DRAM Tutorial

18-447 Lecture

Vivek Seshadri

DRAM Module and Chip

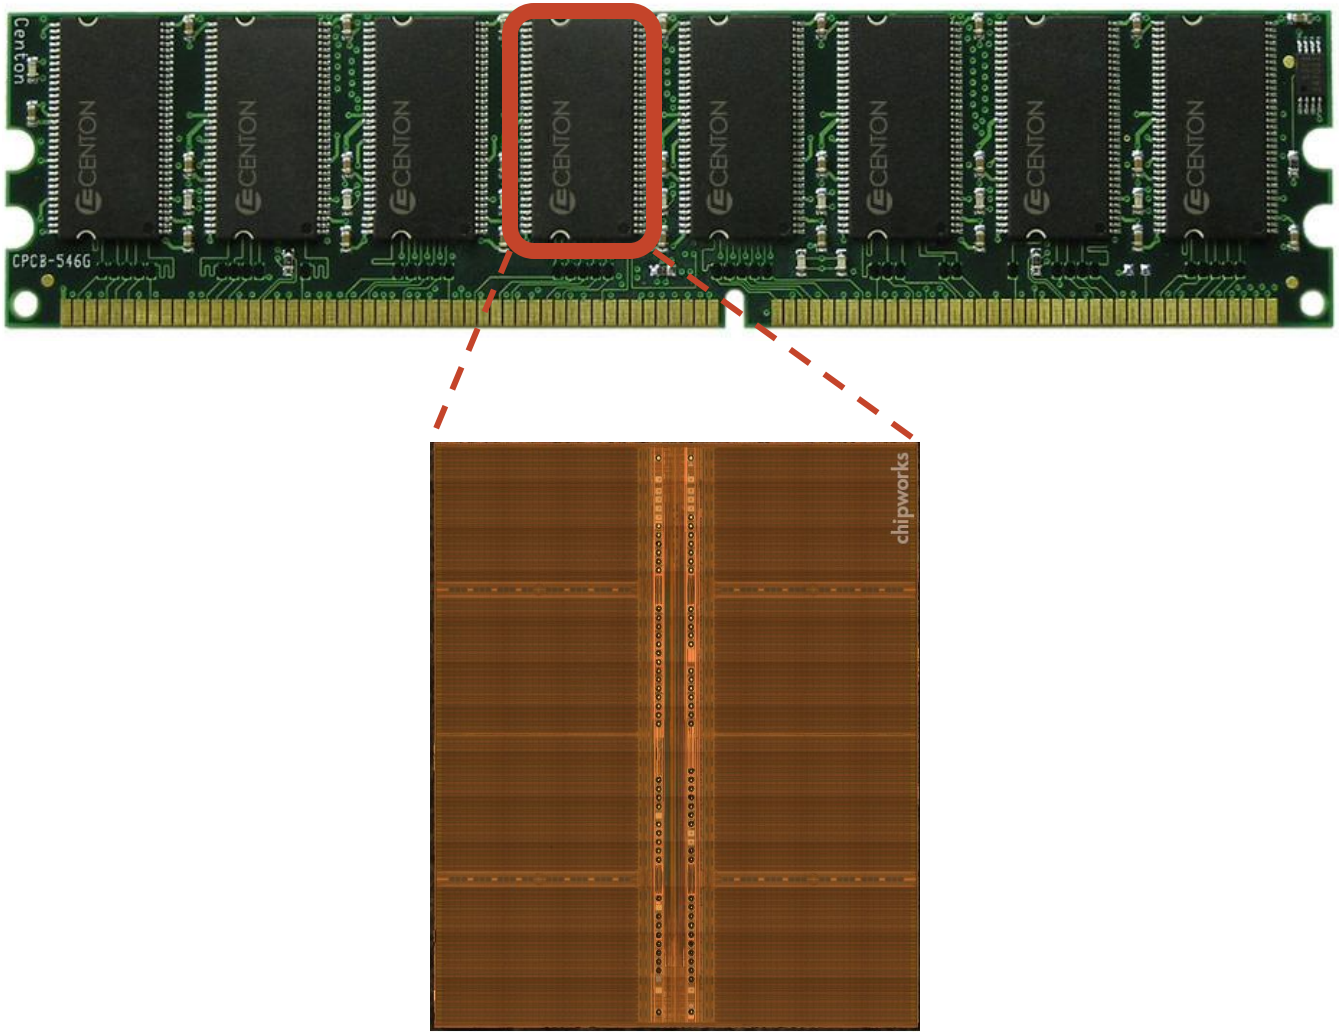

Goals

- Cost
- Latency
- Bandwidth
- Parallelism
- Power
- Energy

DRAM Chip

