A. M. Turing Award Oral History Interview with Fernando J. ("Corby") Corbató

by Steve Webber

Plum Island, Massachusetts

April 6, 2018

Webber: Hi. My name is Steve Webber. It's April 6th, 2018. I'm here to interview Fernando Corbató – everyone knows him as "Corby" – the ACM Turing Award winner of 1990 for his contributions to computer science.

So Corby, tell me about your early life, where you went to school, where you were born, your family, your parents, anything that's interesting for the history of computer science.

Corbató: Well, my parents met at University of California, Berkeley, and they… My father got his doctorate in Spanish literature – and I was born in Oakland, California – and proceeded to… he got his first position at UCLA, which was then a new campus in Los Angeles, as a professor of Spanish literature. So I of course wasn't very conscious of much at those ages, but I basically grew up in West Los Angeles and went to public high schools, first grammar school, then junior high, and then I began high school.

Webber: What were your favorite subjects in high school?

I tended to enjoy the mathematics type of subjects. As I recall, algebra was easy and enjoyable. And the… one of the results of that was that I was in high school at the time when World War II… when Pearl Harbor occurred and MIT was… rather the United States was drawn into the war. As I recall, Pearl Harbor occurred… I think it was on a Sunday, but one of the results was that it shocked the US into reacting to the war effort, and one of the results was that everything went into kind of a wartime footing. And one of the things that happened was that the high school which I was at began to offer, as I recall, an A and a Z period, where the A was an early session and the Z period was a late session, to accommodate students who were going to work in local industry and essentially aid the war effort.

Webber: Did you do any of that yourself?

Corbató: I took advantage and sped up my program and proceeded to graduate from high school in two years rather than the usual three. I think I had to go to a couple summer school classes to squeeze it in, to finish it up. Basically, as I recall, I took solid … – solid geometry – and trigonometry as my two summer

school courses. In any case, I graduated from high school in basically two years rather than three.

Webber: What other things did you do in high school besides study all the time?

Corbató: [laughs] Well, I was… I don't recall anything particularly. I wasn't into athletics, but I took advantage to start UCLA in that following fall. One of the consequences was that I was going to major in physics, although I don't recall… it wasn't very specialized at that early stage.

One of the next pivotal events was some Navy recruiter came by and outlined being able to join the so-called Eddy program, which was a 12-months program that the Navy was organizing to allow students, to train students to be technicians to service all of the fancy equipment that was being deployed in the Navy. It was called the Eddy program. So after seven months of the first year at UCLA, I decided that that was an important step for me to take. It meant I would be able to join the Navy rather than be drafted into the Army, and that seemed preferable at the time. So I did join the Eddy program and proceeded to start on this 12-month program.

Webber: What kind of stuff did they do in the Eddy program? What kind of a…

Corbató: Well, it was a very carefully… It was named after a Captain Eddy thought up the whole idea. The basic idea was to train a cadre of technicians who could service and maintain the very fancy electronic equipment that was being deployed to the fleet. Things like radar, LORAN… I've lost track of the litany at the moment, but it seemed like a great opportunity to get an education and stall off… and also pick my service so I wouldn't have to go in the Army and be potentially some sort of a foot soldier.

Webber: How long did you stay in the Navy then?

Corbató: Well, the Eddy program itself consisted of about 12 months of training. The first two weeks I think were up at Great Lakes training center, where you basically got used to wearing a uniform and being in the Navy. Then the next month was so-called "pre-radio," where you were down outside of Chicago, and where I basically got… It was basically a review of algebra and very early mathematics. [0:10:00] Then it was followed by six months of training at the Del Monte Hotel, which had been commandeered by the Navy, which was in Carmel, California. That was a very famous resort hotel. I spent three months there, learning, basically reviewing elementary mathematics and very straightforward things. Then followed by six months up at Treasure Island, which was an artificial island that had been built off of Yerba Buena Island at the middle of the San Francisco Bay.

The six months at Yerba Buena Island were spent learning, reviewing the detailed circuits and logic of the various pieces of equipment that we were probably going to have to service when we got out deployed into the field. And after the six months at Yerba Buena, I was assigned to a pre-commissioning crew of a destroyer tender which was being built up in Tacoma, Washington. And we were the so-called pre-commissioning crew. I remember we had to load all the supplies on the ship and the ship was being finished as we basically graduated. It was as I… Well, I've lost track of that train of thought.

