ok.

**Hendrie:** Yeah, there is the germ of some ideas but it needs research to turn it into…

**Faggin:** To turn it into a technology and then into a product. So I said, “If you guys don’t want to do it its fine with me, I’m going to go on my own, working with Carver -- not necessarily starting a company.” Just that I wanted to -- because I was very interested in the science -- understand how the brain works and all this other stuff. So I was kind of captivated by all that. And so I said, “I don’t need to work. I’m self sufficient money-wise. I can take a few years off and work on my own, with Carver, trying to develop something.” So the investors said, “Ok, let’s think about it.” And they decided, “Yeah, you are right. We want to invest, you’ll restructure the whole thing.” And basically we restarted the company. Lauren left the company, the UC Irvine team left, and basically it was Carver and I and we decided to hire a few people, go with a few, young people that didn’t know it couldn’t be done. And basically, take our time and do our research and see if there was something there. We did that for six or seven years. And we used very little money. We diluted the money that we received from the venture capital with contracts that I obtained, for example, with the US Postal Service. We made a bunch of studies of pattern recognition for characters, for address labels, and we did a very good job. And it was almost like I had to plot a course of what we needed to do to have a sufficiently robust technology to solve real problems. And when there was a problem -- looking for problems that were on my road map -- then go finding either the company or the institution that needed help, obtaining a contract and then doing the work. And so, in that way, we were not loosing time down a dead end or roads that were not important, simply because we needed to have some money: It was actually part of the roadmap to do that [project]. And so that way we went reasonably straight toward the goal of creating a chip that would do sensing and pattern recognition. The culmination of that effort was actually a number of chips, not just one. One chip was a check reader that was made by one analog chip that had an imager, two neural networks and all the rest of the analog system, plus an interface from that analog system to a micro controller. And so the combination of a micro controller and this chip would do all the tasks of reading the characters in the bottom of a check, by reading them optically. And we achieved five-9 accuracy in that task, which took a lot of time to actually test it out for that accuracy. And it was a very, very, very good piece of engineering work. We also did a chip that was a neural network with floating gate structures -- floating gate synapses -- that you could reconfigure so that you could create a variety of neural networks. And that was done under contract with the US Postal Service. But as we were getting to that point it was clear that, despite the scientific and research value of that [approach], that in fact, the value added that that technology was bringing was not really that great, ok. So at that point…

**Hendrie:** Was a unique way of doing it but there were other ways, and it wasn't doing a job that couldn't be done any other way.

**Faggin:** Yes, in fact, you could do it in software. By doing in software the algorithms that we had tested and found and in some cases discovered, given that in the meantime the progress in digital electronics was tremendous, the whole approach of Carver -- that was analog VLSI, analog computation and so on -- proved not to be viable. Also it lacked the flexibility that you have with digital electronics; every chip is unique and if you change your mind later, too bad, you've got to design a new chip, which is not an easy thing to do. And of course it is much easier to find digital designers and software programmers than it is to find analog designers.

**Hendrie:** Sorry. Analog chips are...

**Faggin:** Are very hard. So when you put all that thing together, that whole approach was really not the correct one. So I blew the whistle and I decided, "hey, good, we have learned a lot, now we have a team that is…"

**Hendrie:** We proved your theory does work, but that doesn't mean it's the best way to do it.

**Faggin:** That's right. And so basically at that point I decided to pursue a different course. Of course, time had gone by and now we are in 92', 93'. And so five years later, five or six years later, it is time to show some return to the investors. We've got to begin to worry about what are we going to do when we grow up. So, at that point Carver was not too happy because he was very much in the direction of analog: analog is the solution and so on. But myself and my team – and the team, really, wasn't mine, it was our team -- was turning against that whole approach, and it was very clear that we had to change direction. In the meantime I had started a project… Because I was on the board of directors of Logitech, I had learned that -- Logitech had done track balls for notebook computers, the early track balls, as navigation devices, as you remember -- they were not very reliable: it is very simple because fingers are greasy and when you move the ball, the ball gets a layer of grease and the ball cannot entrain the little wheels anymore and so [the cursor] skips all over the place. Plus it's bulky and it's heavy and so on and so forth. So I got the idea to try to figure out if we can find a better solution. And out of that, out of a total group effort -- this is actually one of the few times I have experienced where an idea evolved [through group dynamics] to the point where, the final idea, is not possible to tell whose idea it was. Because it went from person to person…

