Design Exploration of Hybrid Caches with Disparate Memory Technologies

XIAOXIA WU

The Pennsylvania State University

JIAN LI, LIXIN ZHANG, EVAN SPEIGHT, and RAM RAJAMONY

IBM Austin Research Laboratory

and

YUAN XIE

The Pennsylvania State University

15

Traditional multilevel SRAM-based cache hierarchies, especially in the context of chip multiprocessors (CMPs), present many challenges in area requirements, core-to-cache balance, power consumption, and design complexity. New advancements in technology enable caches to be built from other technologies, such as Embedded DRAM (EDRAM), Magnetic RAM (MRAM), and Phase-change RAM (PRAM), in both 2D chips or 3D stacked chips. Caches fabricated in these technologies offer dramatically different power-performance characteristics when compared with SRAM-based caches, particularly in the areas of access latency, cell density, and overall power consumption. In this article, we propose to take advantage of the best characteristics that each technology has to offer through the use of Hybrid Cache Architecture (HCA) designs.

We discuss and evaluate two types of hybrid cache architectures: intercache-Level HCA (LHCA), in which the levels in a cache hierarchy can be made of disparate memory technologies; and intracache-level or cache-Region-based HCA (RHCA), where a single level of cache can be partitioned into multiple regions, each of a different memory technology. We have studied a number of different HCA architectures and explored the potential of hardware support for intracache data movement and power consumption management within HCA caches. Utilizing a full-system simulator that has been validated against real hardware, we demonstrate that an LHCA design can provide a geometric mean 6% IPC improvement over a baseline 3-level SRAM cache design under the same area constraint across a collection of 30 workloads. A more aggressive RHCA-based design provides 10% IPC improvement over the baseline. A 2-layer 3D cache stack (3DHCA) of high density memory technology within the same chip footprint gives 16% IPC improvement over the baseline. We also achieve up to a 72% reduction in power consumption over a baseline SRAM-only design. Energy-delay and thermal evaluation for 3DHCA are also presented. In

X. Wu and Y. Xie were supported in part by NSF 0702617, 0905365, 0903432, and SRC grants.

Authors' addresses: X. Wu and Y. Xie, Computer Science and Engineering Department, Pennsylvania State University, University Park, PA, 16802; email: {xwu, yuanxie}@cse.psu.edu; J. Li, L. Zhang, E. Speight, and R. Rajamony, IBM Austin Research Lab, 11501 Burnet Road, Austin, TX, 78758; email: {jianli, zhangl, speight, rajamony}@us.ibm.com.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2010 ACM 1544-3566/2010/12-ART15 \$10.00
DOI 10.1145/1880037.1880040 <http://doi.acm.org/10.1145/1880037.1880040>

ACM Transactions on Architecture and Code Optimization, Vol. 7, No. 3, Article 15, Pub. date: December 2010.

addition to the fast-slow region based RHCA, we further evaluate read-write region based RHCA designs.

Categories and Subject Descriptors: B.3.2 [**Memory Structures**]: Design Styles—*Cache memories*

General Terms: Performance

Additional Key Words and Phrases: Hybrid cache architecture, nonuniform cache architecture, cache hierarchy, Three-dimensional IC design, embedded DRAM, magnetic RAM, phase-change RAM, power, thermal

ACM Reference Format:

Wu, X., Li, J., Zhang, L., Speight, E., Rajamony, R., and Xie, Y. 2010. Design exploration of hybrid caches with disparate memory technologies. *ACM Trans. Architect. Code Optim.* 7, 3, Article 15 (December 2010), 34 pages.

DOI = 10.1145/1880037.1880040 <http://doi.acm.org/10.1145/1880037.1880040>

1. INTRODUCTION

Cache subsystems, particularly on-chip, with multiple layers of large caches have become common in modern processors such as AMD Barcelona, IBM POWER6 and Intel Nehalem. However, the advent of Chip Multiprocessors (CMPs) has increased the pressure on achieving good cache performance. Increased amounts of conventional Static Random Access Memory (SRAM) capacity or the introduction of additional cache levels into the hierarchy are constrained by such factors as chip real estate, the balance between cores and caches, and power consumption. Furthermore, more cache levels may also introduce extra performance and design overheads, including the need for efficient, scalable cache coherence protocols; longer cache miss latency; and addressing the presence of multiple copies of a cache line in multiple levels of the cache hierarchy.

One way to improve the performance of large caches is through the use of a Non-Uniform Cache Architecture (NUCA) [Kim et al. 2002]. NUCA provides a way to manage large on-chip caches where the latency has traditionally been affected by increasing wire delays. In NUCA, a large cache is divided into multiple banks with different access latencies determined by their physical locations to the source of the request. Two main types of NUCA, Static NUCA (SNUCA) and Dynamic NUCA (DNUCA), have been proposed. In SNUCA, a cache line is statically mapped into banks, with the low-order bits of the index determining the bank. In DNUCA, any given line can be mapped to several banks based on the mapping policy. The average latency is reduced by putting frequently accessed data into banks that are physically closer to the source of the request.

Traditional NUCA only utilizes the varied access latency of cache banks, due to their physical locations, to improve performance. The cache banks are typically of the same size, process, and circuit technology. The overall cache size budget is fixed for the same memory technology (e.g., SRAM). Here, we note that different memory technologies may have significantly different properties: density, read/write latency, dynamic/static power consumption, reliability features, scalability, etc. Table I lists important qualitative features of several

Table I. Comparison of Different Memory Technologies

Features	SRAM	eDRAM	MRAM	PRAM
Density	Low	High	High	Very high
Speed	Very Fast	Fast	Fast for read; slow for write	Slow for read; very slow for write
Dynamic Power	Low	Medium	Low for read; very high for write	Medium for read; high for write
Leakage Power	High	Medium	Low	Low
Non-volatility	No	No	Yes	Yes
Scalability	Yes	Yes	Yes	Yes

memory technologies: SRAM, embedded Dynamic RAM (eDRAM), Magnetic RAM (MRAM) [Hosomi et al. 2005], and Phase-change RAM (PRAM) [Chen et al. 2006]. Several observations relevant to this study can be made from Table I.

- PRAM has the highest potential density, but it also has the slowest speed. MRAM and eDRAM also have higher density than SRAM, but both are slower than SRAM. Depending on the design, MRAM read speed may be comparable to that of SRAM.
- MRAM and PRAM have very different read and write features in terms of latency and power consumption, with particularly high write power consumption.
- SRAM has high static¹ power, while MRAM and PRAM have very low static power due to their nonvolatile property.
- eDRAM has a good overall balance of dynamic and static power.

Consequently, a properly designed cache that is made of differing memory technologies may have the potential to outperform its counterpart of single technology. For example, IBM POWER[®] processors have off-chip L3 caches made of eDRAM technology. In addition, even though mixed technologies can also be integrated on the same two-dimensional (2D) chip, the emerging three-dimensional (3D) chip integration technologies may provide further design and manufacture cost benefits for on-chip mixed-technology integration [Dong and Xie 2009].

In this article, we propose and evaluate several Hybrid Cache Architecture (HCA) options to accommodate on-chip cache hierarchies. To fully take advantage of the benefits from varied memory technologies, an HCA allows levels in a cache hierarchy to be constructed from different memory technologies. Alternately, one level of cache can be partitioned into multiple regions of different memory technologies. In addition, we propose techniques such as low-overhead intracache data movement and power-aware policies to improve both cache performance and power in an HCA system. Using a hardware calibrated full-system simulator on a suite of 30 workloads, we show that an intercache-level LHCA design can provide a mean 6% IPC improvement over

¹Static power consumption is mainly due to leakage current. We interchange the notions of static power and leakage power in this work.

a baseline 3-level SRAM cache design under the same area constraint across a collection of 30 workloads. A more aggressive RHCA-based design provides 10% IPC improvement over the baseline. Finally, a 2-layer 3D cache stack (3DHCA) of high density memory technology within the same chip footprint gives 16% IPC improvement over the baseline. All configurations show substantial power reductions.

This work makes the following contributions.

- We present a general approach to constructing on-chip cache hierarchies with differing memory technologies. Such a *hybrid cache* may be either intercache-level (LHCA) or region-based (RHCA), intracache-level. We have studied hybrid caches made of combinations of SRAM, eDRAM, MRAM and PRAM under the same area constraint.
- In the RHCA design with fast and slow regions, we propose a hierarchical NUCA cache with centralized swap buffer, parallel address search, and LRU replacement across cache regions. A cache region itself can be a conventional NUCA of identical cache tiles that differ only by distance. We also discuss a read-write aware RHCA and combine it with fast-slow aware RHCA for further performance improvement.
- We propose improvements for intracache swap operations, including multi-line swap and adaptive swap. We also conduct detailed sensitivity studies for efficient swap operations.
- We propose and evaluate a drowsy hybrid cache technique with significant cache power savings.
- We extend the hybrid cache technique to 3D stacking and evaluate the benefits.

The organization of the rest of the article is as follows: Section 2 provides background information on MRAM, PRAM and 3D chip integration. Section 3 evaluates different memory technologies in terms of performance and power, and motivates the hybrid cache approach. Section 4 discusses the design and simulation methodology, and Section 5 presents a simple three-level on-chip hybrid cache hierarchy. Section 6 gives details of a hierarchical NUCA made of hybrid caches and provides evaluations of power-performance improvements over the simple hybrid cache. Section 7 applies 3D cache stacking to hybrid caches. Related work is discussed in Section 8, and we conclude in Section 9.

2. BACKGROUND

This section provides background information on two emerging memory technologies (MRAM and PRAM) and introduces 3D integration.