Webber: After the Navy, then you went back to UCLA?

Corbató: No. After the Navy, I got out in summer of… about May of 1946, and by then, I had the opportunity to use the GI Bill, which allowed me to go to Caltech, which seemed like an expensive but good choice, but the GI Bill was covering my college tuition. So I started Caltech and I expected to be a… I believe I was majoring in physics and had a nice, pleasant four-year time at Caltech, until I finally graduated in 1950.

My next step was to… I applied to graduate schools in physics, and I remember being admitted to MIT. That sounded like a good opportunity. And, in the summer of 1946, I proceeded to load up my car, jam-packed with all my stuff, and proceeded to drive cross-country by myself.

Webber: So that was in 1950 after graduating from Caltech?

Corbató: That would be 1950 after I graduated from Caltech. Yes, correct. In retrospect, I'm somewhat unnerved by my daring approach, because my car wasn't that… it was a used car and it wasn't that overly reliable. I had to… actually I remember sleeping on the road. I would just pull off to the edge of the road and lay down a sleeping bag and go to sleep. I drove cross-country by myself without… what I would consider now a very hazardous trip. [chuckles]

Webber: It's interesting. I did the same thing going out to my graduate school, and it took me three days to drive across the country all alone, sleeping in the car when I could, same thing. So when you got to MIT, what did you do that summer before entering MIT?

Corbató: Did I do what?

Webber: What did you do in the summer after you arrived in Cambridge I assume or somewhere around there before actually going to…?

Corbató: Oh, it was the end of summer…

Webber: Oh, it was the end of summer. Okay.

Corbató: …that I arrived, and so I had a… I don't know how it got arranged, but I had a… I was assigned to a dorm room in the graduate house and proceeded to share a dorm room with the late Henry Kendall[1], who died in a scuba diving event years later, and Dan Willard and we roomed together for basically the first year. Then I proceeded to… I was persuaded by a fellow graduate student, John Little, to join him in sharing an apartment on Beacon Hill. So I decided I would do that and I proceeded to spend several years in the apartment on Beacon Hill.

Webber: What were you studying at MIT? Was it real physics or had there been any computer stuff yet?

Corbató: No, there was no… [coughs] there was no… there was no particular computer courses at that time. Computers were still something that was evolving. But I did… I did come under the wing of Philip Morse, and he proceeded to get me a fellowship on using computers.

Webber: What was Philip Morse's job? Was he the head of…?

Corbató: Philip Morse, I think he was the executive officer of the physics department. But he was a very well-known figure in physics. And he was a bit entrepreneurial. He proceeded to cadger out of the ONR a series of…

Webber: The ONR, the Office of Naval Research. "ONR" being the Office of Naval Research? Or what does "ONR" stand for?

Corbató: I think it was "Office of Naval Research." But he got this set of fellowships which… and I had one of those. [0:20:00] And one of the things I remember best is that we were allowed to do programming on the Whirlwind computer. The Whirlwind computer was built as a prototype of the SAGE system computers. What made Whirlwind rather unique is that it was a 24-bit parallel computer that is… whereas many of the computers in the early days had been serial computers, which meant that they were considerably slower. But the goal of the Whirlwind computer was in trying to basically create a system which could do real-time computation for air defense. It eventually led to the MITRE system.

Webber: When you were getting into programming the Whirlwind, what did that consist of? What kind of programming was that?

Corbató: Well, it was… programming consisted of creating a punch paper tape, so-called Flexowriter tapes, which would be read into the computer by a fast reader and then proceed to execute on the computer. The main thing that occurred was that because the computer was being used to develop the air defense system that eventually became MITRE, we civil-… the people that were able to use Whirlwind at that… we weren't under any security blanket or… and

[1] Later a Nobel prize-winning physicist. (Ed.)

so we had to use it on the night shift basically from 12 to… from midnight to morning.

Webber: Were you actually programming the eventual real software that was going to be used, or were you just practicing and programming for the fun of it?

Corbató: We were doing a real software. Not for the air defense system. That was all under wraps. We weren't even supposed to know about it. [laughs] But we were allowed to use the scraps basically of the middle of the night. That meant I got used to working crazy hours and sleeping late at night… sleeping in the morning, and getting up later in the day in order to start my midnight shift there.