  

**Hendrie:** People constantly making contributions.

**Faggin:** Contributions, incremental contributions and the end result was very different than what the starting point was -- the starting point that I gave -- which was [just] a starting point. I correctly identified the problem to solve, which is half of the battle, but [as for] how to solve it, my original thought was not as good as what we ended up with. And so, basically, out of that work came the idea of the touchpad for notebook computers. So, because of my connection with Logitech, I went to Logitech and I said, "How about working together? We have this idea. We have tested it out. We showed that it works. It's our intellectual property, but we can develop a chip for you, a mixed-signal chip, that can do the -- mostly analog -- computation that is required to do a touch pad. Are you interested?" And they said, "Oh, yeah, we are interested." Because of course they knew that track balls were a losing proposition. So we entered into a development contract and we agree on a price and so on. And so we started working on this thing. We were working on the chip and they were working on the software and the rest of the system. So now we are in the middle to second half of 93' and we have developed a chip that was, basically, a proof of principle chip, and things worked, so they wanted to negotiate the final agreement to do the production chip. And at that point Logitech started asking prices that were completely [rock bottom], basically there was nothing left for us.

**Hendrie:** They weren't going to… They were going to give you a minuscule royalty.

**Faggin:** Basically it was like if we had been consultants or something. But because the contract that I wrote was very good, the IP was ours because we had done that work already before, so we had full ownership of that.

**Hendrie:** Now is this Pierluigi?

**Faggin:** That was Bobo, actually. So at that point -- we probably by now are in November or something -- I still remember the breakfast that I had with Bobo in the restaurant of the hotel in Menlo Park where, basically, he was leaving nothing on the table [for us]. And I said, "Bobo, sorry, I'm going on my own. I'm going to resign from the board of Logitech from now on… And goodbye!" And that was it. So then I decided that we were going to go down that road: we were going to do a touch pad ourselves. A year later we introduced the touch pad, in 1984 [correction: 1994], at the Comdex -- Fall Comdex -- in Las Vegas. So that was how the transition occurred for Synaptics.

**Faggin**: So, as I said earlier, in November 1994 we introduced the Touchpad for personal computers, for PCs, at the [Fall] Comdex. But the story is a little richer because in May of the same year, the end of May, Apple came out with their PowerBook with a Touchpad. And then we found out, "wow we're not the only ones with this idea." Okay. And then we soon found out that the Touchpad of Apple was licensed from technology of a small company in Salt Lake City, called Cirque.

  

**Hendrie:** How do you spell that?

**Faggin**: C-i-r-q-u-e

**Hendrie:** Okay.

**Faggin**: Okay. And so the story unfolded, eventually, to the following: Apple had been contacted by Cirque. The founder of Cirque was the inventor of this thing and he had been working for a number of years to make his idea work, way before we even started. And he went to Apple and Apple was interested, but there was no VLSI implementation yet. So it was, basically, a breadboard demonstration.

**Hendrie:** Hmm. Hmm.

**Faggin**: And because this company did not have the resources to do a VLSI implementation. So Apple took a license in the technology and developed their own chips. There were three or four chips.

**Hendrie:** Okay.

**Faggin**: Three or four ASICs in this touchpad. And where our implementation was two chips – it was an analog chip -- a mixed signal chip -- and a microcontroller that we were purchasing from the outside.

**Hendrie:** Okay.

**Faggin**: And, of course, our plan was to go to a single chip, but this was the fastest way to market, to test the idea and see if it was working before we would put a lot more energy and effort and money into the project. So it was a staged approach. Because of Apple, though, and the fact that Apple was successful in the marketplace, in other words the Touchpad was…

**Hendrie:** Was accepted.

**Faggin**: Accepted by the marketplace.

**Hendrie:** Yes. You had a market test.

  

**Faggin**: We got lucky. We got lucky! Right? So, thank God for luck every once in a while: we basically had an already made market in front of us. Instead of having to pioneer the market, we went out for the much larger PC market with a solution that was actually better than Apple. It was more robust than Apple's, cheaper than Apple's because we only had two chips instead of three or four, and only one was custom, the other one was a conventional microcontroller.