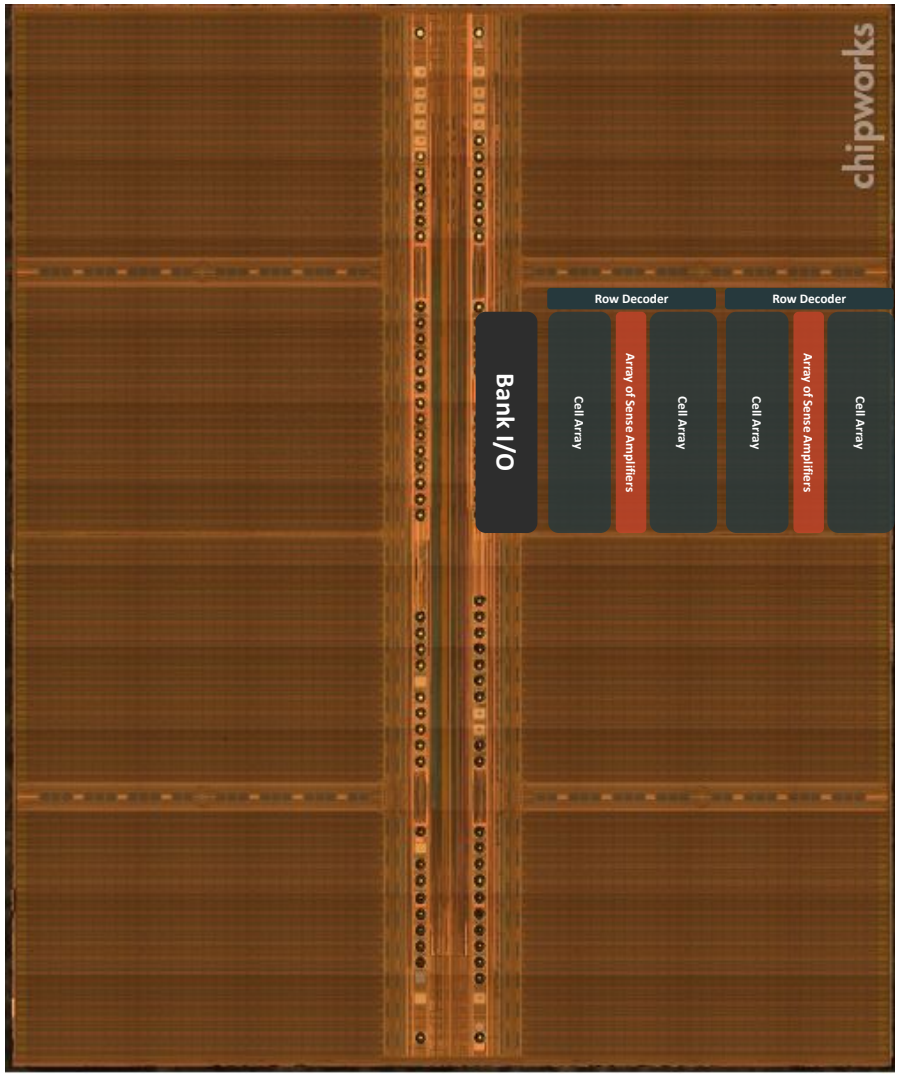
chipworks
Row Decoder
Row Decoder
Bank I/O
Cell Array
Array of Sense Amplifiers
Cell Array
Cell Array
Array of Sense Amplifiers
Cell Array