Webber: After Whirlwind, what was the next computer-related thing you did? I mean how long did you do Whirlwind? What time was that?

Corbató: I did through my graduate work, and I proceeded to do some computations on so-called molecular orbitals. It was really kind of dull. [laughs]

But in any case, by the end of '56, I had gotten my doctorate and I proceeded to… The next thing that happened was that Morse in his entrepreneurial self had convinced IBM to establish what was then a major computer called the IBM 704 at the Computation Center, which was… In fact, he came in late in the game and he got this IBM computer. His deal was that MIT would run this computer, MIT would build a computer room, and would operate it for the benefit of itself and some roughly 40 New England colleges who would have time on the computer. The result was that, my recollection is that Morse had enough clout that he got MIT to essentially build it on the ground floor of a new building, Building 26, which in the style of the day was built on stilts. So Morse persuaded MIT that they had to put in a ground floor covering up the stilts, and the architects were of course incensed, so they retaliated by making the ground floor have no windows…

Webber: [laughs]

Corbató: …and had only transoms, a few transoms, and making it all in blue tile to make it sort of look like it wasn't there.

Webber: So that was MIT's first computer center?

Corbató: Computer science, it was basically a 704. As the computers evolved, they kept… MIT… In those days of course computers were all batch-processed, so MIT basically ran the computer and people would at best come down from the New England colleges and submit jobs, and then go back. But it was not a great arrangement for the New England colleges.

And there were other things that happened. Nobody had counted. One shift was supposed to be for MIT, one shift was supposed to be for the use of the New England colleges. But there was a fourth shift, basically the weekend, which no one had remembered to count in. So we scraped up all the crumbs and had a lot of computer time. But…

Webber: Who is "we"?

Corbató: Basically, "we" was the… we operated the 704 for the benefit of the MIT personnel around the campus who had to learn to program, had to learn to… I guess people did their own keypunching. Basically, it was all very primitive by today's standards. Pretty awkward for the people at the New England colleges, so some of them would come down, stay for a day or two, and then go back. [0:30:00]

One such person was John McCarthy, who was sufficiently interested in trying to use the computer that he proposed the notion of time sharing. That seemed like a very useful idea.

Webber: When was this that he came up with this concept of time sharing?

Corbató: I believe he proposed in a symposium that was run in 1962. Yeah, 1962, I think.

Webber: And was he at MIT then or was he at another school?

Corbató: He came. He decided that it was important enough for him, he moved from Dartmouth to MIT. So he proceeded to… Yeah, and I shared an office with him for a while. But…

Webber: That's interesting. We'll get to it in a bit about CTSS, but Dartmouth had their own time-sharing system called DTSS about the same time. Was he involved in that, do you know?

Corbató: What happened was that there was a propo-… I guess it was proposed to NSF that we build a time-sharing system. The person that was designated to lead that was Professor Herb Teager. Teager was I believe a professor of electrical engineering. So Teager started working on it, but Teager had a somewhat difficult personality. He tended to rebuff… He tended to be a loner. He liked to work alone. He didn't work with other people very easily. So Teager proceeded to play his cards close to his vest and he had started imagining what kind of a time-sharing system he would build. Unfortunately, the more I heard of it, the more grandiose it seemed. He was going to have language translation and all kinds of elaborate ideas that were way beyond the state of the art. And it was about that time McCarthy got fed up at MIT's lack of attention to

his vision of what should happen. So at that point I think, McCarthy went out to Stanford, and so things were kind of in the lurch.

One day I came up with a notion of trying to build a very simple-minded time-sharing system, which we could sneak in and run with the standard computing system with a few caveats, namely that it had to have a smaller amount of memory. I had to set aside 5,000 words for the supervisor program.

Webber: This was on a 7090?

Corbató: I think it was probably a…

Webber: A ’94?

Corbató: No, not yet. It was a 70-… either a ’4 or… It was a 704, or maybe a 709.

Webber: Okay. How much memory did that have of the 5,000? How much more was there?

Corbató: Just a bank of 32K.

Webber: 32K. Okay. And these are 36-bit words?

