

INT: I am real interested in seeing what your perspectives of what Digital's contribution might be in information technologies and what sort of role you played in anecdotes from what you worked on in creating that reality. What was going on in your head?

CT: Several things, I had spent thirteen years at Xerox and it was clear that there was little of the stuff that we had done was going to have any substantial impact on Xerox's position in the financial business or bottom line or whatever. It seemed like it might be a good idea to see if Digital wanted a computer science laboratory. At that time DEC didn't do much research, there was only, there was a corporate research laboratory which had been successively managed by several people.

INT: [INAUDIBLE] it was very much development ---

CT: It was development and actually a fairly deep --- and Digital was doing I guess about 4-5 Billion then and when you're in this business, or at least at that time, you are doing enough business that you believe that your major competition is going to be IBM and the Japanese, you have to look at how they are doing things differently than you are and one thing that was pretty clear was that they were doing a lot more both basic and applied research and so we had talked to several people, several

companies about starting a research laboratory in California and when we talked to Sam it actually didn't take too long to get the whole thing started once the basic decisions had been made. That was kind of nice because it indicated a sort of a way of doing business that was fairly unfamiliar to me. And that 3-4 people could get together and make a decision to basically change the direction of the company in terms of research.

INT: What an opportunity -----

CT: It's worked out that way. I think we've probably paid the bills since we've been here, there are still people in Digital who say that Digital shouldn't be doing research that it should strictly be applying technology but I don't think that's nearly as widespread as it was in '84. The research lab has had a lot to do with that change but we weren't the first to switch labs, the Western Research lab was the first one - that was started by a couple of people from Xerox who used to work for me. But I think that WR [INAUDIBLE] I never quit in the early years, had a much larger impact later.

INT: There were how many workstation projects _____ engineering development and they were more like a level track of working [INAUDIBLE]

CT: That isn't the case at all. In fact _____ and patent project was the thing that I guess sensitized the company in the allergy mean of sensitization. That is they built a system that was considerably faster than any machine that Digital either was making at the time or had on the drawing board. And at that time there was still a large controversy about whether these were good ideas or not and they were basically able to rub the companies nose in the technology. That was good because they got a lot of things started, it was bad because it was done in such a way that was perhaps not the most effective way they could have done it. But, I think one of the things we decided here as a conscious decision when we started _____ was essentially not to do CPU design, because there a large number of very competent, very talented CPU designers in Digital already mostly in Hudson and at MSP and it seemed like we would have very little [INAUDIBLE] at that and so what we decided instead to do was essentially to take advantage of the technology that was developed elsewhere in Digital and use it to build new interesting systems. Rather than glossy views (?)

INT: _____ perspective.

CT: Right and I think that actually worked out very well. A good decision and we did the right thing. I think one of the early research thrusts at that time was

how you could exploit multiprocessors. And the answer has been less satisfying than what I would like and it may well be that multiprocessors are just a bad idea but at least we were able to learn a lot about that and I think help the company and one of the things that I would characterize this place as being for is to help the company eliminate surprises and I think in that area we did quite well because we have done a lot of imitations of multiprocessor technology, we were able to put a large number of our systems in service eventually to use them. People could see how it worked out. The answer was not as well as you hoped, but that is the name of the game in the research biz. Some things don't work out as well as we'd like. I think in the software business we certainly had a similar effect, a lot of the techniques that are now becoming standard in what would be our new operating systems like the OSF-1 or ECE were highly developed here or heavily influenced by work that we'd done here about five years ago.

INT: You'd expect a research center not to expect concrete answers and yet the scores you had the contributions to OSF, to multiprocessing, the entire ability to put stuff in there, it's an interesting role.

CT: A lot of the people here certainly in the early days if less so now, came out of an environment where they had

worked very hard and very long and no effect whatsoever except on the field but not for their company and what we tried to do was get people who were interested in having an impact on what or how Digital did. So it's not surprising that people would try to change things in the directions they thought were appropriate. It's hard to put a dollar value on that, I still don't know whether we've actually paid the bills for this place, I think we have, but I couldn't show you a balance sheet to indicate that.

