



Zilog Oral History Panel on the Founding of the Company and the Development of the Z80 Microprocessor

Moderator:
Michael Slater

Panelists:
Federico Faggin
Masatoshi Shima
Ralph Ungermann

Recorded: April 27, 2007
Mountain View, California

CHM Reference number: X4023.2007

© 2007 Computer History Museum

Michael Slater: We have with us Ralph Ungermann, Federico Faggin, and Masatoshi Shima. I think we'll start. We've had in the previous tapes [oral histories of the Intel 4004 and 8080 MPU projects recorded on April 26, 2007 at the museum] some information on the background of Federico and Shima-San, so we'll start, Ralph, with you. Could you give us a brief summary of your background, your education, and what sort of jobs you had before coming to this project?

Ralph Ungermann: Certainly. I grew up in Southern California. I'm a Berkeley graduate in Double E, and a masters degree in Computer Architecture from UC Irvine. I got out of college and spent a little time in the aerospace industry, and then moved towards the semiconductor business. I joined Collins Radio, a tremendously interesting company to work for. That was in the late '60s, [the company was] completely networked around the whole world. Way ahead of its time. I left there and started to work at Western Digital, a [semiconductor] start-up in Southern California. We did communication chips. Custom communication chips. I left there and went to Intel, because I saw the microprocessor as being the next huge wave, and I was there at the Z80 introduction. I left Intel and joined Zilog together with Federico Faggin. And last, having sold that company to Exxon, and I was at the start-up of Ungermann and Bass and then First Virtual. Today I'm very involved in China start-ups.

Slater: Federico, can you tell us about the genesis of Zilog, how it came to be. How the transition went from Intel to Zilog, and what motivated that?

Federico Faggin: Sure. Back in early '74, there was a major reorganization at Intel, and I became department manager under Les Vadasz. Les Vadasz' job increased, taking over other departments that were not under his control earlier. So in that capacity I had a bunch of groups under me in R&D. The largest one was the microprocessor group that was led by Ralph Ungermann, and Shima was working for Ralph Ungermann at that time. Until that time, Shima was working for me. Then I had the memory group for static RAMs, ROMs and EPROMs. I had tiny circuits. I had custom circuits, because Intel was still involved in some custom chips. And also I had all of the layout services for Intel, for the entire department - both bipolar and the dynamic memories, which were not in my group. But that also had [responsibility for] technology development of the group, and they had a bipolar design group. Intel in those days was still involved in some bipolar chips. And of course, I was working for the Vadasz. So at that time, as my job grew, I moved away from strict concentration on the microprocessor only, and the custom chips only, which were my two prior responsibilities earlier. Ralph Ungermann before was in charge of the custom chip development of Intel. In the middle of 1974, I grew a bit restless. There were many changes at Intel. Intel and the whole economy was in recession. There was a layoff at Intel, about 10 percent of the people were laid off. Sometime a year or so before, Andy Grove had instituted the sign-up sheet. . There was a certain degree of resistance, certainly from me and a number of people about this new way of working at Intel. And I was also not happy about the way that [the] microprocessor was still viewed at Intel at that time. Intel was primarily a memory company. Microprocessors were really important for Intel, only to the extent that they would sell more memories and more chips around them. I

did not feel really appreciated and supported enough by Intel in this area. I felt that whatever I wanted to do, I had to almost put up a fight with Vadasz. For example, the 8080, as was talked about yesterday, it took about nine months before I got permission to do. The 4040 also took a certain amount of convincing to do. And now marketing, microprocessor marketing, under the direction of Bill Davidow was beginning to assert its own rights to develop chips. Now that was an absolutely correct. They should've been defining chips from the Day One, but that was a job that I had done from just about from the beginning after the 4004 and the 8008. And so I felt boxed in at Intel. Also I was working very hard and I felt that I could do better if I started my own company.

Slater: Shima-san, can you tell us about your background. I guess you gave most of your background in the previous talk [see note above], so maybe start with how you got to be at Zilog and what your involvement was with the Project.

Masatoshi Shima: Before the development of 8080, in Japan I used NEC mini-computers, which acted quite similar to Motorola's 6800. But it had two sets of general purpose register, including the program counter and the index register. After I developed 8080 and I supported peripheral chips development, but I wanted to develop the high performance, the next generations of an 8-bit microprocessor. Certainly, the main reason to join Zilog.

Slater: And Ralph, how did you come to make that transition from Intel to Zilog?

Ungermann: The transition was a very smooth one, and a very fun time, I must say. The camaraderie and the spirit at Zilog was very good. And so many of the things that Federico talked about Intel were the opposite at Zilog, and the company basically focused on people as well as the products. My role was primarily around the peripheral chips. I helped define them, and hire people to design them and bringing them to market.

Slater: Federico, so how did the concept for the concept for the Z80 itself then come about? You had started Zilog. Did you have a product plan from the start, and how did you come to formulate what turned out to be the Z80.

Faggin: Let me say a few words more about the formation of Zilog. Basically, as I grew restless, I decided that I was going to leave Intel. I asked Ralph to go out for a drink, and I said, "Ralph, I would like to start a company [to build] microprocessors, are you interested?" And he said, "Yeah! Let's do it!" And that was it. I mean, basically, there was not a discussion of what we were going to do, how we were going to do it and so on. It was just an immediate response. And then we strategized on how to get out of Intel. Then Ralph left first, about September, or maybe even August of '74. And I followed in November. In the end of October, it was Halloween, my last day. When I told my people that I was leaving -- I told Shima that I was leaving -- Shima said, "I want to come!" And I said, "Yeah, I like to have

you, but you have to wait until we got the money, and we figure out what to do.” So with that movement to Zilog then, the first step was to figure out what we were going to do. I did not know. We wanted to do a microprocessor company that was very clear. We had not defined any products before leaving Intel. I did not want to be contaminated with that information. And so we never talked about what we were going to do before.