Corbató: 36-bit words. In order to… I came up with this plan to build a time-sharing system where we snuck out 5,000 words for a supervisor program, which would then become… and the remainder of the memory was for… Having a little trouble retracing… [laughs]

Webber: “For the user image while it was running” or something? What type of hardware additions were required to make it work? I know we needed interrupts and stuff.

Corbató: Very good question. One thing that was required was basically a relocation register, which would allow one to… basically to modify the address that programs could run in, so that you could run programs that were not all at the beginning of memory. A second thing that was required was some sort of a real-time clock so that you could run programs and interrupt them after a certain amount of time. And the key result was, and I forget the chronology here a little bit, but we managed to get another bank of 32K memory so that that allowed us to basically have…

Webber: 32K for the supervisor and then 32K for a user? Something like that?

Corbató: Umm…

Webber: Did IBM make hardware changes for you because you asked for it to make this work?

Corbató: They did. They did for everything but the clock. I think we had to build a clock. One of our more clever engineer clever people managed to hook that up. In any case, we managed to cobble together a system which could basically do swapping of programs. As I recall, the first swapping was we would swap each user out to a magnetic tape, and we could swap in another and run it. It was all very crude and primitive, but we got the system up and running. And I think that was the impetus to get a second bank of 32K words, because we didn't have to swap people out of memory.

Webber: This sounds like the 7094 the CTSS was running on. Was this the system that eventually was CTSS, the one you're describing?

Corbató: We eventually were able to get a sufficient time-sharing system going that Bob Fano saw interactive programming as clearly a way to go, and he proceeded to propose Project MAC. We in turn had… See, Project MAC was first propo-… I think Fano wrote the proposal [0:40:00] in the fall of '63 and then proceeded to…

Webber: What was Bob Fano's role? Was he another professor? Was he…?

Corbató: Bob Fano was a professor in the electrical engineering department who had done outstanding work which I was not particularly familiar with in information systems. But he was sufficiently… sufficiently intrigued by the notion of interactive computing that he wrote the first proposal for Project MAC. And it was to be a… And he chose to start it off by having an invitational summer program for people all around the country to come visit and sort of kick the tires and see what time sharing was like using CTSS. So he proceeded to start up Project MAC, initially on, I recall, the 11th floor or something – or the 10th floor? no, not 10th, 9th floor – over at the…

Webber: 545 Tech Square?

Corbató: …545 Tech Square, where we kind of looked out. There was a… One of the service bureau companies had overambitiously leased out space on the building and they suddenly had hit hard times. So basically Fano was able to get two floors of 545 Tech Square because of their…

Webber: Bad luck? [laughs]

Corbató: Yeah.

Webber: Yes, I remember. I was on the 5th floor for years and years where your office was.

Corbató: We gradually evolved and took over more and more floors of the building.

Webber: Getting back to CTSS, in "Project MAC," what does "MAC" stand for? I've heard a couple of acronyms that…

Corbató: There was a lot of jokes about that. I don't re-… One was "Minsky Against Corby."

Webber: Yeah. [laughs] That would be Marvin Minsky, the linguist.

Corbató: The late Marvin Minsky.

Webber: The one I choose to remember is "Man and Computer," but it was also the first three letters of "machine." And there was a lot of…

Corbató: A lot of, right.

Webber: So Project MAC started in '63, and by then, the CTSS was already there. What happened next?

Corbató: Well, we ran the summer study. We invited a bunch of people to come visit. I didn't of course, but people like Dick Mills, who was Fano's associate director. And we of course … things were happening so fast, we didn't really have a computer. We only had the system over at the Computation Center. So we proceeded to basically borrow the computer time off the Computation Center, which was running normal batch processing, and we proceeded to operate it at a distance using… And as I recall, the modems in those days were terribly clumsy and big boxes.

Webber: Yeah. 110 baud. Incredibly slow.

Corbató: Slow and the like. But…

Webber: When did you get your own computer?

Corbató: It didn't show up until about October of 19-… of that…

Webber: '63?

Corbató: Sixty-…

Webber: And that was put on the 9th floor?

Corbató: And that was put on the 9th… Yes, I think so. Yeah, the 10th floor was what, refrigeration or something?

Webber: Yeah. I think the 9th floor… I remember going up to the 9th floor and actually putting magic numbers into the switches on that computer with Gerry Clancy. Gerry Clancy had written some of the supervisor for CTSS, and if the right pattern were in the switches, it gave Gerry Clancy privileges, all the privileges he needed. So at one time we needed to do something seriously upstairs doing the switches. There was no keyboard or anything.