Sense Amplifier

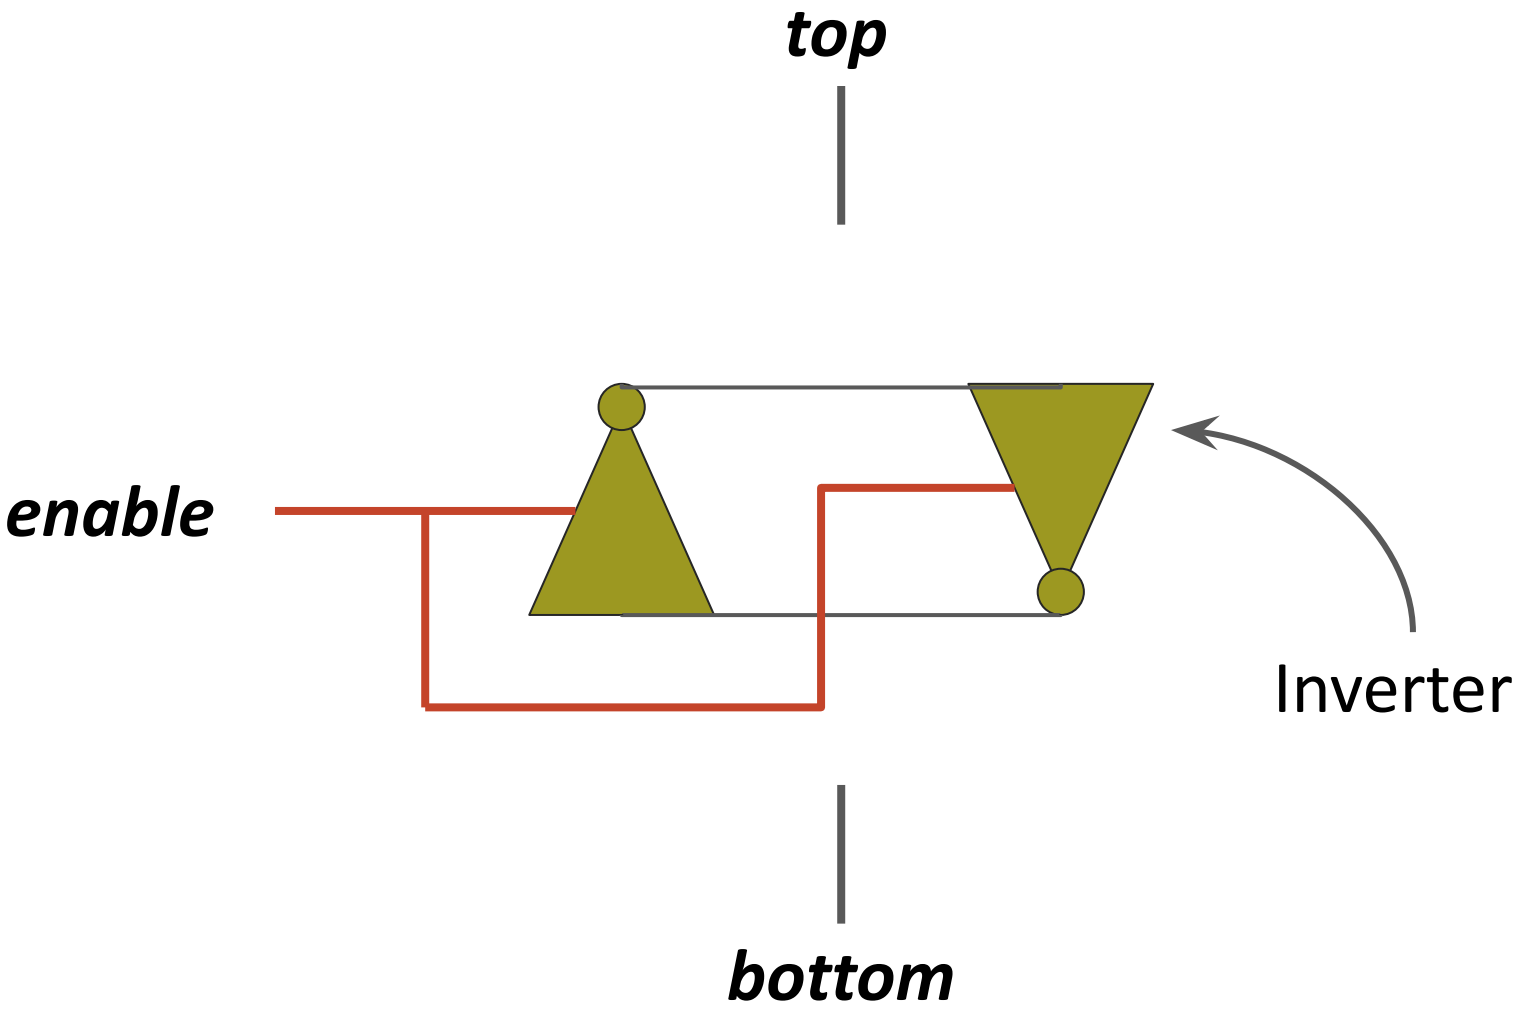

Sense Amplifier – Two