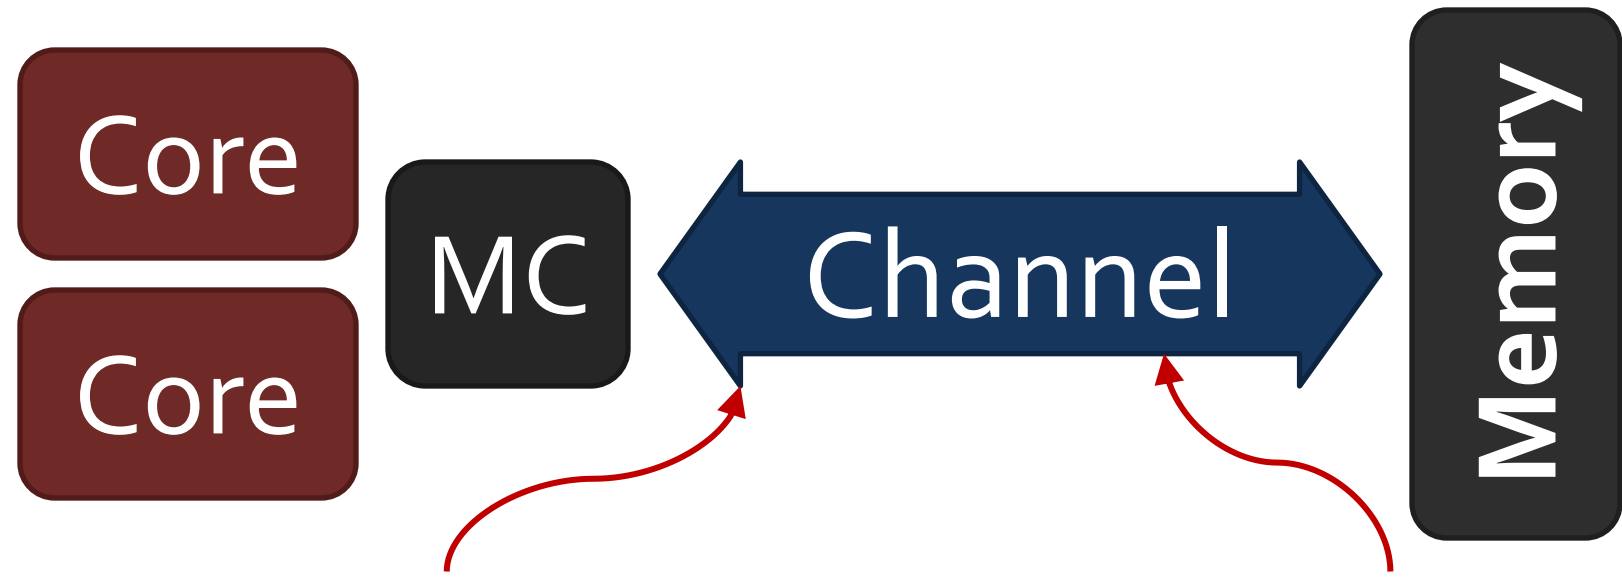


RowClone

Fast and Energy-Efficient In-DRAM Bulk Data Copy and Initialization


Vivek Seshadri, Yoongu Kim, Chris Fallin, Donghyuk Lee, Rachata Ausavarungnirun, Gennady Pekhimenko, Yixin Luo, Onur Mutlu, Phillip B. Gibbons*, Michael A. Kozuch*, Todd C. Mowry
Carnegie Mellon University *Intel Pittsburgh

