Oral History of Michael J. “Mike” Flynn

Interviewed by:
Bob Zeidman

Recorded: December 17, 2012
Mountain View, California

CHM Reference number: X6736.2013

Bob Zeidman: Okay, well I'm Bob Zeidman. And I'm here with Mike Flynn, who's had a long and important career in electrical engineering and computer science. And he's particularly well known for his achievements in computer architecture. Hi, Mike, glad to have you here.

Michael J. Flynn: Glad to be here, Bob.

Zeidman: I'm going to give a little background on you, and then I'll ask you a few questions. So, I know that you worked on the IBM 7090. And you were the design manager for the IBM 360 central processing unit. You were a professor at Northwestern University, Johns Hopkins University, and Stanford University where you were the director of the Computer Systems Laboratory. You were the founding chairman of both the ACM Special Interest Group on Computer Architecture and the IEEE Computer Society's Technical Committee on Computer Architecture.

Flynn: Right.

Zeidman: So, you founded both of those groups. You were the vice president of American Supercomputers, where I worked for you. And you became a mentor to me. And I appreciate how you've helped me out at several key points in my career. And in 1992 you were the recipient of the prestigious ACM IEEE Eckert-Mauchly Award. And you have a whole list of awards. I won't go through all of them here. You're the author of several important textbooks and over two hundred and fifty technical papers. And you're also well known, perhaps best well known at least to people who ask who Mike Flynn is, for the Flynn taxonomy, the SISD, MISD, MIMD, and SIMD classification of computer architecture. So, I'm honored to be able to interview you. And I'm going to start out with a few questions about your early years so we can get some background on who you are and what influenced your decisions in computer science and electrical engineering. So, maybe you can tell me where you were born and where you grew up.

Flynn: Yes. I was born in New York City and grew up there and went to school there. And eventually on graduating, I went to work for IBM in upstate New York.

Zeidman: And maybe you can tell me a little bit about your family growing up.

Flynn: Sure. My family, my mother and father, were immigrants. Well, my father was a special case. He was actually born in Brooklyn, but after this catastrophic flu epidemic where he lost several of his siblings when he was about six years old—his family moved back to Ireland. He returned to the U.S. after his mother died and his father remarried —he didn't get along well with his stepmother. And so, he decided to go back to the U.S. where he met my mother. I have three siblings. And one of them is a professor of computer science, also, at Polytechnic in NYU now. It's part of NYU.

Zeidman: And that's your brother, who's a computer science professor?

Flynn: That's Robert Flynn, yes, right.

Zeidman: Do you ever have any discussions, disagreements?

Flynn: Not particularly, no.

Zeidman: Joint work?

Flynn: No, we've never—he's done some interesting work, but it's never really overlapped with my own. I was more electrical engineering, hardware orientedand he was in a variety of other areas. [However we worked together at ASI in the late 80's, where he consulted on scientific computing; but that project came to an untimely end]

Zeidman: And I know that your Irish roots, I think have influenced you. I know you have a lot of connections to Ireland these days.

Flynn: Well, yes. Actually, as I said, my mother was born there and she actually returned there later in life. After my father died, she would go there for the summers and visit her brother and sisters. So, we actually bought a little place for her there, which we still own, a little cottage on the west coast of Ireland.

Zeidman: Okay. I think you do some work, some research or work with some researchers in Ireland. Is that right?

Flynn: Yes, I do. I have worked there in the past—I was visiting professor at Trinity College, University of Dublin. I was also a visiting professor at Queens University in Belfast at various times. I've consulted for some government organizations there in—basically in the area of information technology and trying to help the country, let's say, adapt and grow its economy in the information technology area.

Zeidman: And so, when you were growing up, what kind of things did you like to do for fun?

Flynn: Well, the thing that really got me into electrical engineering was amateur radio. I did a number of other things, but there were two things that I really liked. One was, probably from the Boy Scouts, the outdoors. I had actually toyed with being a landscape architect at one time. But that wasn't to be. I was sort of restricted in my education to go to a commuter school in New York City. So, I went to a number of colleges in New York City. But I was influenced also by amateur radio. I was an avid member of the amateur radio community, shall we say. And eventually I became the trustee for the Stanford Amateur Radio Club, and the faculty advisor before I retired. So, I did that for a number of years. And that, of course, led me into electrical engineering, which eventually then became computer engineering. I joined IBM and I watched the evolution of the hardware. It was a quite exciting period of time.

Zeidman: So, when you were younger studying amateur radio, if you don't mind because it might give away your age, but around when was that?

Flynn: Oh, this would be in the . . . when did I get my license? It was in the early 1950s.

Zeidman: Okay.

Flynn: Okay.

Zeidman: And I understand you also had a musical interest, that you were a musician.

Flynn: Yes, I was a—I played the French horn and the Mellophone. And I actually did get a small scholarship to Manhattan College to play in the orchestra with the French horn. I also, as part of that scholarship, I had to put away the chairs. So, I don't know whether it was the French horn or the chairs that gave me my little emolument.

Zeidman: So, you were very good at putting away chairs.

Flynn: At least, at least. Yes, yes. And probably as a good a . . . there aren't too many French horn players. It's a very difficult instrument. It's easy to screw it up as I well know. I can always admire people when they really do it well because it's easy to not just quite hit the noteand it's very obvious when you don't.

Zeidman: Do you still play French horn or any other instruments?

Flynn: No, I don't. I haven't done that in years now.

Zeidman: So, one thing I'm curious about is I know that people talk to me about this, about the differences between engineering and the arts. And music, and landscaping, and engineering all seem pretty different. But do you think they're different? Or is there any overlap there?

Flynn: It's a good question. In engineering, what you're trying to do is to apply scientific principle to social need. It's measured in the economy today by stock value and whatever else. In the arts you're also trying to anticipate social interest and perceived elegance, beauty. And of course, a beautiful product of engineering work can be elegant and appreciated, too. So, I suppose there's a continuum. Although, growing up, I looked at these as pretty distinct things.

