Oral History of Butler Lampson

A.M. Turing Award Winner, 1992

Interviewed by:
David C. Brock .

Edited by:
Dag Spicer

Recorded December 6, 2023
Cambridge, MA

CHM Reference number: 2023.0211

Brock: This is an interview with Butler Lampson happening on December 7th, 2023, at the...

Lampson: December 6th, actually.

Brock: Is it the 6th today?

Lampson: It's the 6th.

Brock: Oh. Skipping ahead. Thank you for that correction. Doing great right off the bat. December 6th, 2023. Great. With Butler Lampson at MIT, and I'm David Brock from the Computer History Museum, and let's see if I can get some more facts wrong in my question list. But I believe you were born in Washington, D.C....

Lampson: That's true.

Brock: ...in 1943.

Lampson: That's... yeah.

Brock: Could you tell us a little bit just about your family's background and a little bit about the family you found yourself joining in 1943?

Lampson: Well, I guess my father was in the Army at the time since the Second World War was going on. My mother came from a family of diplomats. Her father had been a fairly high-level diplomatic officer in the U.S. Foreign Service, but he died in the late 1930s. So she was somewhat unmoored. My father came from a Connecticut Yankee family. His father was a doctor who lived in Hartford, and I'm not sure that they would ever have met if it hadn't been for the war that took him to Washington. But he had been studying history, got a Ph.D. from Harvard in English history.

Brock: And then served in the military during the war.

Lampson: Right.

Brock: That brought him to Washington.

Lampson: Yes.

Brock: That's where they met. And your mother's family was from Washington, from that...

Lampson: Well, to the extent that a diplomat's family is from anywhere, yes. They lived in Washington when they were in the United States.

Brock: And then... I think I read somewhere that your father also joined the State Department.

Lampson: That was a little bit later. He joined the Foreign Service, right.

Brock: Yeah. What was his career in that like?

Lampson: Well, his first station was in Ankara, Turkey. So I was taken to Ankara when I was about three and we lived there for about four or five years.

Brock: And then did you return to the Washington area?

Lampson: No, we went to Germany.

Brock: Straight.

Lampson: We lived in... Germany was occupied at the time. We were living in Dusseldorf for a while and then later in Bonn, which was the headquarters of the Occupation, and it was a little bit strange. It didn't seem that way to me at the time, of course, because I was eight or nine. But the American Army apparently had decided that it wasn't healthy for American dependents to associate with Germans. So they built a little American town on the banks of the Rhine. They commandeered a plot of land, and they built some apartment buildings and a shopping center and a school and a church, everything you needed to live, and we lived there in this totally artificial environment. Of course, as I say, it seemed perfectly normal to me because I had never experienced anything else really.

Brock: So you were going to school with...

Lampson: So I was going to an American Army-run school, which was pretty much like a typical midwestern U.S. school, with American teachers and the whole bit, and then a German lady was brought in twice a week to teach us German for 45 minutes. So we didn't learn... needless to say, we didn't learn any German.

<laughter>

Brock: And what about your mother? What were her interests and activities like?

Lampson: Well, she was a diplomat's wife, and in those days that meant that you... and, of course, she'd been a diplomat's daughter, too. So she'd lived in this... what should I say? I don't know if I have much first-hand experience with this, but she was accustomed... her father had been stationed in Hungary, in Budapest, for a while, and subsequently in Montevideo. So she had this experience of living in an upper-class diplomatic environment, but she was, also, because of this, somewhat unmoored from typical U.S. society.

Brock: And sometimes there's sort of a major theme in a household, be it politics, or religion, art, the sciences. Would you have described your household as having some major themes?

Lampson: No, I don't think so.

Brock: Okay. Yeah. Was your father's interest in history at all present in the...?

Lampson: It wasn't particularly obvious, but, of course, I was interested primarily in science and engineering from an early age. I think I was a fairly abnormal child. My guess is, if I were growing up today, I would be diagnosed with being on the autism spectrum. I was not very social, and extremely focused on the things that I was interested in.

Brock: And how did that manifest, that interest in science and engineering that you recall as kind of almost always having been there? Are there some early memories about how that was manifesting?

Lampson: Not when I was that young, no. I don't remember anything about that. I guess I had a chemistry set in Germany, but it was pretty superficial.

Brock: And what about reading? Were you a big reader?

Lampson: I was definitely a reader, yes.

Brock: Could you tell me a little bit more about that, the sorts of things that you were reading and where you were getting the books?

Lampson: I think I typically read books that were written for kids of my age or somewhat older than myself. So I can remember reading books about figures in American history and that sort of thing. I can remember reading books by a man named Henty who was, I think, English. He wrote books typically about teenagers in various exotic places in the British Empire. Although he did spread the wealth a bit more widely, but that was for the most part what it was, and I found those enthralling.

Brock: And was science fiction ever a thing?

Lampson: Definitely, I read a lot of science fiction, yes. I was into that from a fairly early age.

Brock: Were there any particular people that you made a point of reading?

Lampson: Oh, sure. Isaac Asimov and Robert Heinlein, the major figures of 1940s and 1950s science fiction.

Brock: Right, and then after that stint in Germany, did your family return to the States?

Lampson: Yes. My family lived in Washington for five years, between about 1955 and 1960, and I was sent to boarding school.

Brock: Was that because they were going... the possibility that they would go abroad?

Lampson: I'm not quite sure why they decided to send me to boarding school, but I think I failed to get into the highest-level Washington private schools, and I think they weren't happy about the next-level schools. I did go to a school in Washington for a year, but after that they sent me to boarding school, which was the same school that my grandfather had been to.

Brock: And so that was the connection to Lawrenceville?

Lampson: Yes. He actually figured in the famous Owen Johnson's Lawrenceville stories.

Brock: Okay. As an anonymized character?

Lampson: His name was J, Butler Wright. J. stood for Joshua, which he hated, so he refused to be called by that name and insisted on being called Butler, which is how I acquired it as a first name, but in the stories he showed up as Butsy White. He wasn't a particularly major figure, but he definitely was there.

Brock: Well, with your focus on the things that interested you, how did that play out for you, kind of in school before you went to high school? Was that... were you... did you excel at certain things?

Lampson: Well, I had terrible handwriting, but I think I excelled at pretty much everything else academic. I've never been very athletic, so I didn't do sports much.

Brock: Well, so you were at Lawrenceville for, let's see...

Lampson: Four years, from 1956 to 1960.

Brock: Could you just talk about what the school was like in that era? Was it all boys?

Lampson: It was all boys. It went co-ed, I think, about 15 years later, relatively late as the Northeastern prep schools went. Lawrenceville wasn't quite a normal Northeastern prep school, because it was too far south, and it actually had a fair number of students from the South traditionally.

Brock: And it was pretty small at that time, was it not?

Lampson: I think the graduating class was about 150.

Brock: Okay, not that small.

Lampson: Not that small, no, not tiny. It's expanded some since then, and since they started taking in girls, but it's still not very big.

Brock: And how was your interest in science and engineering fostered during those years?

Lampson: Well, it wasn't fostered very well. Looking back on it, I think the science at Lawrenceville was not very good. I think by the standards of 1950s high schools, it probably was excellent. But looking back on it, [it] seems to me I didn't learn much science there, and the instruction doesn't seem to me to have been very distinguished. Mathematics, I studied calculus, and in the 11th grade; by the time I got to the 12th grade, there was no regular class that I could take because I'd already taken all of them. So they actually organized a special two-person class for myself and one other student, who was a fairly good friend of mine. But the chairman of the math department taught

this class, and he actually was not qualified to teach beyond first-year calculus. So we didn't accomplish much, and in fact, what tended to happen was that we all fell asleep in the class, which was the first class of the day.

Brock: The three of you?

Lampson: The three of us, yes, right.

<laughter>

<overlapping conversation>

Lampson: So that was quite disappointing, actually, in retrospect. I missed out on a lot of opportunity that I might have had if I had a better instruction.

Brock: What about for the humanities?

Lampson: That was better.

Brock: Yeah, yeah.

Lampson: I think the history instruction was pretty good, and the English instruction was pretty good. But the science was not very strong.

Brock: Did you have any particularly meaningful mentors to you at that time, people who were encouraging you to pursue your interests in science and engineering?

Lampson: Not especially, no. I was pretty much doing that on my own. I did have a classmate who found an old IBM 650 at Princeton that was way underutilized, and he coaxed the woman that was in charge of it into letting him hack on it and so he and I would go up to Princeton maybe a couple times a week and punch cards and run them through the 650.[1]

Brock: I had read that in another interview, and I just thought it was quite amazing because at that time... so what are we talking about, '58, '59, something like that?

Lampson: '59, yeah. That's right.

Brock: I mean, an electronic digital computer is a relatively rare and expensive thing in the world...

Lampson: Yes.

[1] [interviewee's note] This was **Gillian Richardson of** the Statistical Techniques Research Group**, https://as.amphilsoc.org/repositories/2/archival_objects/969489**

Brock: ...and I was just kind of amazed that they would let you.

Lampson: Well, the 650 was kind of obsolete by that time. It was a vacuum-tube machine, decimal vacuum-tube machine with drum memory. You could get a sizable amount of kit for the 650 from IBM, but the only actual I/O device that this machine had was a combination card-reader punch.

Brock: And then you would take it to a printer?

Lampson: Then you would take it to... and there would be a card, IBM 407 I guess it was, which you fed in cards and it would print the cards. So it was, by even the standards of the 1960s, it was a pretty primitive setup and I guess even then... that was why it didn't have very many users and we were able to squeeze our way onto it.

Brock: Was there a close relationship with the university and Lawrenceville?

Lampson: Not that I'm aware of.

Brock: Yeah, it was just the initiative of your friend?

Lampson: Yeah, I think nowadays, probably, I would guess they have arrangements by which Lawrenceville students could take classes, but there was nothing like that in those days. Fortunately, there was a bus, a public transit service, so we could actually get up there on the bus.

Brock: That was my next question. How did you get there?

