



Oral History Panel on the Development and Promotion of the Intel 8008 Microprocessor

Participants:
Federico Faggin
Hal Feeney
Ed Gelbach
Ted Hoff
Stan Mazor
Hank Smith

Moderated by:
Dave House

Recorded: September 21, 2006
Mountain View, California

CHM Reference number: X3706.2007

© 2006 Computer History Museum

Dave House: Stan, tell us how the 8008 began.

Stan Mazor: Well, in today's personal computers, we have RAM and we have flat panel displays, but in the middle of the '60s, we had core memories on our minicomputers and we typically used either a teletype or a CRT terminal (cathode ray tube). We also nicknamed these "glass teletypes" because they played the role of the teletype, but they were much faster. The key ingredient to the CRT terminal in those days was the shift register; there was no RAM memory in those days. Shift registers were being provided by General Microelectronics, (GME), General Instrument and some other companies. When Intel got started in 1969 in the memory business, we were designing RAM components-- and we really had no customers for them. So one of the things that we did was---to do a custom order for Datapoint (formerly Computer Terminals Corporation) who was one of the makers of a cathode ray terminal, and we made for them a 512-bit re-circulating shift register. So this gave us a ready customer and it was a type of memory product which was in use at that time.

Well, Datapoint, having been in the display [terminal] business, had an idea of going into the business terminal [computer] business; in addition and they had a clever idea, which was to use the shift registers that they were already using inside the display terminal as the main memory of their computer. Now, rotating memories had been used in earlier computers, usually magnetic rotating memories, and so they had come up with an idea of using the shift register that we were making for them inside a business computer [as the main memory]. But inside that computer they needed some registers and they needed something called a pushdown stack, and we had also made at Intel a bipolar memory, 64 bits that was specified by Honeywell. Honeywell said, "We'll buy these from anybody who can make them," and almost everyone in the industry was trying to make this bipolar 64-bit memory.

Well, Victor Poor from Computer Terminals looked at that chip and said, "Well, I could implement the registers of my 8-bit computer and that and I could also implement a pushdown stack if you guys could add a counter to it." So he met with Andy Grove in December of '69 to check on the status of his shift register order, and in that meeting he mentioned to Grove, "Can you make a stack for me?" And Andy Grove really didn't know what a stack was and as an applications engineer, I was called in to be a consultant on that. So I met with Victor Poor around Christmas time in December '69 and asked him what he needed the stack for and he said he was building an 8-bit computer. I asked him a little bit about it, and I had the advantage that I'd been working with Ted Hoff on the 4004 since October (for several months), so the idea of a central processor on a chip was something that we were working on [already]. And I asked him about his computer and how smart it was, and what it had to do. I wrote out in front of him three proposals. One was an 8-bit register set with a stack, and another was a register stack with an arithmetic unit. And then I went on to the third proposal and said, "Well, it's possible we could do the entire 8-bit CPU on one chip." Well, he certainly wasn't about to believe that, but he was interested. I said, "But we need to know more about your computer before we could proceed with that." So in the January/February timeframe, he sent us a programming manual-- sort of an assembler language manual disclosing the instruction set of his computer.

Ted and I had a look at that and one of the interesting things is, [that] it had about the same amount of register memory as we had in the 4004, an 8-deep stack which is about the same amount of memory that we had in the 4004. At first blush, you might think, "Oh, an 8-bit computer's gonna require twice as much

logic as a 4-bit computer and we'll never be able to do that," but actually when you look at it, the instruction decoding logic doesn't need to know how big the data word is. So our preliminary estimate was that if we simplified his instruction set a tiny little bit that we could probably build a [8-bit]CPU on a chip to his specification, and we made a proposal in around the March [1970] timeframe to do that. They came back with the salesman, and then as best as I recollect, the salesman and the buyer had an agreement that we would deliver a hundred thousand chips at \$30 apiece and those would be the CPU chips. Of course, they were accustomed to doing high-volume memory chip deals. Probably computers were being sold from Digital Equipment Company five thousand or ten thousand a year, so a hundred thousand was a pretty ambitious quantity. So we embarked on a custom chip for Datapoint following the 4004 principles in general and that was how we got started in the 8008--- which we designated internally as the 1201.

House: So what happened next?

Mazor: Well, we had to have a design engineer assigned to the project and Hal Feeney came on board and was my officemate for a while. My job was liaison trying to explain to him what we were doing and what the chip had to do. And he was a chip designer; (I was not) so I was there kind of as a helper and translator and we began to work on the general layout. And I think that probably Hal can tell the story best from here on.

House: Ed, did you know that this was happening?

Ed Gelbach: Yeah, I knew it from a different aspect though. I at that point was, I think, the Vice-President of sales for TI [Texas Instruments]. TI at that point was also involved. And I think he had a brother working there, Gus, was the brother's name, and I remember going in and negotiating it.

House: Vic Poor had a brother?

Gelbach: [Yes, He] had a brother.

House: Okay.

Gelbach: And somehow I got involved in negotiating [what] I think the [was] the same product. And I was a little like Andy, I didn't know what it was, but thirty dollars times a hundred thousand sounded like a pretty good product. And, you know, we immediately said, or TI immediately said they could build it and in fact did build it. I don't remember how it turned out, but I do remember they delivered working parts. So I was aware of that when I left there.

House: And you left TI when?

Gelbach: 1971, mid '71.

Mazor: I have a little bit about the part that we knew about and that is-- that our specification was written for Datapoint; we didn't ask them to sign a nondisclosure agreement because it was our customer. It's very likely that our datasheet ended up at Texas Instrument. We believe their [TI's] original proposal was to do it in several chips, three chips in fact; in a way Intel "set down the gauntlet" that it could be done on a single chip. They [TI] were in somewhat disbelief, but they sort of picked up the gauntlet and decided to run with it. So I think TI did build a chip to our specification. It turns out my specification had a "defect" in it and it turns out they built that "defect" right into the chip; and so the chip could never have actually worked. And if you think about it, with all the effort that TI put into that chip in announcing it ahead of Intel, if it had worked, they probably would have made a product with it. To my knowledge, they never did sell it.

Gelbach: I don't think they did either.

Federico Faggin: Well, Vic Poor told me that in fact they tried it and it never worked.

House: So then Hal came on the scene. I understand you were hired to develop this chip. Tell us your story, Hal.

Hal Feeney: I got involved in what was [called] the 1201 at that time. It was Intel's first 8-bit microprocessor and with the 1201 [internal] name. I joined Intel on March 9th of 1970. I remember the date for a variety of reasons, but on March 9th, 1970, I also have the handwritten spec that Stan put together that took the instruction set that Datapoint had defined, translated that into a concept that we could use for developing the microprocessor. So from that concept spec, and from the concept timing diagrams that were there, we looked at the pins that would be needed to communicate with the outside world and went to defining a functional block diagram for the chip itself. We also went into defining what the logic would be to accomplish and set up and decode all of the instructions that were part of the instruction set. So as we went through that process, it took about two months for us to develop the initial preliminary specification and develop the block diagram. And going back to some comments that were made about TI, it was at the same time that TI, based on other knowledge that we've had, TI was designing a 3-chip set and were designing a serial processor. Intel's was a single chip and was a parallel processor. And to that point, CTC, Computer Terminal Corporation at the time before their name was changed to Datapoint, was working with both companies and sharing information back and forth so that they could get the best product that they possibly could. And as it turned out, the Intel product was moving along up through the middle of approximately 1970 where we were finishing all of the functional aspects of the specification, getting the specification approved both internally and then sharing it with CTC. And that carried on up through about July of 1970, and at that point in time, there were some problems. I wasn't really associated with it from the sales side, but there were some problems with whether the customer really was going to follow on and use the chip or was not going to use the chip. As a result, there was about a six-month hiatus. A little bit of work was done on the chip specification, but about a six-month hiatus in terms of the design itself while we went and did some other work at Intel. I also put in some of my time working with Federico on the 4004 and the 4000 series products.

