A Silicon Valley museum-in-the-making showcases a half-century of innovation

The Ghosts of Computers Past

BY PAUL WALLICH
Contributing Editor

The guest book at the nondescript pair of warehouses just across the street from the monumental blimp hangar at NASA's Moffett Federal Airfield in Mountain View, Calif., looks like a *Who's Who* of the computer industry. There's Gene Amdahl, who designed two generations of mainframes at IBM Corp. and another two at his own companies; C. Gordon Bell, who built the minicomputers that powered Digital Equipment Corp.; and Donald Knuth, whose algorithms have set standards in computer science for 40 years.

These and the other pioneers who prowl Moffett come to visit their brainchildren—or to find a home for them. "Anyone who's anyone in the computer industry and is getting on in years" has come to reminisce or to donate, says curator Mike Williams.

Known as the Computer History Museum, the collection here spans a "unique period" in the history of technology, says veteran computer architect and museum board member John Mashey. Fifty years ago, there were essentially no computers; now they are everywhere and becoming just about invisible. Now is the time to preserve the heritage of modern PCs as well, both the artifacts themselves and the detailed information about their development and workings, while so many of the people who created the computer revolution are still available.

Thousands of people around the world maintain collections of old computers—anything from the half-dozen old PCs any given hacker hasn't cleared out of the bedroom to the half-dozen or so Crays in the private museum built in Wisconsin by

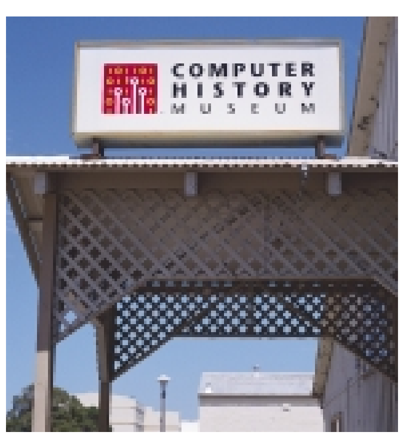

United Parcel Service pilot and computer aficionado James Curry. There are even quirky institutions like the Digibarn, north of Santa Cruz, Calif., where a veteran PC programmer is turning a spur-of-the-moment hobby into a mecca for aging computer buffs to relive past glories [see *IEEE Spectrum*'s visit to the Digibarn at http://www.spectrum.ieee.org].

But the museum at Moffett contains perhaps the most complete collection of groundbreaking hardware and software in the world—from the Hollerith punch-card tabulating machine that rescued the 1890 U.S. census to the LINC laboratory minicomputer to a prototype of the Palm Pilot PDA and an early copy of IBM's gigabyte credit-card–sized disk drive.

Preserving all these artifacts is quite an achievement in a field where last year's top-secret supercomputer is next year's scrap. Just inside the door, Williams points out the JOHNNIAC, which was rescued from the garbage dump not once, but twice. This sole survivor of the original generation of machines designed by computer pioneer John von Neumann was built at Rand Corp. and then donated to a local museum when its useful life ended.

PHOTOS: CLAY MCLACHLAN/GETTY IMAGES

Part of the CDC 6600, the first supercomputer *[see p. 42]*

A few years after that, Williams recounts, the local museum rethought its needs for exhibit space, and the machine's builders came upon their creation in the parking lot, sawed up into sections for pickup by a scrap dealer.

Among the first to recognize the importance of preserving past computers for study was Gwen Bell, former president of the Association for Computing Machinery and the Computer History Museum's founding president. Her first exhibit, of machines that were collected by her husband IEEE Fellow C. Gordon Bell, started out in a glassed-in former coat closet in the mill building at Digital Equipment Corp. (DEC) in Maynard, Mass. The collection later filled a cavernous DEC lobby and eventually formed the nucleus of the Computer Museum in Boston.

In the 1990s, as this museum shifted its focus to public education, much of its collection was transferred to the warehouses at Moffett. When in 1999 it merged with the Museum of Science, Boston, what was left of its historical collection also went west into the Moffett buildings. These artifacts all now make their home at the Computer History Museum.

Visible storage

It isn't really a museum yet. Its board is looking to acquire one of the many office buildings left vacant by the dot-com bust and fill it with museum-quality displays. Meanwhile, the publicly accessible space at Moffett is what curator Williams and his colleagues politely call "visible storage"—

The Hollerith tabulating machine ***[left]*** made the U.S. census of 1890 possible, completing it in less than a year instead of a projected 10 years. It also set the stage for the electromechanical calculating revolution of the first half of the 20th century. The basic punch-card reading mechanism [near right] stayed the same for about 75 years: spring-loaded pins arranged in a grid in the top of the reader could drop down and complete a circuit only in positions where a hole had been punched in the card.

