Fernando Corbató

Editor: David Walden

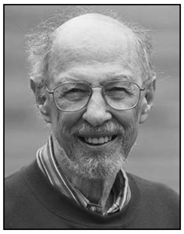

Photo by Jason Dorfman.

Fernando Corbató is a pioneer in the development of time-sharing operating systems.[1] In 1961, he initiated and developed the initial version of the Compatible Time-Sharing System at the Massachusetts Institute of Technology. The CTSS experience influenced the creation of Project MAC, where Corbató and his MIT colleagues developed the Multics (Multiplexed Information and Computing Service) time-sharing system in collaboration with General Electric and initially Bell Telephone Laboratories.

***David Walden*: Please tell me a bit about your early years.**

Fernando Corbató: My parents met in graduate school at the University of California, Berkeley, where they were both graduate students in the Department of Spanish. My father was born in Spain, and my mother, with a Danish-American background, was born in San Francisco. I was born in nearby Oakland after my mother had gotten her masters' degree and while my father was finishing up his doctorate.

After my father got his PhD, he took a position on the faculty at the University of California, Los Angeles (UCLA), which was then just moving from its initial downtown campus to its present campus in Westwood. We lived in Westwood, and I went to a pretty standard set of good public schools.

I entered University High School (UniHi) in the fall of 1941, and in December 1941, the US was catapulted into World War II by the bombing of Pearl Harbor. As a result, UniHi went into extended sessions so some of the students could enter the labor force and work part time while still attending the necessary subjects to graduate. I instead took it as an opportunity to take extra subjects, especially math and science courses, and graduated early in September 1943 so that I could enter UCLA in the fall of 1943.

In the spring of 1944, after about seven months of college, I met a US Navy recruiter who briefed me on an opportunity to join the Navy and entering the year-long Eddy Program (named after the Commander who proposed it). The Navy was faced with a deluge of electronic equipment aboard the ships in its fleet (such as radar, sonar, and fire-control computers) and had initiated a year-long program to train enlisted men to maintain the novel hi-tech electronic gear in the fleet.

I faced the prospect of being drafted as soon as I turned 18 in the summer of 1944, so the chance to join the Navy and enter the Eddy Program looked like a real opportunity. I enlisted in May 1944 and quickly was off to boot camp.

***Walden*: Tell me about your Navy service.**

Corbató: My recollection is that it was about four weeks of boot camp at Great Lakes Naval Station near Chicago. Then there was a month of "pre-radio" in Chicago where we learned elementary stuff to get everyone up to speed. Next we went to Monterey, California, for three months to the site of what is now the Naval Postgraduate School for further basic preparation. And finally, we spent six months at Treasure Island between San Francisco and Oakland, where we had an intensive course focusing on the various pieces of electronic equipment, mostly radar, but a little LORAN and a few other things that were deployed in the fleet. The theory was modest and mostly heuristics, not explanations.

After finishing my year of training, I went on to my only assignment in the Navy, which was to be part of the precommissioning crew of the USS Shenandoah, a destroyer tender. We spent three months in Tacoma, Washington, mostly loading supplies. One week before

Background of Fernando Corbató

Born: 1 July 1926, Oakland, California

Education: California Institute of Technology, BS (physics), 1950; Massachusetts Institute of Technology, PhD (physics), 1956.

Professional Experience: Massachusetts Institute of Technology, 1956 to present.

Honors and Awards: IEEE Computer Society W. Wallace Mc Dowell Award, 1966; American Academy of Arts and Sciences Fellow, 1975; IEEE Fellow, 1975; National Academy of Engineering member, 1976; American Federation of Information Processing Societies Harry Goode Memorial Award, 1980; American Association for the Advancement of Science Fellow, 1982; IEEE Computer Society Computer Pioneer Award, 1982; ACM Alan Turing Award, 1990; ACM Fellow, 1994; NEC Computers and Communications Award, 1998; California Institute of Technology Distinguished Alumni Award, 2003; Computer History Museum Fellow, 2012.

the ship was to be commissioned, the first atomic bombs were dropped and the complexion of the war changed immediately. Instead of being a support vessel in the Pacific and a target for kamikaze planes, we got the ship seaworthy and eventually it went through the Panama Canal. After going up and down the east coast a few times, I ended up being released in Portland, Maine.

That was the extent of my naval experience. I was very fortunate the timing of everything, and I learned a lot about contemporary electronic equipment, which served me in good stead in later years.

