The Evolution of Phase Change Memory

Why PCM is Ready for Prime Time as a Next-Generation, Nonvolatile Memory

Greg Atwood

Micron Technology, Inc.
July 26, 2010

Phase change memory is a term used to describe a class of nonvolatile memory (NVM) devices that exploit the ability of certain materials to rapidly change phase between two stable physical states. Phase change memory blends the attributes commonly associated with NOR-type Flash memory, NAND-type Flash memory, EEPROM memory, and DRAM and is a leading candidate for the next generation of NVM, expanding its use in computing and storage systems (see Figure 1). As traditional electron storage-based memory, such as NOR and NAND Flash, begin to encounter scaling difficulties, PCM is considered to be the best candidate to continue the scaling of NVM.

<table>
<tr><th>Attributes</th><th>PCM</th><th>EEPROM</th><th>NOR</th><th>NAND</th><th>DRAM</th></tr>
<tr><td>Nonvolatile</td><td colspan="4">Yes</td><td>No</td></tr>
<tr><td>Scalability</td><td><10nm</td><td>-4x</td><td>~3x nm</td><td>~1x nm</td><td>~2x nm</td></tr>
<tr><td>Bit Alterable</td><td colspan="2">Yes</td><td colspan="2">No</td><td>Yes</td></tr>
<tr><td>Erase Required</td><td colspan="2">No</td><td colspan="2">Yes</td><td>No</td></tr>
<tr><td>Software</td><td colspan="2">Easy</td><td>Moderate</td><td>Hard</td><td>Easy</td></tr>
<tr><td>Write Speed</td><td>~100 MB/s</td><td>~30 KB/s</td><td>~1 MB/s</td><td>~20 MB/s</td><td>~1 GB/s</td></tr>
<tr><td>Read Speed</td><td>50–100ns</td><td>~200ns</td><td>70–100ns</td><td>15–50µs</td><td>20–80ns</td></tr>
<tr><td>Endurance</td><td>10^{6-8}</td><td>10^{5-8}</td><td>10^{5}</td><td>10^{4-5}</td><td>Unlimited</td></tr>
</table>

Figure 1: PCM—the new class of nonvolatile memory— brings together the best attributes of NOR, NAND, and RAM.

Some of the unique capabilities of PCM include:

Nonvolatile

Like NOR Flash and NAND Flash, PCM is nonvolatile. DRAM, of course, requires a constant power supply, such as a battery backup system, to retain information, resulting in higher power consumption. DRAM technologies also suffer from susceptibility to so-called soft errors or random bit corruption caused by alpha particles or cosmic radiation, an effect not observed in PCM.

Scalability

Existing memory, such as NOR, NAND, EEPROM, and RAM, all rely on charge (electron) storage as their memory mechanism. As scaling progresses to ever-smaller dimensions, the number of stored electrons diminishes to the point where reliability is degraded and continued scaling becomes very difficult. This phenomenon requires the introduction of increasingly complex structures with diminishing returns. PCM does not use electrons, but instead uses a physical phase change as the storage mechanism. Stability of the stored phase has been demonstrated to be at least as small as 5nm. Scalability is one of the major motivations for the development of PCM.

Bit-Alterable/Erase/Software

Like RAM, PCM is bit alterable. Unlike RAM and PCM, Flash technology requires a separate erase step involving the manipulation of a large block of data to change a small amount of information. Bit alterability can greatly simplify the use of the memory, and, combined with the lack of a need to erase, makes software management of the memory much easier. In some usage environments, PCM can be as easy to use as RAM.

Write Speed

PCM is capable of achieving write bandwidth comparable to NAND, but with 100x lower initial latency (time to write the first byte) and with no separate (and slow) erase step required. The lack of a slow erase step before write can greatly increase the overall effective write bandwidth, and the low latency can allow PCM to be used, in many cases, like a direct write memory as opposed to the mass storage disk usage model for NAND. The write speed (bandwidth and latency) of PCM does not match the capability of DRAM although with proper management, some DRAM usage can be displaced for infrequent or managed write applications.

Read Speed

Like RAM and NOR-type Flash, PCM technology features fast random read access times. This enables the execution of code directly from the memory without an intermediate copy to RAM. In contrast, NAND Flash has long random access times—on the order of 10s of microseconds—that prevents it from being used for direct code execution.

The combination of these attributes uniquely positions PCM with an opportunity to provide the next generation of nonvolatile memory with an expanded set of performance capabilities, sitting solidly between today's DRAM and NAND.

The Past–The History of Phase Change Memory

The existence of materials that exhibit a controllable change in phase has been known for many years and has been in use for optical memory applications. Electronic memory based on these materials has recently experienced a resurgence of interest for use as a next-generation NVM for the reasons addressed earlier in this article. Pioneering work conducted by Micron Technology, Inc. and others has moved the technology to the forefront of the memory industry R&D activity with a promise to alter the way NVM is used in memory systems.