Zeidman: And were there people that you looked up to, people you might say influenced your career or that you tried to emulate when you were younger?

Flynn: Not particularly, although certainly the head of the physics department at my high school was an enthusiast for amateur radio. I mean without him that avenue would not have been opened to me. So, he was very exciting and he was dedicated completely to his profession and this interest. He put a lot of effort into it; we had a great radio stationand that was very significant to me. It certainly was an influence.

Zeidman: Do you remember his name?

Flynn: Yes. He was Brother Patrick. I went to a parochial high school. He dedicated his life to—let's say to teaching physics, but with a particular strong interest in radio.

Zeidman: I guess in those days, electrical engineering fell under the category of physics. Is that right?

Flynn: Yes. Well, certainly in high school you approached electrical engineering through the physics program and through the physics curriculum. Physics enables you to understand a little bit about what's going on and how these things work.

Zeidman: And so, what was your first exposure to computers?

Flynn: Through IBM; actually, I chose to go to IBM because of several things. One, it had a continuing education program and I wanted to go on and get a master's degree at least; that was important. It was one of a number of companies that had this kind of program. But in those days, only the progressive companies would have such a program for their engineers. So, that was an important factor. And of course the emergence of the computer, which was dominated by something called UNIVAC, or Unisys today. IBM was more into electromechanical computing company than electronic computing. But they had done some big machines, mostly for the government and that made the whole thing a little bit more exciting.

Zeidman: And I know you mentioned some progressive companies. And I think we talked about there was IBM and Bell Labs. I don't know if there were others when—go ahead.

Flynn: Bell Labs was even more progressive than IBM. For years, they would take their young engineers, and would immediately put them on a program to get a master's degree at a university. IBM had an evening program in conjunction with Syracuse University, where I eventually got my master's degree, so that you could take a course every semester—two courses a year, until you completed your program. So, that opportunity was important to me.

One of the things I should add about IBM in those days, is that you didn't go to work immediately. In that sense, it was similar to Bell Labs. You went into a training program, which lasted for about nine months. In the mornings, you would take courses on logic design, circuit design and programming. I programmed an IBM 650, a machine which they may have pieces of out here [at the museum], an old drum machine. But it was a dominant machine at the time because it worked, and machines were notoriously unreliable in those days. This was a machine that really could reliably do computation. Also from that machine I appreciated the possibilities for compilation because it was a drum machine. So, in assembly language, when you put in an assembly statement in, you could figure out where the next instruction should be by knowing the amount of time it would take to complete the current instruction. So, if you were really clever, you could put it so many addresses further on in the disk. The disk [head] would be there to pick up the next instruction without delay. Well, IBM had a program, which did this automatically. So, the genesis of optimizing compilers was there.

Zeidman: Oh, okay.

Flynn: It was in assembly, though, but it still was an optimizing assembler. But that was one of the things I dealt with when—early on in—I guess it would be the middle to late 1950s.

Zeidman: And if you compare IBM to companies today just in general, do you think—do they have these kinds of programs? Do you see a change for better or for worse?

Flynn: That's a good question. I don't particularly know of them, although Stanford certainly has an ongoing program with a number of companies where Stanford broadcasts master's courses to these companies. You can pick up a course even during company time. So, I suppose, it's actually more convenient today to participate in a program—and it's supported by a number of companies throughout the Bay Area. And I'm sure that nationally the same would be true some sort of way.

Zeidman: So, I know when we talked I mentioned that my parents didn't want me going into computing because they saw it as—they were in contact with people who punched cards all day, people who maybe had a high school education. And they were surprised when I was said I wanted to go into the field of computers. Now, you were twenty years before me, how did your parents react?

Flynn: That was interesting. You see, they had been immigrants, and they gave me total support. I did ask my mother what about my going into electrical engineering. And she said, “Oh God, Michael. I’ll pray for you.” And that was the—

Zeidman: In a good way, though.

Flynn: Yes. Exactly right. I’ll pray that you succeed, in other words. She’s not going to give me any advice at all on this. So, it was my father, who had been a policeman, he insisted that I take the New York City police exam just in case it [engineering] didn’t work out. <laughter> You’ll always have a good job working for the city. But, of course, I didn’t have much interest in that.

Zeidman: But you did take the police exam?

Flynn: I took the exam out of deference to him. Yes, indeed.

Zeidman: And how did you do?

Flynn: Oh, I passed it. I passed it.

Zeidman: Good.

Flynn: <laughter> Thanks.

Zeidman: You’ve got another career in the future if you need it.

Flynn: <laughter> I still had a fall back, yes.

Zeidman: So, and then after that, I know you studied at a number of colleges. Then you already talked about going to IBM.

Flynn: Yes, right, right.

Zeidman: And so, and at the IBM you worked on the IBM 7090, and of course the very famous IBM 360.

Flynn: Yes, I wanted to correct something you said earlier about the CPU. I primarily worked on the Systm360/model 91, their largest machine. Earlier I did work on system 360 architecture in general, when I worked for Fred Brooks, who was really the— the System 360 program manager; I forget his exact title. Then it came time to implement this architecture. Remember the architecture is an instruction set, which spans implementations from small machines to very large machines. Well, the model 90 was the largest machine, and I was the CPU design manager for it. And we were competing with such machines as CDC 6600 designed by Seymour Cray and Jim Thornton, etc. And so, we did build the Model 90 series [Model 90 and 95, later to include the Model 195] which was very innovative. And I was delighted to have the opportunity to work on it.

Zeidman: So, maybe you can tell me about what was particularly unique about them, what you liked about them, what their advantage was.

Flynn: You mean that particular model or the system 360, in general?

Zeidman: Actually both if you could talk about first the model 90.

