

Better Memory


Advances in non-volatile memory are changing the face of computing and ushering in a new era of efficiencies.

SINCE THE DAWN of computing, an ongoing challenge has been to build devices that balance the need for speed and persistent storage. While dynamic random-access memory (DRAM) is fast, it can only hold data as long as it receives an electrical current; when the computing device is switched off, the data disappears. And although storage devices such as hard drives are efficient for holding large volumes of data, they are relatively slow. The result? “A performance or persistence choice that doesn’t give you the best of both worlds,” states David Andersen, an associate professor in the computer science department at Carnegie Mellon University.

Over the last few years, engineers have resolved some of these challenges through solid state drives (SSDs) that contain no disk or other moving parts, yet continue to store data when the devices are switched off. What is more, SSDs use less power and provide higher reliability than hard disk drives. However, they are far from ideal. For one thing, they’re still relatively expensive. For another, while SSD is often an improvement over older technologies and sometimes reduces the need for DRAM, it still does not provide the level of speed, flexibility, and lifespan that users desire.

“There is a desire for more advanced technology, particularly in high-performance computing systems,” says Jim Handy, memory analyst at market research firm Objective Analysis.

All of this is leading researchers down the path to faster and more advanced non-volatile random-access memory (NVRAM) technologies. These technologies—some of them radically different than today’s flash storage technologies—could usher in speed and performance efficiencies that change computing. Unlike DRAM, these systems do not necessarily store the ones



Hewlett-Packard Enterprise Memristor devices on a 300mm wafer.

and zeros of binary code on a capacitor; they instead use memristors, which rely on electrical resistance. This produces efficiency gains, but also energy savings. These technologies could ultimately replace today’s flash, SSD, static random-access memory (SRAM) and dynamic random-access memory (DRAM).

They have names like 3D XPoint, MRAM, MeRAM, Memistor, NRAM, STT-RAM, PCM, CBRAM, RRAM, Millipede, and Racetrack. Much faster persistent memory is a potential game-changer for high-performance clusters and transaction-oriented systems be-

cause checkpoints must become persistent before an operation completes.

Says Andrew Wheeler, vice president and deputy director of HP Labs: “Today’s computer architecture is fundamentally unchanged. It’s the same architecture we’ve been using for 60 years—processors with a fixed amount of local memory, connected to storage and memory over an I/O bus. NVRAM becomes really interesting when you introduce the opportunity to simultaneously reinvent the architecture.”

Flash Forward

The ability to design a more advanced memory architecture would have a profound impact on everything from high-performance computing clusters to smartphones and devices that comprise the Internet of Things. The technology could change basic computing architectures and storage designs, and address issues such as battery life, power requirements, in-memory database (IMDB) designs, and the way applications are coded. “Today, flash (memory) occupies the middle ground between speed and durability. It isn’t

These technologies could usher in speed and performance efficiencies that change computing.

as fast as DRAM and it isn't as durable as a disk drive. The goal is to close the gap further so that it's possible to address the challenges related to large databases and increasingly complex computing problems, as well as consumer devices," Andersen explains.

At HP, for example, researchers are working to develop memristor technology that uses electrons for processing, photons for communication, and ions for storage. "The Machine" creates a vast pool of fast NVRAM, connected to task-specific processors over a high-bandwidth, low-latency photonic fabric. The goal, Wheeler says, is to build a system that better optimizes logic gates while delivering long-term storage. HP refers to the approach as Memory-Driven Computing (MDC). "Every buffer copy or block move that we can design out saves energy, reduces the chances for interception or corruption, and shrinks the security attack surface," he says. The technology, which the company hopes to have commercially available by 2016, would tackle petascale datasets that are beyond reach today.

Memristor technology would consume a fraction of the power of today's memory systems. "At the tiny scale, having tens of terabytes of virtually zero-power memory allows us to build a new class of smart, secure IoT devices that can store their experience to know what's normal and what's novel or to satisfy a query from a neighboring peer or central intelligence," Wheeler explains. "Applications and operating systems that are fully adapted to pervasive non-volatile memory could enable perpetual computing where there is no more 'off switch'. When sufficient energy is present, information is processed; otherwise, the current state is preserved." HP hopes to have The Machine available in a range of form factors over the coming decade, based on price and performance.

