Oral History Panel on the Development and Promotion of the Intel 8048 Microcontroller

Participants:
Henry Blume Jr.
John Ekiss
Yung Feng
Barbara Kline
Howard Raphael
David Stamm

Moderated by:
David Laws

Recorded: July 30, 2008
Mountain View, California

CHM Reference number: X4969.2009

David Laws: My name is David Laws, I'm with the Computer History Museum here in Mountain View, California, and it's July 30th, 2008. And we're here today to record an oral history of some of the key team members who contributed to the development of the Intel MCS-48 microcontroller chip back in the mid to late 1970s. This is an important product because it was Intel's first microcontroller architecture, and its widespread acceptance by the marketplace inspired additional families from Intel and other competitors that established the concept of the fully integrated computer on a chip as one of the most ubiquitous semiconductor products in the world today. So I'd like to start out by asking each of our panel members to introduce themselves. Just if you can give us a brief background on your education, experience before working on this product, and how you came to join this particular product activity at Intel. We'll start on my left with John. Can you tell us how involved you were in this situation?

John Ekiss: Well good morning. My name is John Ekiss. I joined Intel in 1977, which happened to be about two to three years after this 8048 was developed and was initiated. But I came in when the 8048 was just starting into production and had contributed to some aspects of that. Prior to that I had gone to university at what was then called the Carnegie Institute of Technology, now called Carnegie Mellon University; and there's a sign down the road here I noticed saying that they have a presence here in the Valley. I worked at one of the semiconductor company pioneers, the Philco Corporation, starting in 1958. I was there about 12 years, and subsequent to their going out of business I wound up at Motorola in Phoenix, Arizona where I contributed, interestingly enough, to the development of some of their microprocessor, microcontroller products, and then on to Intel, and the story we have for you today.

Laws: Thank you John. Howard?

Howard Raphael: My name is Howard Raphael. I joined Intel in 1973 as product manager for the MCS-4, MCS-40-- that's the 4004 4040 family line, which was called the low end. I have a degree, a BSEE from Rochester Institute of Technology in engineering, and I came to Intel from Singer Business Machines, where I was an engineering manager, and wanted to make a career change to more of a business oriented career, and I was hired to be the product manager at that time; and we evolved the low end product line into the 8048 family.

Laws: Thanks Howard. Barbara.

Barbara Kline: And I'm Barbara Kline and I joined Intel in 1975. I had been a software programmer for two years prior to that, a software engineer, with companies in Seattle and then down in the Bay Area. And my degree is in sociology, I've got a Bachelor of Arts in Sociology from the University of Cincinnati. But it was early enough on that you didn't need a computer science degree. I put myself through school programming computers, and when I got out of school that's where I wound up. Like Howard, I wanted to move from the engineering side of the house into the marketing side of the house, and wound up in the development systems group. My background had given me a lot of good ideas about how to make it possible to develop programs quickly for the new microprocessors, and had introduced Intel's first in-circuit emulator, ICE-80. I had also done some work on Intel's groundbreaking high-level language PLM for the 8-bit processors. And at the time that the 8048 was coming to fruition we divided development systems responsibility, and Brian Halla stayed with the microprocessor line, which was then rolling out the 8085, and I moved and helped develop the products that would help support the 8048 family.

Laws: David.

David Stamm: My name is Dave Stamm. I graduated from Purdue in late 1973 and was part of the first group of college graduates that Intel ever recruited. I started at Intel in January of 1974. I worked on, for various other engineers, on fixing bugs on existing microprocessor chips, like the 4004 and the 4040. Then I developed a chip called the 4308, and then I began what became my dream project. I was single. I lived and breathed the 8048 for two-plus years, night and day, along with a team of another engineer and a draftsperson. So all three of us were-- it was our only responsibility. We were full-time working on this chip and had just a fantastic time.

Laws: Thank you Dave. Henry.

Henry Blume: My name is Henry Blume; I guess I was known as Hank more often. I have degrees from Yale and M.I.T. Came out to Silicon Valley on a lark-- or the Bay Area. In a couple of years I joined Fairchild, literally not knowing that a transistor had three leads, and I was assigned to process development, worked on product manufacturing. I worked in design, spent about eight years in design. So I was at Fairchild for nine years, and I had worked with Grove, Moore, Vadasz, many others, in a bipolar MOS logic memory. I went to National for awhile, Intersil four and a half years. I did a microprocessor at Intersil, and I was the first second source for a number of Intel chips while I was at Intersil, and they actually invited me in and I tested and verified my chips worked, which kind of surprised them because I could get chips out pretty quickly that would match what they were doing. I joined Intel in I think about October 1st, 1974. And they really didn't tell me what they were going to have me do. It turned out that Faggin was leaving and Ralph Ungerman was leaving, and they were just staffing to try to build up some supervisory experience at the time of those losses. And I was assigned to what was called the low end, which was cleaning up some things on the MCS-40. There was a single chip calculator and there was a general understanding that as EPROM technology evolved--it was going to become five volts-- that we should do a controller EPROM chip to exploit the technology, and Intel was hoping to put some logic on all of the EPROM technology. It turned out we couldn't get logic chips out fast enough to do it. But we worked very closely with the 2716 group, and in fact in Intel's timeline of products, the 8748 came out about six months before the 2716, in production, although it passed us in volume immediately. I guess that should cover that part of it.

Laws: Thank you Henry. Yung Feng, I understand you were the first of this group to join Intel. Can you fill us in on your background?

Yung Feng: I was. I joined Intel in 1971, and my first assignment was to work on the 4004, reporting to Federico Faggin. At the time when I joined Intel the 4004 and the family of products were just coming out of fab, and so I helped him to debug those chips and later characterize all those units and figure out how to test it. Prior to coming to Intel I was with a company called Philco Ford, like John Ekiss, although at the time I did not even know him. I have a degree of MSEE from Carnegie Mellon also. That was how I came to Intel. And after the MCS-4 was marketed by Intel, after Intel acquired the right to market the product, we needed to produce it in large volume. I transferred to manufacturing to manage the product engineering effort, and the main idea was to develop test capability. And one of the problems we had at the time was there was no commercially available and capable enough testers to handle these products. The only tester that we had was a very old PW-10. So I came up with the idea of designing our own testers, using the MCS-4 as the controller and to build a reference system to generate the test pattern on the fly. And all we needed to do is to route the data to the device under test, synchronize the two systems and compare the outputs. And we went on to build many testers like that, what we called T-boxes. And well my story is a little long, before I got involved with the 8048. So I might as well just spill it. So with the T-boxes we had a problem of maintaining these boxes, because each one was built separately by a different person, and it became a problem of being able to maintain those boxes in manufacturing. And then came Steve Bissett. Steve Bissett was hired as an 8080 engineer and he was assigned with the T-80, which was the box built for 8080. So he really spent a lot of time with T-80 and learned everything about our T-boxes, and he left the company to build Megatest, and his first product was called the Q-8000. It's a box that basically it's a streamlined T-box and using exactly the same concept but built around the 8080. So he came back, sold the Q-8000 easily to us. And so by 1976 or '77, when the 8048 came, the T-box was very important; I mean the Q-8000 was very important test tool for testing all of the products.

Laws: Okay, thank you. We'll get into more detail on the challenges you faced with the product later on in the talk this morning. So this was late 1973, '74, and you were starting to think about this product. Howard, could you give us a little background in terms of the state of the market for microprocessors and microcontrollers, and who some of the competitors were at that time?

Raphael: Okay. Well the best way to learn a market is to live in the market, and I was at that time trying to promote the 4004 successor, which is the 4040. The 4040 was viewed by Intel as an upgrade to the 4004. It took care of some of its obvious deficiencies, particularly interrupt structure and more instructions, and we were out in the marketplace trying to promote that product. Now there are three currents in the marketplace going on at the same time, at that time. One was a large portion of the engineering community was toggling from combinatorial logic design using digital, DTL and TTL elements, to microprocessor design. So we were educating the market to that aspect. The other thing is the microprocessor market itself was segmenting. So there was a community of people, still a minority of the engineering community, that were designing with microprocessors. But that was segmenting into what became the low end, which is the controller, the 4-bit, the 8-bit single chip processor, which the 8048 was; the 8-bit products and a more high performance product line, which was the bit-slice, and the bit-slice at that time was beginning to become available from AMD and eventually Intel and others. So that segmentation was taking place. And the third area was particularly in the low end, and there we were seeing competition from -- and some of it was vaporware at the time, but it was very much aggressive competition in terms of price and performance -- from other competitors, like Texas Instruments, which was promoting their TMS1000, which was a true controller in the sense that it had the CPU, the ROM, the RAM, the clock, all on one chip, and it was being offered at a very aggressive price, under $10.00. That was a big price point in those days, $10.00. Another competitor that was very strong at that point in time was Mostek and they had a product that they were promoting and then eventually sampled, which was the 3870. And that too was very aggressive and it was a single chip product. So in the context, Intel looked behind the competitive 8-ball, because it's very hard to sell a multiple chipset, which the 4040 was, against a single chip product like the TMS1000, and to some extent the vaporware product at that time, which was the 3870. So we were really beginning to lose designs, and that was the impetus to define a product, which became the 8048. And I did a Products Requirements Document, PRD, which essentially defined a one-chip product, which I felt we needed to be competitive in the marketplace at that time.

