Low Power Hybrid PRAM/DRAM Main Memory

Sungjoo Yoo

Department of Electrical Engineering
Pohang University of Science and Technology
(POSTECH)

sungjoo.yoo@postech.ac.kr

Agenda

- Our problem
  - DRAM refresh energy issue
- Related work
- Proposed idea
  - In hybrid PRAM/DRAM main memory subsystem, DRAM decay with runtime adaptive time out control
- Experiments
- Summary

Our Problem

- DRAM suffers from large refresh energy consumption
- DRAM refresh can consume more than half of DRAM energy in hybrid PRAM/DRAM-based main memory

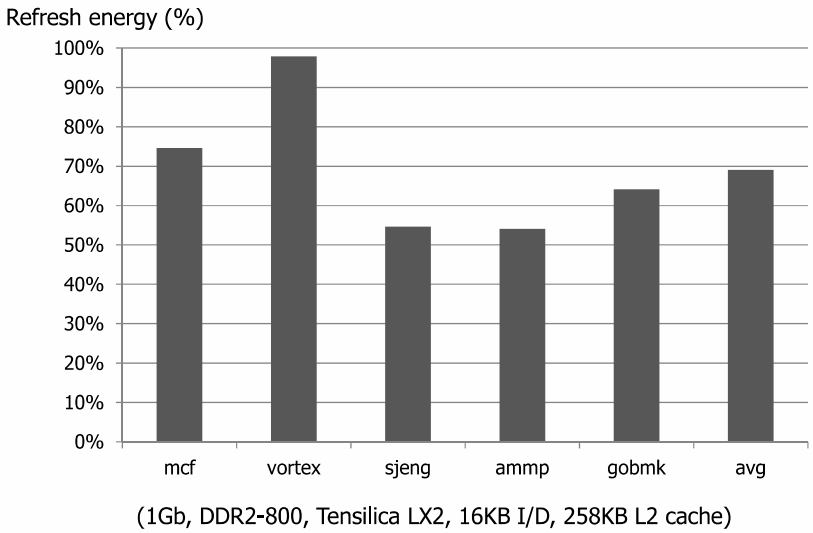

(1Gb, DDR2-800, Tensilica LX2, 16KB I/D, 258KB L2 cache)

Related Work: Low Power Main Memory

- **Dynamic power management utilizing low power modes**
  - Active/standby power down [Micron Technology, Inc.]
  - Power-aware memory access scheduling to achieve longer idle period [I.Hur, HPCA08]
  - Memory controller predicts idle periods in main memory [V.Delaluz, HPCA01] [X.Fan, ISLPED01]
- **DRAM refresh power reduction**
  - Minimizing the number of rows to be refreshed [M.ghosh, MICRO07] [C.Wilkerson, ISCA10]
  - Maximizing refresh period (e.g., 128ms or 256ms) [J.Kim, VLSI03] [R.K.Venkatesan, HPCA06] [T.Hamamoto, '98] [Y.Idei, '98]
- **Non-volatile memory together with DRAM**
  - DRAM + NAND Flash memory, e.g., Pico-server [T. Kgil, JETC08]
  - DRAM + Phase-change RAM (PRAM) [M. Qureshi, ISCA09]
  - **Our proposal in this work**

Hybrid PRAM/DRAM-based Main Memory

- **PRAM: large background main memory**
  - Low standby energy
- **DRAM: last-level cache with decay**
  - Enables lower latency access than PRAM

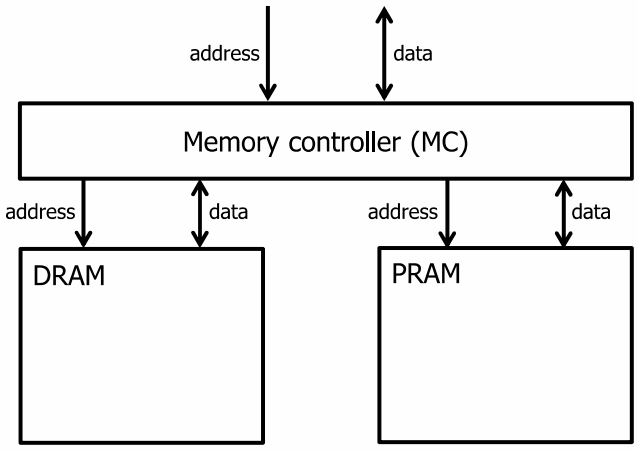

Hybrid PRAM/DRAM-based Main Memory

- **In case of long idle time**
  - All data move from DRAM to PRAM
  - DRAM is turned off
  - Standby energy, e.g., DRAM refresh is avoided