Lampson: That was how we got there.

Brock: And so, previous to that, I'm sure you were well aware of computers through science fiction and other things that you were...

Lampson: Well, to the extent that you could be aware of computers through science fiction, but it was pretty removed from the reality.

Brock: Right, right. When your friend told you, "Hey, I've got access to this computer and we can program," was that something that you were immediately excited about?

Lampson: Oh, yeah, definitely. It was great. Not that I ever did anything very interesting on this machine, but, yeah, it seemed very cool at the time.

Brock: Well, two follow-up questions about that. What was... could you talk about that, your first experience of programming with that? Because it's something that you went on to do for quite some time and be interested in for quite some time. So if you could just talk about what that first experience was.

Lampson: I have no memory of what I... what sort of programs I wrote, what they did or anything.

Brock: But you were captivated?

Lampson: I was captivated, definitely.

Brock: Okay. And then I was wondering, too, just how, as you and your friend were out there programming away on the 650, how unusual was it to have, amongst your kind of peer cohort, to have that experience?

Lampson: Oh, I think it was totally unusual. I'm not aware that anyone else at the school was interested in computing, and I don't know what they would have done if they had been interested, because there was no way that there was going to be a working computer at Lawrenceville in 1959.

Brock: Right, right.

Lampson: Even the smallest computers were fairly expensive.

Brock: Yes. Well, you must have been clicking along pretty well academically.

Lampson: Oh, yeah. I did great at Lawrenceville. It was very bad for me, actually. I got very lazy because I found that the schoolwork was so easy.

Brock: And that didn't... well, I guess that afforded you time to go play with the computer.

Lampson: Indeed.

Brock: But it seems that toward the close of your high school period, physics was looming as a major interest.

Lampson: I had always been very interested in physics, yeah. I didn't have an opportunity to learn much physics at Lawrenceville.

Brock: And do you think it was also... I mean, I'm just thinking of the time period that you're growing up and going to high school. I mean, this is kind of the age of the American physicist, if you will. There's very prominent examples, very prominent on the international stage. Do you think that figured into your interest in physics at all?

Lampson: Well, maybe. I think I was, for the most part, not very conscious of that. But we did have, on one occasion, Robert Oppenheimer came down from Princeton to talk to the science club, and he gave a talk about K mesons. I can remember being absolutely enthralled by it and writing a wildly enthusiastic description of his talk for the school newspaper. I can no longer remember what the content was. It had something to do with the fact that some fraction... if you had a beam of K mesons, some fraction of them would decay as they were traveling down the beam and you could draw some interesting conclusions from that, but I don't remember what the conclusions were. I do remember that Robert Oppenheimer was very... to me, was very charismatic. He's one of the two extremely charismatic people that I've run across in my life.

Brock: Who was the other?

Lampson: The other was a fellow named Ryal R. Poppa, who was the CEO of Data Processing Financial & General (DPF&G), a computer leasing firm which was one of the funders of Berkeley Computer Corporation in the early... in the late 60s and early 70s, and I ran across him because he came to... while BCC was going down the tubes, he came to try to persuade myself and a number of other people not to leave.

Brock: And he was just another person?

Lampson: He had a lot of charisma. Made a big impact on me. Did not succeed in his goal, because it was pretty hopeless at that time. But he made a big impression on me.

Brock: So as high school was coming to a close, was Harvard prominent for you as a place to go, given your father's association with it, having gone there to graduate school?

Lampson: I don't know that had much impact on my decision to go to Harvard. My father went to Amherst as an undergraduate. I don't know. It just seemed like a no-brainer. So I applied and got early decision and got in and didn't think about it much anymore. But two of my good friends also applied to Harvard and got in and, actually, we roomed together for several years at Harvard, which wouldn't be allowed today. But things were much cozier in those days.

Brock: And you were at Harvard from ‘60 to ‘64?

Lampson: That's right.

Brock: So how did this interest of yours in physics develop during those four years?

Lampson: Well, I took a lot of physics courses and I learned a lot of physics. Probably from some points of view, too much physics and not... I definitely missed out on some basic things because I was good at it and I was allowed to take a lot of graduate courses. And so I had a good time. I learned a lot of physics. But I was getting seduced by computing in those days. So that was a big distraction.

Brock: Could you talk about that, too? Because both computing at Harvard and computing in the surrounding communities, where were you really encountering computing?

Lampson: Well, Harvard computing was very weird at the time, because of course Harvard had been one of the pioneers in computing. Howard Aiken was there and he built these giant... well IBM built them actually... but he designed these giant relay machines. There was... and there was a sizable building actually devoted... called the Aiken Laboratory that was devoted to computing. It was a rather strange building. If you went in the entrance, you found yourself in a sizable lobby, which had a glass wall on the right-hand side, and the lobby was rather dimly lit, except there was this receptionist desk at the far end opposite the door. So you could walk over there and get introduced to somebody or whatever.

But on the opposite side of the glass wall were the Mark I and Mark IV relay calculators, which of course were not in use anymore in the early 60s. But they were still there. So you had these banks and banks of relays, paper tape

readers, paper tape this wide [about 2 ft.], which had the programs punched on it, typewriters for doing output, and the whole bit. So all down one side of this, on the opposite side of the wall was the machine room, which was quite large and all down one side of it was the Mark I calculator and all down the other side was the Mark IV calculator, and if you looked way down at the opposite end of the machine room, you could see the only part of it that was lit was the university central computing facility, which was a UNIVAC I.

So you can see that Harvard was pretty primitive. They had been at the forefront of computing in the early 1950s, but that definitely didn't last. In fact, it's quite striking that in the US, if you look at the universities that were into computing in the earliest days, Harvard, Illinois, UCLA, MIT, only one continued to lead in the field in the ‘60s. The others all fell by the wayside. And not that they disappeared; Harvard has moderately respectable computing today, and Illinois certainly does, and UCLA and so forth. But they're not in the forefront, and the schools that came to the fore in the 60s, like Berkeley and Carnegie Mellon, were all different and had not been involved particularly in computing earlier. The only exception was MIT, which continued to lead, maybe not surprisingly. But it's a very striking phenomenon and quite different from what happened in the UK, for instance, where Cambridge and Manchester were the original founders of academic computing in the UK, and they continued to be prominent and are still prominent today.

Brock: Yeah. Or if you think about the University of Pennsylvania.

Lampson: Yes, that's another striking example.

Brock: And it's sort of... yeah, it's interesting because once people leave to commercialize things from a couple of the schools that you mentioned, what is happening on the campus becomes more moribund. But at MIT, there was just this continual activity of making new systems.

Lampson: That’s right. Yep. And part of it, I think, was that MIT was the nucleus of the SAGE system, the SAGE Air Defense System, and that provided a... well, maybe I should say a basis for the university to continue to be involved in computing. In those days, it was very tricky because computers were very expensive and building a computer was extremely non-trivial. So it was hard for universities to continue to have a presence in the field and MIT, I think primarily I suspect because of SAGE, continued to build machines and also do interesting things with them.

Brock: Yeah. That makes perfect sense to me.

Lampson: And, of course, MIT being a leading technical school, it was important to the university, although not that important. It was very striking in the late ‘60s and early ‘70s... for that matter, I guess all the way through the ‘70s and maybe even all the way through the ‘80s, MIT had a major research presence in the form of Project MAC. But there was actually no computing on the campus [except for the computer center]. Project MAC was in Tech Square across the street in leased quarters, and the university for all that time never provided computing with its own space. The Stata Center was the first time MIT actually built any facilities for computing. I asked him... I think it was Joel Moses, who was chair of the CS department at the time, why he thought this was, and he said he thought that the MIT administration thought that computing was a flash in the pan, which I think a lot of people actually thought in the ‘70s.

Brock: It's hard to recover a mindset where you would... it seems unbelievable.

Lampson: It's sort of the opposite today. People take it too seriously.

Brock: Well, where were you finding the opportunities while you were an undergraduate at Harvard? Where were you finding the opportunities to plug into what was happening in computing at MIT and elsewhere?

Lampson: Well, I used to come down to MIT to the Project MAC reading room and I had a limited opportunity to run Fortran programs at the MIT Computing Center, which was managed pretty loosely, so you didn't even need an account. You could just put in your card deck and run your program on the [IBM] 7090 batch processing system. Peter Wegner, who was subsequently, I guess, at Brown, taught at Harvard a course in Fortran programming, which I took, and that's how I learned to program in Fortran.

At Harvard, as I mentioned earlier, when I came there, the university's central computing facility was a UNIVAC-1, which was pretty much a totally obsolete machine by 1960, and I had nothing to do with that. In 1962, they ripped out the Mark I and Mark IV relay calculators, and they ripped out the UNIVAC-1, and they installed an IBM 7094 and that became the university's central computing facility. But I never actually used that much, because I hooked up with an assistant professor of physics [Lou Hand, later at Cornell] who wanted to do computer analysis of spark chamber photographs, and I did a lot of programming for him on a PDP-1, which was attached to the Cambridge Electroni Accelerator, which was where he was doing his work.

Brock: It must have been one of the very first PDP-1s.

Lampson: It was not one of the very first, but it was... I guess, there were never that many PDP-1s.

Brock: And what was that experience like?

Lampson: Well, that was where I really came to grips with interactive computing for the first time, because the PDP-1 was a hands-on machine. So you sat in front of it and flipped the switches, and it had a typewriter for both input and output, and it had a display, a very basic... you told it the xy-coordinate where you wanted a dot, and then lit up the dot, and then you told it in another xy-coordinate and wanted another dot, and it lit up that dot, and if you wanted characters, you had to work it... the bitmap for the character had to be in the machine, and the machine had to interpret it to light up all the dots. So I think you could display maybe 300 or 400 characters on the screen before the flickering got to be too painful and you couldn't really stand it anymore. But this was very advanced for the time. This was more or less the state of computer graphics in the early ‘60s [except for Ivan Sutherland’s work at Lincoln Labs, which I knew nothing about then], and I did a lot of programming for that machine.

