Oral History of Les Davis

Interviewed by: Dag Spicer

Recorded: May 3, 2010
Chippewa Falls, Wisconsin

CHM Reference number: X5796.2010

Dag Spicer: We're here today with Mr. Les Davis. it's May the third, 2010, and Les is one of the eminent computer designers and pioneers in the field who worked at Cray Research, and welcome today, Les, thank you for speaking with us.

Les Davis: Glad to be here.

Spicer: I wanted to start with your life early on. Tell us a bit about where you were born and what your parents did for a living.

Davis: I was born in Minneapolis, my mother was actually an immigrant from-- a German immigrant from Russia, and she and my dad met in Minneapolis, were married. They later separated, and so I was raised by my mother. I went to school in Minneapolis, to a small parochial grade school and then to a Catholic high school, DeLaSalle High School, which is located in Minneapolis. From there, then, after graduating from high school, I struggled to try and find some work at that time, work was a little hard to find. And then shortly thereafter Uncle Sam took care of me by saying “Either come voluntarily or we will come and get you.” So I joined the Navy in 1951 and was in the Navy through 1955. During that time I was fortunate to have been able to go to a electronics school, and basically my training in electronics developed there.

After being discharged from the Navy, I went back to Minneapolis and started going back to the University of Minnesota where I had gone for a short period of time before going into the service. But I needed some money, so I was looking around for work and I found an ad, small ad in the paper for technicians to work in and on computers. It was from a company called Engineering Research Associates. So after an interview there, I liked the environment and went to work.

One of the first projects I worked on was a project headed by an engineer by the name of Seymour Cray. And from there then I was able to take in and start learning about computers, which I really had no basic knowledge of. And that was kind of the beginning of how I got into the computer field.

Spicer: What was it about ERA’s work environment that you liked?

Davis: It was-- the building was a huge glider factory, and it was very non-structured and it just had this appeal that it looked like there was a lot of exciting work going on, the people that were working there were all just generally excited about the work they were doing, and so I-- it just attracted me and I felt that was where I wanted to work. And so I was fortunate to have been able to have a job there.

Spicer: Did you have any mentors at ERA? People who took you along and taught you about computers?

Davis: Certainly, Jim Thornton was one of the big mentors in my time at ERA. And projects that I worked for under him were really very exciting projects.

Spicer: Did you work on any military projects, say for the Navy or--?

Davis: The projects that we had was mostly government projects. One of the projects was called NTDS, Naval Tactical Defense [Data] System. And Jim Thornton was the chief engineer on that project, reporting to Seymour Cray. And yes, that was a military project that we worked on.

Spicer: And did you participate in the 1101?

Davis: No, that was just before my time.

Spicer: Oh, okay, good. So that was your first job, and what did you do there?

Davis: I was a technician doing really engineering work. At that time engineers and technicians really were all struggling to learn about computers. So for those that had come out of school, they really hadn't had a lot of training on computers. But we had to do many of the things, soldering components on the boards, testing boards that had been put together, and generally trying to make work things that the engineers had designed.

Spicer: And tell us a bit about the founding of Control Data and the UNIVAC, Sperry Rand acquisition and how that affected you.

Davis: Well, shortly after I started with Engineering Research Associates, it was acquired by UNIVAC. And I didn't sense many of the things that the upper management of ERA did, this conflict between the East Coast and the Midwest in management. But it was fairly obvious after a while that there was something taking place there. And then when Bill Norris left and Seymour Cray left with him, it obviously left wide open a lot of speculation as to "What will this new company, Control Data, be like?" Certainly the fact that a number of the top engineers that were at quote UNIVAC slash ERA had left, made that quite an appealing place maybe to work. So I did not go with the first wave, I left I believe in 1959 or '58 to join Control Data. And I was able then to work on the first1604s that was being put together.

Spicer: So you were not in that--?