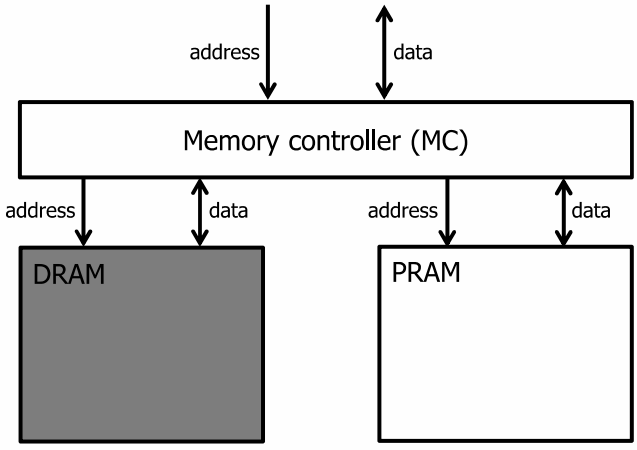

Our Contribution

- **DRAM refresh reduction during runtime**
  - Turn off a part of DRAM by evicting **cold** data from DRAM to PRAM
  - DRAM is still functional providing **hot** data

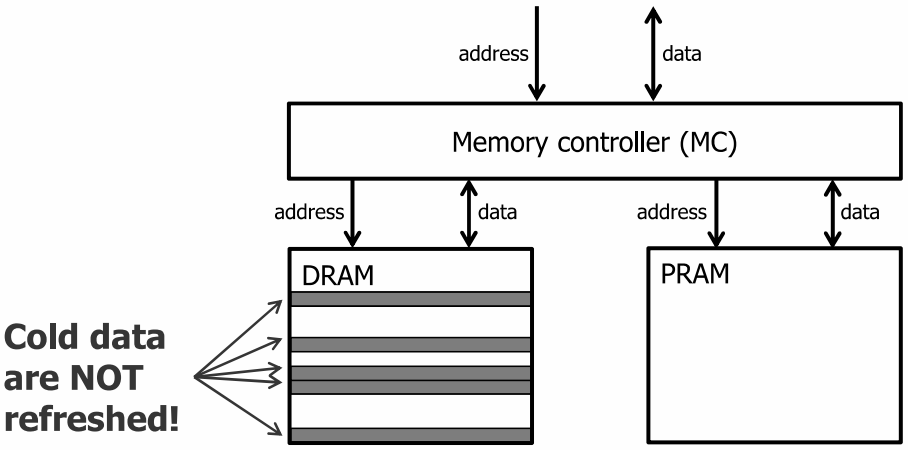

Time Out-based Hot/Cold Management

- On each reference, the counter (per DRAM row) is set to a time out (TO) value
- Periodically, the counter decrements
- If a counter becomes zero, the row is cold and evicted