Brock: And were you doing that as an undergraduate research assistant?

Lampson: Yeah, that's pretty much what it was. I was not really doing research because I was doing what this physics professor wanted me to do. But I got a chance to write a lot of computer programs and even learned something about graphics, because the machine was not very heavily used. So I would work typically at night, and I had lots of time on the machine. So that was where I did most of my programming. The other thing I remember

about computing at Harvard was that Ken Iverson came to Harvard on sabbatical for a year and taught a course on APL. This was in the days before there was an APL implementation. In fact, APL only existed in the form of a book, which was called “A Programming Language” and Ken was actually strongly opposed to the idea of having an implementation, because he was sure that if somebody implemented it, they would make all kinds of compromises and ruin his beautiful language. He was very fortunate later, long after I had left Harvard and he had gone back to IBM, to run across a very talented implementation team, which built a beautiful implementation of APL, which got a lot of use in many places. But this was way before that, when it was strictly pencil and paper.

Brock: And, to me, it seems... and please disabuse me if I'm wrong at that, but the difference between APL and Fortran, that APL itself is kind of like an exercise in mathematics.

Lampson: Absolutely.

Brock: And so, for me, as kind of a more naive observer, it's a different beast than Fortran.

Lampson: Absolutely.

Brock: It's a mathematical structure.

Lampson: As you can see from the fact that there was no implementation.

Brock: Yes. I guess trivially true, but maybe also interestingly true … was that your first encounter of a computer language of that sort?

Lampson: Oh yes, definitely. And there was nothing else like it at the time. Later on, there was the Culler-Fried system, which picked up on the idea of having vectors as a primitive element that you could program with. And the Culler-Fried system, of course, also had a screen. But that was way later. And I think Culler-Fried was not as elegant as a programming language as APL. APL was quite amazing.

Brock: Was that... just in thinking about your work that you get into, not very long after this, in fact, of working on programming languages of your own, how influential was that experience?

Lampson: I don't think it was very influential, because the idea that you could actually have a practical programming language where the primitive elements were not scalars, was... in the ‘60s, you had to make a lot of other compromises in order to make that work and I wasn't willing to make those compromises. So the programming languages that I worked on were all much lower level than APL.

A startling feature of APL was that although you could write loops in APL, the way it's... I'm not sure how the implementation did this, but as long as it was just pencil and paper, the way you had a loop was that you drew a line from the end of the loop back to the beginning with a little arrow on it. But you didn't need loops. That was the amazing thing about APL. I thought I had understood it pretty well, much better than most of the people that were taking Ken's class. But on the last day I discovered that I had actually not understood it that well, because I said to

Ken... we were discussing what APL was good for and what it wasn't good for, and I said it wouldn't be very good for coding a symbol table.

And Ken said, “Oh, no, it would be very easy. You make the symbols in the table as a matrix, where the columns are the characters of the symbol, and each symbol has a row and so you have this big matrix and now if you want to do a lookup, how do you do it? You take the vector that contains the characters of the string you want to look up, and you do equals between the vector and the matrix... that gives you a matrix of ones and zeros, and now you do ‘AND [Boolean logical function] crosswise. That gives you a one for every row in which every character matched, and now you do “find first one.’” Those are all primitive within APL, so there was no loop. And this was a revelation to me. As I say, I thought I had understood the idea pretty well. But, in fact, it was clear that I hadn't.

Brock: So we mentioned your time when you were programming the PDP-1 and having the machine to yourself. Did you have any exposure to timesharing systems during those years?

Lampson: Well, not really. Towards the end of my time as an undergraduate at Harvard, CTSS was working at MIT and I certainly was aware of it, but I never used it. Because you did have to have an account in order to use CTSS. Whereas, as I say, you did not have to have an account to run batch jobs at the MIT computer center.

Brock: Why was this?

Lampson: In fact, it was quite extraordinary. The Computer Center ran FMS [Fortran Monitor System] to run the batch jobs, but there was very little control. In fact, on one occasion, I managed to effectively shut down the computer center because I had written a program that used a scratch tape. So the program wrote some stuff on the tape, and then it rewound the tape and read it back and did some more processing,. The programwas sort of working, and then I made some substantial changes, and I submitted the job and I came back a few hours later to get the printout, and there was no printout. The deck was back, but no printout. That's weird. Usually... this never happens. I submitted the job again, and the same thing happened and after that, I noticed the operation of the computer center seemed to be somewhat disrupted, and it turned out that there was a bug that I introduced when I made changed to the program. I was not rewinding my scratch tape; I was rewinding the batch input tape, and there was nothing to stop you from doing that. So it was running all the jobs again, and it kept doing that until the operator noticed that they'd seen these printouts before.

<laughter>

Lampson: So I ran!

<laughter>

Brock: Yeah, that's right.

Lampson: I certainly wasn't going to admit to having done this.. And since, as I said, there were no accounts, they had no way of easily tracking me down.

Brock: Right. And just thinking about that kind of a program, it sounds kind of like a sophisticated program if you're doing things and writing it to a scratch.

Lampson: I guess it was. I certainly thought it was.

Brock: Was this a program that you wrote to kind of explore the capabilities of a computer and what it could do?

Lampson: I can't remember anything about the program, unfortunately.

Brock: But, in general, what you were running on that machine, were you doing it for a class?

Lampson: I was not doing it for a class. It was definitely exploratory. Yeah, absolutely.

Brock: One thing I also noticed is that as an undergraduate, you joined the ACM.

Lampson: Mm-hm.

Brock: Was that unusual? And was that to get access to journals?

Lampson: Yeah, it was to get access to journals, absolutely. They had a very cheap student rate. I think I joined in September or October of 1960. At that time, if you joined, they sent you the CACM. You joined by the year, calendar year, and they sent you all the... at the time, the ACM only had one publication, which was the Communications, and that was, at that time, I think the main vehicle for technical publication in computing, and they sent you all the issues that had already come out in 1961. Yeah, it was ‘61, I think. I guess I joined it late in ‘61 because I had come to Harvard in 1960.

Brock: Right.

Lampson: And I thought that was really, really, really cool. There were all kinds of interesting things to read about, and I think it cost $20 a year or something like that. Of course, dollars were worth more then, but even so, it was cheap.

Brock: What about any development of your interest in or engagement with the humanities, or were you too busy?

Lampson: No, I took humanities courses. I took a drama course, and I actually took a course from Henry Kissinger in 19th century European diplomacy.

Brock: How was that?

Lampson: It was cool. I took a course from Ken Galbraith. That was really boring because Galbraith's system at that time apparently was, if he was writing a book, which he was at the time... I think it was the book that subsequently turned into “The New Industrial State,” a book about how the U.S. and Russian economies were basically the same. But his system was... each chapter of the book was a lecture. So one lecture per chapter of the

book. Galbraith was apparently a believer in the I-tell-you-three-times principle. So the lectures alternated between being a chapter of the book, and then the next lecture, half of it would be a summary of the previous lecture, and the other half would be a preview of the next chapter. It was very boring because you heard everything three times, and it wasn't that profound. Also, it was completely wrong.

Brock: But a very good technique for writing a book, probably.

Lampson: I guess. But this book was a complete bust, because he completely misunderstood what was going on, which in retrospect is kind of surprising because he was a very smart guy. But he was not a very interesting lecturer, unfortunately.

Brock: And what were your summers like during those times?

Lampson: I was a camp counselor in a summer camp.

Brock: The same place?

Lampson: The same camp, yes, for three years.

Brock: Was that a place that you had gone as a youth?

Lampson: No. No, I got hooked up with them because I went to the university's job center looking for a summer job, and the director of this camp had signed up because he needed counselors, and that was how we got together.

Brock: You must have enjoyed it if you did it on repeat.

Lampson: Yeah, I had a great time.

Brock: Was it also in New England?

Lampson: Yes, it was. Actually, we have a summer place just down the road from it now. That's great. It was on a nice lake.

Brock: Well, could you talk about your... as your time as an undergraduate was coming to a close, [what was] your thinking around pursuing graduate studies?

Lampson: Yes, I assumed I was going to go to graduate school in theoretical physics, and I applied to only two schools, Princeton and Berkeley, and Princeton rejected me, so I went to Berkeley, which was extremely fortunate for me because Princeton was nothing in computing at the time.

Brock: And so you went to Berkeley in ‘64.

Lampson: That's right, and I spent the summer of 1964 at Los Alamos.

Brock: Oh, I didn't realize that.

Lampson: They had an internship program for budding graduate students.

Brock: What was that experience like?

Lampson: It was... in retrospect, it was not really very successful because I got snapped up by a group that wanted a Fortan programmer. So I hacked on a fairly substantial Fortran program for the summer, which meant I didn't really learn much physics, and I didn't learn much of anything else either. But I had a good time.

Brock: What were the computing resources like there?

Lampson: I believe this program ran on a 7094, but they had a Stretch [IBM 7030]. I guess that was before the CDC machines came in, so I think they just had the 7094 and the Stretch. But I never programmed the Stretch.

Brock: And so what was it like... what was your graduate experience like once you got to Berkeley in theoretical physics?

Lampson: Well, it was kind of weird, because I was in the process of being totally seduced by computing at the time In the late fall of 1964 the Fall Joint Computer Conference was in San Francisco, so I went to it, and I ran across a guy named Steve Russell. I can't remember how this happened, but I must have known him to some extent because I remember we chatted fairly freely, and he asked me, “What's going on with Peter Deutsch?” And I said, “Who's Peter Deutsch?” And he said, “Well, Peter's an undergraduate at Berkeley, and his father's a professor at MIT, a physics professor.” Peter had done a lot of programming when he was in high school and a lot of programming at MIT, and that's how he came across Steve Russell. But I had no idea. I'd never heard of Peter Deutsch.

