Data Mapping for Higher Performance and Energy Efficiency in Multi-Level Phase Change Memory

HanBin Yoon*, Naveen Muralimanohar†, Justin Meza*, Onur Mutlu*, Norman P. Jouppi†
*Carnegie Mellon University, †HP Labs

1 Summary

Emerging *Phase Change Memory PCM* technology offers better density over DRAM by its ability to store multiple bits in a single cell. However, MLC PCM exhibits higher access latency and energy compared to single level PCM. The overhead comes from the fine-granularity control of the crystallization process and sensing of the cell that is required to achieve MLC.

We identify characteristics of PCM device physics that allow the multiple bits of an MLC memory cell to be accessed at different latencies and energies. We advocate exposing these differences to the encompassing architecture to improve overall system performance and energy efficiency. To this end, we propose to map the bits of a cell to different logical pages ('decoupling' the bits of an MLC cell). Building on this, we also propose decoupling the row buffers to buffer these logical pages separately. The downside of decoupling is that write backs to the memory array now program more cells, consuming more endurance cycles and reducing PCM lifespan. To mitigate this problem, we propose a data mapping scheme that facilitates the coalescing of writes to the same cells.

To our knowledge, this is the first work to propose alternative data-to-cell mapping schemes for multi-level phase change memory.

2 Proposals

2.1 MLC PCM Bit Decoupling

Program asymmetry and read asymmetry are properties that certain bits of an MLC cell can be programmed and read at a lower latency and energy than the other bits.

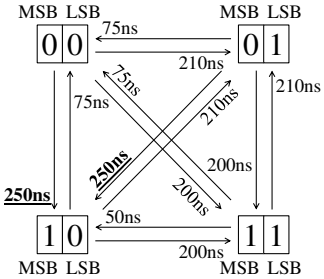

Figure 1: MLC PCM program time.

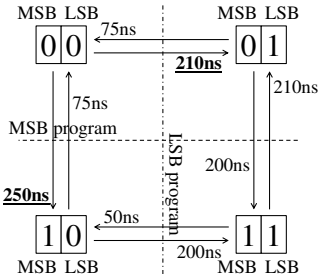

Figure 2: Restraining the alteration of both bits of a cell.

Program asymmetry. Figure 1 shows the latencies incurred in programming any one state of a 2-bit MLC PCM into any other [1]. The programming latency of a row of cells is limited by the highest latency (250 ns) among all the arcs in Figure 1. However, if we impose an artificial constraint that restricts altering both bits of a cell in a single programming operation (Figure 2), then we find that changing the LSB bit alone is bounded by a lower 210 ns latency. This finding shows that the program latency can be reduced for LSB-bit data if the MSB bits are not changed simultaneously.

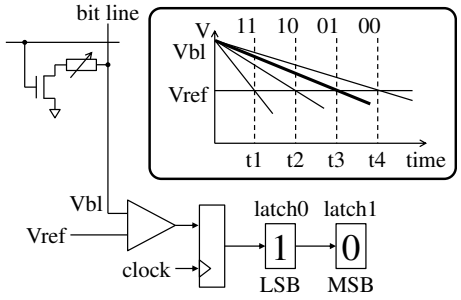

Figure 3: MLC PCM read [2, 3].

Read asymmetry. Figure 3 illustrates an MLC PCM read technique [2, 3]. The cell state is determined by measuring the number of clock cycles taken for the bit line's voltage to drop. For N-bit MLC, the read latency is 2^N clock cycles at the least. However, $(2^N/2)$ cycles after starting the read, it is possible to determine the MSB bit. If the cell state is determined by the $(2^N/2)$th cycle, the MSB bit is a '1', else the MSB bit is a '0' irrespective of the LSB bit. This observation can be exploited to reduce read latency for data mapped to MSB bits, compared to those mapped to LSB bits.

MLC PCM Bit Decoupling. Program asymmetry and read asymmetry in MLC PCM provide strong motivation for logically decoupling the two bits of a cell. By mapping the two bits to separate logical pages as shown in Figure 4, the reduced latencies of LSB bit programming and MSB bit reading can be exposed.

2.2 LSB/MSB Write Coalescing

Although the decoupled bit scheme offers fast access time and higher performance, its downside is that it exhibits worse endurance when compared to the coupled bit scheme. In order to program M bits, $M/2$ cells undergo the physical programming cycle in the coupled bit scheme, whereas M cells undergo physical programming in the decoupled bit scheme.

