# Chapter 1
Introduction

Yuan Xie

Abstract Emerging non-volatile memory (NVM) technologies, such as PCRAM and STT-RAM, are getting mature in recent years. These emerging NVM technologies have demonstrated great potentials to be the candidates for future computer memory architecture design. It is important for SoC designers and computer architects to understand the benefits and limitations of such emerging memory technologies, to improve the performance/power/reliability of future memory architectures. This chapter gives a brief introduction of these memory technologies, reviews recent advances in memory architecture design, discusses the benefits of using at various levels of memory hierarchy, and also reviews the mitigation techniques to overcome the limitations of applying such emerging memory technologies for future memory architecture design.

1.1 Introduction

In the modern computer architecture design, the instruction/data storage follows a hierarchical arrangement called **memory hierarchy**, which takes advantage of locality and performance of memory technologies. Memory hierarchy design is one of the key components in modern computer systems. The importance of the memory hierarchy increases with the advances in performance of the microprocessors. Traditional memory hierarchy design consists of embedded memory (such as SRAM and eDRAM) as on-chip caches, commodity DRAM as main memory, and magnetic hard disk drivers (HDD) as the storage. Recently, solid-state drives (SSD) based on NAND-flash memory have also gained the momentum as the replacement or cache for the traditional magnetic HDD. The closer the memory is placed to microprocessor, the faster latency and higher bandwidth are required, with the penalty of the smaller

Y. Xie (✉)
Pennsylvania State University, University Park, PA, USA
e-mail: yuanxie@cse.psu.edu
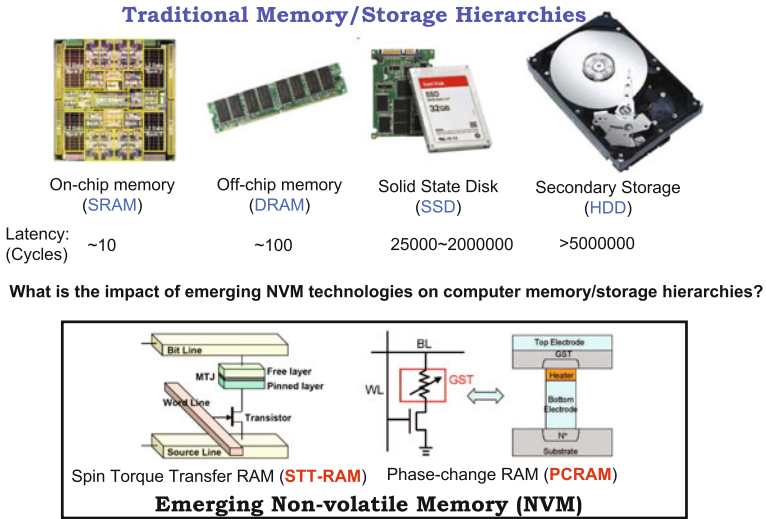
Fig. 1.1 What is the impact of emerging memory technologies on traditional memory/storage hierarchy design?

capacity. Figure 1.1 illustrates a typical memory hierarchy design, where each level of the hierarchy has the properties of smaller size, faster latency, and higher bandwidth than lower levels, with different memory technologies such as SRAM, DRAM, and magnetic hard disk drives (HDD).

Technology scaling of SRAM and DRAM (which are the common memory technologies used in traditional memory hierarchy) are increasingly constrained by fundamental technology limits. In particular, the increasing leakage power for SRAM/DRAM and the increasing refresh dynamic power for DRAM have posed challenges for circuit/architecture designers for future memory hierarchy design.

Recently, emerging memory technologies (such as Spin Torque Transfer RAM(STT-MRAM), Phase-change RAM (PCRAM), and Resistive RAM (ReRAM)), are being explored as potential alternatives of existing memories in future computing systems. Such emerging non-volatile memory (NVM) technologies combine the speed of SRAM, the density of DRAM, and the non-volatility of Flash memory, and hence, become very attractive as the alternatives for future memory hierarchy. It is anticipated that these NVM technologies will break important ground and move closer to market very rapidly.

Simply using new technologies as replacements of existing hierarchy may not be the most desirable approach. For example, using high-density STT-RAM to replace SRAM as on-chip cache can reduce the cache miss rate due to larger capacity and improve performance, on the other hand, the longer write latency for STT-RAM could hurt the performance for write-intensive applications; Also, using high density memory as an extra level of on-chip cache will reduce CPU requests to the traditional, off-package DRAM and thus reduce the average memory access time. However, to

manage this large cache, a substantial amount of space on the CPU chip needs to be taken up by tags and logics, which could be used to increase the size of the next lower level cache. Moreover, trends toward Many-core and System-on-Chip may introduce the need and opportunity for new memory architectures. Consequently, as such emerging memory technologies are getting mature, it is important for SoC designers and computer architects to understand the benefits and limitations for better utilizing them to improve the performance/power/reliability of future computer architecture. Specifically, designers need to seek the answers to the following questions:

- How to model such emerging NVM technologies at the architectural level?
- What will be the impacts of such NVMs on the future memory hierarchy? What will be the novel architectures/applications?
- What are the limitations to overcome for such a new memory hierarchy?

