

Begin interview with Bob Taylor in Systems Research Center in Palo Alto.

BL: --- briefly about the formation then of \_\_\_\_\_, how it began. You had an incomparable research organization, very hard to do and essentially you needed a sponsor. Would you have taken anyone?

BT: Well let's go back to, I didn't join Digital until late November. When I left Xerox in early September I had no idea what I was going to do next, literally. I violated advise that I had given to everyone all my life, never to leave a job until you have another job, but I had no other job lined up when I left Xerox. I thought, what I would do was to take a few months vacation. And then I'm going to go sign up for one of these programs that teaches you how to look for another job, because I hadn't looked for a job for a long time. So when I left the- word got out over the electronic mail network and I began to get phone calls from various companies, sometimes people I knew, sometimes not. And eventually I guess one of those came from Digital and Sam Fuller arranged to come out and talk to me and then Jack Smith and Sam Fuller came and talked with me. But I think that was probably begun in October, maybe not until middle or late October. Because it wasn't until late November that I decided what I was going to do. They first, Sam

anyway, first asked about would I be interested in coming back east. And I said, no I've lived out here for a long time and I think this is where I want to stay. So, then that's what led to the possibility of building a research center out here. But, at that point, with regard to the group that I built at Xerox, only Chuck Factor had left and he had gone to help us start-up and he was involved with that. So everyone else that I had brought together in the Xerox group were still there. Now then after I joined Digital, it became known that I was going to join Digital I guess. The head of corporate research in Xerox unwittingly did me a big favor. He sent electronic mail message out to all of his troops. Saying that he understands that Bob Taylor is going to be building a research group for Digital somewhere in this neighborhood and he wants to assure everyone that Xerox is fully committed to a long-term future of computer science research and that they're very serious about that. He obviously sent this message out as a way of trying to head people off from possibly leaving Xerox. But what he served to do in that case, was to let a lot of people who didn't know otherwise, because I hadn't told them. What it was that I was going to do. So I began to get phone calls from these folks and some other folks and about the possibility of joining Digital and I said, in every case I said, well, if you are interested in joining Digital what I think you should do is write a letter to

your Xerox management saying that you are interested in exploring opportunities elsewhere and that you want to let them know up front so that there is no surprise. And you should send a copy of that letter along with your resume to the following Digital personnel person in Greater Maynard and so that's what people began to do. And so the first, most of the cadre that started this place came from the more than a dozen best people out of the group I put together. So in December, I joined Digital in late November, and during December Sam Fuller, arranged for Digital real estate people come to Palo Alto and get in touch with me and to help me select a building a place to build this enterprise. Actually I had those people contact me before I had signed up to work for Digital which was a really smart move on his part, because it helped convince me that Digital was serious. People contacted me and I would say but I haven't joined Digital yet, and they said, well but, Sam wanted us to get in touch with you and see if we couldn't get working early on this site selection. So we did that and we found a building that was just getting finished on the outside, it was unfinished inside and we used this building here and they said let's sit down with you and an architect and start designing the inside. So we did that. The inside construction began to take place and by the middle of January we had had enough offices on one of the floors finished, several offices and enough utilities

and phones and there were 2-4 offices, there was still construction on that floor and elsewhere in the building but at least those offices were more or less functional so we moved in in '84.

BL: Where were your heads at that time?

BT: Well, it was pretty clear, to me anyway, by this time I had been at Digital a month and a half, by the middle of January, all of December and half of January and I had made several trips back east and I knew some people in Digital anyway from other days. It began to realize that there were probably more people in Digital who understood and appreciated the work that I had done the previous thirteen years at Xerox than there were people in Xerox who understood and appreciated that work. Digital was a computer company, Xerox was a copier company. At Xerox we had maybe two hundred computer trained product engineers to try to transfer technology to interact with and Digital had maybe seven thousand. There is a huge difference, so my initial reaction and the initial reaction of everyone who came here in the early days was one of happiness that we are now working for a company that will regard our work as relevant. Xerox ignores most of the work you did except laser printing, the invention of laser printing and they didn't get the history of that right either. The Xerox view of

the laser printer was interesting because they believe that the optics part of the device was the key and in fact, digital technology was. So people were really relieved to be in a place where their work is appreciated and it's going to be easier to connect with other people in the company. It's going to be easier for us to have an impact.

BL: Where do you think that DEC stands?

