

Quick-freezing alloy accelerates phase-change memory

The demonstration of subnanosecond bit-writing times clears a key hurdle for the next-generation computing candidate.

Our computers rely on a motley mix of technologies to meet their data-storage needs. Document files and other occasionally accessed data typically live on magnetic hard disks or flash drives, which are slow but can store information reliably for decades. Data that must be retrieved more frequently, including processor instructions for calculations, are stored in dynamic random-access memory (DRAM), whose capacitor-based bits are short-lived but manipulable on nanosecond time scales. And for the most demanding calculations, known as cache operations, there is static random-access memory (SRAM). Composed of six-transistor circuits, SRAM bits take up copious chip real estate but can be written or erased in a fraction of a nanosecond.

Shuttling data between the different types of memories—all housed at different physical locations—makes for sluggish processing, especially in applications that handle massive amounts of data. About a decade ago, computer scientists began dreaming of a memory that would do it all—one with the requisite combination of stability, scalability, and speed to serve as long-term storage, cache memory, and all things in between.¹

The dream of a universal memory was inspired in large part by the emergence of phase-change memory, in which bits are stored in tiny alloy fragments that can be switched between crystalline and amorphous-solid states by electric pulses. The bits could easily preserve their state for a decade, and they could be written or erased in tens of nanoseconds—almost as fast as DRAM.

But after years of failed attempts to

substantially improve on those speeds, the dream faded. “It is generally agreed that the vision of a ‘universal memory’ is not realistic,” Stanford University’s Philip Wong and University of California, Berkeley’s Sayeef Salahuddin wrote in a 2015 commentary.²

Now researchers at Xi’an Jiaotong University and the Shanghai Institute of Microsystem and Information Technology (SIMIT), both in China, have demonstrated a phase-change memory cell³ that team member Wei Zhang says “should bring the universal memory concept back into consideration.” The group used a rational-design approach to concoct a quick-freezing alloy that can switch phases in just 700 ps—on par with SRAM write times.

Significant technical hurdles remain

to parlay the advance into a universal memory, but the speed boost immediately makes phase-change memory a more attractive candidate to replace DRAM—and perhaps some SRAM—in next-generation computers.

Mushroom memories

Among phase-change memory’s chief virtues is its simplicity. Each memory cell essentially consists of a semiconducting film—the phase-change alloy—sandwiched between a metal electrode and an insulator. A tiny metal pillar running through the insulator serves as the second electrode, as shown in figure 1. In some space-saving designs, the entire phase-change memory cell takes up an area of just $4F^2$, where F is the dimension of the smallest printable chip feature. (Today’s commercial chips commonly have features as small as 10 nm.) A standard SRAM cell has a footprint of roughly $100F^2$.

The process of writing and reading bits of phase-change memory is also straightforward. A brief, intense electric pulse across the two electrodes switches the alloy film from a fully crystalline state, the set state, to the partially amorphous reset state. The pulse causes the alloy near the pillar to rapidly melt and recool, quenching those atoms in a metastable, liquid-like configuration. In the most common design, pictured in figure 1, the amorphous phase forms a half sphere that tops the pillar like the cap of a mushroom.

The set state can be restored with a longer, less intense pulse that causes the quenched region to warm and cool more slowly, allowing the atoms to resettle in their preferred crystalline arrangement. The alloy’s electrical resistance

is orders of magnitude higher in the reset state than in the set state, so a stored bit can easily be read out with a weak applied voltage.

The bottleneck of the process is the recrystallization step, or set operation. And that step can’t be hurried; it’s limited by the alloy’s intrinsic crystallization rate. Three decades ago, researchers at Matsushita Electric Industrial (now Panasonic) discovered they could perform fast set operations using semiconductor layers that combined tellurium




FIGURE 1. A PHASE-CHANGE MEMORY CELL in the “reset” state, with its semiconductor layer partially quenched in an amorphous phase. The amorphous region appears as a dark semicircle. In the “set” state, the layer would be entirely crystalline. Because the electrical resistances of the set and reset states differ by orders of magnitude, a weak voltage applied across the electrodes can be used to read the cell’s state without disturbing it. The pillar-shaped bottom electrode is about 15 nm wide. (Image credit: IBM Research.)