And so Steve told me that Peter was working on the SDS 940 project, which wasn't called that at the time. It was called Project Genie. He told me how to find it behind an unmarked door in Cory Hall [home of the EECS department at Berkeley]. So after I got back to Berkeley at the end of the conference, I went and looked at this, looked... I went, opened this door, and you opened the door, there was a little airlock. There was another door. You opened that, and then there was a room, about the size of this room probably, with some desks in it and another door at the far end of that room, and I went through the other door, and then I was in a giant room with enormously high ceilings that had a few computers in it. I guess it was an old LGP... no, LGP-30?

Maybe it was an LGP-30. I can't remember exactly, but in the middle of the room there was an SDS 930, and sitting at it was a guy who turned out to be Peter, and he was reading in the paper tape. The whole paper tape was read in fairly quickly, and he got it out and wound it up and read it in again, and I said, “Why are you doing that?” And he said, “It's a two-pass relocatable loader,” and I said, “What?” And he said, "Yes, yes, I know. I know. I'm rewriting it." So that was how I got to know Peter and got introduced to the Genie Project.

Brock: And he was extraordinarily young.

Lampson: He was 16 at the time, I guess.

Brock: Yeah, yeah.

Lampson: 15 or 16.

Brock: And Steve Russell, who you just mentioned, was of Spacewar! fame and one of the early PDP-1 users at MIT. So...

Lampson: That's right. That was how I got to know Peter, who had been hacking on that PDP-1.

Brock: Okay. Who must have been a very young person when he was doing that.

Lampson: Yeah, yeah.

Brock: So that's very early on into your time at Berkeley that you got connected to this project, but what did you learn about it that made you want to be more involved in projects?

Lampson: Oh, yeah. Well, it was clear that what they were doing was very cool. They were trying to build a highly interactive timesharing system. Gosh. Yeah. Yeah. How could anything be cooler than that? So I got hooked up to it right away and that was the end. Yeah. After that I stopped really paying attention seriously to physics.

Brock: Was it after a while that you made a formal switch?

Lampson: Yes.

Brock: Yeah. And did you get a more formal place on that? You obviously did get a more formal place on that.

Lampson: Yes. Yes.

Brock: Yeah. Well...

Lampson: But I had a... what did I have? I had some sort of... I can't even remember what was funding me. It wouldn't have been a National Merit Scholarship because that was for undergraduates. I can't remember what it was. I had some deal...

Brock: That you could transfer.

Lampson: ...that I could transfer, right.

Brock: Well, seeing this highly interactive, timesharing system on a minicomputer as sort of the bee's knees, I was just wondering what the setup for you finding that to be so cool and enthralling? Was there Licklider had been writing by this point? Were you at all influenced by any of the things that...

Lampson: Not at that time, but I think the people that were running the project certainly had been influenced by Licklider and the project was being funded by ARPA.

Brock: Right. Well, maybe you could describe in these years... well, in these years from ‘64 to ‘67, I think you completed your dissertation in ‘67. So in those few years, you really had a number of publications and were very productive on, being involved in a number of software developments. The operating system for this, what became the SDS 940, the Project Genie effort, you helped create the QED editor for it.

Lampson: Actually, Peter Deutsch did all the work. I just wrote the paper because he didn't want to write a paper and I thought there should be a paper, so I wrote the paper.

Brock: Fair enough. But there was also an operating system called CAL that was based on JOSS.

Lampson: No, that was a programming language.

Brock: A programming language. I'm sorry.

Lampson: Yeah.

Brock: Of course. And also a SNOBOL system for the SDS 940.

Lampson: Yeah.

Brock: I was just wondering with all of that, what are kind of the threads that connect all of that body of work together? Did they hang together as a piece for you? How did they fit together?

Lampson: Well, we had this computer that was going to have an operating system that allowed it to have lots of different users that could do anything within the limitations of the SDS 930. The 930 had an address space of 16k 24-bit words., andthe operating system didn't put any constraints on how you could use it. You had character-by-character interaction with the terminal, and there was a file system, and you could write any program you want... you could execute any sequence of machine instructions that you wanted. You had the unlimited richness of 48k bytes of storage, and you could go to town. We wanted to be able to do interesting things, and, of course, we had to have programming languages because writing machine code for this thing, although it was definitely supported, was too hard. If you actually wanted to do a problem as opposed to hacking an operating system, you didn't want to write machine code. JOSS was very cool...

Brock: JOSS?

Lampson: …and I had what I thought were mildly interesting ideas for how to implement it very differently than Cliff Shaw had done originally.

Brock: For people who might not know about it, JOSS was this facility at RAND for people on a timesharing system. Or I guess it ran on the Johnniac, actually <laughs>.

Lampson: Yeah, that's what it was, Johnniac Open Shop System.

Brock: It was a very kind of interactive way for people to use a typewriter or a teleprinter to do mathematics, and very simple, interactive language. I was just wondering how you were familiar with JOSS?

Lampson: I must've read a paper about it. I never actually saw it in action, because you had to go to RAND to do that.

Brock: <laughs> Also, SNOBOL, I don't think people are as familiar with that today. Could you talk about...

Lampson: Well, SNOBOL was the go-to string processing language of the 1960s. It was invented at Bell Labs. If you think about computing, what computing was like in those days, most of it was numerical. The other poles of computing, I think, actually executable computing in those days, were strings, which was SNOBOL, and Lisp. We thought the machine was really too small for Lisp. Of course, Peter Deutsch was really into Lisp, and I don't think he ever built a Lisp for the SDS 940.

Brock: Well, could you...

Lampson: So the motivation for SNOBOL was that we wanted to cover more ground than just numerical computing.

Brock: Right, that makes perfect sense. So, upon the completion of your PhD, it seems like you were almost immediately invited to join the faculty at Berkeley?

Lampson: Yes.

Brock: Is that correct?

Lampson: I was one of Lotfi Zadeh's protégés, and he smoothed my path for me very well.

Brock: What was the department like when you joined it? Was it Electrical Engineering and Computer Science? Can you remember that?

Lampson: Lotfi had changed its name to EECS way early. May have been the very first EECS department. I'm not sure. But he definitely had a vision that computer science was the future.

Brock: How many...

Lampson: But of course, he got a lot of pushback from the rest of the department. So the whole situation, I think, was a little bit uncomfortable because of the pushback. Because the rest of the department definitely was not signed up for this proposition, which is why Berkeley acquired a competitive computer science department.

Brock: <laughs> Yeah. Well, I guess what I was wondering is once you became a faculty member, did things kind of just continue on as they had been with you working on Project Genie, or...?

Lampson: Yeah, pretty much. I taught courses occasionally. Assistant professors were expected to teach one course. I did that. So I taught a programming course, I think. That's what I dimly remember. I think I did it really badly, too.

Brock: <laughs> Why do you say that?

Lampson: Well, I've never really understood what you should teach undergraduates about computers. I'm fairly fearless about teaching graduate students because I figure that they're supposed to be able to take care of themselves, and if I mess up, they should be able to recover. But undergraduates are not supposed to be able to take care of themselves. If you're going to teach undergraduates, you really should have a clear notion of what they need to know. Well, I don't have a clear notion of what undergraduates need to know about computing. I never have had. So I taught programming hacks, basically, which is what I understood best at the time. Looking back on it, I think that was a terrible way to teach undergraduates. But I still find it very scary teaching undergraduates.

Brock: That's really interesting.

Lampson: I don't quite understand why so many people do it so fearlessly.

Brock: <laughs> Yeah. Well, maybe they do it without the same empathic standpoint <laughs>.

Lampson: Well, that's possible.

Brock: <laughs> Yeah. But the focus of your sort of research work was continuing the work on Project Genie, and...?

Lampson: Also, we had the CAL system running on the computer center's CDC 6400.

Brock: Were you becoming also more involved with the activities going on in the computer center?

Lampson: No.

Brock: No? Just...

Lampson: > The Cal system was an operating system research project, which had very little connection, I think, with the Computer Center's service activities. They had gotten some money from the NSF to do research, and so they got a second CDC 6400. They had two 6400s. I think it was one of these deals where you had two CPUs and a big... well, I'm not sure. I don't know how it was arranged. But anyway, they had a second 6400 that was pretty much available for research activities. We were trying to build a timesharing system for it.

Brock: This would be to be building a timesharing system that would be appropriate for a very large computing system? Back to the...

Lampson: Well, not so much that, actually. Our main interest was in pursuing the idea of capabilities. That was the focus... I'm not quite sure how this got to be the focus of the project. It was what Howard Sturgis wanted to do. He was the lead programmer on the project, and I think he must have recruited me to be the faculty supervisor, so to speak. Although I can't remember how it got started. So we built this capability-based system. Capabilities were an idea that had been originated by, I guess, various people, including Jack Dennis at MIT, and we picked up the idea.

I think the system we built was the first capability-based system that actually had users. There were a few that came along later than that, notably the CAP machine at the University of Cambridge. But ours was the first that really had users. Howard was employed as a programmer by the Computer Center, but he was really a mathematician. In fact, he wrote a paper. We have an English visitor named Shepherdson. So the Shepherdson-Sturgis paper pointed out that if you wanted to prove things about what Turing machines could do… writing a program for a Turing machine was fairly difficult, because they're so low level.

But you could write a simulator for a vanilla machine on a Turing machine. Not that you would ever run the simulator. But conceptually, you could write it, and then you could write the program for the real machine, for the simulated machine, and prove what they could do. That would tell you things about... a Turing machine could do, because it was doing those things when it was running the simulation. This seems like an obvious idea today, but it was pretty innovative then. So Howard was a very smart guy with very strong theoretical inclinations. But he also was a programmer, and I think he was basically... Howard provided the intellectual foundations for this project, and I was sort of following along in his footsteps.

Brock: <laughs> It was very much a contrasting approach for what was going on with Multics, for example?

Lampson: It was very much contrasting, yes.

Brock: So it was kind of a rival approach for kind of a...?