This book includes 11 chapters that try to answer the questions mentioned above. These chapters cover different perspectives related to the modeling, design, and architectures of using the emerging memory technologies. We expect that this book can serve as a catalyst to accelerate the adoption of such emerging memory technologies for future computer system design from architecture and system design perspectives.

1.2 Preliminary on Emerging Memory Technologies

Many promising emerging memory technology candidates, such as Phase-Change RAM (PCRAM), Spin Torque Transfer Magnetic RAM (STT-RAM), Resistive RAM (ReRAM), and Memristor, have gained substantial attentions and are being actively pursued by industry [1]. In this section we will briefly describe the fundamentals of these promising emerging memory technologies to be surveyed in this paper, namely, the STT-RAM, the PCRAM, the ReRAM, and Memristor.

STT-RAM is a new type of Magnetic RAM (MRAM) [1], which features non-volatility, fast writing/reading speed (<10 ns), high programming endurance ($>10^{15}$ cycles) and zero standby power [1]. The storage capability or programmability of MRAM arises from magnetic tunneling junction (MTJ), in which a thin tunneling dielectric, e.g., MgO, is sandwiched by two ferromagnetic layers, as shown in Fig. 1.1. One ferromagnetic layer ("pinned layer") is designed to have its magnetization pinned, while the magnetization of the other layer ("free layer") can be flipped by a write event. An MTJ has a low (high) resistance if the magnetizations of the free layer and the pinned layer are parallel (anti-parallel). Prototyping STT-RAM chips have been demonstrated recently by various companies and research groups [2, 3]. Commercial MRAM products have been launched by companies like Everspin and NEC.

PCRAM technology is based on a chalcogenide alloy (typically, Ge_2–Sb_2–Te_5, GST) material) [1, 4]. The data storage capability is achieved from the resistance differences between an amorphous (high-resistance) and a crystalline (low-resistance)

phase of the chalcogenide-based material. In SET operation, the phase change material is crystallized by applying an electrical pulse that heats a significant portion of the cell above its crystallization temperature. In RESET operation, a larger electrical current is applied and then abruptly cut off in order to melt and then quench the material, leaving it in the amorphous state. PCRAM has shown to offer compatible integration with CMOS technology, fast speed, high endurance, and inherent scaling of the phase-change process at 22-nm technology node and beyond [5]. Compared to STT-RAM, PCRAM is even denser with an approximate cell area of $6 \sim 12F^2$ [1], where F is the feature size. In addition, phase change material has a key advantage of the excellent scalability within current CMOS fabrication methodology, with continuous density improvement. Many PCRAM prototypes have been demonstrated in the past years by companies like Hitachi [6], Samsung [7], STMicroelectronics [8], and Numonyx [9].

Resistive RAM (ReRAM) and Memristor

ReRAM memory stores the data as two (single-level cell, or SLC) or more resistance states (multi-level cell, or MLC) of the resistive switch device (RSD). Resistive switching in transition metal oxides was discovered in thin NiO film decades ago. From then, a large variety of metal-oxide materials have been verified to have resistive switching characteristics, including TiO_2, NiO_x, Cr-doped $SrTiO_3$, PCMO, CMO [10], etc. Based on the storage mechanisms, ReRAM materials can be cataloged as filament-based, interface-based, programmable-metallization-cell (PMC), etc. Based on the electrical property of resistive switching, RSDs can be divided into two categories: unipolar or bipolar. Programmable-metallization-cell (PMC) [11] is a promising bipolar switching technology. Its switching mechanism can be explained as forming or breaking the small metallic "nanowire" by moving the metal ions between two sold metal electrodes. Filament-based ReRAM is a typical example of unipolar switching [12] that has been widely investigated. The insulating material between two electrodes can be made conducting through a hopping or tunneling conduction path after the application of a sufficiently high voltage. The data storage could be achieved by breaking (RESET) or reconnecting (SET) the conducting path. Such switching mechanism can in fact be explained with the fourth circuit element, the **memristor** [13–15].

Memristor was predicted by Chua in 1971 [13], based on the completeness of circuit theory. Memristance (M) is a function of charge (q), which depends upon the historic behavior of the current (or voltage) profile [15, 16]. In 2008, the researchers at HP reported the first real device of a memristor in a solid-state thin film two-terminal device by moving the doping front along the device [14]. Afterwards, magnetic technology provides the other possible methods to build a memristive system [17, 18]. Due to its unique historic characteristic, memristor has very broad application including nonvolatile memory, signal processing, control and learning system etc [19].

