

# The quest to save Stephen Hawking's voice

How a Silicon Valley team helped rebuild his distinctive robotic sound

---

By Jason Fagone

Eric Dorsey, a 62-year-old engineer in Palo Alto, was watching TV Tuesday night when he started getting texts that Stephen Hawking had died. He turned on the news and saw clips of the famed physicist speaking in his iconic android voice — the voice that Dorsey had spent so much time as a young man helping to create, and then, much later, to save from destruction.

Dorsey and Hawking had first met nearly 30 years earlier to the day. In March 1988, Hawking was visiting UC Berkeley during a three-week lecture tour.

At 46, Hawking was already famous for his discoveries about quantum physics and black holes, but not as famous as he was about to be. His best-seller, “A Brief History of Time,” was a week away from release, and Californians were curious about this British professor from the University of Cambridge, packing the seats of his public talks, approaching him at meals. Hawking motored into buildings and onto stages in a wheelchair with a seat of maroon sheepskin, zooming around with the nudge of a joystick, grinning as he left journalists and his nurses in the dust.



Photo: PAUL E. ALERS, NYT

---

## IMAGE 1 OF 8

Stephen Hawking gives a talk during NASA's 50th anniversary celebrations at George Washington University in April 2008.

---

When he spoke, it was in the voice of a robot, a voice that emerged from a gray box fixed to the back of his chair. The voice synthesizer, a commercial product known as the CallText 5010, was a novelty then, not yet a part of his identity; he'd begun using it just three years before, after the motor neuron disease amyotrophic lateral sclerosis stole his ability to speak. Hawking selected bits of text on a video screen by moving his cheek, and the CallText turned the text into speech. At the start of one lecture, Hawking joked about it: “The only problem,” he said, to big laughs, “is that it gives me an

American accent.”

Dorsey was with Hawking for part of that trip, tagging along as a sort of authority on the voice, explaining its workings to journalists. He worked at the Mountain View company that manufactured the CallText 5010, a hardware board with two computer chips running custom software.

An upbeat 32-year-old, Dorsey was quiet by nature, but driven. He had joined Speech Plus as an intern, attracted by its mission to help the voiceless and the disabled; now he led a team of engineers, and at least 20,000 lines of his own code were in the CallText, the product that gave voice to the most celebrated scientist of his era.

“We are getting close to answering the age-old questions,” Hawking said at the close of a lecture. “Why are we here? Where did we come from? Thank you for listening to me.”

At the end of his California tour, the physicist gave Dorsey a signed copy of his new book, his thumbprint pressed onto the inside cover.

Hawking returned to Cambridge, Dorsey to his life in California.

Twenty-six years went by before they would cross paths again.

In tech years, that is a millennium. The Internet happened. Silicon Valley boomed, busted, boomed again. Apple, Amazon, Facebook, Google, Uber.

Dorsey, meanwhile, left Speech Plus, which went bankrupt and was sold to a series of other companies. He got married and had kids. He joined a Buddhist temple. He eventually left the field of speech technology altogether, becoming an engineering leader at DVR maker TiVo.

Tech, he'd learned, moves so fast. “There's a new iPhone every year,” Dorsey says. “Everything just kind of gets buried in the dustbin of history very, very quickly.”

That's why, when an email from Cambridge University arrived out of the blue in 2014, Dorsey was surprised. It came from Hawking's technical assistant, Jonathan Wood, who was responsible for Hawking's communications systems.

Wood explained something so improbable that Dorsey had trouble understanding at first: Hawking was still using the CallText 5010 speech synthesizer, a version last upgraded in 1986. In nearly 30 years, he had never switched to newer technology. Hawking liked the voice just the way it was, and had stubbornly refused other options. But now the hardware was showing wear and tear. If it failed entirely, his distinctive voice would be lost to the ages.

The solution, Wood believed, was to replicate the decaying hardware in new software, to somehow transplant a 30-year-old voice synthesizer into a modern laptop — without changing the sound of the voice. For years, he and several colleagues in Cambridge had been exploring different approaches. What did Dorsey think?

*Thirty years old? He thought. Oh, my God.*

It wouldn't be easy. They might have to locate the old source code. They might have to find the original chips and the manuals for those chips. They couldn't buy them anymore, the companies don't exist. Solving the problem might mean mounting an archaeological dig through an antiquated era of technology.

But it was for Stephen Hawking.

"Let's get it done," Dorsey said.

---

The poet Longfellow once wrote that the human voice is "the organ of the soul." More than any other part of us, our voice expresses who we are, and the smallest fluctuations swing meaning in ways that are hard for computers to understand. You speak a sentence, and the intonation rises or falls depending on whether you're making a statement or asking a question. You do it without thinking, but a computer has to make a guess.