When I joined -- when Zilog was formed -- by the way, Ralph and his wife set up the original office in State Street in downtown Los Altos, so I already had an office and a desk when I arrived at Zilog. Of course, Zilog in those days was not called Zilog yet. And then we had an interview with a reporter from the *Electronic News*, which was a weekly newspaper in the electronics field, was very well-known. And this reporter reported that there was this company that had been formed with Faggin, Ungermann, and they were going to be some kind of microprocessor company. And that caught the eye of Exxon Enterprises. Exxon Enterprises in those days was a subsidiary of Exxon Corporation, the oil company. And they were making venture capital investments in the R&D information technology. So somebody called me up from Exxon and wanted to visit. And I told them that we were not ready yet, that we had not decided what we were going to do, but if they would come around later on, we would gladly see them. So then they called me soon after saying that they were going to be in the area ten days later. That's when I got busy trying to figure out what to do.

And so the first idea that I came up with was the idea of a single chip microcontroller, which I called the 2001. And it would have been the equivalent, more or less in power to the Intel 8048 that came out a couple of years later. And I did a bit of work on that. Enough to have a sense of what it would look like. Ralph and I put a business plan together very quickly, so that when the Exxon guys were come, they would have already a sense of what we were going to do. I think we hit it off right at that point, and that was the beginning of a dialogue with Exxon Corporation, that led eventually to the financing of the company in June of the next year, of '75. I might add that in those days, because of the recession, and because of the bubble of the early '70s in the stock market, there was no venture capital money. In the entire 1975 only \$10 million of venture capital were invested in high tech. And that was an absolutely very low number even for those days. The industry was hundreds of millions of dollars a year, and it was really a disappointment. So we were lucky to actually find Exxon Enterprises interested in us at that time. So then as I was thinking more and more on the 2001, since we did not have a fab, we were a fabless company, we did not have any plan to have a fab soon, because there was no hope for us to raise the amount of money required to have a fab right away from the beginning.

I identified Synertek as a foundry to use for our product. Then as I talked to them and got the idea of the costs, and so early on it was very clear to me that we could not really compete with a company with its own fab selling a single chip microcontroller. That would've been a highly price-sensitive product. So I was trying to find out something else that would be the first product for the company to do. And it wasn't until December that I came up with the idea of the Z80. What in those days I called Super 80. Basically what I wanted to do then, what came up as sort of a gestalt concept to me was the idea of a seamless chipset with a CPU and a number of peripheral components designed around a five-volt bus. In those

days the 8080 had three power supplies, plus 12, plus five, minus-five. The bus structure was a reasonable bus structure, but it was not designed for the entire system with all the bells and whistles that you wish. Particularly the vector interrupt that was possible if you had done a good CPU and peripheral design together.

I defined the basic architecture of the Z80 in the following several months. Basically I wanted to combine the best of the 8080 I wanted it to be machine-core compatible with the 8080. And add a lot of the functions that were on the 6800, which at that time I considered a very, very nice microprocessor. So I wanted to have a couple of index registers, more 16-bit operations, a better interrupt structure. The whole idea of doubling the number of registers. And I could exchange the register with an exchange instruction, the whole register set. That was an idea that I had used already on the Intel 4040. So that one could serve it up very fast if that was a necessity. And on and on. So that architecture was pretty much defined, probably 90 percent of what ended up being the final architecture was defined before Shima came in February of 1975. Now we had a number of problems at that time to get money, even with the interest of Exxon Enterprises, from the very beginning. And we had a hard time to find the acceptance at Exxon Enterprises of the idea of second sourcing. Second sourcing was a necessity for us, because being a start-up company, we could not be able -- we couldn't see how we could get into large accounts where people had to bank on a small company, invest in software development, and then [risk that the] a company disappear. So it was essential for us to have a second source. And it took a while to figure out the strategy how to do that. And that [task] was weaved in, with the design of the Z80 that started in detail when Shima joined the company.

Slater: It's a good segue for Shima to talk about the role that you played, and tell us a little bit about the how the project was broken up, who played what role, and what were the roles that you played, and some of the challenges that you faced.

Shima: Let me talk about the 8-bit microprocessor. The 8080 brought computing power to the system user. Personally the goal for Z80 was to develop the ultimate high performance 8-bit microprocessor. Accordingly {bringing} both the 8080 user and the 6800 user to the Z80 campus and also expanding to the new applications. At that time both 4K bit DRAM and floppy disk [controller chip] came out into the market. Therefore, it was expected that it's now or later, that a disk operating system is coming out. Then we made some interface for DRAM memory. And I became familiar with all of the design activity in the microprocessor development through the project of 4004, 8080, and sometimes it is said that a big success is obtained when developing the challenging new product three times. And Z80 became the third microprocessor for me. So I was able to concentrate on the functional specifications, the hardware architecture design, and all of the design activities. I was able to design whatever I wanted. And personally I wanted to develop the best and the [most] wonderful 8-bit microprocessor in the world. Also sparing enough time for the hardware architecture design.

Slater: Thanks. And Ralph, what was your role on the team.

Ungermann: Well, as I mentioned we wanted to have a complete series of peripheral chips to make it easy to put these into systems. And so I headed up the activity of the peripheral chips. In addition, we built a simulator to check the logic, and also be the basis of the development system that would be required to run. We brought in some software people. I started -- headed up a small group with Charlie Bass. Doug Broyles and I built a simulator of the Z80 at breakneck speed, trying to keep up with Shima-san's design. I remember one event that sticks out very clearly, and that is that we were just ready to the put the mylar into the layout, when we found a bug in the simulator, in the software -- I mean, in the hardware. And we basically went to Shima and said, "We have to stop, there's a bug in the chip." Shima looked at us carefully, and thought for a couple seconds, and said, "No, no problem, we just have to put a contact here." And it was the most amazing thing that he would take this logic error and turn it into a physical fix in just a few seconds. And sure enough when it was fixed, and when the chip came out, it was bug-free.