2.1 Magnetic RAM (MRAM)

The basic difference between MRAM and conventional RAM technologies (such as SRAM/DRAM) is that the information carrier of MRAM is a Magnetic Tunnel Junction (MTJ) instead of electric charges [Hosomi et al. 2005; Kawahara et al. 2007; Zhao et al. 2006]. Each MTJ contains *two ferromagnetic layers* and *one*

Fig. 1. A conceptual view of MTJ structure. (a) Antiparallel (high resistance), which indicates “1” state; (b) Parallel (low resistance), which indicates “0” state.

Fig. 2. An illustration of an MRAM cell.

tunnel barrier layer. Figure 1 shows a conceptual illustration of an MTJ. One of the ferromagnetic layers (the reference layer) has a fixed magnetic direction while the other one (the free layer) can change its magnetic direction via an external electromagnetic field or a spin-transfer torque. If the two ferromagnetic layers have different directions, the MTJ resistance is high, indicating a “1” state (the antiparallel case in Figure 1(a)); if the two layers have the same direction, the MTJ resistance is low, indicating a “0” state (the parallel case in Figure 1(b)).

The MRAM technology discussed in this article is called *Spin-Transfer Torque RAM* (STT-RAM), which is a new generation of MRAM technologies. STT-RAMs change the magnetic direction of the free layer by directly passing a spin-polarized current through the MTJ structure. Compared to the previous generation of MRAMs that used external magnetic fields to reverse the MTJ status, STT-RAMs have the advantage of scalability, as the *threshold current* to make the status reversal will decrease as the size of the MTJ becomes smaller.

In the STT-RAM memory cell design, the most popular structure is composed of one NMOS transistor as the access controller and one MTJ as the storage element (“1T1J” structure) [Hosomi et al. 2005]. As illustrated in Figure 2, the storage element, MTJ, is connected in series with the NMOS transistor. The NMOS transistor is controlled by the word-line (WL) signal. The detailed read and write operations for each MRAM cell is described as follows.

— *Write Operation*. When a *write operation* is performed, a positive voltage difference is established between the source-line (SL) and bit-line (BL) for writing for a “0” or a negative voltage difference is established for writing a “1.” The current amplitude required to ensure a successful status reversal

is called the threshold current. This current is related to the material of the tunnel barrier layer, the writing pulse duration, and the MTJ geometry.

- Read Operation.* When a *read operation* is desired, the NMOS is turned enabled and a voltage ($V_{BL} - V_{SL}$) is applied between the BL and the SL. This voltage is negative and is usually very small ($-0.1V$ as demonstrated in Hosomi et al. [2005]). The voltage difference will cause a current to pass through the MTJ, but it is small enough to not invoke a disturbed write operation. The value of the current is determined by the equivalent resistance of MTJs. A sense amplifier compares this current with a reference current and then decides whether a “0” or a “1” is read from the selected MRAM cell.

2.2 Phase-Change RAM (PRAM)

PRAM is another promising memory technology [Atwood and Bez 2005; Bedeschi et al. 2005; Chung 2008; Dae-Hwan et al. 2003; Lee et al. 2009; Qureshi et al. 2009; Zhou et al. 2009]. It has a wide resistance range, which is about three orders of magnitude; therefore, multilevel PRAM allows the storage of multiple bits per cell. Two to four bits per cell have already been demonstrated [Chung 2008]. The basic structure of a PRAM cell consists of a standard NMOS access transistor and a small volume of phase change material, GST ($\text{Ge}_2\text{Sb}_2\text{Te}_5$), as shown in Figure 3. The phase change material can be switched from an amorphous phase (reset or “0” state) to a crystalline phase (set or “1” state), or vice versa, with the application of heat. The read and write operations for a PRAM cell is described as follows.

- Write Operation.* there are two kinds of PRAM write operations: the *SET* operation that switches the GST into crystalline phase and the *RESET* operation that switches the GST into amorphous phase. The *SET* operation crystallizes GST by heating it above its crystallization temperature, and the *RESET* operation melt-quenches GST to make the material amorphous [Chung 2008]. These two operations are controlled by electrical current: high-power pulses for the *RESET* operation heat the memory cell above the GST melting temperature; moderate power but longer duration pulses for the *SET* operation heat the cell above the GST crystallization temperature but below the melting temperature. The temperature is controlled by passing through a certain amount of electrical current and generating the required Joule heat.
- Read Operation.* To read the data stored in PRAM cells, a small voltage is applied across the GST. Since the *SET* status and *RESET* status have a large variance on their equivalent resistance, the data is sensed by measuring the pass-through current. The read voltage is set to be sufficiently strong to invoke detectable current but remains low enough to avoid write disturbance. Like other RAM technologies, each PRAM cell needs an access device for control purpose. As shown in Figure 3, every basic PRAM cell contains one GST and one NMOS access transistor. This structure has a name of “1T1R” where “T” stands for the NMOS transistor and “R” stands for GST. The GST

Fig. 3. An illustration of a PRAM cell. When phase change material GST is in an amorphous phase, it indicates “0” state; when GST is in a crystalline phase, it indicates “1” state.

in each PRAM cell is connected to the drain-region of the NMOS in series so that the data stored in PRAM cells can be accessed by wordline controlling.

As described, MRAM and PRAM memory technologies are made of different materials than SRAM and eDRAM and have different read/write operations. However, caches constructed from these technologies have similar structure from a logic designer’s point of view due to the similarity of the peripheral circuits.

2.3 3D Integration

The increasing number of devices and functionality on a single chip leads to more complexity in interconnecting devices with a large number of metal layers. With continued technology scaling, the on-chip interconnect has emerged as a dominant source of circuit delay and power consumption [Joyner and Meindl 2002; Kim et al. 2002; Xie et al. 2006]. The reduction of interconnect delay and power consumption is of paramount importance for deep-submicron designs. 3D ICs have recently emerged as a promising means to mitigate these interconnect-related problems [Ababei et al. 2005; Davis et al. 2005; Joyner and Meindl 2002; Xie et al. 2006]. Several 3D integration technologies have been explored recently, including wire bonded, microbump, contactless (capacitive or inductive), and through-silicon-via (TSV) vertical interconnects [Davis et al. 2005]. TSV 3D integration has the potential to offer the greatest vertical interconnect density, and therefore is currently the most promising one among vertical interconnect technologies. In 3D ICs that are based on TSV technology, multiple active device layers are stacked together (through wafer stacking or die stacking) with direct vertical TSV interconnects [Xie et al. 2006]. 3D ICs offer a number of advantages over traditional two-dimensional (2D) designs [Xie et al. 2006].

- Global interconnects are shorter because vertical distances (or the length of TSVs) between two layers are usually in the range of $10\ \mu m$ to $100\ \mu m$ [Xie et al. 2006], depending on manufacturing processes.
- The reduction of average interconnect length and increased available bandwidth due to die stacking improves performance.
- Wiring length reduction (with accompanying reduced capacitance) leads to lower interconnect power consumption.
- 3D enables higher packing density and a smaller chip footprint.

Table II. Cache Parameters of Differing Memory Technologies (45nm)

Cache	Density	Latency (cycles)	Dyn. energy/op. (nJ)	Static power (W)
SRAM (1MB, 4 banks)	1	8	0.388	1.36
eDRAM (4MB, 16 banks)	4	24	0.72	0.4
MRAM (4MB, 16 banks)	4	read: 20, write: 60	read:0.4, write: 2.3	0.15
PRAM (16MB, 64 banks)	16	read: 40, write: 200	read:0.8, write:1.5	0.3

Fig. 4. Performance (top, a) and power (bottom, b) comparison of SRAM, eDRAM, MRAM, and PRAM L2 caches under the same area constraint.

—Most relevant to this work, 3D naturally supports the use of mixed-technology chips as each die can have different circuit or process technology.

3. MOTIVATION

In this section, we compare the simulated performance and power of caches individually made by pure SRAM, eDRAM, MRAM, or PRAM. We consider a 2-level on-chip cache, where the L2 can be a 1MB SRAM, 4MB eDRAM, 4MB MRAM, or 16MB PRAM under the same area constraint as determined by the appropriate density ratio. We defer the specification of cache parameters (see Table II) and our simulation methodology to Section 4.

Figure 4(a) shows the performance comparison as measured by instructions-per-cycle (IPC) of eDRAM, MRAM, and PRAM normalized to SRAM over 30 workloads under the same chip area constraint. Since the memory technologies have significantly different densities and access delays, the cache capacity and latency under the same area budget are also very different. Across the workloads, some applications prefer shorter cache latency over larger capacity, while others do better with the opposite. Some workloads are both cache latency and capacity sensitive, while others are CPU intensive and changes in cache latency and capacity do not affect their performance. As a result, we observe significant IPC variations between SRAM and the other three memory technologies. The rightmost bars are the geometric mean of the workloads. The

Fig. 5. L2 hit rate comparison of SRAM, eDRAM, MRAM, and PRAM L2 caches under the same area constraint.

Fig. 6. L1 hit rate comparison of SRAM, eDRAM, MRAM, and PRAM L2 caches under the same area constraint.

SRAM-based cache mean IPC is similar to that of a cache constructed from either eDRAM or MRAM. Although the eDRAM and MRAM caches are slower than the SRAM cache, the resulting performance loss is offset by their increased cache capacity. MRAM's relatively long write latency does not greatly hinder its performance. PRAM's average performance is not promising due to its slower read and write access latency, although it does offer the largest capacity. However, we do observe that PRAM is the best choice for a few workloads. Indeed, PRAM has recently attracted much research interest from the architecture community [Lee et al. 2009; Qureshi et al. 2009; Zhou et al. 2009].

