

A Preliminary Study on System-level Impact of Persistent Main Memory

Taciano Perez¹, Ney Laert Vilar Calazans², César A. F. De Rose²

¹Hewlett-Packard R&D, Porto Alegre, Brazil

²PUCRS, Faculty of Informatics, Porto Alegre, Brazil

taciano.perez@hp.com, {ney.calazans, cesar.derose}@pucrs.br

Abstract—For almost 30 years, computer memory systems have been essentially the same: volatile, high speed memory technologies like SRAM and DRAM used for cache and main memory; magnetic disks for high-end data storage; and persistent, low speed flash memory for storage with low capacity/low energy consumption requirements such as embedded/mobile devices. Today we watch the emergence of new non-volatile memory (NVM) technologies that promise to radically change the landscape of memory systems. This work presents system-level latency and energy impacts of a computer architecture with persistent main memory using PCRAM and Memristor. Our experimental results support the feasibility of employing emerging non-volatile memory technologies as persistent main memory, indicating that performance penalties should be mild, and energy improvements should be significant, up to 45.5% less when using PCRAM and 72.4% less when using Memristor.

Keywords— Non-Volatile Memory, Persistent Main Memory, PCRAM, RRAM, Memristor

I. Introduction

For almost 30 years the memory hierarchy used in computer design has been essentially the same: volatile, high speed memory technologies like SRAM and DRAM used for cache and main memory; magnetic disks for high-end data storage; and persistent, low speed flash memory for storage with low capacity/low energy consumption requirements, such as embedded/mobile devices [1].

Today we watch the emergence of new memory technologies that promise to significantly change the landscape of memory systems. Non-Volatile Memory (NVM) technologies such as Phase-Change RAM (PCRAM), Magnetic RAM (MRAM) and Memristor will possibly enable memory chips that are non-volatile, require low energy and have density and latency closer to current DRAM chips [2]. The creation of byte-addressable, non-volatile solid state memory could make a significant amount of persistent main memory available to computer systems, allowing for consolidating these two different levels of the storage hierarchy — main memory and persistent storage — into a single level, something that has never been possible before (at least not in mass scale).

The advent of main memory as the primary persistent storage can deeply affect the complete computing stack, including application software, operating system, busses, memory system and their interaction with other devices, such as processors and I/O adapters [3]. In order to fully as-

sess system-wide impacts in latency, energy, heat, space and cost, it is required to take into account different layers when modeling or simulating a computer system with persistent main memory.

This work presents an evaluation of application workloads in a hypothetical computer with persistent main memory, through the use of experimental models and simulations, aiming to identify the major system-level impacts of persistent main memory in latency and energy.

II. Limitations of Current Memory Technology

The memory subsystem has become one of the most important topics in computer design, being one of the main factors impacting system-level performance [1]. The high level of sophistication attained by modern memory systems is largely derived from the process predicted by Gordon Moore in 1965 [4], known as Moore's Law, which states that the number of devices that can be integrated on a chip of fixed area would double every 12 months (later amended to doubling every 18–24 months). This behavior has made the prediction of near-future product developments extremely reliable because the underlying device physics, materials, and fabrication processes have all been scalable, at least until now [2].

An important issue for the near future concerns DRAM approaching physical limits that might restrict its growth in the next decade, creating a “power wall”. DRAM must not only place charge in a storage capacitor but must also mitigate sub-threshold charge leakage through the access device. Thus capacitors must be large enough to store charge for reliable sensing and transistors must be large enough to exert effective control over the channel. Given this context, predictions state that DRAM will face scaling limitations as feature size continues decreasing [5]. DRAM is also increasingly affecting the energy footprint of computer systems, being responsible for as much as 40% of the system energy budget in certain cases [6].

Due to these limitations, there is extensive research to create new alternatives of memory technology that can address these problems and prevent a “power wall” from being reached.

III. Emerging Memory Technologies

There are several new Non-Volatile Memory (NVM) technologies under research today. One study [7] lists 13 such technologies: FRAM, MRAM, STTRAM, PCRAM, NRAM, RRAM, CBRAM, SEM, Polymer, Molecular,

Racetrack, Holographic and Probe, in different stages of maturity. This study is limited to two of these technologies: Phase-Change RAM and Memristor (a type of RRAM), since they are among the most mature technologies being considered as possible replacements for DRAM as main memory.