Flynn: Let me reflect on [system 360] first because it was an important epic in the evolution of computers — I'll go back [before system 360] to the 7090. There was the 709 machine, which is a tube machine, the last of the great IBM tube machines. And I'm sure there are fragments of it [the 709] in this museum [Computer History Museum] here. All of these tube machines were expensive machines and difficult to maintain, etc. IBM put a lot of effort into making a reliable big scientific machine. And as I said earlier, they had already had some success with the 650, which is more of a commercial machine. There was this machine class distinction, between a decimal machine and a binary machine: binary for the scientists and decimal for everyone else.

Zeidman: When you say decimal, can I ask you, you're talking about decimal encoding of the—?

Flynn: Yes, the—all the data formats were decimal data formats.

Zeidman: In the hardware?

Flynn: Yes, in hardware as well as software.

Zeidman: Right.

Flynn: So, the plan was, in those days, that they would move to this new generation of transistorized machines. This would be epitomized by something called the Stretch, or the IBM 7030, which was a bridge too far, actually. They did an enormous amount of work in standardizing on circuit cards and creating a methodology for design. I have a whole little talk I give sometimes on that moment where the methodology of machine design and the methodology of manufacturing changed. IBM was really advanced, not so much in its architecture or anything else, but in its ability to make things, its ability to manufacture and manufacture reliably, and take designs and preserve the integrity of the design as well as create the ultimate product. So, the government—the Department of Defense had come to IBM and asked it to produce, I think it was like twelve or thirteen, transistorized versions of the 709 [the tube machine]. So, this became then a technology tour de force. It was not changing the architecture nor changing the logic particularly. It was simply mapping the logic over into this new technology. I had been a circuit designer in these days, and so, naturally, it had come to me to manage a little prototype of this machine and to prove its feasibility and we did it. This was the first instance of using a bipolar technology called emitter coupled logic, which became a standard for years thereafter for high speed switching technology.

Zeidman: Right.

Flynn: That was developed in the group that I worked in at IBM. And one of those circuits would fit on a small card; there was a little cap you could cut to change a circuit configuration slightly. So, with a series of about twenty different kinds of cards, and the ability to cut these caps, you could do all the logic needed. IBM automated the production of these small cards; It was called SMS, standard modular system. The cards would fit into these larger panels. I think they were like twenty-eight [card slots] by twelve rows, twenty-eight columns. And each slot would take one of these cards; they [the panels] were wire wrapped to connect the wires on the panel. This could all be done by these huge Gardner Denver pneumatic machines. The process was managed by punched cards. There was a whole design methodology called ALDs, automated logic design system: each block on a transparent sheet represented a card [circuit] and the designer would connect these. Then overnight, somebody would take the designs, and create the punch cards that would be run [after processing] through the equivalent to what we would now call a place and route routine. It would all be documented and that documentation drove the manufacturing process. When the government asked IBM how much they would charge for these machines, IBM priced it based upon the way they were building the old 709s. But by the time they were really built, it was done with this automated process. What could have taken months before, would take—you could build one of these machines practically overnight. The same thing happened with memory, with the core memory. And the profits were extraordinary in the core memory. IBM, first of all, when they produced these machines, they thought they were only going to produce them for the government. But it turned out that the competition had so much trouble with reliability and other things that this was the only reliable alternative, and it dominated the field. I think three hundred or more of these machines were made. And the price wasn't reduced particularly because there was no need to do that.

Zeidman: And about what year was this?

Flynn: This would be delivered first in ’59 and delivered—in variations, through the middle ‘60s because after that system 360 would take over. But I have to credit IBM. The profits on that machine—they sold for about three million dollars a piece. These are big million dollars, too [adjusted for inflation]. And three hundred of them, you can see this is not a bad number—and the profit margins on the memory were over ninety percent because the memory had been priced on the basis of using oil cooling to keep the temperature of the cores constant, but it turned out that air cooling was fine; so, this expense vanished, and so, it was a very, very profitable machine. But I credit IBM for having the imagination to turn its back on all those profits. They could have just continued making these machines, but they came out with something new, which was system 360. It was regarded as a very risky venture at the time: throw away all these other machines, each one of which was quite profitable, and organize them now into a single architecture, a single instruction set architecture. It was a bold and imaginative thing to do. So, I think later on – in the ‘80s or so- when it came time to do something similar, IBM [management] wasn’t quite up to it.

Zeidman: So, before you go on to the 360, I had a couple questions about some things you said. You talk about cutting caps on the ECL transistors.

Flynn: That’s right.

Zeidman: You’re talking about physically cutting them?

Flynn: No, no, these are—yes, it would be done by the machine of course. But these were little like rails, little—I have one somewhere. And you would cut the rails to reconfigure the circuitry.

Zeidman: Okay.

Flynn: You would attach a load or no load, if you were going to dot OR the circuits together, or whatever, through this mechanical manipulation of the rails. You’re just cutting them off to isolate this circuit from that circuit.

Zeidman: And then you also talked about taking—I think you were talking about schematics, and they were going to punch cards. Was that a manual process or an automatic? Was somebody actually looking at it, or was there some automatic process?

Flynn: The designer would put the card and the card type into a box and draw the line to another box representing the other card it was going to be connected to. Now, that process then would be transcribed overnight by a support staff and that would be put into punched cards. Then whatever was entered would be produced for the next iteration of the design. Then the design could be printed out as a hard copy version of what you had entered in pencil before. It was backed up with the punched card file.

Zeidman: And was that unique to IBM? Or was that done elsewhere?

Flynn: At the time, yes [it was unique]. As far as I know, there was nothing like it. And in fact, many of the other companies, rival companies, in those days approached the design of a computer much as they approached the design [and production] of a TV set: to fix an "oops" you go this far down the [production] line and fix this. So, you fix it and forget about the other earlier machines. IBM's approach wasn't that way because, in fact, they owned all the machines. They were only renting them to customers. So, they had this complete dedication to service because it was their own machine.