In order to better understand the behavior of different workloads we also illustrate L2 cache hit rate (global hit rate) for these design cases, shown in Figure 5. It indicates that more than half of the 30 workloads (especially *bzip2*, *mcf*, *omnetpp*, and *cg*) require larger than 1MB L2 cache size since the hit rate increases when the cache capacity becomes larger. The hit rate remains almost the same for some workloads with increased cache capacity, indicating that these workloads have relatively small working set (e.g., *clustalw*). However, for some workloads, the hit rate is even lower with increased cache capacity, which seems to be contradictory with the intuition (e.g. *blast*, *barnes*, and *dedup*). Therefore, we also examine L1 data hit rate and find that L1 hit rate actually varies under different design cases for those workloads. L1 hit rate result is shown in Figure 6. In our simulation, L1 hits include L1 hits in the cache (L1-hit) as well as L1 hits in the load miss queue (L1-hit-M). The purpose of load miss queue is to hide the load misses so that the CPU can work on the following independent instructions. Normally, merge operations occur if the requested data is less than one cache line size. With merged requests, the buffer entries and retries may be reduced to decrease cache contention. The L1 hit rate in the cache and load miss queue varies due to the difference

of read/write latencies for different memory technologies and imbalance of read/write latencies for MRAM/PRAM. For example, some loads, which are separated by an intermediate store to the same cache line, may be disrupted with different store latencies. Due to the timing difference, L1 hit rate may be different (which is shown in Figure 6) so that L2 hit rate also varies depending on the behavior of the workloads. However, although the geometric mean of L2 hit rate increases when MRAM and PRAM are used, the performance is not improved due to the longer latency of these technologies.

Figure 4(b) depicts the static and dynamic power comparison for each option for all 30 of the workloads studied. Note that no power-aware design is applied to any of the studied memory technologies in these results. Due to its large static power consumption, SRAM consumes significantly more power than the other three options despite also having the smallest cache capacity. The power consumption variations among eDRAM, MRAM and PRAM can also be seen due to their differing dynamic read, dynamic write, and static power requirements. One may apply customized power-aware techniques to reduce the power consumption of these memory technologies. For example, since MRAM and PRAM are nonvolatile, the static power for their memory cells are negligible. Hence, one only needs to reduce the static power of the peripheral circuitry. The static reduction techniques for SRAM and eDRAM tend to be more complicated. On the other hand, the large write energy of MRAM and PRAM, in particular, suggests the need for special power-aware techniques beyond those required for SRAM and eDRAM.

In summary, the divergent latency, density, and power of differing memory technologies may affect the power and performance of a cache depending on what choice is made for the underlying technology. Due to diverse workload characteristics, no single memory technology in consideration has the best power-performance for all workloads under the same chip area constraint. SRAM appears to be much more power hungry than its three counterparts, absent any power-reduction techniques, and consequently gives emerging memory technologies performance leeway in a power constrained design environment. This motivates us to study hybrid caches, which combine the advantages of all these memory technologies in a synergistic fashion for better overall cache power-performance.

4. METHODOLOGY

In this section, we describe our simulation and design methodology. Note that the power and performance comparisons in Figure 4 (Section 3) also use the simulation methodology described here.

4.1 System Configuration

We based our parameters on searches of appropriate literature [Atwood and Bez 2005; Bedeschi et al. 2005; Chung 2008; Dae-Hwan et al. 2003; Hanzawa et al. 2007; Hosomi et al. 2005; Kawahara et al. 2007; Pellizzer et al. 2004; Zhao et al. 2006] for typical density, latency, and energy numbers for the studied memory technologies, and then scale these to 45nm technology. All cache parameters

Table III. System Configuration

Processor	Eight-way issue out-of-order, 8-core, 4GHz
L1	32KB DL1 + 32KB IL1, 128B, 4-way, 1 R/W port, latency part of LSU
L2/L3/L4	See corresponding design cases
Memory	400 cycles latency, memory controller vs. core speed 1:2, 16MB large page

Table IV. Workloads

Workload	Applications and Program size
SpecInt06	reference input: astar, bzip2, gcc, gobmk, h264, hmmer-sp, libquantum, mcf, omnetpp, perl, sjeng
SPECJBB	JVM version 1.1.8, 16 warehouses
NAS	Class C: cg, lu, mg, sp, ua
BioPerf	reference input: blast, clustalw, hmmer
PARSEC	native input: dedup, fluidanimate, freqmine, streamcluster
SPLASH2	barnes(64K particles), fmm(128k particles), fft(4M points), lu(2048*2048 matrix), ocean(2048*2048 grid)

used in this study were obtained either from CACTI [Muralimanohar et al. 2007] or its modified versions [Dong et al. 2008; Mangalagiri et al. 2008] and are shown in Table II. For all the memory technologies, the cache associativity is 16, the block size is 128B and the bank size is fixed to be 256KB. Since MRAM and PRAM are emerging memory technologies, the projection of their features tends to be more varied than the ones for established technologies such as SRAM and eDRAM, however, we have chosen cache parameters in-line with other researchers' assumptions.

We assume an 8-core CMP system with eight-way issue out-of-order cores representing future PowerPC-based processors. The experiments are conducted using a full system simulator [Bohrer et al. 2004] that has been validated against existing POWER5 hardware [Sinharoy et al. 2005]. In this article, we keep the configurations of processor core, L1 caches, on-chip interconnect, and memory system the same, and only study the design of different low-level caches (e.g., L2, L3 or L4) under the same chip area constraint or the same footprint in the case of 3D chip stacking. Table III gives our system configuration.

4.2 Workloads

The benchmarks we used in this study are chosen from a wide spectrum of workloads: SpecInt2006 [SPEC 2006], NPB [Bailey et al. 1991], SPLASH2 [Woo et al. 1995], PARSEC [Bienia et al. 2008], BioPerf [Bader et al. 2005], and SpecJBB [Morin et al. 2005]. Four PARSEC workloads covering the range of memory footprints of the whole PARSEC suite are selected. Table IV gives the problem size and other parameters of the benchmarks. For all workloads except SPLASH2, we use either sampled reference or native input sets to represent a real-world execution scenario. For SPLASH2, we increase the input set of some of the workloads such that the total number of dynamic instructions are more than 10 billion.

Fig. 7. Overview of hybrid cache design methodology.

All results, except those shown in Figure 14, examine workload performance with a single thread of execution. In the case of multithreaded simulations, we run one thread per processor core. In order to reasonably evaluate large cache designs, we construct each simulation in three phases with decreasing simulation speed: (1) we fast forward to a meaningful application phase specified by in-house workload analysis, which may take 10s–100s billion of instructions; (2) we warm up the caches by 10s billion of instructions; and (3) we simulate the system cycle-by-cycle for a few billion of instructions and collect simulation results. Both performance and power statistics are collected from cycle mode execution. Our cache power model adds the static and dynamic power of the caches used by a workload in the simulation. The static power is obtained from CACTI or its modified versions, as shown in Table II. The dynamic power factors in the number of read and write accesses and their corresponding per-access energy values are given in Table II.

4.3 Design Methodology

Throughout the hybrid cache studies presented here, we assume the chip area, or the chip footprint in the 3D integration scenario, is fixed for all the design cases. Figure 7 provides an overview of our design methodology.

In our 2D baseline system, each processor core has three levels of private caches. All three levels of caching are comprised of SRAM. This configuration

serves as the *baseline configuration* in this work. One approach to a hybrid cache is to replace SRAM L3 with eDRAM, MRAM or PRAM for larger on-chip cache capacity (Scenario A in Figure 7). This is an intercacheLevel hybrid cache, or LHCA, which is evaluated in Section 5.

We consider a core and its caches a *chiplet*. In a 2D chip design scenario, one can merge L2 and L3 to form a hybrid, coarse-grained NUCA cache with L2 fast- and slow-regions made of SRAM and eDRAM/MRAM/PRAM, respectively (Scenario B in Figure 7). The cache regions are mutually exclusive. This is an intracache-level or Region-based hybrid cache, RHCA. We discuss this scenario in Section 6. In the same section, we also evaluate a simple power-aware HCA design, called *drowsy hybrid cache*.

In a typical 3D chip stacking design scenario, an L4 cache can be stacked onto the original 2D chiplet. We consider adding a PRAM cache as an L4 due to its high density (Scenario C in Figure 7). Two 3D hybrid caches are shown. One merges all L2, L3 and L4 caches to form a fast, middle, and slow L2-regioned cache with SRAM, eDRAM and PRAM, respectively (Scenario D in Figure 7). The other combines L2 and L3 to form an L2 cache with fast and slow regions comprised of SRAM and eDRAM/MRAM, respectively, with an additional PRAM-based L3 cache (Scenario E in Figure 7). These three design points embody the *3D Hybrid Cache* (3DHCA) and are evaluated in Section 7.

We expect a few benefits of such hybrid cache designs.

- By applying memory technologies of higher density than SRAM, the effective cache size can increase significantly under the same chip area constraint.
- Because static power typically dominates in caches, applying non-volatile memory technologies can reduce cache power significantly without extra design overhead.
- By merging multiple cache levels into one, the multiple cache regions can be checked in parallel as opposed to sequentially testing each level in the cache hierarchy. When the fast region returns valid data, the search signal to the slow region can be canceled. Additionally, such timing overlapping also offers opportunity for power-aware designs.
- Due to the mutual exclusivity between the cache regions in a cache level, there are no duplicate cache lines, increasing the effective cache capacity.
- By reducing the number of cache levels, the coherence traffic between levels is reduced.
- Within a cache level, the status bit array can be accessed in coordination for better overall allocation of demand (load/store) or prefetch cache lines as well as cache line replacement.
- Performance can be improved by placing faster cache regions closer to the cache controller and employing mechanisms to maintain hot cache lines in fast regions as appropriate.
- Power-aware HCA offers extra power-performance benefits to hybrid caches.