***Walden*: How did you end up at the California Institute of Technology after the Navy?**

Corbató: A boyhood friend had gone to Cal Tech as a part of the Navy V-12 program and had been stationed there. He liked the place and recommended it to me. Also, the GI Bill was then in effect, and it covered the tuition. Cal Tech in those days had, I think, four three-hour entrance exams. That was a pain, but I did get in and spent four very enjoyable years there.

***Walden*: How did you end up at MIT for your graduate work?**

Corbató: I had majored in physics at Cal Tech, and I applied to several schools. MIT accepted me and furthermore offered me an assistantship. I didn't know much about the place, but I loaded up everything in my 1936 Ford car, drove across country to Cambridge, and started my graduate student career at MIT.

***Walden*: Somewhere along the line you switched from physics to electrical engineering.**

Corbató: That was after my PhD. In my graduate work, I got involved with molecular physics under John Slater's direction and that involved a lot of computation. In the process of doing that, I ended up using the Whirlwind computer, one of the first fast, real-time computers.[2] It was a real eye-opener. I really enjoyed the notion of being able to create a program where, although you put decision points in your program, you didn't quite know how it was going to really operate until you ran it. That was where I learned a lot about the basics of computing.

Phil Morse was another professor in the physics department and, to some extent, my mentor. He had persuaded IBM to install an IBM 704 computer at MIT, where MIT would manage it and IBM would maintain it. MIT would be entitled to use one shift. In turn, MIT would administer the computer for the benefit of roughly 40-odd New England colleges, which would use another shift of time. IBM would get to keep a third shift of time for itself. IBM was generous, and it was a good deal. Obviously, they were trying to encourage people to get used to using their equipment.

Phil Morse was in the process of staffing the center, and he asked me after I received my PhD to become a research associate in charge of an assistantship program he had arranged for from the Office of Naval Research—about a dozen research assistantships from various departments around MIT. The research assistants would be partially committed to working at the Computation Center. Initially, I was a research associate, then an assistant director for programming research, and later an associate director.

***Walden*: So, you didn't really start in the Electrical Engineering Department, but rather in the Computation Center as part of the research staff.**

Corbató: That's correct, and it's there that I became interested in time-sharing.

John McCarthy, who had become a faculty member in the Electrical Engineering Department and was also a member of the center's research staff, first proposed the idea of time-sharing and was an eloquent advocate for it.[3]

In those days, everybody was using those big computers in batch-processing mode—because computer time was considered so precious you didn't waste it. Jobs were batched onto a magnetic tape, which would be run through the main computer and generate an output tape—mostly core dumps. The minimum time to get a job run was typically two to four hours; usually it was even worse—once a day. Trying to write a big program perfectly was nearly impossible. Time-sharing was viewed as a potential remedy for debugging.

McCarthy proposed to Morse that we ask IBM to make a few key changes to the IBM computer (by that time it had been upgraded to an IBM 709) to allow the possibility of doing time-sharing on the machine. Herb Teager, another faculty member in the EE Department, decided he was going to do that. Unfortunately, Herb's plan for doing time-sharing was to do everything from scratch with a new operating system and new mechanisms. He was well intentioned, but hopelessly ambitious, in my view.

So as kind of an interim step, I came up with the notion that maybe we could do an interim time-sharing system as a demonstration that the ideas were worthwhile. It would run on the equipment that Teager had been trying to organize, and I later dubbed it the Compatible Time-Sharing System, but it came to be known as CTSS.[4] It was a sufficiently successful venture and caught enough attention that Bob Fano encouraged me to join the EE Department. So I went and talked to Peter Elias, who was then the department head, had my interview, and soon became a faculty member in the EE Department.

***Walden*: Was there any explicit computer science component when you joined the EE Department?**

Corbató: The EE Department had a history, going back to the 1950s, of being a broad department with varied interests. There were factions of the department who were into biology, factions in material sciences, and factions in computing, in the sense that these were people who were teaching programming, who went into electronic circuits, and so forth. However, it took a couple decades before the department formally relabeled itself the Electrical Engineering and Computer Science Department. Besides teaching various programming subjects for the department, my role was leading the CTSS team; that became a key part of starting Project MAC.[5]

***Walden*: Please say something about learning to be a computer science person yourself and teaching others.**

Corbató: With my physics background, I had been using Whirlwind intensively—I learned programming by doing. Also, one of the tricks I picked up along the way for learning about large systems was to make some modest system changes. This didn't require the burden of doing the whole system over but required understanding the system in order to change it intelligently.

For example, in an "Introduction to Programming" course we taught, I developed an assembly language program called the Classroom Assembly Program for generating machine language code from symbolic instructions.[6] The students were given a book including the listings of this entire program. (I remember writing the program over a long weekend—it was only about 1,000 lines.) As an exercise, the students were asked to add features to this system, and by forcing the students to figure out how to add something, they had to understand most of it. It was a pretty successful technique.