INT: If you put enough in and get new technology up to speed then some of those come in and you'll have a 3-5 year jump on the competition, if you implement them.

CT: We hope that that's true, we are beginning to see now some of that coming out and if the company is aggressive and vigorous to exploit technologies then ---

INT: Let me ask about one of the contributions to multiprocessing which I imagine needed to be investigated---

CT: Well, most of our early programs were very conscious, they were established to look at areas where digital wasn't working very much and that we knew a fair amount about. I can still remember one day when one or

two days, Bob and Butler and Sam and I sat around in the Comfort Inn and figured out what SRC was going to be doing over the next 2-3 years and that included the language development. The fact that we would build an operating system that was UNIX-like but supported some of these other things that a multiprocessor might be useful _____ those things all came out of specials.

INT: How did you go about perceiving?

CT: How did we choose new projects?

INT: Yes, using this as an example; first you'd need a multiprocessor workstation and you couldn't buy one so you made one.

CT: And similarly for the software. We couldn't really buy them unless you had strange modular pattern so we took one that had been developed and made modifications to it and add the features and language that we thought would be important for building larger-scale distributed systems and spent two years doing that work and to some extent we are still doing quite a bit of it. We're not building our own computing hardware anymore and the reason it that all computers now are built with the same chips and we can't do substantially better than Hudson can. And so, what we try to do is get, influence what

the product people are doing and make the product better but so they could be more useful for us in our research. So we will now use the product line. That was an opportunity that we took three years ago and that's not come to fruition yet, but we're hoping that it will. That was an area where we did the same sort of thing - we looked around at what Digital was doing in Network and didn't like it for several reasons. Primarily it didn't seem to be a good basis for building computer distributed systems because the program on there is procedure call and the latencies are [INAUDIBLE] We felt it might be hard to deal with, a high performance system running on a large operating machines so that coupled with the availability and the semiconductor market of some of the chips that we use which led us to start backward.

INT: Potential applications in the network.

CT: That's right, that's what we're are doing now. Our major highway project right now is to build Digital's version of that. If you work closely in interview to try to grease the skids for it's product ____ it was a little too late. The Autonet was a little too late. 100 megabit networks, well FED has added from before most things our customers deal with. If it's good enough to make clusters work it's good enough to build these resister clusters where you can put two machines, fifty

kilometers apart so we don't see the original Autonet going anywhere but we learned a great deal, primarily about the system aspects of the network. But I think, in fact, that that project is probably the best example that I know of of a project that is particularly good for a place like SRC to do. Because it has everything that, everybody can participate essentially. The theory group worked on the algorithms for configuring the network, the network itself is pretty significant distributed system but the pieces of the network talk to one another, exchange messages about the help through network and so on and do this reconfiguration thing which is the real strength of the network, basically hit it with an ax and it will continue to work. That's presumably what people want in modern networks and are not getting right now. They are not getting from anybody but us at the moment. And then we have the hardware people who actually build the thing and install it so we were able to measure it's performance then and find out in real operation how the thing works. So if we can find a few projects like that where there is a vacuum within DEC and a need in the world, and people here who have the necessary strengths to build the thing, that's a great opportunity.

INT: Tell me about what kind of linkages there are with engineering, it seems as if it's not formalized. It just happens like people like water going downhill falling

where they weave their tenuous way.

CT: What's true is that people, a lot of people here knew a lot of people in engineering, not a whole lot, but enough to make a difference. For instance, Doug Clark, who was in MSP and was the architect of the, forgotten the name of the machine and subsequently for the R-VAX, Argonaut. I have known Doug for a long time. He worked one summer at Xerox park so we already had a few people with whom we had worked in the past and generally not superdirectly but we knew them because we go to the same conferences and so forth. So there was a little bit of a catalyst there and I think quite a few people worked pretty hard to develop those relationships into ongoing things that would allow us to have access to that. It's a little hard when you're in California.

INT: On the other hand, I'm sure the isolation helps to some extent, you can use it when you want it.

CT: You can say, oh sorry I wasn't quite up to getting on an airplane this week, but a lot of people spend a fair amount of time on the East coast, I spend about a week a month.

INT: Do you find yourself under much pressure to put to work on more development projects?