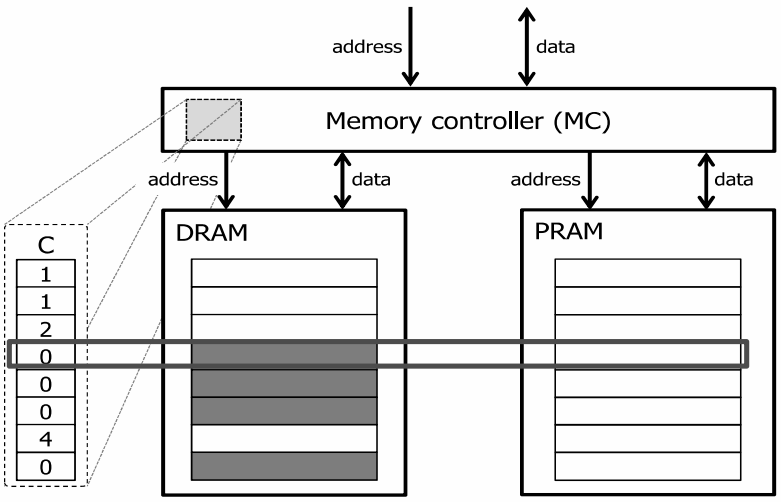

Runtime-Adaptive Time Out Control

- **TO vs Energy**
  - TO increase ➔ DRAM energy↑, PRAM energy↓
  - We can find TO_{best} which gives the minimum total energy, E_{min}

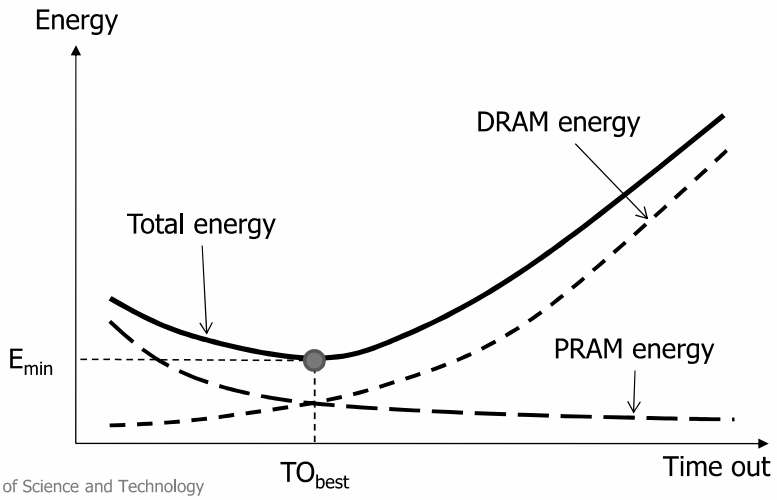

Runtime-Adaptive Time Out Control

- **TO vs Performance overhead**
  - Given performance bound, OV_{max}, we need to find the time out, TO_{best_ov} for minimum energy, E_{min_ov}

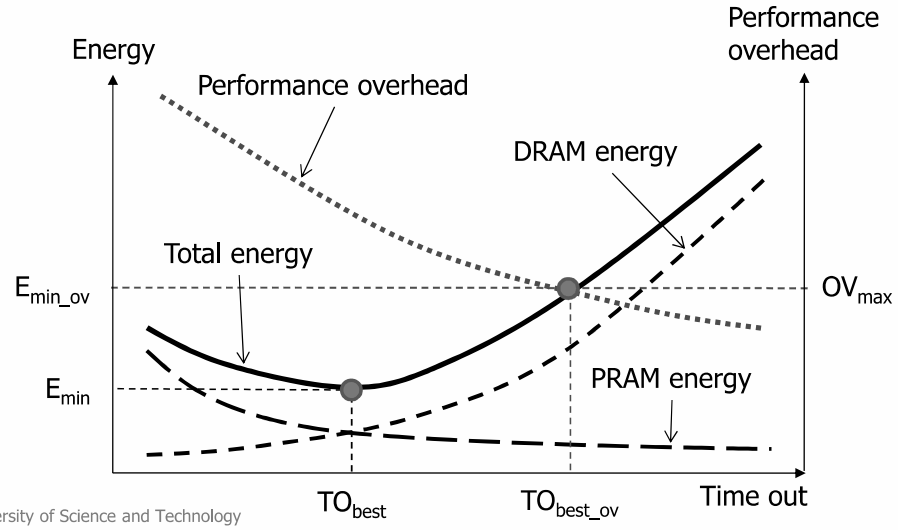

Sampling-based Time Out Control

- **Time out is adjusted to dynamically changing behavior**
- **Check group (without DRAM decay)**
  - Gives baseline performance
  - If the main group gives performance violation, then TO is increased
- **Alternative groups (TO+1 and TO-1 groups)**
  - Provide the energy gain of alternatives → periodical comparison of energy consumption, and select the best TO for the next period