Memory Bandwidth Bottleneck



Limited Bandwidth High Energy

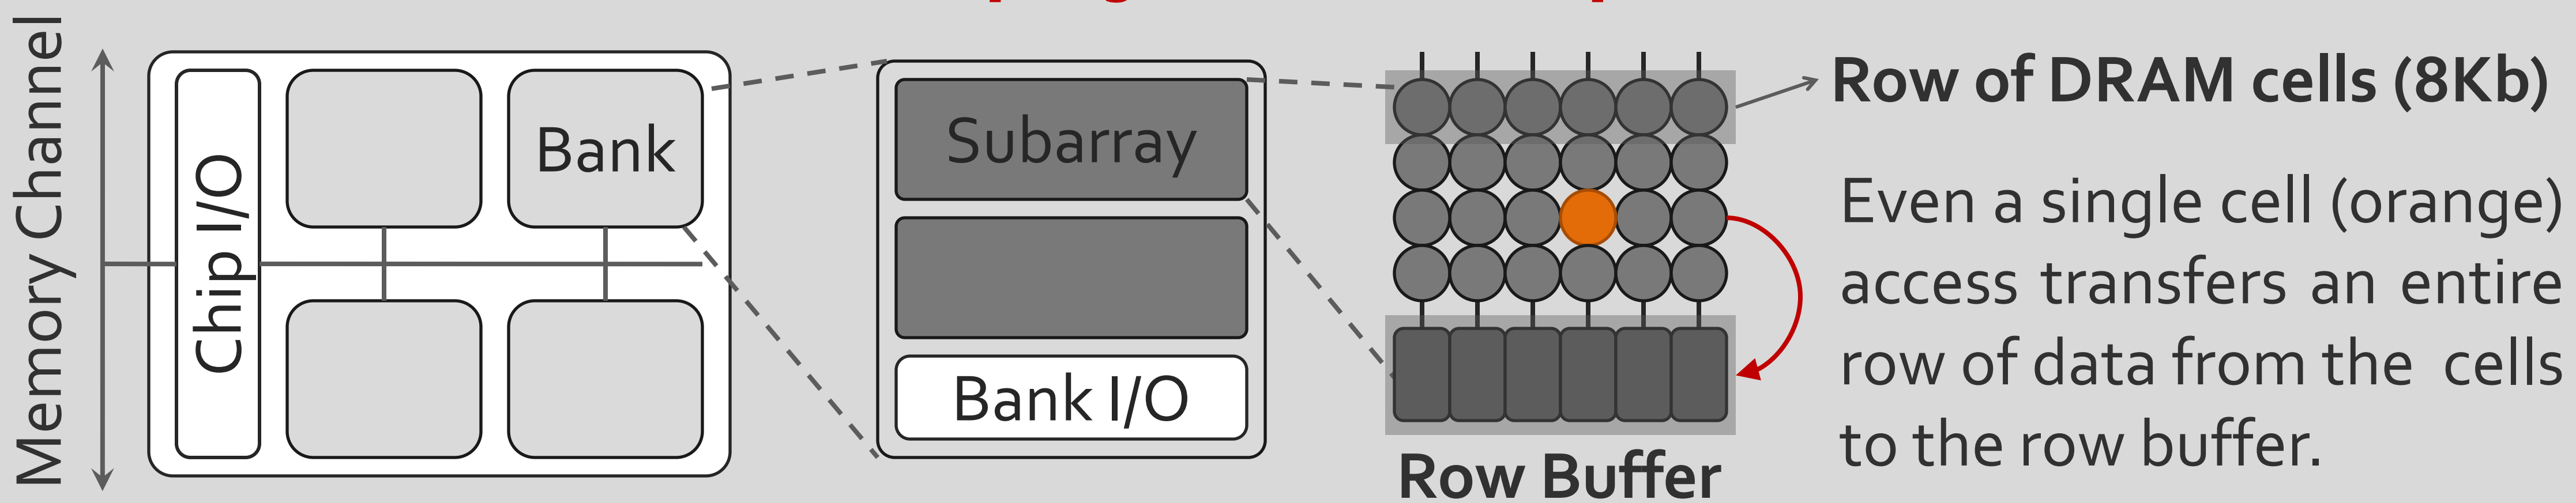
Bulk Copy and Initialization



- › Triggered frequently by many applications
- › Consume high latency, bandwidth, and energy
- › Do not require any computation