Slater: So the layout work was all done by Shima-san, and were there other people on the team? How many people were involved in building this chip?

Faggin: Well, I managed the layout of the Z80 basically, Shima and I hired Wong Sha who was a senior layout draftsman. And then we hired a more junior layout draftsman, which I don't remember the name. But they were very slow. I had promised the board that we would have samples of the chip on May 9, I believe. March 9, sorry, of '76. And I was starting to get very worried about the delay, you know, the time it was taking to build the layout. Remember the Z80 had about twice as many transistors as the 8080, and at the rate at which we were going, we would take much longer to layout than even the 8080. The 8080 was in layout for about six months, so it would've been a long, long time. So I got into the act, and I decided to actually do layout myself, manage the two draftsmen, and layout the rest of the chip myself. I was working 80 hours a week for three-and-a-half months. And I drew personally about two-thirds, if not three-quarters of the chip, because I was working much, much longer hours. And I was much faster than the draftsmen, and I was able to do that layout in that short time. Obviously, Shima did all the logic design, the circuit design, that part was totally Shima's doing, but the layout was essentially my doing.

Slater: And was there software work going on at the same time?

Faggin: Yeah, basically the way we divided up was Ralph, with Doug Broyles, Dean Brown and Charlie Bass was worrying about the development system and the software. Dean Brown was the head of software, and Charlie Bass worked for Dean. And myself, and Shima, and two draftsmen were worrying about the Z80 CPU. The entire Zilog in those days had 11 people. And we stayed 11 people until January of '76. The time after we saw the first Z80, and saw that the first Z80 was almost working, and had only a couple of bugs. Then I decided to hire one more person. It was Bernard Peuto, the computer architect to actually begin the process of defining the next generation processor.

Slater: For any of you who want to answer, what do you think were the biggest challenges in creating this design? Shima-san?

Shima: Before that, let me talk about the schedule. It took one year and one month from the kick-off meeting to fully working silicon. The Z80 project started in the middle of February, 1975, and it was completed in March, 1976. The product development took about two months with my memory. And it was completed in April. Next, the hardware architectural design took about two months. I started in the beginning of March, and I completed the hardware architecture in the beginning of May. Then it took about one month to generate a circuit design handbook, and also layout the design method. The detail logical design started in the latter half of May. And it took about six months of designing from logical design to tape out. First the development of the logical design of the register file and address update section was completed on June. The first version of all remaining logic was completed on July 19. But there are lots of iterations. The second version of logical design, except system bus control logic, was completed on August 7. The remaining system bus control logic was completed on September 16. Still there are many, many minor logic mistakes to be fixed. And finally the tape out was completed in November, 1975, where it took about two months for mask making and wafer processing. Next, can I talk about my opinion for product definition?

Slater: Sure.

Shima: I describe my opinion regarding how Z80 product specification was defined. Our critical step at first it was decided that Z80 keep binary code compatibility with 8080, in order to take the market share of the 8080 greatly and quickly. And secondly, two sets of index registers were added in order to go into the market of Motorola's 6800. One index register was not enough. Thirdly, in order to support the high-speed task switching, in the beginning I asked to complete two sets of register files including the program counter. But it was too complicated for customer. Then we gave up on [the idea] of the two sets of general purpose registers. Fourthly, the powerful and flexible system bus was introduced, and it allowed a relatively slow memory. Fifthly, a one-phase clock instead of two-phase clock was chosen for easier system design. Sixthly, in order to bring both 8080 and 6800 surely, the abundant powerful instruction set were added enhancing a critical set, enhancing the 16-bit data transfer, 16-bit calculation, and also adding the bit to data type, and byte to string data type. And reliable single mode for branch and index address single mode for data memory. But new instruction sets were selected very carefully, because the remaining instruction code of 8080 was only twelve. At first six instruction codes were used for the reductive branch, and conditional reductive branch, and branch for loop, in order to get high speed and reduce the program size. Next, the value instruction code was used for newly added instruction set group. Though the instructions are bit instructions, rotating instructions and shift instructions. And two instruction codes were used for the exchange, the set of general purpose register, and exchanging the set of accumulator and flags. Also those support the high-speed task switching. And one instruction code was used for another newly added instruction set. Such as string operations and 16-bit to data transfer, 16-bit addition, and 16-bit substructure, and a newly defined input and output instruction, which provide the register indirect addressing mode, and so on. And lastly, are two instruction code were used

for prefixing instruction. But at that time we didn't say prefix the instructions, but maybe meant they broke down easy to understand the prefix instruction in order to support index addressing mode, and index register operations. I will talk about that one later.

Slater: Thank you. Frederico, can you talk about challenges of getting the chip fabricated and what your strategy was about getting the chips built, building the fab, and getting second sources.