Laws: Where did the initiative for you going to define this document come from? Did it come from the field-- "We need this." Did it come from top management? Did it come from engineering?

Raphael: Well as I recall, Intel was, at that time, very focused on the 8-bit products and the memory world. The area that I was in, which was called the "Low End," the MCS-4, the MCS-40. Although the MCS-4 was the progenitor of the computer business, it was not viewed as the most competitive business segment in Intel obviously. The 8-bit [products] were very promising. In those days we were getting $360.00 for an 8080, and the low end just didn't have the sizzle and the pizzazz that the 8-bit stuff had. Plus a lot of very big customers were interested in the 8-bit stuff and we were still trying to convert the market into microprocessors for the 4-bit. So anyway, to make a long story short, I was living in the marketplace and understanding the marketplace, and I really felt that we should do something more aggressive to try to be competitive in the 4-bit [area] -- that was my business; although I didn't have profit and loss responsibility I took it upon myself to be my business. And we had the engineering people who were fixing the 4040, and we had a great crew, and I just did this products requirement document, presented it to management, and then management took that and began, as I recall, a taskforce. And I think Hank Blume was in charge of the taskforce. And the taskforce was then assigned to define at a microscopic level what the 8048 should be. It was really a great time, both in my career and in the industry, to be doing this, because Intel had such a pizzazz; people were so interested in anything that Intel did, because even then it was viewed as a technology leader. And Intel was very focused on microprocessors and memory that getting entrée into large customers was so easy. And frankly in those days there was not a lot of input from customers really. You had to sort of read the customers' requirements, and they were mostly I/O related, and interpret those, and that's what we did in terms of defining products.

Laws: So Henry, you were reporting to Vadasz at this time?

Blume: I was reporting to Jean Claude Cornet. They promoted him about a month after I arrived. But he pretty much left me alone. He was worried about the 8-bit and I was worried about the microcontroller area.

Laws: And so who asked you to head up the taskforce to define the architecture for this product, do you remember? How did the taskforce come about?

Blume: Even before I arrived in '74, there was a group under Ted Hoff, and Sam Schwartz was the guy who said this is what a low end controller should have. But it was fairly wishy-washy, there was nothing specific about it. So I think there was an understanding that Intel was going to do a low end product [and] that it was going to use the EPROM technology. And the EPROM technology had made some increments to make five volts out of it. They put in depletion mode transistors; they put in buried contacts for us. So we made most of our decisions pretty much by osmosis. I think we pretty much agreed on most decisions, to retain the Harvard architecture where the program addressability is separate from data addressability; they're different fields. We were going to have 8-bit transfers. We were going to use all the peripherals from the 8080 line. And I think really the one decision we had to make, aside from getting project approval, was whether it was to be 4-bits or 8-bits, because we came from the 4-bit world. We really couldn't decide whether it was 4-bits or 8-bits, so we set up a taskforce to make that decision. It turned out the taskforce got overruled by Ed Gelbach saying 4-bits didn't have enough sex appeal. So we voted to use 8-bits. And then we were pretty much left alone. We defined it ourselves, we kept on going. We weren't going to have compilers; we were just going to have assemblers, which were viewed as a reasonable systems problem. And I think our definition proceeded as the 8748 went on. So we did get project approval, but I think we were working on it before project approval. And I think there was a basic understanding that Intel was going to do something.

Laws: This was a big decision to make the first product in this line a combination of EPROM and logic on a chip. Was that a hotly debated decision?

Raphael: No, it was a very strategic decision that was made. It was based on business considerations and the nature of the microprocessor business. And it was perceived that, number one, the 8748 should be the first product because it put forth Intel's best technology forward. But the most important thing was that it allowed people to do designs on the fly. And although we had development systems, there was resistance in the marketplace to buying development systems because they were expensive, and EPROM sort of fit in with the rest of Intel's technology. We also felt that actually if we got that out first it would take the pressure off of masked versions of the product which were expensive to turn around and mask, and at that time Intel didn't have a great capacity for doing masked ROM. So we were very much interested in doing the 8748 first. And one of the things that was debated was -- that was a very aggressive architectural move, on an engineering level, and the taskforce sort of had to come to grips with the comfort of doing that. The other part of the 8048 that gets forgotten, that I think should be mentioned, is the fact that we defined really two versions of the product, one with an I/O structure that went out, like a normal microprocessor, and the other structure was an 8041, I believe it was called, which had the I/O structure coming in, which could be used as a peripheral. And the impetus for that was the 8080 family was taking off and there was a lot of need for specialized peripherals. So we figured if we could program a microcontroller to be a peripheral and make it look like a peripheral that would be a handsome addition to the 8080 line. And so the taskforce really came up with both of those products. They were all variants of the 8048 but with a different I/O structure. And I believe that they were both done concurrently, but I defer my memory to perhaps Dave or Henry on that.

Laws: So Dave at what point did you get presented with this challenge of designing this animal?

Stamm: Well I sort of lucked into the project, from my perspective. I was working on the 4040 with Howard, and I loved working with Howard and Hank. And so this project was being run by the same group that was doing that project, and I said, "Hey, I think I can take the lead on this. What do you think?" And they gave me a shot at it. So, but this notion of combining the EPROM and the logic on the first chip, it scared the hell out of me, but I would never admit it because it was such a sexy thing to be building what was the first programmable, reprogrammable single-chip processor. And I was also enamored by the idea -- I knew how long it would take to debug these kind of programs that people would write, that would run on the 8048, that we could charge lots and lots of dollars for the reprogrammable version and have sort of a bait and switch where they'd have a low price for the ROM version, but they'd never get there because they were continuing to tweak their program and they basically would say, "Well let's go ahead and go in production with a reprogrammable one, until we build our first 1000 units, and then maybe we can go to the ROM one." And I don't know to what degree that actually worked out where we sold a lot more of the programmable ones for production use, compared to the ROM mask version. But that was such an intriguing idea, that I was the last to admit this is going to be technically a very difficult project because not only do we have to build a new instruction set, a new architecture, logic design using a technology that was designed for EPROMs and incorporate that all on one chip, deal with multiple voltages -- because it couldn't be just 5 volt because of the programming voltages associated with EPROM-- this is going to be hard.

Laws: Interesting. Hank, you had a point.

Blume: Yes, as I recall one of the few ground rules we had in project approval was that we had to have at least 256 bytes of memory on the chip. And Howard and Dave and I talked about it a couple of minutes and we said we're going to have 1000 bytes or nothing; a token product wasn't going to do. So we're going to put 1000 bytes on, and no one kicked at that. They really left us alone to do what we had to do once we passed the 4-bit/8-bit crisis. And then we moved on. Then it was a matter of implementation. I think the peripheral interface, reversing the BUS, I think that was one of my ideas to try to get project approval and to get the 8080 people to think more kindly of us as something that would help them rather than diverting resources. I don't think it ever paid off quite enough but it was part of the approval, and something we did later. There were a couple of other engineers that took our basic artwork and changed it, and that produced some internal problems. But we got it done, we got it out.

Laws: So with this challenging project ahead of you Dave, what kind of tools did you have available?

Stamm: Not much. Well first of all I was on the hook to come up with a die-size estimate, and especially since I was in the camp that said it really had to be 8-bit, one of the counter-arguments to 8-bit is well what does that mean for the die-size? And then, as Hank mentioned, we wanted to put at least 1K of memory on the chip. So certainly, as Howard said, the 8086 family and the high-end processors were sexy, from a development perspective. But the challenges in the single chip world, the low end world, were as difficult, in my view almost more difficult, because we had to figure out how to take like ten pounds of potatoes and get it in a five pound bag and then charge a ten pound potato price for it.

Raphael: Five pounds of potatoes.

Stamm: Yes, five, sorry, five pounds, yes absolutely. So die-size was everything. And as far as the tools we had, we had a circuit simulator that ran on a mainframe computer, and that was it. We did all the logic design and logic checking basically by hand, and there were no tools to do the actual physical layout of the chip, that was all done on vellum with draftsman type skills. There were no design rule checking tools to ensure whether the various spacing between layers or components were proper. There wasn't much. What's interesting on the circuit simulation side is that -- and Intel at that time had one computer system and it ran the Finance Department's month-end close as well. And so when they were trying to close their books they basically said, "Sorry, you guys in engineering, you can't run your circuit simulation programs because we've got to close our financials and report to the street." Those days are long gone! And ultimately that was part of the impetus on why I decided to start one of the early CAD companies after I left Intel because I really felt that engineering teams needed their own dedicated individual workstations, and part of that came from my experience working on the 8048 where I got shut down because it was time to close the books.