Many companies are working on ReRAM technology and chip design, including Fujitsu, Sharp, HP lab, Unity Semiconductor Corp., Adesto Technology Inc. (a spin-off from AMD), etc. And in Europe, the research institute IMEC is doing independent research on ReRAMs with its partners Samsung Electronics Co. Ltd., Hynix

Table 1.1 Comparison of different memory technologies [21]

Features	SRAM	eDRAM	STT-RAM	PCRAM	ReRAM
Density	Low	High	High	Very high	Very high
Speed	Very Fast	Fast	Fast for read; slow for write	Slow for read; very slow for write	Slow for read/write
Dynamic Power	Low	Medium	Low for read; very high for write	Medium for read; high for write	Medium for read; high for write
Leakage Power	High	Medium	Low	Low	Low
Non-volatility	No	No	Yes	Yes	Yes

Semiconductor inc., Elpida Inc. and Micron Technology Inc. The main efforts on ReRAM research devote to material and devices [10]. Many circuit design issues have also been addressed, such as power-supply voltage and current monitoring. Recently, Sandisk and Toshiba demonstrated a 32 Gb ReRAM prototype in ISSCC 2013 [20].

Table 1.1 shows the comparison of these three emerging memory technologies against the conventional memory technologies used in traditional memory hierarchies.

1.3 Modeling

To help the architectural level and system-level design space exploration of the SRAM-based or DRAM-based cache and memory, various modeling tools have been developed during the last decade. For example, CACTI [22] and DRAMsim [23] have become widely used in the computer architecture community to estimate the speed, power, and area parameters of SRAM and DRAM caches and main memory.

Similarly, for computer architects to explore new design opportunities at architecture and system levels that the emerging memory technologies can provide, architectural level STT-RAM-based cache model [24, 25] and PCRAM-based cache/memory model [26] have been recently developed. Such architectural models provide the extraction of all important parameters, including access latency, dynamic access power, leakage power, die area, I/O bandwidth, etc., to facilitate architecture-level analysis, and to bridge the gap between the abundant research activities at process and device levels and the lack of a high-level cache and memory model for emerging NVMs.

The architectural modeling for cache and main memory built with emerging memory technologies (such as STT-RAM and PCRAM) raises many unique research issues and challenges.

- First, some circuitry modules in PCRAM/MRAM have different requirements from those originally designed for SRAM/DRAM. For example, the existing sense

amplifier model in CACTI [22] and DRAMsim [23] is based on voltage-mode sensing, while PCRAM data reading usually uses a current-mode sense amplifier.

- Second, due to the unique device mechanisms, the models of PCRAM/MRAM need specialized circuits to properly handle their operations. For example, in PCRAM, the specific pulse shapes are required to heat up GST material quickly and to cool it down gradually during the RESET and especially SET operations. Hence, a model of the slow quench pulse shaper need to be created.
- Finally, the memory cell structures between STT-RAM/PCRAM and SRAM/ DRAM are different. PCRAM and STT-RAM typically use a simple "1T1R" (one-transistor-one-resistor) or "1D1R" (one-diode-one-resistor) structure, while SRAM and DRAM cell has a conventional "6T" structure and "1T1C" (one-transistor-one-capacitor) structure, respectively. The difference of cell structures directly leads to different cell sizes and array structures.

In addition, where to place these emerging memories in the traditional memory hierarchy also influences the modeling methodologies. For example, the emerging NVMs could be used as a replacement for on-chip cache or for off-chip DIMM (dual in-line memory module). Obviously, the performance/power of on-chip cache and off-chip DIMM would be quite different: When a NVM is integrated with logics on the same die, there is no off-chip pin limitation so that the interface between NVM and logic can be re-designed to provide a much higher bandwidth. Furthermore, off-chip memory is not affected by the thermal profile of the microprocessor core while the on-chip cache is affected by the heat dissipation from the hot cores. While higher on-chip temperature has a negative impact on SRAM/DRAM memory, it may have a positive influence on PCRAM because the heat can facilitate the write operations of PCRAM cell. The performance estimation of PCRAM becomes much more complicated in such a case.

Moreover, building an accurate PCRAM/MRAM simulator needs close collaborations with the industry to understand physics and circuit details, as well as architectural level requirements such as the interface/interconnect with the multi-core CPUs.

Chapter 2 of this book introduces a modeling tool called NVsim by Dong et al. This tool is widely used by research community as an open-source modeling tool for emerging memory technologies such as STT-RAM and PCRAM.