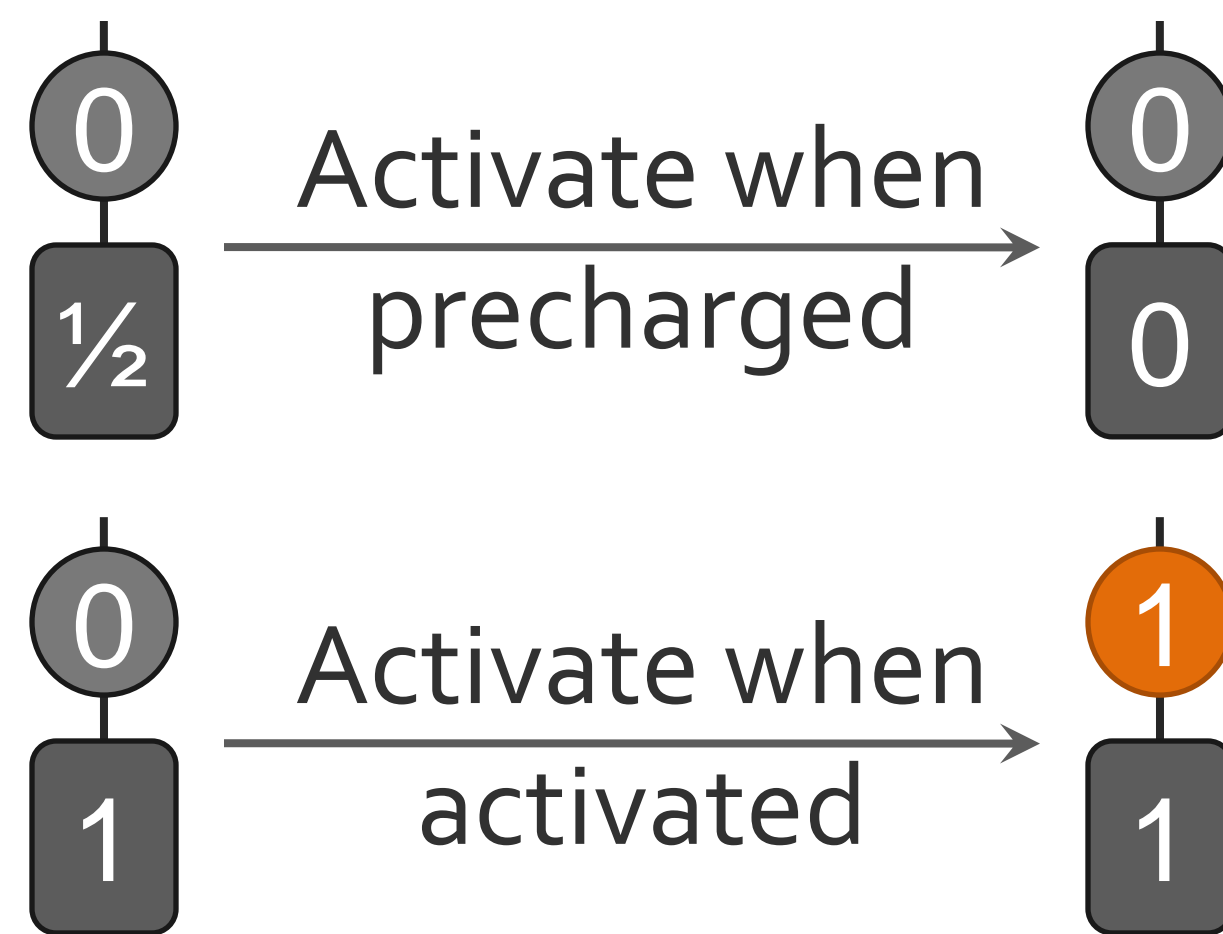

Our Approach: Perform them in DRAM

DRAM Chip Organization and Operation




RowClone – Fast Parallel Mode (FPM)

2-step row-to-row copy



+ **11.6X** latency reduction, **74.4X** energy reduction
– src and dst in same subarray, only full row copy

Pipelined Serial Mode (PSM)



- › Bank-to-bank cacheline copy
- › Overlap read/write using shared bus
- › 1.9X latency, 3.2X energy reduction

Overall DRAM area cost = **0.01%**

System Design


- › ISA: memcpy and meminit
- › μ Arch: manage coherence
- › OS: smart page mapping