A. Phase-Change RAM (PCRAM)

Phase-Change Random Access Memory (also called PCRAM, PRAM or PCM) is currently the most mature of the new memory technologies under research. It relies on some materials, called phase-change materials, that exist in two different phases with distinct properties: an amorphous phase, characterized by high electrical resistivity, and a crystalline phase, characterized by low electrical resistivity [8]. These two phases can be repeatedly and rapidly cycled by applying heat to the material [8], [2].

The principle of phase-change memory is known since the 1960s, but only recent discoveries of phase-change materials with faster crystallization speeds led to the possibility of commercial feasibility. The most important materials are chalcogenides such as $Ge_2Sb_2Te_5$ (GST), that can crystallize in less than 100 ns [8]. In a memory cell, the SET operation is achieved by crystallizing the material and RESET by making it amorphous [2], [5].

PCRAM has been demonstrated to work in 20nm device prototype and is projected to scale down to 9nm [8], [5]. The SET latency is the longest and determines the write performance. Latencies of 150 ns for SET and 40 ns for RESET operations have been demonstrated. Write energy is determined by the RESET operation, which dissipates 480 μW , while SET dissipates 90 μW . The read latency is 48 ns and dissipates 40 μW . Both read and write latencies are several times slower than DRAM, although only by tens of nanoseconds [5].

Endurance is bound to the number of writes. This happens because when current is injected into a phase-change material, thermal expansion and contraction degrade the contacts, so that currents can no longer be reliably injected into the cell. The current write endurance varies between 10^4 and 10^9 writes, but it can be conservatively assumed 10^8 as a reasonable reference value [5], [9].

B. Memristor

Despite the fact that PCRAM also uses resistance variations to store bit values, the term **Resistive RAM** (RRAM or ReRAM) has been applied to a distinct set of technologies that explore the same phenomenon. There is a long list of RRAM technologies [2], [7], [10]. This work concentrates on the **memristor** [10], the most mature example of RRAM being proposed as replacement for DRAM as main memory.

Since the 19th century, three fundamental passive circuit elements were known: the resistor, the inductor and the capacitor. In 1971, Leon Chua theorized the existence of

a fourth passive circuit element, which he called the memristor [11], but no actual physical device with memristive properties could be constructed. In 2008, a group of scientists reported the invention of a device that behaved as predicted for a memristor [12]. Later the same year, an article detailed how that device could be used to create nonvolatile memories [10].

A memristor is a two-terminal device whose resistance depends on the magnitude and polarity of the voltage applied to it and the length of time that voltage has been applied. When the voltage is turned off, the memristor remembers its most recent resistance until the next time it is turned on. The property of storing resistance values means that a memristor can be used as a nonvolatile memory [13].

This first memristor device created consisted of a crossbar of platinum wires with titanium dioxide (TiO_2) switches. Each switch consists of a lower layer of stoichiometric titanium dioxide (TiO_2), which is electrically insulating, and an upper layer of oxygen-deficient titanium dioxide (TiO_{2-x}), which is conductive. The size of each layer can be changed by applying voltage to the top electrode. If a positive voltage is applied, the TiO_{2-x} layer thickness increases and the switch becomes conductive (ON state). A negative voltage has the opposite effect (OFF state) [12], [10], [13]. Several oxides other than TiO_2 are known to present similar bipolar resistive switching, and there are multiple research projects in motion to explore these other materials for similar memory device implementations [2].

A cell size of 10nm has been achieved and a size of 4–5nm is predicted for the next few years [14], [13]. Access speeds both for reads and writes are expected to be in the same order of magnitude as DRAM speeds, within about a factor of two, likely around 100ns [15], [3]. Initial memristor prototypes demonstrated limited endurance, in the order of 10^7 write cycles [3], but is expected to improve as research progresses [13].