The counters on the back of the machine could then be wired to correspond to any combination of punch-card data. So if you wanted to know, say, the number of Swiss-born farmers in Ohio with holdings of less than 50 acres, income of more than US $100 a year, and at least one child in public school, you could simply wire up the pins and run the cards. Plugboards for "programming" such selections remained in use for punch-card readers through the 1970s.

computers are laid out in long rows with as much concern for what fits where as for chronology or theme. That nearly endless bank of vacuum tubes on the left, for instance, is a tiny chunk of the AN/FSQ-7, otherwise known as SAGE (Semi-Automatic Ground Environment), whose 27 dual-processor installations protected the United States from bomber attack from its full deployment in 1963 through the early 1980s.

At the near end of the same aisle is Gene Amdahl's Wisconsin Integrally Synchronized Computer (WISC), the computer he built to do the calculations for his Ph.D. thesis, and at the far end of the aisle is Seymour Cray's CDC 6600, the high-speed machine that inspired IBM Corp.'s Thomas J. Watson to complain that his research labs had been bested by a company of "34 men, including the janitor."

In the room next to the mainframes, Williams walks through a huge door into the real storage area. Here on the right stretch rows of shelving three or four meters high. On the left is "Area 51" (so named because no one is quite sure what lurks behind its metal-mesh partition), where old mainframes and super minis are jammed so close that volunteers have to climb over layers of hardware to examine half-buried machines. Williams pulls a crate off one of the shelves and finds a mechanical computer for calculating the area under a curve. Nearby is another crate about 60 cm long containing a cathode-ray-tube memory element. Williams shakes his head fondly, recalling that CRT memories were so sensitive that just the static electricity from combing your hair could crash a machine that used them.

The **WISC** (Wisconsin Integrally Synchronized Computer) was Gene Amdahl's first machine, designed in 1950 to perform the calculations for his Ph.D. in physics. Amdahl never worked as a physicist; instead he went to work for IBM Corp. in 1952, designing more computers, and the WISC was completed in 1955. Amdahl designed mainframes for a succession of his own companies into the 1980s.

The WISC [right] acquired the bullet holes in its console [upper left] long after Amdahl left Wisconsin, according to the Computer History Museum's curator, Mike Williams. The machine and the University of Wisconsin computer center manager retired at about the same time, and the manager kept it in his basement where his teenage son used it to hold up targets for pistol shooting. "I was giving a tour and telling that story," says Williams, "when I noticed someone smiling; it was Gene. Afterwards he explained to me, based on the pattern of the misses, all the things that you could learn about the shooter, for example, that he had to be right-handed."

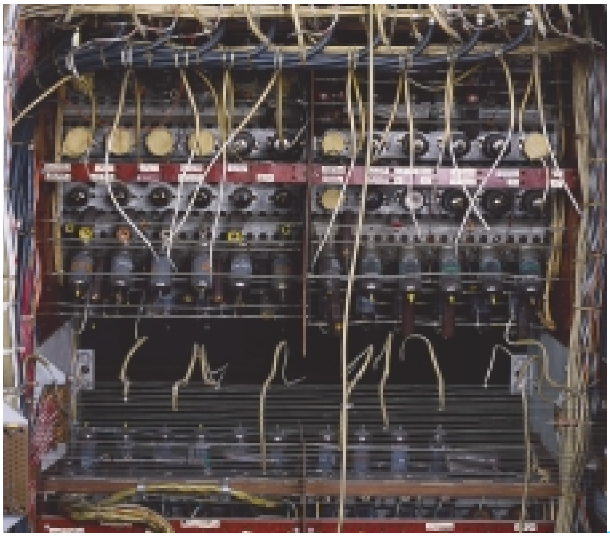

Cut wires show where the **JOHNNIAC** (Rand Corp.'s first computer) was chopped apart for scrap after being discarded from a museum exhibit. Built in 1953 and decommissioned in 1966, the JOHNNIAC is the only survivor of close to a dozen machines built during the early 1950s using blueprints supplied by John von Neumann. Among them were the ILLIAC at the University of Illinois, the WEIZAC in Israel, the CILIAC in Australia, and Nicholas Metropolis's MANIAC at Los Alamos National Laboratory in New Mexico.

Most first-generation computers were designed by companies that kept architectures proprietary, but von Neumann, then at the Institute for Advanced Study at Princeton, would even send technicians to help other organizations. As a result, he is widely credited with the "von Neumann architecture," the standard for modern computers, even though he had nothing to do with the detailed design of the machines. As a mathematician, "he wouldn't have known a resistor from a vacuum tube," notes Computer History Museum curator Mike Williams.

Meanwhile, the collecting, cataloging, and restoration continues. Once a week Williams and his staff go through the accumulated e-mail from computer centers whose managers are hoping to get a white elephant off their hands.

Failed machines, he notes, can be at least as important for understanding the history of computer science as the successes. That's one reason that a KSR-1 (perhaps the only unit of the cacheless multiprocessor design ever sold) and a 16384-processor Connection Machine sit cheek-by-jowl with Xerox's so-called D-machines that helped inspire the PC revolution.

Other noble commercial failures include the ETA-10 supercomputer (intended to be cooled by liquid nitrogen) and the Cray 3, which compressed a 16-processor supercomputer into a volume only slightly larger than a shoebox (albeit the ancillary equipment required to supply power and remove heat generated at an average rate of more than 10 W/cm^3 took up most of a room).