Faggin: Well, let me take a step back first, because we did not cover one important aspect of the Z80, and it was the choice of the process technology to use. The technology that I decided to use was the depletion-mode, five volt N-channel with silicon gate. And that technology was actually one that I developed for the 2102A at Intel back in the middle of '74. It's a technology that allowed a large improvement of speed and power performance for N-channel. In fact, I built with Dick Paschley the 2102A, which is a fast static memory, 1K static memory. And with that technology we got 80 nanosecond access time against the best that could be done was 400 [ns?] with the 2102. So there was a factor of five in improvement in speed-power product for that technology. That technology, by the way, was at around the same time the technology the Mostek had for their static RAMs. And there was another company, Synertek that had that same technology at about the same time. So it was not an invention of mine. It was simply something that was already being used by other companies, but Intel was resisting to use that depletion load in those days. And I recognized that that was the correct technology to do the next generation microprocessor. So before Shima came, I developed the methodology of design of the depletion load, which was different than the methodology of design of high-voltage N channel. And I basically brought up the methodology that I developed originally for the 4004, and it was used pretty much for all of the other subsequent generation microprocessors. I updated for the five volt depletion load device, and that was what Shima got when he joined the company. Then the other thing was to be able to have a relationship with somebody that would build the wafers. So I got into a contract with Synertek to build wafers for us early on for the Z80, and at a price that would allow us to be in business. And by the way, it was interesting because at some point Synertek turned against us. They wanted to -- you know, the president called me up, and when he found that he was building the Z80 for us, after we announced the product, called me up and he said, "Well, if you don't give me the Z80 as a second source, I'm not going to build wafers for you." And I said, "What! I have a contract with you! You are saying that you are a wafer foundry, and now you turn around and you say that you want my product. No way!" Okay, and that forced us also to go to Mostek, because Mostek was the only other company that had the five volt process in those days. But that choice was the correct choice. In fact, that technology in its variations over time, shrinking and so on, was a technology that was used easily for the next five to ten years, before CMOS supplanted that technology.

Slater: You didn't initially intend to build a fab, but you did end up building one?

Faggin: Yeah, it isn't that I didn't intend to build a fab. I didn't believe that we could possibly get the money to build the fab. You know, even if in those days fabs were not very expensive, it would still take

several million dollars for the equipment, whereas now, of course, it's several billion. But in those days with the three-four million of equipment, you can actually build a small fab. So I didn't think we could get the money. But then after the success of the Z80, Ralph and I were able to convince Exxon Enterprises to actually fund a fab. Also it was absolutely essential if we had a second sourcing arrangement with Mostek as we did, that we were independent in manufacturing in order to survive as a company. So it was an imperative to have a fab at that point. Let me say a few more words about the development process, because as you know, from the overview, we basically got \$500,000 the first financing of the company. And I promised to Exxon that we would get on March 9 that we would get the first working CPU, and that I would spend \$400,000. And Shima was able to get the Z80, the first working Z80 on March 9. And that was an incredible feat. And we spent \$400,000. So we did the CPU, the development system, the software and everything else with \$400,000. Which, in this day, would be considered totally impossible. But even for those days was probably one-tenth, or at least between one-fifth and one-tenth of what it would take anybody to do.

Slater: Let's talk now about some of the challenges the team faced in creating this product. Ralph, on the part of it that you were working on, what were the biggest challenges that you faced?

Ungermann: That's a hard question. I don't remember many hard challenges, other than just a lot of tremendous, a lot of hard work. But in the software side we had to do a lot of basic work. We did a real-time operating system, and that is another tremendous project done by very few people very well. And the software was a unique and very valuable part of all of Zilog's sales.

Slater: And Federico, from your perspective, what were the biggest challenges with the Z80?

Faggin: Well, the biggest challenge was to do it in a short time, because that compressed everything. That forced our hand, in a way, because I was afraid at that time that Intel would come out with another -- an improved 8080, and we'd be -- and I didn't know what Intel had planned to do. And we had already lost a fair amount of time, you know, raising money and so on. So that we did not have much time advantage over Intel. In fact, we had time "disadvantage" over Intel. So that was a big challenge for me. The money was a big challenge. You know, we had half-a-million and we had to have this thing working so that we could get some more money so that we could build the company. The other big challenge for me was really working through this issue of getting a second source. Who are we going to partner with? How're we going to make money, to make a business out of it? I was the CEO of the company. Ralph, in those days, was the executive VP, and you know, and I wasn't just worried about the Z80 CPU or the development system, I was worrying about, "Can we build a company here?" So a lot of my challenges had to do also with how we're going to make a business down this road. And so Ralph and I were a little bit alone on that in the sense that I wouldn't necessarily share those concerns with others, but we would have long discussion, the two of us, trying to understand how we could actually make a business out of this. So I would say that those were the challenges the top of my mind for that time.

Slater: And what were some of the solutions to those challenges?

Faggin: Well, for the timing of the product out, that was put three-and-a-half months of my life at 80 hours a week doing the layout. I mean, that was the only way I could solve that challenge. We couldn't find a layout draftsman. They weren't experienced enough or fast enough. And that was the absolute, the limiting time in terms of getting the product out. Because the design, Shima pretty much was doing it all the time it was allotted, but the layout was just not getting done fast enough. And so that's why I had to do it. You know, a CEO doing layout draftsman job was not something that would be normal, but that's what I had to do, and I did it. So that's how I solved that problem. And that certainly saved many, many, many months for the project. As I said, I did two-thirds to three-quarters of the chip drawing it myself. And business challenges were really solved only when Exxon agreed to let us build a fab. So the fab was the key to giving us independence, so that we would not be in a position of underdog to our second source. So having done all the work, and then the second source basically being able to take the business away from us. So it was important also in the business deal that I negotiated with Mostek to negotiate in a way that it would not give them so much advantage that they could take advantage of us, because they were a bigger company, with a lot more feet on the street, and we were just a start-up. I thought they believed that they could get us in the marketplace. And but because of the good work that we did as a team, and also the fact that we did not give them, for example, the development systems and software, and those tools that were essential for the designing. And the fact that we built a fab very fast, we were able to actually not lose much of the business. We probably ended up with a 60/40, 70/30 market share in Z80 at the end of the day, you know, many years later.