***Walden*: CTSS,[5] Project MAC,[6] and Multics[7] are all well known and extensively documented, but tell me a bit about the progression from one to the next.**

Corbató: One of the things that helped launch Project MAC was that we had CTSS as a working prototype time-sharing system running on standard equipment with a few hardware modifications. The plan was that using CTSS as a working system and launching place, we would then move on to create a more suitable time-sharing system. That next stage came to be known as Multics.

CTSS was only ever replicated on two machines—one in the Computation Center and one at Project MAC—but it was a good demonstration, and we were able to use all the programs that had been written for the 704 and 709 with minor modifications. Users could write real programs. We had a Fortran compiler already available, and we had Bernard Galler, Bruce Arden, and Robert Graham's MAD [Michigan Algorithm Decoder] compiler. It was a useful computer to get work done. We had enough online storage so that people could keep their programs online. They didn't have to mess with pesky punched cards anymore. So it was a tremendous introduction to using an interactive computer.

When Project MAC was formed, under Bob Fano's leadership, there was a summer study to get it off to a running start. We invited about 100 or so prominent computer scientists from around the country who we thought were good candidates to come and learn about time-sharing, using CTSS as the vehicle to show it off. People came and spent anywhere from one to three weeks, talking to each other and exploring the ideas of what it was like using it. It was a good experience, and people got to use it directly. I remember Don Knuth came, among many others.

Once we announced our plan to create a new time-sharing system, we went shopping around to look for hardware vendors and ended up picking GE. Bell Labs joined in because they wanted to be part of the action. We had the situation where we had to create

I've had the fun of watching the computing field evolve at a breakneck pace—we've gone from gigantic mainframes to pocket computers and phones.

a time-sharing system to run on this customized hardware with three different parties participating.

We then wrote our plans for Multics and the classic set of six papers that were presented jointly at one of the computer conferences.[8] In some sense, we over reached. We laid out an ambitious plan, in hindsight. What made it really hard was that we hadn't quite thought it through in terms of the fact that we would have no tools to work with. When we built CTSS, we had compilers, assemblers, everything was already there. We had changed the architecture so drastically on the GE 645 that the assembler program was not even working. We had a tremendous effort to try to start from scratch; it was like starting a computer company. Because we weren't wearing managerial hats, we didn't lay out a schedule and say that these are deadlines and we've got to meet them, and if we don't meet them, here's what we'll do instead. We just sort of said this is our goal and we'll try to do it.

As a result, the project took significantly longer than we had hoped, and this put a strain on everybody—from people seeking funds from ARPA to Bell Labs, which had gone into the venture expecting to have a new research computer for use at their labs (they, probably a little naively, thought of it as being product-level quality right off the bat). We had a tough three or four years to get rolling. We used CTSS as a staging area; we programmed everything on CTSS at first until we could finally get Multics to run and then we could use it to work on itself.

***Walden*: Still, quite a few Multics systems were sold, and it had considerable impact on the world of computing in terms of lessons learned and ideas implemented in follow-on systems.[9]**

Corbató: I think our ideas were sound. We didn't factor in pulling it all together in our planning, but we gradually laced it together. There were a few scary times when the project nearly got killed, but we got through those.

***Walden*: During the CTSS/Multics era, there were other time-sharing systems being built—for instance, the SDS930-based system at Berkeley. Did Multics draw anything from those other systems?**

Corbató: We were so busy solving our own problems that I don't think we picked up any particular ideas from anybody else. But we were aware of other efforts, and from our vantage point, they were all compromised efforts where they had given up various aspects of our grand plan. Our grand plan was a little too grand, but eventually we pulled through. I think Multics sales were mostly to the government, at first. Honeywell (GE sold its computer division to Honeywell) charged plenty of money for the system. I think there was a maximum of about 200 sites.

***Walden*: What direction did your research take after Multics?**

Corbató: I never did get back into direct research very much but began to play various roles in departmental administration.

I was asked if I wanted to head Project Athena, which was then forming, but I was too tired at that point. My wife had just died, and I wasn't ready to do that. Fortunately, Jerry Saltzer, who had been one of my doctoral students, did a wonderful job of picking that up and going on.[10]

Bob Fano not only led Project MAC, but when Louis Smullin, who was then the head of the Electrical Engineering Department, decided he needed two associate heads, Bob Fano became the first associate head for computer science. When Fano later went on sabbatical, I became the next associate head for computer science. That was the beginning of my serving the department in various administrative roles including a second stint as associate head for computer science as well as doing my share of teaching at the undergraduate level. Also, as a consequence of my involvement with CTSS and Multics, I was elected to the National

Academy of Engineering. As a member, I was on various NAE committees involving different issues for a period many years.