Mazor: Hal, if I could interject, just a couple of things on the [8008] specification. Ted [also] worked on the specification; we worked on it together. And just a couple of interesting things about that specification. The Datapoint instruction set did have a branch-on-bit instruction which we decided was not needed, so we removed it. Basically we were building a computer for our customer so we didn't challenge their instruction set, or test it out in any way. But Ted realized that they were sorely missing an "increment" and "decrement" instruction-- so that was added based on Ted's idea. So that was one of the enhancements [we made]. Another minor point--[is] that the jump address in the 8008 [Jump instruction] was 16-bits, but because we were so short of pins, we only used 14-bits of address--- and we thought that sixteen thousand bytes of [main] memory was a "huge" amount of memory. Who would ever need more? So although the instruction had room for it, we only used 14 of the 16 [address] bits. And lastly, the original design for Datapoint... what they wanted was a [bit] serial machine. And if you think about a serial machine, you have to process all the addresses and data one-bit at a time, and the rational way to do that is: low-bit to high-bit because that's the way that carry would propagate. So it means that [in] the jump instruction itself, the way the 14-bit address would be put in a serial machine is bit-backwards, as you look at it, because that's the way you'd want to process it. Well, we were gonna built a byte-parallel machine, not bit-serial and our compromise (in the spirit of the customer and just for him), we put the bytes in backwards. We put the low- byte [first] and then the high-byte. This has since been dubbed "Little Endian" format and it's sort of contrary to what you'd think would be natural. Well, we did it for Datapoint. As you'll see, they never did use the [8008] chip and so it was in some sense "a mistake", but that [Little Endian format] has lived on to the 8080 and 8086 and [is] one of the marks of this family.

Feeney: The design hiatus that we had gave us an opportunity to go from that July timeframe to the following January and accomplish several things. One, the original design was intended to go into a 16-pin dual inline package or 16-pin DIP. The reason for that was because Intel was in high volume production with memories at that time and production management wanted us to stay with the most popular package that was being used at the time. Good fortune had it that over the next few months, the 1103 was released into the marketplace. 1103 was a 1,000-bit dynamic random access memory. The 1103 used an 18-pin package.

Faggin: It [the 8008] originally used a 16-pin package.

Feeney: Right. It originally used a 16-pin package and it became very clear and very urgent that another pin was needed. As a result, [when] the 18-pin package became quote, a standard package, within Intel that gave us an opportunity to take a look at potential limitations with the 16-pin design that was being done for the 1201 and with that, the issues that Stan brought up in terms of [the need for] a few new instructions and a few more pins. Two more pins gave us the ability to get more information to the outside world. Most important, the restart instruction was designed just to start the machine but it wasn't originally designed to push all of the information in the machine onto a stack. The additional pins gave us the opportunity of doing that. So by the time the January timeframe rolled around, we had the additional pins, we had additional instructions and enhancements, and some additional timing things that came in from inputs from both Stan and Ted in the December/January timeframe.

Ted Hoff: I might point out that one of the reasons for changing the restart instruction was that Stan and I had figured out it might be nice to add interrupt to the processor. And our first thought would be to make

use of a call to some routine that would be loaded instead of whatever was coming from memory when you serviced an interrupt. And then we looked at the amount of logic that it would take to implement that outside the processor and it looked ferocious. But we had this restart instruction which was one byte instead of three and we figured it would be a lot easier if that were a call. And it turned out the way the processor was, it looked like it didn't cost us anything in terms of utility if we made restart a call. And so that became one of the reasons for the change.

Feeney: Right. And it became a natural fit into the instruction set and very, very true, no overhead cost at all because we were able to deal with it in exactly the same manner that we were dealing with other elements in the processor.

House: So you basically faked an interrupt by forcing in a jump instruction.

Hoff: Exactly.

Feeney: Yeah, but it was a one-byte...

Hoff: It's got to save the contents.

Feeney: Yeah, save the content rather than just moving. But it was a one-byte instruction. And with three bits then directing the user or directing the program to one of eight different locations in memory where a legitimate call would then be inserted into the system. But the one thing that's very important with the 1201 (the 8008) in contrast to the 4000 series, [is that] the 4000 series had its own special memories, special ROM, special RAM, and special serial memory device. The 8008 was planned in such a way that it could be used with any standard type of memory because as Stan pointed out, we had this exorbitant amount of memory, 16K, available that could be any mix of standard ROM for program, standard RAM for storage and that could be partitioned in any way between all RAM or a mix of RAM and a mix of ROM. And using the standard products that Intel had or perceived at the time, and then as it turned out, (and we'll be talking about this I think a little bit later,) when the EPROM became available, that was really the item that provided the greatest jumpstart for the microprocessor.

But back to the design side of things, we then went with our design starting - or restarting in the January timeframe of 1971 and carrying it through to having our first devices in the late fall of 1971. So with that, we implemented a relatively small chip and if I recall, it was something like a hundred and twenty-four by a hundred and seventy-one mils on a side. It had to fit in the 18-pin packages that we had which was quite a constraint. In contrast, Ed mentioned that TI had a design and TI's design was basically something about two hundred and twenty mils on a side and they had the luxury, in quotes, of going in a 24-pin package. And TI did announce its chip and had a photograph of its chip available in Electronics Magazine in approximately July of 1971. But, from any information that we were able to obtain, it was never fully functional, but it was a very good-looking chip and it was a nice wall size to use. And I might mention, and we can go into it in a few minutes, but from the history point of view, we had our photographs of the 8008 in the October/November timeframe as we were setting out to debug it and get

the chip ready for production. But with that, we had a lot of experience at Intel of developing new chips and putting a two or four-page datasheet around every chip that came out and equipping our salesmen with the whole stack of datasheets to hand out to the prospective users, be it designers or buyers.

In the case of the 1201, the major challenge that we had was that internally our one-page preliminary spec that we were restricted to having for every single product done within Intel, had to be on a single sheet of paper and every place in this spec it effectively said, "See instruction manual. See instruction manual. See instruction manual." And the real spec that went along with it was about thirty pages at the time with genesis from Stan's original document and it turned into a full-blown book. Intel got into the publishing business with the announcement of the product. Ed probably never knew or never expected as he came to Intel that we'd be selling books. Sometimes I think more books went out than chips.

Hoff: In fact, we made more money from some of the literature because we charged for the literature.

House: So you sold that data manual and the program manual.

Mazor: Datapoint did announce their computer. They called it the 2200 and they ended up building it entirely of TTL, and they ended up building a parallel version of it. So by the time we had our [8008] chip ready, they said our chip was slower than the computer that they were [already] producing. And I was in constantly talking [communication] with one of their [Datapoint] architects, [Harry Pyle] and they started to enhance their architecture. So over time, the Datapoint 2200 evolved and the instruction set got a little bit away from what we had done [on the 8008]. And also we didn't worry too much about software at the beginning of the [8008] project because it was for our customer, and he had the system and we assumed that they'd be writing all the programs. But later on as we got into the project, we had to start thinking about: "how are we gonna program this and what are we gonna do with it?"

Feeney: Right. And that became very significant from two points of view. One is looking at the customer, but the other side of it is internally. We had to test this thing. We were testing a computer. And so basically what we had to do was set up a system whereby we used a computer to test a computer. And we didn't have a lot of tools. We certainly didn't have any history on this because the 4004 and the 8008 were the first two microprocessors done by Intel and we basically had to set up a system in such a way that we would test the individual devices and present the full range of programming to them and exercise every element within the computer using our own homebuilt test devices. But as an example, I guess we can show the result of this effort.

House: Turn it [the die photograph] towards the camera.

Feeney: When each of these chips was built, we took a photo micrograph of the die and then we were able to go through the design and set up and partition a block diagram on here. In general, we had a bus going around the chip; we accepted all of the instructions and decoded them. We had this whole mass of random logic on one side to go through and set up all of the operations that had to take place from the arithmetic unit and from any other random logic work that had to be done. On the other side, we had a

very regular section of memory. You look at this in contrast to the chips that are developed today and you see how much more regular they are in their design, with massive, massive amounts of memory. But this was the very first attempt and first embodiment of doing this kind of thing and trying to do it in what turned out to be a relatively small chip then. Compared to today's chips [it is] a small chip.