BT: A lot of people in DEC said they were really glad to see DEC do research investment. Sam was just beginning to do that. He had a year and half earlier begun a small group in Palo Alto called the Western Research Lab. And the chap who headed that group worked for me at Xerox and Forrest \_\_\_\_\_ is his name. Probably was instrumental in helping Sam to recruit as perhaps an advisor to Sam. So by this time Digital was a four billion dollar a year in revenue company and it for the first time began to afford a significant outside product engineering research investment. If your much smaller than that it's pretty difficult, you won't find companies in computers that use revenues less then four billion a year who have any serious research investment.

BL: And as a research organization stands, the growth of several research groups but they were all developed.

BT: Well, I don't know about all, but there was a lot of \_\_\_\_\_ in some of their previous, so called research activities. So that's it then Sam has gone on to build two other research groups which some of the rest of us have been able to help with, in Paris, outside Paris, and one in Cambridge, next to MIT and we are beginning to partially form yet another promising group in Cambridge, so it's gradually growing into a significant research group.

BL: What were some interesting research problems?

BT: I think that when any of us took and probably still today, a long-term view at this, crystal ball kind of view, that attempted to reach out, let's say ten years, that even then or now - no one of us saw a revolutionary new theme in computing. In the sense that time-sharing had been a theme that took us beyond batch processing, or in the sense that the Argonaut and then later DECNet gave us yet another way of connecting that community as a time-sharing system. Or in the sense that personal distributive computing which we pioneered at Xerox, had in some sense, subsumed time sharing. Nobody then nor today knows what the next overriding theme of how we use computers is really going to do. But we did know back in '83-84 that there were others who had been involved in

building the first personal distributed computer system. Client server model, giving the personal work station sitting on the desk connected through a local area network to a collection of servers that were sitting around in a building on a campus like arrangement. Those of us who had been involved in the first of those knew that there was a lot that we didn't know about personal distributive computing and so to us that seemed then and still seems that there is a large amount of ignorance that we have to overcome. A large amount of land to be plowed and crops to be sown and harvested. So, you could characterize that kind of prognostication as sort of a consolidation with, years of consolidation, of ethic to go through to try to understand how to build reliable, available robust and serviceable distributive computers. Distributed information packets. There is still an awful lot that we don't know about how to do that and that's our major theme here. To try to evolve over time as modern as to the computer environment. As we are capable of doing in order to learn how to do that. We are guinea pigs with our own inventions in some sense. We try to build what we design and we try to use what we build.

BL: I wonder if you were pushed into \_\_\_\_\_ how you model such a complex system.

BT: Well, we don't really model it in any formal way

because we don't know how. No one knows how to build a formal model, if you mean how do we decide what to do next, then we simply look at what we think are weaknesses or the opportunities for more performance or more service. Or more reliability or any gain that we think we might make along the whole host of dimensions and if we have a good idea about trying something that we think will give us that gain, then we try it, we will design it, build it and use it. We either find out that yes we now have a set of capabilities we didn't have before or we find that we are worse off than we were before in which case we try to learn what we can and throw it away and go back to the drawing board, so it's very much a trial and error process, it's not a process that's amenable to formal models. Not at least at the global systems level. Now there are some pieces of it that with formalisms can be a great help. That's why we have here a dozen or so purists, to help us in doing what we do. but we don't know how to model the whole system.

BL: Looking at it there are interesting continuities but there is also just a vague question of how [INAUDIBLE]

BT: That's kind of a continual process here. The best ideas, the best technical ideas in computer science research come from the computer scientists. They don't come from managers. A good manager will try very hard to

hire people who know a lot more about computer science than he does. If he is the best you can do, as a computer scientist, then his operation is probably not going to be very good. So given that it would be silly, having hired these people, or trained computer scientists with outstanding track records, having hired them and not to pay attention to what they tell you about what the opportunities, technological opportunities are that would also be silly. So this is understood all the way up through Ken Olsen, Ken who is himself a technical person, at least he once was, he doesn't try to tell us what research areas to dig into, you wouldn't - he knows that that's not a thing he ought to do and likewise, Jack Smith doesn't try to do that and Sam Fuller doesn't try to do that. Same Fuller has a Ph.D. in computer science from Stanford. But, he also knows that what we've got be doing is hiring the best computer scientists who are active in research as opposed to managing, that we can hire and then listen to them, you support them. Now where I might have to come in occasionally is if we have two things to do and not enough time or people to do them and then some hard choices have to be made, but I don't sit by myself in my office on the top floor of the building and make those choices, I get out and stir up the waters, I get people to argue about it. With me and with one another, about these trade-offs. Ultimately in any difference of opinion I am seeking what I think of as