IV. Persistent Main Memory

The creation of a byte-addressable, non-volatile solid state memory technology could make a significant amount of persistent main memory available to computer systems, allowing a collapse of two different levels of the storage hierarchy — main memory and persistent storage — into a single level, something that has never been possible before. The advent of main memory as the primary persistent storage may deeply affect most computing layers, including application software, operating system, busses, memory system and their interaction with other devices, such as processors and I/O adapters [16], [17]. In order to fully assess system-wide impacts in latency, energy, heat, space and cost, it is required to take into account all these different layers when modeling or simulating a hypothetical computer system with persistent main memory.

Such a system, with persistent storage collapsed into main memory, would have impacts in several system aspects:

- **Timing:** the latency for reading/writing NVM cells

tends to be higher than DRAM. On the other hand, without the need to access high-latency block devices, the overall performance should be dramatically increased for most applications.

- **Energy:** DRAM today accounts for about 30% of the energy consumption of a server, and disk is responsible for another 10% [18]. These components consume both dynamic power (used to change data) and static power (used for data retention and component availability). Static power used by DRAM is consumed mostly for memory refreshes, and can account for more than half of the total power consumption [9]; in disks, most static power is used to keep the disk spinning and ready to answer requests with low latencies. Static power consumption has also the disadvantages of increasing with the memory size and not being energy-proportional, i.e., it is constant over time instead of proportional to the workload [18]. Typically, NVM technologies have higher dynamic power consumption than DRAM, but negligible static power consumption. Some studies project that NVM can increase overall memory energy efficiency up to 65% [19]. The precise order of magnitude of the energy savings obtained through the removal of disks and replacement of DRAM by NVM is not yet known, but based on these facts, it is expected to be considerable.

- **Heat:** heat and energy are two intertwined subjects. Heat emission tends to be proportional to energy dissipation. Today, air conditioning is the dominant solution to cool datacenter servers. Air conditioning, in turn, is another major power consumer. Its components, including CRACs (Computer Room Air Conditioning), chillers and humidifiers, account for about 45% of the total datacenter power consumption [18]. Roughly speaking, in a datacenter today for each dollar spent to power a server another dollar is spent to cool it. Since energy and heat are so interconnected, thermal efficiency should also be very positively affected by NVM adoption.

- **Space:** space is one of the less studied aspects related with the adoption of NVM. The volume occupied by the internal components of servers is likely to be significantly impacted. Magnetic disks have higher areal density (bits/square inch) than solid state memory, but require several components such as spindle, head, actuator and motor that make disk drives bulky units. The volume taken by storage media is relevant both for mobile devices and datacenters with high server density. The authors are not aware of any existing study estimating the impact of NVM on space.

- **Cost:** cost of NVM-based systems is another aspect not studied so far. Specifically for datacenters, Barroso and Holzle [18] proposed a cost model that considers the Total Cost of Ownership (TCO) of the datacenter as being composed of datacenter depreciation, datacenter operational expenses, server depreciation and server operational expenses. Their study demonstrates that the electricity bill is one of the major costs of a datacenter operation, and the superior energy efficiency of NVM should have a positive effect on it. On the other hand, the cost/bit of solid-state

memory today is much higher than disks: the cost of 1 MB of DRAM is about \$0.05, while 1 MB of disk costs about \$0.0003 [20]. There are not yet available market prices for NVM, but it is expected that their cost/bit will also be higher than disk, possibly by a significant factor. This is another area pending study.

Looking at all discussed aspects, it becomes clear that the impact of NVM-based persistent memory cannot be adequately evaluated at component or subsystem level, but requires a holistic assessment of computing systems. The next section explores different system architectures that can be considered.

The first proposals for the application of NVM technologies evaluated the individual replacement of existing memory hierarchy levels (such as processor cache, main memory and persistent storage) by NVM counterparts, with gains of performance and efficiency at subsystem level [19]. Most of these proposals do not imply a radical redesign of computing systems as a whole, but localized changes to specific subsystems. The advantage of this approach is avoiding disruption on existing standards, with the limitation of not exploiting the full potential of these technologies.