Fig. 8. Performance (top, a) and power (bottom, b) comparison of eDRAM/MRAM/PRAM L3 caches and their SRAM baseline.

5. INTER-LAYER HYBRID CACHE

In the three-level SRAM cache baseline, we use a 256KB L2 cache and a 1MB L3 cache. We assume the cell density ratios between eDRAM, MRAM, and PRAM to SRAM are approximately 4, 4 and 16, respectively, as shown in Table II. Therefore, a 1MB SRAM, 4MB eDRAM, 4MB MRAM and 16MB PRAM L3 cache all have the same chip area. The type of memory technology in the L3 cache is the only difference between the 3-level SRAM cache baseline system and the LHCA systems with MRAM, eDRAM, or PRAM in our studies.

Figure 8(a) shows that the balance between latency and capacity is best stuck by the eDRAM L3 cache configuration. This system delivers the best average performance compared to SRAM, with a 6% improvement due to larger capacity for the same chip area although for some workloads which prefer shorter latency, the benefit of larger capacity is offset by the longer latency. MRAM and PRAM also perform better. Therefore, we choose 256KB L2 + 16-bank 4MB eDRAM L3 as our best 3-level hybrid cache configuration for LHCA. We compare other designs with this configuration in addition to the SRAM-only option in the rest of this article. Note that even though we use larger numbers of banks in larger caches to maintain a constant 256KB per bank as in previous NUCA work [Kim et al. 2002], our experiments showed little or no performance gain with increasing the number of banks beyond four.

The power comparison is shown in Figure 8(b). Similarly, to Figure 4(b) we observe significant power savings by replacing SRAM with other memory technologies. Interestingly, the MRAM power results show the SRAM-MRAM hybrid cache to consume the least normalized power, in contrast to the results in Figure 4(b). This is because the insertion of a SRAM L2 that is much larger than the L1 removes many accesses that would have otherwise hit in the MRAM L3. Therefore, the MRAM power consumption, particularly write power, is reduced significantly. Recall in Table II that MRAM's read and static power is very low.

Fig. 9. L2 plus L3 hit rate comparison of eDRAM/MRAM/PRAM L3 caches and their SRAM baseline.

However, due to the addition of a 256KB SRAM L2 to all system configurations, the relative static power consumption of all configurations increases compared to Figure 4(b). Note that no power-aware design is applied in these experiments. Power savings come entirely from introducing low-power memory technologies.

Figure 9 depicts the L2 plus L3 hit rate for these four design cases (L2 hit count plus L3 hit count and then divided by the total access count from the processor). With increased cache capacity the hit rate either remains the same or increases depending on the cache requirements of the workloads. The result is slightly different with the L2 hit rate in Section 3, in which some workloads have lower L2 hit rate with larger cache capacity. The reason is that in LHCA the inserted 256KB SRAM-based cache removes many accesses that would have otherwise hit in L3 cache, which has difference/imbalance of the read/write latencies. Therefore, with same read/write latencies in L2 cache, L1 hit rate difference caused by the difference/imbalance of the read/write latencies is significantly reduced.

6. REGION-BASED HCA

In this section, we examine the performance and power consumption of the RHCA caches consisting of one fast region made of SRAM and one slow region made of eDRAM, MRAM, or PRAM.

6.1 Hybrid L2 Cache

We propose to flatten the cache hierarchy by merging the eDRAM or MRAM or PRAM L3 into the SRAM L2. The resulting L2 cache thus consists of one small fast (SRAM) region and one large slow (eDRAM or MRAM or PRAM) region, as shown in Scenario B in Figure 7. The large hybrid L2 cache has the potential of providing fast-region access time and large-region capacity simultaneously. Fully exploring this potential requires proper cache line replacement and data migration policies.

6.1.1 Cache Line Migration Policy. Figure 10 depicts the cache line migration policy we use in our RHCA design. This policy uses one sticky bit for each line in the fast region and a two-bit saturating counter for each line in the slow region to control data movement between regions. The saturating counter records the access frequency and the bit indicates if a line cannot be moved, thus *sticky*. We consider a line whose access frequency counter is over the

Fig. 10. Cache line allocation and migration policy in the RHCA.

predetermined threshold “hot”. A “hot” line is a candidate for migrating into the fast region. Unless noted otherwise, we use 06/10 as the threshold value, which is equivalent to the MSB bit of the two-bit counter being 1.

On a cache miss, a new line is loaded into the hybrid cache. The new line always replaces the LRU line regardless of the region where the LRU line is located. If the new line is inserted into the fast region, its sticky bit is initiated to be 1. If the new line is inserted into the slow region, its saturating counter is initialized to 0.

On a cache hit, if the corresponding cache line resides in the fast region, its sticky bit is always set. If the line resides in the slow region, its saturating counter is incremented by one. If the counter reaches the threshold value, a swap operation to move the line into the fast region will be attempted. The LRU line in the fast region within the same set is selected as the potential destination. If the sticky bit of the selected line in the fast region is 0, we swap the two lines, set the sticky bit for the line moved into the fast region, and reset the saturating counter for the line moved into the slow region. If the sticky bit of the selected line in the fast region is set, we clear the sticky bit and cancel the swap operation. The role of the sticky bit is to protect a line in the fast region once and therefore effectively delay the swap once to avoid unnecessary swapping of lines.

We adopt a simple mapping scheme from Kim et al. [2002] for our RHCA design. For example, if the cache has 16 ways and there are 16 banks then each cache way can be mapped to any of these banks. In this case, each bank holds one way of the cache data. In order to perform quick search, as we mentioned in Section 4, the multiple cache regions can be checked in parallel. When the fast region returns valid data, the search signal to the slow region can be canceled.

6.1.2 Hardware Support. The hardware support for the swap operations is shown in Figure 11. The fast region and the slow region each has a tag and status array (left) as well as data array (right) allocated on both sides of the address decoder. The address decoder is replicated to meet timing demands. The trade-offs between the number of replicated decoders and bank partition is

Fig. 11. Block diagram of the proposed RHCA. Structures with slash patterns are new components.

beyond the scope of this article. The main additions are the saturating counters, the sticky bits, and the swap buffer.

A swap operation involves reading out two cache lines from two regions and writing each to the opposite region. Because of the speed difference between two regions and the contention on the cache arrays, a line read out of a region may not be able to go to the opposite region immediately and therefore must be temporarily buffered elsewhere. To simplify logic, we propose to utilize a swap buffer and serialize the swap operation as follows. First the data in the slow region is read out and placed into the buffer. Then the data in the fast region is read out and written to the slow region. Finally, the line in the swap buffer is written to fast region. Each of the three steps may take multiple cycles. The swap buffer contains multiple entries and allows multiple outstanding swap operation. Note that the first step is already being done as part of the process of loading the line into the upper level cache. We simply need to save the line in the swap buffer before it can be written to the fast region.

An alternative approach is to read both lines in parallel. In this approach, the swap buffer is either double-ported or specially arranged to allow two writes simultaneously. We have evaluated the sensitivity of swap latency and swap buffer size. The swap buffer is snooped for coherence operations. A snoop hit in the swap buffer will result in a retry response in our simulated system. Our simulation results indicate that such scenario rarely happens and is not a concern for performance degradation. More detailed analysis of swap buffer is shown in Section 6.3.

The base RHCA design uses a saturating counter for every line in the slow region. A saturating counter can be implemented with about 20 transistors, which is a very small overhead relative to the much larger tag and data arrays (hence very small power overhead). To further reduce the overhead of the saturating counters, a saturating counter can be used for a group of cache lines. In such a case, the saturating counter will record the access frequency

Table V. Fast-Slow Region Hybrid Cache Parameters in L2

RHCA (fast+slow)	Fast region (latency)	L2 size (slow region latency)	Bank number	Associativity
SRAM+eDRAM	256KB (6 cycles)	4MB (24 cycles)	16	16
SRAM+MRAM	256KB (6 cycles)	4MB (read 20; write 60)	16	16
SRAM+PRAM	256KB (6 cycles)	16MB (read 40; write 200)	64	64

of the corresponding group of cache lines. A more important goal for such placement is to see if there is a constructive alivior between the accesses of the cache lines in the same group. In other words, if the cache lines in the same group have similar access patterns, the saturating counter will be able to trigger cache line swap to the faster region earlier and help improving the performance. Furthermore, a placement of one saturating counter per group of cache lines also facilitates *multiline swap*, which moves a group of cache lines at one time. The multiline swap has the “effects” of intracache prefetching because it pushes some lines into the fast region based on access patterns of other lines. We discuss the relation between the saturating counter placement and multiline swap in Section 6.3.

6.1.3 Drowsy RHCA. The coordination support among the regions of a hybrid cache and the parallel address search among cache regions also opens new directions for power-aware designs. One simple yet effective approach is to keep the slow region of a hybrid cache in drowsy mode [Flautner et al. 2002]. For a slow region made of eDRAM, the refresh operations can be operated at a slower rate with lower voltage. In the case of a slow region made of a nonvolatile memory, drowsy mode can be power-gating the nonvolatile memory cells and/or corresponding peripheral CMOS logic. A cache line in a slow region is woken up only when it needs to provide data or execute swap operations.