House: So Federico had been busily designing the 4004 and [had] come up with some modules and methodologies in silicon technology that wound up being applied. Federico, could you tell us about how you moved from the 4004 and your involvement with the 8008 and your contributions to the 8008?

Faggin: Well, I was hired in April of 1970, just about a month after Hal. In fact, I believe it was April 3rd. And I was the leader of what was called the Busicom project in those days. It had no other designation, and [was] essentially four chips that consisted of the 4004, the CPU, 4-bit CPU and ROM and RAM and a shift register. Soon after I joined, I also found out that there was another microprocessor project going which was the 1201, the 8008. And I met Hal, not as soon as I got there because Hal was getting married and he was on a honeymoon, and so when he came back from the honeymoon, I met him.

Feeney: We met because we were sharing pullout trays on the same desk.

Faggin: Yeah, we had very, very narrow accommodations. But anyway, I found out that the 8008, 1201, was also another project going on. Actually I felt a little jealous of Hal because Hal had one chip to design, I had four. The 4004 was going to be the last one of the four. And so I thought that Hal was going to have the first microprocessor on the market. But, you know, I had enough to worry about in doing my own stuff and so I soon forgot about it. But basically, with the design of the MCS-4, the Busicom project, as it was later called, I had to create the entire methodology and design with silicon gate [technology] for random logic. There was nothing there. Intel had never designed random logic chips. They only had designed memories up until that point. Memories are much simpler from a design point of view, from a layout point of view, and from a testing point of view and so on, and there was not even an understanding of what it took to design random logic chips. For example, in the schedule that Vadasz had given Busicom, the layout of the 4004 was supposed to take four weeks, exactly just like the ROM or the RAM. And in fact when I got to Intel, the customer expected to have samples of the entire set by September that year.

I joined in April, [and found] that the company was supposed to start the project in October of that [prior] year, but nothing got started and so I found that I was six months late before I even started it. And the nine months that were required to do those four chips were really a tall order to start with. So it was a very difficult time and I had to do a lot of things with a very short time and essentially without any help. I was by myself. I didn't even have a draftsman and I had to basically invent my way out. Which is what I did. And so I developed the entire methodology, how to do it, the circuit blocks, the way you organized the entire layout, and the way you do the simulation, using graphics design so that you would quickly be able to size the transistors. We were not supposed to use computers because they were expensive in those days. And so I had my slide rule and graphs that I had taken from actual transistors, normalized, and developed a methodology to do that. Fortunately, I had invented the bootstrap load, which nobody thought was possible. That was actually what made possible the speed/power required to make the

microprocessor work. In those days without that you could not do it. You would have had either too many gates or too much power. And then of course there was the other thing that was necessary.

Another invention of mine was the buried contact, which was the ability to connect directly the silicon to the junction so that you would have like two-layers of interconnections. And so I got busy working and when the first run of the 4004 came out, it was in January of 1971, it pretty much worked. At that point, it was clear that we could do it; so Vadasz asked me to supervise the 8008 project. Hal Feeney of course did the detail design using the 4004 as a model, because the layout of the 4004 was very similar to what was required for the 8008. Many of the circuits that were used in the 4004 were used in the 8008 so that became a much, much more manageable project. It was not a team project by any stretch of the imagination, but it was, much easier because, there was a model that worked. If you see the layout of the 8008 and the 4004, unless you know what you are looking at, you cannot tell them apart.

Feeney: Being able to piggyback on many of the tools was just incredibly helpful for getting design moving along. And the other thing you point out is the fact that being limited on pins, limited on power supplies, whatever, that the bootstrap load became very, very critical.

Faggin: Yeah, it was essential. We worked together, for the first month or so. Basically, I transferred my experience to Hal and then Hal pretty much went on his own. Occasionally he needed some help in some issues, but fundamentally, Hal did the entire design by himself. When it came out toward the end of the year, toward the end of '71, everything was working except for a few errors. And by that time, Hank [Smith] was after me because he wanted Hal in marketing. He was coming to me every week saying, "You know, Hal has to come over. Hal has to come over." And I said, "Well, you know, he's got to finish the project. Then he can come over." And eventually...

Hank Smith: I did win finally.

Faggin: He won. Hal went to marketing before the 8008 was totally completed and characterized. The only hard time that I got with that was that after Hal left, I found a problem with the memory. That got me some sleepless nights - basically the memory was losing its memory. It was a dynamic memory and it was charge injection that I had to figure out a way around. Finally, I fixed that and that was it. The 8008 then was put into production, so after that, it was marketing's baby.

Feeney: And with that, just for the sake of this history, the 4000 series was announced in November of 1971 and then the 1201 then was dubbed the 8008 to be consistent with and to show that it was greater and bigger, etcetera than the 4004. That was announced in April of 1972.

House: Federico, tell us about the naming. These names didn't fit Intel's normal naming algorithms, did they?

Faggin: Under Intel normal algorithm, which consisted of the following, the first digit referred to the technology, so one was P-channel, two was N-channel, three was bipolar. And then the second digit referred to the type of device, one was RAM, two was shift register, I believe.

Feeney: No, it was random logic.

Faggin: Right, right. Two was random logic. Three was ROM and so on.

Feeney: Four was shift register.

Faggin: Yeah, four and five were shift registers. So following that logic, the Busicom set would have been probably called 1202, 1203, 1204, 1205, which I didn't like. So I said, "We should call it a 4000 family, 4001, 2, 3, 4," but Les [Vadasz] didn't want to hear about it. He just didn't want to hear about it. I was breaking the system. But I was insistent. I liked 4001, 4002, 4003, 4004. By the way, I found out that 4004 is the computed age of the universe if you look at the Bible. Of course, I didn't know at that time, but it was very interesting to find out. Eventually I wore down Vadasz and he gave in. He said, "Okay, okay. Okay, do whatever you want." But my logic was that it was a family, so it should have a sense of family, number one. Number two, there was a ROM and a RAM and a shift register and a CPU so they would have to be called [something different] according to the system, with names like 1202 for the CPU, 1305 for the ROM, and, 1107 for the RAM, or whatever. , It would have been a mishmash and logic finally won the day.

House: When did the 1201 get converted to the 8008?

Faggin: I think it was in marketing. I don't know who. Probably you had the idea.

Smith: Yeah, and then also the MCS-4 and the MCS-8. That moniker was put on to try and tie all of them together and not give the impression they were individual devices, but a family of devices that worked together. As a matter of fact, for the MCS-8, we kind of bastardized that because there wasn't really a family.

Feeney: We created it.

Smith: We created it and made that also a family of devices, although there were standard memory products in there.

Feeney: Following on that, because of the [low] speed of our processor and the demands of it, it didn't demand the highest performance RAM; it didn't demand the highest performance ROM, so it gave us the opportunity of designating certain memory devices that fit perfectly with the 8008 processor. And it gave

us the opportunity of selling memory devices that weren't quite as fast and EPROMs that weren't quite as fast.

House: So yeah, the 2102 and the...

Feeney: 8102 ...

House: ...was the static RAM. Actually, the 8004 speed selection probably charged more.

Mazor: When we were working on the 4004, I went to Gordon Moore and I said, "We're going to need an assembler for it." Well of course an "assembler" is someone that you use on the manufacturing line to build stuff. So, I ended up writing an assembler for the 4004 on the PDP-10 computer and [I] wrote it using the [PDP-10's] macro assembler also. But then for the 8008 [software], we ultimately went out to a contractor [Sandy Goldstein, of Gray Computer Co.], and wrote a contract, to have an assembler program written in FORTRAN for the MCS-8. So we started in the early days to do a "little bit" of software for the microprocessor's support.

House: Ted, tell us how management and the salespeople looked at the microprocessor in those days.