<crew talk>

Spicer: Okay, so just before the break there we were discussing you just joining Control Data. And can you tell us what the state of computers was at the time? Did many people know what they were or how they could be used?

Davis: I think the general public didn't have a good appreciation for it, but certainly the military did. And that was-- and the scientific community--they understood the value of computers and they were the pushers, you might say, of the need for computing power. But for the average person, I don't think the computer was of much significance at the time.

Spicer: So you began at Control Data in about 1959 I think?

Davis: Yes.

Spicer: And so what did you do there?

Davis: I was hired as an Assistant Engineer at Control Data. And the first project that I worked on was Serial #1,1604 and there I was involved in debugging the computer, and basically checking out the design of the computer, and got a little bit of exposure to design on some of the what we call the functional units at that time.

Spicer: So you were doing designing as well?

Davis: Well, it was pretty low level in comparison to what a lot of people were doing, but it gave me the exposure to that.

Spicer: What were the most serious problems facing someone who wanted to build-- a company that wanted to build a computer? Was it the financing or getting customers, the technical challenges?

Davis: It probably was all of those, certainly finances, you needed money to acquire the parts and the engineers to do the design work. You needed customers to support or buy the machines. And I think just having the knowledge of what it is to design to meet the needs of the customer. In this case it was a scientific customer. And that’s the case where Seymour really shined, he understood what it took to solve mathematical problems and he was able to take and play that into designs just very effectively.

Spicer: You and I were looking at some of his notebooks a bit earlier, and you pointed out to me that really the whole design of the computer is expressed in terms of Boolean logic equations. There are no electrical schematics or even rudimentary block diagrams perhaps, but barely, can you comment a bit about that and what that was like for you coming in for the first time and seeing that approach?

Davis: Well, it was certainly different in the sense of not having block diagrams or schematics as I was accustomed to. But as we started working with it, you found the simplicity of the thing and the effectiveness of being able to express functions that a computer would do in what amount to mathematical terms. And it turned out that people did adapt to it rather quickly and found it very effective. And I thought in many cases [it] eliminated a lot of the errors that you might have had if you had gone through the process of drawing schematics, not ending up getting lines in the wrong places. Once you had written it as an equation, unless you had made a mistake with the equation, it was pretty clear as to where the input would go or where an output would go as well.

Spicer: What other projects did you work on at Control Data?

Davis: The 6600 project, the 7600 project and then the 8600 project that never quite made it to manufacturing.

Spicer: Can you tell us a bit about the 8600 because as you pointed out, it didn't ship as a product, but it I think planted the seeds for later designs, is that right? Maybe for the Cray 1-to some extent?

Davis: Well, the 8600 still used what we called discrete components as opposed to having integrated circuits. And as a result of that, it requires very large numbers of discreet parts. But in order to achieve the speeds that Seymour was striving for, we had to take and push the parts very close together. And that raised two issues. One is you had a large number of parts but you had a lot of heat generated in a small area, and that put quite a demand on the cooling engineers to take and figure out how to take and remove the heat from the components so that they would operate within their operating temperature range. And the problem with the 8600 was there were so many parts, and the discrete parts were not able to be cooled effectively. And as a result of it, the failure rate was just too high to be able to manufacture the machine and deliver to customers. So it did lay the basis for and revealed a very strong shortcoming that with discrete parts, unless you can lower the heat requirement of the parts, you're not going to make it. And in any event, you probably need a very effective cooling system to cool these parts. So those are things I think that were carried forward in Seymour's mind to the Cray 1.

Spicer: Did he-- I think there was a refrigeration engineer named Dean Roush?

Davis: Dean Roush.

Spicer: Dean Roush.

Davis: Dean Roush is one of those unsung heroes, I mean, he was just marvelous and a genius in his own right in coming up with cooling techniques. So a lot of the work that was done on the 6600, the 7600, the 8600 and then the Cray 1 were Dean's-- were his work. And he worked in conjunction with Seymour and between the two of them they came up with very unique schemes. But it was Dean who was the one that was able to implement the thought process or the idea and actually make it a manufacturable product. But Dean is one of the people that just never got the recognition he should have.