***Walden*: Do you have any summary words about your era of computing?**
Corbató: I've only recently begun to reflect on it. I was very, very lucky. When you think of modern computing, it only really began around the late 1940s and early 1950s, and I happened to step in just at the right time. I've had the fun of watching it evolve at a breakneck pace—watching the evolution of the computing field itself as we've gone from gigantic mainframes to pocket computers and phones, which are really computers too. It's an amazing evolution that is beginning to impact the sociology of populations in ways that are still unclear to me.

I can remember how we were excited when we first began to do email and file transfers and use computers at a distance. Even though we had the Internet in the early 1970s, it wasn't until 20 years later, the early 1990s, when the World Wide Web sprang to life. The potential was there all the time, but it took that seminal set of contributions to revolutionize everybody's lives, including those of their spouses and children.

I guess the only comparison I have is from talking to my father about his early days. He talked a little bit about how he'd seen his first movie brought to his small town in Spain with a travelling fair when he was just a boy in the 1890s. The movie, which was projected on a sheet and had no sound, was of two trains coming together and having a head-on collision. When the trains collided, someone banged some oil cans together. From these crude beginnings he saw modern movies and television evolve. He also saw the emergence of electric lights, automobiles, radio, telephones, and air travel to mention only the most obvious changes we now take for granted. I've seen a lot of change too in my own life, and I wonder, what will my children and their successors see?

***Walden*: Thank you very much. It's been an enormous pleasure hearing about your life in computing.**

References and Notes

1. This interview was started by email in February 2011. It was finished on 21 July 2011 in Professor Corbató's office at MIT. It was edited down from its original length of 6,970 words to the size of this column with the approval of the interviewee. The editor added the citations. At least two scholarly oral histories of Corbató exist: A.L. Norberg, "Interview with Fernando Corbató," 18 Apr. 1989 and 14 Nov. 1990, Charles Babbage Inst., call no. OH 162, http://purl.umn.edu/107230; S. Webber, "Oral History of Fernando Corbató," 1 Feb. 2006, Computer History Museum, reference no. X3438.2006, http://archive.computerhistory.org/resources/access/text/Oral_History/102658041.05.01.acc.pdf. See also K.A Frenkel, "An Interview with Fernando Jose Corbató," *Comm. ACM,* vol. 34, no. 9, 1991, pp. 83–90.
2. T. Green, *Bright Boys: 1938–1958, Two Decades that Changed Everything,* AK Peters, 2010.
3. *IEEE Annals,* special issue on time-sharing at MIT, vol. 14, no. 1, 1992.
4. F. Corbató, M. Merwin-Daggett, and R. Daley, "An Experimental Time-Sharing System," *Proc. Spring Joint Computer Conf.,* vol. 21, 1962, pp. 335–344, http://larch-www.lcs.mit.edu:8001/~corbato/sjcc62; D. Walden and T. Van Vleck, eds., "The Compatible Time Sharing System (1961–1973): Fiftieth Anniversary Commemorative Overview," IEEE Computer Society, 2011, http://www.multicians.org/thvv/compatible-time-sharing-system.pdf.
5. *IEEE Annals,* special issue on MIT's Project MAC, vol. 14, no. 2, 1992.
6. F.J. Corbató, J.W. Poduska, and J.H. Saltzer, *Advanced Computer Programming: A Case Study of a Classroom Assembly Program,* MIT Press, 1963; http://www.bitsavers.org/pdf/mit/ctss/AdvCompPgmg_May63.pdf.
7. F.J. Corbató and V.A. Vyssotsky, "Introduction and Overview of the Multics System," *Proc. 1965 Fall Joint Computer Conf.,* pp. 185–196. See also T. Van Vleck's extensive website of Multics history: http://www.multicians.org.
8. Multics was described initially in a set of six papers presented at the 1965 Fall Joint Computer Conference: *Proc. AFIPS* 1965 *Fall Joint Computer Conf.,* vol. 27, Spartan Books, pp. 185–248. The papers explained what was intended, which in some cases was later revised in light of experience.
9. F. Corbató, "On Building Systems That Will Fail," ACM Turning Award Lecture, 5 Mar. 1991, http://larch-www.lcs.mit.edu:8001/~corbato/turing91.
10. The Athena project provided computing and communications for a new (circa 1980s) educational environment at MIT: http://web.mit.edu/Saltzer/www/publications/faculty.pdf and http://web.mit.edu/6.933/www/final.pdf.