Hoff: When the product first became available, it was still considered custom for these two customers [Busicom and Datapoint] although we had one or two others for the 8008. But there was a very negative feeling within the company toward making it a standard or full-blown product. The argument was [that] Intel was a memory company and if we go into the computer business our customers, the ones using our memory chips to build mainframe computers, are going to see us as competition and they'll be looking to get their memory chips from somebody else. So it could hurt the memory business. That was one negative. The other was that it was hard enough to train people in marketing to go out and deal with customers for the memory chips. Computers were seen as an order of magnitude more difficult and the training and the support would be terrible. And I remember I had a meeting with, I believe it was Bob and Gordon, but it was Bob Noyce I was talking to primarily, sometime I believe in the early summer of 1971. And Bob said to me something like, "Well, we're not ready to make a decision to announce this product. We have, you know, a lot of other..." And I said basically, "Every time you delay announcing it, what you're really doing is making a decision not to announce." And I said, "Eventually, somebody else is gonna open up this business and we'll be left behind even though we're the first ones to have products." So it really wasn't until Ed joined Intel that I found a positive attitude in marketing toward offering these as a product. Up until that time, there was a lot of negative feeling.

House: Hank, you were there before Ed joined, so tell us about your involvement with this product. You were obviously involved with both products at the same time.

Smith: I joined Intel in early '71. My initial responsibility was to market and support the memory chips, the 1103, the 2102, the 3101; whatever the right nomenclature was for bipolar. That was my initial

responsibility. When the microcomputer was developed and we were talking about marketing the microcomputer, I got very much involved in developing all of the support documentation that was required naming the family, putting together the packages that were gonna be required. Very early on we understood that this was very different than selling a memory device. Microcomputers were gonna require a lot of service, a lot of support. Engineers typically were not used to programming. This was an entirely different way of designing. In following on what Ted said, being first was extremely important because you developed a very loyal customer. Once they put their effort into developing the software and the instruction sets that were required, it was very difficult then for them to move and try somebody else's device. They were locked in. So our initial marketing focus was to get this thing designed into as many applications as we possibly could. This was when we were selling the MCS-4 family and the MCS-8 family.

House: So how did the transition get made from this is strictly custom product for two customers to a general purpose computer chip for the general marketplace?

Smith: The MCS-4 family was introduced in November of '71 as a standard product. As far as I was concerned, that's what we were selling. It was not a custom product and the decision had already been made at that point.

Feeney: But if I could interrupt though, at that time, there was a restriction or a requirement that any customer using that product sign an agreement that they would not use that product in a calculator-oriented environment.

Smith: Yeah, I think that was part of the original agreement with Busicom.

Feeney: I think it was standard.

Smith: Yeah, correct.

Gelbach: [Later] we renegotiated that out, by the way.

Smith: We also found that the traditional marketing organization wasn't correct for marketing a product like this, and Ed set up the Microcomputer Systems Group. And I don't remember when you actually got there, Ed.

Gelbach: I think in June 1971.

Smith: '71? We formed the group in early '72 and Microcomputer Systems, which I was given responsibility for, had marketing, which was Hal, had software, had engineering, Phil Tai, we had our own manufacturing. It was a complete autonomous entity. And we initially developed what we called

simulator cards, the SIM-801 and the SIM-401, to make it easier for a designer to design in the product. It then occurred to us also that we really needed to have a high level language to develop these products and I brought in Gary Kildall.

House: How did you find Gary?

Smith: You know, I really don't remember what the sequence of events was. Do you remember?

Gelbach: I dealt with Gary either with you or after you because we were gonna buy the operating system from him. I forget what happened to it.

Smith: Well, he developed PL/M at Intel. He had an office there where he developed PL/M.

House: Did he develop ISIS as well?

Feeney: No, PL/M was the one that he developed.

House: Where did ISIS come from?

Mazor: - We were using PDP-10s in the company [Intel] and it had quite a nice operating system so we developed internally the [8008] operating system, which had the look and feel quite a bit like the PDP-10. I'm talking about the commands the user would see like "DIRectory", "DELeTe", "COPY file", and so on.

Smith: But the interesting part is, is that PL/M eventually evolved into CP/M and Gary started Digital Research. There was a real battle at that point as to whether CP/M or MS-DOS was going to be the operating system for the PC.

House: But PL/M was really a language and CP/M was an operating system, correct?

Smith: PL/M was based on *PL/I* and then CP/M was based on PL/M. And if things had worked one way or the other, Gary could have had the operating system for IBM's [personal] computer.

House: So PL/M was an operating system?

Smith: The forerunner.

Mazor: PL/M-8 was a programming language which was written as a cross-compiler [for the 8008]. And the 8008 was a “pretty primitive” assembler language, and PL/M-8 was a “pretty primitive” high-level language, actually. Gary [Kildall] worked alongside us at Intel and one of the things we did is provide him with a floppy disk as a form of payment, and for the floppy disk he developed an operating system. And I think our development of ISIS inside [Intel] sort of paralleled the work that he was doing. And that operating system [of Gary Kildall] eventually developed to become CP/M at that time also had a [PL/M] compiler that ran on it [his operating system]. And it's quite possible that the operating system [itself] was written in PL/M. I don't remember that myself. We [Intel] had a tight relationship with Gary but of course he founded Digital Research and never did join Intel, although we approached him on numerous occasions [to become an Intel employee].

Smith: What followed after that was the development of the Intellec systems, the Intellec 4 and the Intellec 8 which were basically primitive personal computers, if you want to call them that. They were the forerunner of the MDS system and then of personal computers. They were designed in our group and were marketed as a separate product.

House: Did they have any similarity to the SIM-8 and SIM-4? Did they use the same electronics inside?

Smith: There was a lot of commonality. Yeah, absolutely.

Gelbach: Actually, it was the same boards and we just put a box around it.

Feeney: Well, no we made two different moves there. We had the SIM-801 in its own box and then later on, we translated it to, I believe, boards that were similar in format to what the memory systems division had so that we could utilize some of their memories. Those devices were critically important, mainly because they were used both as PROM programmers or sometimes standalone. The PROM was really the unsung hero for the development and the growth of the microprocessor business. Customers were able to develop their programs, ship their initial systems, be able to use those systems and sometimes [chose] not even to move over to ROM, but just continue to use the PROMs as their shipping device as they changed the software.

Smith: We were really the first [company] to have the simulator boards, to have the high-level language, to have the Intellec [type] systems. None of this stuff existed and so we pioneered those products and actually had a pretty good business selling those products. And it was critical I think to the success of the microcomputer in being able to provide that kind of support to the engineers out there who were trying to design these products in.

Gelbach: We also started the application engineering group that was knowledgeable in some of the computer areas. Prior to that, it was more discrete memory. Then we changed the concept of the sales force, but more predominantly, the application group to could go in and talk computer language to all of our customers, which most of us at that time could not.

Feeney: I think our first application engineers were hired in the '72 or 1973 timeframe.

Gelbach: I thought it was earlier.

Hoff: I think it was earlier than that, yeah.

Gelbach: I had kind of started the concept of application engineering at TI and [had] hired them there as component people. And I think I started it [at Intel] relatively soon after I joined. So I would have guessed probably early '72, mid '72 maybe at the latest. It was a necessity as no one in the company could really speak the language except the microprocessor group.

Smith: But I think, as is true of a lot of successful inventions or processes, when you're doing it initially, you don't have any idea what the implications are going to be long term. It was just fortuitous that we had the Busicom deal and the 1201 that this evolved. It would be nice to say that we all knew this was something really significant and it was going to have a significant impact, but I don't think any of us really appreciated what was going to happen in the future.

Hoff: Well, I know one of the discussions I had with Bob Noyce was how would we support this? And this was early '71 or early summer. I think I said I thought that my group could maybe support a dozen or so key customers. Now, that was a good number because that was the number of key customers we had for memory chips. And then I said that for rest we would provide literature, but they'd be on their own. Now, I kind of figured in the back of my mind if enough of them came out, we would have to do something [else] to support them. But at least it was reassuring.

Feeney: The thing that's very important with that as Hank and everybody were pointing out is that development systems carried the company for a while - or carried the microprocessor business for a while. We were certainly making greater revenue on development systems than on the individual devices that were going out. And at that same time, the number of manuals that we were printing was an incredible number and Stan and Ted were going out doing seminars for customers around the U.S. and then gradually it became around the world. But backing up a little bit before that, there was a pushback in early '71 where customers like Seiko were saying, "We want to learn more about what you're doing in microprocessors. We have an interest in using those processor products." There were customers from Europe and customers from Japan that were expressing interest very early on, even before we had the device laid out.