**Hendrie:** Right.

**Faggin**: So we got designed-in immediately. We had samples in the July timeframe. A few months after Apple's [announcement], actually, we already had samples of the Touchpad; we already had started the design-in activity [with customers]. So, by January, a couple months after the official announcement, we were in production. And by August of that year, we were making money, we were actually profitable as a company. So it was a very, very rapid success, because we were attacking a relatively large market, [therefore] the ramp-up was quite fast and...

**Hendrie:** Yeah. And there weren't any good solutions.

**Faggin**: It's right, there were no good solutions.

**Hendrie:** IBM's little stick and that was more or less it.

**Faggin**: That's right. Yeah. The IBM stick had been introduced… I think it was November the year before.

**Hendrie:** Yes.

**Faggin**: November '92. No, November '93. So, there we are: we now have a product; we now have a market. And that was the beginning of the success of Synaptics as a company. Of course, as we were developing the Touchpad, we went to the end of every project that we were in the middle of, the research project that we had. Some of them with the U.S. Postal Service, some with other organizations, so that we cleaned up all that, so that we could focus entirely, move the company entirely in the [new] direction -- which was done. And that was the beginning of a successful history for Synaptics. But it was interesting because, in the meantime, Logitech decided to also do their own version [of touchpad] and go down that same road and compete with us, which, of course, we expected. But they made two fundamental mistakes. The first mistake was: they went all the way to a single-chip because they knew we were going to have a single chip…

  

**Hendrie:** They said, "We'll take a short cut and we'll get out ahead of them."

**Faggin**: Right. So, number one. That, of course, made the project much more difficult. Second error that they made was that they decided to use an implementation with ROM instead of with what we used, which was OTP, One-Time Programmable ROM. So, we could much more rapidly evolve the firmware.

**Hendrie:** Yes, of course.

**Faggin**: Okay. And of course, had we not done that, we would have been dead because there were so many things that needed improvement and that you find out only when you're in high volume [production], because it's not a digital system, so when people use it, you have all kinds of situations that occur that you will never in your life predict.

**Hendrie:** Right.

**Faggin**: And the only way around it is to change the code.

**Hendrie:** Right. Of course. Well, hopefully that is the way around it as opposed to changing the analog chip.

**Faggin**: Yeah. And that tells you that the analog chip is not… You cannot do anything with it. You couldn't possibly evolve it [fast enough]. Okay. So, we basically went down our path and the fact that we had OTP was the blessing, and the fact that we had lots of problem was the curse. But we stayed with the game; we never lost a customer or lost a beat, although we were [going] from crisis to crisis during the first couple of years, which delayed our design of the single chip, because the design of the single chip was supposed to start immediately after the introduction of the Touchpad. We could never get there because we kept on having issues and problems, the engineers were always thrown back into solving problems on a crisis mode, which was [requiring] a lot of energy from everybody. And so in the meantime, about a year and a half after our introduction, Logitech showed up with their own single chip [touchpad]. It was not particularly good. Clearly they had long ways to go because in the meantime we had evolved our technology quite a bit beyond the initial one, but they started bombing prices. Their strategy was to basically take it away from us.

**Hendrie:** Buy the market.

**Faggin**: Buy the market. You know, selling below cost and…

  

**Hendrie:** Cause they had a whole bunch of other products and they financed it with the mouse.

**Faggin**: That's right. So, initially the industry did not take them seriously because they were not working very well but they kept on improving and so on. And two years later -- so we are now, let's see… By '97, they began to have working Touchpads.

**Hendrie:** Now had they taken away any of your major OEMs?

**Faggin**: Nothing. Nothing yet.

**Hendrie:** Nothing yet.

**Faggin**: Nothing, nothing yet. But they were…

**Hendrie:** They were out there trying.

**Faggin**: They were forcing…

**Hendrie:** And it's forcing your prices to be lower.

**Faggin**: Forcing our prices to come down faster than otherwise would have happened, even if they were not able to deliver. Although it was not too bad because we were simply holding our own saying, "Okay, buy from them," knowing that they couldn't deliver. So, by '97, we were now designing the single chip. Finally, we got around to doing that. But there was a period of almost a year of overlap where we didn't have a single chip yet and they were actually delivering product at prices that were really giving us a lot of hard time. And they were gaining market share, they were beginning to take away customers.

