Enhancing Phase-Change Memory via DRAM Cache

Brian Pak, Lincoln Roop

Goals

- Increase the lifetime of Phase-Change Memory by reducing the number of writes to PCM
  - PCM cell write/erase endurance: ~10^7
  - Limits PCM from being sole main memory
- Guaranteed lifetime for the worst case (e.g. wear-out attacks)
  - Protects PCM lifetime with minimal performance degradation

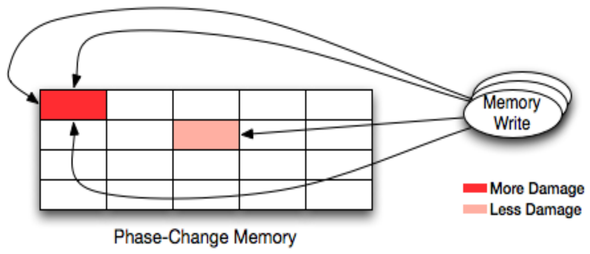

PCM suffers wear-out problem with multiple writes to the same location

With DRAM write-back cache, unnecessary writes to PCM can be avoided

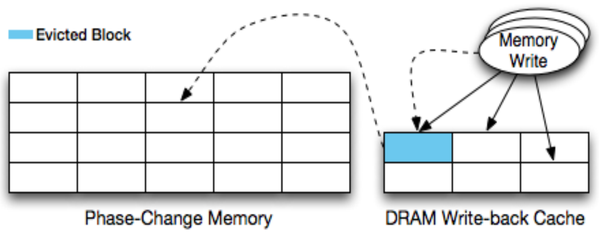

Approach

- Introduce an extra level of indirection: DRAM Cache
- More robust and wear-out-resistant cache replacement policy
  - *Access Density Decay (ADD)* Policy
    - Keep track of the <u>number of accesses</u> per block (A) – *Access Density*
      - A = (α x # Writes) + (β x # Reads), where $\alpha = weight_{write}, \beta = weight_{read}$
    - *<u>Decay</u>* access density by half for every *N* instructions
    - For *eviction*, a block with the <u>least access density</u> is evicted to PCM
- Employ wear-leveling to evenly distribute the writes across PCM (when a block is evicted from the cache)

Experiments and Results

- Experiment I
  - PCM Lifetime comparison
    - PCM Only
    - PCM + DRAM Cache (LRU)
    - PCM + DRAM Cache (ADD)
- Result

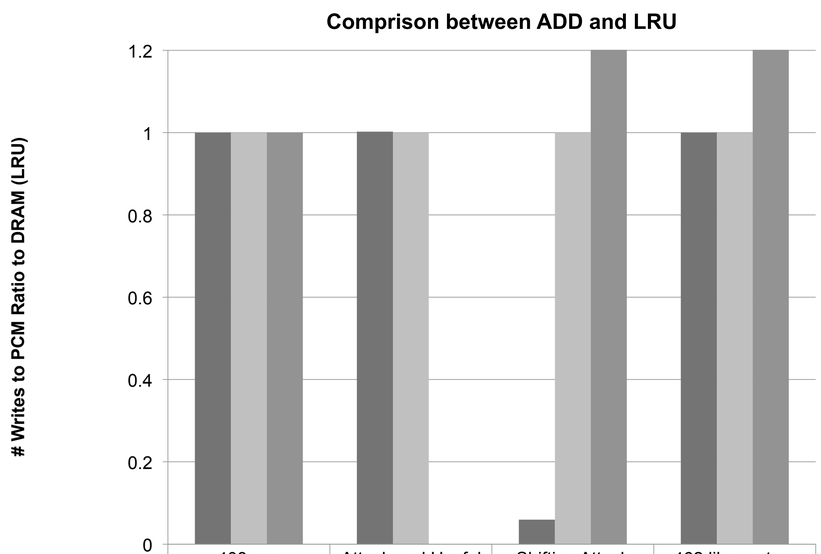

	403.gcc	Attack and Useful	Shifting Attack	462.libquantum
DRAM Cache 1MB (ADD)	1	1.00254399	0.059665871	1
DRAM Cache 1MB (LRU)	1	1	1	1
ALL PCM	1	0	3468.854893	43.21399062

- Experiment II
  - ADD Parameter (decay cycle) comparison
    - Decay Cycle: 200K, 500K, 1M, 2M instructions
- Result

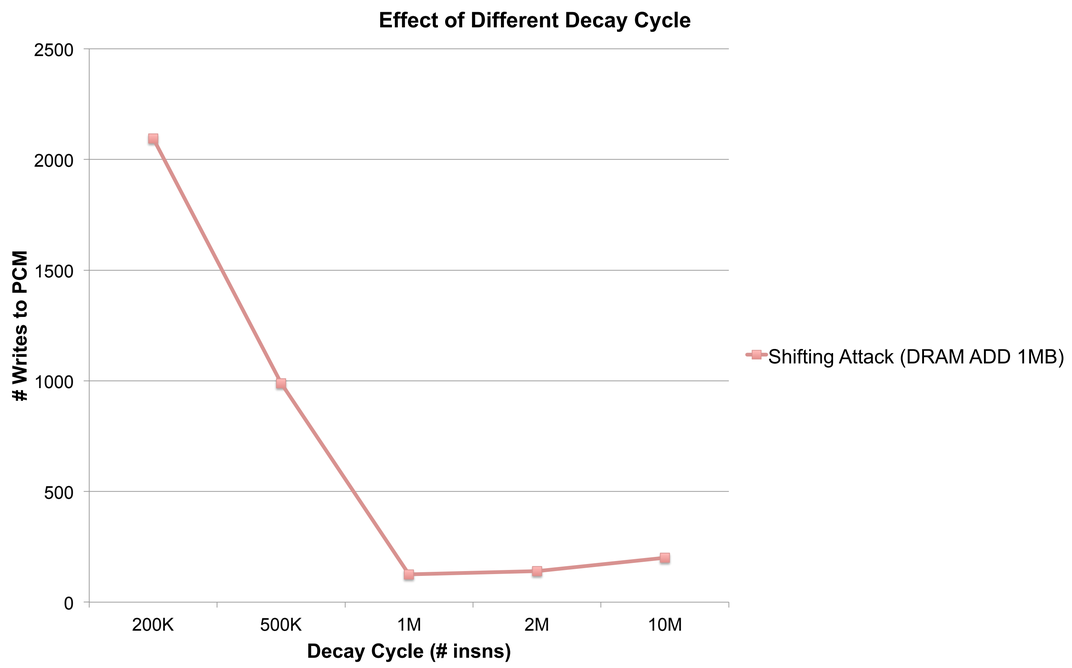

Carnegie Mellon University