Anyway, so you got your own computer in the fall of ’63, and you immediately, I assume, brought CTSS up on it?

Corbató: Yeah. Well, it was actually October-ish. At that point, we could run… we could basically let the Computation Center have all of its machine back. Actually, we may have run it for a while more. But in any case, the… the computer system, the one thing that was sort of noticeable about it was, in order to keep our sanity, the computer of Project MAC had red panels and the computer back at the Computation Center had blue panels. So we used to refer to them as the red and the blue machines.

Webber: Sounds a little political these days. [chuckles]

Corbató: Yeah, right.

Webber: Go on. So…

Corbató: No, that’s okay.

Webber: Now we have a compu-… Is it true that eventually the Computer Center machine ran CTSS all time as well?

Corbató: I don’t remember…

Webber: Maybe, yeah.

Corbató: I don’t remember that it did. It had an extra bank of core memory. So it possibly could have.

Webber: So when you got your computer in ’63, did it have disks? Or was it all mag tape, or…? I think you might have been swapping out the disk by then.

Corbató: Yeah. We had begun… We had somehow managed to start storing programs on disk rather than one tape per user. That was of course extravagant, just as a demo, to use a tape per user.

Webber: And IBM was working with you on all of this, trying to get all their stuff involved?

Corbató: Yes. Good point. IBM was working with us, and one of the things that we shortly began to plan was, it was decided that Project MAC should plan a new machine meant explicitly for time sharing. IBM sort of thought they had us in their pocket. But in fact, we went and visited a lot of manufacturers and the like. We…

Webber: This was in probably ’64? I mean the Fall Joint Computer papers about Multics came out in ’65, so I assume it was a year or two before that.

Corbató: It was… Let’s see. Boy, it’s hard to keep track of dates.

Webber: Yeah. It’s a long time ago. [0:50:00]

Corbató: We went on a series of trips visiting different vendors to see if they were interested in building a machine more explicitly for time sharing. And the… I forget who was actually part of the group that visited the different vendors. Certainly I was. And we visited CDC, we visited… We ended up visiting General Electric. Well…

Webber: What about UNIVAC and Burroughs and all of those? So you went to all of…?

Corbató: We hit all of the big ones, if they were interested in building a computer explicitly for time sharing. It turns out that there was a rather daring and one might have said… GE had a computer division out in…

Webber: Phoenix.

Corbató: …Phoenix, and we proceeded… they had an engineer who basically was much too amenable to change, and he basically promised to do anything we wanted. And…

Webber: Who was that?

Corbató: John Couleur.

Webber: John Couleur. J.F.C.

Corbató: And what we didn’t appreciate at that time was that GE was a very large company, but it basically was run with each division running its own shop and basically answering… Not “answering.” It wasn’t as integrated as a company like IBM, which had a single management. GE had multiple managements. So Couleur basically promised to do anything we wanted, and it sounded pretty

good to our naïve selves. So we elected to work with John Couleur to design what came to be the 645.

Webber: Right. GE already had a 635 that was running GECOS batch processing all the time. So he was going to modify that.

Corbató: He was going to modify it all to hell. And one of the consequences of that was that IBM, who had sort of viewed us as being likely to pick an IBM machine, was secretly building… had basically not unveiled the fact that they had a different notion entirely of how to build computers, namely building a family of computers, all of which had the same architecture, which ran faster or slower depending on how much you paid for. So the result was that when they heard that we had agreed to work with GE, IBM totally panicked and proceeded to… Well, I think that they did some terribly drastic things. They embarked on a time-sharing system of their own which was called Model 67[2] as I recall.

Webber: Yes. And it would run CP/CMS eventually, which was a virtual machine to run time-sharing users. Right.

Corbató: And they did a crash program of trying to assign a thousand programmers to work on it at once. They came up terribly lame.

Webber: Right. Brooks' book, *Mythical Man-Month*…

Corbató: What's that?

Webber: Brooks' book, *The Mythical Man-Month*, explained how IBM kept putting more and more people on it, which was slowing it down more and more, because you don't throw people at it, you throw intelligence.