Slater: Shima-san, from the engineering work that you were doing, what were the big challenges that you faced, and the solutions that you came up with?

Shima: Yeah, the functional specification became excellent. Therefore, it became necessary to spend lots of time on the hardware architecture design at the time of product definition. It was quite interesting the development job, but it was very severe work. There were seven challenges of designing Z80. How to get higher performance than both Intel's 8080, and also Motorola's 6800. How to expand the number of instruction set. The next one is a most important job. How to differentiate Z80 logic from 8080 logic. Next one is usage of one-phase clock. Then chip size to be small as much as possible, and no trouble in the layout; I had a lots of trouble in layout of the 8080. The last one was how to stop the copy of layout. That is my seven challenges.

Slater: You want to tell us about your solutions to those challenges?

Shima: Yes. In order to develop the 8-bit microprocessor, as I said, I spent about two months for hardware architecture design, and planning the chip plan in advance. I didn't want to repeat same mistake in the layout of 8080. So I introduced a unique design technology as follows: At first in order to expand the number of instruction set with small number of transistors I introduced two prefix in the

instruction. One for index register X, one for the index register Y. If the prefix instruction was added to certain instruction, including the old 8080 instruction, and the new instruction set, the prefix instruction is able to redirect the selection of each register to index register X or Y. And this was our pride for not only 16-bit index register, but also the lower or higher of index register. But we didn't publish that one. And also I had to use this prefix method before I developed the desktop calculator in 1968. And secondarily, in order to differentiate Z80 logic from 8080 logic, I introduced many unique design technologies as follows. One-phase clock instead of two-phase clock, and 4-bit arithmetic instead of 8-bit arithmetic unit. The storage of both accumulator and flags into the register file. Those were more than enough to differentiate Z80 logic from 8080 logic. At first I introduced the pipeline 4-bit ALU. Next, I introduced the pipeline among the instruction set and the instruction executions, and minimum number of the clock for instruction was four clocks. And I use a four clock for instruction execution. But still there was some overlap between execution and instruction fetching. Also that was quite different logic from 8080's. Then I introduced a very powerful, flexible unique 8-bit internal data bus. The considerably large size transistors were placed on the internal data bus like the gate of a canal. Please look at this internal system bus. [shows diagram] And here are these big gates, here. And it was divided into the four sections. Let me go back to here. Usually it works as one set of internal data bus. However, if required its internal data bus is divided into four kinds of internal bus in order to perform four different types of operation concurrently. The first internal bus is used for the instruction fetch. And second internal data bus is used for the accumulator-related operations. The next set internal data bus is used for the flags related operations. Fourth internal data bus is placed inside of ALU to be used for not only the 8bit operation, but also digit operation, bit operation, rotate operation, and the shift operation with very small number of transistor. With this internal shift and data bus architecture, both accumulator and the flag are placed into the register file. In addition to it, the considerably large size of transistor was placed inside of register file. Its gate was off. The general purpose register including accumulator and the flag are separated from the group of program counter and stack pointer, thus differing their operation related with general purpose register, accumulator and the flag. The other register is able to be updated concurrently. Those design technology brought so many advantage as follows: the significant difference from 8080 logic, and significant chip size reductions, and no performance problem in ALU. And also lots of layout space for flag logic. That was a big trouble in 8080. Still performance of the Z80 is almost same as 8080 performance, and is the same clock frequency. There was not any trouble to increase the clock frequency by 25 percent, because the different type of transistor that was used. And thirdly external system bus was improved. I designed the logic of 8080 by using simple logic, because there was a lack of manpower and a very tight schedule. As a result in 8080, the time of half-clock was usually wasted at the time of address output and also instruction-fetch time. In Z80, address information always was read out from the register file, putting up to output of address by half-clock up. Also instruction was fetched to the input buffer. Then it was stored to the instruction register at next state of time. And that means even if Z80 used two-clock instruction fetch cycle, instead of three-clock in 8080, Z80 was able to provide same memory access time as 8080.

Slater: Thank you.

Faggin: This is like doing algebra in public.

Slater: Yeah.

Shima: This is why we made a good Z80. Without this one nobody can make a Z80.

Slater: Do you recall, any of you, the key statistics in terms of how many transistors and what the process technology was.

Faggin: Well, yeah, the process technology I mentioned earlier was 5-volt depletion mode.

Slater: What was the dimension?

Faggin: The dimension, roughly 200 x 200 mils. Do you remember, Shima, exactly?

Shima: I don't know.

Faggin: About 200 x 200.

Slater: The process size?

Faggin: So it would be 5 mils -- the process size was about -- the process size actually is an important point, because I spent a lot of time defining common design rules, because we did not have a fab. Any fab has its own process with its own specific design rules. Because we didn't have a fab, I had to generate a composite of design rules that would allow us to use the same tooling with different fabs. Actually I spent a fair amount of time to devise that. And we lost some chip size because we couldn't optimize like a specific fab can. But that gave us a much more robust design at the end.

Shima: But also the Z80 is much more compact than the 8080.

Faggin: I think that the minimum line was six microns still. So we -- by the way, in those days, the whole idea of scaling was just beginning to emerge. In fact, Intel was kind of late in adopting that technique. I found out that Mostek was much more advanced than Intel in applying scaling and shrinking, and that kind of idea into their products. And so what happened, I used the design rules to include those ideas, so that when we adopted Mostek we were able to shrink the Z80 by at least 20 percent and making it more cost effective.

Slater: Where did you start out in clock speed? And what did you ultimately achieve?

Faggin: Two megahertz. And we had a pretty good selection of four megahertz. We had already four megahertz from original mask set, we probably had 20 percent distribution at four.