Today's synthesized voices, like Apple's Siri, rely on prerecorded libraries of natural sound. Voice actors record huge libraries of words and syllables, and software chops them up and reassembles them into sentences on the fly. But 30 years ago, computers could only produce a "stick-figure version" of a human voice, says Patti Price, a speech recognition specialist and linguist in Palo Alto.

Back then, she worked as a postdoc in the Massachusetts Institute of Technology lab of Dennis Klatt, a tall, thin, opera-loving scientist originally from Wisconsin. Klatt is the godfather of Hawking's voice. He blasted his own throat with X-rays to measure the shape of his voice box as he articulated certain sounds, and then developed a software model of speech, the Klatt Model, based on his own voice.

Speech Plus took Klatt's model, improved on it, and commercialized it in various products, including the CallText 5010. One of Dorsey's contributions was to write an algorithm that controlled the intonation of the voice, the rise and fall of words and sentences. Speech Plus would sell thousands of CallText systems, though many customers complained that the voice sounded too robotic.

But Hawking liked it.

True, it was robotic, but he appreciated that it was easy to understand: “noise-robust,” as Price explains. The shape of its waveform was more like a series of plateaus than the steep mountain cliffs of human voices, which fall off more sharply. The flattish slope of Hawking’s voice made it cut through noise in amphitheaters and lecture halls. He often began his speeches with the same line — “Can you hear me?” — and the audience would respond with an enthusiastic “Yes!”

“It’s got a ring to it that sticks out,” Price says.

“It’s very intelligible,” Dorsey says. “You can listen to it for a long time, and it’s not irritating.”

Hawking’s only complaint was that it didn’t have a British accent.

Over the years, as synthetic voices grew more natural, taking advantage of faster chips and cheap storage, Hawking had chances to upgrade. In 1996, a Massachusetts speech technology company called Nuance, which had acquired the remains of Speech Plus, upgraded the CallText with evolved software code that made the voice sound fuller and faster, less robotic, with shorter pauses between sentences — to the engineers, an obvious improvement.

They sent Hawking a sample of the new voice, thinking he’d be pleased. He was not. He said the intonation wasn’t right. He preferred the 1986 voice, the one modulated by Dorsey’s intonation algorithm. Hawking would stick with that one.

“I keep it because I have not heard a voice I like better,” he once said, “and because I have identified with it.” He could change to a smoother voice, but then he wouldn’t sound like himself.

“To Stephen, his equipment is like a part of his body,” said Wood, his chief technical aide. “To upgrade him to a new piece of software or a new piece of hardware ... he’s having to change a physical part of himself.”

---

Starting around 2009, Wood and several others at Cambridge began trying to separate Hawking's voice from the failing CallText hardware. The group would include Peter Benie, a computer guru at the university; Pawel Wozniak, a local engineering student; and Mark Green, an experienced electrical engineer with Intel, which had a long relationship with Hawking.

One option they considered was tweaking a modern synthetic voice like Siri to sound more like Hawking. But Siri-type systems rely on the vast computer power of Internet clouds, and Hawking couldn't be constantly tethered to the Internet. Benie also tried a completely different approach. He wrote a software emulator for the CallText — essentially a program that would fool a modern PC into thinking it was actually the old CallText. But the samples it produced didn't sound faithful enough for Hawking's taste.

By the time Cambridge reached out to Dorsey in 2014, they were investigating a third avenue: track down the old CallText source code, now owned by Nuance, and port it to Hawking's laptop, transplanting the old voice into a fresh new body.

Was it possible? Dorsey had no idea. It depended on whether he could find the source code, or, failing that, information that would let him reverse-engineer the source code.

He started emailing colleagues he hadn't seen in 30 years, asking if they had any CallText bric-a-brac still laying around: boards, chips, manuals. One guy found an actual CallText board in his garage. Others located dusty schematics.

It had the feel of a mad scramble through an earlier era of technology. But people everywhere leaped at the chance to help. "The goal is to save his voice," Dorsey said. "Once you go to somebody — 'I need you to help save Stephen Hawking's voice' — they immediately wake up."

His closest collaborator in Palo Alto soon became Price, the speech technologist who had once studied with Klatt, the godfather of Hawking's voice. She was a master at analyzing audio samples, comparing one to another and using their audio fingerprints to reverse engineer how they must have been created.

Dorsey's archaeological quest for old code turned out to be a frustrating one. No one at Nuance was able to find the source code from the 1986 version of CallText. They did, however, find the code for the upgraded 1996 version of the voice, on a backup tape in an office in Belgium. After a few months of work, Nuance engineers got the code up and running and sent a series of audio samples to Hawking's team, adjusting the program to try to match the 1986 voice.

It didn't quite work. For one thing, the match was close but not perfect. Hawking flagged subtle differences others had trouble discerning. "It's like recognizing your mother's voice," Price said. "When you hear them over the phone, they say two syllables and you know if that's right or not."

The other issue was that Nuance owned the code, not Hawking. The famed physicist had always been intent on controlling the use of his own voice. If the team avoided using proprietary software, Hawking was likely to have more control.