Raphael: It should be mentioned that at that time there were two buildings to Intel. One was the Bowers Avenue building and the other was the Livermore fab, as I recall. So actually it was a pretty small company in those days.

Laws: And they'd moved out of the Middlefield Road facility in Mountain View by then, I presume.

Raphael: Yes, this was Bowers Avenue.

Kline: We were in the basement.

Raphael: Yes, and everybody was crammed into Bowers Avenue.

Stamm: My desk was actually in the hall.

Raphael: Yes, so Intel was kind of bursting at the seams a bit in those days.

Laws: That was going on for a long time I understand.

Blume: I had another thought, to back up what Dave Stamm was saying. For the circuit design he had come from the P-channel background. There was P-channel circuitry. Some of it was used on a different technology for the 8080, but basically there were no design libraries to work with. It was all on a blank sheet of paper.

Stamm: Yes.

Blume: And if you do everything on a blank sheet of paper you're apt to make a mistake.

Laws: Sure.

Stamm: I don't know if I told you what I did then. I took Shima, who developed the early microprocessors -- his English wasn't good, he certainly couldn't have taught a training class to explain to people how he had done it. I took all of his designs -- there was no full documentation pack in those days -- I just took them home and I went over every single circuit, every design, and just poured over all of his schematics, and then I would go back to him and ask him a few questions. So my learning process was just studying other people's designs. That was the only way there was to do it at the time, and Intel was on the leading edge of it, so where else could I go? That was the only place.

Laws: And Shima's work was all logic. How mature was the EPROM technology by that point in time?

Stamm: It was beginning to be fairly well developed, but 5-volt EPROM was brand new.

Laws: And what kind of design tolerances were you using?

Stamm: In terms of spacings?

Laws: Yes spacings.

Stamm: Boy I don't recall. Do you?

Blume: 10 micron.

Laws: 10 micron.

Raphael: Yes, big chips.

Stamm: Yes.

Blume: I think it was more like five microns.

Laws: Okay.

Blume: But it was in multiple microns.

Raphael: It was changing in those days. It was getting smaller.

Laws: And wafer sizes were two inches maybe.

Stamm: Three.

Raphael: Three, three inch wafers. And Dave alluded to another impetus for the 8048 and that was the single voltage chip. As I recall that was also a big market demand. People just wanted one voltage supply versus multiple voltage supplies, which the P-channel stuff I believe had, if I recall.

Blume: There was an N-channel EPROM, a 2708. It was 8K bits. I think it worked with a +5, +12 volt power supply, and we were very uncertain whether it be self-erasing or not. We tried to life test them on the roofs of the buildings, and doing a bunch of other things. So we were working on getting a denser technology, a 5 volt only technology, and buried contacts and depletion mode transistors so it would work with 5 volts, at about the same time. So we worked very closely with the group. Our mask designer laid out the first test chips for the process. So we had a very intimate working relationship with the process people at the time. It was a beautiful working relationship.

Laws: Were you all crammed into that one building?

Stamm: Yes.

Laws: They were just around the corner or right next door. It makes it a lot easier.

Stamm: Yes. It George Perlegos who ultimately started Atmel who was one of the process people, and we were meeting with them regularly.

Raphael: I just want to mention that Hank did a great job in organizing that thing. During the development we had regular progress meetings, and where necessary marketing participated. The other thing that was very important during that development was the MCS, the microcomputer systems test equipment that was being developed, the development system stuff that was developed for the product, Barbara worked on, and we worked with Barbara to do that, and that was a very important part of the whole development as well.

Kline: Yes, I think it's important to put that piece in its perspective, because as you heard Howard say earlier, there was kind of a dynamic tension between the development systems group and the chip group. I came into this from a programmer's perspective and was very well aware of some of the challenges of debugging, both in design and in the field, of a microprocessor based system. And the development system strategies, the things that we worked on there, were just so exciting because you were taking from a place where everything was a hard design and you'd put wires together, to a place where there was this logic swirling around inside that you could only kind of track. To create tools that made it easier to do that-- and I mean a lot easier to do that-- was really exciting. I think one of the challenges of the development systems products, back in those days, was that it took a different mindset to design with those kind of products than it did to do a classical engineering design, and that was one of the challenges that I worked on with other people in the development systems group is helping people to see just how powerful these tools were and what a difference they could make.

Laws: At what point in the development of the architecture and the work that Dave was doing in design did you start to get involved, and did you have to affect the design in any way in order to build the right kind of tools?

Kline: For the 8048 we didn't. We were pretty new in terms of in-circuit emulation, which I think is one of the critical tools. We had to develop a PROM programmer and that's an important piece and it's pretty well understood. We had to develop assembly language and that was pretty well understood. In-circuit emulation was something new that had come out with the 8080, and we were working in tandem on the ICE85 and on the ICE48. And I think that, as Howard will say, the in-circuit emulator wasn't on the critical path to getting the chip out. So it came out in kind of a secondary role. And as we learned with ICE48, trying to get inside and see what the code was doing and emulate it took maybe more effort than we had thought. And we did really rework the process when we came to the 8051 in that not only did we develop a bond-out chip, which made it easier to develop an in-circuit emulator, but that was also one of the first chips, one of the first designs that came out of that whole process, so that we got the in-circuit emulators in there earlier for people to do their designs.

Laws: Okay, interesting.

Blume: So we talked about the lack of CAD tools. Dave developed what was an almost a CAD tool, but we did have a TTL functional equivalent breadboard.

Stamm: Oh yes.

Blume: And as we were debugging the chip he would find ways of breaking metal lines on the chip, hooking it up to the breadboard, part chip, part breadboard, to see what parts worked, what parts didn't work. It was a heroic effort, but that was the state of CAD at the time. I don't think we ever worked with a systems group in this break it out thing, or I don't think they ever used it. But it helped us in getting to functionally good parts.

Stamm: I'd forgotten about the TTL breadboard that was a major part of the project.

Laws: There were a lot of TTL breadboards built in the back in those days.

Stamm: Yes definitely.

Raphael: I just want to say that because the chip had the ability to break out the memory from the CPU, it made it very amenable to the in-circuit emulation, and of course that was done not only, largely for testing purposes; it was to insert the emulation. But that was part of it.

Laws: Yung, at what point did you get invited to test this beast? Were you involved at the design stage at all, to give inputs like -- if you do this it would be easier to test? Or were you handed a finished product and said okay.

Feng: We did have some input to that. We insisted on testability. We probably insisted on the ability to access externally the ROM and the CPU separately. Also I may add, when you ask the importance of EPROM, I remember the EPROM was designed by Dov Froman, the very first one, about the time I joined Intel in 1971. And when we marketed the MCS-4, it turned out the MCS-4 uses a ROM, it was ROM-based, and very quickly we found out that EPROM was important. So I actually designed the interface chips-- the 4008 and 4009 to facilitate interfacing the 4004 with EPROMs and external memory. And I think that had something to do with helping to sell the MCS-4. And probably that was how we learned the importance of EPROM to work with the CPU, and that may have something to do with defining the 8048, 8748.

Laws: You talked earlier Dave about this being two years of your life. Can you give me some ideas about what you did during those two years? What were the challenges, who were you working with, and how did the phases of design work out for you?