Slater: And this was when the 8080 was at what?

Faggin: The 8080 was at 1.5 and two I believe.

Slater: So if we can shift gears a little bit to talk about some of the business and marketing side. What was the competition like? Who were the big competitors you were facing, and how were the customers choosing amongst their alternatives?

Faggin: Well, maybe you want to say something?

Ungermann: Yes, the 6800, obviously, was a good competitor. And the Intel 8080 was really taking off. And so we were able to go poach their design wins, because of the software compatibility, which was the key issue here in terms of being able to go after that market. People liked the Z80. They liked it a lot, and it was a very easy sell, because it was simple.

Slater: And then the 8085 came out a little after the Z80.

Ungermann: 85. It came out after I left.

Faggin: The 8085 came out about a year after. That was Intel's Super 80.

Ungermann: Oh yeah.

Faggin: That was Intel's version. That was basically a five volt version of the 8080 with a couple of instructions added.

Slater: And did that provide much competition for you?

Faggin: Not that I remember, no. But Intel was still formidable because of their presence in the marketplace, you know, the development systems created a stickiness for Intel. And all the peripheral components that Intel had at that point or was beginning to have were started under our tenure. And so we ended up poisoning a little bit of our [own] well.

Slater: You had some concerns about Japanese companies, in particular copying your design. How did you address that concern?

Faggin: Yes, we were concerned about others copying the Z80. So I was trying to figure what we could do that that would be effective, and that's when I came across an idea that if we use the depletion load the mask that doesn't leave any trace, then I could create depletion load devices that look like enhancement mode devices. And by doing that we could trick the customer into believing that a certain logic was implemented, when it was not. Then I told Shima, "Shima, this is the idea how to implement traps. Put traps, you know, figure out how to do the worst possible traps that you can imagine," and then Shima with his mind, that was steel mind, was able to actually figure out a bunch of traps that he could talk about.

Slater: You want to tell us a little about that Shima?

Shima: I didn't count [on] talking about that mostly. I placed six traps for stopping the copy of the layout by the copy maker. And one transistor was added to existing enhancement transistors. And I added a transistor looks like an enhancement transistor. But if transistors are set to be always on state by the ion implantations, it has a drastic effect on very much. I heard from NEC later the copy maker delayed the announcement of Z80 compatible product for about six months. That is what I got from NEC. And finally a total transistor of Z80 became 8,200 while a total of transistor of 8080 was 4,800.

Slater: Who were the primary customers early on, and what was the pricing strategy? You want to talk about those aspects, either of you?

Ungermann: We went after the PC business as we know it today. And we brought our prices down to allow us to do that very quickly. When we came out with the product, we priced the Z80 at \$200 a chip, and it took us just a few days to understand from our customers that they needed a lower price to be able to go after the market. However, we also had a number of system companies that were thought that they could get the IBM software to run on the Z80. They could crack a lot of IBM markets, but that of course, didn't ever happen, because the Z80 was not designed to compete with a mainframe.

Slater: Were there important embedded applications also?

Ungermann: Yes, there were many embedded applications.

Faggin: Yeah, one of the first customers that we had was Cromemco. In fact, Roger Mellon came to the office and I remember personally selling him one Z80 for \$200, and Ralph didn't say it, but NEC was actually our first customer, and I think you sold the Z80 to the competition. <laughs>

Ungermann: When we were packaging the product up for shipment, somebody knocked on the door of the office and it was an NEC guy with four \$100 bills for two chips. And then right after that a very successful Japanese computer company came in with four \$200 bills. That was just the start of the copying.

Faggin: We also sold quite a few Z80s to Tandy, Radio Shack, for the TRS-80.

Ungermann: Trash 80, as it was called.

Faggin: Well, Trash 80 later on. So the TRS-80 was very successful, what today would be called a personal computer, home computer -- hobby computer in those days. And we also were adopted by Videc for their word processor. Videc, was an Exxon Enterprises company; by Quiz - for the edit-drawing typewriter; by Wang for their older word processor line. The Z80, because of the string instructions, was very popular with word processing. Because in one instruction you could move a block of data. And we also had a large volume customer in the gaming area, Fidelity Electronics, or something like that. They built a chess game, computer chess game. And I remember that I was given one that had voice synthesis, and you know, when you turn it on it would say, "I am Fidelity, your computer opponent. Select your move." And then you proceed to play with the computer. We were selling between 100 and 200,000 units a month for that.

Slater: Just to that one customer.

Faggin: Just to that customer, yeah.

Slater: What did the volume peak at for the total product?

Faggin: Well, the volume peaked much after my time. I figure, because now I'm chairman of Zilog, and I went through a little bit of data, I figured there were over two billion Z80s sold by the various sources. But many Z80s are actually used in embedded application, so we don't even know how many of those there are. For example, NEC is still using Z80s in many of their embedded ASICs.

Slater: So it's turned out to be an extremely long-lived product.

Faggin: It's still in high production today.

Slater: What was your expectation about that? Were you thinking these were short life-cycle products? Did you have any sense that this was something that people were gonna still be using in 20 years, 30 years?

Faggin: No. It far surpassed my expectations. I thought in my -- you know, when I was putting the business plan together for Zilog in those days, and in my calculations that I was making predictions of what might happen and so on, my calculation was that it would peak at around '81-'82, and then it would sort of peter out as the new generation, the 16-bit generation was going to take over. In reality, it kept on growing for at least until the mid-'80s, if not longer. And then it came down very slowly. And in fact, many new applications that are embedded appear, and those we can't keep track of.

Slater: And the relationship with Exxon -- how did that evolve? Was that an asset? There was obviously an asset to you financially. Was it an asset in the marketplace?