HP is not the only player in the space. Intel and Micron are collaborating on a technology called 3D XPoint memory, which the companies claim is 1,000 times faster than the NAND flash storage used in current memory cards and in solid state drives. The dual in-line memory modules (DIMMs) are designed to be compatible with today's DDR4 SDRAM (double data rate fourth-generation synchronous dynamic

random-access memory) but deliver a fourfold capacity increase. The proprietary solution offers performance gains without modifications to the underlying operating system or applications. However, the platform would require a re-designed central processing unit (CPU) and new extensions in order to take advantage of the 3D XPoint technology. Analysts say the technology would benefit organizations running large numbers of servers in a datacenter. 3D XPoint, for instance, would anticipate when data is required and transfer it in advance to the 3D XPoint component.

Other technologies are emerging as well. For instance, Crossbar has produced a working test chip for its RRAM (resistive random-access memory) technology. The company claims the system delivers 100 times lower read latency than NAND flash storage, along with 20 times faster writes without any block erase design constraints. It also delivers up to 1 terabyte of storage on a single chip, in an architecture that is 3D-stackable and scalable to sub-10 nanometers. Within the chip, each cell surrounds an insulating switching medium between electrode layers. When electrodes receive voltage nanoparticles in the switching medium, they form a conductive filament. The design supports stacking and it can be scaled down to fabrication nodes smaller than five nanometers.

Another technology, MeRAM (magneto-electric random access memory), replaces the electrical current of spin-transfer torque (STT) with voltage to

Memristor technology, using electronics for processing, photons for communication, and ions for storage, would consume a fraction of the power of today's systems.

write data. This nanoscale approach results in 10 times to 1,000 times greater energy efficiency.

"At this point, nobody knows which of the horses in the NVRAM game will win or how things will play out, but the bottom line is that the technology will very likely make a big impact on computing," Andersen says.

Making it All Compute

The practical benefits of next-generation NVRAM could be profound. Los Alamos National Laboratory began using NAND flash storage for high checkpointing and other high-performance computing in its Trinity system in September 2015. The National Energy Research Scientific Computing Center (NERSC) Cori system also will utilize the concept. Trinity holds nearly 2 petabytes of DRAM in main memory and 4 petabytes of NVRAM to support an I/O enhancement—essentially a new storage layer—known as a burst buffer. Gary Grider, division leader for the Los Alamos High Performance Computing Division, says more advanced versions of the technology will be incorporated into future CORAL (Collaboration of Oak Ridge, Argonne, and Livermore) supercomputers that will tap into the knowledge gained from the new tier of storage in Trinity.

Grider says NVRAM advances will have a major impact on supercomputers and also on consumer devices, including laptop computers, smartphones, and cameras. As prices drop and the technology advances, "It will become far more ubiquitous."

He also believes next-generation NVRAM could make today's data storage hierarchies obsolete. He points to the Intel-led U.S. Department of Energy Storage Fast Forward DAOS (Distributed Application Object Storage) Project, which targets HPC applications with scalable, transactional, versioned, and end-to-end reliable exploitation of multiple tiers of non-volatile storage. It will exploit non-volatile storage on compute node, in system burst buffer nodes, and on remotely attached parallel disks systems.

In addition to I/O use cases, next-generation NVRAM could be harnessed by some applications for out-of-core direct use in order to tap into a slower but larger memory pool.

Andersen believes NVRAM could help make future devices smaller and cheaper, as well as speeding start-up times for certain types of devices and sensors. “A suspend-resume mode has a lot of advantages for sensors and actuators that are part of the Internet of Things. The goal for these devices is to be insanely cheap and efficient.”

NVRAM promises to deliver benefits at an equivalent or cheaper per-giga-byte price point as today’s flash technology, he says. It could also deliver improvements to today’s battery-backed database technology and in NOR flash, which is often used in mobile phones because it consumes minimal energy during the write process.

Marc Staimer, president of Dragon Slayer Consulting, says next-generation NVRAM will introduce new functions and capabilities “that will be developed over time.” He believes the technology, like early flash technology, will initially “show up at the consumer end and prove itself out” before enterprises and others begin using it for high-end data center requirements. “You will likely see it in smartphones, tablets, and laptops before you see it in servers and storage systems on a widespread basis.”