Spicer: And then without his efforts, the machine would've just melted or overheated?

Davis: They would-- yeah, the failure rate would've been so high that you couldn't have made it work.

Spicer: How significant was the 1604 to launching Control Data?

Davis: I think it was the thing that was the product that made Control Data, because at the time it was the fastest scientific computer that was available on kind of like a commercial basis. It wasn't built to order as a one-of-a-kind. And it's really the product that launched Control Data.

Spicer: Did you get any feedback through colleagues maybe at other companies or-- on the reaction of other computer companies to what you were doing? Say IBM and what they thought of your efforts, and I think they were quite surprised that you could do so much with so few people.

Davis: Yeah. Well, the classic one I believe is where IBM commented on how we were you able to develop the 6600 with so few people when we [IBM] had this very huge project in our own right and were not able to make it a successful project. So it got that kind of feedback, but most of us who worked for Seymour did not have intimate contact with other engineers at other companies.

Spicer: Can you tell us-- this brings us around to Seymour's-- or the engineering style of Seymour Cray and eventually Cray Research, which carried through to Cray Research, how would you describe that method of working?

Davis: Probably unique, small numbers of people, very little bureaucracy. Seymour liked to work by himself, but he also realized that he needed help in putting things together. But he pretty much pioneered how he envisioned a module should be built and then he would pass that on to the engineers to fully implement and get into development. But he quite often would come in at night because he liked the quietness of the evening hours, and then come in during the day and go over the results of the day's work. But not work with a large number of people.

Spicer: And I want to mention a quote, which I hope won't embarrass you, but I'm sure you're familiar with it, which is "What Seymour threw out the window, Les Davis caught." Have you heard that one?

Davis: A little different version of it, but I have.

Spicer: Okay, what's your version?

Davis: Well, just the fact that I would take and be able to build things that he had designed and that-- so there's some truth in that, but it's not totally true.

Spicer: I think you're being very modest, but we'll leave it at that. What was the most challenging part of working let's say, on the 6600, which was a really landmark product?

Davis: Just being able to achieve the speeds that was felt necessary to make it a competitive product. And the first attempt at it failed. It's when we were in Minneapolis, and usually what you did, it was kind of evolutionary, where you would take the previous generation of product and enhance it somewhat. And the first attempt was using a modified 1604 module. And as we got into building it, we found that that module did not perform, and part of the problem was that the speeds that we were trying to attain, you had to treat wires as transmission lines, as opposed to just treating them as a wire. And that wasn't recognized at the time. And the 1604 module was not able to accommodate transmission lines. At that point is where Seymour kind of threw everything out the window, and I didn't catch anything that time, and started with a new module that basically was more dense in assembly, what they called a "cordwood" package. And it also then treated the wires as transmission lines and took the speed on the wire into account and actually used that. And it was about at that time then when the project moved from Minneapolis to Chippewa Falls. But there was quite a severe setback in that first effort to get the 6600 built.

So moving to Chippewa I think was for two reasons, one was get away from a little bit of the interruption you would have from marketing types, and I think Seymour felt from management types, and put it in a more isolated environment. And so then we moved to Chippewa and we were able to complete the project with a new design and deliver it to Livermore.

Spicer: This is the 6600?

Davis: The 6600, yes.

Spicer: And how was the checkout process with Livermore? Was it accepted relatively easily or were there some back and forth and debugging that needed to be done?

Davis: Well, with the early machines it was a rather unique way, and I'm sure by today's standard it looked pretty clumsy, but we depended very heavily on the customer codes to help detect any problems with a machine. We developed a set of diagnostics generally that were not as stringent as the customer codes ended up being. So in the case of Serial #1 6600 we actually had programmers from Livermore living in-house for a number of months running an unclassified version of their code to make sure we could get it operating. And then once we got to a certain point where they felt it was acceptable, we took it and delivered the machine. And then we had to have basically on-site customer engineers 24/7 to keep the machine running and that. But nevertheless it was-- despite the problems that we encountered, it was very successful in being able to solve or meet the applications that the customer wanted, in this case it was Livermore.