Stable States

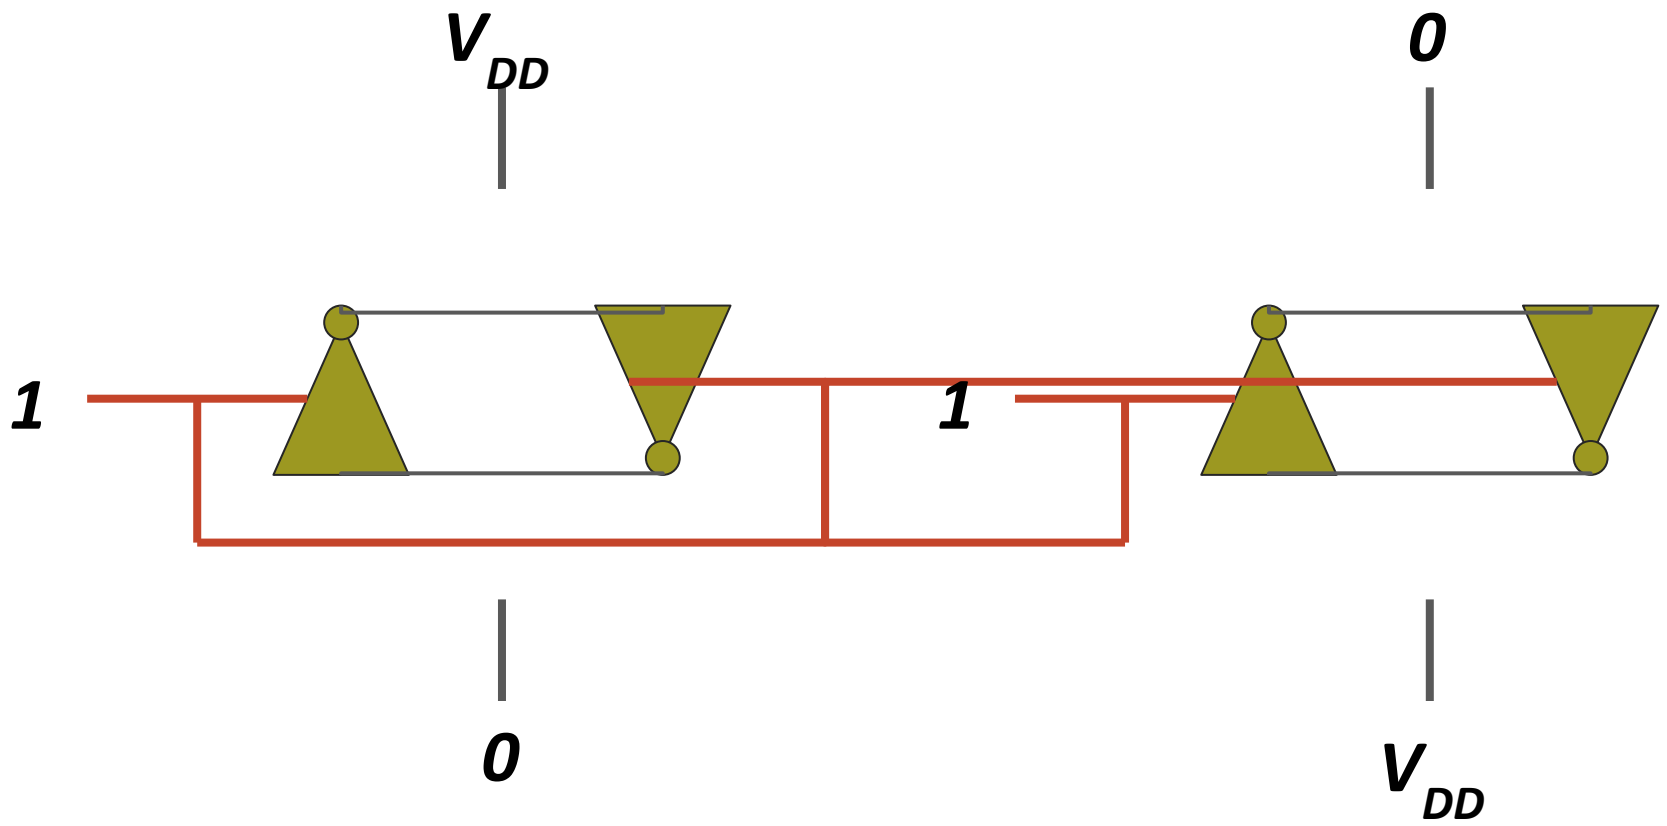

Logical “1”

Logical “0”