Faggin: Exactly. The decision to restart the 8008 was based on Seiko's interest. As you remember Ed, you and I, visited Japan.

Gelbach: That was [the incident] with the keys?

Faggin: With the keys. Yeah.

Gelbach: We're both leaving together to go to Japan. Federico was waving out at the airplane and I said "What are you waving for?" And he says, "I have my car keys in my pocket and my wife is probably trying to figure out where I am right now."

Faggin: But, anyway.

House: You were going to see Seiko and they had an interest in the 8008?

Faggin: They really wanted the 8008 at that time. They were very serious. They already had pretty much decided to use it, and the reason why the 8008 was restarted in January, as opposed to being simply abandoned because Datapoint was not interested in it anymore, was that wanted an 8-bit micro to do a scientific calculator. That was what got that project restarted.

Feeney: In a very early advertisement brochure that we put together, we had a number of different applications for Intel's microprocessors. One of the applications was the device that Seiko had. It was very similar to the large calculators from Wang and also from Hewlett-Packard in the 1970, '71 timeframe.

House: So what were the other applications that we had envisioned at the time?

Feeney: Well, we've got a number of applications here [in the advertising brochure] and these weren't applications that were envisioned, these were ones that customers actually did. We put together this brochure of Intel microcomputer [applications] of the Seiko calculator, a blood analyzer, a business machine, -and I think this was from Comstar, a very large mechanical device or mechanical control device. We had a point of sale machine from Staid, a communication management device from Action Communications, another point of sale device from Pitney Bowes-Alpex...

Smith: Point of sale was a major application.

Feeney: Also an Inforex check processing machine. So these were all early thoughts, early concepts of how things could be done. It was amazing how quickly the customers latched onto this. The date on this brochure is April of 1973, so it's basically a year to year and a half after our announcement of the products that these devices [the finished products] were in customers' hands.

Hoff: We worked very closely with Regis McKenna on developing the advertising and support material. And they came up with some pretty clever...

Feeney: There was a traffic light ad, right.

House: I remember talking about traffic lights and blood analyzers.

Gelbach: Traffic lights were a big application at the beginning because we needed [something] to talk about and everybody understood it.

Smith: Talk about anything, yeah.

Feeney: It was hard to talk about industrial control and some of those things, but the things that would touch people's lives [were helpful].

Feeney: Toledo Scale was a big early customer.

Smith: Computers in a teacup was something that one of Regis' guys came up with. The brochure was tiny also. It was done just about the time of the National Computer Conference. This was in '73 as well.

House: Tell us the story about going to the National Computer Conference for the first time and the reaction you got from management. This was a new era for Intel.

Smith: Yeah. Well, Ed and I both were in the room with Gordon Moore and Bob Noyce. We were outlining what we were doing and what our plans were.

Gelbach: And we had all the displays on the table, I remember. And there were the standard blue boxes [development system], but there were also brochures and posters.

Smith: Oh yeah.

Gelbach: And we were really excited. We were going after this business and we had worked on it six months maybe, a year, whatever it was. And it was the future of Intel. And we had no doubts this was it. And as Hank's explaining it, Gordon's face is getting longer and longer.

Smith: And redder and redder.

Gelbach: And at the end of it, or somewhere I must have stopped and said, "What's wrong?" And Gordon just made it very clear we weren't gonna do that, you know, and so "Okay."

House: But you went anyway.

Smith: We did. We modified some of our claims and we took off some of the more obvious references to computers, but we went anyway and we introduced it at the show. But I have never seen Gordon either before or after as angry as he was then.

House: And what was he angry about?

Smith: Well, it was the fact that we were going to be competing against our primary customers [as far as Gordon was concerned]. We were not going to introduce a computer. This is not a computer.

Hoff: I had an impression after we announced the products, and I ended up going out with marketing people to visit a number of customers, particularly large computer companies, they had been developing some concepts for controllers and their peripheral devices. And when we came out with the microprocessor, they were able to use the microprocessor and it ended up [that] they didn't have to do all of this elaborate development that was going to cost them a lot. We were offering almost everything they needed. So it seemed to me that we ended up complementing the big mainframe guys in some of these peripheral needs.

Smith: [The microprocessor] became the peripheral controller. And of course history bears that out; that we never really were competing with those customers.

Mazor: However we did have a chance to visit [customers], for example, I visited Digital Equipment Company [DEC]. And you have to understand that if you have a component which is a central processor and you go to a CPU designer, he's rather insulted. You know, "How dare you. Who do you think you are?" And also I think Federico will comment that we had a lot of criticism. "Well, why did you do "that", and why didn't you do it this way?" because obviously as CPU designers, they knew and they cared [about CPU design].

Faggin: Yeah. My first visit to a customer after the visit to Seiko with Ed [was] a few months later and we went in Europe. And I visited Nixdorf and Plessey. I remember they were incensed about the idea that a chip company could actually have microprocessors. Really, they did not like it at all.

Faggin: Nixdorf was absolutely anal about it. I mean they were just, you know, "How did you do that?" and, "You cannot do that. That doesn't work." And, "That's bad. That's bad." The contrast with other customers where you had customers that wanted to-[use it was incredible]. The wanted to make a blood analyzer in Denmark. I forgot the name of the company, and they were elated because they finally got something that they could use to solve their problem. And so you could see the difference in mindset. But at the same time, it was useful to get that feedback because in some sense, they [Nixdorf] had some good comments. Some of the things that they wanted to see in a computer that were not there became ideas for the next one.

Smith: A lot of these customers found that they could develop a family of products fairly easily by just changing the instruction set and they didn't have to redesign it.

Feeney: The other part of it was that the customers at one point found that the only alternative that they had was going and putting a minicomputer in this blood analyzer or in this scale or in this point of sale device. And it's probably in some respects lucky for us in the microprocessor industry that we weren't competing directly with the minicomputer manufacturers at that time because it allowed us to get in in somewhat of a stealth way. The initial customers were all not on the tip of Intel's tongue as far as major customers go. In fact, I recall a discussion I had with Ed one day, he was looking at our customer list and he was saying, "Hal where are all the big name customers? We're dealing with UNIVAC, we're dealing with Burroughs, and we're dealing with IBM in the memory side of the business, customers [who] rank in who's who. I look at your customers and it's always who's that?" But it gave us a baseline to sell lots of microprocessor development systems, and then gave us the way of developing the base for all of these different computer applications. On top of that, when the CDCs, the UNIVACs, the IBMs did begin using Intel devices in volume, they were using them in their peripheral devices where they wanted compute power and they couldn't afford or it wasn't practical to get the compute power from the big iron computers that were available at the time.

Mazor: I think it's worth saying here that the real application of the microprocessor was in control applications-- inside instruments or systems of various types, but not as a general purpose computer. We had three or four companies that wanted to build something that you could more or less call a "personal computer". (We tended in those days to call it [that application] a "hobby computer".) But the limiting factor, one has to recall, is that printers were very expensive, that magnetic storage was very expensive, and so someone might build a personal computer, but it really wouldn't have any capabilities [without these expensive peripherals]. So what held back the personal computer (as we know it today) was the development of these [low cost] peripherals. And we shouldn't underestimate the requirement [need] for those peripherals which we now use today and that are very inexpensive.

Hoff: When we looked at the market, we saw what is now called embedded control. It seems like the media is not even aware of the existence of embedded control. It's always the desktop or the notebook computer, [that gets the attention] but it's still a pretty significant market, but it does not get nearly the publicity that the other machines get.

Gelbach: Orders of magnitude, higher volume in terms of number of units.

Feeney: The earlier concept with selling the microprocessor, in traveling with Bob Noyce, Bob always tried to describe the microprocessor the same way you think of a fractional horsepower motor in your house. You've got your furnace motor, you've got your washing machine motor, your dishwasher, your garbage disposal, you've got motors in all of your clocks and whatever. And at that point, it seemed like Bob was certainly an optimist in terms of the way the microprocessor would become pervasive. But on the other hand, if you look at it in retrospect, [he was] probably a pessimist because when you look at the BMW with fifty or sixty microprocessors in a high-end version, it's just mind boggling. I don't think any of us could have seen how far this would go.