Stamm: There were three of us who were full-time on the project: myself, I was in charge of the instruction set and the logic design and the overall chip schedule; David Buddy, who was in charge of the complex tasks associated with integrating the EPROM technology and all of the EPROM programming logic and the sensing logic for the EPROM components, as well as most of the circuit design; and then Dwayne Hook who was responsible for all of the layout of the chip. And from my perspective the unsung heroes in those days were really the layout designers because it was such a difficult, tedious process to manually lay out every single component, every little contact, every little nook and cranny of a logic design that in our head was in terms of a circuit diagram, to actually put that into a physical, multi-layer design was really tough, and Dwayne Hook was excellent at that. So the three of us were working night and day on this project. The first phase was really the instruction set design. So here I am, I'm out of college a year, had a Bachelor's degree, thinking to myself what business do I have developing the next generation instruction set for these components? And really nobody was looking that carefully over my shoulder, but I said, okay, fine I'll go for it. Luckily in college I had taken assembly language programming in a language called COMPASS, which was part of CDC, which is no more. But I learned a lot about assembly language programming there, and then I'd worked on the 4040 and I looked at all the other instruction sets in order to try to come up with what kind of instruction set might make sense. The challenge was [that] you had to consider the instruction set in terms of how much complexity in die-cost it would be. So, for example, subtractions and compares would be really valuable instructions, but they added dramatically in additional chip area. So I decided we really couldn't afford those, especially since we were going to go into an 8-bit design. I had been a big proponent of 8-bit. One of the counter arguments was cost and die-size. So during these design steps I did everything I could to jettison features and functions that I thought was going to add extra die-size. Looking back at it now I think I probably went a bit overboard -- although at the time it would've been difficult to have made that conclusion -- because I eliminated the ability to -- since we had no subtraction unit, if you did jumps you had to jump to an exact address; you couldn't jump on a relative address. Memory addressability, we only went to 4K, and to go to 4K you actually had to manually set a bit to go -- I want this 2K bank, I want this separate bank; you had to actually use a full instruction to set a bit to go to those addresses. So basically it was like a bank structure. So there were lots and lots of limitations that existed, primarily in order to keep the die-size down. So there was probably about a three month period where I was focused on these kind of instruction set and design issues, and Dave Buddy was focused on working with the processing group, the technology group, in order to figure out how to incorporate this 5 volt EPROM process, which was coming out for the first time in the 8748, on the chip with the logic design.

Laws: As you were making these decisions to eject some of the baggage, who was looking over your shoulder? Henry were you seeing this going on, or Howard?

Blume: No, I was. I was probably looking over Dave's shoulder more than he thought. But I had about 10 or 11 benchmarks. Some of it came from the MCS-40 manual. Little things: traffic light controller, A/D converters, things of that sort, and I would run by Dave's instruction sets and see how much ROM space did it take and how much time did it take? Although I don't think time was that important. And we thought that we were perhaps 40% more code efficient than the 4040.

Stamm: That's right.

Blume: We could get it done in about 40% less code than the 4040. But we also knew that the code being used in 4040s was going up about 40% a year. So one of the blockages we had was a basic 2K addressability at a time when we thought the average code used was going to be 1½K, and it was going up 40% a year. So we kind of knew we were in a little bit of trouble on the memory addressability. But I think everything used a lot less ROM code, and then the instructions he took off really didn't cost much.

Stamm: Yes, I'd forgotten about that. In fact I think we ended up, there were about 90 instructions, and 90% of them were just a single byte. So we really worked very hard to -- but I'd forgotten about those benchmarks.

Blume: I also want to point out Dave's two famous - or his favorite instructions were SEX and SIN, which stood for set external mode and set internal mode. And then when the people in systems took over they removed SEX and SIN.

Laws: They just renamed them or dropped them altogether?

Blume: No, no, they renamed them.

Laws: Howard, were you aware some of the baggage was being thrown overboard here, and was it of concern to you?

Raphael: Yes, at some point in time we were. We had certain customers we were showing the preliminary information on this product to, to see if they liked it. I recall showing the instruction set to some customers. And as I say there was much more interest in the price performance points and the technology, particularly at the engineering management level. Mostly I interfaced at the engineering management level in corporations. So there wasn't a lot of technical input really, in the controller area. They were much more interested in the performance and the price of the product. And there was no question about it. The definition of the 8048 exceeded marketing's requirements because it had everything - it had a lot of things that the 4040 didn't have. And as I say again, it was very well orchestrated. That whole development effort, from a marketing point of view, was very well orchestrated. As I recall we had sampleable product. It may not have been the full EPROM product. I think after second silicon we had parts, you guys really did a great job in terms of turning that product around, getting it out. Sometimes in previous developments I was involved in you'd have to wait for three or four rounds of silicon, but they got it out pretty good.

Laws: So continuing with the design process Dave. So you've defined the instruction set, now what was the next challenge?

Stamm: Well one last point on the instruction set. So I think it's kind of interesting, we were building one of the world's smallest computers but I learned, well I learned from one of the world's largest computers, the CDC, in order to build the right instruction set for that chip. So that was fun. Well the next phase really began to be the TTL breadboarding. So we built this large breadboard, which in and of itself was yet a whole other design project, using the same type of combinatorial logic that we were hoping to eliminate through the development of the 8048. So we were using TTL and DTL devices. And the breadboard was huge as I recall. It was probably five foot tall by two or three feet wide, just completely covered with wiring on the back. And we also did test chips on the EPROM development and began doing the physical layout of the chip. So that phase probably lasted about six months or so.

Laws: And the layouts then transferred onto Mylar in those days, and these were hand stripped Mylars or--

Stamm: I think it still was Rubylith, wasn't it?

Blume: Rubylith. We drew on Mylar and then we cut-- Rubyliths were laid over.

Stamm: That's right.

Blume: So it was drawn on Mylar, which tolerated erasing pretty well. It was in pencil on Mylar with a lot of erasing and redraw.

Stamm: And what I worked hard to get was the best Rubylith peeler, which actually made a huge difference on your project, because once the Rubylith was cut, which was a red film on top of a plastic backing, then someone would have to peel out all the portions that were supposed to be removed, and if they didn't do that very carefully then they could introduce errors or the problems would be difficult to discover. So you really wanted the best person on your project, all the way.

Laws: It's a lost art, I imagine, a Rubylith peeler.

Stamm: Yes, so Rubylith peeling was extremely important.

Laws: Do you remember who that was?

Stamm: I can see her, I can't recall her name. I can still visualize her but I don't recall her name.

Laws: Sure, an important team player. And so now you've finished the layout, you've created the Rubyliths. You go to mask making. Was mask making in-house in those days?

Stamm: I don't believe so. I think we used a…

Raphael: No I don't think Intel was doing that.

Stamm: …third party to do the mask development and production.

Laws: And so the plates go to fab. Were they ready for you? Were there any process tweaks required along the way?

Stamm: I don't believe so. The frustrating thing for me on the project was the length of time it would take. I subsequently went into software development. In software development you can change a line of code and within seconds now actually you can retest it. But when you're doing this kind of chip design, especially in those days, you had relatively little tools to check out your design ahead of time and then you had to wait multiple months before you could determine whether what you had taped out -- it was called taped out when you sent out the tapes for manufacturing -- actually matched. That frustration was compounded because in the same lab where I worked was a group of engineers working on Intel's DRAM memory chips, and DRAM had so much priority at Intel, and there was so much pressure on their project, that if they wanted to make a change, they would try three design changes simultaneously, and they'd send out three copies of tapes simultaneously, and they would have their chips back well before mine. And then they'd send out another three and another three. So I felt that they were just kind of churning in terms of how much production work they were generating for the fab, while my project had to sit and wait for them. So that was something that was certainly frustrating in those days. But you would wait-- and I don't recall how long a wait it was; was it 12 weeks or--

Blume: 15 weeks was standard.

Raphael: Yes, that's what I recall. Red tags were like eight.

Blume: Something that happened, as we did the first test chip for the EPROM, we learned some things. We learned we had to change the design rules, particularly dealing with high voltage, and we saved the 2716 iteration because we had done this learning. So I think they were ready to give us a red tag and run through pretty quickly our first chip. Then we got it and we had to hook the chip up, half to the breadboard, half to the chip itself, to see what worked, what didn't work, where the errors were. So I don't think fab limited us as much as it might have on some other chips because it was a new process and a new process group and they needed something more meaningful to work with.

Laws: So when did you see first silicon Dave, do you remember?

Stamm: The month and year-- no I don't recall.

Blume: The date I key on is that we introduced the 8748 in the fourth quarter of '76, and we had a key result to ship 1000 revenue units in the first quarter of '77. I think we actually got 770, which I didn't consider a bad miss; I don't think anybody did.

Stamm: So it was fall of '76 then probably.

Blume: But I think we were working with silicon probably from the second quarter of '76, to the fourth quarter.

Stamm: Yes, I think you're right, I think it was mid '76 when we had our first components back.

Laws: Sure. And then you tested it by comparing it to the TTL emulator?

Stamm: Yes, and by writing actual programs and determining whether those were working properly. Plus we had to do quite a bit of testing that Dave Buddy was concentrated on. And we were a great team. Dwayne Hook, Dave and I, we just had so much fun working with each other. But he had to do all of the reliability testing on the EPROM programming and do all of the sensing work, and we had to do temperature tests and things like that, on the EPROM technology.

Raphael: As I recall, on the critical path to getting this product, in particular the low cost version, was packaging technology. There was packaging technology required for the EPROM version because of the cavity size of the die; I think the die was about a half-inch square. And also remember that marketing was very interested in having that under $10.00 chip solution, which required the plastic version, and I believe there were some packaging requirements involved in that. Yung Feng might recall more about that than I do. But that was a very critical part of this project as well.

Laws: What size package was it Yung Feng? Do you remember how many leads?