1.4 Leveraging Emerging Memory Technologies in Architecture Design

As the emerging memory technologies are getting mature, integrating such memory technologies into the memory hierarchies (as shown in Fig. 1.1) provides new opportunities for future memory architecture designs. Specifically, there are several characteristics of STT-RAM and PCRAM that make them promising as working class memories (i.e., on-chip caches and off-chip main memories), or as storage class mem-

ories: (1) Compared to SRAM/DRAM, these emerging memories usually have much higher density, with comparable fast access time; (2) Due to the non-volatility feature, they have zero standby power, and immune to radiation-induced soft errors; (3) Compared to NAND-Flash SSD, STT-RAM/PCRAM are byte-addressable. In addition, different hybrid compositions of memory hierarchy by using SRAM, DRAM, and PCRAM or MRAM can be motivated by different power and access behaviors of various memory technologies. For example, leakage power is dominant in SRAM and DRAM arrays; on the contrary, due to non-volatility, PCRAM or STT-RAM array consumes zero leakage power when idling but a much higher energy during write operations. Hence, the trade-off among using different memory technologies at various hierarchy levels becomes an interesting research topic. In addition, if these memory are used as on-chip cache or main memory rather than as storage, the data retention time for non-volatility is not that important since data are used and overwritten in a very short period of time. Consequently, data retention time can be traded for better performance and energy benefits (as demonstrated by Chap. 7).

In this book, Chaps. 3–9 covers different design options of using such emerging memory technologies at different level of memory hierarchies. Chapter 10 proposes a design space exploration framework for circuit-architecture co-optimization for NVM memory architecture design. Chapter 11 describes a prototyping effort that fabricated an NVM-based processor design.

1.4.1 Leveraging NVMs as On-Chip Cache

Replacing SRAM-based on-chip cache with STT-RAM/PCRAM can potentially improve performance and reduce power consumption. With larger on-chip cache capacity (due to its higher density), STT-RAM/PCRAM based on-chip cache can help reduce the cache miss rate, which helps improve the performance. Zero-standby leakage can also help reduce the power consumption. On the other hand, longer write-latency of such NVM-based cache may incur performance degradation and offset the benefits from the reduced cache miss rate. Although PCRAM is much denser than SRAM, the limited endurance makes it unaffordable to directly use PCRAM as on-chip caches, which have highly frequent accesses.

The performance/power benefits of STT-RAM for single-core processor were investigated by Dong et al. [24]. The research demonstrated that STT-RAM-based L2 cache can bring performance improvement and achieve more than 70 % power consumption reduction at the same time. The benefits of using STT-RAM shared L2 cache for multi-core processors were demonstrated by Sun et al. [25]. The simulation result shows that the optimized MRAM L2 cache improves performance by 4.91 % and reduces power by 73.5 % compared to the conventional SRAM L2 cache with a similar area. Wu et al. [21] studied a number of different hybrid-cache architectures (HCA) that are composed of SRAM/eDRAM/STT-RAM/PCRAM for IBM Power 7 cache architecture, and explored the potential of hardware support for intra-cache data movement and power consumption management within HCA caches. Under the

same area constraint across a collection of 30 workloads, such aggressive hybrid-cache design provides 10–16 % performance improvement over the baseline design with a 3-level SRAM-only cache design, and achieves up to a 72 % reduction in power consumption.

In this book, Chaps. 6 and 7 give details on the evaluation of using NVM as on-chip cache, and the mitigation techniques to overcome some limitations such as performance/power overhead related to the write operations. device-architecture co-optimization can also be applied to achieve better performance/power benefits.

1.4.2 Leveraging NVMs as Main Memory

There are abundant recent investigations on using PCRAM as a replacement for the current DRAM-based main memory architecture. Lee et al. [27] demonstrated that a pure PCRAM-based main memory architecture implementation is about 1.6x slower and requires 2.2x energy than a DRAM-based main memory, mainly due to the overhead of write-operations. They proposed to re-design the PCM buffer organizations, with narrow buffers to mitigate high energy PCM writes. Also with multiple buffer rows, it can exploit locality to coalesce writes, hiding their latency and energy, such that the performance is only 1.2x slower with a similar energy consumption compared to the DRAM-based system. Qureshi et al. [28] proposed a main memory system consisting of PCM storage coupled with a small DRAM buffer, so that it can leverage the latency benefits of DRAM and the capacity benefits of PCM. Such memory architecture could reduce page faults by 5x and provide a speedup of 3x. A similar study conducted by Zhou et al. [29] demonstrated that the PCRAM-based main memory consumes only 65 % of the total energy of the DRAM main memory with the same capacity, and the energy-delay product is reduced by 60 %, with various techniques to mitigate the overhead of write-operations. All these work have demonstrated the feasibility of using PCRAM as main memory in the near future.