Zeidman: Did the other companies also rent their machines? Or did they sell them?

Flynn: They largely would do either.

Zeidman: I was just going to say I see a lot of companies still taking the non-IBM approach of just fixing something, getting it to work, and throwing it out there.

Flynn: Yes. Yes, yes, yes.

Zeidman: In fact, in software unfortunately, there's pushes to do that kind of software. That's considered good development. Why spend time planning and debugging if you can put something out and get your customers to debug it for you?

Flynn: Yes.

Zeidman: Unfortunately. I think unfortunately.

Flynn: Unfortunately, yes.

Zeidman: So, then the IBM 360, and you mentioned that it was a big gamble. So, tell me about what was the gamble and why did IBM want to do this.

Flynn: Well, I described the early system to you, the ALD system, the design system. I described the technology with all of these Gardner Denver machines and wire wrapping. Basically, IBM threw all that out and started over again less than five, six years after developing this technology, which allowed it to begin to dominate the computer business. Threw it all out and started over from scratch. Of course, they had everything they had learned, but do a completely new system; the technology was new; the architecture was new; the software was certainly new: it was a big gamble. They were ceasing, basically, the development of follow-on products [for the older product lines]—of course they had to do that, for these other machines, the 1401—which I think there's one here in this museum, the 7090, and all of these other lines, the scientific lines, the commercial lines. The development of follow-on products ceased. And all effort went into this 360.

Zeidman: And so, did they abandon products that were partly through their process?

Flynn: No, they simply—remember they owned these machines. So, they just simply, basically, ceased development on them. Although in some cases, IBM was certainly not adverse to having two [development] programs for the same market and just compete internally and then at the end…pfftt…close one down.

Zeidman: Did people mind that? I know there was a lot of loyalty to IBM.

Flynn: Yeah, it was a—I'd been in that position several times. And sometimes you won. And sometimes you lost. And it wasn't any fun to lose.

Zeidman: And about what year did IBM start on the 360?

Flynn: Oh, sure. That would be in about 1960; well, it probably had its genesis earlier. But certainly by 1960, it was conceptually under development. IBM had sent me to Purdue for my PhD and I came back in '61. I was immediately working on the architecture, the system 360 architecture.

Zeidman: So, how did you get sent to Purdue? What was that?

Flynn: Oh, Purdue had asked IBM to send an instructor there to teach logic design and computing. I had finished my master's degree and was looking around [for a PhD program]. In fact, I had asked if I could take a leave or something to go on and get a PhD. So, they said, oh this is an easy one. One and one make two. So, we'll send me over to Purdue, and you teach a couple of courses over there, and you got a couple years to get your PhD.

Zeidman: Oh and about how long did it take from conception to getting out the door for the IBM 360?

Flynn: Well, you see it came out in pieces because—and I suppose our machine, the model 91, was the last piece. And that came out in—if I'm not correct about—first one was delivered in about '67. So, from about '63 through '67, the various models came out.

Zeidman: And, if I remember correctly, you were heavily involved in the floating point unit. Is that right?

Flynn: Well, I was the design manager for the system. And there was a lot of novelty in both floating point, but also things like—now, I'm certainly not claiming responsibility for all these. It was a wonderful team of people there. For example, Bob Tomasulo developed "Tomasulo's algorithm", which he called—"the common data bus". But I called it the Tomasulo algorithm in a paper that was written contemporaneously with this [machine development] and that name stuck. This was an early version of data flow. Other ideas: the speculation on branch, the out of order execution of instructions, all were—they were very early developments in computing, incorporated in this machine. As far as floating point was concerned, the pipelined floating point adder, the use of multiplicative divide, a Newton-Raphson type divide, were all, I think, relatively novel. I'm not going to swear that each one was unique to that time, but it was all quite early on. And so, it was a quite impressive accomplishment.

Zeidman: Now I know that we hear stories about companies like Apple, and some of the modern, more recent companies doing development on the Macintosh at Apple, for example, or I'm sure there are others—Facebook in the dorm rooms, where people have a feeling of camaraderie. And it feels like you're going to change the world. Of course there's a lot of stories about people who have tried that and didn't end up changing the world.

Flynn: <laughs> Yes.

Zeidman: But in the IBM 360 did you have that feeling? Did you know that you were creating something that was going to be big and significant and feel that kind of camaraderie?

Flynn: I— within the group, yes. There's always challenges from outside the group. So, there's a degree of camaraderie. There was a saying in IBM that you didn't mind if the [external] competition would find this out, but you certainly didn't want this rival group [within IBM] to find out whatever it was because they could hurt you. The [external] competition…pssht… they weren't going to go anywhere. <laughter> So, of course it was the measure of the dominance of IBM at the time.

Zeidman: And I know that I've heard stories about IBM and Thomas Watson, I guess Sr. and Jr., creating an atmosphere of pretty significant loyalty to them and to IBM. Did you feel that way? Did you see that at IBM?

Flynn: Yes. Yes. In fact, I got into trouble. And one of the reasons I got transferred, I started with IBM in Endicott, and one of the reasons I got transferred to Poughkeepsie, which was probably a better move for me anyway, was a dinner that was—I was on the training program, as I said. And the great man, Mr. Watson the elder came to dinner. And we all stood up and had to sing this song to him, you see, which—with our hearts—you know, whatever else. And so, I didn't stand up. And I said I stand up for the country and for the deity, but not for Mr. Watson. So, that was probably—it showed that I wasn't ready for the paternal aspect of IBM. So I went over to Poughkeepsie, and I think I prospered there.

Zeidman: Well, they must have realized your contributions to the company if they kept you despite all that.