Feng: I don't really remember. I think 40, 40 pins.

Laws: 40 pins.

Stamm: Yes, 40 pins.

Blume: It was a 40 lead package. The die-size was 248 equivalent. It was rectangular but it was the effective area of 248 mls. And the first package had a plastic cover so that you could shine UV light through it and erase it.

Laws: Quartz.

Blume: Quartz, okay, it was something that you had this-- I don't think this was totally hermetic either. There were some issues.

Raphael: It was a CERDIP package, as I recall.

Blume: Yes, 40-lead CERDIP. What was it, 100 mil spacings on the leads? On the EPROM technology you programmed every bit, you baked it, you tested it so every bit was still there. You erased every bit. You tested to see if you erased every bit, then you programmed some and erased some more. There was a whole process flow to go through to get a pretty reliable product out. And I know we had to do it, I know we were involved with it, and then it took care of itself. Well I guess product engineering put some time and effort trying to smooth the flow a little bit.

Laws: So at what point did you start getting involved in testing the devices Yung Feng? The very early testing would've been done by the engineers.

Feng: Well the design engineers usually they do all the debugging, and when the product is considered production worthy they try to transfer the product to the product engineers. That's when the product engineers take over. And product engineers are responsible for developing the production test capabilities. And for the products like the 8748, and the 8048 later, we basically used the Q8000 that I mentioned earlier. And for the 8748 it was quite an elaborate test because we had to test the EPROM. We had to program it and erase it and then program it again, to make sure that everything can be programmed.

Laws: Okay, so now we're at the point where you ship the first products in late '76, early '77. A lot of other things must've been prepared to go along with this shipment. There must've been marketing materials; there must've been customer tools to use the product. So what did marketing do in preparation for this?

Stamm: If I can insert there, there was one additional milestone, that was probably the most important one to me before we went to that stage, and that was when I had to demonstrate the working chip to Les Vadasz, Andy Grove, Gord Moore. So Dave Buddy and I came up with the idea. I was actually counting cards at blackjack in those days. I loved going to Lake Tahoe and gambling. And so we decided that we were going to try to write a blackjack playing program, all within the 8048. And all we wanted to demonstrate it was two single chips. So we wanted just the chip connected to a dumb terminal, which is a black and white terminal with its own integrated keyboard that could put up little strange little characters so that you could actually form playing cards. And we built this, we were able to do it, we could just fit it in. And I realized some of the limitations of my own instruction set, unfortunately, when I was writing this program to try to do this blackjack and finally we got it working. And then we had this huge power supply, so this one little silicon chip, a huge power supply, at this dumb terminal -- what could we do about this power supply? So I think Dave came up with the idea of why don't we make a homemade battery? So we took like a lemon or an orange, I don't remember, with two dissimilar metals, and we stuck it in there. We hooked that up to the power supply input and then we had the single chip device connected to the dumb terminal and you could sit there and say hit, stand, double down -- actually I don't think we did double down because that would've meant more complexity and I think I was out of programming space. But that was a big success and that was kind of the highlight.

Laws: So it was appreciated by the audience that you built it for.

Stamm: Yes, yes, that was a big success.

Laws: Great. Do you remember when that was? That would be late '76.

Stamm: I think that was mid-'76, I think it was mid-'76 when we did that.

Raphael: You're doing great.

Stamm: Thank you.

Blume: I think Dave Stamm left the project in maybe November '76 to move on to the 186. So I think it was before that time. Dave you mentioned, and I think Howard should take over from me, we had to do an assembler, and I don't remember if we did it or we contracted it out. We also wrote a user's manual. I remember writing several chapters of the user's manual, before somebody rewrote it and make it better English. And I was saying, how does this instruction work? And particular in the page boundaries, what happens if you exceed the page boundary?

Stamm: Yes.

Blume: So we'd argue a little bit about how the instruction worked, and he did the breadboard to go and see this, what really happened.

Stamm: Yes, yes, there were-- especially the page-- some of the limitations we put in were troublesome when it came to actually doing the programming, and I'm sure that produced problems within circuit emulators and others. The other challenge we had was this 40-pin package. That's all we had, and so we started having to double up some of the pins, because we had to have an external, what was called external address mode, so that you had to have a way to test the CPU separately from the memory. And we actually took certain pins, and I don't think this had been done before, where they would have multiple states; 0 would mean one thing, +5 would mean another, +12 would mean a third thing and +18 would mean a fourth thing, because there just wasn't enough pins to do everything that we needed to do. So that was another sort of interesting innovation. So I'm sorry so.

Laws: No, that's fine, thank you for addressing that.

Stamm: That was sort of a milestone that then led to the okay, “Now let's figure out how to turn this thing into production.”

Raphael: Well we had a marketing introduction plan, which was very comprehensive. It involved public relations and it involved advertising, we had advertising done. In those days we were using Regis McKenna. And we did a variety of promotional products. One was a brochure, a very slick brochure, on the 8048. In those days computer product lines had these great slick brochures, we did one of those.

We did a user's manual, data sheet, and of course we did schedules for tradeshows. And the biggest thing was press tour and a tour to the sales force to introduce the product, which was done usually at a sales meeting. They had regional sales meetings and annual sales meetings. So there was a very comprehensive effort to introduce the product. And I just want to say that [in] marketing, we love to pre-sell things, and so there was a lot of pre-selling of this product prior to its coming out, because we were feeling competitive pressure even going into it. So you know how disciplined marketing people are. Anyway, so we pre-sold it. So the critical path, the thing that I remember worrying the most about is getting first silicon. And that's why I say it was such a wonderful thing to have it come out; first silicon and then they had a second run and a second iteration -- as I recall it was sample-able. And I do recall the 15 or 16 week turnaround time for getting things through fab. And that was another reason why the 8748 was such a key product because customers were not willing to wait 16 weeks to get a ROM iteration on their product. So it was a very well demanded product, from that point of view.

Laws: What was the response from the press, and then from customers?

Raphael: Well as I say in those days the trade press was totally key to anything that Intel did, and particularly in the microprocessor area. Because the press more than anybody understood the change in design trends in the United States going from combinatorial, parallel, non-programmed state machines to microcomputers. So anything with microprocessors in it was big news. And I wrote, over the year's time that I was doing this, I wrote like 20 or 30 articles on microprocessor design, some form of microprocessor design. I must say some of those articles were variations of the others, but it was purely a PR effort to get get microprocessors in the magazines. And in those days we thought that an article was equivalent to an ad. In those days a full page advertisement in trade magazines was $5000.00. So if you can get an article in a magazine that was equivalent, in terms of public relations and promotion, to an ad. So it was a very valuable thing. And of course I did it for the money because I was getting paid by Intel to write the articles, they'd give you a bonus for writing an article, plus the magazines in those days paid a stipend. So I made a lot of money that year, writing articles. I'm not ashamed of it.

Laws: Were sales forces as excited as the press, or did you have some pioneering to do to show them how to sell it?

Raphael: Well the nice thing about Intel, as I learned when I moved on to other companies, is that Intel is very focused. So the sales force only had essentially two things to sell, memories and microprocessors. So they were very receptive and open minded to anything, and the sales force was probably pretty technical too, as sales forces go. We had application engineers as well as sales people. So there was no problem in terms of reception and they were very excited about it because the market was demanding that product.

Laws: And do you remember a few customers early on that made a difference in its acceptance?

Raphael: Well I recall we did a lot with companies, like gas pump companies like Gilbarco; I remember Tektronix. Gee I remember some of the automotive people were very interested in this, particularly Chrysler was very open minded to new technology, perhaps more so than the other two domestic suppliers. I remember there was a lot of interest in this product in Europe, a lot of interest in this product in Japan. Europe, there was a very fertile sales effort going on. Tom Lawrence was in charge of Europe in those days and he was really good. And in Japan there was a lot of interest in it, but at the end of the day there wasn't a lot of sales in Japan. But it was certainly fun for me to travel in Japan in those days.

Laws: Henry, I understand you were quite involved in helping sell the product or promote the product.

Blume: Oh I wouldn't say that much. Tom Lawrence arranged for me to spend a couple of weeks in Europe traveling around in the second quarter of '77, talking at a number of conferences about this new kind of thing, but also talking to the sales force and having a guy from the factory update them on what's going on. They always felt the factory guys had more credibility. I think the FAEs were extremely key in selling the product, once they learned how to design it, because they did a lot of customer designs, or at least blocked it out for the customer so the customer could do the detailing, but they would set up the overall concept. So I think the FAEs were great in that. In what timeframe some of this happened, I really shouldn't say. I key more on things that happened in '77, which I think a lot of the introductory stuff was in the fourth quarter of '76.