Lampson: Well, we didn't really think of it that way. The way we thought of it was that we wanted to explore. The practical justification for using capabilities has always been they would help with reliability. This has never been demonstrated, but intellectually, it's a very appealing idea. As far as I know, every time people have tried to put it into practice, it's been a bust. So I think it was Jim Morris who said, " the capabilities operating system's an idea of the future and always will be." Many of the bugs that show up in programs happen because some part of the program runs outside of the boundaries of what it's supposed to do, and tramples on some piece of the state that's supposed to belong to some other part of the program. It should never have gone there. So the whole idea of capabilities is to make that either very difficult or impossible by constraining what parts of the state the program can manipulate at any given point. The difficulty with that idea is that you have to plan out very carefully what parts of the state the program needs to touch. If you want to be hard-nosed about it, so that you can be confident that it's actually impossible for the program to escape its boundaries, you have to have a very clear idea of what it is that you're trying to accomplish.

Usually, you don't have such a clear idea. So the result is that the program is much harder to write because the standards are much higher. Typically, people have found that the game isn't worth the candle. So that's what we learned on the Cal system, I think, was that we carried out this capability idea pretty comprehensively, and we

actually documented it fairly carefully. The conclusion, I think we all came to it, certainly I came to, was that it was a bust. But intellectually, very interesting. I think there are still things that one can learn from the design of the **CAL** system. Many of them, unfortunately, about how not to do things. But subsequent attempts to build capability-based systems have, in my opinion, not paid enough attention to the **CAL** system, and have suffered by...

Brock: By repeating these...

Lampson: They've made mistakes that were closely related to the mistakes that we made.

Brock: <laughs> Right. Well, we wanted to ask you in those first years after you joined the faculty, Project Genie as such, as an ARPA-supported program, was that kind of running its course?

Lampson: No, we wanted to build a second-generation system, and we tried hard to do that. The reason that Berkeley Computer Corporation was started was we finally came to the conclusion that we would not be able to do this successfully in the university, because it was just too hard to hire people and buy things.

Brock: Were you talking to ARPA about wanting to do a second, to continue the work, to do another system?

Lampson: Yeah, and ARPA was fine with that, I think. I was not involved in those discussions. There was no indication that I know of that ARPA was reluctant to continue the project. But we just didn't think we could do it in the university.

Brock: That would be because of staffing, resources, all of these?

Lampson: All of the above.

Brock: Yeah, and that raises a whole <laughs> kind of interesting set of questions about creating new systems in an academic context, the timeframes, the staffing, the cost. So was the idea that maybe this new firm to create the second system would get an ARPA contract to help support the work?

Lampson: I don’t know that was ever seriously considered. I'm not aware of it. No, the idea was that we would get venture capital. Of course, there was no venture capital industry at the time, it was just getting started. So we were one of the first venture capital failures. But I think there were four companies funded by venture capital to build timesharing systems around 1970, and they all failed, not too surprisingly, because building a computer was actually fairly hard in those days.

Even minor details like trying to make sure that you could print from the computer system were hard, because there was no third party. Well, there were third party printers that you could buy, but they didn't actually work. The only printers that actually worked were the ones that were made by big computer firms, mainly IBM. So for example, the way the BCC system solved this problem was we actually leased a system, IBM 360 Model 30, so that we could print from it and use its tape drives, because we had had painful experiences on the Genie Project. Alternative strategies for getting printing done were all very unsuccessful. Getting data from the BCC system to the 360 was a pain, but at least we knew the 1403 printer worked.

Brock: <laughs> Well, a couple questions about that. I know that SDS kind of commercialized...

Lampson: They commercialized the...

Brock: -the SDS 940.

Lampson: Kicking and screaming, they commercialized the SDS 940, right.

Brock: The first Genie system, if you will.

Lampson: Yes.

Brock: But it did seem like many of the companies that were offering timesharing computing as a commercial service adopted the machine?

Lampson: Yes. Well, it was the first machine you could buy, that...

Brock: I guess that's true.

Lampson: -allowed you to run machine-language programs in... the competitor was BASIC, and several people had BASIC systems that allowed you to run BASIC programs under timesharing. But you couldn't run machine-language programs on those systems, because...

the security wasn't good enough. So if you ran a machine-language program on one of those systems, it could step all over the system, and if it had bugs, it was bound to step on something critical and crash it.

Brock: I did an interview once with Ann Hardy, who worked at Tymshare and worked with the system, <laughs> and she was describing how it was just the users trying everything that they could to crash the system. <laughs> Really exploring the machine, really seeing how far they could go.

Lampson: That was a chuckle for users. But even, I think, from a practical point of view, if users were trying to behave as well as possible, and not deliberately trying to crash the system, people were bound to make mistakes, and it didn't take very many mistakes to cause serious problems. The MIT AI group built a system called ITS for Incompatible Timesharing, by contrast with CTSS, and they put in a command to crash the system in order to make it less interesting, so that people would give up this attempt to find bugs.

Brock: That's hilarious. So it was such a known phenomenon that...

Lampson: Oh, yes. Absolutely.

Brock: Oh, that is really...

Lampson: Absolutely.

Brock: Well, I want to talk a little bit more about Berkeley Computer Corporation. But just a few questions before that, if I could? This is a very interesting time in US history, and with a lot of political and cultural changes going on around Berkeley in this period. So I was just wondering how all of that was kind of being felt by you in your life?

Lampson: Well, I tried to ignore it as much as possible, because I wasn't at all interested in the politics. I regarded what was going on on the campus as a big distraction. For the most part, I succeeded, I think. When I came to Berkeley in 1964, I hooked up with two other physics graduate students and a chemist. Jointly, the four of us rented an apartment on Blake Street, just to the west of the university campus. But two of my roommates were arrested in the first free speech demonstrations in the late fall of 1964. But I wasn't interested in that at all.

Brock: <laughs> Right. As you were getting so deep into these efforts to create this kind of new world of timeshared computing, interactive computing, I was wondering if you were finding other thinkers, other people from computing, or outside of computing, that were influencing your thinking, or your vision about where this whole thing was going? If you could just speak to that? Other people who you were thinking were doing... or just other people you thought who were doing important kind of work in this direction?

Lampson: Well, that was the Multics group at MIT, of course. I certainly paid a lot of attention to them. I think it took me quite a while to conclude that their project was way too ambitious. I guess it was only after BCC that I really figured that out, because, of course, BCC was way too ambitious too, for somewhat different reasons. The Multics people had visions about what computing ought to be, and some of those visions have still not been realized. They set out to realize their visions, and I think they just lost track of how hard it was going to be.

Brock: Was it on Project Genie that you met your future wife, Lois?

Lampson: Yes.

Brock: Could you talk about her involvement with the project and how...?

Lampson: She was hired as a sort of junior programmer-gopher. Actually, I think Mel Pirtle hired her to provide me with a girlfriend. He never told me that, but I think...

Brock: <laughs> You think that's true?

Lampson: I think that’s true, yes. He hired...

Brock: Wow.

Lampson: -Lois and Dana Angluin, who subsequently became a theoretician at Yale. Dana was for Peter Deutsch, and Lois was for me, and Lois was very success... it worked. It didn't work with Dana and Peter.

Brock: <laughs> I know subsequently, I believe Lois went into...

Lampson: She became an immunologist.

Brock: Immunologist, like a clinical immunologist?

Lampson: No, not clinical.

Brock: Research?

Lampson: She didn't have an MD, a researcher.

Brock: Research, yeah. That work began...

Lampson: That was her career, yeah.

Brock: -graduate studies at Berkeley?

Lampson: Yes, I think so.

Brock: Yeah. Well, could you also talk about your experience and participation in the ARPA community as a graduate student and as this kind of fresh professor? Just what that experience was like with you, interacting with the other people who were ARPA contractors?

Lampson: Well, ARPA used to run these get-togethers for graduate students and young faculty on ARPA projects, and they were enormously stimulating. I went to several of them, where I ran into people like Alan Kay and some of the Multics people, and so forth. I don't have clear memories of how those things went, but I think they were cool. That's about all I can remember.

Brock: <laughs> Okay. In the formation of Berkeley Computer Corporation, 1969, we talked about the initial goal was to build the next <laughs> system, but it was to implement some new ideas that...

Lampson: Yeah.

Brock: -everyone had about an interactive timesharing system?

Lampson: Yes, that's right.

Brock: I was wondering how unusual it was at the time for you, as a faculty member, to be starting a venture-backed company? It's become the...

Lampson: Well, I left Berkeley to... I took a leave from Berkeley...

Brock: Oh, I didn’t realize.

Lampson: -to start that company. Oh, yeah. I couldn't have continued to take a paycheck from Berkeley while I was working at BCC. It was a full-time job.

Brock: Okay, I hadn't realized that. Do you think it was just because now, just in contrast, now many people who are...

Lampson: Yeah, no, it wasn't like that at all.

Brock: Yeah, very different.

Lampson: Yes. If PARC hadn't come along, I might have gone back to Berkeley, but I never seriously thought about it.

Brock: Because it seems like you were advancing very rapidly, like getting a promotion, and everything seemed... so, could you describe what was happening to the company around 1971? The project had been going for...

Lampson: A couple of years.

Brock: -two years, something. Well, could you just describe what was happening?

Lampson: Yeah. Two things were happening. One, we were running out of money, and two, we actually had succeeded, somewhat improbably, in building a working system, but it was increasingly clear that it was not going to meet the original goals of supporting lots and lots of users. It became pretty clear, by late 1970, that it wasn't going to work. I heard subsequently from Mel and Wayne Lichtenberger, and so forth, that when they were doing the initial planning for BCC, they worked out how much they thought would be realistic. How much funding they would need to produce a viable company with a viable product. It was way more than we raised initially. I think at the time BCC was started, as I recall, the economy was doing pretty well, and raising money was not that hard. By 1971, it wasn't so great, and I think there just wasn't any realistic prospect of raising the necessary money. So that was one problem with the company. It was never really adequately funded to achieve its goals. Although, of course, I was blissfully unaware of this, being a young engineer entirely focused on technical problems. But in retrospect, looking back on it, the whole project was way too ambitious. It was a classic second system phenomenon. Not as bad as Multics, mind you, but <laughter> yeah, we had all these visions of what we were going to achieve that were not well-founded in reality.