CT: no. There are a lot of competent engineers. When I do this I refer to the role as being a technological troubadour. I am supposed to trot out alternatives and sing a song about how wonderful modern technology is and get people excited. But actually, I don't think I personally have ever directly worked on computer development projects. There are people who have, Gale Schwartz, spent quite a bit of time working on stuff

_____.

INT: You were trying to develop a useful personal computer in a Network distributed setting. Then you came up with the firefly for what purpose?

CT: The basic idea was if you could find a way to program these things well you would have a performance jump on people who didn't build systems that way, that you could maintain over time. If you could really make use of processors with a given level of technology you could be eight times as fast as without. Now, that's changed a bit in the last few years as the underlying silica technology has gotten so much faster and it's gotten faster, faster, doubling in performance every two years. Now the issues are a little different because things are changing fast enough with the low level components that the system platforms to make use of those

components have to change at a much higher rate than they used to. That makes it hard to maintain any kind of a lead. Without continually having to redesign for these newer and higher speed interfaces. As the chips provide. I think the six thousand series was an example of a machine that was very successful because the underlying platform that is the extra line, lived through several generations of processors. It's not at all clear that we'll be able to have that kind of good fortune in the future because it looks like now you don't want to build a platform that will support a factor of twenty in performance because the low performance systems pay extra cost for that so at a time when performance was going up fairly gradually, forty or fifty percent a year, rather than doubling, you could envision building systems that had a five year lifetime but we can't with product lifetime of two years and specific problem because the development cycles are longer than that. So we have to figure out what to do about that. And that's one of the big challenges for engineering and to the extent that we can figure it out and make that better we are working on it.

INT: So it's more effective now instead of to try to boost that technology curve even further, to rather pace the turnaround time between using the technology ----

CT: What you want to do it find out how to design things rapidly. That's the thing. Design and get it out to the market. And it's not likely that from one generation of technology to the next you'll be able to carry nearly as much of the rest of the system forward.

INT: So you have to redesign?

CT: A lot of it will be designed. So we are working here a lot harder on design tools in order to make that possible.

INT: How did you explore the ways to make it operational?

CT: Well, we built the hardware and the system to make use of it. The software system. This is a software problem, it's not a hardware problem at all. The problem is finding problems that you can decompose well enough that they can run together on multiple processors or, for some things multiprocessors are fairly good and it's not hard to figure how to use them. Timesharing for example, it does it's own processing and that can be quite a cost effective way to boost the performance of your timesharing system but we were more interested in things like we could make a window system fast, so Greg Nelson upstairs spent a lot of time working through various

alternatives of using Multithreaded implementation of window system. The kind of software I do is mostly CAD tools, but I build a simulator, for instance, that could get linear speed up on a multiprocessor. And I think for several years that occupied a lot of people trying to figure out how to use these new capabilities.

INT: {INAUDIBLE}

CT: RPC was indebted to Xerox. Bruce Nelson wrote a thesis and he did a lot of work as a summer student at Xerox and all the RPC _____ .

INT: Where did thread come from?

CT: Well we had threads at Cedar which was the programming _____, I am not sure whether, certainly I don't think there were any full-blown implementations of threads for that, there were some things that were sort of like threads, but the idea in those days was lightweight processors. For that part we were all _____ processors contract.

[END OF SIDE ONE]

[BEGIN SIDE TWO]

CT: Threads were a good program to turn on; they help you structure your program into cooperative activities and then if you have the multiprocessor hooked up so they run parallel that's just gravy.

INT: Did any of these ideas end up in Digital's massively parallel architecture?

CT: That's Mass publishing and those machines are quite different, they are single instruction multidata stream machines, quite a different idea rather than using multiple instruction, which is what the firefly is about.

INT: What was it like implementing that?

CT: Implementing the firefly? It was hard work, primarily because we didn't have very good computer tools to do it with, about half the work in building hardware is building the tools to build the hardware and if you design computers for a living, then the way you think about the computer is that it's going to be the thing that runs the tools that allow you to build the next one. The firefly was tough because it was '85, six years ago, it's a much more manual process than you use today and much less highly automated, less simulation, much more debugging. Much less predictability in terms of ultimate performance of the design or whether it ---

INT: How did the refinement of tools affect system design? Or the actual engineering of the design?