Sense Amplifier Operation

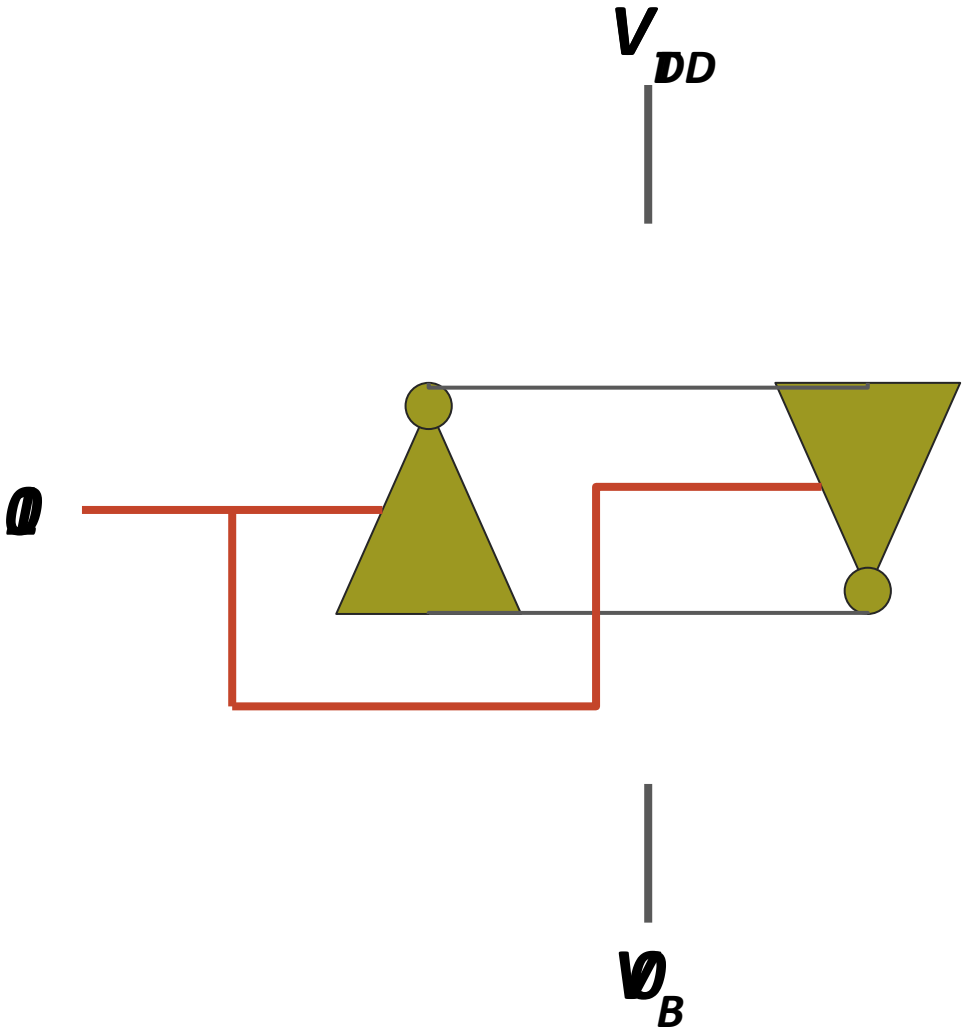

$V_T > V_B$

DRAM Cell – Capacitor

Empty State

Logical “0”

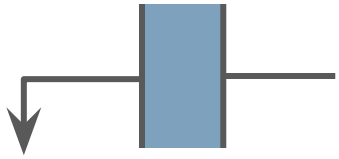

Fully Charged State

Logical “1”