Flynn: Right. Yes, well I'm sure I was on a short leash. In other words, I wasn't going to get two of these [incidences], but it really didn't much matter to me. You know when you're young like that you have a lot of confidence in yourself—it'll work out. And it did in the end.

Zeidman: And I know on the IBM 360, you worked with a lot of people who also went on to become pretty well known. And wonder if you could tell us about them? Some of the people that you worked with on that project.

Flynn: Well, the overall management of the project and leadership was in the hands of a man named Bob Evans, who I guess was lab manager at the time. He had a number of different titles, but the heart and soul of system 360 was Bob Evans and he pushed it. He—and of course he was, for me, a lifelong friend and I worked with him in a number of projects later even after we left IBM. Of course I mentioned Fred Brooks, I worked for Fred Brooks who is well known to all. I worked with Gene Amdahl; we wrote papers on early versions of the model 90, and—Gene was a contributor to this machine. He certainly reviewed all the designs and made helpful suggestions.

Zeidman: So, I'll just mention that Bob Evans—so he became vice-president of IBM.

Flynn: Yeah, that's right. He was the division president for a number of divisions and then eventually vice-president.

Zeidman: And Fred Brooks. I know him from authoring the book the "Mythical Man Month."

Flynn: Yes, right.

Zeidman: And of course, I actually didn't realize it. At some point I found out that he worked on the IBM 360. I knew him as a guru of engineering management.

Flynn: No, actually he was [360 architecture manager], the architecture—distinct from all of these other things that were going on, the technology and everything else, Bob Evans had the complete charge. Fred had more or less the instruction set and the decomposition into various models, which model would do what, and whatever else, so the sort of the planning of the implementation, rather than the implementation. Eventually he took over the OS 360, the big project for the operating system. So, that's where he gets the [insights for the] mythical man month.

Zeidman: Okay.

Flynn: Because he certainly saw that first hand, the difficulty in producing reliable, large software programs.

Zeidman: Okay so, Mike I was going to ask I think just one more question about the 360.

Flynn: Okay.

Zeidman: And that is, I know it's considered a success, so I assume you would agree it was a success.

Flynn: My God, to this day, they're still making these machines. I think now it's called the Z series or something like that. I would have never, never projected or predicted that it would last more than 50 years, as an instruction set and as an evolutionary series of machines. So surely it was a success. In fact, it was one of the reasons IBM ran into trouble, my speculation—I'd left the company by the 1980s or so—but it ran into serious problems in the late '80s and it was because people become blinded by success and profit. You have this one line of machines, which is making an enormous amount of money for you. So you're very reluctant to move off to some other machine or to try something new. That's why I give Watson a lot of credit, when he moved— when he agreed to go to the System 360 and abandon 7090 and 1401 and all the other machines, because they were very profitable lines. And when the test came again, to abandon 360 perhaps for something else— and in fact, IBM had produced what we called the RISC machine, all the optimizing compiler work, networking. I mean, you could go down a litany of things. The PC, everything, was done either by IBM's research division, or by some commercial division. Yet, when it came to bet on any of these things, no; it was blinded by the profits it was getting from the 360 or its follow-ons, 370, 380, whatever they were. It was my speculation that was one of the problems that it had. Here's a company with all the technology, innovation that you could want, all the money you could want and still unable, from a management sense, to put it all together and to make some transition. So later on, I think, under new leadership in the '90s, it did put itself together.

Zeidman: In the '80s, are you talking about one particular thing, like the personal computer or just in general?

Flynn: Just in general because IBM's personal computer was the first one, but rather than making it itself, it chooses this little company, Intel, to make the devices, and it chooses another little company, an unknown company, Microsoft, to do the systems software for it. So yes, it illustrates the concept of the product was IBM's, but its ability to look at reduced profit margins, that it couldn't deal with, it seemed to me anyway. I was an outsider at this time, but I can see the two transitions. I can see the movement to System 360 on the one hand, and the difficulty it got into with the progress of the technology and the movement to personal computing in the '90s, say, early '90s, late '80s.

Zeidman: It's interesting that during those '80s, I remember that IBM was applauded for going to outside companies, to assist it in the manufacture and software of the processor. But I think history's shown that maybe that was not the best choice for them.

Flynn: Yes.

Zeidman: You were one of the pioneers in the field of computer architecture, as far as the study, and then turning that into a science. So how did you get interested in that in the first place? In other words, how did you get interested in not just designing hardware at a place like IBM, but then turning it into a field of study?

Flynn: I left IBM, as I said, with mixed feelings, because I really liked IBM, but I wanted to have a freer hand. It was clearly going to be a very constrained environment, in development anyway, and so I went to a university and there was a lot of interesting questions. You had freedom at the university to look at these questions and to consider them. And of course, I think the most interesting question, in those days anyway, was, what role do various pieces of computing play in the optimization of the computing process? There are so many tradeoffs, such a large design space, and the architecture plays such a pivotal role in determining which element of the design space is to be emphasized and which dimension of the design space is really relatively unimportant. So it seemed to me to be sort of the basis for putting together the better computation.

Zeidman: What year was this, you left IBM?

Flynn: I left IBM in, I think it was '66.

Zeidman: That's when you went to Northwestern University?

Flynn: That's right, basically, right.

Zeidman: I know when I've worked in the computer field, designing processors, there was always this sexiness about designing certain parts. For example, designing the CPU was more interesting than designing the memory controller. I found that people wanted to emphasize that. Did you find that when you were either at IBM or doing your research in architecture, that there was a tendency for people to want to work on some parts, even though improving another part might actually affect the overall performance more?