More recently, proposals for more radical system redesigns started being published. A good example is the architecture of Nanostores [17], that proposes parallel systems with a massive number of low-cost processors co-located with non-volatile data stores. This system is targeted for data-centric workloads, such as search, sort and video transcoding. Systems such as these have the potential for more sophisticated uses of NVM technology, but depart more radically from existing standards, which probably will result in later adoption.

In the present study, it is used a conservative approach of a simple commodity system where DRAM is fully replaced by non-volatile memory, either Memristor or PCRAM. No other significant changes will be applied. The system aspects being analyzed are timing and energy.

V. Persistent Main Memory

Our target configuration is a single-processor Pentium IV computer with 256 MB of main memory, simulated using Virtutech Simics [21]. The main aspects of the experimental setup are described in Table I (the latency values were derived from [5], [22], [15]). This setup enables us to exercise variations in the memory technology parameters in a very simple commodity system. We believe that future systems with persistent main memory will have different characteristics, such as a much larger memory size (comparable to current hard disks) and different physical memory organization, since JEDEC's DDRx does not support hundreds of Gigabytes. Our purpose is not to propose a radical depart from current architectures, but to have a first-order approximation of the impact of persistent main memory in current designs. In addition, a system without radical changes will probably be a necessary step in the evolution of NVM adoption. Such a simple system allows

TABLE I
EXPERIMENTAL TARGET CONFIGURATION SETUP.

Processor	Pentium IV 100 MHz (single-core)
L1 Cache	Size: 16 Kb (D-cache) + 16 Kb (I-cache) Associativity: 4-way (D-cache), 2-way (I-cache) Penalty: picoseconds Replacement policy: LRU
L2 Cache	Size: 512 Kb Associativity: 8x Penalty: 10 ns Replacement policy: LRU
Main Memory	Size: 256 MB Penalty: technology-dependent
Disk	20 GB
OS	Fedora Core release 5 (Bordeaux) Kernel: 2.6.15-1.2054_FC5

us to understand the changes with a single element variation: the memory technology used in the main memory. In future works, we intend to simulate machines with much larger amounts of main memory.

A set of workloads representing short common tasks was exercised in three different technology scenarios: DRAM, Memristor and PCRAM. The main memory latency and energy model was changed in order to match these technologies, and all other simulation aspects remained unchanged between the different scenarios. The technology parameters for each scenario are depicted on Table II.

The workloads consisted in:

- **gcc** - compilation of a small application.
- **gzip** - compression of a text file using GNU/Linux gzip application.
- **sort** - sorting of the contents of a text file using GNU/Linux sort application.

In order to estimate the overall energy consumption of each memory technology scenarios, it was used an energy model that considers two separate elements:

1. **Dynamic energy** - the energy consumed in order to read/write memory addresses.
2. **Refresh energy** - the energy consumed to keep the memory contents alive. It is only relevant for DRAM, since PCRAM and Memristor don't need refresh due to their own persistent nature.

Subthreshold power leakage is a third important component of the total energy consumption. NVM should have leakage at least similar to DRAM, or probably better, since idle memory banks can be turned off without losing its contents. This study haven't considered leakage due to the lack of published information on the leakage of memristor memory devices at this moment.

Table II displays the input parameters that were employed for each technology. These parameters were obtained us-

ing CACTI [23] (for DRAM) and NVSim [24] (for PCRAM and Memristor).

Every scenario was executed three times, and the average results are presented in the next section.

TABLE II
TECHNOLOGY PARAMETERS FOR DRAM, MEMRISTOR AND PCRAM.

	DRAM	Memristor	PCRAM
Read Latency (ns)	50	100	50
Write Latency (ns)	48	100	150
Read Energy (nJ)	3.4539	0.5729	27.1833
Write Energy (nJ)	3.4475	1.3475	1.3501
Refresh Power (mW)	0.0867	0	0
Feature Size (nm)	45	45	45

VI. Results and Analysis

TABLE III
EXECUTION TIMES FOR EXPERIMENTAL WORKLOADS (CPU TIME), IN SECONDS. TIME DEVIATION OF MEMRISTOR AND PCRAM RELATIVE TO DRAM ARE SHOWN AS PERCENTAGE (%).