FIGURE 2. A SEEDLING CRYSTAL, simulated with molecular dynamics. Under conditions typical of the amorphous-to-crystalline transition in a phase-change memory device, a precursor crystallite of scandium (red) and tellurium (blue) efficiently seeds a bulk crystalline phase of Te and antimony (yellow) that quickly fills the roughly 10 nm³ simulation volume. Without such precursors, nascent SbTe crystals would typically disintegrate before growing large enough to stabilize. (Adapted from ref. 3.)

with germanium, antimony, or both. Since then, most phase-change memory devices have used some variation of that recipe, known as GST. But for decades the physics of GST's fast crystallization was a mystery.

In the mid 2000s, researchers began to lift the veil. A key was the alloy's rock-salt crystalline structure: Cations of Ge and Sb and anions of Te occupy alternating sites on a cubic lattice, like a three-dimensional checkerboard. A series of experiments and simulations found evi-

dence that fragments of the checkerboard—four-atom rings termed ABAB squares—persisted even in the amorphous phase.⁴ It followed that those squares might serve as prefabricated building blocks that hasten the assembly of fledgling crystals.

Under the heat of a set operation, however, the ABAB squares are ephemeral. They are constantly being formed and broken by thermal fluctuations, and as a result, most nascent crystallites disintegrate before growing large

enough to nucleate a thermodynamically stable crystalline phase. So when the team at Xi'an Jiaotong and SIMIT recently set out to speed up the alloy's phase changes, they started by asking, How can we make ABAB squares more durable?

Lords of the rings

For clues, the researchers revisited previous work by team members Feng Rao and Zhitang Song. At SIMIT, Rao, Song, and their coworkers had improved on GST's write speeds by doping SbTe not with Ge but with titanium, a transition metal.⁵ Measurements suggested that the metal formed durable bonds with Te at angles that could, in theory, give rise to ABAB squares. Inspired by those results, the Xi'an Jiaotong-SIMIT collaboration decided to screen the remaining transition metals for those that might bond with Te to form even stronger rings.

Too much metal would spoil the alloy's electrical properties, so the plan was to dope the mixture just enough to create a sparse dispersion of crystalline precursors—durable, transition-metal-rich assemblies of a few dozen ABAB

Nanopositioning Systems

Piezo Nanopositioning Systems

UHV Nanopositioners

Low Noise PicoQ[®] Sensors

AFM, NSOM & Microscopy

High Precision Micropositioning

Custom Solutions



MCL
MAD CITY LABS INC.

madcitylabs.com

Explore What's Possible Pure Indium and Indium Alloys

Unique properties to achieve:

- Low-temperature alloys: solid or liquid
- Glass-to-metal soldering
- Cryogenic or high CTE interfaces

Indium Corporation delivers:

- Verified purity, batch analysis
 - Custom metal forms and compounds
 - Application support:
- Over 80 years of R&D with indium metal

Contact our engineers today: jhisert@indium.com

Learn more: www.indium.com/applicationsPHYI

**From One Engineer
To Another[®]**

©2018 Indium Corporation



INDIUM
CORPORATION[®]

squares. The precursors would then seed the growth of a bulk SbTe crystalline phase. To make sure the bulk crystal could grow seamlessly from the precursors, the ABAB squares would need to closely match the geometry and spacing of SbTe's rock-salt lattice.

Guided by density functional theory calculations, the team whittled the three dozen contenders down to one, scandium. Molecular dynamics simulations confirmed that under typical set-operation conditions, ScTe squares survived 10 times as long as their GST counterparts and, as shown in figure 2, could effectively seed the growth of SbTe crystals. But the proof was in the experimental pudding: The researchers fashioned a prototype memory cell from their scandium-antimony-tellurium concoction, dubbed SST, and showed that they could perform set operations with the device in just 700 ps.