CT: I think they've made project outcomes much more predictable. You can know fairly early on in the course of the design that what you are doing is the right direction. The other thing is that having models in the form of software of the machine that you are building is a much better way of specifying that. There are more than one person working on the project it's great, nothing beats a program for explaining how something should work. No amount of English prose can say it as well and, in fact, in the new ___ we did for the Alpha demonstration, the PD bus, the funny little specification for that is a prose and the English prose part is ten pages, if you look a typical DEC product spec these days, and this is something I have been trying to change, it looks like this, and the problem is you can't run this and it's filled with errors and inconsistencies whereas, if you specify something by building a model you can execute it and find out how fast it will run, you can have a very precise, formal specification for it. maybe one of these days we will be able to prove that fact, hardware to be implemented is something that needs a specification. The technology isn't quite there yet but it is better than it was as little as five years ago.

INT: I also want to get into Autonet the system aspect of the network sounds interesting. I haven't read the paper yet.

CT: That was the goal and that was an interesting example of how a fine research place like this we go about things differently. We really look, we had a set of goals for the network that were quite a bit different than most of the people doing network "research" these days. We knew it had to be high band and it's got to have low range too. Everybody thinks that, but we also said we want monkeys to be able to install the thing. Plug it together in any way and it would still work. Providing everything is connected. And that was a reaction to problems that our customers have with these large networks that we install. You plug a new component into the network and the changes are very good that it will break because you have violated some rule and if you are a network manager you are very careful to change or expand your network because if you break it, you have several hundred users very angry at you. So we said, let's build a network that has a property that it will be very available, not necessarily reliable, customers don't give a hoot about reliability, what they care about is availability so we said let's not worry too much about building highly reliable hardware, but let's make the

system such that if that hardware fails the network can still function. I haven't seen that as an explicit goal in any of the networks that people have been trotting out over the last few years. That's one of the nice things about doing that type of work, you get to see the problem that real people have using systems.

INT: And the ability to standardize it eventually?

CT: Well, we hope so standardizing things these days is a little tough. I guess my own believe is that the whole standardization process is flawed these days. What tends to happen is that people standardize things they have never built and that they don't know will work. And, in fact, if you look at the SDI standards DEC worked very hard on that, but a tremendous amount of the work had to do with finding ways to make this standard actually function. And it did for quite a while. We did a lot of fairly heroic things that people in ___ work on that in order to make SDI a reality. The future bus [INAUDIBLE] is an example of that. Tremendously thick standard, nobody has ever implemented it. We are doing the first implementations. What you tend to get is people who like to go to standards committee meetings in nice places and they will sit on standards committees for their whole professional careers and that's why it's flawed. It's a good thing for companies to standardize things, but my

belief is that you should not be able to make a thing in Triple A standard if you can't show a working implementation of it. Then you can negotiate about changes to make it better or more universal or whatever you like, that's the way the Ethernet standard was put together. We had an implementation and had it for five years and what wound up being standardized was quite different. But, there was enough technical reality there that people didn't wander off into la-la land. Which is what tends to happen these days in standards committees, though I am, people believe that it is not possible to bring products to market in low level network space unless there are standards. I don't believe that there are a few people who would do it, we are one of them, IBM is another, but there was a time when Digital basically had to own the network business and I think it's quite possible that we could do that again. _____ has been displaying over the last four or five years.

INT: It's time to create that. Like cross-functional communications. You were talking about that, what are the interesting problems arising from systems research.

CT: Well, that's interesting. Some of the things that we are beginning to think about are questions like what is it going to be like when travel is no longer necessary. When you have a sufficient communication

infrastructure throughout the country or throughout the world that it will be possible to have desk to desk multipurpose teleconference and you won't have to go anywhere. That will be a big change. I am not sure what extent keratin ___ have on the upper atmosphere, but I am sure it can be positive. You can't do that right now because bandwidths are not available, but we can simulate it. And do it on a perhaps, modest scale.

INT: And find out what the problems are.