Corbató: So that didn't work at all. But one of the results was that one of the programmers who had been with us initially had gone to work for IBM, Bob Creasy, and he basically saved the bacon for them by implementing a… as I recall, a simpleminded time-sharing system actually called VMS I think for them. And basically IBM sold people these… People had bought these computers basically on faith that IBM would provide them with an operating system, when in fact they had no such thing.

Webber: Right, I don't think VMS ever really made the light of day. The only thing that saved them at all was the CP/CMS thing that ran on the 67.

So IBM was upset. So John Couleur and GE is now building you the 645. When did that arrive?

Corbató: When did that what?

---

[2] IBM 360

Webber: Well, before that, in the fall of '65, you and several other wrote these great papers in the Fall Joint Computer Conference about something called Multics, your new, next-generation time-sharing system. What does "Multics" stand for? "Multiplexed Information Computing Service." And the key thing about that was it was going to be a service just like the telephone dial tone – it's always there. So Multics was born with those papers I believe. Tell me how those papers came to be and how you had to get a team together to do all that work, because there was a lot of design work before those papers were published.

Corbató: Well, I think it was Fano that insisted we write in advance what we wanted to do. So we wrote these papers, which were a little grandiose and fanciful, as I recall.

Webber: We did most of it.

Corbató: [laughs] And we weren't able to pull off all of the things we thought we would. But I'm losing my thread of thought here.

Webber: So we were talking about the Fall Joint Computer papers. Who were some of the other authors of those papers besides you?

Corbató: Well, as I recall, I think Fano and Ed David wrote one I think.

Webber: Was Saltzer involved in those, Jerry Saltzer?

Corbató: Saltzer? Umm…

Webber: He was involved about that time I know. I know Stan Dunten wrote one on I/O.

Corbató: What's that again?

Webber: Stan Dunten wrote one on the I/O system that Multics might have, with maybe Vyssotsky or something.

Corbató: I've lost track.

Webber: Yeah. Anyway, so those papers kind of laid the groundwork for what became Multics. What happened next? You start hiring people to do this big work? I know the hardware wasn't there yet.

Corbató: You know, I guess we started hiring people. I've lost track of exactly the chronology of everything occurring. But we clearly developed a staff of… [1:00:00] all on the 5th floor there.

Webber: I remember I was hired in about seven months later, and there were people already there by the time I got there. How big did the group get on the 5th floor that was doing this work?

Corbató: Well, we had two or three people per office, and basically we had one-half of the 5th floor. So I've lost track of the headcount.

Webber: Right. It was probably in the order of 30 or so. I know that the IBM effort had a thousand or something. They had way too many people trying to do it, and that was part of their problem.

Corbató: IBM went at it all wrong. And, as I mentioned earlier, Bob Creasy saved their bacon by implementing VM, which was a very simpleminded swapping system, which… Because IBM had already sold the hardware to people and customers all over the place had this equipment sitting idle, not doing what it was supposed to do.

Webber: Right. The development for Multics started on CTSS. For a year or two, they were using that as the development platform?

Corbató: Well, we started on CTSS, but… and we continued that way for quite a long time, as I recall. Because Multics initially, when we first tried to run the system, it was a lame duck. Nobody had proceeded to time out what was going to happen when, and we realized that we had a monster on our hands. As I recall, we did a very drastic slashing of our goals, things like the hardware supported a possible 64-bit page size and a 1024. And we just scrapped the… We basically had a change-making arrangement. We scrapped that as overdesigned by a lot. Things like that we had to sort of cut out with a meat ax.

But we basically got the system up and running. And I forget exactly the chronology of how people…

Webber: Yeah. As I recall, I came back in the spring of '68 and it had been booted by then, but it was slow because it was so complex, as you mentioned. But within the next six months, it was running and people were using it, but not for development yet. It was still CTSS and we would make tapes on the CTSS machine and bring them over to boot on the 645 at that time. Did GE provide all the features you had asked for?

Corbató: Well, they were only thinking of hardware, and yeah, that was there. Couleur delivered on the hardware, but he had no idea what the software was doing. He just had sort of grandiose visions and he was way too… in some sense, he was way too compliant. So there was a lot of naïveté started with Couleur and also with IBM when they went racing off to try to counterpunch.

Webber: What were some of the features of Multics that were really unique and that you were really proud of and whatnot?