6.2 Results

6.2.1 Single-Threaded Performance. RHCA can be SRAM-eDRAM, SRAM-MRAM or SRAM-PRAM based, similar to the LHCA design points. The RHCA cache design parameters for the proposed hybrid L2 cache are listed in Table V. For all the design cases, the cache has one read/write port and the block size is 128 Bytes. We compare the RHCA designs with the 3-level SRAM-only baseline as shown in Figure 7 and the SRAM-eDRAM based LHCA (the best LHCA in Section 5). Note that in Table V, each RHCA configuration is 256KB less in total size compared to its LHCA counterpart. This is to avoid complicated indexing schemes often associated with odd-sized caches. Nonetheless, even with the slightly smaller capacity, the RHCA still outperforms LHCA. We also compare our counter-based data migration design with the generational promotion approach first proposed for DNUCA by Kim et al. [2002]. Generational promotion moves a line to a closer bank on each hit. Note that in the DNUCA configuration, the mapping, search, and replacement policies are the same as in our RHCA.

Figure 12(a) shows the performance of SRAM-eDRAM RHCA. We observe that different benchmarks have different results because they have various cache requirement: some benchmarks prefer large size and some prefer small

Fig. 12. Performance of SRAM-eDRAM (top, a), SRAM-MRAM (middle, b) and SRAM-PRAM (bottom, c) RHCA.

latency. Basically, there are three categories: 1) benchmarks that require large cache size, have relatively high miss rate, and can obtain benefit from cache locality (e.g., *bzip2*, *mcf*, *omnetpp*, *cg* and *mg*). For these benchmarks, RHCA offers high performance improvement; 2) benchmarks that require modest cache size (e.g., *hmmer* and *sp*), the benefit of RHCA may not outperform data movement overhead so that the performance of RHCA is slightly degraded compared to 3-level SRAM baseline and LHCA; 3) benchmarks that are not memory intensive (e.g., *clustalw*), the performance is almost the same for the four comparison cases. Although different benchmarks show various performance trend, we still observe that the RHCA design has a geometric mean performance improvement of almost 8% over the SRAM-only design. It is 2% faster than LHCA and 6% faster than DNUCA. RHCA outperforms DNUCA because of the difference in the speed of data movement. RHCA moves frequently-accessed cache lines directly to the fast region from any bank in the slow region. DNUCA moves cache lines one bank at a time. Because of the large number of banks in the slow region, it often requires many more hits before a line moves into the fast region. In addition, since the swap frequency is high in DNUCA it results in higher swap overhead compared to RHCA. Figure 13 depicts the hit rate comparisons for 3-level SRAM baseline, best LHCA, DNUCA, and SRAM-eDRAM RHCA (L2 plus L3 hit rate for the first two cases, L2 hit rate for the remaining two cases). We observe that although the hit rate for the last three cases are similar, the

Fig. 13. L2/L2 plus L3 hit rate comparison.

Fig. 14. Performance comparison for multi-core design.

performance varies due to hit count difference in different cache regions/levels, which have differing latencies.

Figure 12(b) depicts the performance for the SRAM-MRAM RHCA design. It also shows that SRAM-eDRAM RHCA has similar, albeit slightly reduced, performance than SRAM-MRAM RHCA. This is interesting, because it means that MRAM's relatively slower write latency does not hurt its overall performance due to its faster read operations. Figure 12(c) shows that SRAM-PRAM RHCA has a 5% performance degradation relative to SRAM-eDRAM LHCA due to its long write latency. Nonetheless, it has slightly better performance than the 3-level SRAM baseline configuration.

In summary, SRAM-eDRAM RHCA can achieve a larger performance improvement over the SRAM-only cache design and also outperforms SRAM-eDRAM LHCA. The SRAM-PRAM RHCA design is not very promising for an L2 cache due to the long latency of PRAM. However, we will show that it is a promising technology for lower level caches due to its high density in Section 7. SRAM-eDRAM RHCA and SRAM-MRAM RHCA have similar performance. We focus on the analysis of SRAM-eDRAM RHCA in the following sections except explicitly mentioned sections.

6.2.2 Multithreaded Performance. Previous experiments focused on the RHCA evaluation for single-threaded runs of the workloads. In this section, we examine the RHCA designs on an eight-core CMP. We adopt the same line migration policy in the single-core design. Figure 14 shows the performance comparison between the SRAM-eDRAM RHCA design, the two baselines, and the DNUCA design for multithread benchmarks. It shows that RHCA still outperforms the baseline, LHCA and DNUCA. As a matter of fact, RHCA has

Fig. 15. Performance (top, a) and power (bottom, b) comparison for SRAM-eDRAM RHCA with different wake-up latencies.

modestly better speedups for multi-threaded runs than for single-threaded runs compared to LHCA and DNUCA because RHCA in the CMP moves hot lines and their tags closer to the cores more efficiently, effectively reducing the overhead of coherence traffic. We expect a multi-core oriented RHCA design will improve multi-threaded workloads further.

6.2.3 Power Savings with Drowsy Hybrid Cache. To generalize the evaluation of the hybrid drowsy cache option yet maintain a reasonable power comparison, we assume fixed static power of zero in the drowsy mode and only vary the wake-up and sleep transition time. In addition, we assume linearly-scaled transition voltage of wake-up and sleep transitions. The wake up time from sleep mode to fully functional mode is set to be 1/4/8/16 cycles in order to study the performance loss and power saving under different wake-up times. The performance and power comparison among the baseline and different wake-up latencies are shown in Figure 15(a) and Figure 15(b) (the six bars in Figure 15(b) are the power for LHCA, RHCA, and power aware design with four different wake-up cycles). The results indicate that up to 8 cycles wake-up latency incur less than 1% loss in IPC. With 16 cycles wake-up latency, the performance degradation compared to HCA is less than 3%. In exchange, average power saving is around 50%. Note that the power saving is beyond what the LHCA design has achieved (53% in Section 5). In Figure 15(b), the power consumption is partitioned into three segments: static power, dynamic power for normal data access and dynamic power for data migration (swap operations). The remaining static power of the drowsy hybrid cache mainly comes from the fast region, which is not drowsy. The configuration of the plots here uses 2-hit saturating counter threshold. If using 1-hit threshold, one would expect more swap power. Due to the swap power overhead, the power consumption of RHCA is slightly higher than that of LHCA.

Fig. 16. Performance (top, a) and Power (bottom, b) comparison for different threshold value of saturating counters for SRAM-eDRAM RHCA.

6.3 Discussion

6.3.1 Threshold Sensitivity of Saturating Counters. The default threshold of 2 means that a line swap is triggered if it has been hit twice. A larger threshold makes more conservative choices and may miss more opportunities. A lower threshold makes more aggressive choices and may thrash the fast region.

To better understand the effects of the threshold, we vary the threshold from one hit to three hits in our experiments and show the performance and power results in Figure 16. From the performance perspective, the results show that most of benchmarks are insensitive to the threshold. A few benchmarks, *bzip2*, *cg*, *lu*, and *specjbb*, prefer three hits (reduce cache pollution), while *mg* and *fft* prefer one hit (one possible reason is that lots of data are used twice or three times). The performance difference between different thresholds is less than 1%. These results suggest that a one-hit threshold is sufficient for most applications, which effectively can be replaced by the existing LRU bit in the status array. However, we also find that higher threshold values tend to help prefetched data more, and to reduce cache pollution to the fast region. From the power perspective, many benchmarks consume more swap power when the counter threshold is one because there are more swap operations. Some benchmarks have different trend since the performance is also considered when calculating the power consumption. More power may be consumed with even lower energy consumption due to the delay difference. For some benchmarks, the swap power for 1-hit is very small. The reason is that in our simulation, before the swap occurs, the cache line is checked to guarantee that it does not reside in other queues (waiting for some requests). Therefore, the swap operations may be much less than the hit counts if many cache lines are in queues before swapping. Another observation is that more swap operations offer higher performance improvement in RHCA (e.g., *bzip2*, *mcf* and *cg*),

Fig. 17. Buffer entry and latency study.

indicating that swap helps to reduce the latency for frequently accessed data for these benchmarks.

6.3.2 Sensitivity of Swap Latency and Buffer Entries. As discussed in Section 6.1.2, a swap may be done in three steps or two steps. We evaluate these two swap options with their corresponding access latency, buffer sizes, and data bus width. Particularly, we compare the performance of a 16-entry buffer, a two-entry buffer and a single-entry buffer in RHCA. Both the 16-entry and two-entry buffers deploy two-step swap operations, while the single-entry buffer three-step.

For a design with two swap buffers, we assume that one swap operation takes 38 cycles. It includes one read/write latency for the fast region (6 cycles), one read/write latency for the slow region (24 cycles), and two bus transfers. For a bus width of 32 bytes, transferring 128 bytes cache line data from the cache/buffer to the buffer/cache takes 4 cycles. On the other hand, for a design with one swap buffer, we conservatively assume that one swap operation requires 76 cycles ($6 \times 2 + 24 \times 2 + 4 \times 4$ cycles). We find that many workloads perform better with the two-step swap process and 16 entries are sufficient for all the benchmarks. The result shown in Figure 17 indicates that it gives 2% performance difference between the 16-entry buffer and the 1-entry buffer. There are two benchmarks that have better performance with fewer buffer entries. One possible reason is the contention caused by more buffer entries. In this work, 16-entry buffer is used for other experiments.

6.3.3 Replacement and Insertion Policy. In the previous study, data replacement is based on LRU bits to choose the LRU line in both regions. We have also evaluated three other replacement policies: (1) Replace the line in the fastest bank and evict this line to lower level; (2) Replace the line in the slowest bank and evict this line to lower level; (3) Put the line in the fastest bank, move the existing line in the fastest bank to a randomly selected bank in slow region, and evict the line in the selected slow bank. As expected, we observe that LRU-based policy performs equally well with option 3 for most workloads, both of which are better than the other two options. These results are shown in Figure 18.

6.3.4 Adaptive Swap. From the comparison between LHCA and RHCA, we observe that swaps can help many benchmarks. However, they can also pollute the cache and hurt performance, even if the frequency of swaps is not

Fig. 18. Different replacement policies comparison for SRAM-eDRAM RHCA.