Spicer: And it became mass produced, by supercomputer terms, I think over 100 were made.

Davis: When you include all of the variants of the 6600 that we had, yeah.

Spicer: And then so as production went on, I'm sure reliability improved and testing improved and so on.

Davis: It was a learning curve as you built it, and we didn't have the luxury of building a machine and then running it on the floor for a couple of years to find all the problems. And the customers really were demanding in many cases and were willing to accept this risk of some unreliability for getting some very high performance in say, a fairly large period of the operation of the machine.

Spicer: Let's move onto Cray Research, what was it like when you started-- tell us about how you got there and what it was like when it was just starting?

Davis: Basically Seymour and a number of his friends from Control Data felt that it would be appropriate because of the reluctance of Control Data to support the high-end machines. And Seymour had one goal in mind, to build the fastest machines. And I think there was a reluctance on Control Data's part at that time to fully support a project, in this case a project beyond or even the 8600 project because of the money it required. And so Seymour with some of his friends was able to raise enough capital and go off on his own to take and basically build fast machines. The environment was really extremely exciting, also kind of scary because there was very little financial support in terms of what you really needed to do it. And the numbers of people that we had were very small, so-- but as far as working on the project and the people that were involved, it was a tremendous opportunity and extremely exciting.

Spicer: So I think Cray Research was founded in '72 and the Cray-1 came out in '76 so there was a good period of four years or so when there was no income and fairly heavy expenditures.

Davis: Yes. So there were periods of "Will we get a paycheck, will we not?" And again, some of the early investors, the Frank Mullaneys, the George Hansons, Noel Stone, people like that, they were just tremendously supportive of the company. And I think it's thanks to their effort that we were able to stay in business long enough to build a machine. But we were fortunate in one other respect, from the technical side, and that was that the industry was producing integrated circuits, very un-integrated compared to today's standard, but two gates per package. And the power densities were much lower than the discrete parts. And then at the same time there was a memory chip that was based on integrated circuits, so you didn't have to rely on core memories. And so it's kind of like the perfect storm where certain things came together at just the right time to help build a machine. And it was Seymour's design in that it just all kind of meshed together very well.

Spicer: How easy was it to work with him sort of one-on-one? Did you have long walks in the fields talking about ideas or was it pretty straightforward?

Davis: We got along I think very well. I understood that he needed his time and that he really was the one that had the ideas. So we were able I think to exchange ideas with very little interaction. But in the evenings he would come in and do his work, and then in the morning come in and maybe spend a half hour or 45 minutes reviewing what has to be done and what was done. And then I'd kind of go my way, and he'd go his way and do his thing. I got along extremely well with him, even on a social basis and that, but no, there were none of these long talks and stuff like that, no.

Spicer: So let's start with the Cray-1 now. And what were the goals of the Cray-1? What was the idea behind the design?

Davis: Well, the goals were to build something significantly faster than what was on the market. At that time, the 6600 was the-- or the 7600 was the machine to beat. And by the same token, Control Data at that time had a vector machine that was called a "Star" machine. So those were the-- kind of the goals to take and get performance levels that would exceed those machines. And I think IBM at the same time had Stretch [IBM 7030] and other perhaps machines and designs, but again, the goal was to be significantly faster than any of these machines and be able to make it a manufacturable machine. And that’s where I think Seymour had in his mind the value of vector processing, but he wanted to implement it differently than Control Data had or other companies had at that time. And as opposed to using their approach, he took this approach of-- where you had register-to-register operations and you could operate with very short vectors instead of having it to be a very long stream vector. And those are the things-- and it kept the scalar performance, which was important at the time as well. So he married those together and that’s-- along with the clock speed improvements, we were able to make-- made a significant performance improvement over any other product on the market.