Faggin: It was an asset in the beginning, because we got money when otherwise we would not have been able to start a company under the financing venture capital climate that existed in 1975, and the window of opportunity was that there was no other time that a microprocessor company could really start - a general purpose microprocessor company, like the way I saw it. So it was essential to have Exxon as an investor, given their willingness to invest in us. But then later on it became a liability, because they became a competitor in a sense. Or viewed as a competitor in the marketplace by many of our customers. Exxon had an intention, which we were not privy to, they did not tell us that in the early days. They had a master plan to create a major information technology company that was going to compete with IBM. And in fact, they started doing advertisement, presenting their company in advertisement in the '78-'79 timeframe. Basically presented themselves as a challenger to IBM. They had a company called Videc that was making word processors, Quiz that was making electronic typewriters, Quip that was making fax machines. There was Zilog that was making, of course, chips. And they had another 20 companies in the wings that many of them they had started themselves in all areas, from printers like Qume, in the area of displays, and so on and so forth. So they really presented us with a big problem. That's one of the reasons why we were not designed-in as a company in the IBM PC. Only one of the reasons, but certainly an important one.

Ungermann: The early design-ins it was a very, very important thing. We could say, "We've got the money of Exxon, and the brains of Intel." And it was very easy then to get a customer to shift from an 8080 to a Z80.

Faggin: Plus we had a second source, which is very credible, Mostek being a credible second source, where Intel did not have a credible 8080 second source. And so the combination of all three factors gave us a very rapid momentum in the marketplace.

Ungermann: So the big question here is, "Why not Zilog instead of Intel?" What happened?

Slater: And what do you think the answer is?

Ungermann: Well, I think that we had a winning strategy, and we were winning in the marketplace, but I think we tried to do too many things, other things, and do memories, and we expanded our horizon beyond our capacity.

Slater: So it wasn't focused enough.

Ungermann: Yeah, that's my view.

Slater: Federico?

Faggin: Definitely that is one reason for that. I also feel that we were inexperienced. Both Ralph and I had never run a company, and we had really negative guidance from the board. The board of Zilog was really three excellent people, but they were the boss of the boss. You know, the boss, and then his report, and then his report. So anytime there was a decision to be made, the lower guy on the totem pole would look up, the other guy would look up at the next, and Ben Sykes who was the senior guy would say, "Yes," or "No." And that was the decision process at the board. So we did not have a diverse enough board that would allow us to bring other perspectives. Also the perspective of Exxon was basically later acquiring us, and making us part of their empire. So they had no interest in their view to really see to it that we would really be sort of out there in the minds of people as a business. They were interested to see us to gain market share, but you had to do more than that. You had to be in the minds and hearts of people in the business community, and we were not. We couldn't do that. Also Ralph and I had a falling-out that did not help. We got into some ego battles, I guess.

Ungermann: Set off by Exxon.

Faggin: Yeah, they were set off by Exxon, but we were not mature enough to handle that in a mature way. So that certainly didn't help. I can say that for myself. And basically toward the end of my period at Zilog as CEO, I was spending more time in New York keeping Exxon out of the way, than I was spending time with customers. You know, having customers design with our product. So, that's how bad it got. So what was truly an asset at the beginning became a liability at the end.

Ungermann: I agree with that completely.

Slater: Well, I think our final question is what unique memories each of you may have of the project, some story you'd like to tell. Something you think was an interesting event. Shima-san, do you want to start with that?

Shima: Yes, I enjoyed Z80 project very much, because I got all of the experience of microprocessor before Z80 development. And Z80 was transported to the production in March, 1976, where I left Z80 project. And then Zilog changed the second source to Mostek. Then the wafer came out in May 1976 it didn't work completely. At that time I took vacations, and I went to the company in the afternoon. There was nothing wrong in both our tester and the test vectors. Then I started to debug its wafer, one step-by-step of testing its wafer with many combination of instruction. And finally, I was able to locate the bug. Next I checked its location carefully with microscope. Then I found out the thin shadow under the metal line. It looks like the poly silicon under the wide metal line of power supply had been cut, but nobody believed what I said. And later, I was told that there was one defect on poly silicon mask. And all of debug were completed at 4:00 a.m. And it took more than half a day to reach to the conclusion. And then I'm lucky. I got the heavy cold, and lost one week of vacation. That is my unique memory. Quite unique. Same as 8080.

Faggin: Well, I have many, many unique memories, frankly. I have the starting with the memory of how we came up with the name of the company. It was Ralph and I at the end of the day would sit around and try to figure out names. And we had a list of names like Electronic Semiconductor, Integrated, etcetera, and etcetera. And tried combinations, and we would come up with a name, and after many glasses of wine, then go home, and then in the morning, I would ask Ralph, "Ralph, do you remember the name that we liked yesterday?" He said, "No, I don't." And I didn't remember it either. So it was very clear that was not a good name. Eventually, I was playing with the word I-Log, I-Log for Integrated Logic. And Ralph says, "Zilog!" And I said, "Wow! That's good!" You know, I said, "Z for the last word of Integrated Logic." And then we went home after another couple glasses of wine. And in the morning, we both remembered the name! And so that's how Zilog stuck. And Zilog became the name of the company. Then the other memory that I have is the fact that we were all going at lunch to the local Safeway store that was not far from where we were. Remember, we're in State Street in downtown Los Altos. So we would go there and buy food. I don't know if it was you, or somebody else, would always buy mackerel. You know, it was smelly as hell! And I would buy acma in camembert, and that's how I got my cholesterol level much higher, and it hasn't come down from that time. And we had a fun time eating together, you know, in a collegial way, you know, in the company. <laughs>

Ungermann: You hit some big ones. Good memories. Many, many good memories. We worked together around the clock, and as I said before even, around one or two days of very serious work, but we always took time out to have fun.