Brock: Would it have required working on a third system to sort of pare things back in order to... if you had been able to raise money to go forward?

Lampson: Possibly. I never really thought about what it would take to turn the work that we did at BCC into something viable.

Brock: So could you talk about how the opportunity to go to Xerox PARC came up as you were weighing return to Berkeley, or continue with the company? <laughs>.

Lampson: Well, continuing with the company didn't seem realistic.

Brock: Yeah, I guess that was...

Lampson: It was pretty clear the company was going to go down the tubes, and it was even clearer that the system we had built, that was actually running and had users, would require a huge amount of work to make it viable, and there was no realistic prospect of doing that work. So I never seriously considered continuing with the company, and Bob Taylor, whom I had known from my ARPA days, came along and scooped us up, and told a very convincing story about how it was going to be great, and I didn't give it too much thought.

Brock: <laughs> So, what proportion of the team involved with the company went to PARC? Was it almost everyone, or...?

Lampson: It was almost all the top engineers. Chuck Thacker, Peter Deutsch. Charles Simonyi didn't quite go to PARC, but close. And a couple of other people, I think. Jim Mitchell. Those were the essential people. So it was a very strong nucleus.

Brock: Was that the initial nucleus for the Computer Science Lab story?

Lampson: Yes, absolutely.

Brock: There was...

Lampson: No, actually, Bob Taylor brought a couple of people from Utah with him when he came. These were the people who set the direction for CSL, initially.

Brock: In Bob Taylor's pitch to you all to come join him, this idea of "The Office of the Future", that that's what was going... the role of <laughs> computers in the office of the future. I was just interested to know, obviously, it was a successful pitch, but how did that resonate with you? Did you think that was a really interesting thing for computing?

Lampson: Well, nobody knew what "The Office of the Future" meant, of course, but that was okay. It had always been a dream of mine to encompass much more of the physical... when I did timesharing in the '60s, the only way you could get stuff communicated from the users to their machine and back was by Teletype. That was pretty bad compared to what you could do in an ordinary office, where you had at your disposal much higher quality typewriters, and pens and paper. Amazing <laughter> old technology, but you could do a lot of stuff with it that you could not do with a Teletype.

So at the Genie Project, we had a couple of graphics terminals and we thought that was important, but they were not very... by modern standards, they were not very good graphics terminals, and they were also very expensive. There were ways to use computers to at least do typesetting so that you could print... in principle, you could print beautiful stuff, and it had been a dream of mine to be able to do that sort of thing. But we never could do that at the Genie Project. It definitely seemed like "The Office of the Future" was a good slogan to use to overcome these obstacles and make the computer a A much less limited participant in the whole enterprise. When we heard about Gary Starkweather with his SLOT [Scanning Laser Output Terminal], of course, that seemed like a no-brainer to overcome a lot of these problems. However we were going to get around the graphics obstacles was not yet clear in 1971. But the goal seemed like an excellent one.

Brock: And one that would...

Lampson: Would not limit us in any way.

Brock: Yeah, I mean, it would serve for programming, for doing all sorts of literary tasks, but these facilities would be of use to people doing scientific or engineering computing also...?

Lampson: Of course, we were not really focused on scientific computing because people don't do scientific computing in offices.

Brock: True.

Lampson: We never really worried about that. In fact, a big gap in the work that we did at PARC was... a gap was filled very successfully by Unix. We never really thought about how you were going to... how users were going to program these systems. You couldn't program Bravo to any significant extent, and you couldn't program the drawing programs or any of those systems. But that seemed okay to us because, after all, people don't do programming in offices. So we didn't worry about that. But we pretty much covered the ground for people actually doing offices with them. One interesting exception, we didn't have spreadsheets, and I think that was because we didn't have any use for spreadsheets, because we didn't get along at all well with the accounting department at PARC. So it was definitely a potential barrier there.

Brock: <laughs> Yeah, to providing them with useful tools. I guess one thing I wanted to ask you about is, of course, at PARC there's this early episode with making MAXC and providing yourselves with a good computer for doing the kind of computer science work that everybody was interested in doing. But I think that it's this move from interactive timesharing to interactive computing in the guise of networked personal computing, what I think you have called distributed personal computing. That does seem like an interesting shift of conception for how interactive computing would be realized or delivered to people. Can you talk about that as a kind of conceptual shift, if it indeed was one?

Lampson: Well, there were always fundamental problems with the idea of doing interactive computing with timesharing. It was Jim Morris, I believe, who said, “The nicest thing about the Alto is that it doesn't run faster at night,” and that sums up one of the issues pretty well. Prior to the conception of the Alto, it had never really occurred to anyone that you could have a personal computer, and people who came to PARC after we had Altos were just blown away by that idea. We'd been developing it, so we were not so much blown away by it. But it was a radical idea. It has always seemed to me that the evolution of computing paradigm has been gated by hardware.

The Alto couldn't have been done any sooner than it was. Why? Because memory was too expensive. It was Chuck's brilliant insight that memory had gotten cheap enough that you could afford to give up half a megabit of memory for the screen. I remember when I was an undergraduate at Harvard and they acquired their IBM 7094, for some reason the price list for the 7094 showed up in the reading room, and I read it, and the price of memory on the 7094 was a dollar a bit.

Brock: A dollar a bit?

Lampson: A dollar a bit, a 1964 dollar for a bit.

Brock: Yikes.

Lampson: Right.

Brock: Wow.

Lampson: Wow, right. So, the Alto screen was half a million... if you tried to use 7094 memory to make an Alto screen, it was half a million dollars.

Brock: ‘64 dollars?

Lampson: Right. Maybe two or three million dollars of today's money. So that was a different universe.

Brock: That was also, in essence, this paradigm’s being gated by hardware. It's also gated by your imagination of where the hardware is going to go?

Lampson: That made it feasible, that made it not crazy to do this. Right, absolutely. You had to believe in Moore's Law.

Brock: Yeah, and I guess you were already seeing... and I know that was a big calculation going on inside Intel, was about that competition or cost per bit. They were constantly worried about cost per bit, versus cores, versus what they <laughs> were doing, and...

Lampson: Yes.

Brock: Wow, that's fascinating. While you were at PARC, you were at the center of these hugely influential developments, that ranged from operating systems to programming languages, word processors to printing protocols, editable document protocols to system modeling, and all within this framework of networks to personal computing, to distributed personal computing. It's kind of an incredible <laughs> bundle, and I was really interested to know what you think is most important for other people to understand about those technologies that you helped develop at PARC during this period? What's their real meaning for people today, or what should people know about that body of work?

Lampson: I don't know, it's in front of them every day. I don't understand the question.

Brock: <laughs> Well, I guess that is the answer to it. I have a question on this list about... I wrote an article recently for IEEE Spectrum about the 50th anniversary of the Alto, where I argued that we're still mostly living in the Alto world.

Lampson: Absolutely.

Brock: So, I mean...

Lampson: What's changed since the Alto? The answer is... well, leaving AI aside, the answer is the web. That's the only answer that I can see.

Brock: So what do you make of that? Because that's an incredible longevity for that set of technologies that were kind of really instantiated together. Why do you think they're so enduring, and why do you think there hasn't been something else? Because I imagine you never imagined that it would be that enduring when you were building these things.

Lampson: Well...

Brock: Or maybe you did.

Lampson: -our goal, really, was to put the power of computers behind... a lot of technology that had been around for centuries. Paper and ink is a very powerful technology, and our goal had been to computerize that. So we were lucky then, we were able to do that. But the goal is pretty straightforward. You could almost say we didn't need to have any vision beyond just that...

Brock: Because there was this accumulated centuries old kind of traditions of working with text, working with image...

Lampson: Right.

Brock: -and what was going on was just bringing that world, making the computer something that you could participate in that tradition with?

Lampson: Exactly, and some of it was quite hard. But it's not surprising to me at all that it lasted, because it was just the natural evolution of those technologies, which, as I said, had been around for centuries, and were extremely refined and powerful, and very deeply embedded in our culture. So we were just lucky that we were pretty successful in doing that.

Brock: Again, it's this idea that you're talking about... you were talking about, about these sort of paradigm shifts in computing being hardware gated? Or, yeah...

Lampson: Well, that's what made it possible.

Brock: That's what made it possible to make the computer into this tool for reading and writing and communication.

Lampson: Right, right.

Brock: I had a question about that, which was, this project into making the computer into a much better tool for writing and reading. Now, that appears to me to be a crucial aspect as we were just talking about what was

happening at Xerox PARC, but I was wondering if you think this project of making the computer into, which for shorthand I'll call a literary machine, do you think that was playing off of Ted Nelson? <laughs>

Lampson: Yes.

Brock: But a tool for reading and writing and communicating. Was that particular to PARC or do you think that's a larger effort in the history of computers?

Lampson: Oh, it's definitely a larger effort. Much of the work that's been done in graphics was directed to that goal. All the work that's been done on font design. The work that's been done on publishing technologies. All those are things that lot of people have worked on. Some of them, going back to well before PARC got started, and some of them, I think, certainly PARC's work, was very much inspired by Doug Engelbart's work. Some of the things we did were pioneered by him, and of course, Peter Deutsch and I both worked for Doug as graduate students, so we got to know his work pretty well. You might say that Doug developed many of these ideas first. He was very limited in the hardware that he was able to use, but many of the details of what he did couldn't be carried on. Also, Doug had a cult, which I think prevented, made it hard, for people to pick up on his... It certainly made it hard for him to evolve his ideas. I think after about 1968 he didn't really evolve them significantly because of the cult.

Brock: Not enough challenge there?

Lampson: Well, yeah. Doug had a vision that he’s going to augment human intellect, and I think that vision, I guess it was motivating initially, but I think after a while it got to be limiting and distracting.

Brock: And it seems like your approach to making systems has increasingly emphasized simplicity and then an open-endedness of the use. I don't know. At least that's how it seems to me.