Raphael: Yes we, Intel had distributors too in those days and we did a lot of distributor seminars, and I gather, from those seminars, there was a lot of interest generated, design interest. Because the beauty of the product was you could get the product and because it was EPROM capability you could immediately start producing a prototype and that really was probably the 8048's family's best feature. And so anybody who had one of these things, an EPROM programmer, could essentially start doing a product; and people did.

Stamm: And they could speed up their own product introduction scheme.

Raphael: Time to market. And the other thing, when we consciously decided to do the 8748, we knew that designs have a gestation period. In other words from the time somebody gets interested in your component, until they design it in and they go through a pre-production process, that's a several month process at best, and in large corporations it could be years in those days. So we had plenty of time to design the ROM version if we had the 8748, and of course the development tools. But as I said the market broke down in two places: people who had the money to buy the development tools, which were pretty expensive in those days, relatively speaking, and people who didn't want to spend the money on the development tools because that was too much of a commitment. I never could understand that because here's somebody committing to a new project and they won't spend three or four thousand dollars for a development system. But anyway that's the way people thought, and so we had a lower cost solution for them, essentially what was a PROM programmer; I think we called it an evaluation board or something. But it solved the problem of getting 8748 designs going.

Laws: Barbara, what was the complement of tools that you developed to support the product, and which do you think was most important in terms of making it successful in the marketplace?

Kline: Well I think all of the tools were important. You had to have an assembler, you had to have the PROM programmer, which as I recall was specific to the 8748, because we had done some other ones as well. Being an old programmer I really felt that in-circuit emulation was an important piece of what was going on, and it was one of my campaigns the whole time I was at Intel to get people to understand the power of that system. I think I had mentioned earlier, it's a different way of designing for the engineer. Somebody that's coming in from the programming side, if you actually had a separate person doing the programming, could pretty quickly see the value of an in-circuit emulator. But if you were coming in from the other side it was a little bit more difficult, even though debugging [using ICE] was incredible. I remember going out -- this was on the 8080 which was the first in-circuit emulator -- and finding problems that had stymied an engineering group for months in about a half an hour. I think part of that was frustrating [to design groups] because it was almost too easy, I think it [ease-of-use] had some of its own challenges. But I really felt in-circuit emulation was a very important part of what Intel brought to the market in terms of development system tools. And it was interesting sitting on the other side of the fence from Howard because all of the publications loved the articles he was writing, and if you tried to do something on the development systems and why they were important-- I had somebody tell me that until we got a second source they weren't going to cover them, and it was kind of--

Laws: Electronic Design magazine I think that was.

Kline: Yes exactly, and I'm sitting there going, okay, we have more education to do here on what a development system is. But I think it was really insightful of Intel to move in that direction. Because I was designing with the 8008, before I came to Intel. I had developed some of my own tools to figure out how to do it. But it was a totally different level from what you could do when you actually had a group that was focused on creating design tools. I was really excited to be a part of that program and I loved the opportunity to work with the microcontroller group, because it was in many ways a different problem from the microprocessor group.

Laws: What would a reasonable complement of tools cost a customer in those days, do you remember, say a starter kit?

Kline: I'm thinking -- if I recall correctly the development system itself was like about five or six thousand.

Raphael: Yes.

Kline: And then I think by the time you got ICE and a PROM program, which was pretty much a standalone system -- I think the assemblers mostly kind of were free. So it would be like eight, nine-thousand dollars. It was always an interesting thing in the sales aspect, going back to -- the opportunity and the challenge of being the development systems part of this program is the chips had sex; there wasn't a whole lot of sex in a development system. And getting people to understand, getting our sales guys to understand what an important component of the sale that was, because it was great to go in and sell all the chips, and everybody would get excited about the chips, but this piece of it that was so important to getting them out quickly and effectively was really viewed as kind of an unnecessary expense. And believe me, coming from the software side, I knew absolutely that it was not, because I could've chewed up that amount of time doing a software design in one single project.

Raphael: Yes, we used to tell the sales force to procure - to mature a design win, if you sold a development system you were in. And the sales force was -- it was easier to sell development systems on the high end products than the low end products because people, for some reason, were willing to spend money on the high end product developments. They were usually people like NCR or HP who had a vision in terms of the worth of the investment in the development systems. The low end people were not quite as amenable to that. And one of the things that I did with all the product is in those days there were standard PROM programming manufacturers; in other words they were people who had PROM programmers which had personality cards. And I don't know if I was supposed to do this or not -- but I recall trying to encourage one of them to develop a PROM programming personality card for the 8748.

Laws: Data I/O probably.

Raphael: Yes, it somewhat conflicted with the development system business.

Kline: Yes somewhat.

Raphael: But I figured the more the merrier.

Laws: Sure. And a lot of these low end customers, as you call them, were not terribly sophisticated in terms of designing computers. They were designing small control systems with TTL. So these were hardware people.

Raphael: They were hardware people.

Kline: They were.

Laws: You had to train them in the understanding of how to use these tools.

Raphael: Almost everybody I spoke to in those days were doing their first microprocessor design and were going away from combinatorial logic. I had a presentation I used to do which discussed the difference between combinatorial logic design and microprocessor design, and that was very popular in those days. As I said, the number of people who were designing with microprocessors in the engineering community in those days was very small, but fortunately for our industry they were very influential. And the microprocessor trend had caught the eye of the engineering management community, and that's really what it needed to toggle this thing. So it was very important from that point of view.

Blume: I think a large number of customers bought from distributors -- we didn't know who they were -- and this was let's say first half of '77. Six to nine months later they started ordering. So we got a whole bunch of orders that we didn't know were there. We also had made parts that the EPROM didn't work but the CPU worked. So we called it 8035. We sold it as CPU only. And on the street they got the word that you could buy an 8035; if it's in the right package you could actually program it. So there was a lot of street smarts out there on using this product.

Stamm: That's right.

Blume: They were programming some of them. The bits weren't guaranteed by us but they were good enough for their prototypes.

Stamm: That's right, I remember that.

Blume: So there was a whole underground of people using this that we really didn't know that much about.

Laws: Barbara, I believe you traveled in Europe fairly extensively, didn't you, showing people how to use these tools and training.

Kline: Right. Yes, it was interesting. I went over with Henry in early '77 and we spent -- well I was over there I think for about five weeks. Tom Lawrence was really excited about having headquarters people come over and work, and we spent a lot of time explaining what was going on in the development systems group for me and some of the other products. And actually a decent chunk of the microcontroller group wound up in Europe, because Ben Franklin, who had worked with Howard, went over there sometime early in '77, and I followed him later on in that year and was part of a program that Intel put together where they were rotating headquarters people through Europe in two and three year cycles, with the idea of transferring more current information. The idea was that if you had worked at headquarters you'd be able to get what you needed faster, because you were a known quantity, than if you were somebody brand new. Europe was a wonderful experience, I enjoyed it tremendously and I think I learned a lot from getting that different perspective.

Laws: Okay, we're now into 1977. Howard, you've moved on to another company by now; I think you went on to National. And John, you recently joined Intel.

Ekiss: Right.

Laws: Can you give me a feeling as to how the product was perceived in the company when you joined, and any involvement you had.

Ekiss: Okay. I joined in the fourth quarter of 1977. Remember, I had had a background at Motorola running a division that had been developing, not a microcontroller but a microprocessor that would compete with à la the 8080, and the 8085 actually I think more specifically. And I came in as the manufacturing manager reporting to Les Vadasz, who was in charge of the microcomputer complement division. And it was a new organization. In fact my introduction to the company was to come to work the first day, all bright and shiny and expecting to talk to my new boss, and he wasn't around, but Henry came out to greet me and spent some time showing me around the company and introducing me to my direct reports. Anyway, I was really excited because I had seen ads on the 8048, before I even when to Intel, and when I was going to be involved in it I was really excited because this clearly was a whole new market area, product area. And it didn't take long until I got very deeply involved in it, because it was getting to the point where you were going to have to actually demonstrate shipping pretty large volumes and getting the costs in line and worry about testing it in volume. So I proceeded to get involved from that point of view. I initiated, at Vadasz's request of course, a cost reduction program; developed a model, working with Yung Feng, to what we had to do to get to the cost structure that we needed, which we had initially targeted at somewhere in the low $2.00 range, for the manufacturing cost, which meant we would come in pretty much under the $10.00 goal that had been set. And but of course we had to execute, of course, against that plan, and we did, and I'll let Yung Feng talk about that. Also I got involved in the cycle time issue, the time to actually manufacture a part. If we were going to go into the business of supplying products that were hardwired -- namely the read only memory in the memory is programmed with the mask -- we would have to do a better job of number one the absolute time to manufacture the part and also the reproducibility of that time. Because when you told somebody, who was now going to go into production, that he was going to start getting his production parts in ten weeks, or whatever date you gave him, he was ready, and you better be ready to ship, because if you didn't, or if the parts were defective or the yield wasn't right or you had to make a mask change, you were in for another long cycle time. And I did early on have to spend a fair amount of my time working on those kinds of issues. So we did put in a program to do that. So it was a really interesting time and actually it got more interesting as time went on. But I'll turn it off at this point, and I think we should ask Yung Feng to talk about some of the details of things he got involved in.