What machines is Williams collecting now? There are still a few traditional supercomputers out in the wild, especially from companies like NEC, which has developed six models of super machines over the past 20 years. But most of the world's fastest machines today are simply enormous collections of commodity microprocessors and memory. For those, a rack or two will do to demonstrate the structure—much as a few dozen connected modules from the ENIAC or the SAGE can stand in effectively for the whole.

The era of rapid innovation in high-performance computers like the Cray 2 and Cray Y-MP may well be over, says Williams, because the military and industrial demand that supported companies like Cray has waned. Innovative applications of relatively ordinary CPUs

AN/FSQ-7, or **SAGE**, first deployed in 1958, was the heart of North American air defense during a long period of the Cold War, with 27 installations, each in its own custom building, throughout the United States and Canada. The racks here represent perhaps 5 percent of a single processor; a full installation would occupy a sizable part of the blimp hangar behind the museum. The pull-out trays were crucial to the system's performance, with each CPU typically running for less than a day before one of its 50 000 vacuum tubes needed replacing. Each installation consisted of two identical computers, one running, the other being fixed.

LINC was the first laboratory minicomputer [above]. Initially built in 1962, it contained analog-to-digital and digital-to-analog converters to record data and control experimental equipment. It also boasted what was for a while a generous 2048 words of memory. Back then, the notion that ordinary scientists could have their own computer to record and analyze data was almost preposterous. Within a few years, though, Digital Equipment Corp.'s LINCs were running animal psychology experiments, recording evoked potentials from neurons, and measuring arterial blood flow. More than a thousand machines based on the original design were eventually built; only this one and a handful of others survive.

are where the action is now, he notes. On one wall in the relatively empty back room, for example, hangs the first working prototype for the Palm Pilot. Across the aisle is the IBM microdrive, capable of holding a gigabyte of information nestled comfortably within the hub of its 10-MB great-great-grandparent.

Going virtual

Impressive though the hardware may be, the old machines tell only the skeleton of a story without software and documentation to flesh them out. Both the actual texts of the machines' workings and the recollections of the people who built and used them are invaluable. The museum is collecting as many oral histories as possible, along with videotapes showing SAGE and other important computers while they were still being used. And a select few machines are being carefully brought back to life, so that they can run the programs that made them famous.

That restoration poses all kinds of problems. No one makes ferrite core memory anymore, for example, so the museum's IBM 1620 uses a tiny DRAM module instead. But many other machines will never be restored to working order, says Williams. The Crays, in particular, says museum board member Mashey, were "handbuilt by elves" and hard enough to keep running even when they were new.

For those machines, simulations will have to do: today a bit-perfect emulator of an old machine can be run by a PC in software much faster and more reliably than the hardware ever worked, and the old operating system and software can be run on top of that.

Indeed, once you've started transferring parts of the museum to cyberspace, why stop? C. Gordon Bell argues that the entire

The **CDC 6600** infuriated IBM Corp. head Thomas J. Watson when it came out in 1964. Why? Because Control Data Corp (CDC), a tiny Minneapolis company, had just made a computer faster than any model produced by his cadre of thousands of R&D engineers. CDC's machine also set the tone for the many other supercomputers Seymour Cray would design by combining simple architecture and aggressive implementation [see also blowup opposite title page].

The **Cray 2** supercomputer [left], first built in 1985, is a comfortable fit for Computer History Museum curator Mike Williams. Seymour Cray's trademark torus helped simplify the machine's design by equalizing the distance electric signals had to travel. Circuit boards are visible through transparent panels in the top of the machine, which was filled with fluorocarbon coolant when in operation.

Modular plumbing is perhaps the most obvious feature of the 1988 **Cray Y-MP** [right], designed by supercomputer architect Steve Chen. Manifolds distributed liquid to cooling channels embedded in each circuit board—blue is the cold side, red is hot. A quick-disconnect fitting on each tubing segment let workers swap out boards without draining and refilling the entire machine, a distinct improvement on the Cray 2.

kit and caboodle should be "cyberized" so that hundreds of millions of people could visit every year instead of mere thousands. His vision, based on his ongoing research work in telepresence, would also help dissolve the bottleneck most museums face in being able to display only a tiny fraction of their collection at any given time. In 50 years, he says, almost none of the pioneering machines that now take up center stage in visible storage will still be out on the museum's exhibit floor. But in cyberspace, each of them can take up as much room as it needs.

By then, the massively parallel processors, mixed analog and digital processors, and "smart dust" that researchers are now touting as the present and future of the computer revolution will be as much a part of ancient history as the vacuum-tube behemoths of primordial computing are now. And for those who need to study the past to know how to build the future, all the different generations will still be accessible. ●

Tekla S. Perry, *Editor*

The wall of PCs includes groundbreaking machines like the **Apple Macintosh** [top shelf, left of vertical support], **Osborne Compaq** [second shelf, below Mac], and the **Xerox Alto** [floor, below Osborne], along with now obscure PCs like the **Apple Lisa** [left of the Mac].