Lampson: Yeah, well, pencil and paper is very open-ended, and we were definitely following that model. No question about it. Absolutely.

Brock: Well, let me just glance at my question list quickly. One of the things I thought was... so we talked about some of these larger contextual things about the work at PARC, the plausibility of devoting so much hardware to the screen, et cetera, and having this centuries-old model <laughs> of working with the printed word, or in written word, as something to bring into the computer, or to make a computer something that could participate in that cultural tradition, sort of big-scale context. I think there is a question that I have, and certainly I think other people will have about, to understand how it took place. What do you think were some of the ingredients that allowed PARC, Xerox PARC, in this time period, to be a place where these things came together? The way people worked together, the mix of people? What do you think was in that brew to allow this to occur?

Lampson: Well, if you ask, “What were the components of it?” I think the answer is Gary Starkweather, Chuck Thacker, and Bob Taylor, because Gary provided the foundation that made it possible to get the stuff on paper. Chuck provided the ability to develop hardware that met our need, whatever it was, and Bob Taylor made everyone work together on this enterprise that was pretty wide-ranging.

Brock: And I guess that's the sort of, I forget who I read described him as the impresario or something, but I guess that's one thing that is sort of lacking, let's say, in a university department. There would be a revolution if somebody tried to be <laughs> an impresario like that.

Lampson: It's very difficult to mount this kind of big project in a university. That's right.

Brock: I was also interested that at a certain point during your time working in PARC's computer science laboratory, you were actually moved to and were working remotely from Philadelphia.

Lampson: Yes.

Brock: Is that right? And so did you have an Alto connected to PARC...

Lampson: I had two Altos, actually.

Brock: Two Altos. <laughs>

Lampson: And we had a\t leased 9,600 baud telephone line that was connecting my two Altos to Palo Alto, and that was my computing setup. It worked fairly well, surprisingly, actually.

Brock: <laughs> And your wife Lois at that time had an appointment at Penn? Is that why you were...

Lampson: Penn, that's right. Yes, that's right.

Brock: And for how long were you living there?

Lampson: Between 1979 and 1987.

Brock: This next question kind of relates to that, which is I believe that the advances in computer networking and internetworking, including but also beyond the Ethernet made at PARC while you were there, are maybe underappreciated historically, and I'm thinking particularly about the PUP protocol.

Lampson: Yes, it's absolutely underappreciated.

Brock: And this large Xerox corporate internetwork that was built using PUP. Well, I'm glad you agree with that, <laughs> but could you talk more about that?

Lampson: For some reason, the PUP guys were not allowed to publish. I don't understand why, but they weren't. So the story goes that at that time, Vint Cerf was running a seminar at Stanford developing the internet protocols, and the PUP guy would go to that seminar and they weren't allowed to actually describe PUP, but they were allowed to ask questions. Supposedly, at one point after John Schoch or somebody had asked enough questions, Vint looked up at him and said, "You've done this before, haven't you?" Which they had. Absolutely.

Brock: So I think that's just not well appreciated at large in the history.

Lampson: I think that's absolutely true. I say it whenever I get an opportunity, but I don't have a big reputation as a networking guy. I'm not a networking guy.

Brock: But that predicated your working from Philadelphia for those years?

Lampson: Well, we depended on PUP for sure, yeah.

Brock: Well, I think it was in 1983 that there was this real rupture in the Computer Science Laboratory, and a whole group of people, yourself included, Chuck Thacker, ended up at a new laboratory at Digital Equipment Company. Could you just talk about your experience of that whole thing?

Lampson: Well, in some respects, it wasn't very disruptive for me because I was disconnected anyway, and I was able to continue working with the people I'd been working with before and it didn't make that much difference to me that they were working for DEC. The main difference for me was that I got to interact with lot of people in DEC Engineering. For the most part, I didn't interact very much with Xerox Engineering, so that was interesting, but it wasn't a big difference for me since I wasn’t involved in the management aspects of the thing.

Brock: Did you have to return your Altos or...

Lampson: Yes.

Brock: You did.

Lampson: I was sad.

Brock: What did you get to replace them?

Lampson: First, I got a VAX 11/730, which was not a big success, and then later on, I got a MicroVAX. That was better, so for couple of years, I suffered withdrawal symptoms, for sure. But then the Mac came along and that was okay because it had Word, and Word was Bravo, and the LaserWriter came along and so I was pretty comfortable after that.

Brock: It was a pretty significant group of you who went to DEC.

Lampson: Yes. Almost everyone who hadn't already left for Adobe.

Brock: So it was that.

Lampson: Yes.

Brock: It really cleared out.

Lampson: Very few people. Very few people were left.

Brock: There was certainly a negative story about what had happened in the laboratory to give an impetus out, but what was the positive story from DEC? How was this lab going to fit into what DEC wanted to do?

Lampson: We're going to do cool stuff. We'd already done a lot of cool stuff and were going to do a lot more cool stuff. Of course... Sorry. It wasn't nearly as cool as what we did at PARC, but I don't think that was really anyone's fault. We were very lucky at PARC. DEC supported the lab very well, I think. It didn't constrain it in any significant way. But the condition... There just weren't the same opportunities. In fact, if you look at what was going on worldwide between 1985 and 1995, nothing very interesting was happening. A lot of things got commercialized that had already been thought up. But yeah, nothing...

Brock: The start of the web?

Lampson: Well, 1995 was about, well, I think Bill Gates' Internet memo was written in early 1995, so that's why I picked that date. The web was the next interesting thing that happened, and it was at the end of that period.

Brock: But you were with DEC in almost just that period.

Lampson: Yes.

Brock: '83 to '95, but...

Lampson: That's why I picked it.

<laughs>

Brock: But I think you were, again, to my read, kind of ahead of things because you're moving into distributed systems. So clearly networked computing, the idea of what you could do with distributing computing over a network, this was becoming something that was concerning you or interesting you more. Is that fair to say?

Lampson: Yes, definitely, and I'm not sure whether we achieved this, but probably by 1995, distributed computing was still not really working. Leslie Lamport, you probably know Leslie's comment, that "a distributed system is one in which I can't get my work done because some computer has failed that I've never even heard of."

Brock: Yeah. <laughs> That's a good definition.

Lampson: And that's not true anymore. But yeah, was still probably pretty true in 1995.

Brock: Well, I was just wondering if you could share about, you know, was there one particular experience or what was the insight that led you to shift in that direction? What made you see that this was the direction in...

Lampson: Distributed computing?

Brock: Distributed computing. <laughs>

Lampson: Well, we had a lot of computers, so they had to communicate, and A and B. Let's see. We had a lot of computers. We wanted the systems to keep working. There's no way to do that except with a distributed computer, because if it's not a distributed system, most likely if some part of it breaks it's going to stop working, and then, of course, later on, the demise of Moore's Law has given a third reason for. If you want a system to do more computing it has to be distributed because you can't make individual nodes faster.

Brock: Right. You have to concatenate them.

Lampson: You have to get them to work together somehow. But I think the availability is the crucial thing.

Brock: And so in 87, you moved up here to Cambridge; is that right?

Lampson: Mm-hm.

Brock: And did that bring you into closer, did you work in a DEC office?

Lampson: No.

Brock: Still working from home?

Lampson: Well, no, that's not true, actually, because shortly after that DEC started its Cambridge Research Lab. So I had an office I could come to every day.

Brock: And so did that lead you to kind of more collaborations with more people within DEC, or were you...

Lampson: To some extent, but I was doing a lot of that before because I was coming up and visiting, for example, visiting the network group and talking to them about their problems, making suggestions, and I think that had a fair amount of impact. DEC had this thing called the Systems Task Force, which was chartered to understand the engineering budget and make recommendations for, which was moderately successful, I think, and I was on that so I would come up here pretty regularly for those meetings, and I got to know a lot of the senior engineers pretty well that way. But, of course, the trajectory of DEC between 1984 and 1985 was a sad one. But I put most of the blame for that on The Wall Street Journal.

Brock: How so?

Lampson: Well, every couple of months in the mid-'80s, they would run a front-page article explaining how great DEC was and how DEC was going to blow IBM out of the water, which was nonsense. But unfortunately, the DEC management believed those articles, and they... in a way, it's ironic because when Xerox bought SDS, they had the same idea, that they were going to blow IBM out of the water, with much less justification. But both ideas were crazy.

Brock: Yeah, I wonder if an alternate strategy would have even been allowed? To just say, “No, we're going to specialize in a different area for DEC.”

Lampson: For DEC.

Brock: Yeah, rather than trying to expand it.

Lampson: Oh, yeah. If you think about what happened starting in the mid-'80s, workstations and networking, those areas that DEC should have been able to own, because DEC networking was way better than anyone else's at the beginning of that period. Unfortunately, DEC got sucked into the OSI madness, and I never completely understood why. The only explanation that I could come up with was that, well, aside from this Wall Street Journal thing, which was a huge distraction for the company... The other explanation I have for it is that they lost Gordon Bell, and I think that in the late-'70s, Gordon provided the focus for DEC Engineering, and he was quoted as having said that he wouldn't try to stop bad projects unless it was absolutely essential to have their resources for good projects.

But I think it was really Gordon's vision that was responsible for the success of DEC in the early-'80s and the lack of that vision was fatal. There was no one who could replace Gordon, because no one else had the vision and the credibility.

Brock: And maybe the institution had grown up around having him there, so there’s a void.

Lampson: Well, that could be. I wasn't at DEC, of course, while that was happening.

Brock: Well, it seemed like around '95 that there was, again, a group of people who had been associated with the Systems Research Center at DEC, and maybe some other places at DEC who went over and got involved, then joined Microsoft's research.

Lampson: That was later.

Brock: That was later? Okay.

Lampson: Yeah, Microsoft resisted starting a research lab in the Bay Area for many years and it wasn't, I think, around 2000, maybe. I didn't realize that. Yeah, because I joined Microsoft in 1995, and I repeatedly pushed for this with no success.