Primitives and Applications Accelerated by RowClone

Copy-on-Write



Checkpoint



Page zero

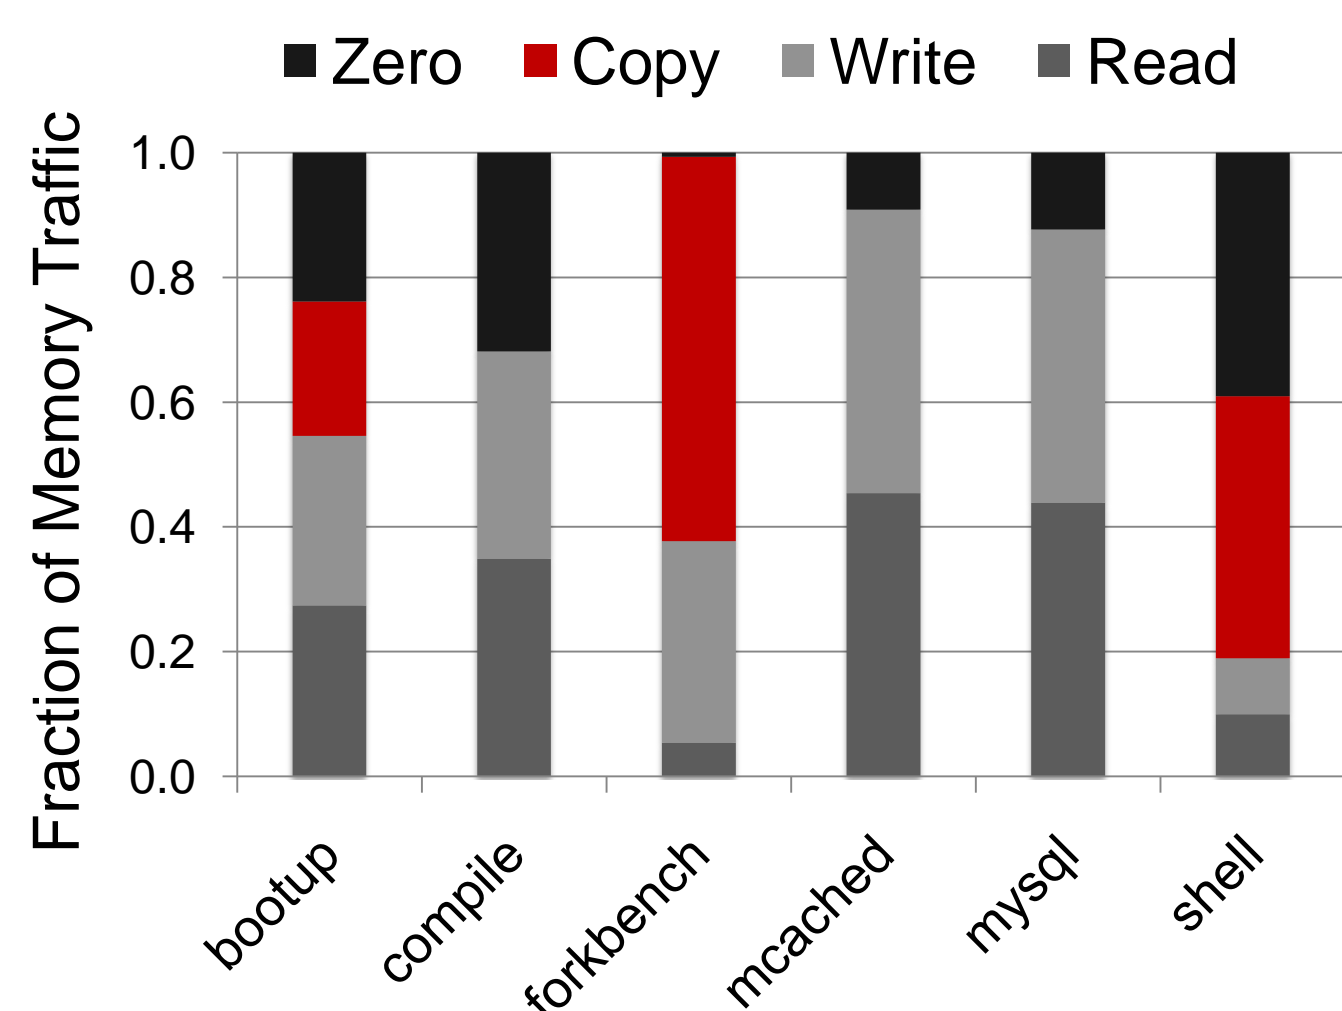