Flynn: You know, you have to look at where the expense is and where the performance is. If you're limited by memory, I mean if memory is just a brick, if you will, something that can't be particularly refined in any way, then if there's not much you can do with it, well, you just have to live with it. But if memory is the most expensive part of the system and you can do something to support memory, then of course, it's up to you to do that. So it's really understanding the role of that component or that subsystem and its importance to the overall computation, or corpus of computations that you're trying to serve that makes your decision. So I think that the architect has to put it together and it has to be rooted in an understanding of what the applications are. The biggest criticism I have of the field, in a sense, is that someone coming up and saying, "Well, I have a new architecture," and you say, "Well, what is it good for?" They'll [the users will] have to figure out how to program it. Unless you look at the need, unless you look at what the problem is, or at least a series of problems, depending upon the scope of what applications you're trying to serve, you can't really get a notion of this optimality. You can't do it. To say, "I have a new idea," that's fine. It's nice to have a new idea, but does that idea have any value? That's a much more difficult question, and for whom does it have value? That's why I always liked to work and I also admire early designers who were working together with numerical analysts, etc., to try to create answers to questions, rather than create answers to issues that there are no questions for.

Zeidman: You were obviously a pioneer in the area of computer architecture, but who else would you consider a pioneer, people that you worked with or people that you are familiar with?

Flynn: Oh my goodness, there are so many of the people that had good understanding of these issues. Certainly a man I admired greatly was Seymour Cray and his colleague Jim Thornton at CDC. I worked with Jim Thornton later on a national academy committee. They had good understanding. First of all, they understood how the machine was going to be used and for whom it was intended and then they went to serve that need. So it was large scientific computation. But there are many others who I won't even begin to try to give you a litany of their names.

Zeidman: I remember at American Supercomputers. I talked to some of the people working on the memory controller. It never occurred to me that things like interleaving and interleaving banks and looking at what kinds of instructions were running would actually improve performance.

Flynn: Right.

Zeidman: So that was interesting. How did you come up with the Flynn taxonomy, the SIMD, MIMD?

Flynn: That was done after I left IBM, shortly after I left. I wrote a paper, "Very High Speed Computing" and I just did that to organize my thoughts. I was talking particularly about the Model 91 in this paper, among other things. I looked around the field. This was for a special issue of the Proceedings of the IEEE on computing. So I looked around the field. Well, there are all sorts of parallel things going on here. How do I distinguish them, how do I talk about them? And so the taxonomy became clear to me then. Contemporaneously with this, there was a movement to build the ILLIAC IV. The designer behind that was a professor at the University of Illinois, one of the great, I believe, people of the day and sadly he died at a young age. His name was Dan Slotnick. He had a great debate in 1967 with my colleague at IBM, Gene Amdahl. Amdahl's position was that there's no need to get into this parallel stuff because if you just wait a little bit, things are going to get faster and better, and then you don't have to change your programs at all. And Slotnick's point was, well maybe so, but in order to take advantage of parallel processing, you have to put a lot of effort into it, but if you put the effort into it, you can achieve really spectacular results, because you were able to use smaller elements, simpler elements, but many of them and achieve the much better performance. There's a famous quote of his that I can't repeat offhand, but basically, he says that if the premium is for ease of programming, forget parallel processing. He put it in a nicer way. But of course, nowadays, you see, you can't avoid parallel processing. So I believe his insight is really applicable today, more than in 1967. This notion that you can have the speed up without putting a lot of effort into it I think is illusionary.

Zeidman: You mentioned Gene Amdahl. You had worked with him on the 360.

Flynn: That's right.

Zeidman: And then he went to form Amdahl Computers at some point.

Flynn: That's right.

Zeidman: Were you involved with that?

Flynn: No.

Zeidman: Do you have any thoughts about it?

Flynn: He actually went and helped IBM form an experimental laboratory in Menlo Park. He and a group of people, some of them from the Model 90, went up to Menlo Park to build a grand new supercomputer. As in many cases, it failed. IBM disbanded the effort after a while. But Gene was enamored with many of the ideas and he left IBM and took those ideas and made it the basis of Amdahl Computing.

Zeidman: OK.

Flynn: I did talk to him from time to time over the years. I haven't seen him in a number of years, but I believe he's still around and still with us.

Zeidman: The next thing I was going to talk to you about is your university research. You started at Northwestern and you taught there, I have from 1966 to 1970.

Flynn: That's correct.

Zeidman: Nineteen seventy to '75, you were at Johns Hopkins University.

Flynn: That's right.

Zeidman: And then '75 to '99, when you became professor emeritus, you were in Stanford.

Flynn: That's right.

Zeidman: Maybe you could tell me a little bit about what you did and what you thought, first at Northwestern.

Flynn: Well, at Northwestern, we did a number of things. I'm trying to remember back that far now, but I was particularly interested in microprogramming and in various kinds of— well, I've always had an interest over the years in arithmetic, computer arithmetic. I had a number of different interests in the computing arena. In those days, I think that I was primarily concerned with parallel processing and trying to understand how these [computation] pieces could interact together in a positive way, and to try to enhance it, I suppose. Throughout these entire years, I was always trying to find out the best way to do computation, the most reliable way, ways which could be done with ease of design, and ways which would provide the best cost =performance.

Zeidman: What would you say you liked about Northwestern? Was there anything you didn't like? What led you to leave and go to Johns Hopkins?

Flynn: In each one of these cases, there are just accidental things that happen. I went to Johns Hopkins because they had just developed a computer science department. It was a small department and it seemed to have a lot of opportunity, and I would become a full professor there. I reluctantly left Northwestern. I was perfectly happy there, it was just that this was a promotion and it seemed to be a broader opportunity. Now at Johns Hopkins, this rumor may be apocryphal, but the Penn Central railroad failed while I was there and their endowment suffered. In one of these academic internal struggles, some of the academics decided they had to do something. We had to close down some departments. So even though the computer science department there actually had been funded by the National Science Foundation on an independent grant, it was chosen for closure. So I went over to electrical engineering. I didn't lose my job or anything. It was nothing like that, but I was discouraged by their, let's say, shortsightedness. Later, of course, they did re-form the computer science department.

