Microelectronics Journal 45 (2014) 211–216

Contents lists available at ScienceDirect

Microelectronics Journal

journal homepage: www.elsevier.com/locate/mejo

System-level impacts of persistent main memory using a search engine

Taciano Perez [a,*], Ney Laert Vilar Calazans [b], César A.F. De Rose [b]

[a] HP R&D, Porto Alegre, Brazil
[b] PUCRS, Faculty of Informatics, Porto Alegre, Brazil

ARTICLE INFO

Article history:
Received 2 June 2013
Received in revised form
1 November 2013
Accepted 6 November 2013
Available online 26 November 2013

Keywords:
Non-volatile memory
Persistent main memory
PCRAM
RRAM
Memristor

ABSTRACT

Computer memory systems traditionally use distinct technologies for different hierarchy levels, typically volatile, high speed, high cost/byte solid state memory for caches and main memory (SRAM and DRAM), and non-volatile, low speed, low cost/byte technologies (magnetic disks and flash) for secondary storage. Currently, non-volatile memory (NVM) technologies are emerging and may substantially change the landscape of memory systems. In this work we assess system-level latency and energy impacts of a computer with persistent main memory using PCRAM and Memristor, comparing the development and execution of a search engine application implementing both a traditional file-based approach and a memory persistence approach (Mnemosyne). Our observations show that using memory persistence on top of NVM main memory, instead of a file-based approach on top DRAM/Disk, produces less than half lines of code, is more than 4× faster to develop, consumes 33× less memory energy, and executes search tasks up to 33× faster.

1. Introduction

Since the beginning of mainstream commercial computing, memory hierarchy used in computer design has been employing volatile, high speed memory technologies (SRAM and DRAM) for caches and main memory, and non-volatile, low speed technologies (magnetic disks and flash memory) for secondary storage [1].

Currently, non-volatile memory (NVM) technologies are emerging and may substantially change the landscape of memory systems. Non-volatile memory (NVM) technologies such as phase-change RAM (PCRAM), magnetic RAM (MRAM) and Memristor promise to 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 secondary storage – into a single level.

In the previous work [3], we have already assessed some of the implications of NVM, considering elements as computing time and power consumption. This work extends it, presenting an evaluation of a search engine application in a hypothetical computer with persistent main memory. Through the use of simulation, we aim to identify the major system-level impacts of persistent main memory in latency and energy. To the best of our knowledge, this is the first study evaluating both PCRAM and Memristor using a programming interface specific to persistent main memory, while considering timing, energy, and impacts on code development effort and complexity.

2. Emerging memory technologies

There are several new non-volatile memory (NVM) technologies under research today [4]. This study focuses on two of these technologies: phase-change RAM (PCRAM) and Memristor, since they are among the most mature candidate technologies for DRAM replacement. Table 1 compares the main properties of traditional memory/storage technologies with PCRAM and Memristor. Data was obtained from [2,5,4,6–9].

2.1. 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 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 [10]. These two phases can be repeatedly and rapidly cycled by applying heat to the material [10,2].

* Corresponding author. Tel.: +55 51 21213915.
E-mail address: taciano.perez@hp.com (T. Perez).

0026-2692/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.mejo.2013.11.001

Table 1
Comparison of memory/storage technologies.

Maturity	SRAM Product	DRAM Product	Disk Product	NAND Flash Product	PCRAM Adv. Dev.	Memristor Early dev.
Read latency	<10 ns	10–60 ns	8.5 ms	25 μs	48 ns	10–100 ns
Write latency	<10 ns	10–60 ns	9.5 ms	200 μs	40–150 ns	10–100 ns
Static power	Yes	Yes	Yes	No	No	No
Endurance	$>10^{15}$	$>10^{15}$	$>10^{15}$	$>10^{4}$	$>10^{8}$	$>10^{12}$
Nonvolatility	No	No	Yes	Yes	Yes	Yes

2.2. Memristor

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 retaining resistance values means that a Memristor can be used as a nonvolatile memory [11].

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) [12,13,11]. 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].

3. Persistent main memory

These novel non-volatile memory technologies can potentially make a significant amount of persistent main memory available to computer systems. It would allow a collapse of two different levels of the storage hierarchy – main memory and persistent storage – into a single level, something that has never been practically feasible 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 [14,15]. In order to fully assess system-wide impacts on 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.

Initial 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 [16–19]. Most of these proposals do not imply a radical redesign of computing systems as a whole, but localized changes to specific subsystems.

More recently, proposals for more radical system redesigns were published. A good example is the architecture of Nanostores [15] 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, and particularly suited for scale-out server environments.

In the present study, we explore 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.

Table 2
Experimental target configuration setup (guest).

Processor	x86-64 (Hammer) 20 MHz (single-core)
L1 Cache	Size: 16 Kb (D-cache)+16 Kb (I-cache) Associativity: 4-way (D-cache), 2-way (I-cache) Penalty: 100–1000 ps Replacement policy: LRU
L2 Cache	Size: 512 Kb Associativity: 8 × Penalty: 10 ns Replacement policy: LRU
Main memory	Size: 4096 MB Penalty: technology-dependent (see Table 3)
Disk	20 GB
OS	Fedora Core release 5 (Bordeaux) Kernel: 2.6.15-1.2054_FC5

4. Experimental setup

4.1. Simulation environment

The simulated system (guest) is a single-processor x86 64-bit (Hammer) computer with 4 GB of main memory. It was simulated using Virtutech Simics [20], a full-system simulator. The main parameters of the simulation are described in Table 2. 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 to have a first-order approximation of the impact of persistent main memory in current designs.

We executed a set of computing tasks (described in the next section) in three different scenarios: DRAM, PCRAM and Memristor. For each scenario, the memory latency and energy parameters were set as shown in Table 3, using a customized version of the trans-staller module in Simics. The latency values at device-level were derived from [21–23]. A low clock frequency was used in order to expedite test execution. As validation, we observed the ratio between in-memory/file-based using DRAM in both the simulated computer and real ones, and found them to be consistent. We believe that the number of memory loads/stores,