	DRAM	Memristor	PCRAM
gcc	6.732	6.740 (0.12%)	6.780 (0.71%)
gzip	9.293	9.325 (0.34%)	9.409 (1.25%)
sort	30.998	31.262 (0.85%)	31.242 (0.79%)

The execution times for each scenario are shown in Table III. We notice a mild degradation of performance in Memristor and PCRAM compared to DRAM. Although Memristor is considered here 2x slower than DRAM and PCRAM write latency 3x slower than DRAM, the overall system performance impact is below 1%. This is consistent with the results published in similar studies [25], [6], [5], [22], [9]. The main factor behind this phenomenon is the high rate of L1 and L2 cache hits in a typical workload, which in our experiments were consistently above 90% for L1 Read and 84% for L2 Read.

A consequence of these observations is that main memory using Memristor or PCRAM cells still need processor caches in order to avoid severe performance penalties. Design proposals for persistent caches are explored in [26], [27], [28], [29].

The main impact of non-volatile memories is on the energy footprint of main memory, which is significantly smaller in non-volatile technologies than DRAM. In Figure 1 we can see that PCRAM consumed 10% more dynamic energy than DRAM, and Memristor consumed 45% less dynamic energy than DRAM. But the more dramatic difference comes from the static consumption made by DRAM for refreshing its contents, which are not necessary in Memristor and PCRAM due to their non-volatile nature. In the workload with larger duration (sort) the energy spent

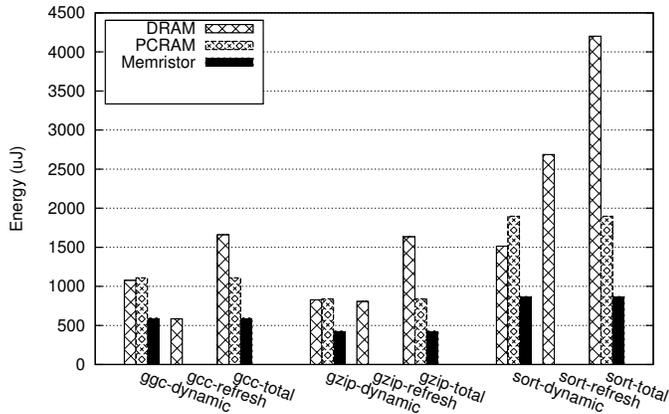


Fig. 1. Energy comparison by workload/technology, in μJ .

on refreshes surpassed the dynamic consumption of energy. Considering the total energy budget, in our experiments PCRAM consumed in average 45.5% less energy than DRAM and Memristor 72.4% less energy than DRAM, for the same workloads. The detailed energy consumption results for each workload/technology can be seen on Table IV.

TABLE IV
ENERGY CONSUMPTION OF EXPERIMENTAL WORKLOADS, IN μJ .

		DRAM	Memristor	PCRAM
gcc	Dynamic	1076.73	595.58	1109.77
	Refresh	583.66	0.00	0.00
	Total	1660.39	595.58	1109.77
gzip	Dynamic	827.57	428.54	841.88
	Refresh	805.70	0.00	0.00
	Total	1633.27	428.54	841.88
sort	Dynamic	1514.05	870.22	1897.64
	Refresh	2687.53	0.00	0.00
	Total	4201.58	870.22	1897.64

These results support the feasibility of employing emerging non-volatile memory technologies to craft persistent main memory. Many architectural changes are required in order to replace a three-level storage hierarchy (caches, main memory, disk) by a two-level hierarchy (caches, main memory), including memory devices, memory subsystem organization and operating system support, but our preliminary study, considering only changes in the memory technology, indicates that performance penalties should be mild and energy improvements should be significant. The next step proposed in our study is to make changes at the operating system and application software level in order to avoid access to high-latency storage devices such as magnetic disks and flash memory. All data would be stored in main memory, what is likely to yield expressive performance improvements at system level.

Comparing Memristor and PCRAM, both are feasible to

be used. Memristor has performance advantages over PCRAM due to more uniform read/write access latencies, and benefits from a considerably lower dynamic energy.