1.4.3 Leveraging NVM to Improve NAND-Flash SSD

NAND flash memory has been widely adopted by various applications such as laptops and mobile phones. In addition, because of its better performance compared to the traditional HDD, NAND flash memory has been proposed to be used as a cache in HDD, or even as the replacement of HDD in some applications. However, one well-known limitation of NAND flash memory is the "erase-before-write" requirement. It cannot update the data by directly overwriting it. Instead, a time-consuming erase operation must be performed before the overwriting. To make it even worse, the erase operation cannot be performed selectively on a particular data item or page but can only be done for a large block called the "erase unit". Since the size of an erase unit

(typically 128 K or 256 K Bytes) is much larger than that of a page (typically 512 $\sim$ 8 K Bytes), even a small update to a single page requires all the pages within the erase unit to be erased and written again.

Compared to NAND flash memory, PCRAM/STT-MRAM has advantages of random access and direct in-place updating. Consequently, Chap. 3 gives details on how to use a hybrid storage architecture to combine the advantages of NAND flash memory and PCRAM/MRAM. In such hybrid storage architecture, PCRAM is used as the log region for NAND-flash. Such hybrid architecture has the following advantages: (1) the ability of "in-place updating" can significantly improve the usage efficiency of log region by eliminating the out-of-date log data; (2) the fine-granularity access of PCRAM can greatly reduce the read traffic from SSD to main memory; (3) the energy consumption of the storage system is reduced as the overhead of writing and reading log data is decreased with the PCRAM log region; and (4) the lifetime of the NAND flash memory in the hybrid storage could be increased because the number of erase operations is reduced.

1.4.4 Enabling Fault-Tolerant Exascale Computing

Due to the continuously reduced feature size, supply voltage, and increased on-chip density, computer systems are projected to be more susceptible to hard errors and transient errors. Compared to SRAM/DRAM memory, PCRAM/STT-RAM memory has unique features such as non-volatility and resilience to soft errors. The application of such unique features could enable novel architecture design for applications that can address the reliability challenges for future Exascale scale computing.

For example, checkpointing/rollback scheme, where the processor takes frequent checkpoints at a certain time interval and stores them to hard disk, is one of the most common approaches to ensure the fault-tolerance of a computing system. In the current peta-scale massive parallel processing (MPP) systems, such traditional checkpointing to hard disk could incur a large performance overhead and is not a scalable solution for future Exascale computing. For example, Dong et al. [30] proposed three variants of PCRAM-based hybrid checkpointing schemes, which reduce the checkpoint overhead and offer a smooth transition from the conventional pure HDD checkpoint to the ideal 3D PCRAM mechanism. With a 3D PCRAM approach, multiple layers of PCRAM memory are stacked on top of DRAM memory, integrated with the emerging 3D integration technology. With a massive memory bandwidth provided by the through-silicon-via (TSVs) enabled by 3D integration, fast and high-bandwidth local checkpointing can be realized. The proposed pure 3D PCRAM-based mechanism can ultimately take checkpoints with overhead less than 4 % on a projected Exascale system.

1.5 Mitigation Techniques for STT-RAM/PCRAM Memory

The previous section presents the benefits of using these emerging memory technologies in computer system design. However, such benefits can only be achieved with mitigation techniques that can help address the inherited disadvantages that related to the write operations: (1) Because of the non-volatility feature, it usually takes much longer and more energy for write operations compared to read operations; (2) Some emerging memory technologies such as PCRAM has the wear-out problem (lifetime reliability), which is one of the major concerns of using it as working memory rather than storage class memory. Consequently, introducing these emerging memory technologies into current memory hierarchy design gives rise to new opportunities but also presents new challenges that need to be addressed. In this section, we review mitigation techniques that help address the disadvantages of such emerging technologies.

1.5.1 Techniques to Mitigate Latency/Energy Overheads of Write Operations

In order to use the emerging NVMs as cache and memory, several design issues need to be solved. The most important one is the performance and energy overheads in write operations. A NVM has a more stable mechanism for data keeping, compared to a volatile memory such as SRAM and DRAM. Accordingly, it needs to take a longer time and consume more energy to over-write the existing data. This is the intrinsic characteristic of NVMs. PCRAM and MRAM are not exceptional. If we directly replace SRAM caches with PCRAM/MRAM ones, the long latency and high energy consumption in write operations could offset the performance and power benefits, and even result in degradation when the cache write intensity is high. Therefore, it is imperative to study techniques to mitigate the overheads of write operations in NVMs.