Hoff: When the three of us were recognized for the 4004 at the Inventors Hall of Fame, they told the audience, "You might have ridden here in an automobile that was designed with the aid of a notebook computer. They totally missed the point. Almost everybody probably drove there in an automobile that was running under computer control.

Mazor: I need to mention that Ted had some very innovative ideas when we were playing with the microprocessors in the early days and [he] started several projects. One of them was a video game that was played on a television set [using MCS-4] and we actually had [built] a primitive "Space War" game with objects moving on the screen. And in an applications [engineering] group, the idea is-- can you find new applications that you'd kind of like to go into business with. Well, management, not being familiar with that application, thought it was "pretty silly" to play a game on a television set. And similarly, another application we talked about was building a scientific calculator-- I'm talking about a handheld small scientific calculator. And I recall an upper-management man had a slide rule on his belt and said, "Well, why would you ever want one of those?" And so I would say we [Intel applications research] had some visions and some of those visions were subsequently fulfilled by other companies and our customers.

Gelbach: The microprocessor was a boon to Intel in non-obvious ways. Number one, I think we looked at it [because] the customers were very interested in it. We would get more requests to talk about it. Nobody wanted to talk much about memories. They understood the memory sequence was going to go from one K to two K to eight K, whatever it was, but the microprocessor tickled the imagination of all of the engineers. Every customer that we ever had for whatever product always wanted to talk microprocessors. It was the entry door. We used to give seminars that we eventually charged [for]. No one had ever charged for them before and we just came up with the idea, "Hey, let's charge twenty-five bucks apiece and at least it'll keep the riff-raff out." So it turned out to be, and I'm gonna say this, a little simpler on the sell than it sounded because it was such a new product. But everybody except computer engineers wanted to know what it was. We used to push it that if you don't understand it, you're gonna be left out. It was really clear that this is the next generation. So was it really lucky or was it skillful.

Faggin: Or both.

Feeney: Boy, part of it was a matter of creating our own luck or creating the opportunity for it. I think Intel did it by having an environment where it could happen. Certainly not intentionally getting into this business, but by creating the environment where this could happen and ultimately change [the direction of] their business.

House: If you look at microprocessor history, there were a number of companies that came out with processors in the early '70s and mid '70s. Few of them, with the exception of one we'll talk about in a subsequent interview, Zilog's Z80, ever actually became successful products. I mean RCA, Intersil. You've all seen the microprocessor history poster. There's a whole bunch of products that came out, but [they] had no successors.

Mazor: Well, I can speak about a few of them [competitors] because we [Mazor, Faggin] had been at Fairchild. Fairchild then went on to [build] a calculator chip which was [bit] serial and focused on

arithmetic and didn't focus on the peripheral [control] problem at all. So I think Fairchild did just what we [at Intel] thought was the "wrong thing" to do. Interestingly enough, I think it was RCA that built a version of the PDP-8 and neglected the fact that the [PDP-8] subroutine return mechanism wouldn't work in read only memory [ROM]. And one of the big applications was control applications with read only memory-- in which case you have to have something like the stack [for subroutines]. And I remember people saying, you know, how ridiculous it was of us [Intel] to invent a new architecture. Why didn't we, you know, go with an existing architecture, not realizing that those architectures were not designed around [program stored in] read only memories and might fail for these reasons. And another one, that I like very much is, AMI, which was a very successful "custom house" in Silicon Valley. They thought the 8008 was a terrible design and they could make one [CPU] that was much better. When they completed the masks for it, there was no table large enough in the company to hold the artwork for it, and [so the chip was so big that] they were never able to make that chip (called AMI 7200). So, you know, competent people could say, "This is what you really need to do," but if you weren't concerned about and didn't realize the limitations of the package and of the power and of the silicon, if you didn't have the expertise like guys, like Federico and Hal, then they would do the wrong thing; and they wouldn't succeed in the market (or they couldn't build it at all).

Faggin: On the other hand they were successful from the point of view of working microprocessors, for example the PPS-4 of Rockwell. It was announced probably a year after the 4004, maybe even less than that.

House: And that was successful?

Faggin: That was successful, meaning it worked and it had some applications. But what Rockwell did not do, they did not do nearly as good a marketing job as Intel did. They basically were selling it them as chips. There were more like custom things, so they would support you, but they would not have tools, they would not have the paraphernalia that we had developed at Intel. And as such, they had very few applications and that was it.

Feeney: Many of the other companies also didn't have the potpourri of products to go around the microprocessor. When we were first marketing the processors, we'd look at the ratio of memory to the processors and very often the value of the processor set going out was weighted toward other chips in the Intel portfolio and not to the microprocessor itself. That was important as far as being able to take in funds for developments like this. I want to go back to what Stan was saying for a minute where he talked about the AMI chip and the size of the chip and the tables and tools. When we did the 8008, it was I think the last chip that was drawn at five hundred times actual size and the drawings were all done by hand. It was also the last chip at Intel that didn't have any computer layout support with it. We did it by hand and we actually had to do the drawings in two separate halves.

House: When you say do it, you mean the mask-set design?

Feeney: Yes, doing the mask-set design. I can't remember the exact width; it was either four feet or five feet wide, of the vellum-style paper that we were using. So we could only do one-half of it at a time. We

had to set it up in such a way that we could then later paste it together on a table that was the size of a ping pong table. So everything after that was done at about two hundred times the actual size, not five hundred times the actual size.

House: Tell me about the use of computers in the design of the 8008.

Mazor: I'll make a comment about it and that was that-- Ted had written a logic simulator for the PDP-8 and we also had subscribed to a time-sharing service, a company called Applicon (in Boston) that had a time-sharing logic simulator. I was Hal's officemate and when we were doing the logic design or he was doing the logic design, part of my job was to do logic simulation [for Hal] and verify that the logic simulated okay-- as a way of verifying it [the design]. And other than that, we did not have a lot of resources for computers. What do you remember about it, Hal?

Feeney: Well, we used a time-sharing service for some resources. At best we could do one or two or three transistors to calculate propagation delays and calculate the size of the gates, size of some of the resistors and whatever, and to generate speed for the devices. We would go through and model small circuits as part of the design and then try to replicate or reuse those small circuits as much as possible throughout the design. I think I ran up something like a fifteen thousand dollar bill at one point for time-sharing services and I was told that we shouldn't be spending that much money on design tools. What a contrast now to the design teams and the millions and millions of dollars that are put into it. It's just a totally different viewpoint. We were bootstrapping our self into that business and from a management perspective had no concept of the business we were getting into.

Mazor: I think if you ask Federico what tools he had, I think he's had three colored pencils.

Faggin: _____. Yeah. But to go back to the question of why was Intel all alone. There was another aspect that's important to note and that is that Intel had silicon gate technology. All the other micros that were done were all using metalgate [technology] in the early generations. So they were way too expensive. They were not as cost-effective as silicon gate. The 6800 it was the first device of Motorola that was silicon gate. Also the F8 that was done by Fairchild was silicon gate. So only with that generation, about 1973-74, where the competition had a competitive technology to what we had. Then you started seeing competing devices. Before that, they couldn't compete in speed, never mind in size or cost.

Feeney: Well, the other side of that is because the EPROM [erasable programmable [read-only memory](#)] came out within maybe a month or two before the 8008 came onto the marketplace. It was initially very difficult [to use]. It was I think, a program once-only type of device. Sometime after that they came out with the quartz lid so it could be erased by UV [ultra violet light]. Even then it was very difficult to program. So how do you make it easy for the EPROM to be programmed? You use a microprocessor to do it. With the ability to use a microprocessor, the programming was done. Also the programming algorithm could be changed as needed as we learned more about how to do it. It was also less expensive. So then customers started buying the microprocessors to do that. There were other third-party stand-alone program manufacturers for EPROMs and those third-party groups were selling in

competition with Intel, but our devices were less expensive. So you've got a convergence of technologies that synergistically supported each other. And then just the growing awareness that sales pushed from the development system perspective made it happen.

House: Let me turn to Stan.