1. Small – Cannot drive circuits
2. Reading destroys the state

Capacitor to Sense Amplifier

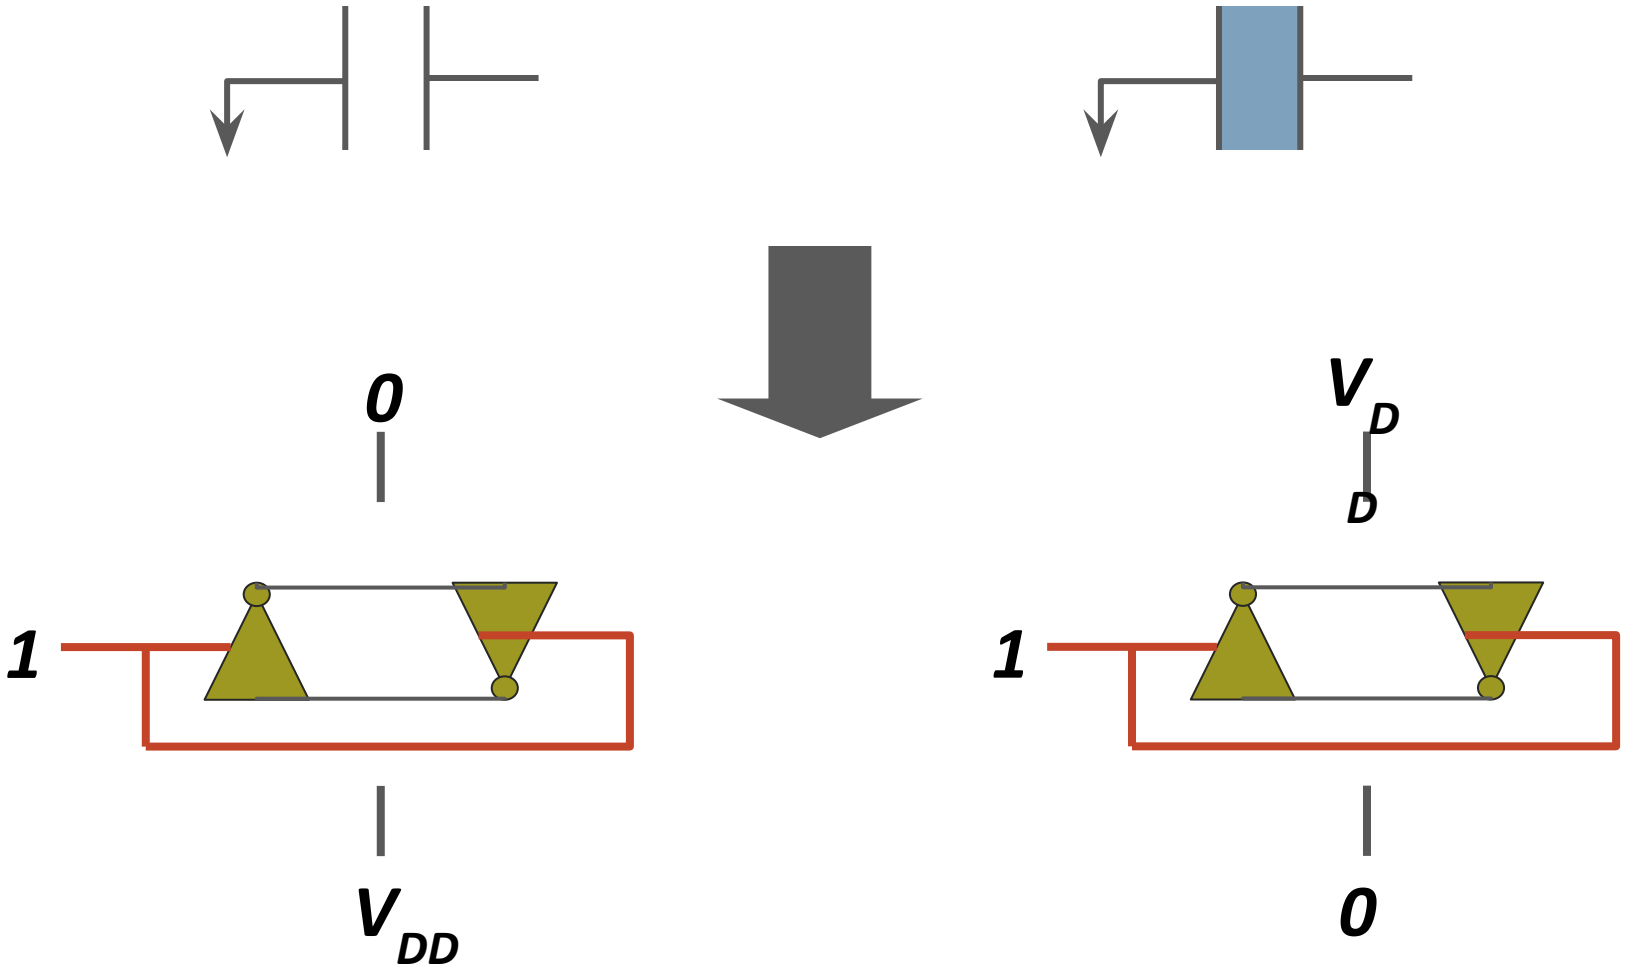

DRAM Cell Operation

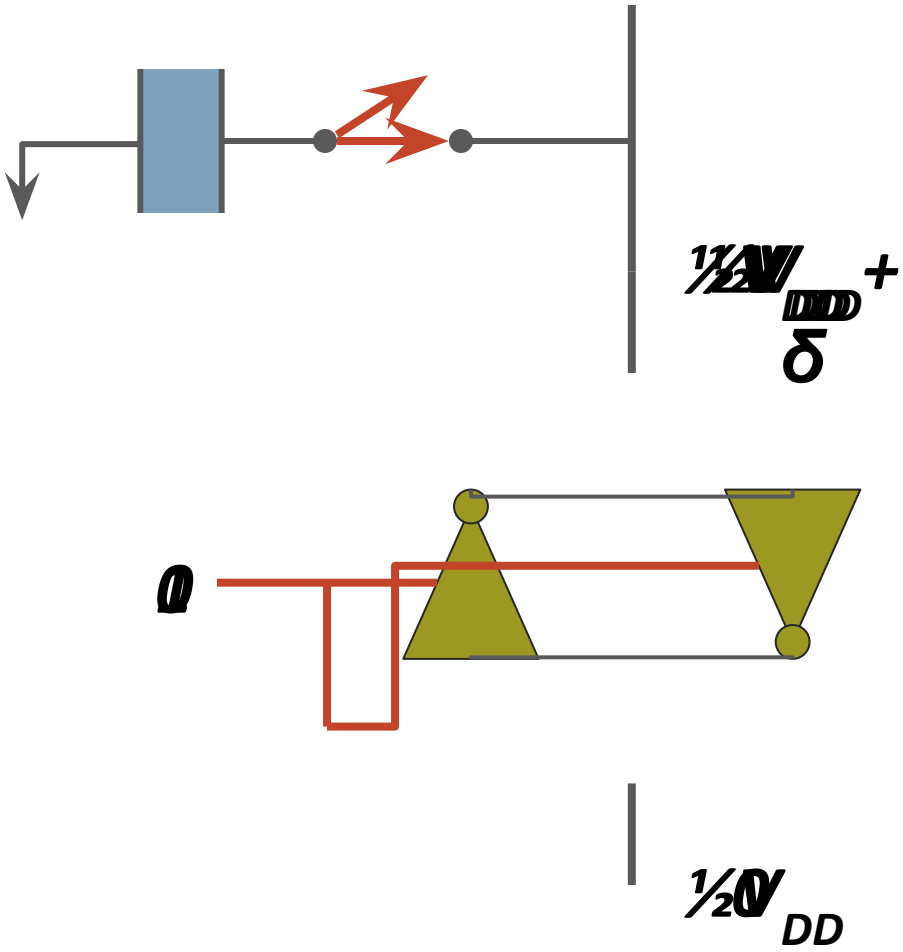