Spicer: And I know the word “balance” meant a lot to Seymour Cray in terms of system design. Can you tell us a bit about that?

Davis: I'm sorry, the--?

Spicer: The word “balance.”

Davis: Balance.

Spicer: Balanced system designs, can you tell us about that? And what things are you balancing, why is it important?

Davis: Well, scalar performance was extremely important because that-- so many of the codes had been written at that time. I think Seymour was an excellent programmer so he understood what programmers could and could not do. And I think in that case, he realized the value of being able to have a vector capability that would be compatible with scalar capability and not require extremely long vectors in order to take advantage of the vectors. Because so many of the codes could not implement very long vectors. So the Cray-1 with it's design allowed for the use of good scalar performance as well as vector performance, either very long vector performance or short vector performance. And I think it's that difference that our competitors did not achieve.

Spicer: How was the Cray-1 greeted when it was first announced and shown to the world?

Davis: I think there was-- as I kind of remember, it had a lot of hoopla about it. It certainly had a neat appearance, there was a lot made of the cushions and the circular design, the compactness of this big machine compared to what other supercomputers had at the time, where they would take up a roomfull of cabinets. So I think it was meant with a lot of enthusiasm, especially in the scientific area. And then once they saw performance numbers based on some of their codes, it-- they were quite excited about that.

Spicer: Well, tell us about the shape and why it's important that it is the shape that it's in.

Davis: The idea was to-- again, communicating from module to module in a machine was one of the things that actually was a detriment or slowing the systems down. So ideally if you could take-- and either put all the parts on one board like they are doing today with integrated circuits, that would be the best way. But that was not achievable at the time of the Cray-1. So we ended up requiring modules, and in order to get these modules to communicate to one another, the wire lengths became extremely important. And so a circular design ideally gives you the shortest space from a module to a module. And there was some hope at one point that we could have enclosed even the open area and then have had the wires running across the boundaries as opposed to being routed around the circular part. But the whole idea was to minimize wire length.

Spicer: So the fact that it's also a kind of industrial sculpture is sort of incidental?

Davis: Incidental, but I think on purpose as well, and that again was one of the attributes that I think Seymour and Dean had is they could make something industrial but also make it attractive. And so the treatment that was made of covering up the front parts of the machine and putting leather on the base where the power supplies were located, those were not just incidental, that was where the artistic part came in.

Spicer: And was it felt that the machines-- I think Seymour Cray at one time said these are very expensive machines and they should sort of look the part in a way, they should look powerful and convey a sense of power and capability.

Davis: Yeah, I think that's a true statement.

Spicer: How about building a Cray-1, tell us a bit about that, how long it took and how many people were involved and what they did.

Davis: Well, there are two parts of it. One was where you had to take and put the mainframe together, assemble all the cold bars or the chassis parts, ordering the parts for those. Then the other was building the modules. Once you had the mainframe assembled and all of the module connectors in place, then it required probably a good three months of wiring time, in which ladies would take and route the wires from one module to another module. And if you look at the back side of a Cray-1 you can see how dense that wiring is. And it's amazing that these ladies were able to do the work that they did as well as they were able to. But that period took about three months. At the same time, we would build the modules so that those two efforts were going on in parallel. And then there was anywhere from a two to three month period of--

<crew talk>

Spicer: Okay, sorry about that. We're back from our break. And we were talking about the process of building a Cray-1 and what was involved. And I think we had got to the point where the assembler women were hand-wiring the backplane of the Cray-1. So where would we go from that stage?