Mazor: Well, I'm an engineer, not a marketing guy, but I did reflect on Jim Lally who was leading the Development System Group for quite a while and quite successfully, and he was quite good at analyzing the market. He believed that the development system would be sold to engineers, and that market was relatively small. In that in that kind of marketplace, which is not elastic with respect to price, then the strategy is to raise the selling price by adding features, which we did over time, e.g. by adding in circuit emulators and other features. I've conjectured that had he made the opposite decision, conjectured that this was an elastic market; he would have found ways of reducing the selling price. And then if that was the case, the [elastic] market would have evolved much larger and Intel might have found itself in the personal computer business. We certainly know that we were selling the [MDS] product not as a general purpose computer, but as a microcomputer development system [MDS]. I will further add that [even] Digital Equipment Company didn't sell computers; they sold Programmed Data Processors [PDP] because they also didn't want to compete in the "normal business channels" of selling computers. And we [Intel] might have been through the back door into the personal computer business as we know it today.

House: So the development system had a keyboard, it had a display, it had DRAM, it had floppy disk, and later got a hard drive. It had all of the things that were in a personal computer. What stopped it from being the personal computer?

Mazor: We even had customers who would tell us that if they wanted to buy a computer in their company, it was a really big deal. They had to have permission from their "computer Czar" but they could buy a [Intel] development system and they could quote, use it as a computer. That was an easier thing for them to do regardless of the price.

Feeney: And a lot of customers bought it at that level and when it had an operating system, it had WordStar on it and other tools. I think there was some version of Multiplan that was on it too. And so they had tools available to them to use it as a personal computer.

House: Later we made a portable version called the personal development system, PDS, which looked like a Compaq. Later Compaq came out with their luggable and that was sold with a word processor and a spreadsheet on it. And so it was about as close as we got. In fact, Jack Carsten tells the story that [Rod] Canion came and visited him and he showed him the PDS. It wasn't too long after that that Compaq was created with a product that was quite similar. At least [Jack] Carsten believes that was the genesis of the Compaq luggable.

Mazor: And that's not to say we would have succeeded in the business as we know it now, but it still raises a question.

House: I understand management was afraid to go into the computer business, again from [fear of] competing with the customers.

Mazor: Yes.

Faggin: But I also believe that the potential with the personal computer was not seen at all definitely at Intel. I remember Gordon being quoted as saying, "Well, what do you do with a personal computer?" Basically, it was not understood, it came out of a different milieu. It came out of young people that were enthusiastic. They saw [the] potential [although] it was not clearly articulated. In fact, the first applications of personal computers were really silly applications that eventually developed into the mainstream technology. So I think that there are other aspects to it, other than just the lack of vision of Intel management. I think that it was ...

House: An idea before its time.

Faggin: Yeah, before its time, and by different people.

Smith: I was at Venrock when Apple, got started. As a matter of fact, we were investors in Apple. We had no way to evaluate the initial business plan it because there was nothing even remotely close to it that had been done before and the projections were ridiculous. You just had to do it on the faith of the people that were involved. And you're absolutely right; the first applications that they talked about were, doing your menu planning in the kitchen and, these bizarre, ridiculous things. And actually, Apple didn't really take off until there was a real application which was VisiCalc. That really is what got that product going.

Mazor: The so-called killer applications were certainly probably: text preparation and printing and then also VisiCalc

Faggin: VisiCalc was because there were word processors before Apple. You did not have any way to do spreadsheets except with a large computer and even those were difficult to use.

Hoff: We found a number of people in our marketing departments buying Apples and so on just for things like VisiCalc so that they could to their [sales] projections that way.

House: Ed, I understand that when we first started selling microprocessors, it was really coupled with EPROM. Tell us that story.

Gelbach: One of the early products that Intel came out with was a programmable or erasable programmable ROM. The intent was always that the customer would use it as a design aid. Use a couple of 'em and then order in large quantities a fixed mask-ROM. As it turned out, the microprocessor was the perfect tool to fit with an EPROM because customers always wanted to change the programs or somehow the way it worked. They would start off and say, "Well, we only need about a hundred of 'em. We go into production, we're gonna change it to a hardwired ROM." And as it turned out, none of the customers ever changed. And the cost, or at least the selling price of an EPROM was probably in the hundred dollar range and the cost of a ROM was in the three dollar range. So we were more than enthusiastic about just letting the customer design and use all the EPROMs he wanted. It was a significant profit margin for Intel; in fact, probably was the largest of any product Intel has ever made, including microprocessors, from a margin viewpoint. So it turned out to be a really good strategy because EPROMs or programmable products, they aren't EPROMs anymore, have a large market.

House: So you were more excited about selling the EPROMs than the microprocessors.

Gelbach: In the beginning, absolutely. And you get smarter as the days go on.

Feeney: Relative to the comment about getting smarter, I'd like to relate something that I found as I moved from engineering into marketing. I was kind of on the frontline as an application engineer working in the factory dealing with literally thousands of different customers calling in. And instead of looking at the microprocessor through my own pair of eyes, [or Ted's or Stan's or Federico's] you're looking at it through a thousand different pair of eyes. "Can we do this with it?" "Can we use it in this mode?" "What about these instructions that you have in it or what about some instructions that aren't being used?" "What happens with the processor if we apply those instructions to the processor?" "Can we use it this way or change things with memory?" or whatever else it may be. There were a myriad of different customers looking at it in different ways. Going back to Federico's comments about dealing with the customers in Europe and some of them saying, "You know, it's got difficulties here and problems here," and so forth, you'd go and talk to the same kind of customers in Japan and you would find that the Japanese would even be a little bit more critical and say, "There's a problem here, there's a problem here, but we can fix this problem, we can fix that problem," and they would race off into production with a product and not have any concern at all. So that was more of an attitudinal thing. Of racing ahead of the market, getting products into the market using new, innovative technology. These customers, the visionary ones, were the ones that were driving all of this and that ultimately drove the success that led to the devices that out there today.

Mazor: I think upon reflection, we [Intel] were a memory company and we were really paying a lot of attention to two things: one was the small die size and the other was the number of pins. If we look back on the 8008, something that handicapped it was probably the number of pins in the package and had we been more aggressive in that regard, we would have had a perhaps more successful and better product. Interestingly enough, TI, when they did their chip design, they used some automated tools and not all of the area of the die was "active" so the yield is not therefore so badly impacted if you get a defect in an inactive portion of the chip. So perhaps in our general methodology we might have achieved faster turnaround times in subsequent designs if we hadn't had the mentality of a memory company. But that just raises some questions. Obviously, we were successful in what we did.

Feeney: But, Stan, let me go back to a fact of life at that time. The fact of life at that time is that we did not have the pizza-sized wafers that we're dealing with today. We were dealing with two-inch wafers and we had to get as many devices on a two-inch wafer as we possibly could. And as a result, to have a number of good die candidates, and I'm not sure how many were on there, but we're talking about something in [the range of] one hundred, two hundred devices on a wafer. So from a production point of view, getting the small chip was really critical. I mean we used every allowable square millimeter that we possibly could get inside of the package. In fact we were right on the edge of filling the cavity [of the package] with as much as we could. That was the defining size for us. But you're right. If we could have gone to 24-pin packages or something larger, it would have made a difference, but it may also have made a big difference in the overall cost or the productability.

House: The cost was never a limit though on this product. There were not enough of them made where the yield was a significant impact. This was the launch pad that got you to the 8080.

Feeney: Right. We talked before about numbers and I'm not sure what Ed recollects, but 8008 certainly was produced well into- or at least available to customers well into the late '70s, early '80s. And we're talking probably about hundreds of thousands of devices being produced, not millions.

Gelbach: We figured out how to price it 'cause we thought it would be eventually a replacement for the IBM 360. I think we started the first pricing at three hundred and sixty dollars.

Feeney: No it was the 8080 that went to three hundred and sixty dollars. The 8008 started at a hundred and twenty and we had a speed-enhanced version of it that was at one hundred and eighty. And so for the 8080, since it was better, it had to be higher in price and that's where the three sixty came from.