- *Hybrid Cache/Memory Architecture.* To leverage the benefits of both the traditional SRAM/DRAM (such as fast write-operations) and the emerging NVMs (such as high density, low leakage, and resilient to soft error), a hybrid cache/memory architecture can be used, such as STT-RAM/SRAM hybrid on-chip cache, which is described in details in Chap. 6, or PCRAM/DRAM hybrid main memory [28]. In such hybrid architecture, instead of building a pure STT-RAM-based cache or a pure PCRAM-based main memory, we could replace a portion of MRAM or PCRAM cells with SRAM or DRAM elements, respectively. The main purpose is to keep most of write intensive data within SRAM/DRAM part, and hence, to reduce write operations in NVM parts. Therefore, the dynamic power consumption can be reduced and performance can be further improved. The major challenges to this architecture are how to physically arrange two different types of memories and how to migrate data in between.

- *Novel Buffer Architecture.* The write buffer design in modern processors works well for SRAM-based caches, which have approximately equivalent read and write speeds. However, the traditional write buffer design may not be suitable for NVM-based caches, which feature a large variation between read and write latencies. Chapter 6 will give details on how to design a novel write buffer architecture to mitigate the write-latency overhead. For example, in the scenario where a write operation is followed by several read operations, the ongoing write operation may block the upcoming read ones and cause performance degradation. The cache write buffer can be improved to prevent the critical read operations from being blocked by long write operations. For example, a higher priority can be assigned to read operations when competition happens between read and write. In an extreme condition when write-retirements are always stalled by read operations, write buffer could become full, which can also degrade cache performance. Hence, how to properly deal with read/write sequence and whether this mechanism could be dynamically controlled based on applications also need to be investigated. A similar write cancellation and write pausing techniques are also proposed in Ref. [31]. In addition, Lee et al. [27] also proposed to redesign the PCRAM buffer, using narrow buffers to help mitigate high energy PCM writes. Multiple buffer rows can exploit locality to coalesce writes, hiding their latency and energy.
- *Eliminating Redundant Bit-Writes.* In a conventional memory access, a write updates an entire row of the memory cells. A large portion of such writes are redundant. A read-before-write operation can help identify such redundant bits and cancel those redundant write operations to save energy and reduce the impact on performance [32].
- *Data Inverting.* To further reduce the number of writes to PCRAM cells, a data inverting scheme [32, 33] can be adopted in the PCRAM write logic. When a new data is written to a cache block, we first read its old data value, and compute the Hamming distance (HD) between the two values. If the calculated HD is larger than the half of the cache block size, the new data value is inverted before the store operation. An extra status bit is set to 1 to denote that the stored value has been inverted.

1.5.2 Techniques to Improve Lifetime for NVMs

Write endurance is another severe challenge in PCRAM memory design. The state-of-the-art process technology has demonstrated that the write endurance for PCRAM is around 10^8–10^9 [29]. The problem is further aggravated by the fact that writes to caches and main memory can be extremely skewed. Consequently, those cells suffering from more frequent write operations will fail much sooner than the rest.

Techniques that proposed in the previous sub-section to reduce the number of write operations to STT-RAM/PCRAM can definitely help the lifetime of the memory, besides reducing the write energy overhead. In addition to those techniques, the following schemes can be used to further improve the lifetime of the memory.

- *Wear leveling.* Wear leveling technique, which has been widely implemented in NAND flash memory, attempts to work around the limitations of write endurance by arranging data access so that write operations can be distributed evenly across all the storage cells. Wear leveling technique can also be applied to PCRAM/MRAM-based cache and memory. a range of wear leveling techniques for PCRAM have been examined [27–29, 32, 34] recently. Such wear leveling techniques include: (1) Row Shifting. A simple shifting scheme can be applied to evenly distribute writes within a row. The scheme is implemented through an additional row shifter along with a shift offset register. On a read access, data is shifted back before being passed to the processor. (2) Word-line Remapping and Bit-line Shifting. Bit-line shifter and word-line remapper are used to spread the writes over the memory cells inside one cache block and among cache blocks, respectively. (3) Segment Swapping. Periodically, memory segments of high and low write accesses are swapped. The memory controller keeps track of the write counts of each segment, and a mapping table between the "virtual" and "true" segment number. Chapter 9 of this book will also cover more details about the wear-leveling techniques, including new considerations of intra-set and inter-set write variations when NVM is used as on-chip cache.
- *Graceful degradation.* In such scheme, the PCRAM allows continued operation through graceful degradation when hard faults occur [35]. The memory pages that contain hard faults are not discarded. Instead, they are dynamically formed pairs of complementary pages that act as a single page of storage, such that the total effective memory capacity is reduced but the the lifetime of PCRAM can be improved by up to $40\times$ over conventional error-detection techniques.

1.6 Conclusion

This chapter reviews recent advance in memory architecture design with such emerging memory technology, discusses the benefits of using STT-RAM/PCRAM at various level of memory hierarchy, and also presents the mitigation techniques to overcome the challenges of applying such emerging memory technologies for future memory architecture design. These recent architectural-level studies have demonstrated that emerging memory technologies like STT-RAM/PCRAM/ReRAM have great potentials to improve future computer memory architecture design, and enable novel applications such as novel checkpointing techniques for future Exascale computing.