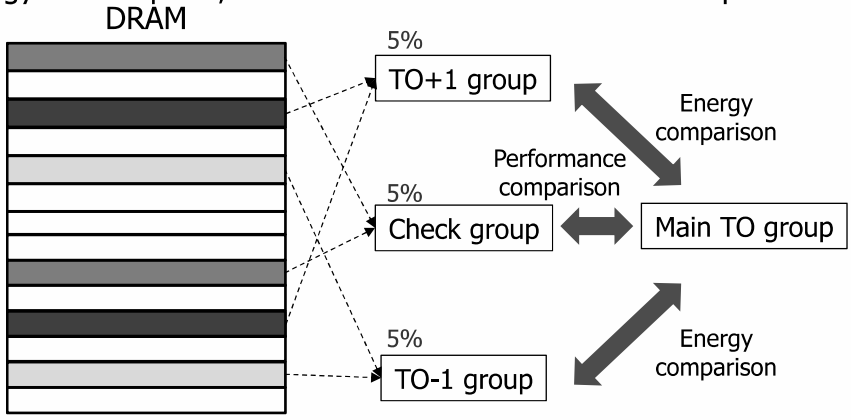

Target Architecture

- Cycle-accurate SystemC models
  - L2 cache and hybrid main memory

Component	Details
CPU core	Tensilica LX2 (7 stage pipeline), 32b address, 64b data, 4-way 16KB I/D, 400MHz
L2 cache	16-way 256KB, I/D shared, 64B cache line, 400MHz, 5 pipe stage (input buffering, tag retrieval, tag match, data access, and output)
Memory controller	FR-FCFS policy, closed scheme
DRAM	1Gb, 32b DDR2-800, $t_{CL}/t_{RP}/t_{RCD}$ = 12.5ns/12.5ns/12.5ns
PRAM	1Gb, $t_{CL}/t_{RP}/t_{RCD}$ = 12.5ns/160ns/55ns, Row buffer size = 64B

- Benchmarks from SPEC2000 and SPEC 2006

Experimental Results: Static TO

- Trade-off between energy and performance overhead
- 23.5% ~ 94.7% reduction in the energy consumption

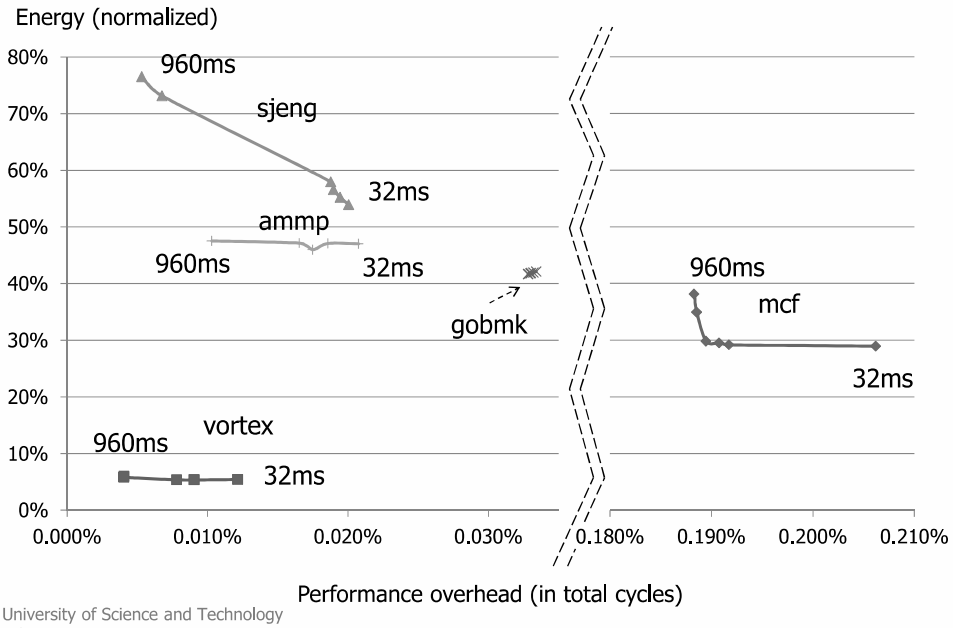

Sampling-based Time Out Control

- TO_{best_ov} is tracked during runtime under performance bound (5% in memory access latency)
- Fast tracking of program phase change