We propose mitigating this increase in cell wear by coalescing the writebacks to the two bits of a cell, such that the cell may be programmed only once instead of twice. We achieve this by interleaving cache blocks between the LSB

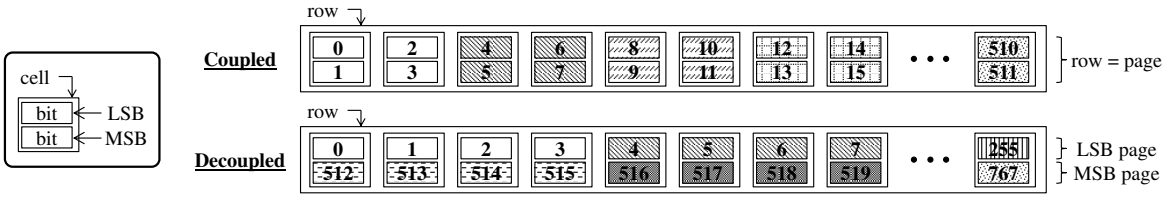

Figure 4: Coupled and decoupled bit schemes. For illustration, cache blocks the size of 4 bits are highlighted in different shades. We use 64 byte cache blocks in our evaluations.

and MSB pages of a row (every X blocks, which we denote as LMI-X). The spatial locality in writebacks is exploited to facilitate the coalescing of cache block writebacks to the same cells.

2.3 Split Page Buffering

By decoupling the bits of a cell and mapping them to separate logical pages, we can treat the row buffer as separate page buffers, as shown in Figure 5. One of the two latches for every cell is grouped together to form a page buffer, and the other latch for every cell is grouped together to form another page buffer. We find that this fine-granularity organization presents an opportunity for an increased row buffer hit rate (the fraction of memory accesses that are row buffer hits) compared to when using a single large row buffer, thereby achieving higher performance and energy efficiency.

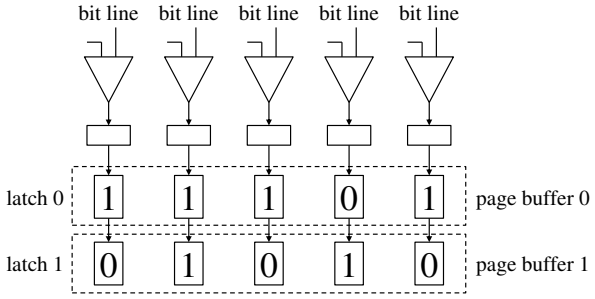

Figure 5: Treating the row buffer as two separate page buffers.

3 Results

Our evaluations show improved performance and energy efficiency for a system that employs MLC PCM memory in the decoupled bit scheme, compared to the coupled bit scheme. For SPEC CPU2006 benchmarks executed on a single core system, bit decoupling reduces average memory service time by 19.9% and increases IPC by 12.6%, while reducing memory energy consumption by 13.1%. Split page buffering further increases the performance improvement to 30.3% and memory energy reduction to 31.5%. The lifetime of the LMI-1 scheme comes within 24.7% of the coupled bit scheme. An 8-core evaluation (Figure 6) shows 12.7% improvement in weighted speedup and 12.9% improvement in performance per Watt in the case of the decoupled scheme relative to the coupled scheme.

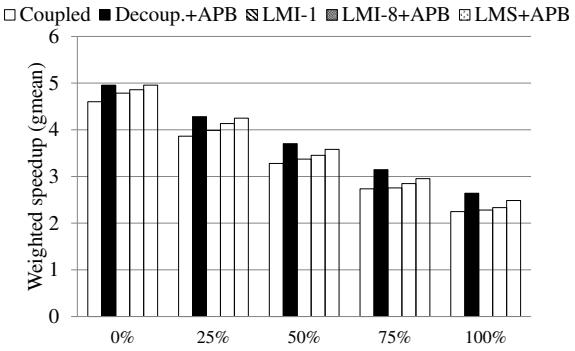

(a) Weighted speedup.

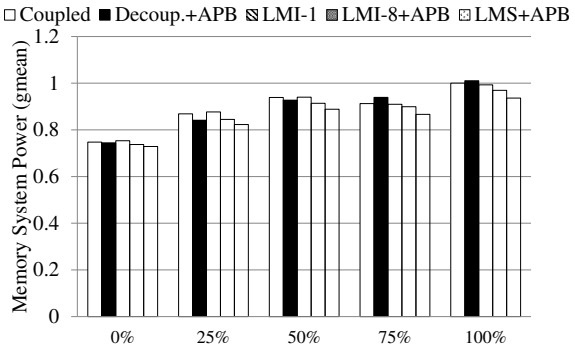

(b) Memory system power (normalized to the coupled bit scheme at 100% mem-intensity).

Figure 6: 8-core system evaluation. The x-axes indicate the proportion of memory-intensive (last level cache misses per kilo instructions > 5) benchmarks in a workload. LMS (LSB/MSB-Serial) is LMI-X with the parameter X at maximum. APB = Split page buffering.

References

[1] M. Joshi et al. Mercury: A fast and energy-efficient multi-level cell based Phase Change Memory system. In *HPCA*, Feb. 2011.

[2] J. Li et al. A Novel Reconfigurable Sensing Scheme for Variable Level Storage in Phase Change Memory. In *Memory Workshop (IMW), 2011 3rd IEEE International*, 2011.

[3] M. K. Qureshi et al. Morphable memory system: a robust architecture for exploiting multi-level phase change memories. ISCA '10, New York, NY, USA, 2010. ACM.