Corbató: Well, I think we had… I've lost track of the details at this point, but we had the notion of access control where not just anybody could reference a file or not all memory was equivalent. Basically, you only had the right to that which you were authorized to get at. That sort of idea. We kept the crucial ingredients, but we abandoned some of the fine structure of small pages and things like that.

Webber: Right. I mean the Fall Joint papers mentioned dynamic linking, which now everyone accepts; hierarchical filesystems, which now is in every system out there. These came out first in Unix. Virtual memory, which now everyone has, although it wasn't first on Multics. It was possibly the first commercial system. I know Atlas had it back years before. But all the paging and segmentation were brand new in Multics. There were many things that had to come together for that.

Who were some of the key people that were developing it? I know Bob Daley was there, Jerry Saltzer was a good influence. Umm… It was only 60-70 years ago. [laughs] Bob Graham.

Corbató: Well, you were among them.

Webber: I was there, yeah. Bob Graham.

Corbató: Bob Graham.

Webber: He did the dynamic linking. I think he might have written the paper in the Fall. There were so many new ideas that it was very impressive to me. I mean we were creating computer science back then because it hadn't yet been done, a lot of these things. And you were leading that effort.

I remember you would go down to Washington every year, get another round of financing to keep it going another year. Can you tell me a little bit about all those efforts?

Corbató: Well, it was kind of scary. A few key people… Licklider eventually ended up down in Washington as a program manager who had…

Webber: For DARPA?

Corbató: Of AR-… Well, I guess…

Webber: The Defense Advanced Research Projects Agency. Right.

Corbató: Yeah. It was… Actually, and Licklider, who had seen time sharing on a PDP-4, which was a very modest experiment done up at BBN, somehow he was really very naïve about a larger system.

But one key person who played a pivotal role in maintaining our viability was Ed Fredkin. He was a close friend of Licklider and he had his ear. At one point, Licklider was about to shut down the project. It was not reaching its goals fast enough, or maybe never. And Fredkin lobbied hard for us and basically saved the day. So Fredkin was sometimes a very wild-eyed guy, but he had a remarkable good head on his shoulder when it came to doing things right.

Webber: Another name that just came to my mind was Ted Glaser, who was working on it [1:10:00] early on too, I think in the hardware design. Was he working on that?

Corbató: Ted Glaser was… he had been… he was a remarkable person to begin with. He had been blind since he was about six, yet he had played a key role in the Burroughs 5000 I think. And he came and joined the project because he saw it as a great chance to make some important steps, as I remember.

Webber: I remember he had a seeing eye dog, but he would lead the dog around, the dog didn't lead him around. He was an amazing person.

Another thing that occurred to me was one of the key things about Multics was that it was the first large system written in a high-level language, PL/I.

Corbató: Yeah, that was a rather grandiose decision, [laughs] which fortunately… I forget who did it, but we had to strip it down to the core elements, so it was basically a much more simpleminded language. The real true PL/I was a nightmare of features.

Webber: Well, eventually Multics had the full PL/I, by the end of the first…

Corbató: Well, yeah, Bob Freiburghouse was the one that put that through. Yeah. It was amazing because it was too grandiose a language to implement it in. I don't know.

Webber: One of the other things about the Multics project, it was actually three different organizations. There was GE, Bell Labs, and MIT. Bell Labs eventually pulled out and a couple of the people, Dennis Ritchie and Vyssotsky and a couple of others, did Unix, which is a play on words for "Multics."

Corbató: Yes. Well, Ken Thompson was… I don't know that we can claim any credit for it. He was a very, very savvy programmer to begin with, and he basically saw Unix as a chance to simplify life immensely. So he did.

Webber: Right. The one thing that Unix borrowed from… well, it borrowed a lot of stuff from Multics, but it was also written in a high-level language. In C. It wasn't that high-level, but it was a high-level language, no longer assembly, and that made it portable to a lot of platforms, because high-level language, all you need is a compiler. So they learned that from Multics as well.

What other things can you think of that spun off of Multics that the whole industry took hold of? I've mentioned a few, but you might have some others.

