Enabling Efficient and Scalable Hybrid Memories Using Fine-Granularity DRAM Cache Management

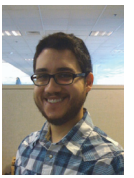
Justin Meza


HanBin Yoon


Onur Mutlu

As feature sizes continue to shrink, future chip multiprocessors are expected to integrate more and more cores on a single chip, increasing the aggregate demand for main memory capacity. Satisfying such a demand with DRAM alone may prove difficult due to DRAM scaling challenges. To address this problem, recent work has proposed using DRAM as a cache to large non-volatile memories, such as phase-change memory (PCM), which are projected to be much more scalable than DRAM at comparable access latencies. A key challenge in scaling hybrid main memories is how to efficiently manage the metadata (such as tag, LRU, valid, and dirty bits) for data in such a large DRAM cache at a fine granularity.

Most prior hardware-based approaches toward large, fine granularity DRAM caches have either (1) stored metadata for each cache block in a large SRAM structure, limiting scalability and increasing cost, or (2) stored metadata in a contiguous region of DRAM, requiring additional accesses and reducing performance. Our goal is to achieve minimal performance degradation compared to large on-chip SRAM metadata structures, but with orders of magnitude lower hardware overhead.

Based on the observation that storing metadata off-chip in the same row as their data exploits DRAM row buffer locality, this paper reduces the overhead of fine-granularity DRAM caches by only caching the metadata for recently accessed rows on-chip using a small buffer. Leveraging the flexibility and efficiency of such a fine-granularity DRAM cache, we also develop an adaptive policy to choose the best granularity when migrating data into DRAM. On a hybrid memory with a 512MB DRAM cache, our proposal using an 8KB on-chip buffer can achieve within 6% of the performance of (Fig. 1), and 18% better energy efficiency than (Fig. 2), a conventional 8MB SRAM metadata store, even when the energy overhead due to large SRAM metadata storage is not considered.

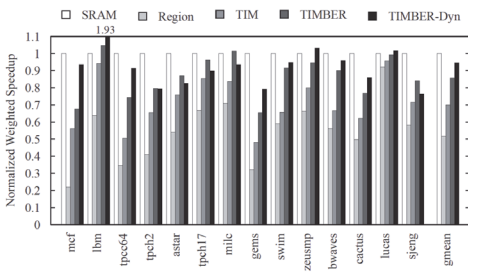

Fig. 1: Multi-core performance (8 copies of each application).

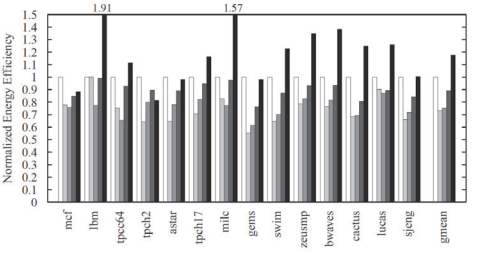

Fig 2: Multi-core energy efficiency (performance per Watt).