VII. RELATED RESEARCH

Other studies have evaluated the usage of emerging non-volatile memory technologies as main memory in commodity systems, using PCRAM or hybrid DRAM+PCRAM designs.

Lee et al. [5] propose memory subsystem enhancements in order to craft PCRAM main memory overcoming the latency and endurance limitations of the technology. These enhancements include buffer reorganization and partial writes (writing only modified data). The work of Zhou et al. [9] tackles similar problems proposing a 3D die stacked chip multiprocessor that puts together several processors and main memory on the same chip. They also propose removing redundant bit-writes, wear leveling through row shifting and segment swapping. The authors of both these papers later made a proposal that combines their approaches [22]. These works focus on memory subsystem enhancements in order to improve endurance and latency of PCRAM main memory. Our work considers Memristor as an alternative technology together with PCRAM, and focuses on system-level issues, building upon memory subsystem designs such as those proposed in the referenced studies.

Qureshi et al. [6] described a main memory system using a combination of PCRAM and DRAM. As a way to address PCRAM low endurance and relatively slow access, they propose a hybrid architecture where a DRAM buffer is placed in the front of the main PCRAM storage. In order to exploit this architecture, they propose mechanisms of lazy-write organization, line-level writes, fine-grained wear-leveling and page-level bypass for write filtering. Mogul et al. [30] proposed a hybrid design for main memory combining non-volatile memory with DRAM called FLAM, with two variants: flash+DRAM and PCRAM+DRAM. The motivation is a main memory that has less cost per bit, lower power consumption and higher density. According to their proposal, the CPU would not be able to directly access the non-volatile memory area, and a migration mechanism (triggered by the operating system) allows for data written in the DRAM buffer region to be migrated to the non-volatile region. Condit et al. [31] proposed a file system (called BPFs) designed for byte-addressable, persistent RAM (BPRAM), a term they use to designate NVM technologies such as PCRAM. They argue that with such technologies the file system should use direct-mapped memory on the memory bus instead of block devices through the I/O subsystem. They assume that PCRAM is presented to the system as DDR-compatible DIMMs available through the memory bus. A DRAM buffer is present, and both PCRAM and DRAM addresses are directly mapped by the CPU for both reading and writing. It is also assumed that the PCRAM memory controller

implements wear leveling and write failure detection mechanisms, such as the ones described in [6], [9]. Dhiman et al. [25] propose an hybrid PCRAM/DRAM main memory system called PDRAM, consisting of hardware enhancements to the memory controller in order to manager access to PCRAM pages, and software enhancements to the operating system page manager to perform wear leveling by page swapping/migration. These works concentrate on hybrid PCRAM+DRAM designs in order to improve PCRAM latency and endurance as main memory. Our work considers uniform main memory using a single technology (either Memristor or PCRAM) from a system-level perspective, and does not explore detailed memory subsystem improvements.

A system architecture called Nanostores [15], [3] has been proposed, describing parallel systems with a massive number of low-cost processors co-located with non-volatile data stores. This system is targeted for data-centric workloads, such as search, sort and video transcoding, and represents a more significant depart from current designs.

To the best of our knowledge, no previous work compared, from a system-level perspective, both PCRAM and Memristor as uniform replacements for DRAM as main memory technology in a simple commodity computer system.

VIII. CONCLUSION

This work presented a preliminary evaluation of workloads in a hypothetical computer with persistent main memory, through the use of experimental models and simulations, aiming to identify the major system-level impacts of persistent main memory in latency and energy.

It was observed that main memory using Memristor or PCRAM cells still need processor caches in order to avoid severe performance penalties. Overall performance of non-volatile technologies in main memory is similar to current DRAM results.

The main impact of non-volatile memories is on the energy footprint of main memory, which is significantly smaller in non-volatile technologies than DRAM. In our experiments, PCRAM consumed on average 45.5% less energy than DRAM and Memristor 72.4% less energy than DRAM, for the same workloads.

The great challenge of using such technologies as main memory is their low endurance. Wear leveling techniques such as those described in [5], [22], [6], [9] are expected to contribute positively.