a class two disagreement, a disagreement where each party to the other satisfaction describe the other's point of view. If you can get people to reach a class two disagreement they can work together on a common project, whereas, people who have standard kind of agreements where neither can describe to the other's satisfaction the other's point of view and so you get people in those positions to work together. So I might be a catalytic agent with regard to what we're working on and occasionally I might try to get people interested in the theme of some sort, for example, I am generally unhappy with Digital's electronic mail system. I think it has a promise of being a lot better. It's a difficult problem because there are different pieces of electronic mail system for which different sets of groups and companies are responsible and it's a system that has all the elements of system complexity represented in that one phrase, electronic mail, ranging from user interface- of which there could be a variety- the product user end to the deepest transport mechanism types of issues down at the company network. Anyway this is an area that I think we have a lot of room for improvement and occasionally I try to get people interested in thinking about that. But I don't tell people what projects to begin or how to proceed on a particular problem.

BL: What changed the roll of creative imagination vs

some sort of structural direction, and what's the role of the individual in pursuing these themes.

BT: Distributive systems is another variation on that theme, in a sense, communication was a critical element to distributive systems, but there is creativity comes in many different forms, in the birth of a concept, in the the implementation of a concept, in the management of an enterprise or a project. One of my favorite quotes is from Will Rogers - he says we are all ignorant, only on different subjects. There is room for creativity in all areas.

BL: Can you give me an example of individual teamwork - over the last few years.

BT: We are known as the Systems Research Center. The reason is because we primarily do systems research as opposed to device research or component research, or hardware research or software research. Now systems research is software intensive but it involves often hardware innovation as well, but moreover systems research cannot be done by a single individual so system research by it's very nature - it's essence is teamwork. And consequently this is not - at this research center if a researcher came and sat in his office behind doors and worked with either his pencil and paper or his white

board or his computer even, in isolation, after a while people would begin to bother him, not let him keep his door closed, pull him out and get him some help on this, that or the other and if his wallflower tendencies persisted, eventually he would be out of here. We do a lot of peer review of projects, of people and of performance and an isolationist here could not survive the peer review process for very long. So teamwork is critical to what we do and there are a lot of challenges associated with the right mixture of a team, the right size of a team. The team could be too large or too small, too large a team can penalize you in much more than too small a team can. So we tend to have small teams, we tend to take a problem and use a divide and conquer process, break a problem among groups, have a small group tackle this or that piece, sometimes an individual may work part-time in one group on one piece and part-time with another group and another piece. And ultimately we have to bring these pieces together because what we are doing here is evolving as advanced a computer environment for our own experimentation as we are capable of doing. So not only do these small groups of people as a group have to construct a piece that works and satisfies whatever the objectives of that project were but when that piece reaches a certain point it then has to be merged into the rest of the pieces that are already existing or the new pieces and work as a whole. So

that's what system research really is all about. Time-sharing and the research that led to the first timesharing systems has this property it is a good example of a systems research work. The Argonaut research had this property, it's another example of consistent research and personal distributing computing and laser printing and the development of Ethernet, those are all examples of system research and they all have this same set of properties, depending enormously on teamwork.

[END OF SIDE ONE BLANK SIDE TWO]

[BEGIN SIDE THREE]

BT: There are mechanisms for connectivity. We ask senior technical people to come hear a review of what we are doing every year. For two half days. We ask them to help us with making contacts throughout the course of the year, which many of them have done. We try to go back and give seminars back east, we also have some engineering groups here in Palo Alto and we try to work with them as well. So the connectivity is never complete and the isolation is never totally removed. There is a sense in which the isolation can't be totally removed because there are certain areas of technology that if you are going to work in them by definition you are going to be isolated because - the example I gave a moment ago -

there was no existing multi-processor work station and there was no operating system to run it on so since you built your own operating system, you are now an isolated world. By definition. So you can't if you try to innovate in base systems technology then you are going to have to have periods where you are isolated there is no way around it and you will take criticism from some sources that you are too isolated, yes, I know but we've got to work this out. They think the \_\_\_\_\_ was \_\_\_\_ at the Systems Research Lab for building these risk machines. They are too isolated. But research has to be a thorn in the side of engineering on occasion. That's part of it.