CT: Right, we plan to have desk to desk teleconferences with NCRA, we hope as far as Paris, but definitely between this place and Cambridge, within a year. You can do that on a T1 line we're actually beginning to talk with a set of people, including Sprint and the Minnesota Supercomputer Center and the University of Kansas to do another of these data base test networks. And our contribution to that would be automated installation at the University of Kansas. What we get for that is we get the print with 45 megabit line from Sprint for both talking to this database network in the midwest and experiments with it. That will make things quite different. I was talking to our systems manager today who actually needed to build a piece of software, bits and pieces of it were here and bit and pieces were in Cambridge, and he finally decided that it was going to be

too difficult to transfer the files so he just mounted the Cambridge filesystem on one of our systems and did the work over the net. We couldn't conceive of that a couple of years ago, but now we can and it works, so what are we going to do? Another question we are asking is, what will things be like when a place has one terabyte of information in the base. We have about 70 megabytes of spinning storage now. And it makes us careful about how we manage our storage. We spend a lot of time working on things like these, the VESTA system which you may have heard of, which is a revision control system for software for large software contracts. We think that's a very good idea. A major underlying idea is this repository that maintains the illusion that you keep everything and a lot of hard work and careful design investment went into making that repository, that illusion work efficiently. Perhaps if you could really [INAUDIBLE] presumably things would be a lot simpler. Another thing you might be able to do and we've been thinking of this trying to figure out whether it's time. Digital libraries. If you have high performance implementation then it seems perfectly sensible to have information in only one place. What would it look like if everybody in Digital has access to all of the Digital libraries over a network. So as a modest scale experiment we are beginning to think about whether we could take out library and digitize all of it. That's about 100

Kegabites of information, you can't quite do that today. But, there aren't that many orders of magnitude in between what's in our library and what's in the Library of Congress. So, it's pretty certain that within our lifetime we will have that kind of a system volume available. Big problems there are non-technical things like copyright laws.

INT: How would that affect apart from easy infringement of access--- or finding similar things ---

CT: A question that I want to be able to ask is, tell me all the paper that were published in any technical journal in the last two years that had the word local network there may be ten thousand and I can filter that down very rapidly and it would make a substantial change in the way we worked because it will now - subscribe to a few journals and more than the library subscribes to, but if you had them all, I think that would make a fundamental change in the way the scientific research is done. So those are some of the wild things we think about. We can't do any of it today, a big problem is storage. Surprising that it's not communications, except maybe when you think about the transcontinental aspects. But, with some of the new mega- optic storage devices now being developed a terabyte is only 100 10-megabyte discs. Maybe I'm a piker maybe I should be thinking of a hundred

kilobytes but I think it would be quite interesting. If we can figure out how to do some of these things. And provide good interfaces between human and information and one of the things we have now to whet our appetites about the system about NI, which basically has a full-word index of every message that has ever been sent to the bulletin board and every message that appears on the internet school for the past month. So it is a rolling thing, it can go into NI and say, tell me every message that referenced the word, radiation, and you get two hundred articles and you filter through, but it's current information. So if people could do that with their library it would be so much nicer. We do that filtering out like walking around, the fact that the Dewey decimal system works the way it does. Certainly a lot of times when that fails.

INT: I read an article about the databases and what's interesting about them.

CT: Right now databases are not nearly as useful as they could be, there are difficult to use except by professionals. I don't do my own literature searching, not like the librarians do, they're specialists. But, they have the potential to be a lot better. Particularly if you have a lot more storage and a lot more speed on the part of the thing.

INT: Interesting directions, keep following them where they may lead.

CT: We'll see if Digital can make a buck off it.

---They are number one in the whole world. Then there are a few people in the number two spot and we are in there with them. And the problem is that when you've got that you are committed, you've got this tiger by the tail, you have to maintain the investment strain to keep reaping the benefits and it is not clear that you can afford a major diversion of a few hundred million dollars a year that it takes to keep a place actually on the cutting edge of their technology, whatever it is, and still do the job. What we've seen is that investments in silicon can pay back better than investments in rocket science packaging. Every time I go to a meeting I have a slide that I can pull up and slap down on the view-graph at a moments-----

[END OF INTERVIEW]