Zeidman: I was going to ask that, yes.

Flynn: Yes, but I left when I got an opportunity. When Stanford asked me if I would consider leaving, I said, "Sure."

Zeidman: And you came all the way to the west coast at that point.

Flynn: Yes, right, exactly right.

Zeidman: You were married at the time and brought your family here?

Flynn: Yes, yes, yes, and each one of these was disruptive, but my dear wife was tolerant and went along with the joke, if you will, with the ride.

Zeidman: Doesn't seem to be too bad a place to go to.

Flynn: Yes. Exactly right.

Zeidman: I have some overall questions. Maybe you've answered this already, but what do you think was the most exciting project in your career?

Flynn: Boy, you know, that's a tough question, because each project you work on provides a different outcome and a different challenge. We talked briefly about American Supercomputers. That was a grand challenge and a very disappointing one in the end. On the other hand, it illustrated something. There was this notion, and there still is this notion, that you can go out to the market research people and they will tell you what the market's going to be like in five years. They're going to poll 341 little people, make some phone calls. They're going to fill out the forms, then tell you. Now in those days, and what year would this be now? This is in the '80s. I don't remember exactly the years, but the conventional wisdom was from all of these market researchers, that there's this emerging market, mini-supercomputer market, which is just ready to explode. In fact, there were one or two attempt: Convex, I think, was a relative success at this, but it was clear that this was going to be the next bull [market]. And market research was all there to show you this. We, as you know, formed this company, that with Cray's permission would do a small version of a Cray machine. It seemed to be a reasonable thing. We had some good ideas and good technology and we would put it together. Now we failed to get second round funding, which we needed. Any of these projects are not cheap to do, but looking back on it, we were fortunate. We were fortunate because that market was never there. What was happening, and no one could see it, was that the microprocessors were aggressively improving their performance. So by the time that these other machines would have been mature or ready for delivery and ready for full production, the microprocessors were coming along and yes, they might have been a tenth the performance but they were a hundredth or a thousandth the cost. And so the marketplace for this cost -performance, scientific computing, started to evaporate. And even the successful players failed. That was very important—very disappointing of course—but very important to see that the future is very difficult to see. It's the understanding of the movement of technology. This is a message for all of the architects too. You cannot design the machine for the technology of today. You have to design the machine for the technology of the time when your production is expected to reach full peak [production]. And that technology is quite different and can be sometimes maybe even unknown now, but you have to have a really good feel for it and a good way to transition to something new as it comes out. So that was certainly an important lesson.

Zeidman: I was trying to calculate in my head. I think it was 1987 or 1988; somewhere around there.

Flynn: Yes, that sounds about right.

Zeidman: And I do remember we ran a simulation at the end when we were trying to get funding. I think some people were excited that we got good performance on a word processor application or something like that. We could beat an Intel processor by a factor of four, but of course, our cost was a factor of 100? 1,000?

Flynn: Yes, right, right, right, right.

Zeidman: If you could do anything different in your career, is there something you look back on and say, "I wish I had done this differently?"

Flynn: That's a tough question. Actually, looking back on things, if I knew I was going to fail with American Supercomputers, I wouldn't have done it. So sure, there are things that I wouldn't have done if I had known that it was not going to work out. I certainly think my years at Stanford have been wonderful years, and my years at IBM were wonderful and my years in academic life were fine. I have no regrets about that path. I can't go beyond that, I guess.

Zeidman: OK. Fair enough. What are you working on these days?

Flynn: I'm trying to help a couple of young companies and one of them is Maxeler where I'm chairman of the board. It's again in high performance computing. They're using a dataflow approach, which goes back to my old taxonomy, MISD, where the data flows through a fixed set or a fixed graph of instructions. They seem to have been very successful and I'm delighted to see them succeed. They're based mostly in London, so I get over there a couple of times a year. So far, so good. It's like American Supercomputers in a sense. I don't want to wish that result on it, but you never know until you're really successful what the result will be.

Zeidman: And I guess, of course, there are different measures of success. You'd always like to have the financial success, but I know that I learned a lot, and I think a lot of people at American Supercomputers learned a lot.

Flynn: Sure. You can learn from failure too.

Zeidman: Although it's more pleasant to learn from success.

Flynn: <laughter> Amen, amen.

Zeidman: Are you still utilizing any of the concepts from the IBM 360 at this new company?

Flynn: No, not particularly, although I must say that it reminds me in many ways of what IBM was in the beginning, which is a service company, meaning that from the time that you've identified a customer and an application, you stay with that customer and provide service to them for as long as they need it. And whatever service that is, you provide it. In fact, in this case, you're designing a machine for an application. You're given an application and you're implementing this application on a large FPGA [field-programmable gate array] or series of FPGAs. So it's a custom designed machine for the application.

Since the pipelines can run 500 or 600 deep, you can have 2,000 operations going on at the same cycle, every cycle. This is what gives you the computational power. It's the massive size of these FPGAs that enables you to get the computational power. Even though it's slower per cycle, it's the massive amounts of parallelism that's achievable, that gives you the performance improvement. So I find that very interesting. This is a sort of approach which hasn't attracted as much [recent] attention, this dataflow approach. In the old days, dataflow, of course, goes back to Jack Dennis at MIT and Arvind and others. But in those days, there was no way to interconnect the computational nodes in a sensible way. You could have these nodes, but how do you get from this node to this node for this application and this one to this one for another application. Well, you could conceive of some enormous switch, but practically speaking, it was unrealizable. But the elegance of the FPGA approach is that you do have this massive interconnection and it does enable you to create very large data flow graphs and then to realize the potential of that approach. So for me, it's interesting because it's an approach which really hitherto hasn't been much explored because of the limitations of the technology. And yet now we have technologies which have emerged simply because they're getting larger and larger, which can provide a vehicle to realize this kind of computation. So I'm delighted to be able to see this, even if financially it doesn't turn out to be the next Intel, it's okay. Which reminds me, many years ago—and this is when I was at Northwestern. I'd just left IBM—I knew, of course, Bob Noyce from my time at IBM, because we had in fact given subcontracts to Bob Noyce [at Fairchild Semiconductor] and some people at TI, [Jack] Kilby in fact, and some people at Motorola, to build these [16-bit register integrated circuit] chips for us for our Model 90. He had left Fairchild to form this little company called Intel and he asked me, would I like to join him as one of the founders? It's another one of these mistakes, right?