Faggin: My frustration back in those days was basically the fact that Intel was a memory company. Microprocessors were not really seen as something by itself that needed to be really pushed. My frustration was lock-in pin count packages. When I came up with the idea of the 8080, it took me nine months to get [Les] Vadasz to give me permission to do it. I sent him a memo back in early '72 to do the 8080. It wasn't until the end of '72 that I could start. And the fight that I had to put [up] to get a 40-pin package, you cannot believe it, because 40-pin package was not acceptable. Fortunately, we had a custom project. It was a single-chip calculator that had to be in 40-pin. So I said, "Now, Vadasz, you tell me. Why can we do a calculator on a 40-pin package and not a microprocessor on a 40-pin package?" I nailed him with that one.

House: He had to give up then.

Faggin: Yeah, he had to give up. But it was always a fight. For meat Intel it was always a fight to stay on the crest. We had started the business and we were pulling back instead of going all the way. For example, to get programmable peripheral devices, okay? They were basically memories with a different name - very small, simple chips. To have a coherent family of chips, there was no interest. It was too much of an investment and so they didn't want to go down that road. That is why I ended up starting my

company because of that frustration. And the Z80 was conceived from the very beginning as a family, as a family of components that would be seamless, would work together well. It would make sense. So I need to register because it was that frustrating for me.

Smith: Well, I think we got that.

Hoff: I think one of the problems was that so many people involved in Intel were out of the traditional semiconductor business and that bias carried through. There was a few times where [a] new package caused problems. At one point, they were gonna save a few cents on lead frames and they went to a punched lead frame instead of an etched one and it caused reliability problems. So there was a real negative view toward anything different in the way of processing or packaging even though the business of the microprocessor might have easily supported the research and everything necessary to make it work.

Mazor: On the positive side, with the EPROM we went to a quartz lid which was quite a revolutionary thing and we had ceramic packages which were more expensive as well, albeit they carried a high average selling price.

Faggin: Yeah, but it was 24 pins.

Mazor: I'd like to say one other thing and that is in November '71, we announced the 4004, and in April '72, just four or five months later, we introduced the 8008. So we really had a "one, two punch". And I would liken it later to having the "motorbike", the 4004, and the "station wagon", the 8008. 'Cause the 4004 was for control applications with a very small part count, three or four chips and doing 4-bit arithmetic, and the 8008 required some other chips around it, but could use general purpose memories and could run out of RAM in particular--- and so it was the more general purpose solution. So here we were in the early days with two complementary product lines, that were in no way compatible in terms of software or hardware features otherwise. We were in a very interesting starting point with these two lines.

House: Which one sold better?

Faggin: I think the 4004 sold better.

Smith: I think the actual volume, the 4004 did.

Feeney: It got into cash register-type applications. NCR was a large customer for a very long period of time. The 8008 got into some fairly high-running terminal applications. But those certainly didn't run at the same rate or same volume as either a scale or a cash register application. But want to go back to the comment that was made about having special things and needing special things, and Intel's view. This is a very early 8008 user's manual of fifty, sixty pages on how to use it, how to get started with an 8-bit

microprocessor. We worked to meet the launch date and to get the user's manual out so we had all of the things that the customers would need to answer their questions. At launch, we didn't have a four-page datasheet available in Intel blue. And there was criticism because we didn't have the standard datasheet that everybody was accustomed to but there wasn't a recognition of the desperate need for documentation and more documentation. And this little manual grew to be a much bigger manual as time went on. And we all know the story today about the volume of books that are sold with every microprocessor that goes out.

Smith: So we sold a bunch of those books too, as I recall.

Gelbach: I remember we were the first one in the industry to charge for literature, which was unheard of. But it worked.

Smith: I want to go back to Federico's frustration and Stan's comments. You have to put it into perspective that Intel was founded solely as a memory company and its sole objective was to design and produce semiconductor memories. We kind of stepped into this microcomputer [business] by chance. It was never planned to be that way. Everybody at Intel was really steeped in memory design and manufacturing so it's not hard to understand that there would be frustration and that there wasn't the commitment. And of course nobody had any idea how large this was gonna get. So it took a lot of pushing to get the microprocessor going.

Hoff: When Intel was first formed, Bob Graham and I published an article predicting that we would get the cost of memory down to a penny a bit. It was figured at that point we'd compare with cores. I found a price list from, I think, September of 1972. 1103s were selling in like hundred quantities for around seven dollars, so it was well under a penny a bit, so we met the goal. In fact, one thing that I thought was kind of interesting when we looked at the sales of the 8008, and I was once asked after the 8080 came along, didn't the sales of the 8008 just evaporate? And I saw a plot of it about a year after the 8080 had come out. What had happened is that the 8008 had ramped up in production and then just kind of leveled off with the 8080 and it just stayed constant and then the 8080 started to ramp up. So apparently, the sales continued on well after the 8080 came out.

Feeney: One of the earlier customers was Kodak in a large photocopying machine. Kodak was not an electronics company. This was really one of the first electronic ventures that they got into. Once they got it designed, they didn't have the money nor did they have the inclination to go in and totally redesign it. And there were many, many applications like that. So it's only when the applications totally turn over to something new [that the design is changed]. Some of the first electronic cash registers in McDonald's utilized the 8008. And eventually it was replaced, but the initial design [was never upgraded to the 8080]

Hoff: That happened in spite of the idea that the 8080 was intended to be upward compatible from the 8008.

Feeney: Certainly upward compatible, but at the same time, it required a reinvestment and the customers want to get on with their business of solving their problems and not going back and reengineering things. And that's something that we didn't understand at the time.

House: It wasn't socket compatible?

Feeney: Correct.

Mazor: I'd like to reflect that some of the best years I enjoyed were at Intel working with these fellows. We were a small team and we had a lot of the correct ingredients. Engineering was able to meet with the customer and hear their requirements, interpret them, and move pretty quickly into engineering products. And working together; for example, myself in applications engineering with a chip designer in the same office. This wasn't going through different departments and through e-mails and so on, but, you know, [we had] a lot of close synergy. And so I think part of our success was that we were young and small and quite able to move [quickly] as needed.

Hoff: The other thing too is that all of us here were not specialists in one narrow area, but had a broad range of interests. And the idea that Federico had, his prior experience with computers and calculators made the transfer over just so great. You know, [with] a different team [it] might have been much more difficult to accomplish that.

Mazor: In particular, Ted and I, early were worried about the 4004 being too aggressive, and specified another chip which has never been heard of called the 4005. This was a very much simpler CPU than the 4004. And with our agreement with MIL [Microsystems International Limited], the company in Canada, we had arranged for Ken Au, a designer, to come down and be my officemate. The idea was that they [MIL] would develop that processor; we'd develop the ROMs and RAMs. Well, Ken was my officemate and I was transferring the [4005] architectural ideas to him and he was working on the design. I turned to him one day and said, "And how are you gonna test it?" And he said, "Test it? That's not my problem. That's somebody else's."

House: It's another department.

Mazor: Exactly, another department. And one of the things that Federico specifically did as he was designing, the 4004, was putting in some features to make the chip [more] testable. Because if you don't design a chip to be tested then, you know, possibly it can't be tested... and what good is it?

Faggin: Surely it can be tested.

Mazor: Thank you. So this reflects, in this case of the Canadian company, a larger company with more departments with a more narrow focus and, you know, it's much more problematic. We of course had success with the 4004, so we eventually abandoned our efforts and did not join in with MIL in the 4005.

Feeney: Well, as part of this too, we're talking about how all of us got started and were involved in making it happen, but there were so many other people involved in this product all the way along through the whole production side of Intel; the mask designers, the people in production engineering and whatever to make it happen. I was going through some notes this morning and I ran across the names of Rodney Sayer [ph?] and Julie Hendrix, the designers on the product, and Bob Holmstrom [ph?] who was the product engineer that actually got the thing through and got it into production. So there's a whole handful of people. And Hank pulled out a photograph earlier today of the early microprocessor group with about two dozen people in it, all the people that it took really to get this thing into the market and make it really successful.

House: Does anybody have anything else to add? Well, I think we can wrap it up then. Thank you, everybody.

END OF INTERVIEW