The rest of this book will give more details for different perspectives that are introduced in this chapter. With all these initial research efforts, we believe that the emerging of these new memory technologies will change the landscape of future memory architecture design.

References

1. International Technology Roadmap for Semiconductor, 2007.

2. Honjo, H., Saito, S., Ito, Y., Miura, S., Kato, Y., Mori, K., Ozaki, Y., Kobayashi, Y., Ohshima, N., Kinoshita, K., Suzuki, T., Nagahara, K., Ishiwata, N., Suemitsu, K., Fukami, S., Hada, H., Sugibayashi, T., Nebashi, R., Sakimura, n., & Kasai, N. (2009). A 90 nm 12 ns 32 Mb 2T1MTJ MRAM. *IEEE International Solid-State Circuits Conference (ISSCC)* (pp. 462–463).

3. Kawahara, T., Takemura, R., Miura, K., Hayakawa, J., Ikeda, S., Lee, Y. M., et al. (2008). 2 Mb SPRAM (SPin-transfer torque RAM) with bit-by-bit bi-directional current write and parallelizing-direction current read. *IEEE Journal of Solid-State Circuits*, *43*(1), 109–120.

4. Raoux, S., et al. (2008). Phase-change random access memory: A scalable technology. *IBM Journal of Research and Development*, *52*(4/5), 465–481.

5. Chen, Y.C., Rettner, C.T., Raoux, S., Burr, G.W., Chen, S.H., Shelby, R.M., Salinga, M., et al. (2006). Ultra-thin phase-change bridge memory device using gesb. *Proceedings of the IEEE International Electron Devices Meeting* (pp. 30.3.1–30.3.4).

6. Osada, K., Kotabe, A., Matsui, Y., Matsuzaki, N., Takaura, N., Moniwa, M., Kawahara, T., Hanzawa, S., & Kitai, N. (2007). A 512 kb embedded pram with 416 kbs write throughput at 100μa cell write current. *IEEE International Solid-State Circuits Conference (ISSCC)* (p. 26.2).

7. Cho, W.-Y., Kang, S., Choi, B.-G., Oh, H.-R., Lee, C.-S., Kim, H.-J., Park, J.-M., Wang, Q., Park, M.-H., Ro, Y.-H., Choi, J.-Y., Kim, K.-S., Kim, Y.-R., Chung, W.-R., Cho, H.-K., Lim, K.-W., Choi, C.-H., Shin, I.-C., Kim, D.-E., Yu, K.-S., Kwak, C.-K., Kim, C.-H., Lee, K.-J., & Cho, B. (2007) A 90 nm 1.8 V 512 Mb diode-switch pram with 266 Mb/s read throughput. *IEEE International Solid-State Circuits Conference (ISSCC)* (p. 26.1).

8. Pirola, A., Marmonier, L., Pasotti, M., Borghi, M., Mattavelli, P., Zuliani, P., Scotti, L., Mastracchio, G., Bedeschi, F., Gastaldi, R., Bez, R., De Sandre, G., & Bettini, L. (2010). A 90 nm 4 Mb embedded phase-change memory with 1.2 V 12 ns read access time and 1 Mb/s write throughput. *IEEE International Solid-State Circuits Conference (ISSCC)* (p. 14.7).

9. Barkley, G., Giduturi, H., Schippers, S., Vimercati, D., Villa, C., & Mills, D. (2010). A 45 nm 1 Gb 1.8 V phase-change memory. *IEEE International Solid-State Circuits Conference (ISSCC)* (p. 14.8).

10. Wong, P., Lee, H., Yu, S., et al. (1951). Metal-Oxide ReRAM. *Proceedings of the IEEE*, *100*(6), 2012.

11. Kozicki, M.N., Balakrishnan, M., Gopalan, C., Ratnakumar, C., & Mitkova, M. (2005) Programmable metallization cell memory based on ag-ge-s and cu-ge-s solid electrolytes. *Non-Volatile Memory Technology Symposium* (pp. 83–89).

12. Inoue, I.H., Yasuda, S., Akinaga, H., & Takagi, H. (2008). Nonpolar resistance switching of metal/binary-transition-metal oxides/metal sandwiches: Homogeneous/inhomogeneous transition of current distribution. *Physical Review B*, *77*, 035105.

13. Chua, L.O. (1971) Memristor—The missing circuit element. *IEEE Transactions on Circuit Theory*, *CT-18*, 507–519.

14. Tour, J. M., & He, T. (2008). The fourth element. *Nature*, *453*, 42–43.

15. Strukov, D.B., Snider, G.S., Stewart, D.R., & Williams, R.S. (2008) The missing memristor found. *Nature*, *453*, 80–83.