The history of phase-change materials can be traced back to work starting in the 1950s by Dr. Stanford Ovshinsky who was researching the properties of a class of glassy materials that exhibited the ability to easily and stably change between two phases. By the late 1960s, he had reported that certain of these materials exhibited a reversible change both in resistivity and reflectivity when changing between an ordered (polycrystalline) state and a disordered (amorphous) state. It was recognized that this effect could be exploited both for optical memory as well as electronic memory. In a September 28, 1970 issue of *Electronics*, Energy Conversion Devices (ECD), a company formed by Dr. Ovshinsky, in collaboration with Intel's Gordon Moore, reported the world's first electronic phase change memory array, a 256-bit semiconductor device.

Nearly 30 years later, Ovonyx, a joint venture between ECD and Tyler Lowrey, the former CTO/COO of Micron Technology, was formed. In February 2000, Intel and

Ovonyx announced collaboration and a licensing agreement that spawned the modern age of research and development in PCM. In December of 2000, STMicroelectronics (ST) and Ovonyx also began a collaboration. By 2003, the three companies had joined forces to accelerate progress on the technology by avoiding duplication in basic, precompetitive R&D and through expanding the research scope. In 2005, ST and Intel agreed to codevelop a 90nm PCM technology.

In 2008, ST and Intel combined their NOR, NAND (ST's NAND), and PCM business to form a new Flash company called Numonyx (now part of Micron). The formation of what was Numonyx further accelerated progress in the development of PCM, resulting in the first commercial PCM product at the end of 2008. In the intervening years since that first significant work in 1970, much progress has been made in semiconductor manufacturing technology, enabling the practical development of PCM both for optical and electronic storage devices. Phase-change materials have been in use for many years for high-volume rewritable CDs and DVDs. With the start of production of phase-change materials for electronic memory by Numonyx and others, PCM begins to deliver on its promise to expand the usage of nonvolatile memory.

The Present–The Technology Behind Phase Change Memory

The PCM technology uses a class of materials known as chalcogenides (prounounced kal-KOJ-uh-nydes). Chalcogenides are alloys that contain an element in the oxygen/sulphur family of the periodic table (group 16 in the new style or group VIa in the old-style periodic table). Micron® PCM uses an alloy of germanium, antimony, and tellurium (Ge2Sb2Te5), known more commonly as GST. Most companies performing research and development in PCM today are using GST or closely related alloys.

Phase-change chalcogenides exhibit a reversible phase change between the amorphous phase and the crystalline phase. As illustrated in Figure 2, in the amorphous phase, there is an absence of regular order to the crystalline lattice. In this phase, the material demonstrates high resistivity and low reflectivity. In contrast, in the polycrystalline phase, the material has a regular crystalline structure and exhibits high reflectivity and low resistivity. As shown in the Sample column of Figure 2, these changes in phase are observable.

Structure	*Order*	*Resistivity*	*Properties*
Amorphous	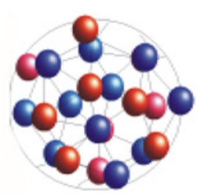		Short-range atomic order High reflectivity High resistivity
Polycrystalline	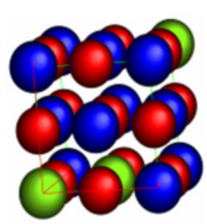	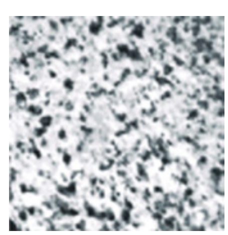	Long-range atomic order Low reflectivity Low resistivity

Figure 2: Phase change chalcogenides exhibit a reversible phase change between the amorphous phase and the crystalline phase.

Theory Implementation

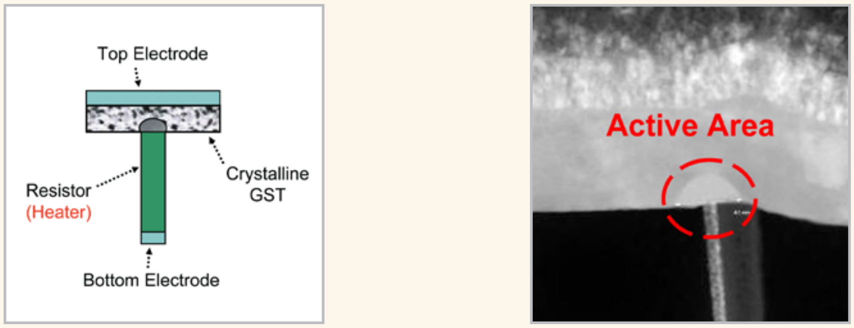

Figure 3: A graphical representation and image of a basic PCM storage element.

In PCM, we are exploiting the difference in resistivity between the two phases of the material. This phase change is induced in the material through localized Joule heating caused by current injection. The final phase of the material is modulated by the magnitude of the injected current and the time of the operation.

Micron is currently in production with two 128Mb, 90nm PCM parts (P5Q serial PCM and P8P parallel PCM) developed specifically for embedded applications, such as EEPROM and SPI NOR consolidation, software simplification, high endurance, or battery-backed SRAM replacement. The parts are attractive for these applications specifically because they deliver the features discussed above, including byte alterability and high endurance.