---

At this point, they switched tacks and returned to one of their original ideas: to emulate the CallText in software, similar to how PCs can emulate old Nintendo games that aren't sold anymore.

The CallText, of course, was a more intricate beast than a Nintendo, driven by two obsolete and complexly interacting chips, one made by Intel and the other by NEC. Building the emulator demanded heroic feats of programming, intuition and high-tech surgery. The chips had to be removed from a spare CallText board with tweezers and a screwdriver. An emulator for the Intel chip had to be written from scratch, by Benie. A separate emulator, for the NEC, was borrowed from an open-source Nintendo emulator called Higan.

Then all these disparate pieces had to be glued together. It was a little like doing a jigsaw puzzle in a dark room. One chip was passing a mysterious packet to the other every 10 milliseconds. Why? What was in it?

For a while, it was tough going. Some of the audio samples were so poor that no one dared play them for Hawking.

The breakthrough came just before Christmas 2017, when the emulator finally started producing sounds that resembled the familiar voice they had been chasing. It had some minor glitches, but according to Price, the voice was an acoustical match to Hawking's, the waveforms virtually identical. The only perceptible difference was a lack of analog buzz. "It's like a clean and shiny scrubbed-up version of his voice," Price says.

When Benie heard it for the first time, coming out of a computer instead of Hawking's voice box, he thought it sounded more American than Hawking's voice. It was just an aural illusion. Benie realized that perhaps, whenever he saw Hawking speak, he had been mentally adding a hint of Britishness.

Audio of the test text run through the emulator. Media: San Francisco Chronicle

---

Over the next weeks, in Cambridge and Palo Alto, the team members continued to debug the new voice, feeding it snippets of old Hawking speeches and sample texts full of random commas, listening to the results.

On Jan. 17, the team felt ready to demonstrate the new voice for Hawking. Wood, Wozniak, and Benie went to Hawking's home in Cambridge and played him samples on a Linux laptop. To the team's relief and happiness, Hawking gave his blessing. It did sound like his voice.

They still needed to port the voice to the PC, so temporarily, Wood loaded a version of the voice onto a miniature hardware board known as a Raspberry Pi. He thought Hawking might want to evaluate the voice in everyday life, and the Pi was the quickest way to get him up and running.

On Jan. 26, Wood took the Pi along to Hawking's house and asked if he'd like to try it out. Hawking raised his eyebrows, which meant "yes."

The team put the Pi in a tiny black box, attached it to Hawking's chair with Velcro, and plugged it into the voice box. Then they disconnected the CallText. For the first time in 33 years, Hawking was able to speak without it.

Final version of Stephen Hawking's voice emulator from 2018 Media: San Francisco Chronicle

---

Wood watched eagerly for Hawking's reaction.

“I love it,” Hawking said.

For the next few weeks in private conversations, Hawking continued to speak through the emulator and the Raspberry Pi, chatting happily with friends and colleagues. Wood said, “It was a pleasure to be able to give him something like that, that so many people have worked on for so many years.”

All that remained, the final step in the project, was to get the PC version, still a bit buggy, working smoothly. But after a few more code revisions, it was finally bug-free.

“We had pretty much completed all the technical hurdles,” Dorsey said. “Everybody felt, finally, this is it, it’s going to work, this is done.”

And that is when Hawking got sick, in February.

---

According to Wood, Hawking continued using the emulator until his final days. He was able to talk with his loved ones and caregivers with the new software on the Raspberry Pi. The last words he spoke while wired into his chair, whatever they were, he spoke with a version of his voice that lives only in code, potentially deathless bits and bytes.

Everyone on the project understood that Hawking might not live long enough to get much use out of the emulator. He had been sick before, but always recovered. In 2014, when Wood first contacted Dorsey, Hawking was 72. They decided, though, that his CallText boards could keel over in six months, while Hawking might live to 80.

Along with sadness at Hawking’s death, Dorsey can’t help feeling some disappointment. He and the team had raced for years to build a complicated thing that had worked beautifully, but now sat idle.

At the same time, the project brought him back to his younger self, the guy who wanted to use engineering to perform good deeds and help people. All those years ago, working on the intonation algorithm in the CallText, he couldn’t imagine that it would end up helping define a genius of science to the world, and even to himself.

#### Hawking's voices

Compare Hawking's 1986 voice with the one he rejected in 1996 and the emulated version from 2018

1986    1996    2018

Tech changes fast. Most machines end up as dust, and when we die, our voices die with us. Hawking’s voice is different. The original CallText boards have passed to his estate, to use as his family wishes. So has the new software, the CallText emulator, which can be ported to future platforms as they are invented.

Hawking was, famously, an atheist, skeptical of the afterlife; “We have this one life to appreciate the grand design of this universe,” he once said, “and for that, I am extremely grateful.” But there is no longer any physical reason his voice can’t live forever.