The experimental results support the feasibility of employing emerging non-volatile memory technologies as persistent main memory. Many architectural changes are required in order to replace a three-level storage hierarchy (caches, main memory, disk) by a two-level hierarchy (caches, main memory), including memory devices, memory subsystem organization and operating system support, but our preliminary study, considering only changes in the memory technology, indicates that performance penalties should be mild and energy improvements should be signif-

icant. As a future step, we propose to make changes at operating system and application software level in order to avoid access to high-latency storage devices such as magnetic disks and flash memory and store all data in main memory, what is likely to yield expressive performance improvements at system level.

ACKNOWLEDGEMENTS

The authors would like to thank Greg Astfalk for reviewing an initial version of this article and providing valuable comments and feedback.

REFERENCES

- [1] B. Jacob, S.W. Ng, and D.T. Wang, *Memory systems: cache, DRAM, disk*, Morgan Kaufmann Pub, 2007.
- [2] G.W. Burr, B.N. Kurdi, J.C. Scott, C.H. Lam, K. Gopalakrishnan, and R.S. Shenoy, "Overview of candidate device technologies for storage-class memory," *IBM Journal of Research and Development*, vol. 52, no. 4, pp. 449–464, 2008.
- [3] D.A. Roberts, *Efficient Data Center Architectures Using Non-Volatile Memory and Reliability Techniques*, Ph.D. thesis, The University of Michigan, 2011.
- [4] G.E. Moore, "Cramming more components onto integrated circuits," *Electronics*, vol. 38, no. 8, pp. 114–117, 1965.
- [5] B.C. Lee, E. Ipek, O. Mutlu, and D. Burger, "Architecting phase change memory as a scalable dram alternative," *ACM SIGARCH Computer Architecture News*, vol. 37, no. 3, pp. 2–13, 2009.
- [6] M.K. Qureshi, V. Srinivasan, and J.A. Rivers, "Scalable high performance main memory system using phase-change memory technology," in *Proceedings of the 36th annual international symposium on Computer architecture*. ACM, 2009, pp. 24–33.
- [7] M.H. Kryder and C.S. Kim, "After Hard Drives - What Comes Next?," *Magnetics, IEEE Transactions on*, vol. 45, no. 10, pp. 3406–3413, 2009.
- [8] S. Raoux, GW Burr, MJ Breitwisch, CT Rettner, Y.C. Chen, RM Shelby, M. Salinga, D. Krebs, S.H. Chen, H.L. Lung, et al., "Phase-change random access memory: A scalable technology," *IBM Journal of Research and Development*, vol. 52, no. 4.5, pp. 465–479, 2010.
- [9] P. Zhou, B. Zhao, J. Yang, and Y. Zhang, "A durable and energy efficient main memory using phase change memory technology," *ACM SIGARCH Computer Architecture News*, vol. 37, no. 3, pp. 14–23, 2009.
- [10] J.J. Yang, M.D. Pickett, X. Li, D.A.A. Ohlberg, D.R. Stewart, and R.S. Williams, "Memristive switching mechanism for metal/oxide/metal nanodevices," *Nature nanotechnology*, vol. 3, no. 7, pp. 429–433, 2008.
- [11] L. Chua, "Memristor-the missing circuit element," *IEEE Transactions on Circuit Theory*, vol. 18, no. 5, pp. 507–519, 1971.
- [12] D.B. Strukov, G.S. Snider, D.R. Stewart, and R.S. Williams, "The missing memristor found," *Nature*, vol. 453, no. 7191, pp. 80–83, 2008.
- [13] R.S. Williams, "How we found the missing memristor," *IEEE spectrum*, vol. 45, no. 12, pp. 28–35, 2008.
- [14] D.L. Lewis and H.H.S. Lee, "Architectural evaluation of 3D stacked RRAM caches," in *IEEE International Conference on 3D System Integration, 2009. 3DIC 2009*. IEEE, 2009, pp. 1–4.
- [15] P. Ranganathan, "From Microprocessors to Nanostores: Rethinking Data-Centric Systems," *IEEE Computer*, vol. 44, no. 1, pp. 39–48, 2011.
- [16] H. Volos, A. J. Tack, and M. M. Swift, "Mnemosyne: Lightweight Persistent Memory," In 9th USENIX Symposium on Operating Systems Design and Implementation (OSDI '10), Poster, October 2010.
- [17] D. Roberts, J. Chang, P. Ranganathan, and T.N. Mudge, "Is Storage Hierarchy Dead? Co-located Compute-Storage NVRAM-based Architectures for Data-Centric Workloads," Tech. Rep., HP Labs, 2010.