The Future–Looking Ahead at Phase Change Memory

Figure 4 shows a 1Gb, 45nm, state-of-the-art, PCM die currently in development. PCM is on a trajectory to converge with DRAM for lithography capability with the next generation of 32nm, which is due in the next few years.

PCM may eventually be used in ultra high-performance memory subsystems to achieve solid state drive (SSD) performance and reliability that is unachievable with NAND and at power consumption levels (and nonvolatility) that cannot be achieved with RAM. As system

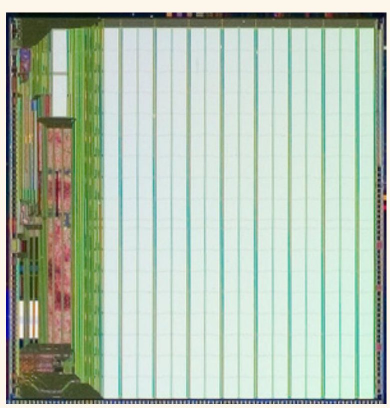

Figure 4: 1Gb, 45nm PCM die.

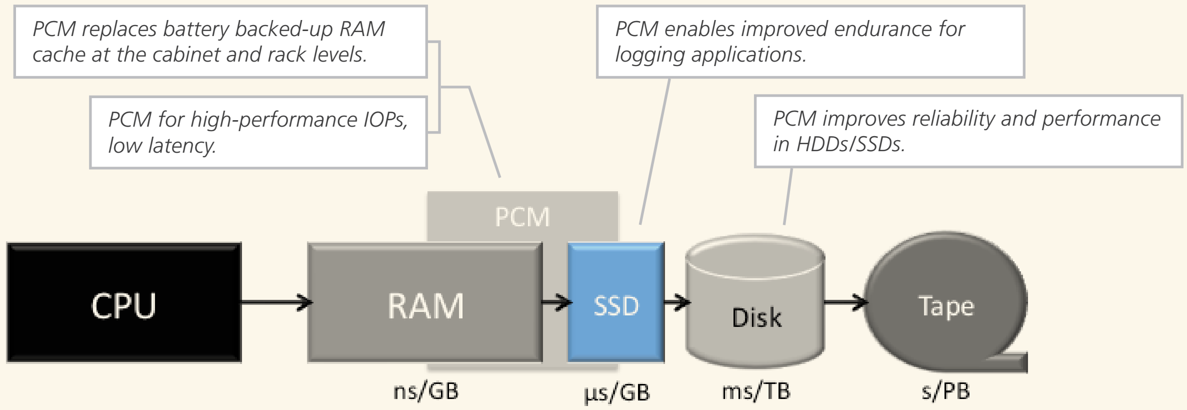

Figure 5: PCM in the memory hierarchy.

software evolves to fully exploit the capabilities of PCM, it will be able to move from an I/O drive-based usage model to a directly memory-mapped main-memory-based usage model, further increasing its performance value. This migration is shown in Figure 5. As the PCM cost structure continues to decline through scaling at a faster rate than DRAM, the number of applications will increase. However, it is not envisioned to directly challenge NAND on a cost basis.

Looking further into the future, a derivative of PCM, called PCMS (stacked PCM), is being researched today. PCMS enables the stacking of several layers of PCM memory arrays on top of each other and promises to deliver costs that are lower than NAND, but with the scalability of PCMS. Should the research be successful, PCMS could provide a natural evolution of the phase change memory concept.

Phase Change Memory— The Next Generation of Nonvolatile Memory— Ready Today

The road for PCM has been a long one. It started with the basic concept in the 1950s, evolved through early feasibility demonstration in the 1970s, and moved on to a resurgence of activity in the 2000s, resulting in products with the potential to provide breakthrough capabilities in the 2010s. The evolution of the semiconductor industry has enabled the manufacture of structures small enough to exploit the unique storage capabilities of PCM, while, in parallel, volume usage of PCM materials for optical memory has driven the rapid learning of the material properties and physics. This trend, combined with the 10 years of development by Intel, ST, and Micron, now has PCM ready for prime time as a next-generation, nonvolatile memory at a time when traditional electron storage memory is beginning to encounter difficulties. Expect PCM to leverage its unique capabilities as a storage class memory to find a new position in the memory hierarchy.

Greg Atwood *is a senior fellow at Micron Technology, Inc. He received an M.S. degree in physics from Purdue University in 1979, joining Intel Corporation in the same year. At Intel, he worked on numerous technology development programs including Logic, SRAM, EPROM, E2PROM, Flash, Multi-level Flash, and phase-change memory. In 2008, he joined Numonyx B.V. as their first senior fellow with a focus on phase-change products; Numonyx was acquired by Micron in May 2010. Since 2002, Greg has been primarily focused on developing phase-change memory and bringing the technology to market. He is the author of numerous articles and papers and holds more than 40 patents.*