Corbató: Well, the… I think one of the notions we had which I think has spun off more was we tried to insist that people described what they were going to do before they did it, which forced you to become more honest about what you really plan to do. Because in the early days of programming, people just started writing code and just seeing what came out, which… I think the notions of access control — we originally had a multi-ring fantasy that there would be all kinds of layers of an operating system. I think we simplified that to the inner core and the outside. [laughs]

Webber: Right. But then eventually that made it into the hardware in the next version of the 645 computer. By then, GE had been bought by Honeywell, so Honeywell kind of took over the hardware line, but it was still John Couleur down in Phoenix building the hardware and whatnot. I don't know if you recall, but he tried to design another generation beyond the 6180 and did not understand software, as you mentioned before, and his scheme was not going to work, but…

Looking back on your long career, obviously you did a lot of teaching – professor at MIT. What are some of the things that you remember most dearly about your career there? I mean you met a lot of important people.

Corbató: Well, initially I guess I… you might say more of my teaching was in leading the group rather than in more formal classroom teaching. I have done some recitation instruction, instructing, but I didn't get too deeply in the groove of that. My primary role was leading a group of research programmers, which I would argue that was just done by the seat of the pants. I didn't have a great plan. I didn't go to leadership school. [laughs]

Webber: At one time though, were you head of the Computer Center?

Corbató: I was… Fano kept trying to avoid a split in the department, and so he argued that there should be an associate head for computer science and a counterpart, an associate head for electrical engineering, both reporting, both working with the department head. So I did a couple of stints as associate head for computer science, once with Paul Penfield and once with… He went out to Oregon. I've lost track of his name.

Webber: Is there anything that you would like to mention for the many people that are going to listen to this as a history of computer science as well as your life, that they should know fifty, a hundred years from now? What would be interesting? I mean you've seen so much change in your life, as we talked about earlier.

Corbató: Well, one of the things I don't think we foresaw so well at the time, and yet it was totally predictable, was computers were… or the microprocessors were improving by basically a factor of two every two years or something.

Webber: Yeah, right. Moore's law.

Corbató: We used to have some… But one of the things that inevitably happened was that eventually a CPU could be put on a chip. Once that happened, we had basically a profound change on the computing industry in the sense that all kinds of gadgets began to develop – iPhones – and once you could have that much computing power within a very modest physical component, you had new possibilities developing all over the map. I don't know that people totally foresaw the evolution of computing. [1:20:00] There used to be a factor of two every two years, I think was kind of the rule of thumb in almost everything. Obviously, that had to stop somewhere or somehow. And we basically are seeing today a proliferation of devices and functions and specialized equipment, which people didn't foresee at the time. How could you?

Webber: I remember one time I was talking to Jerry Saltzer and, just to let people know, the Multics computers were the size of a dozen very large refrigerators. You needed a big air-conditioned room, a big room to hold it all. And at one time, Jerry Saltzer told me – this was probably in 1970 – that you could get all of that computing power into a briefcase. And I knew he was wrong, but boy was he correct. Now it's on your phone, because these phones we have are more powerful than those Multics computers were. It is really amazing. It's been an awful lot of change since we've been in the business. It's an impressive evolution. The speed and size changes you refer to, commonly called "Moore's law," who predicted it so many years ago, it's only now slowing down. And they're talking about quantum computers, which are going to extend it even more. And it's just amazing, the fact that not many people could see the eventual place where these computers would go.

Corbató: Yeah. I have trouble understanding exactly what a quantum computer does. [laughs]

Webber: Right. Well, I'm not sure very many people in the world do, so yeah, we're not alone there.

Corbató: Yeah, yeah. I mean I'm not totally skeptical, but I'm totally mystified by the nomenclature that's used to describe them and things like that.

Webber: Well, being a physicist, the whole concept of quantum mechanics is hard to believe, the conclusions you can draw from it, so it's amazing.

Any last words you'd like to impart to the history of the world?

Corbató: Well, I don't know if I have anything profound to say, except that I think that … up to now, the evolution of computing has been sort of a horizontal one into more and more specialized applications around specialized fields. I think… I don't quite know where that ends. Everything seems to be computerized.

One thing that does sort of give one pause is how ubiquitous computing will be. Will it be in every device we touch? We certainly won't think of programming them so much as designing them to do something. The… But there was some comedian who once said, "The future lies ahead." [chuckles]

Webber: Very good. Good words to end on. Thank you very much, Corby.

Corbató: Okay. Thank you.

[end of recording]