- [18] L.A. Barroso and U. Holzle, *The datacenter as a computer: An introduction to the design of warehouse-scale machines*, vol. 4, Morgan & Claypool Publishers, 2009.
- [19] T. Perez and C.A.F. De Rose, "Non-Volatile Memory: Emerging Technologies And Their Impacts on Memory Systems (TR-060)," Tech. Rep., Faculdade de Informática, Pontificia Universidade Católica do Rio Grande do Sul (PUCRS), 2010.
- [20] G. Graefe, "The five-minute rule twenty years later, and how flash memory changes the rules," in *Proceedings of the 3rd international workshop on Data management on new hardware*. ACM, 2007, pp. 1–9.
- [21] P.S. Magnusson, M. Christensson, J. Eskilson, D. Forsgren, G. Haallberg, J. Hogberg, F. Larsson, A. Moestedt, and B. Werner, "Simics: A full system simulation platform," *IEEE Computer*, pp. 50–58, 2002.
- [22] B.C. Lee, P. Zhou, J. Yang, Y. Zhang, B. Zhao, E. Ipek, O. Mutlu, and D. Burger, "Phase-Change Technology and the Future of Main Memory," *IEEE Micro*, vol. 30, no. 1, pp. 143, 2010.
- [23] S. Thoziyoor, N. Muralimanohar, J.H. Ahn, and N.P. Jouppi, "CACTI 5.1," *HP Laboratories*, April, 2008.
- [24] X. Dong, N.P. Jouppi, and Y. Xie, "PCRAMsim: system-level performance, energy, and area modeling for phase-change ram," in *Proceedings of the 2009 International Conference on Computer-Aided Design*. ACM, 2009, pp. 269–275.
- [25] G. Dhiman, R. Ayoub, and T. Rosing, "PDRAM: A hybrid PRAM and DRAM main memory system," in *Design Automation Conference, 2009. DAC'09. 46th ACM/IEEE*. IEEE, 2009, pp. 664–669.
- [26] X. Dong, X. Wu, G. Sun, Y. Xie, H. Li, and Y. Chen, "Circuit and microarchitecture evaluation of 3D stacking magnetic RAM (MRAM) as a universal memory replacement," in *Design Automation Conference, 2008. DAC 2008. 45th ACM/IEEE*. IEEE, 2008, pp. 554–559.
- [27] C.K. Koh, W.F. Wong, Y. Chen, and H. Li, "The Salvage Cache: A fault-tolerant cache architecture for next-generation memory technologies," in *Computer Design, 2009. ICCD 2009. IEEE International Conference on*. IEEE, 2009, pp. 268–274.
- [28] X. Wu, J. Li, L. Zhang, E. Speight, R. Rajamony, and Y. Xie, "Hybrid cache architecture with disparate memory technologies," in *ACM SIGARCH Computer Architecture News*. ACM, 2009, vol. 37, pp. 34–45.
- [29] X. Wu, J. Li, L. Zhang, E. Speight, and Y. Xie, "Power and performance of read-write aware hybrid caches with non-volatile memories," in *Proceedings of the Conference on Design, Automation and Test in Europe*, 2009, pp. 737–742.
- [30] J.C. Mogul, E. Argollo, M. Shah, and P. Faraboschi, "Operating system support for NVM+ DRAM hybrid main memory," in *Proceedings of the 12th conference on Hot topics in operating systems*. USENIX Association, 2009, pp. 14–14.
- [31] J. Condit, E.B. Nightingale, C. Frost, E. Ipek, B. Lee, D. Burger, and D. Coetzee, "Better I/O through byte-addressable, persistent memory," in *Proceedings of the ACM SIGOPS 22nd symposium on Operating systems principles*. ACM, 2009, pp. 133–146.