Zeidman: <laughs>

Flynn: I said, "No, I don't think so. I think I'll stay at the university." <laughs> There are many ways of making mistakes, but in this business you can still do quite all right, even if you made a number of mistakes.

Zeidman: Mike, you've done some work for the government, I understand and maybe you'd like to talk about that also?

Flynn: I've been a consultant for a number of agencies and had a long term relationship actually with NASA, both here at Ames—well, NASA Ames and the Department of Energy in particular, Argon Laboratories, when I was at Northwestern. I found both of the research facilities to be really fascinating places and it's certainly very hopeful to me and understanding better the needs of the, let's say, computational community. There was one other case which I thought was really fascinating, where I was eventually the chairman of the National Academy of Science Committee for Ballistic Missile Defense. And the committee had been formed before I was a member, actually. It was charged with answering this question: we're going to spend all this money on antimissile missiles. These are the ballistic missiles that eventually actually were built in North Dakota.

Zeidman: About what year was this?

Flynn: This would be during Carter, let's say the late '70s. It was called Sentinel at one time, but it had all these different names. With each President, there would be a new version and it would change its name again. Safeguard it was called, under Nixon, I think it was. But the underlying question was, and raised by a number of people, was, would it ever work? I know it was a million lines of basically assembly code. The dimensions of it are extraordinary. But in any event, the question was, would it work? I think McCracken was the original chairman of this committee and basically, he said, "We don't know." So the committee continued and I became a member of the committee and we were looking at various technologies that the army was playing with to do this, that and the other. I can remember going in these bowels of some mountain on a train, to see this simulated war room where psssht—the missile would come up and bang, knock down the enemy missile. One of my colleagues on the committee was Jim Thornton of CDC. But one of the things that struck me, it seemed to us we were existing to make it look like we were continuing to answer this question about the viability of the system, when it fact we weren't. So once I saw that, I decided we would close the committee down, so we did close it down. We wrote this report and the report, which was secret of course, would criticize the army, and criticize it for—it's probably still very secret, but basically it criticized it because of the waste, the implied waste, that much of this could be accomplished with commercial equipment, and they had developed a complete, unique strategy. They had done everything from scratch. The circuits were different. All of the approaches, the language they used, were different. Everything that you could imagine about the approach to computing, to develop a machine, were done from scratch. And it was done in those days by Bell Labs and built by Western Electric, their subsidiary. So we criticized it for that. Of course, the army refused to accept the report. The National Academy pointed out that, "This is secret. Only you guys can read it." The army replied: "We don't like what you're saying." So eventually, they had to accept the report, but it was interesting. Dealing with especially the defense establishment, you can say what you want sometimes, but it's not going to change anything. It certainly in my mind gave a lot of credence to Eisenhower's remarks about the relationship between the military and the commercial arena is a delicate one. I think we're going to see that issue again in years ahead, dealing with that issue.

Zeidman: Have you been involved with the government since then?

Flynn: Not particularly. No, in fact, I think I lost a lot of flavor.

Zeidman: They didn't invite you back. <laughs>

Flynn: Yeah, they didn't invite me back.

Zeidman: Just to clarify, it sounds like you were criticizing the techniques, but not necessarily the technology? Or did you also—

Flynn: No, I was criticizing the expenditure. Really, it comes down to the expenditure. Why do all of this? In fair play, later the notion of COTS [commercial off-the-shelf], easy, commercially available, became very fashionable in DoD [United States Department of Defense]. But in these days, it was anything but that. This is what we were criticizing, basically: this need to do everything from scratch, everything, and have no standardization. There were a number of reports, which are available online now, which review the history of the Safeguard project as it was called. Just to give you a [hint of the project’s] breadth, I think there was something like the equivalent of 5,000 people working for five years on the programming and the development of this system, basically. An enormous amount of money went into this effort and some of it had to be, but a lot of it, I thought, was wasted.

Zeidman: I know when I was at Stanford, there was a lot of controversy about whether an antimissile system would even technically be feasible. I don't know if you weighed in on that controversy.

Flynn: No, I didn't. They had some successes. They built three of these systems and one of them was installed in North Dakota, which was the operational one, protecting the missiles there. One of them was in Colorado and one of them was in Meck Island [near Kwajalein], in Micronesia. They did some tests in Micronesia and they claimed to have some success. Sometimes it would work, sometimes it wouldn't.

Zeidman: I just wanted to leave with a final question about what advice you would give to someone who was getting into the field of computers now. What do you think they should do? Should they do that? Should they even get into that field, and if so, how should they do that?

Flynn: You know, computing is amplification of human intelligence, among other things. In my mind, certainly, the human intelligence could use a lot of amplification. Insofar as a field with growth possibilities, I'm sure it's there. I believe it will become, over the years, increasingly important. But just as we've seen with the movement towards social interactions, basically using a computer for social media, sometimes it produces some side effects which are totally unpredictable and unforeseen. But I guess if I were giving advice, I still think that the notion that you can provide some enhancement to the ability for the human species, to get its act together in better ways, to resolve our situation on the planet and to improve interactions among our societies through the computer and communications among computers; this is the amplifier for intelligence.

Zeidman: Thank you very much. It's been a pleasure interviewing you.

Flynn: Thank you, Bob.

END OF INTERVIEW