**Hendrie:** Was there any reason in your view to believe that they would have lower costs once you got your single chip?

**Faggin**: Oh no.

**Hendrie:** But with two chips, they could, they probably did, if it worked, it had lower cost.

  

**Faggin**: We were also learning how to reduce the cost, we were moving the manufacturing to Asia, we were doing a bunch of stuff to reduce the cost, of course. Right? They already had manufacturing in Asia, in China, and so on and so forth. Basically it was an ongoing company, they didn't have to go through the startup phase that we were going through. And we were going through it all with our own money, without any additional funding, basically using the money that we were creating. Fortunately we were cash flow positive even despite the rapid growth and we were not [cash] starved and we could invest but obviously it takes time too. So, for a period of time between mid '77 to [the] early part of '78…

**Hendrie:** You mean '97?

**Faggin**: Yeah, sorry '97 to '98. So, for about nine months to a year period of time, we suffered a bit. We actually lost some money, not much. We lost some money because we had to meet prices not to lose the business, and they were too low to have enough margin to pay for all the R&D that we were doing to develop the chips and so on. But our strategy was to develop a proprietary chip instead of having, like Logitech did, Zilog develop a chip for them. And so our costs would have been much lower than what they would have ever paid at Zilog. We also developed a very specialized microcontroller that would be much better to the task that we had in front of us than using a conventional microcontroller.

**Hendrie:** So, you didn't use a conventional core?

**Faggin**: We didn't use a conventional core. Also that [strategy] would allow us to have proprietary [code]: nobody could ever tell what we were doing because the firmware was in OTP, One-Time Programmable, so there were no physical traces of the code. Of course, we would program the chip. then basically, lock the system and throw away the key so that you could never read out what was there anyway. The only way to tell what was there would be to actually look and see if you could see any differences between ones and zeros in the OTP.

**Hendrie:** Yeah. Exactly.

**Faggin**: But, you can't.

**Hendrie:** You can't? Okay.

**Faggin**: You can't see electrons and, of course, then you have to figure out what instruction set we have in the [microcontroller]…

**Hendrie:** Exactly.

  

**Faggin**: I mean it would be just a huge job. You would be much better off starting from scratch.

**Hendrie:** You're much better off to pay one of your engineers a million dollars to have him sneak the information out. I'm sorry, but…

**Faggin**: Yeah. That's another way of getting it. So, for all those reasons we took a longer road, but the only road that we had that would give us a cost advantage over them. And at the end, as you will hear in a second, it paid handsomely. But we suffered for that interval of time when they were pushing hard and trying to break our back and we were not ready, we were still developing the single chip. By the middle of '98, however, we had the single chip and we were beginning to ramp-up production of [new] models because when you are in production with a certain model, you got to go through the lifetime of that model because you cannot re-qualify a different model with different electronics; it doesn't work that way. So, basically, the introduction of lower cost versions with the single chip had to come with new models.

**Hendrie:** Yes. You aren't going to retrofit a model. That's a fool's errand for the manufacturer, for the computer maker.

**Faggin**: That's right. So, basically, it took about a year to go from when we begun to have Touchpad with a single chip -- it was in the early part of 1998 -- until we were completely converted: It was at the end of '98.

**Hendrie:** Okay. Now who was your biggest OEM at this time that's using this?

**Faggin**: Well, we pretty much had all the big U. S. guys, Compaq and Dell and, essentially, in many of those we were a single source. So, we were doing okay. But there was a lot of pressure by the ODMs, Taiwanese ODMs, to lower the price. We didn't want to and so they were designing-in Logitech. Then Logitech tried to do one more thing. We were a little reticent to do custom designs because it's a lot of work, generally you don't get paid for it, and you don't even pay for your NRE and so basically it's a lot of effort.

**Hendrie:** Yes.

**Faggin**: But because Logitech had a major operation in Taiwan, they could work very closely with the customer and so it was a way to sneak around us and get locked in with a custom design. It was actually a good strategy. And when the idea [of a custom module] would come up from time to time, basically pushed by the customers, our people were against [it] and I was ambivalent: I saw the strategic value of that but it was a lot of work too, so I was ambivalent. But what happened was that, in June of '98, a customer came in with a Touchpad -- it was a Logitech Touchpad -- they had to go into production in July and the Touchpad didn't work and they were in panic. They said, "Can you help us?"