Laws: Sure. So at this point you were trying to build larger volumes, Yung Feng. You had to get the test times down and more efficient and more accurate. How did you go about this?

Feng: We saw two challenges at the time. One was testing the 8048 on the Q8000. The Q8000 testing requires a personality board. Basically it's based on the T-box concept that for each new product you design a system that is based on a known good unit, and you use that to generate all the test patterns required. And for the 8048, it has a ROM on board, and we potentially could have multiple numbers of patterns; which of course we did. So in order to accommodate the 8048 onto the Q8000, a engineer named Jean Pierre Masbou, working with Joe Louie, came up with the idea that for each ROM code we would run through a algorithm to generate a very short code, which we called a "checksum". So for each pattern, when the customer pattern comes in, we would develop this "checksum", together with the customer pattern; we'd put it on the ROM. So when the product came out of manufacturing, we would-use the tester to extract the code from the silicon and then compare with what it's supposed to be to identify each customer pattern. This way we eliminated the need for a reference unit on the personality board for each customer pattern.

Laws: I presume this is at the wafer sort level, before you packaged them.

Feng: At the wafer sort level, yes. And you would do this on a number of dies to convince yourself that you had got a good die and this is the right one. And of course before we did this, Jean Pierre had developed a very elaborate mathematical analysis to demonstrate that this code or this "checksum" would be unique for that code, that the probability would be very, very high, that this would work. And so once we found the code we would actually read the ROM and extract the pattern and use that as a reference to test the rest of the lot. And that basically is how we accommodated the 8048 on the Q8000. It was a big breakthrough at the time. Of course and without that we wouldn't have been able to test the 8048 with such a low cost tester. And with this, eventually as the volume grew, we were able to also transfer these boxes to offshore test facilities quite successfully. The other challenge was, as John mentioned, the cost of the product. It was in the beginning of 1978 when we foresaw this huge volume that was going to come in, and also associated with it was a very challenging ASP. I remember the ASP forecasted that by the end of '78 it was going to be like about $4.50, and at the time the cost of the product was about $8.50. So during that year, 1978, we had to figure out a way to reduce the cost to something more reasonable, and we set a goal of something around $2.50, for the end of the year. And we went out and we looked at every possible opportunity. The main things were yield improvement. At the time the yield was atrocious and-- which was the number one costing factor. And then we went on to convert the product onto four inch wafers, and 8048 was probably the first, or among the first product to be converted onto the four inch wafers, in the Livermore plant. And in the assembly area we tooled the lead frame and used the stamped lead frame and converted it to an epoxy die attached process from silver, and then we streamlined the product flow, and eventually we were very successful. By the end of the year I recall specifically the product was down to $2.53.

Laws: Which year was this, do you remember?

Feng: 1978.

Laws: '78.

Feng: So 1978 this cost reduction program was our number one project.

Laws: So you essentially cut it in half.

Raphael: No three.

Feng: More than that.

Laws: Hank?

Blume: I'd like to talk more about 1977, particularly a time before John got there in December. At the start of 1977, the first of the year, Dave Stamm, Dave Buddy had both moved to other areas. I'd hired Gene Hill, who had done some other chips, but was a superior talent, and I got promoted to a bunch of other responsibilities, as well as single chips. So I promoted Gene, and he and Bob Wikershine [ph?], who I think will be in part of this on the 8051, and a couple of other engineers, and they literally went through the 8748 and they looked at every node to see if it was fast enough. They looked at every node to see if there were non-fatal design rule violations, and we found some non-fatal design rule errors; we found slow nodes. Well we shipped almost 1000 units the first quarter of '77; I don't think any of them were at advertised speed, they were all slower. So we had to speed the product up, get the yield up. I'd invented an I/O chip called the 8243, which was so small and so stupid that Intel didn't know how to make it; they literally couldn't get something with yields so high, so they didn't schedule it very often. But we finally got that out so the customer would get a complete kit. So we had a clean product that the speed-- we could ship products to speed and ship with confidence. The second thing that happened is well technology was very benign in the development. In 1977 they moved manufacturing from Santa Clara one to Portland. I think it was the first process put in the Portland fab plant, and the product came out half as fast, and they didn't know why because it seemed to meet the test chip requirements, but it was just half as fast. So we had to diagnose that. Manufacturing wanted some design rule changes, they wanted more spacing from metal onto something else, and first layer poly was there, and they wanted to put a lot of first layer poly on the die so they could etch it away and the etch would work better. And it's easy to do on EPROM but it's not so hard to do on a very complicated logic chip where the basic ruby that had been used for a lot of other things. The second thing that happened is there was some sales resistance; particularly Mastec was very aggressive with 3870, and Bill Davidow remarked that we won every design that went under 5000 units and Mostek won every design that went over 5000 units; which kind of made him happy because it was 8748s. But we set up a low end taskforce, and I think it had Jeff Miller, Ben Franklin, Don Philips, Lionel Smith and Gene Hill, and they went through application by application and said how much ROM code, how much I/O, what else do you need? Kind of to vector the sales force into more profitable kind of pursuits, but this also resulted in a requirements document for the 8051. It became clear sometime in '77 that the addressability was limitation, and the need for the 8051 came out a little earlier. At any rate, from some sales resistance in the first and the second quarter, to I think an overwhelming acceptance by the beginning of the fourth quarter, we had a cleaner chip and we could ship the whole kit. We had the low end taskforce direct in the field. Mostek just disappeared. They stopped quoting 3870s. And I think it's just a wonderful world if a reputable competitor validates the concept by advertising they have it and then disappears, and you're there, the only guy in town.

Laws: The kind of thing you can't plan for.

Blume: That's just a nice thing. And then Intel made a management change. It used to be David Allen and Vadasz, two in a box for the components division. So Vadasz had microcomputer components division, David Allen had the systems division. And when Vadasz got there he took some volume ROM orders, and that got the field's attention, because the salesmen would talk to each other and find what's selling and what's not selling, and if their friends told them it's not selling, the factory couldn't push them. But if the friends tell them it was selling, they'd go out and sell it. They'd have more confidence that if their fellow salesman could sell something they could sell it. And there was a lot of increase in interest in the 8048 and the whole family by that time. So I think it went from a technical achievement, a superb technical achievement, but prototype quantities of parts that weren't full speed, to something we could make with some confidence; not at the right cost, but we could make a speed product with confidence and know it was there.

Laws: And the move to the Oregon plant, did that get a cost reduction or did it just provide more capacity?

Stamm: That was the EPROM technology.

Ekiss: That was the EPROM technology.

Blume: I don't know if the cost per wafer went down in that move at all. The yield went down -- the yield crashed for a while until it got up and its performance went down until it got up. Then I think it was a wash. I mean, it clearly was new people in fab. Oregon hadn't done fab before, it hadn't done anything more than token fab at Tektronix before that.

Ekiss: I think the availability, by '78, of the 8748 was always good.

Laws: And the volume designs were turning on then.

Blume: Well all these guys in the garage shops were starting to reorder. If you buy one from a distributor in the first quarter, in the third and fourth quarter you need to buy 20 or 50, and all of a sudden the 1000 we didn't know about turned into some reasonable numbers.

Laws: Now by this time Howard you'd moved on to National.

Raphael: Yes, I'd moved on to National.

Laws: You obviously had some confidence in this chip.

Raphael: Yes, I was moved onto National as marketing director, and then I was promoted to group director of operations. And one of the things that I did there is introduce the segmentation, the market segmentation that Intel had developed, which was right on the money. And of the course the project that I had national second source was the 8048. So we did an NS8048. It was a non-sanctioned second source. So it was essentially engineered from scratch. But there was no question about it, the market was there and it was a good little business for National in those days; plus we developed other things at National, higher end microprocessors, and we copied Intel and set up a design center in Israel, and that was to develop more comprehensive computer products. But the low end was still very much in the focus.

Laws: How big a difference did you find in the sales force at National versus Intel? One is very commodity oriented, the other is very technically oriented.

Raphael: Yes, well National's strength was very much production. They had a wonderful and competent production machine. But the difference between Intel and National, and why Intel was so much more successful in this area, in my opinion, was the focus of the organization. Intel, both its strength and weakness was its product line was very narrow, it just had memories and microprocessors, but it did a wonderfully superb job in marketing those two things. National was very broadly diversified. We had linear, we had digital, we had all different types of semiconductor product lines, including discretes; and that was a strength. But from a standpoint of a microprocessor product line it was actually a weakness, because the sales force would want to sell the products that had the path of least resistance, and microprocessors was a very technical sale. So it was very, very hard for me to get the attention of the sales force to really make the effort to sell those- sell microprocessors. That was a really big difference between the two companies.