BL: Which ties into the Xerox example where they couldn't figure out corporately whether they were really an office information architecture company or what. You wonder if there are comparisons of Digital in the eighties to Xerox in the seventies in terms of floundering about and their goals and commitments.

BT: You mean around VMS and UNIX or ULTRIX or the workstation wars or what?

BL: I'd like to tie all of them together into one. A strong lack of definition.

BT: There were a lot of different forces in the company that were turning in different directions and there still are some but I think we are the realities of competition of too low profits finally brought some sanity to it slowly. But it's a thirteen billion dollar a year company in revenue and it's going to be hard to kill. If you were a CEO of such a company and for some perverse reason you decide to kill the company, you'd have a hard time doing that. It's very important.

BL: How did your perception of Digital change as you became in the company and saw, you were no stranger to the art of organization.

BT: Do you have any evidence that my perception of Digital did change? Because I have had a lot of respect for Digital from when I first learned anything about a computer company. My computer history began in the late fifties, early sixties. Digital began in 1957. The group that Ken worked with at MIT and at Lincoln had in it people who later became friends of mine and some of the machines that Ken had something to do with like TX-2 at Lincoln Lab is a machine that I had a lot to do with when I was at AARCO sponsoring the computer research at Lincoln Labs. So there has been a lot of different kinds of connectivity between me and Digital from the early days. When Digital moved from selling modules

principally to selling computers, principally I began to get more and more connected with Digital. JC had a large influence on me and he was the head of the group that Beau Frank and \_\_\_\_\_ had bought the PDP-1. That group convinced Digital at the time to put some memory protection hardware on that machine and they did that, and that opened up the whole country and timesharing world. So that was a critical element in the history of the timesharing. Our AARCO was really the primary sponsor of timesharing research and that work had gotten underway by Lichlighter, at AARPO I came to them in 1965 so I was there from '65 to late '69 and during that time was timesharing before it evolved I bought up a lot of PDP-10's. I am sure I bought the PDP-10' but the money came through my office so think of them as I bought them. I am sure that AARPO contractors bought more PDP-10's than any other single source, I don't know how many - I never counted, but we bought PDP-6's first, we also bought a fair number of PDP-8 and some other equipment along the line, well Digital paid more attention to the work that the AARPO funded groups were doing than any other computer company. Part of my job was to get this computer research that AARPO was funding out into the real world so that the Defense Department the U.S. government more largely could take advantage of it. If you can't get it out into the commercial world then people cannot take advantage of it. So Digital helped me

a lot. I am sure that Digital wasn't aware or even care about helping me but the fact of the matter is that they did, by virtue of the fact that a lot of those people were tracking at least, roughly, sometimes carefully work that was going on in various research places that I and AARPO were supporting. Then in 1966 I started a project called the AARPOnet and Digital by '69 we had the first three nodes of the network up and running and Digital began to pay attention to that as well. Then by 1973, Digital launched DECnet which borrowed a lot from AARPOnet because AARPOnet by that time was a working system and so Digital could see that this stuff works and that made it easier for them to launch DECnet. By that time in '73, I'm working for Xerox so my connections with Digital become less, until the late '70's when the Xerox product development group goes to Digital, Bob Metcalf, who had worked for me in the Xerox Computer Science lab before he went over to research and development. The Xerox development group approached Digital and Intel to negotiate an Ethernet license, well Ethernet was admitted in my lab and so there was that connection there that began to come back again between individuals so I can't really say that given all of that early history with digital I can only say that my model of Digital changed very much over these years. I believe that there have been two things that are responsible for my not working for Digital earlier and one was that Digital, as I said,

about the time I joined Digital they couldn't afford a significant research investment. Because their revenues couldn't have supported it. It's interesting to note that Xerox was a four billion dollar a year company when they launched their big research investment, internal. So was Digital, maybe there is a magic number there someplace, but at any rate, it's reasonable to understand that it takes a fairly sizable revenue before you can afford research investment. So that's one aspect of it and another aspect of it might be, this is speculation, might have to do with Gordon Bell, I always got the sense the Gordon Bell liked to run things out of his inside coat pocket and I've heard people tell me in Digital, that Gordon Bell not only recruited them for research but he drove them out of research after a while because of short term nature that he approached things with. It's highly likely if that's the way that "research" was dealt with when he was in digital. Re: development. That it's not really research it's advanced development, maybe sometimes it's not so advanced development. I wouldn't have worked under those circumstances. And Gordon wouldn't have wanted me either, so either one of those factors could have been operational and certainly the first one about being able to afford serious investment, is the one we have most definite evidence of. But I've always thought of Digital as a company that to it's credit runs more on technical peer evaluation, the