**Hendrie:** Who was this customer? Can you say?

**Faggin**: I think Compal or Arima. Yeah, but I'm not sure though.

**Faggin**: One of those guys. And so at that point we realized the extent to which the damage was being made by Logitech. It had not taken place yet, but they had a bunch of designs out there going on that we didn't know because, of course, their customers didn't want to say anything so it was really kept very much in stealth mode. So we would have been surprised later on. But I took that opportunity. I said, "Okay, yes, we're going to do it and we'll give you a Touchpad, an engineering prototype, in a week, but you have to commit to this price." And it was a price higher than what they wanted but, hey, you know, you want it or you don't? So they agreed to the price and we delivered a week later a working prototype. We made few mods, mostly in firmware, and it went in production. And at that point, as we realized what was going on, we went around everywhere and we converted every single customer. And I made a decision that we were going to be in the custom [business] because at that point it was clear that we had to do it. So, we went around and got rid of every single custom design because, see, the Chinese talk to each other, so other guys got scared because of what happened to this customer: they had a Touchpad that didn't work and they were about to go [to production]. It was a crisis. So, we got all the designs back and we completely killed the business of Logitech. Within a few months, they got out of the business. My estimate is that they poured about $40 million into this whole thing to try to kill us and in the end we were exiting [the confrontation] by lowering prices and increasing gross margins, as we were converting to the single chip all of our products. And in the end of '98 we were at 40 percent gross margin, in a business that normally gets 20 percent.

**Hendrie:** Yes. I was going to say that's remarkable for a semiconductor product.

**Faggin**: Yes. That's the story. And that was a major satisfaction for me because it was really brutal. Basically, they took advantage of us in many ways and then, of course, we could have gone through legal means but then nobody would have won anyway. So, what's the point. So, first you win the battle and then, the rest, who cares.

**Hendrie:** And then they can say when they really reflect, maybe we should have paid a fair royalty and we would have had the whole market.

**Faggin**: That's correct. And by the way, you see that V500 cordless mouse there [on top of the credenza] which has one of our touch strips. So now Logitech is a customer of Synaptics and that gives me a lot of satisfaction.

  

**Hendrie:** Is the CEO still there?

**Faggin**: There's a new CEO.

**Hendrie:** There's a new CEO.

**Faggin**: Yeah. Bobo -- it's the nickname of Daniel Borel -- who was the CEO at that time and also a co-founder of the company. He is Chairman, I believe, at this point. But, anyway, so that's the story.

**Hendrie:** Oh, that's a wonderful story.

**Faggin**: And then at that time a number of us were pooped. I certainly was pooped. It had been a pretty long struggle.

**Hendrie:** Very intense.

**Faggin**: Very intense. From '93 until the end of '98 they were very, very intense six years for me and I needed to take a rest. I had started about a year before to look for a CEO to replace me and finally I was lucky to find one and we made a deal at the end of the year and he joined in January of '99.

**Hendrie:** Who was that?

**Faggin**: Francis Lee.

**Hendrie:** Okay.

**Faggin**: And he has been President and CEO since that [time]. I kicked myself upstairs to Chairman and I've been Chairman, working a day a week for a number of years and then, since I took over Foveon, as CEO, now I'm simply…

**Hendrie:** You just go to the board meetings.

**Faggin**: I go to the board meetings. And I go to lunch once a week or twice a month with Francis, to stay up on how are things going.

  

**Hendrie:** On what's going on.  And did you sell your stock or did you…

**Faggin**:  Of course not.  I sold a little bit but most of the stock I still have and is doing quite well. Thank you.

**Hendrie:**  Very good. A total reversal of the Zilog story.

**Faggin**:  That's right. That's the way it ought to be.

**Hendrie:**  Innovative product.  It won huge market acceptance.  And you made money.

**Faggin**:  Well, I made money also at Zilog, but not as much, and certainly the potential that was there… In fact, it's amazing, the potential of Zilog was 20 times,100 times the potential of Synaptics and where Synaptics is realizing its full potential, Zilog didn't – which is interesting. But, hey, that's how you learn.