Davis: After the wiring was completed, then the process would be to plug in the modules and start the checkout period, which amounted to running diagnostic codes that our engineers had developed. And in many cases, this included segments from customer codes that they would allow us to use as diagnostics. And we would go through that process, and net: the total could vary I'd say anywhere from six to nine months, depending on how far in advance you could have ordered parts like integrated circuits, mechanical parts of that sort. And then once the machine had passed these diagnostics that we had developed, then quite often a customer would come in, and he in turn would take and run codes that he had developed to make sure that he felt comfortable that it was ready to ship.

Spicer: And how many Cray-1s were built do you think?

Davis: I've lost track of that number.

Spicer: Would it be 100?

Davis: That’s one you'll have to look in the record books on. There was-- it's a question of the Cray-1 and the variants of it, the X-MP, but the Cray-1 I would say in the fifties at least, but I'm not sure of that number.

Spicer: Now, I think at about this time or a few years later, Steve Chen was brought on board. And how were the relations between the two design groups?

Davis: Well, at that time Seymour had a design group and then Steve Chen’s design group reported to myself, Seymour’s group reported to him. And we basically had what I would call a “skunkworks” project, and that was the Cray-2, the company was betting on the Cray-2 at that time to be the next system. We went ahead and we took the Cray-1 and made some modifications to the circuits, various minor-- we improved the clock speed somewhat, we improved the packing density by packaging the module slightly different. But it was basically what we call today a skunkworks project. And Steve was the lead engineer on that project, and as a result of that, we were able to take and meet a very short time schedule. And in the meantime the Cray-2 had slipped it's schedule so that the Cray X-MP ended up kind of fitting in very nicely from a corporate standpoint of having a new product available in a fairly timely manner, even though it didn’t have all the exotic capabilities of the Cray-2. But performance-wise it worked very well.

Spicer: And was it the two groups worked in isolation I guess?

Davis: Yes, yeah, there was no interaction between them other than some of the hallway chatter, but we had our separate buildings and facilities as well.

Spicer: So how did-- your role must have changed then from being Seymour's right-hand man to having management responsibilities of your own. Did that happen?

Davis: Yeah, there were some changes there in the sense that had at that time manufacturing, all of the engineering groups, with the exception of the Cray-2. And I did take and maintain support for the Cray-2, but not on quite as close a level as I had with the Cray-1. And at that point, the goal was, first of all, make sure the Cray-1s out in the field were operating successfully. And then I had a follow-on project, and I kind of, like I say, felt that that was, in my view, an important thing to have a backup to the Cray-2. So yeah, my goals changed a little bit, but nevertheless I kept very close communication with Seymour. But at the time, he really didn't have the need for-- they were still developing the immersion cooling, and so we had some splits there with the engineering folks. But it was a little different role.

Spicer: And the Cray-2 as it turned out as maybe not as successful as Cray had hoped, is that a fair statement?

Davis: I think that might be true. I believe it met it's goals but it took longer than expected. And the costs were probably a little bit higher, but for certain applications it was *the* machine to have. But for the general Cray-1 user for a follow-on machine, the X-MP was perhaps a easier fit because there were no software changes required, you could basically move your code from the Cray-1 to the X-MP with a little effort.

Spicer: As Cray Research grew in size--it went from I think 100 people, not even, to over 5,000 at it's peak--how did the company change? You were there at the beginning and then you were there at the end. So how did you see that change? I think earlier you mentioned it was like the lumber industry here in Chippewa Falls or something.

Davis: Yeah, well, there was a lot of rapid growth. And we ended up spending a lot of time building buildings as well as building computers, but you had to pay attention perhaps more-- or a lot more than you would had liked to, just to some of the administrative type of thing. But the other exciting thing was is the number of people that we attracted, talented people. We had engineers, programmers, IC designers, it was really very exciting to watch and see that taking place in a small community like Chippewa Falls, it's something I would have never have imagined. And as I stated before, we ended up being pretty much vertically integrated where we built our own cooling bars or frames, integrated circuits, our printed circuit boards, doing all of-- a good share of it internally. And so that part was really quite exciting, being able to accomplish that. But it also had perhaps it's share of frustrations.