Bulk Zeroing



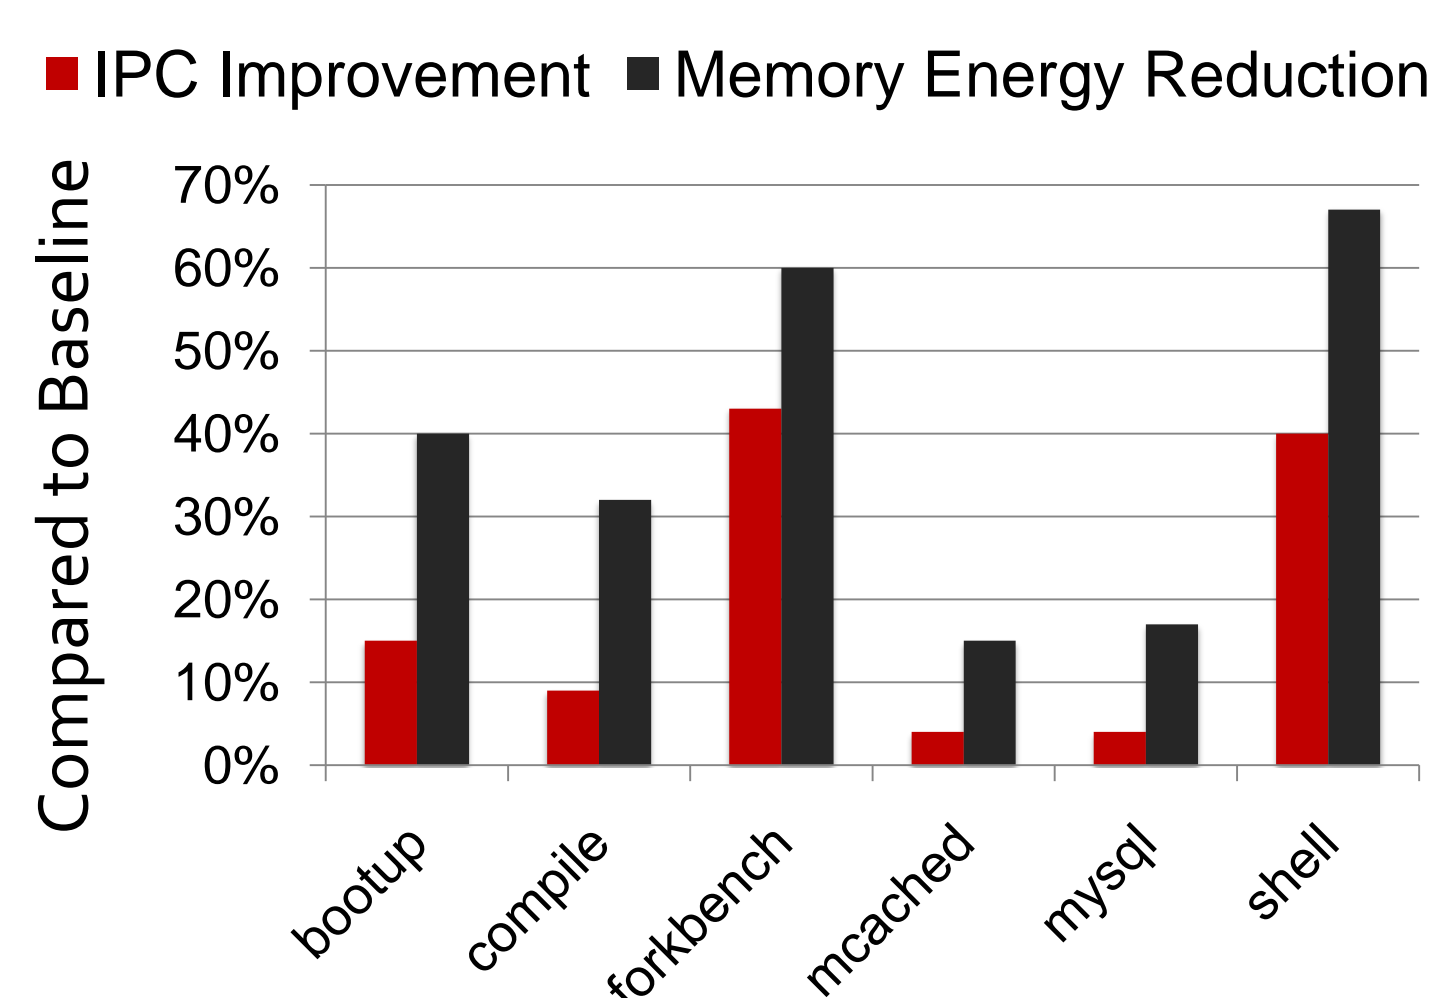
VM Cloning



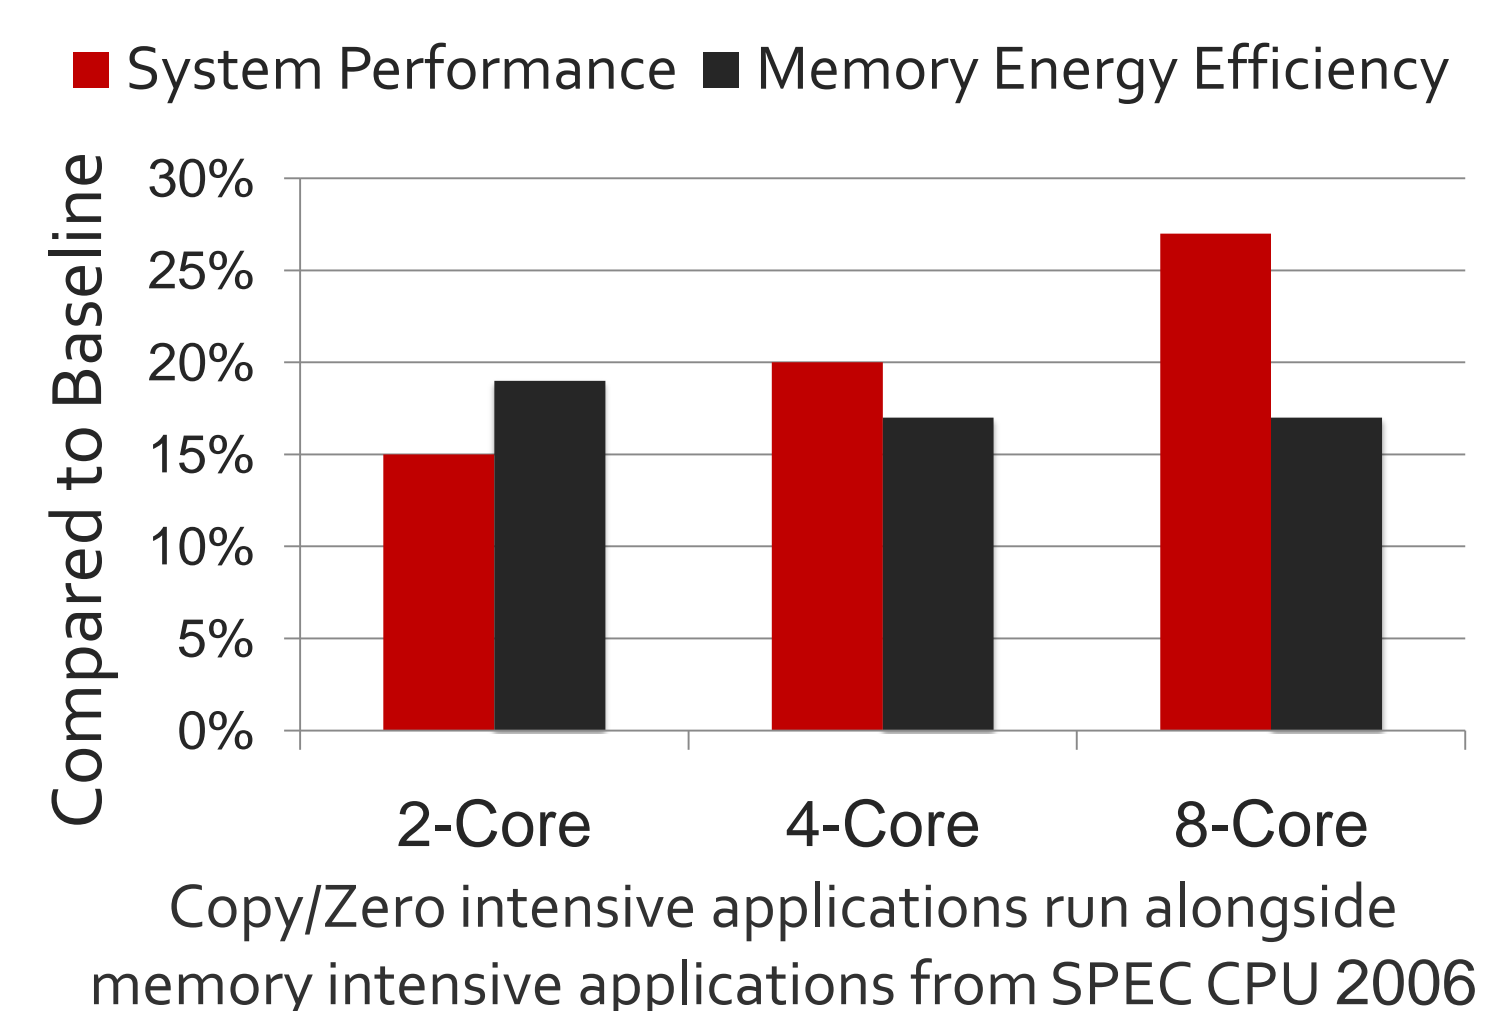
Copy/Zero Intensive Apps



Single Core Results



Multi-Core Results



Copy/Zero intensive applications run alongside memory intensive applications from SPEC CPU 2006