**Hendrie:**  And it seems to me this is a story that repeated itself over and over again that large corporations are not capable of making a success out of a start-up and I think that is because the management has a different skill set to move up in a large corporation and a different psychology and to just apply it to a start-up it kills it.

**Faggin**:  So now I am running Foveon.

**Hendrie:**  So tell me about Foveon. How did you get involved in Foveon?

**Faggin**:  Foveon is a spin-off of Synaptics and National Semiconductor.  So Foveon started by virtue of work that was started at Synaptics and that continued after the decision of Synaptics to go after the touch-pad.  Carver [Mead] was still interested in pursuing imagers and the silicon retina and I encouraged that, and he kept a low level of effort in that direction. Then, toward the end of '94, we had National Semiconductor invest in the company [Synaptics], at the time that we were introducing the touchpad.  So they came in and they actually took most of that round of financing, it was money required to ramp up the business as we were going in production shortly [afterwards], mostly for working capital needs of the company. National was interested in buying the company at that time, but we decided not to sell, we decided to stay at the level of that relationship [as an investor].  And the interest of National was also in many other projects that had been going on in the past at Synaptics: trying to get something useful out of all that energy that was spent and the results we had achieved in analog design and so on.

**Hendrie:**  Were they a foundry for the touchpad chips?

  

**Faggin**: No they never were. So sometimes later, after they came in with the investment, we decided to start a joint project and Carver was leading the project, looking at CMOS image sensors. So Carver worked for a number of years and basically we had a few guys at Synaptics but they were really working for Carver -- for convenience, we were hosting them -- and Carver was mostly working at National with some of the National guys helping to develop new technologies and then it was very clear that, unless that activity was really spun off, that nothing would come of it and so a decision was made by both Synaptics and National to actually create Foveon, and that happened in mid ’97. Carver was the initial CEO of the company and I was a founding board member and there were members of the board from National. So it was National, us -- Synaptics -- and Foveon, and that’s it. And, I think, though, there is not enough time to go through the Foveon story, so we may have to continue some other time. But we have at least 15 minutes that we can go.

**Hendrie:** Shall we do the remaining 15 minutes we have on the tape?

**Faggin**: Yes, so the original idea of Carver was to develop a high-end camera, a camera for photo studio using a prism and 3 imagers so that you could split the color. The prism would split the color, would split light into RGB and each of the sensors would sense one of the colors.

**Hendrie:** These are different physics of sensors than CCDs?

**Faggin**: Using CMOS sensors.

**Hendrie:** Using CMOS sensors, not the CCD principle, the shift registers…

**Faggin**: Not the CCD principle. In those days there were a number of people and companies working on CMOS sensors so it was, by a number of people, perceived as a potential new avenue to make sensors. It’s like random access [memory] versus serial memory; it is a standard technology so that you could integrate other things in that chip; it is a technology which is more widely available and evolving more rapidly, so all of those factors make it attractive and it was a way for new entrants to get into a business that at that point was totally in the hand of Japan Inc. There was a little bit of Kodak, but Kodak was not much of a player, and so it was mostly the Sonys, the Sharps and the Sanyos, etcetera, of the world. But Carver wanted to have a camera that did not have the limitations of conventional Bayer pattern, the color filter arrays: red, green, blue and green. It is a pattern of four [pixels] where [there are] two greens, one red and one blue and that pattern repeats and is called, Bayer pattern, and that’s the pattern of color filters which is on top of the sensors of a classic [design] both CCD or CMOS, that’s a classical way to get color [information].

**Hendrie:** Same scheme that is used on VIDICONs.

**Faggin**: Yes of course. And basically, there is more green because that's the luminance channel, etc., etc., but that's based on solid psychophysics. But, of course, you have all kinds of sampling artifacts because at each point of the image you only measure one of the three colors, and the other two [colors] need to be guessed by looking around at the colors around that location. So, if you have a sharp edge? What is the color?

