Using STT-RAM to Enable Energy-Efficient Near-Threshold Chip Multiprocessors

Xiang Pan and Radu Teodorescu
Computer Architecture Research Lab, The Ohio State University

Near-Threshold CMP Design Challenges

Problem: Increasing core counts in CMPs → "Power Wall" → "Dark Silicon"

Solution: Near-Threshold Computing

Challenges: **(1)** performance degradation; **(2)** amplified process variation; **(3)** leakage power dominates.

Leakage power dominates in NT-region

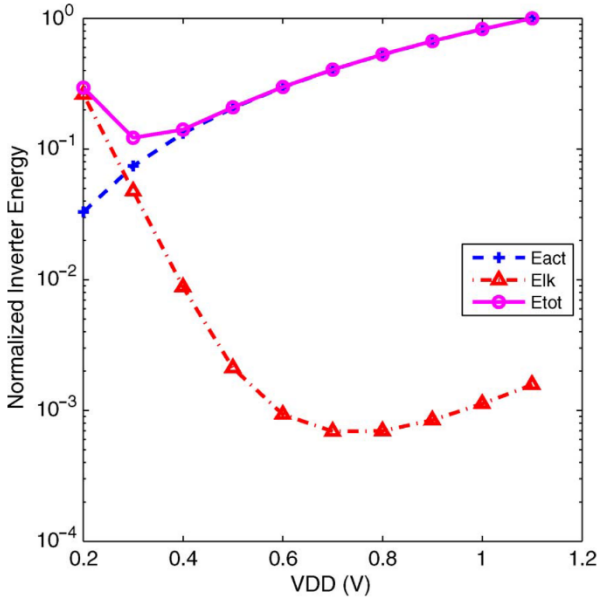

Cache power is mainly leakage dominated

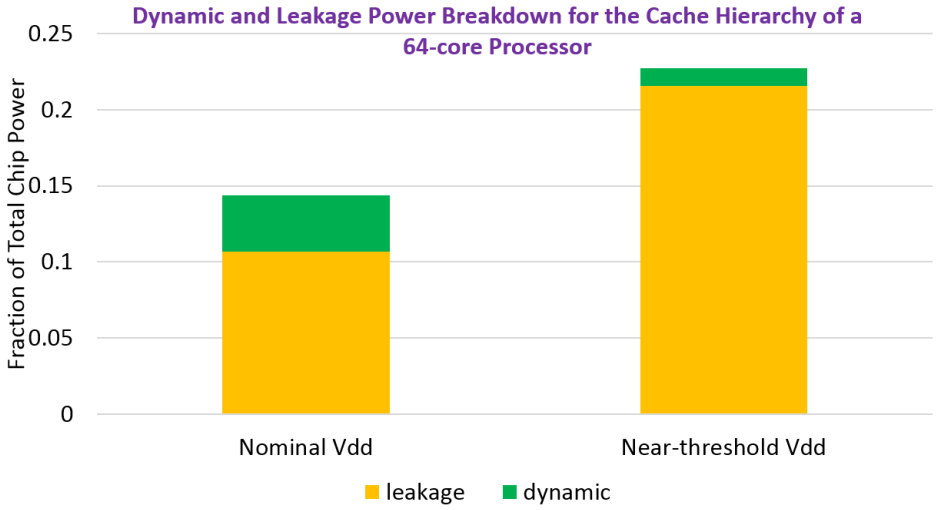

STT-RAM Very Good Fit for NT-CMP Designs

STT-RAM (Spin Transfer Torque RAM) Structure

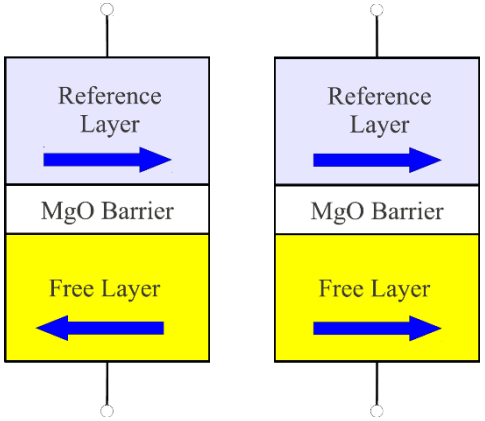

MTJ (Magnetic Tunnel Junction)

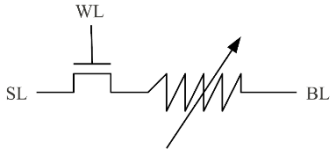

STT-RAM Cell (1T-1MTJ)

NT-CMP	STT-RAM
High leakage power	Near-zero leakage power
Low operating frequency	Long latency writes
Large numbers of cores requiring high cache capacity	High density (~4x denser than SRAM)

STT-RAM Based NT-CMP Design

Key Ideas:

- Use STT-RAM for all caches:
  - STT-RAM @ High-Vdd, cores @ NT-Vdd → fast cache reads.
  - Clustered-CMP in which the L1 caches are **shared** within each cluster, removing coherence costs.
- Address variation-induced performance heterogeneity:
  - Shared L1 cache enables low overhead within-cluster thread migration.
  - Periodically migrating threads between "fast" and "slow" cores achieves homogeneous CMP performance.

Proposed design reduces power consumption by 7%,
improves performance by 17%, and reduces energy by 21%.