Fig. 19. IPC comparison for prefetch effect on SRAM-eDRAM RHCA design.

high, as mentioned in Section 6.2.1. Based on this observation, we implement an simple adaptive swap scheme in the RHCA design. When the ratio between the number of swaps and the number of accesses to the slow region is less than a threshold (in our evaluation, the threshold is set to be 15%), we disable swap operations and make the slow region become a victim buffer of the fast region. Such an adaptive swap scheme offers almost 2% mean IPC improvement over the simple swap policy across the 30 workloads.

6.3.5 Prefetch and Multiline Swap. Stride prefetch technique utilizes repeatable memory access patterns to reduce cache miss rate [Jouppi 1990; Iacovovici et al. 2004]. Since L1 and L2 misses often show repetitive access patterns, current software and hardware prefetch engine use miss patterns to predict future misses. Specifically, hardware prefetch engines observe the stride between two recent misses and verify the stride using following misses. If the stride misses reaches a threshold, the prefetch engine initiates accesses to the following cache lines before being referenced by future instructions. In this section, we evaluate L2 prefetch effect on our RHCA design and compared it with RHCA and baseline without L2 prefetch (all cases have L1 prefetch). Our prefetch technique is based on an aggressive Power5-like hardware prefetch engine. Figure 19 shows the performance comparison of SRAM-eDRAM RHCA design with and without L2 prefetch. The result shows that RHCA design with L2 prefetch has 5% overall performance improvement over RHCA without L2 prefetch.

Multiline swap can be combined with prefetch since multi-line swap is more like a prefetch within the cache level. We can use a counter to keep track of accesses to a corresponding group of cache lines instead of each individual line. There are two ways to update counters in this case. The first is to update the counter based on hits to a given line in the group, but the movement is group

based. The second is to update the counter on hits to any line in the group. When a swap is triggered, multiple cache lines, two in our experiments, can be moved in tandem. Our results show that group line movement gives about extra 1% mean IPC improvement over the prefetch result.

6.3.6 Read-Write Aware RHCA. Previous design cases focus on fast-slow region based RHCA, for example, SRAM is fast region and eDRAM/MRAM/PRAM is slow region. As we mentioned in Section 1, MRAM and PRAM have very different read and write features in terms of latency and power consumption. On the other hand, we observe that different applications have differing read and write behaviors, and read/write access patterns also vary along the time line for one application [Wu et al. 2009]. In addition, reads are usually more critical than writes since reads are in the critical path. Therefore, reads and writes may be distinguished in order to achieve better power-performance. Consequently, if we combine read-write region with previous fast-slow region based RHCA design it may offer performance/power improvement. As a case study, we reexamine the cache line migration policy and hardware support for fast-slow and read-write region HCA (RWHCA) and illustrate the performance and power results for our SRAM/MRAM based RWHCA design.

In our SRAM/MRAM based RWHCA design, SRAM region is considered as fast/read region and MRAM portion is acted as slow/write region. In our new cache line migration policy, the only difference with Figure 10 is that before the swap, if we detect it is a write operation then the swap is canceled and the sticky bit is cleared, otherwise the swap operation resumes. The purpose of this policy is to reduce the write operations in the fast/read (SRAM) region and prevent the pollution to the faster read operations due to the high criticality of reads. The hardware support for RWHCA remains the same except there are extra control bits enabling the swap. Figure 20 depicts the performance and power comparison for SRAM/MRAM based RHCA and RWHCA. The results show that a number of benchmarks (e.g., *dedup*, *lu_splash*, and *ua*) obtain substantial performance benefit from RWHCA since they have large amount of write operations and the reduction of the writes to the fast/read region allows more reads to reside. However, some benchmarks (e.g., *bzip2* and *hmmersp*) do not get benefit since the swapping of writes in RHCA, which would provide performance improvement, is disabled in RWHCA. Figure 20(b) depicts the power comparison between RHCA and RWHCA. Consistent with theoretical analysis, the result shows that less swap power is consumed in RWHCA but more writes hit in the MRAM region, which has higher write power, resulting in higher normal dynamic power. Overall, RWHCA achieves slightly better performance than RHCA but also consumes slightly more power.

7. 3D HYBRID CACHE STACKING

3D cache stacking enables the addition of more cache levels without sacrificing the number of cores. These extra cache levels should be at least a few times larger than the cache level above it in the cache hierarchy to effectively reduce miss rate. We assume the 3D cache layer has the same footprint as its

Fig. 20. Performance (top, a) and Power (bottom, b) comparison between SRAM/MRAM based RHCA and RWHCA.

corresponding 2D chiplet, which consists of a processor core and its original caches. If a memory technology of the same density is used, then multilayer 3D cache stacking is anticipated. However, multilayer 3D stacking may incur mounting problems in power delivery, cooling, and TSV efficiency. Therefore, we expect a denser memory technology to be an alternative approach to multilayer 3D cache stacking. In this article, we consider PRAM (Scenario C in Figure 7). Besides its high density, PRAM also has very low static power, which further helps address the cooling issues with 3D. We scale the latency and power parameters of PRAM as shown in Table II for the 3D cases. We assume the processor and memory domain clock frequencies of 3D are the same as its 2D counterpart.

Similarly to the approach in Section 6, one can combine some or all of the L2, L3 and L4 caches to obtain a region-based hybrid cache (RHCA). The design issues of this extension are also similar to what we have discussed in Section 6. We also study two more design scenarios: combining all L2, L3 and L4 to obtain a three-region hybrid cache of fast, middle and slow regions (Scenario D in Figure 7); and combining L2 and L3 as we did in Section 6.1 and logically upgrade L4 to L3 (Scenario E in Figure 7). As a result, we have three 3DHCA designs for comparison: (1) Four-level caches with 256KB SRAM L2, 4MB eDRAM L3 and 32MB PRAM L4 (3DHCA-C, Scenario C in Figure 7); (2) A three-region RHCA with 256KB SRAM fast region, 3.75MB eDRAM middle region and 28MB PRAM slow region (3DHCA-D, 32MB in total, Scenario D in Figure 7); (3) A 4MB SRAM-eDRAM RHCA L2 (Section 6) and 32MB L3 cache (3DHCA-E, Scenario E in Figure 7). In 3DHCA-D, the frequently used cache line in the PRAM-based slow region can be mitigated to fast region as well as middle region (LRU-based replacement policy) in order to prevent thrashing the fast region. The frequently used cache line in the middle region is mitigated to the fast region, which is the same as two-region RHCA.

Fig. 21. 3DHCA performance comparison.

7.1 Performance and Power Evaluation

Figure 21 illustrates the performance comparison of these three design cases with 3-level SRAM baseline, SRAM-eDRAM LHCA (Scenario A in Figure 7) and SRAM-eDRAM RHCA (Scenario B in Figure 7). The results show that all three design cases exhibit large improvements over 3-level SRAM baseline and SRAM-eDRAM LHCA. In addition, they achieve on average 1%, 3% and 6% IPC improvement over SRAM-eDRAM RHCA. Among them, Scenario E has better performance than both Scenario C and D and it achieves 16% IPC improvement than the pure SRAM baseline. We observe that although the total cache capacity of 3DHCA-D is smaller than that of 3DHCA-C, it has on average better performance, indicating that the multi-region hybrid cache is efficiently used to take advantage of the latency and capacity tradeoffs. Since 3DHCA-E has best performance in 3D design cases, it indicates that RHCA performs better in middle level caches, such as the L2, where both latency and capacity issues affect sensitive applications.

Figure 22 illustrates the power comparison of these three 3D design cases with 3-level SRAM baseline, SRAM-eDRAM LHCA, and SRAM-eDRAM RHCA (the six bars in Figure 22 are the power for 3-level SRAM, LHCA, RHCA, 3DHCA-C, 3DHCA-D, and 3DHCA-E). We observe that the leakage power of 3D design cases are much larger than that of LHCA and RHCA because of the peripheral circuits for the extra PRAM layer. Among the three 3D design cases, 3DHCA-E assumes more power than 3DHCA-C and 3DHCA-D. The power of 3DHCA-D is slightly larger than that of 3DHCA-C although the capacity of 3DHCA-D is smaller. This is because the swap operations in 3DHCA-D incur extra power overhead. The results indicate that all three design cases have higher power consumption than LHCA and RHCA in 2D case but still have lower power consumption than 3-level SRAM baseline.

We also illustrate Energy, Energy-Delay Product (EDP), and Energy-Delay² Product (ED^2P) results for all design cases in Figure 23(a), Figure 23(b) and Figure 23(c), respectively. The energy metric applies all weight on the energy consumption of applications, which is important to evaluate battery based systems, where energy efficiency affects battery life. On the other hand, the

Fig. 22. 3DHCA power comparison.

Fig. 23. Energy (top, a) and Energy-Delay Product (middle, b), and Energy-Delay² Product (bottom, c) comparison.

ED^2P metric applies more weight on performance, which is useful to evaluate high-end systems. The EDP metric falls in the middle of the former two.

We observe that RHCA has the best (lowest) overall energy, EDP and ED^2P across the workloads, with LHCA comes closely next. However, each individual application can exhibit drastic different characteristics over the three metrics. For example, *libquantum* favors 3DHCA configurations with EDP and ED^2P , thanks to the larger performance improvement with large 3D caches. When performance improvement with large 3D caches is negligible as with *hmmer* – *sp*,

Fig. 24. Peak temperature comparison between RHCA and 3DHCA-E.