"I was surprised that such fast crystallization is possible," says SangBum Kim, a research scientist at IBM. The write-speed improvements are especially astonishing, he adds, because they come at no loss to data retention time. "Usually there's a trade-off."

Indeed, the Xi'an Jiaotong-SIMIT collaboration effectively lowered the energy barrier for the transition between logic

states, and intuitively, that should also raise the odds of unintended bit flips. But SST's estimated data-retention time is actually slightly better than that of GST.

Zhang chalks up the counterintuitive result to competing kinetic effects: Although the scandium-laced precursors lower the energy barrier for nucleation, they also increase viscosity, which stifles diffusion and thereby slows proliferation of the crystalline phase. At the elevated temperatures that prevail during a set operation, the energy-barrier effect wins out. But under the much cooler conditions at which bits are stored, the viscous effect wins.

Big step for big data?

Although the switching speeds of the SST-based device are suitable for a universal memory, the endurance currently is not. The scandium-doped cells become unreliable after about a million logic cycles, too few to stand up to the rigors of cache operations. Zhang and his colleagues think they can improve endurance with tweaks to the device setup and material composition—and there's no fundamental reason they can't. But they'll need to do so by orders of magnitude to make the technology a viable cache memory.

Even if phase-change memory never fully supplants SRAM, it is already durable, stable, and fast enough to contend with magnetoresistive random-access memory and resistive random-access memory in the competition to supplant DRAM as the workhorse memory of next-generation computers. Unlike DRAM, phase-change memory is nonvolatile: It retains information even without a power supply. That and its high storage density make it especially attractive for energy-intensive big-data applications such as machine learning and neuromorphic computing.

"If you use these nonvolatile memory devices with neuromorphic computing, you can save a lot of energy and power," says Kim. "And you could accelerate computations by orders of magnitude."

Ashley G. Smart

References

1. M. Wuttig, *Nat. Mater.* **4**, 265 (2005); J. Åkerman, *Science* **308**, 508 (2005).
2. H.-S. P. Wong, S. Salahuddin, *Nat. Nanotechnol.* **10**, 191 (2015).
3. F. Rao et al., *Science* (in press), doi:10.1126/science.aao3212.
4. S. Kohara et al., *Appl. Phys. Lett.* **89**, 201910 (2006); J. Hegedüs, S. R. Elliott, *Nat. Mater.* **7**, 399 (2008).
5. M. Zhu et al., *Nat. Commun.* **5**, 4086 (2014).

Attosecond measurements reach electronvolt precision

With a train of ultrashort light pulses, researchers disentangle energetically similar photoionization channels—and solve a seven-year-old puzzle.

If you know where you are, you don't know where you're going, and if you know where you're going, you don't know where you are: The Heisenberg uncertainty principle fundamentally limits how precisely position and momentum can simultaneously be known. Another formulation of the principle, linking the uncertainties of energy and time, would seem to suggest that ultrafast measurements—such as the attosecond experiments that probe electronic motion in atoms, molecules, and solids—can't be both temporally and spectrally precise.




FIGURE 1. PHOTOIONIZATION OF NEON probed by a train of attosecond pulses (green) combined with an IR wave (purple). The ionization signal exhibits constructive or destructive interference, depending on the photoionization time delay and the timing offset of the two waves.

But that's not quite true. The uncertainty principle does place a lower bound on the product of the time and energy uncertainties of an isolated electromagnetic pulse. An experiment's temporal precision, however, is not limited by the duration of the pulses used in it; rather,

it's often determined by the degree of control over the relative timing of two pulses. Extreme UV (XUV) pulses with durations on the order of 100 attoseconds can nevertheless make measurements with few-attosecond precision.

Spectral precision is another matter.