Laws: And I think Barbara you mentioned when we were talking yesterday that as part of this process Intel had an extraordinarily good understanding of the segmentation of the market into a number of different areas.

Kline: Right.

Laws: Did you feel that at the time or is that looking back?

Kline: No I felt that way at the time. It was really interesting because when I had joined Intel the development systems idea was very small. As a matter of fact my first job there was helping customers debug Fortran code to create programs for the 8008. And as time evolved and we came into these other tools -- I was also the person that introduced the high level language for the 8-bit microcomputer family, microprocessor family. And what we were doing was groundbreaking work. Developing microcontrollers led to different ways of thinking about where micro-computing could go. At the same time, the 8-bit microprocessors moved on into 16 bits and above that. I just found it fascinating and exciting to be involved in trying to figure out how to really design with this brand new tool [development systems] and when the microcontrollers came out it required a different approach. And so we had to really look at how we could divide those two groups. And yes-- in a way I think Intel kind of fell into things. There were a bunch of really young, very smart, very creative people, and at the same time we had a perfect storm in terms of developing some really exciting things back then. It was a wonderful time to be there.

Laws: John, now you were responsible for manufacturing this product in high volume, and did you get involved with many customers and some pricing issues relative to their needs?

Ekiss: Well when we were a functional organization and I was in charge of manufacturing I didn't get involved in that. I was aware of the pricing and I knew what the task was in terms of costs, and so I executed against that. It wasn't until later when the Vadasz's microcomputer division made one last split and a separate entity for prosecuting the microcontroller business was set up and was migrated down to Arizona, and that happened in 1979.

Laws: Which is a whole other story.

Ekiss: Which is a whole other story which we can get to later, an upcoming event.

Laws: Okay, I think we're getting towards the end of the session here. Hank, you have a comment here?

Blume: Yes, I'd like to add one postscript. In the fourth quarter of '77 we started working to define the 8051. And the 8748 would have been temporary; it would have been a nice thing for a few years. The 8051 is now 31 years since definition and they're still making new variations of it, they're still adding new development systems for it. It has the greatest longevity of anything in electronics, except batteries maybe.

Laws: Maybe the 741 [op-amp] is a little longer but comparable in some way.

Blume: The 741 yes. Anyway John Horton did the instruction set one weekend and it was never changed, and I had two engineers working for Gene Hill, they were working on alternate instruction sets, and we had an agonizing week. We met at 7:00 every morning to compare the instruction sets, compare the cost, and finally I said, "Fighting another battle, charter battle, whose instruction set do we use, just isn't worth it. This is good enough. The cost is low enough, we can use it as it is." So we avoided a big political battle and went on with the 8051. We lost time in moving the operation to Phoenix because only one engineer that I'm aware of from the design group moved. But the 8748 story I think has to end at the start of the 8051, and a good architectural start, because the low end taskforce, John Horton's brilliant instruction set architecture-- just the fact that it lasted 30 years shows just how well that was done.

Laws: This is a nice segue to a future oral history we'd like to do. So what I'd like to do is to ask each of you for any final comments you wanted to make, and if you've got any particular story that was important to you as part of this project, we'd love to get it on record. John, we'll start with you.

Ekiss: Well I have a limited number of stories in the timeframe that you're -- leading up to the end of '77 or early '78 because of the position I had there. But I guess I would say that late 1977 when I came, and through 1978, was a very big learning experience for me, one because I learned how a company like Intel operated, and it was really night and day compared to the two companies where I had spent a fair amount of my career. So I was really learning a lot and I was very happy about that, and I really enjoyed coming to work every morning, and I remember that. And it was just even better when Intel decided to send a group down to Arizona and I was selected to head up that group. So I have really fond memories of that, and I still have friends, people I was friends with from that time period.

Laws: Thank you. Howard.

Raphael: Okay, I just want to say two things. One is it was a pleasure to work at Intel, and when I left Intel to go on, I left it for all the human weaknesses that people leave companies. I went on for more money and power, but I certainly didn't abandon the program because it was not a good program; I perceived it was a great program. I also want to say for posterity what a pleasure it was to work with these people during that period of time. And it really was. You don't know it at the time but in hindsight, when you look back on your life, it really was a wonderful part of my career and it was truly nice to work with these folks.

Laws: That's great. Barbara, last comments.

Kline: I thought that I was very lucky to be assigned to the microcontroller program. It was a different type of mindset than what we were doing in microprocessors. There were a lot of exciting things there, and one of the things that I really remember was traveling with Howard a number of times to go back to Detroit to talk to the automotive industry, and there's lots of stories about Intel's relationship with the automotive industry. But at the time, while we were back there saying, "This is something that is really exciting, it can help with fuel economy," and this and that and the other, I was thinking forward and what this was going to be like when these chips actually got installed in the cars. I recently purchased a new car, which was kind of a lemon, and I had taken it back about five or six times, and I had decided that there was a logic problem. The guy at the shop was trying to explain to me that this was a motherboard and that motherboards were perfect and never did anything go wrong with motherboards and that I just didn't know what I was talking about. So I reminded him that I had been there at the very beginning of this, and I got a new motherboard and the car works now.

Laws: Nice story. And your car is also full of microcontrollers.

Kline: Yes, as he was explaining to me there are no mechanics anymore. You push down on the gas pedal and it sends a signal off someplace. But it's fun to see 30 years later what evolved from back when we were developing this.

Laws: Dave, do you have anything to tell us? Did you have any perception as to where this thing was going to go in the big scheme of things in 20 years time?

Stamm: Well I thought I was in the beginning of something that could be huge in the whole microprocessor revolution. Having just graduated from college, spend a few years working on a chip like this was just a story come true for me. And I loved the team. I loved working with Hank, who would keep me grounded when I would try to go off with some crazy ideas. And Howard, my first exposure to, hey wait a second, somebody has to ultimately buy this; you have to market it, it's not going to sell itself; what's involved with that? And learning all of those components from him. I look back on it, what I learned for my career in those years was just phenomenal. It's also amazing to me that it was such a small group, just myself and one engineer, and one person had the entire responsibility for this component. If you look at the size of teams that are required today to do microprocessor development, or even the developments that were done three or four years later than that, they had teams that were at least five times that large, for most of those designs. So when I look back at it now and say how could we possibly have accomplished that? It sort of eludes me. I don't know. I think it was a great group of people. We had clear focus on what we were trying to accomplish, and we knew the schedule was important, and had all the parameters of cost and functionality, and all those were very clear and straightened out. And so my subsequent experience in having run a lot of engineering projects now is that if those goals are clear and everybody's aligned, you can get a lot done, and certainly that was the environment at Intel at the time. I didn't know how rare it was but it was a great group and I had a fantastic time.

Laws: Henry, this was an interesting time for you. The world changed.

Blume: Yeah, yeah, I was probably so deep into the forest I could just see trees and I couldn't see forest. But maybe a couple of quick stories. About the time I was phasing out of the single chip one of the automotive customers wanted two million ROMs the next year, and Intel calculated the value per wafer for what they'd get and told the customer they couldn't have any. That's not the way to treat automotive customers. Twenty years later I played tennis with a guy who was responsible for a design team at Signetics designing variations of the 8051 -- this is a long time later. But the story I have is really kind of a curious tidbit. I'm lying on bed Sunday morning reading the San Francisco Chronicle Magazine section, maybe 1982, '83, and there's an article translated from De Veldt in Hamburg. And Volvo was talking about a whole new chip that's going to revolutionize the consumer world and all the electronics in Europe, and I looked at the picture and it's my chip from five, six years ago. And I was going to write The Chronicle a nasty letter saying they'd been hooked by putting this thing in for something that was done in Silicon Valley, and I said oh “what's the bother, we've moved beyond that.” But we pioneered the concept, but the concept had a long way to run, not only in variations of it and in improvements of it, but just in spreading the concept. It was the start of a long-running, expanding world.

Laws: Yung Feng, what's your favorite memory of the project?

Feng: Well to the product engineers, we were always fire fighting, and as much as we tried to get out of that mode, but we seemed to be always facing some kind of fires -- the cost is too high, the yield is too low, the customer is not happy with something.

Stamm: That's the way it is.

Feng: But in retrospect to think about the time and the history that we have made, I am happy to say that I was part of it. I enjoyed it very much.

Laws: Well thank you everybody for your time and your memories. In 200 years someone's going to really enjoy this show. Thank you.

END OF INTERVIEW