3DHCA generally is worse (higher) across all energy, EDP and ED^2P metrics. Such energy-delay tradeoffs can be used to guide the cache subsystem design that targets certain workloads or application markets. The fact the RHCA sustains the best overall power-performance across all three metrics exhibits its advantage over the other design choices.

7.2 Thermal Evaluation

Power and thermal issues have become a primary concern in traditional 2D IC design. The stacking of multiple active layers in 3D design leads to even higher power densities than its 2D counterpart, which in turn exacerbates the thermal problem. Therefore, it is important to conduct thermal evaluation for our 3DHCA design.

In this section, we evaluate the temperature of the 3D design cases as well as 2D design cases and examine if our 2-layer 3D stacking designs cause severe thermal problem. In the evaluation, the bottom layer (the layer the processor resides) is allocated right next to heat sink to help heat dissipation. PRAM layer is stacked on top of the bottom layer and further away from the heat sink. To run the thermal analysis, we use Hotspot [Huang et al. 2008] for the temperature estimation and compare the peak temperature for 2D and 3D cases. The power of processor cores is obtained using a core power model which is along the line with Wattch [Brooks et al. 2000]. Depending on different workloads, the processor power varies from 5 Watt to 9 Watt. The power of cache is calculated using cache parameters from CACTI and cache access statistics from the simulator. The areas of the core and caches are along the lines of POWER5 processor [Sinharoy et al. 2005], which we scale to 45nm technology node.

The results show that the peak temperatures for different 2D cases within one benchmark are similar due to relatively small power difference (same for different 3D cases). Therefore, we only illustrate the peak temperature comparison between one 2D case (RHCA) and one 3D case (3DHCA-E) for different workloads, shown in Figure 24. In our evaluation, the ambient temperature is set to be 45°C and the thermal interface resistivity is 0.5W/mK. The figure shows that the peak temperature in 3DHCA-E is, on average, 5°C higher than that of RHCA, which is consistent with previous work on 3D stacking thermal evaluation [Loh 2008; Madan et al. 2009] and indicates that the temperature of our 2-layer 3D stacking design is acceptable. In the evaluation, the drowsy

mode cache is not considered. If drowsy cache technique is used in 3D cases, the peak temperature is expected to be lower.

8. RELATED WORK

A lot of research has been done in caches. In the following, we only list a few that are very related to our work.

There are several NUCA studies for single core and chip multiprocessors (CMP) in the literature [Beckmann and Wood 2004; Chishti et al. 2003; Chishti et al. 2005; Huh et al. 2005; Kim et al. 2002]. Kim et al. [2002] propose the novel NUCA concept for large caches and compare several SNUCA and DNUCA designs in which data movement is based on generational promotion. Subsequently, distance associativity based NUCA, called NuRapid, is proposed in single core and multi-core designs [Chishti et al. 2003; Chishti et al. 2005]. NuRapid decouples data placement from tag placement by separating it from set associativity. In Beckmann and Wood [2004], transmission line based NUCA is presented for multi-core design and a prefetch scheme is evaluated for performance improvement. Recently, the sharing degree of shared L2 NUCA caches in CMP design was examined [Huh et al. 2005]. However, in these NUCA designs, the access latency differences are mainly from interconnect delays. In our RHCA design, the latency as well as power differences are from disparate memory technologies. Additionally, our RHCA is a hierarchical design. At a high level, RHCA is made of cache regions of different sizes with differing memory technologies. At a base level, a cache region itself can be a conventional NUCA.

Recently, 3D has enabled mixed-technology integration and offers advantages of lower global wire delay and smaller area. Previous work focuses on the performance improvement and power reduction by stacking cache or main memory on top of processors [Bryan et al. 2006; Ghosh and Lee 2007; Kgil et al. 2006; Li et al. 2006; Liu et al. 2005; Loh 2008]. Typically, these stacked caches are SRAM-based, as opposed to our LHCA design. In addition to LHCA, we also evaluate RHCA and 3DHCA, which may consist of a level of RHCA cache, to fully explore the HCA design space.

In the context of stacking DRAM memory in 3D, Loh [2008] proposes an interesting 3D design of on-chip main memory to boost chip performance. Recently, a reconfigurable cache in 3D stacking is proposed, in which the baseline cache is made of SRAM and a reconfigurable eDRAM based cache can be turned on/off based on the cache size requirements [Madan et al. 2009]. Another work studied SRAM-MRAM based hybrid cache to improve performance of pure MRAM based cache [Sun et al. 2009]. Read-write aware hybrid cache design combining SRAM with nonvolatile memories MRAM/PRAM is also evaluated [Wu et al. 2009]. In this work, we study eDRAM/MRAM/PRAM in our HCA cache design, which can be applied in both a 2D and 3D context.

9. CONCLUSIONS

In this article, we have presented a hybrid cache architecture to construct on-chip cache hierarchies with differing memory technologies. We have proposed

both inter- and intracache level hybrid cache designs. We have studied hybrid caches made of combinations of SRAM, eDRAM, MRAM and PRAM under the same area constraint. In addition, we have proposed and evaluated low-overhead intracache data movement power-aware policies and their hardware support to both improve cache performance and reduce power. For a collection of 30 workloads, the geometric mean of simulation results based on a hardware calibrated full-system simulator show that an intercache-layer HCA design can provide 6% IPC improvement over a baseline 3-level SRAM cache design under the same area constraint. A more aggressive RHCA-based design provides 10% IPC improvement over the baseline. Finally, a 2-layer 3D cache stack (3DHCA) of high density memory technology within the same chip footprint gives 16% IPC improvement over the baseline. We also observe a power reduction of up to 70% across all configurations. Furthermore, we evaluate read-write regions for RHCA, and discuss energy-delay and thermal evaluation for 3DHCA.

Overall, we have shown the potential of applying hybrid caches to rebalance the cache subsystem design, and we have discussed a design direction to further improve hybrid cache power-performance. As an initial study, we have mainly presented hybrid cache performance in the context of single-threaded execution. Multithreaded workloads on chip multiprocessors opens further avenues for exploration beyond the initial results presented here.

REFERENCES

- ABABEI, C., FENG, Y., GOPLEN, B., MOGAL, H., ZHANG, T., BAZARGAN, K., AND SAPATNEKAR, S. S. 2005. Placement and routing in 3D integrated circuits. *IEEE Design Test Comput.* 22, 6, 520–531.
- ATWOOD, G. AND BEZ, R. 2005. Current status of chalcogenide phase change memory. In *Device Research Conference Digest.*, Vol. 1, 29–33.
- BADER, D. A., LI, Y., LI, T., AND SACHDEVA, V. 2005. BioPerf: A benchmark suite to evaluate high-performance computer architecture on bioinformatics applications. In *Proceedings of the IEEE International Symposium on Workload Characterization.* 163–173.
- BAILEY, D., BARTON, J., LASINSKI, T., AND SIMON, H. 1991. The NAS parallel benchmarks. Tech. rep. RNR-91-002 revision2. 453–464.
- BECKMANN, B. M. AND WOOD, D. A. 2004. Managing wire delay in large chip-multiprocessor caches. In *Proceedings of the International Symposium on Microarchitecture.* IEEE Computer Society, Los Alamitos, CA, 319–330.
- BEDESCHI, F., BONIZZONI, E., CASAGRANDE, G., GASTALDI, R., RESTA, C., TORELLI, G., AND ZELLA, D. 2005. SET and RESET pulse characterization in BJT-selected phase-change memories. In *Proceedings of the IEEE International Symposium on Circuits and Systems.* Vol. 2, 1270–1273.
- BIENIA, C., KUMAR, S., SINGH, J. P., AND LI, K. 2008. The PARSEC benchmark suite: Characterization and architectural implications. In *Proceedings of the 17th International Conference on Parallel Architectures and Compilation Techniques.* ACM, New York, NY, 72–81.
- BOHRER, P., PETERSON, J., ELNOZAHY, M., RAJAMONY, R., GHEITH, A., ROCKHOLD, R., LEFURGY, C., SHAFI, H., NAKRA, T., SIMPSON, R., SPEIGHT, E., SUDEEP, K., HENSBERGEN, E. V., AND ZHANG, L. 2004. Mambo: A full system simulator for the powerpc architecture. *SIGMETRICS Perform. Eval. Rev.* 31, 4, 8–12.
- BROOKS, D., TIWARI, V., AND MARTONOSI, M. 2000. Wattch: A framework for architectural-level power analysis and optimizations. In *Proceedings of the International Symposium on Computer Architecture.* ACM, New York, NY, 83–94.
- BRYAN, B., MURALI, A., NED, B., JOHN, D., LEI, J., GABRIEL, H. L., DON, M., PAT, M., DONALD, W. N., DANIEL, P., PAUL, R., JEFF, R., SADASIVAN, S., JOHN, S., AND CLAIR, W. 2006. Die stacking (3D)