When the technology does move into data center systems, it will not make NAND flash storage immediately obsolete, just as NAND flash storage did not make hard drives immediately obsolete. “There will be a cost differential that will place these new NVRAM technologies (in a) high-performance solutions tier with a higher cost,” he says. “Over time that will gradually change, and variations of the new NVRAM technologies will move downstream to lower tiers, squeezing out slower NAND flash storage.”

For now, Handy says the industry must begin to define standards for how these memory technologies will communicate with standard programming interfaces. Meanwhile, system manufacturers will have to place their bets on which new NVRAM technology makes it to market first with the desired characteristics.

Nevertheless, the writing appears on the wall or, perhaps, in the chips. Says Staimer: “These technologies are not just an iteration of existing technology; they are a breakthrough.

“Nobody knows which of the horses in the NVRAM game will win or how things will play out, but the bottom line is that the technology will very likely make a big impact on computing.”

This is not just another generation of NAND flash; it is a significant leap forward in performance and wear-life well above today’s flash. It will change computer architectures, break down the barriers between memory and storage, and ultimately change how we do computing.”

For more on non-volatile memory, see the article by Nanavati et al. in this issue on page 56. **C**

Further Reading

Pelley, S., Wenisch, T.F., Gold, B.T., and Bridge, B., Storage Management in the NVRAM Era, Proceedings of the VLDB Endowment, Volume 7 Issue 2, October 2013 Pages 121-132. <http://dl.acm.org/citation.cfm?id=2732231>.

Oh, S., and Ryu, Y., Multi-core Scheduling Scheme for Wireless Sensor Nodes with NVRAM-Based Hybrid Memory, Ubiquitous Computing Application and Wireless Sensor, Volume 331 of the series Lecture Notes in Electrical Engineering, pp 45-52. http://link.springer.com/chapter/10.1007/978-94-017-9618-7_5#

Van Essen, B., Pearce, R., and Gokhale, M., On the Role of NVRAM in Data-intensive Architectures: An Evaluation, Center for Appl. Sci. Comput., Lawrence Livermore Nat. Lab., Livermore, CA, USA, http://ieeexplore.ieee.org/xpl/articleDetails.jsp?tp=&arnumber=6267871&url=http%3A%2F%2Fieeexplore.ieee.org%2Fxppls%2Fabs_all.jsp%3Farnumber%3D6267871.

Samuel Greengard is an author and journalist based in West Linn, OR.

ACM Member News

ABTRUSE REASONING IDENTIFIES SOLUTIONS TO PRACTICAL PROBLEMS



Alan Bundy is a problem solver.

A professor of Automated Reasoning in the School of Informatics

at the University of Edinburgh, Bundy utilizes abstruse mathematical reasoning to find solutions to practical problems in the development and maintenance of computing systems.

Author of more than 270 publications and recipient of more than 60 research grants, Bundy’s research focuses on the automation of mathematical reasoning, combining artificial intelligence with theoretical computer science.

His current research interest involves the need for automated mechanisms of representational change in autonomous agents “that have to cope with a changing environment to find a mechanism that aids in efficient solutions to problems,” he says.

Bundy is proudest of his work on automating the conjecturing of intermediate lemmas—a side-theorem required to prove a main theorem. Bundy’s lemma illustrated that the failure of an initial proof attempt of the main theorem could sometimes be diagnosed to suggest the form of the missing lemma. “It’s like creatively plotting a cross-country journey,” he explains. “We developed an idea of a plan; when it went wrong you could figure out what you needed.”

That “generic theorem” resulted in an entire family of proofs that took Bundy over five years to construct.

A fellow of the Royal Society of Edinburgh and of the Royal Academy of Engineers, Bundy was named a Commander of the Most Excellent Order of the British Empire (CBE) by Queen Elizabeth in 2012. “I was nervous and forgot the proper form to address Her Majesty. She’s a game lady and pretended not to notice,” he laughed.

—Laura DiDio

Copyright of Communications of the ACM is the property of Association for Computing Machinery and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.