Brock: So which part of Microsoft did you join?

Lampson: That's a funny story. I did not join the research organization. I was hired to be Paul Maritz's technical associate or advisor or something. At the time that I was hired, Paul Maritz had a staff job, but shortly before I actually showed up, Microsoft had a big reorganization, and Paul went from having a staff job to having half the company reporting to him, and so it didn't seem that having me 2500 miles away and trying to give Paul Maritz technical insights was going to work, because he had so much management responsibility. So they gave me to

[CTO] Nathan [Myhrvold], and I was perfectly happy with that. I started kibitzing around and spending a few days a month in Redmond and... I'm sorry, I lost the thread.

Brock: But then was it other people from the DEC lab came into Microsoft around 2000? So after you.

Lampson: No. Microsoft, actually, decided that they would have a Bay Area research lab, but I think it was around 2000 that they started Microsoft Research Silicon Valley, because they started the Cambridge lab and the Beijing lab before that.

Brock: And is that when...

Lampson: Might even have been a little later than 2000.

Brock: And is it around then when people like Chuck Thacker and Leslie Lamport and...

Lampson: Well, Chuck joined Microsoft in order to help start the Cambridge lab.

Brock: I didn't realize that.

Lampson: And he spent a couple years in Cambridge.

Brock: And that's where you two collaborated again on the tablet PC; is that correct?

Lampson: No, that was later. That was after Cambridge. Yeah, after that Chuck went back to the Bay Area, but he's still working for Microsoft.

Brock: And I believe, from my reading, that seemed like a significant project for that you two collaborated on.

Lampson: Yes, we did.

Brock: Yeah, what did that project mean for you? What did you find most meaningful about it?

Lampson: I wanted to have a tablet PC for many years, and we actually did a project, Chuck built a primitive version of the tablet in DEC, which it was completely impractical. It was basically only a reading device. I guess Bill Gates had been interested in the idea of a tablet PC, and so he chartered a Microsoft manager named Dick Brass to run a tablet PC project, and Chuck and I piled into it, and Chuck built the prototype tablet PC.

But the Windows group never got engaged in it, and the problem was that at that time, a tablet PC was going to be sufficiently expensive that it pretty much had to run Windows. If it was any feebler than that nobody would buy it because it would cost too much, and so the result was that although the tablet PC stuff was originally developed by an independent group, but building on Windows, subsequently the effort was folded back into the Windows group, and they didn't take it very seriously.

It was all part of Microsoft's failure to understand that handheld devices, and notably phones... What's the right way to say this? Microsoft was into phones from early on, but at that time, the major cell phone carriers were gatekeepers for phones and Microsoft got in bed with them. Well, you do not want to be in bed with a cell phone carrier because they're brain dead, and the phone companies have been brain dead about computers ever since the '60s.

They've always had visions. AT&T got into computers and they bombed out completely, and so the Microsoft phone people got captured by the carriers because the carriers were basically their customers, and so they were unable to, they were unable to conceive of anything like the iPhone. The only reason the iPhone was successful was that AT&T was desperate. They were losing out in the cell phone business in a big way, so they were willing to cut a deal with Apple that effectively gave Apple control of the device; they would never have cut that deal with Microsoft because they were terrified of Microsoft, and they were not terrified of Apple.

Brock: So interesting.

Lampson: Oh yeah, they were convinced that Microsoft wanted to take over their business.

Brock: Right. That's interesting.

Lampson: Of course, they should've been terrified of Apple, but Apple wasn't a very big deal in those days. Basically, Apple had the iPod and that was about it, because Macintosh was very successful in certain limited domains, but globally it was not a big deal.

Brock: Right, in terms of market share, yeah.

Lampson: Right, exactly.

Brock: Yeah, that's fascinating. Well, and so would you say that during your tenure at Microsoft, outside of this tablet PC effort, were you, did your concentration,remain in issues around distributed computing and computer security? Would you say that was the...

Lampson: Yeah, pretty much.

Brock: The emphasis?

Lampson: Pretty much. Towards the end of that period... I have written a tome. In the early '80s, I wrote a paper called "Hints for Computer System Design," which I didn't take too seriously, but a lot of people really liked it and it won one of the first Test-of-Time awards at SIGOPS. For a long time after that I wasn't particularly motivated to change that paper. As I said before, in the '80s, nothing much interesting happened in computing, but later on, more interesting things started to happen, and I was motivated to write a new version of that paper. I put a lot of effort into that.

Brock: And it's quite long, isn't it?

Lampson: It is.

Brock: It's some hundreds of pages, or a hundred, two hundred pages, something like that?

Lampson: Yeah. I think the long version is about a hundred pages of text. There is a short version that's still longer than the original paper, because, of course, I know a lot more now. I tried to put everything that I know into that paper. Maybe that was a mistake, but that's what I did.

Brock: Well, I was looking at, I think, the revised version of that, which I thought was very nice, that was freely accessible, and it's a very interesting mix of perspectives, or at different registers, I guess, the sorts of advice and perspective that you're giving, from kind of a broader level to some more detailed levels. I don't know. I was interested in how it ranged through those.

Lampson: Yeah, well, system design involves all those levels.

Brock: That's true. <laughs> Well, I wanted to say that... This is something we talked a little bit before the camera turned on, but I noticed that a phrase conspicuously absent from your CV is artificial intelligence.

Lampson: I've never worked on that, no.

Brock: <laughs> I was wondering what was behind that.

Lampson: Why have I never worked on it?

Brock: Yeah, was there something...

Lampson: Because I never had any ideas.

Brock: Okay. <laughs> That's very simple. I did want to ask you a question that does relate to your career-long interest in system design, which was an article that was published recently in Science Magazine with, I think, you and some other co-authors…

Lampson: Yeah, Charles Leiserson. Right.

Brock: ...about the future of computing performance after Moore's Law, and I was just wondering if you could just talk about that vision for the future of computing and performance gains coming from the top rather than the bottom of a system. Could you explain sort of what the concept is there?

Lampson: Well, the way I like to think about it is if you look at today's PC, it's about 100,000 times more capable than the Alto in every dimension. Processing cycles, main memory, storage, networking, across the board. It does maybe, what, 50 times as much? So that leaves a factor of 2,000 unaccounted for. Where did that go? Well, I don't know where some of it went. Some of it went, you know, into things like high-quality worldwide typography, it

consumes a lot of cycles. A bunch of things that, but there's no question that a lot of it's pissed away. So there's a lot of room for improvement. That's the top line.

Brock: I thought it was interesting in the article, too, that saying a couple things. That while for almost all of your career in computing, the community could basically count on this cadence of improvement in semiconductor technology, the same computing for much less cost, more computing for the same cost. It just was on a kind of a metronome. But I thought the observation in the paper that these computing performance improvements at the top of architecture from software, that one could expect that to not be as regular, that we should foresee more punctuated or spiky sort of improvements, and I think, to me, that seems interesting because it's not a situation that the community has really faced maybe since the early '60s. I just was wondering how you thought, how big of a change you think that might be, that irregularity.

Lampson: Well, I don't know. The fact that there's this factor of 2,000 says there's no particular reason to believe that for anything concrete that you want to do, the availability of hardware is likely to be a constraining factor. Availability, in particular, availability of computing cycles, is unlikely to be constraining. So other things are going to be much more important, and the fact that we might have to think more carefully about things, that we've gotten away with murder for 30 years, thanks to Moore's Law, and with relatively few exceptions. Even things like deep learning, where there's some definite things, are relatively smaller, people are focused on squeezing as much performance out of the hardware as possible. Even there I think the biggest gains have come from getting more insight into what it is that you're actually trying to accomplish. So things like ResNet and transformers have been much more important for the... But certainly, if you want to... So I don't see the, lack of Moore's Law is really that, I don't think... raw performance is just not that critical. I think other things are much more important.

Brock: That's fascinating. Well, I did want to, you know, we talked a little bit about joining the ACM back in 1961.

Lampson: 1961.

Brock: and I was just wondering if you could tell us a little bit about how being a member of that professional community has served you and what it was like for you to receive the Turing Award.

Lampson: Well, of course, receiving the Turing Award is a big deal. It's a bigger deal now than it was when I got it in 1992. I guess it wasn't a big surprise when it happened, because thinking back on what we'd accomplished at PARC, it seemed like a natural thing to be rewarded for, and in fact, four people from that community have gotten the Turing Award, Chuck Thacker, Alan Kay, and Bob Metcalfe, as well as myself, and they could've given it to Warnock and Geschke too. Too late now.

Brock: Unfortunately.

Lampson: But <pauses> so yeah, it was cool, but it wasn't a big surprise.

Brock: Well, I wanted to maybe close by asking you to talk about something that I was lucky enough to talk with you a little bit about before the camera was rolling with some of the work that is happening today in deep learning

and these enlarged deep learning systems like large language models, that these kind of systems may involve a change in the character of computing and computing research.

Lampson: Yeah, it seems to me that we don't really understand how they work. This is a new phenomenon in computing. In the past we've always understood how things work, and it's always been possible to burrow down through the levels of abstraction all the way to the gates and see exactly what's happening and why. With deep learning, with large language models, or even small models, I think that's no longer true. We don't understand why they work. Yeah, it's true that we can print out all of the billions of weights of the net that's been constructed, but they don't mean anything, and we don't understand what the consequences of tweaking one of them is, for example. So I think it's fair to say that we really don't understand how it works, and if that continues to be the case, it seems to me computing will become much more like biology than it has been historically. Of course, another possibility is that we will understand much more how they work, and then maybe that won't happen. It's up for grabs.

Brock: That seems like a big, open question in the future. <laughs>

Lampson: I think so. Very big, open question.

Brock: Yeah, yeah. Well, thank you very much for agreeing to do this and for being here with me today for this interview.

Lampson: Well, it’s always interesting to think back on what happened 50 years ago and what it means today. So I'm grateful for the opportunity.

Brock: Well, thanks again.

END OF THE INTERVIEW