- microarchitecture. In *Proceedings of the International Symposium on Microarchitecture*. IEEE Computer Society, Los Alamitos, CA, 469–479.
- CHEN, Y., RETTNER, C., RAOUX, S., BURR, G., CHEN, S., SHELBY, R., SALINGA, M., RISK, W., HAPP, T., MCCLELLAND, G., BREITWISCH, M., SCHROTT, A., PHILIPP, J., LEE, M., CHEEK, R., NIRSCHL, T., LAMOREY, M., CHEN, C., JOSEPH, E., ZAIDI, S., YEE, B., LUNG, H., BERGMANN, R., AND LAM, C. 2006. Ultra-thin phase-change bridge memory device using GeSb. In *Proceedings of the International Electron Devices Meeting*, 1–4.
- CHISHTI, Z., POWELL, M. D., AND VIJAYKUMAR, T. N. 2003. Distance associativity for high-performance energy-efficient non-uniform cache architectures. In *Proceedings of the International Symposium on Microarchitecture*. 55–64.
- CHISHTI, Z., POWELL, M. D., AND VIJAYKUMAR, T. N. 2005. Optimizing replication, communication, and capacity allocation in CMPs. *SIGARCH Comput. Archit. News* 33, 2, 357–368.
- CHUNG, L. 2008. Cell design considerations for phase change memory as a universal memory. In *Proceedings of the International Symposium on VLSI Technology, Systems and Applications*. 132–133.
- DAE-HWAN, K., DONG-HO, A., KI-BUM, K., WEBB, J. F., AND KYUNG-WOO, Y. 2003. One-dimensional heat conduction model for an electrical phase change random access memory device with an $8F^2$ memory cell. *J. Appl. Phys.* 94, 5, 3536–3542.
- DAVIS, W. R., WILSON, J., MICK, S., XU, J., HUA, H., MINEO, C., SULE, A. M., STEER, M., AND FRANZON, P. D. 2005. Demystifying 3D ICs: The pros and cons of going vertical. *IEEE Desi. Test Comput.* 22, 6, 498–510.
- DONG, X., WU, X., SUN, G., XIE, Y., LI, H., AND CHEN, Y. 2008. Circuit and microarchitecture evaluation of 3D stacking magnetic RAM (MRAM) as a universal memory replacement. In *Proceedings of the Design Automation Conference*. 554–559.
- DONG, X. AND XIE, Y. 2009. System-level cost analysis and design exploration for three-dimensional integrated circuits (3D ICs). In *Proceedings of the Asia and South Pacific Design Automation Conference*. IEEE Los Alamitos, CA, NJ, 234–241.
- FLAUTNER, K., KIM, N. S., MARTIN, S., BLAAUW, D., AND MUDGE, T. 2002. Drowsy caches: Simple techniques for reducing leakage power. *SIGARCH Comput. Archit. News* 30, 2, 148–157.
- GHOSH, M. AND LEE, H.-H. S. 2007. Smart refresh: An enhanced memory controller design for reducing energy in conventional and 3D die-stacked DRAMs. In *Proceedings of the International Symposium on Microarchitecture*. 134–145.
- HANZAWA, S., KITAI, N., OSADA, K., KOTABE, A., MATSUI, Y., MATSUZAKI, N., TAKAURA, N., MONIWA, M., AND KAWAHARA, T. 2007. A 512KB embedded phase change memory with 416kB/s write throughput at 100uA cell write current. In *Proceedings of the IEEE International Solid-State Circuits Conference*. 474–616.
- HOSOMI, M., YAMAGISHI, H., YAMAMOTO, T., ET AL. 2005. A novel nonvolatile memory with spin torque transfer magnetization switching: Spin-RAM. In *Proceedings of the International Electron Devices Meeting*. 459–462.
- HUANG, W., SANKARANARAYANAN, K., SKADRON, K., RIBANDO, R., AND STAN, M. 2008. Accurate, pre-RTL temperature-aware design using a parameterized, geometric thermal model. *IEEE Trans. Comput.* 57, 9, 1277–1288.
- HUH, J., KIM, C., SHAFI, H., ZHANG, L., BURGER, D., AND KECKLER, S. W. 2005. A NUCA substrate for flexible CMP cache sharing. In *Proceedings of the International Conference on Supercomputing*. 31–40.
- IACOBOWICI, S., SPRACKLEN, L., KADAMBI, S., CHOU, Y., AND ABRAHAM, S. G. 2004. Effective stream-based and execution-based data prefetching. In *Proceedings of the International Conference on Supercomputing*. 1–11.
- JOUPPI, N. 1990. Improving direct-mapped cache performance by the addition of a small fully-associative cache and prefetch buffers. In *Proceedings of the International Symposium on Computer Architecture*. 364–373.
- JOYNER, J. W. AND MEINDL, J. D. 2002. Opportunities for reduced power dissipation using three-dimensional integration. In *Proceedings of the Interconnect Technology Conference*. 148–150.

- KAWAHARA, T., TAKEMURA, R., MIURA, K., AND ET AL. 2007. 2Mb Spin-transfer torque RAM (SPRAM) with bit-by-bit bidirectional current write and parallelizing-direction current read. In *Proceedings of the IEEE International Solid-State Circuits Conference*. 480–617.
- KGIL, T., D'SOUZA, S., SAIDI, A., BINKERT, N., DRESLINSKI, R., MUDGE, T., REINHARDT, S., AND FLAUTNER, K. 2006. PicoServer: Using 3D stacking technology to enable a compact energy efficient chip multiprocessor. *SIGARCH Comput. Archit. News* 34, 5, 117–128.
- KIM, C., BURGER, D., AND KECKLER, S. W. 2002. An adaptive, non-uniform cache structure for wire-delay dominated on-chip caches. In *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems*. 211–222.
- LEE, B. C., IPEK, E., MUTLU, O., AND BURGER, D. 2009. Architecting phase change memory as a scalable DRAM alternative. In *Proceedings of the 36th International Symposium on Computer Architecture*. 2–13.
- LI, F., NICOPoulos, C., RICHARDSON, T., XIE, Y., NARAYANAN, V., AND KANDEMIR, M. 2006. Design and management of 3D chip multiprocessors using network-in-memory. In *Proceedings of the International Symposium on Computer Architecture* 34, 2, 130–141.
- LIU, C., GANUSOV, I., BURTSCHER, M., AND TIWARI, S. 2005. Bridging the processor-memory performance gap with 3D IC technology. *IEEE Des. Test Comput.* 22, 6, 556–564.
- LOH, G. H. 2008. 3D-stacked memory architectures for multi-core processors. In *Proceedings of the International Symposium on Computer Architecture*. 453–464.
- MADAN, N., ZHAO, L., MURALIMANO HAR, N., UDIPI, A., BALASUBRAMONIAN, R., IYER, R., MAKINENI, S., AND NEWELL, D. 2009. Optimizing communication and capacity in a 3D stacked reconfigurable cache hierarchy. In *Proceedings of the IEEE International Symposium on High Performance Computer Architecture*, 262–274.
- MANGALAGIRI, P., SARPATWARI, K., YANAMANDRA, A., NARAYANAN, V., XIE, Y., IRWIN, M. J., AND KARIM, O. A. 2008. A low-power phase change memory based hybrid cache architecture. In *Proceedings of the Great Lakes Symposium on VLSI*. 395–398.
- MORIN, R., KUMAR, A., AND ILYINA, E. 2005. A multi-level comparative performance characterization of SPECjbb2005 Versus SPECjbb2000. In *Proceedings of the IEEE International Workload Characterization*. 67–75.
- MURALIMANO HAR, N., BALASUBRAMONIAN, R., AND JOUPPI, N. 2007. Optimizing NUCA organizations and wiring alternatives for large caches with CACTI 6.0. In *Proceedings of the 40th Annual IEEE/ACM International Symposium on Microarchitecture*. IEEE Computer Society, Los Alamitos, CA, 3–14.
- PELLIZZER, F., PIROVANO, A., OTTOGALLI, F., MAGISTRETTI, M., SCARAVAGGI, M., ZULIANI, P., TOSI, M., BENVENUTI, A., BESANA, P., CADEO, S., MARANGON, T., MORANDI, R., PIVA, R., SPANDRE, A., ZONCA, R., MODELLI, A., VARESI, E., LOWREY, T., LACAITA, A., CASAGRANDE, G., CAPPELLETTI, P., AND BEZ, R. 2004. Novel utrench phase-change memory cell for embedded and stand-alone non-volatile memory applications. In *Proceedings of the Symposium on VLSI Technology*. 18–19.
- QURESHI, M. K., SRINIVASAN, V., AND RIVERS, J. A. 2009. Scalable high performance main memory system using phase-change memory technology. In *Proceedings of the 36th International Symposium on Computer Architecture*. 24–33.
- SINHARROY, B., KALLA, R. N., TENDLER, J. M., EICKEMEYER, R. J., AND JOYNER, J. B. 2005. Power5 system microarchitecture. *IBM J. Resear. Devel.* 49, 4/5, 505–521.
- SPEC. 2006. Standard Performance Evaluation Corporation. <http://www.spec.org/cpu2006/>.
- SUN, G., DONG, X., XIE, Y., LI, J., AND CHEN, Y. 2009. A novel architecture of the 3D stacked MRAM L2 cache for CMPs. In *Proceedings of the IEEE International Symposium on High Performance Computer Architecture*. 239–249.
- WOO, S. C., OHARA, M., TORRIE, E., SINGH, J. P., AND GUPTA, A. 1995. The SPLASH-2 programs: Characterization and methodological considerations. *SIGARCH Comput. Archit. News* 23, 2, 24–36.
- WU, X., LI, J., ZHANG, L., SPEIGHT, E., AND XIE, Y. 2009. 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*. 737–742.
- XIE, Y., LOH, G. H., BLACK, B., AND BERNSTEIN, K. 2006. Design space exploration for 3D architectures. *J. Emerg. Technol. Comput. Syst.* 2, 2, 65–103.

- ZHAO, W., BELHAIRE, E., MISTRAL, Q., CHAPPERT, C., JAVERLIAC, V., DIENY, B., AND NICOLLE, E. 2006. Macro-model of spin-transfer torque based magnetic tunnel junction device for hybrid magnetic-CMOS design. In *Proceedings of the IEEE International Behavioral Modeling and Simulation Workshop*. 40–43.
- ZHOU, P., ZHAO, B., YANG, J., AND ZHANG, Y. 2009. A durable and energy efficient main memory using phase change memory technology. In *Proceedings of the 36th International Symposium on Computer Architecture*. 14–23.

Received May 2009; revised December 2009; accepted February 2010