management of the company, many managers have technical backgrounds and that makes it easy for a technical person to explain what he's doing and why he's doing it and that's a big benefit and they have a technically oriented baggage system, they respect directness, straightforwardness, excessive diplomacy gets in the way. It's an engineering oriented culture. And I like that. I've always believed that about Digital and I still believe it although sometimes I think in some ways we are getting too far away from it. If you want to make something happen at Digital, you find out who the people are that agree with you and who are strong technically and you can usually make something happen, now that's a blessing and a curse, when I first joined the company and was exploring and trying to learn what was going on I found seven different display controller projects scattered through engineering and no one of the seven knew about the other six. So you say, that's wasteful, well, yeah, at some level but what are the alternatives, you would have some czar or committee, which is worse, sitting in charge and making judgements that sometimes cut out wheat as well as chaff. So in some ways although Digital is a very large company, it operates sometimes in a beneficial way, like a small company. Go talk to so and so and see what you can make happen as a result of the two or three or six of you getting together. In a lot of companies to try to get cooperation across groups

you have to go do some miniature version of the U.N. and sit around building institutional understandings of institutional agreements which none of the working people are ever involved in. So consequently the objectives and methods of operation are laid down by the people who aren't actually going to be doing the work and therefore don't know what proper objectives or methods of operation ought to be and it's just not very satisfying. So when I look at the two models that we've been talking about here - I might be prejudicial.

BL: Once they get wonderfully inspirational. That's consistent with all I've heard. Going back from the earliest history we really see the evidence of RPO being sucked in by every Digital engineer and taken as this is the model. There were just a few disconnects that I never figured out how they could have missed the user interface stuff. The company never had any user interface? There were thirty-four work station projects in the '82- '83 time zone.

BT: It got a little out of control. The workstation dilemma and the user interface dilemma are really different dilemmas. They are related slightly in the sense that without a bit map display you don't have the ability to do any more than, you don't have a lot of user interface choices. With bit-map display now you've got a

medium that is highly moveable and shapeable but shapeable through software, not hardware. And the people that you needed to have pull together to make the right workstation revolution were primarily hardware people. People that you move together to get a reasonable utilization of a bit map display are software people. It's also worth noting that - this is sort of a thesis of this vignette is that old ideas die hard and old habits die hard, new habits are very difficult to take on when we build the first bit map display work station in Xerox computer science lab, it's called the Alto, the people who built it and used it day in and day out and most of them were software oriented people, used that bit map display for three years like a glass teletype. Before they figured out- we can do windows and icons and all kinds of other stuff here.

BL: I see Digital's success with the VT-100 as being maybe the thing that stifled any kind of innovation.

BT: Well, there are people at SRC who loved their little keyboard editors, they love to type in character strains, long character strains, control X, Y, Z, 3 lower case, 2, 6, they did great investment of their nervous system in learning conspicuous long serious of codes that they had to type in to edit something or to write a program, but to make computer do something and by God, they're fast

at it and good at it. Never mind the fact that it was dehumanizing to have to learn that crap in the first place, so when you say, you can take this mouse and point at that icon and it'll just pop that thing right up. Oh, I don't have to take my hand off the keyboard to do that. Well here it is 1991, so it's not too surprising that these user interface area progress is so slow. you have tons of people inside Digital who, as you say, because of internals were so successful and they learned all these keyboard commands and thought why move into the twentieth century, my model-T runs just fine. It's not only the new people at Digital, we've got new, young people who have the same sickness. So it's going to be hard to get people to change those things quickly.

BL: That's one of the most interesting challenges in systems research. But you think that by using the power to commute in computers is a tool of the model ---- DO you think that that has the seeds of it's own solution?

BT: It takes longer than anyone thinks.

BL: Is it great finally to work for a computer company?

BT: yes, it is. The first week I told everybody who might come on you'll really enjoy working for a computer company.

DEC - BOB TAYLOR, TAPE 1-2-3  
PAGE 23

[END OF TAYLOR INTERVIEW]