Spicer: Now the reason that you became vertically integrated was of necessity as much as out of corporate desire? Is that right?

Davis: That's right. We simply weren't able to buy the parts anyplace else in the world. We had tried that and not succeeded, so we ended up from a timing standpoint… because as you were building machines, there was a pressing need almost every four years to come up with another generation of machines in order to accomplish that in a timely manner, it almost forced us to do the work internally.

Spicer: And how did Chippewa Falls change with the arrival of Cray and 5,000 more people to the area?

Davis: I think it handled it pretty well, it certainly made changes. We ended up with large numbers of people with different backgrounds, [different] than Chippewa perhaps had experienced before, highly skilled people. On the other hand I think many of the people that worked here had an opportunity that was kind of a once in a lifetime opportunity, because it was really the thing to do if you could is to work for Cray Research during those days. And we made a lot of money, profit sharing was very good. And I think it spilled over into the community that it was extremely helpful to the community.

Spicer: Is there anything else you want to say about working at Cray Research?

Davis: It was a once in a lifetime experience. I never thought that I'd have an opportunity like that. But even as some of us get together nowadays and we talk about it, we say “Boy, those were really great days, weren’t they”

Spicer: That’s wonderful. One more question, which is the Seymour stories that we hear, and I wanted to get your feedback just quickly on whether these are true or false. There are just a couple of them I have. One of them is that he built a lot of tunnels?

Davis: He built some tunnels. He had a small home or cottage on the lake close by Lake Wissota. And he would take in-- I think just to get some physical exercise and that and to take and kind of test himself, he would actually dig some tunnels. He would get support lumber and build a real tunnel. But it was not with the idea of doing something unique. I think it was as much to get some physical exercise and to give himself some time away from his normal routine. So yeah, he did do that. But he didn’t build tunnels, but it was a couple, maybe a tunnel that he had built there on his lakeshore.

Davis: The other one is the story of his sailboat, which he built a new sailboat every year and then burned it down in the fall.

Davis: He'd build some new sailboats, but I think the reason he burned them, he didn’t think anybody wanted them and wanted a way to get rid of them. But yeah, he would-- again, it's kind of his perfec-- he wanted to perfect and build another one that was better than the previous one he had built. So it follows along with his doing things well.

Spicer: And starting with a clean sheet of paper, I guess.

Davis: Always a clean sheet of paper, he's not one for clutter.

Spicer: Was that a problem for customers who wanted compatibility among generations of machines? Was that difficult to do that with a new architecture every time?

Davis: Sometimes it was, but again, if the-- it was often a question of the only way to get a performance difference is by having something quite different. But probably one of the classic examples there was the X-MP versus the Cray-2, where the architectural differences perhaps were a little bit of a stumbling block for some customers in moving to the Cray-2, where the X-MP was a much smoother path for them to follow. But in many cases, like moving from a 6600 to a Cray-1, the only way you could achieve the big difference was by introducing some architectural differences. But again, Seymour was very clever there, he kept the scalar performance of the 6600s, 7600s while introducing vector processing. So he never quite abandoned the customers completely and he kept something open for them.

Spicer: And supercomputing itself as a field was pretty much defined by Cray, would you agree?

Davis: I'd say--

Spicer: Cray Research?

Davis: it had a big impact on supercomputing, yes.

Spicer: And then part of the explosive growth of Cray was probably the market itself was expanding, the applications for supercomputers.

Davis: Yeah, as soon as the applications were developed for the Cray, it made a significant difference then, because it opened up the aircraft industry as-- we were able to get the auto industry, so as opposed to just having the military involvement, we were now able to get to oil companies and many commercial customers where they were able to use the machine either in their design or exploration and were able to do experimentation without having to do the actual work itself. So you could simulate wind tunnels, you could simulate drilling, you could simulate car design. And a simulation was the big thing, but we needed applications, and we were able fortunately to attract a large number of software developers to write applications for the Cray.