16. Chua, L. O. (1976). Memristive devices and systems. *Proceedings of IEEE*, *64*, 209–223.

17. Pershin, Y.V., & Di Ventra, M. (2008). Spin memristive systems: Spin memory effects in semiconductor spintronics. *Physical Review B: Condensed Matter*, *78*, 113309.

18. Wang, X., et al. (2009). Spin memristor through spin-torque-induced magnetization motion. *IEEE Electron Device Letters*, *30*, 294–297.

19. Chen, y., & Wang, X. (2009). Compact modeling and corner analysis of spintronic memristor. *IEEE/ACM International Symposium on Nanoscale Architectures 2009 (NANOARCH'09)* (pp. 7–12).

20. Liu, T., Yan, T., Scheuerlein, R., et al. (2013). A 130.7 mm^2 2-Layer 32 Gb ReRAM Memory Device in 24 nm Technology. *Proceedings of International Solid State Circuits Conference* (pp. 210–211).
21. Wu, X., Li, J., Zhang, L., Speight, E., Rajamony, R., & Xie, Y. (2009). Hybrid cache architecture with disparate memory technologies. *36th International Symposium on Computer Architecture, ISCA '09*.
22. Wilton, S. J. E., & Jouppi, N. P. (1996). CACTI: An enhanced cache access and cycle time model. *IEEE Journal of Solid-State Circuits*, *31*, 677–688.
23. Wang, David, Ganesh, Brinda, Tuaycharoen, Nuengwong, Baynes, Kathleen, Jaleel, Aamer, & Jacob, Bruce. (2005). DRAMsim: a memory system simulator. *SIGARCH Computer Architecture News*, *33*(4), 100–107.
24. Dong, X. , Wu, X. , Sun, G., Xie, Y., Li, H., & Chen, H. (2008). Circuit and microarchitecture evaluation of 3D stacking magnetic RAM (MRAM) as a universal memory replacement. *Proceedings of Conference on Design Automation* (pp. 554–559).
25. Sun, G., Dong, X., Xie, Y., Li, J., & Chen, Y. (2009). A novel architecture of the 3d stacked mram l2 cache for cmps. *IEEE 15th International Symposium on High Performance Computer Architecture* (pp. 239–249).
26. Dong, X., Jouppi, N., & Xie, Y. (2009) Pcramsim: System-level performance, energy, and area modeling for phase-change ram. *International Conference on Computer-Aided Design (ICCAD)* (pp. 269–275).
27. Lee, B.C., Ipek, E., Mutlu, D., & Burger, D. (2009). Architecting phase change memory as a scalable dram alternative. *Proceedings of ISCA* (pp. 2–13).
28. Qureshi, M.K., Srinivasan, V., & Rivers, J.A. (2009) Scalable high performance main memory system using phase-change memory technology. *ISCA '09: Proceedings of the 36th annual international symposium on Computer architecture*, New York, NY, USA, 2009. ACM (pp. 24–33).
29. Zhou, P., Zhao, B., Yang, J., & Zhang, Y. (2009) A durable and energy efficient main memory using phase change memory technology. *Proceedings of ISCA* (pp. 14–23).
30. Dong, X., Muralimanohar, N., Jouppi, N., Kaufmann, R., & Xie, Y. (2009) Leveraging 3d pcram technologies to reduce checkpoint overhead for future exascale systems. *International Conference on High Performance Computing, Networking, Storage and, Analysis (SC09)*.
31. Qureshi, M., Franceschini, M., Lastras, M. (2010). Improving read performance of phase change memories via write cancellation and write pausing. *Proceedings of International Symposium on High Performance Computer Architecture (HPCA)*.
32. Joo, Y., Niu, D., Dong, X., Sun, G., Chang, N., & Xie, Y. (2010). Energy- and endurance-aware design of phase change memory caches. *Proceedings of Design Automation and Test in Europe*.
33. Cho, S., & Lee, H. (2009). Flip-N-Write: A simple deterministic technique to improve PRAM write performance, energy and endurance. *Proceedings of International Symoposium on Microarchitecture (MICRO)*.
34. Qureshi, M., Karidis, J., Franceschini, M., Srinivasan, V., Lastras, L., & Abali, B. (2009). Enhancing lifetime and security of phase change memories via start-gap wear leveling. *Proceedings of International Symoposium on Microarchitecture (MICRO)*.
35. Ipek, E., Condit, J., Nightingale, E., Burger, D., & Moscibroda, T. (2010). Dynamically replicated memory: Building reliable systems from nanoscale resistive memories. *Proceedings of International Conference on Architecture Support for Programming Languages and Operating Systems*.