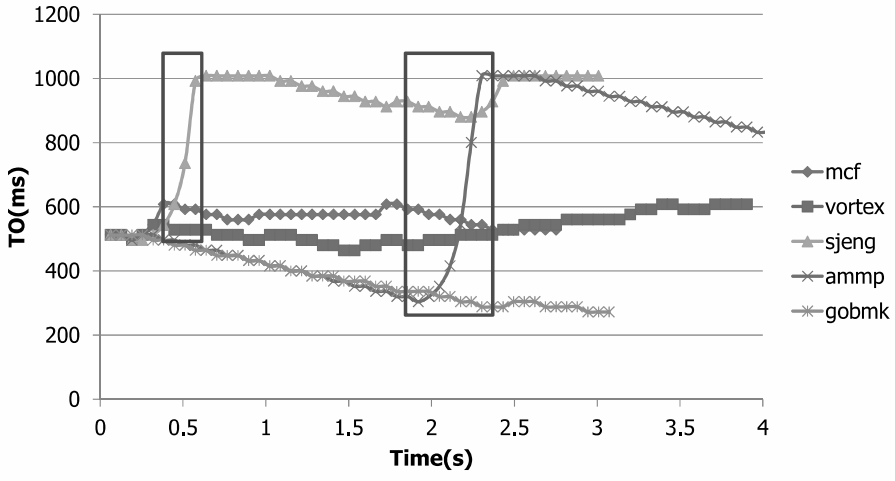

Experimental Results

- Runtime adaptive TO gives comparable results to static TO which requires design-time optimization

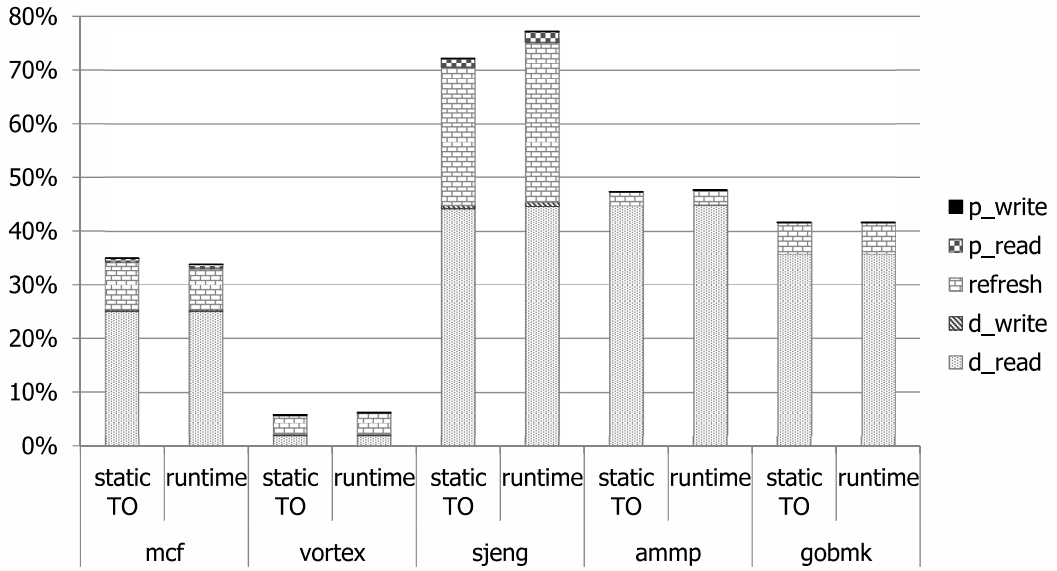

Summary

- Problem
  - **Refresh energy in DRAM** occupies a significant portion of total energy consumption in main memory
- Our idea
  - In hybrid PRAM/DRAM main memory subsystem
  - **DRAM decay based on hot/cold management**
  - Runtime-adaptive time out control to obtain the minimum energy consumption
- Result
  - Reduction in energy consumption of main memory by 23.5%~94.7%