Spicer: If a youngster came up to you and asked you what supercomputing is and what is it good for, what would you tell them?

Davis: It'd be hard to impress the kids today because they have so much computing capability available to them, whether little iPods or little desktop computers. So I'd say that it would be hard to explain, but you'd have to go back historically to look at what supercomputing did do to kind of push ahead the development of science.

Spicer: A lot of people speak of a “third paradigm” in science, which is-- the first two are Theory and Experiment and the third is Modeling or Simulation, which is really come into it's own because of supercomputers, the ability to model nature and ask questions that you couldn’t do otherwise.

Davis: Yeah, and that's-- yeah, I think that was the key, especially during the nineties and that-- being able to do the experimentation without actually having to physically do it.

Spicer: I mean, that's really kind of magic in a way.

Davis: Yeah, it is.

Spicer: Simulating a wind tunnel inside of a machine or a weather system or-- as we sit around in this museum here and you're surrounded by your old friends, how do you feel? A little nostalgic or proud ...

Davis: I'm proud of that Cray had done and I think what I found especially satisfying is the number of people that we had and their involvement in developing the products. So Seymour Cray was our guide, but we had so many of these other workers that saw to the fact, whether they were in sales, whether they were test engineers, whether they were IC engineers. But it takes a lot of people to build a company that Cray Research was, and hopefully the remnants of Cray Research will continue to be successful.

Spicer: Just finally, when Seymour Cray passed away, how did that affect the company and did it really lose it's soul or--?

Davis: Well, Seymour had kind of moved away.

Spicer: He'd left beforehand?

Davis: Yeah, he'd moved out to Boulder. And I think the Cray-3 project had kind of separated the two companies. But I think so many people back here really felt that something was lost and probably, like you indicated, the heart and soul was lost. Because I don't think anybody ever-- even though Seymour did leave, they kind of felt he was still around.

Spicer: Perfect. Thank you, Les. That's a wonderful interview.

Q: Mike [ph?], I have this one little thing, if you wouldn't mind.

Spicer: Yeah?

Q: You had talked about the kids today and the iPods and the desktops and things like that, a lot of the Hollywood magic that we see today is really the result of the supercomputer and the animation aspects that you guys started. Actually The Last Starfighter I believe was done with the Cray-1-S.

Davis: Yeah, there was an interesting period there when Apple had purchased a Cray to do their design. And at the same time Seymour was using an Apple computer to do his design. So that kind of-- interesting different of application. You know, here you're using a big machine to design a little machine and you have someone else like Seymour using a little machine to design a big machine.

Spicer: That's great.

Davis: It's interesting.

Spicer: Do you know how he used his Mac by the way? How he used his Apple?

Davis: I think mostly for-- he would take and write your equations for the modules and that. So instead of having to necessarily write it out in pencil and paper, you could do it on a machine. And the other thing is that Seymour was a tremendous writer. He wrote the first instruction manuals for the Cray-1. And he was just terribly picky as to how they should be written. He was not happy with some of the tech writers and he was getting so frustrated because he [had to] explain things. So he ended up writing his own reference manual, and he spent maybe six months-- so it was interesting to watch the man work. He'd be idle for a while doing things in his mind, developing the thought processes, at least this is the way I saw it, and then he'd start implementing it, putting it on paper. And then actually while it was-- and developing the circuits and that, and while it was being built, in the case of some of the machines, he was actually writing software. And he had written operating systems and a Fortran compiler. And then he has this-- what he might say "Why do you want to do this, this job of writing a reference manual", which was very good. So we spent many hours trying to get the right-- when he was using a typewriter, the-- I forget what you call those balls on the type--

Spicer: They call it the typeball.

Davis: Oh yeah, I mean, he just spent many hours getting the exact print quality that he wanted. So he was a perfectionist, it kind of carried across his whole lifespan. So interesting.

Spicer: That's great.

END OF INTERVIEW