**Hendrie:** On each side of the edge?

**Faggin**: On each side of the edge. So when you do averaging to find the blue, say, it is very clear that you have all kind of artifacts, you also have two sampling frequencies because two out of four are greens, one out of four is blue and one out of four is red, therefore you have different sampling frequencies so you have horrible artifacts in the color. That's why digital photography, particularly in those days, produced pictures that were not very good compared to conventional photography. So, by sensing at each point all three colors, by the way of using a prism, it was possible to reduce, to remove these artifacts and that was the basis of that camera. And again the idea of splitting the light with a prism to do this was already done before so there was nothing new there. The novelty was to have CMOS using this embodiment and building a camera. The camera came out, it produced certainly better picture than could be produced in those days with the best digital cameras, but it was way too expensive: the prism is expensive; three imagers instead of one; it was bulky -- the whole design of the camera, basically, was a lens and an imaging system bolted against a notebook computer, because all the computation was done in a notebook computer, and also you had a much better, much larger screen to see the image, to have the controls and so on. So from the ergonomics point of view it was interesting, but it was hardly a portable camera. So it was against the grain, really, and against the direction of the industry. But in the course of that work, the engineering people here at the company learned a big lesson, they learned a lot about how to build cameras, how to do image processing, how to build CMOS imagers, and so that basic expertise was quite valuable. And then around the year 2000, one of the engineers here, Dick Merrill, had the idea of sensing all three colors and each pixel position by burying, in the silicon, photodiodes at different depths, by taking advantage of the fact that silicon is a filter for light. In other words, a photon of high energy, a blue photon, is absorbed near the surface of the silicon; green travels a little deeper, half a micron to a micron; red deeper yet. So if you put "catchers" of photons at different depths in the silicon you can sense all three colors, using totally solid state [means], instead of plastic filters on top of the sensors, achieving a much higher density of information capture than is possible with conventional techniques. That was the winning idea of Foveon and that started a new direction for the company. Obviously, it is not a simple task to put three sensors at different depths in silicon. So, between the basic idea and actually doing it, it took a while.

**Hendrie:** When did you join and sign on as the CEO?

**Faggin**: One and half year ago.

**Hendrie:** So they were still working on this.

**Faggin**: They actually had a product already out and a product that was being designed. The first product saw the light of day in 2002. That was the Sigma camera. Sigma is a Japanese company new to digital photography. So we actually ended up designing the entire firmware and the whole electronics and the sensor. They did the focusing system and the camera and the optics and everything else, and the camera controls. But all the fundamental digital imaging we did for them, with the strength of the work that we already had done in the early camera, but with a major effort. And the camera got great, rave reviews. The camera produced superior pictures, under normal lighting conditions. There was one problem: the low-light sensitivity was much, much worse than CCD because, obviously, CCD has had 25 years of evolution and this technology needs a lot more evolution to be competitive with CCD. So the camera was well received, the technology acclaimed for its novelty and brilliance, but the company was not making money, it was actually loosing a lot of money, because the camera was not a success in the market. They [Sigma] sold a reasonable number but not even close to be able to feed our company. Then the company decided to go after the consumer DSC [digital still camera] with a smaller chip, a 4.5 megapixel chip -- the Sigma had 10 megapixel.

**Hendrie:** Really high end. Professional

**Faggin**: A very large chip. Prosumer.

**Hendrie:** Yeah, prosumer, okay. Pricing of the Sigma camera, roughly?

**Faggin**: About $1300, without the lens, just the body. That's how SLR are priced typically, anyway. So then the company decided to go after the consumer DSC, which was a fundamental error, a complete error because it was doing more of what was wrong. Because, basically, until you solve the sensitivity problem, you cannot go after CCD.

**Hendrie:** Can't compete against it.

**Faggin**: Particularly in a market which is controlled by the Japanese. So, strategically it was a very wrong thing to do. So the company began to be in difficulty and that's when I stepped in. There is not enough time to go through the story, but actually the story may be interesting to tell a few months from now.

**Hendrie:** That's very good, yes. This is an ongoing story and we should do it after maybe… a new chapter of this.

  

**Faggin:** Because this is an ongoing story and hopefully the story will have a happy ending in about six months. So why don't we adjourn six months from now so that I'll give you the happy ending!

**Hendrie:** That sounds wonderful! Thank you very much Federico for taking the time out of your busy schedule to do this oral history for the Computer History Museum. We really appreciated it.

END OF INTERVIEW

  