Faggin: Yeah, I had a still collection of corks of a number of champagne bottles that we opened up, you know, when we got the financing, when we got the first chip out, the first Z80 out, when the first Z80 really

worked, and in fact, on the first Z80 that really worked it was in early March, and this really is indelibly in my mind is that Shima has this, you know, micro-manipulator with the wafer of the Z80, and so we say let's try with the development system. So we had actually prepared a socket that would go from the micro-manipulator for the wafer tester, directly into the socket that would plug into the board of the development system that these guys had done. And so we plug the socket, type control/C on the teletype, and we got "ch-ch-ch-ch-chump" in the prompt bat. And that meant that pretty much everything was working! Because just from that -- just to do that you would have to exercise a good 80 percent of the chip.

Ungermann: First wafer, first chip, we got this message back, and the first thing we did was run down to the grocery store and buy champagne.

Faggin: That's right! And we celebrated!

Ungermann: And we got drunk and didn't finish the testing for several days.

Faggin: Yeah! <laughs> No, Shima was there next day.

Ungermann: <laughs> Yeah, that's right.

Faggin: But we were drunk for two days. <laughs>

Faggin: But you may want to -- there are so many -- I want to say a few more words, because those early days of Zilog were so magic that we really need to capture them here. And I really think that because there isn't any other opportunity to do that. So maybe we need probably another five, ten minutes that to-- you may want to think of some other, you know, important...

Ungermann: Just getting into it now.

Faggin: What?

Ungermann: I'm getting into it now.

Faggin: Yeah, good!

Slater: Okay, well, let's hear some of those. Federico, you have some more memories of the early days of Zilog you'd like to share?

Faggin: Well, one that is particularly poignant is that when the decision was made to have our own fab, that at that point I hire Len Perham from AMD to create the fab and run the fab. And so he came in around May of 1976. So we decided that we're going to build a new building, build a fab. And of course develop the process in that fab for five-volt depletion load. And Len Perham was really pretty much in charge of that. And he was able to starting in June, we had selected a builder, which was Carl Berg, and starting from the selection of the builder in an empty lot, in June of '76, by January '77 we had completely working Z80s in the fab that was in a new building. And that was an incredible feat that goes to the credit of Len Perham. I remember the first day that I got into the building we actually started -- we got into the building in October, so the building started in June, and we moved in October of '76. Up until that point we were mostly in State Street in a number of offices at that point. The morning that I moved in, I was so moved that I actually called up my father on the phone and say, "Dad, I have my own building now. My company has a building." And I was so proud of the accomplishment that we all had made, and of this situation.

Slater: Anything come back to you, Ralph, you want to share?

Ungermann: I'm going to go back over a couple of things that we've mentioned briefly. But I'll never forget Shima-san's ability to take a logic equation and convert it to physical layout of a mask. And I remember many meetings when we were working on the instruction set. And Federico would say, "It really needs this one," and Shima would say, "No! It adds two more microns." He would completely analyze the instruction set in real time, and veto or approve the instructions based on his ability to perceive a different world than I work in.

Slater: So you'd be talking about instruction sets and he's visualizing in his head the logic that he needs to implement it.

Shima: Well, I have a super-computer in my brain.

Ungermann: Yeah, and as Federico mentioned it's not often the CEO would go off and lay out a chip, but he did. And I remember the morning when he was halfway through that layout, and he decided to tear it up and start over, and then ...

Faggin: Yeah, you remind me of that. Basically, the worst nightmare of when you do layout is at the end you cannot fit everything on this chip size that you have staked for yourself. And toward the end of it, as I was turning the last corner of the chip and finishing, it was very clear that there were too many transistors, and I couldn't fit them. I had to really compact it more. So it was a very sinking feeling. I

pondered for about half-an-hour, and it was obvious that I couldn't fit it! And so I took an eraser, and I erased three weeks of my work, or two weeks of my work. And I start all over again, and did it, and at the end, I think I had no more than a couple mils left of space at the end. I had forgotten about that, but that was an incredible experience.

Slater: So was this towards the ends of the days of hand layout?

Faggin: Well, not really. No, even the Z8000 was hand-laid out. The computer-aided design didn't start to be used in earnest until the early '80s. But in the early days there was magic in Zilog. I mean, we were working all together well. There was an esprit de corps that was incredible. We could almost do nothing wrong. In fact, it was even too easy in a way. That's why we undertook to do too many things, like Ralph said earlier. But the enthusiasm and the energy that was in the company were incredible! I mean, I don't have any other experience in my life like those early two years of the company. It was magical. It really was magical! And we had tremendous number of good people. In fact, many of the people that we hired ended up becoming CEOs of companies in the Valley. I mean, there are more companies that were spin-off of Zilog that you can imagine, you know. So we had a lot of good people that learned good and bad lessons that were important for their life in our company.

Slater: Well, anything further? I guess to close, Shima-san asked me to read a statement that he wrote to ensure it's as clear as it can be for anybody listening. So I will read this now. These are Shima's words. "Development engineers involved in creative development projects can be likened to explorers who go into unmapped territory without a compass. Having hope for success, and fear of failure. Also, an engineer must be armed with the firm belief that his mission is nothing but development, and must be determined to go his own way, never following another's tracks. Microprocessors evolved from 4-bit to 64-bit microprocessors by strong demands from applications, including computer technology, such as pipeline, super pipeline, super scaler, VLIW, cache memory, and virtual memory system. In the 20th Century microprocessors were used for increasing power of intelligence. On the other hand in the 21st Century, microprocessors will evolve into tools to bring forth wisdom for all mankind."

Shima: Thank you.

END OF INTERVIEW