DRAM Subarray – Building Block for DRAM Chip

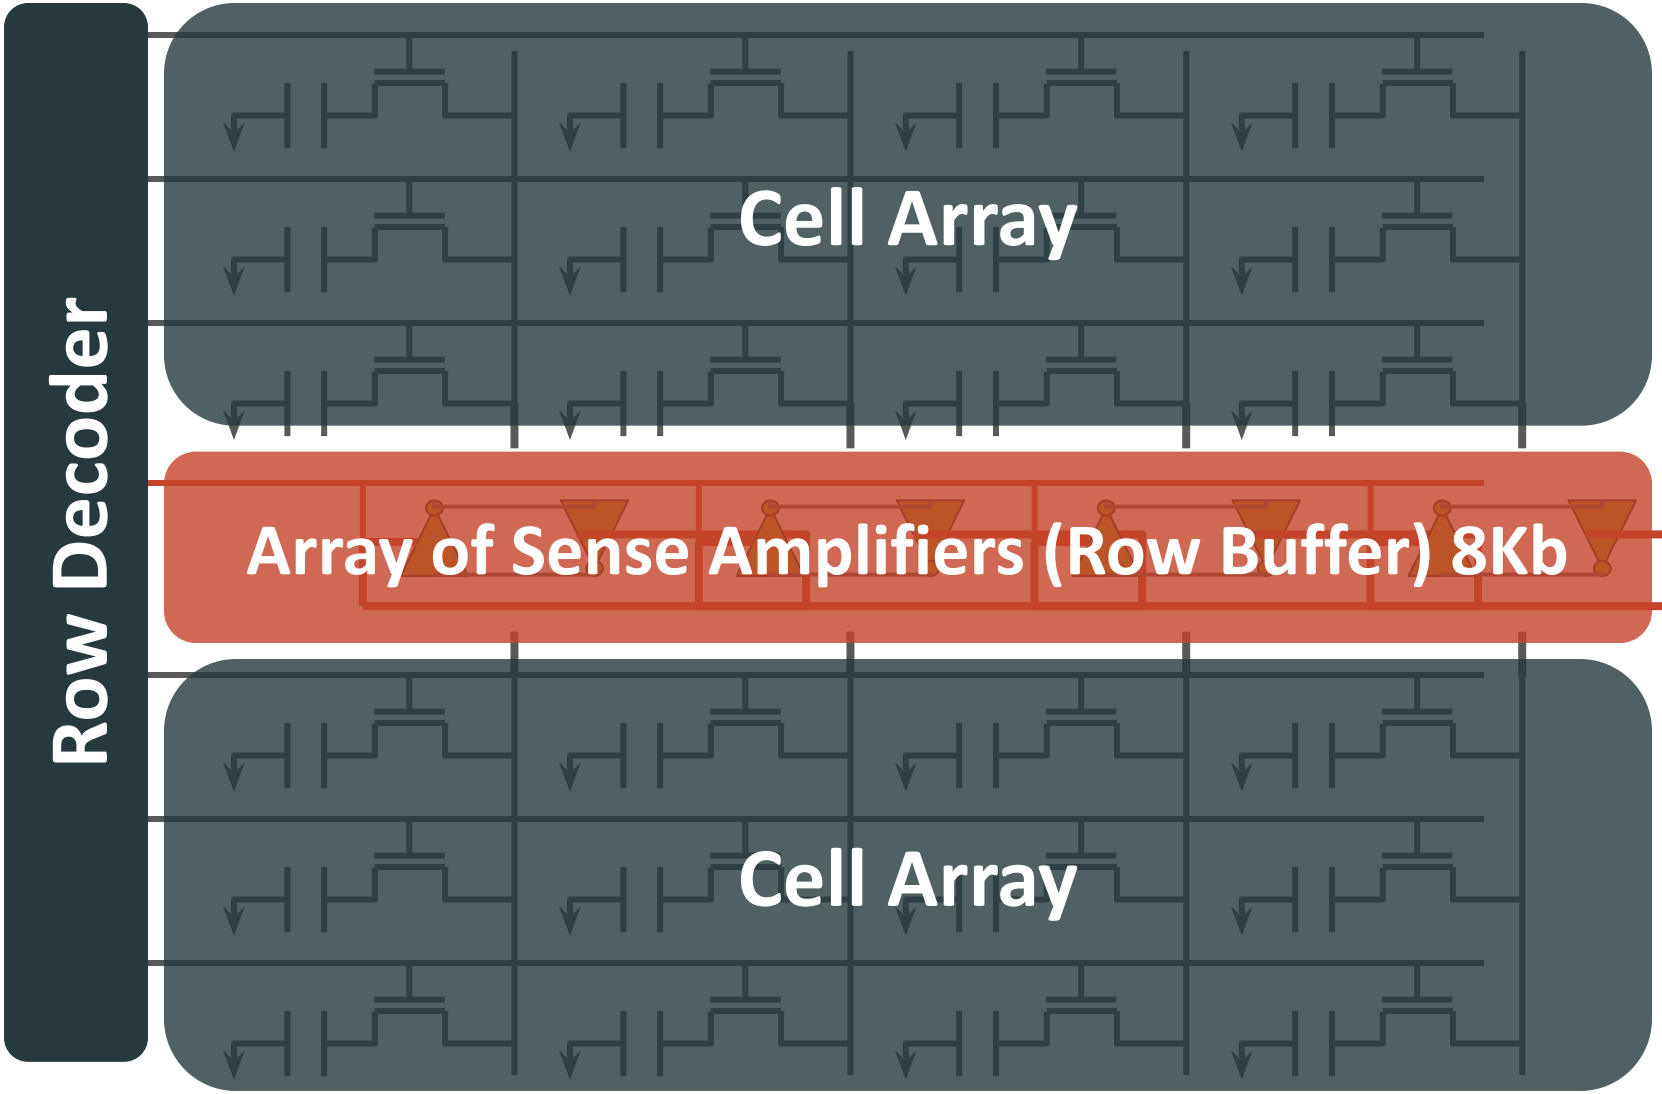

DRAM Bank

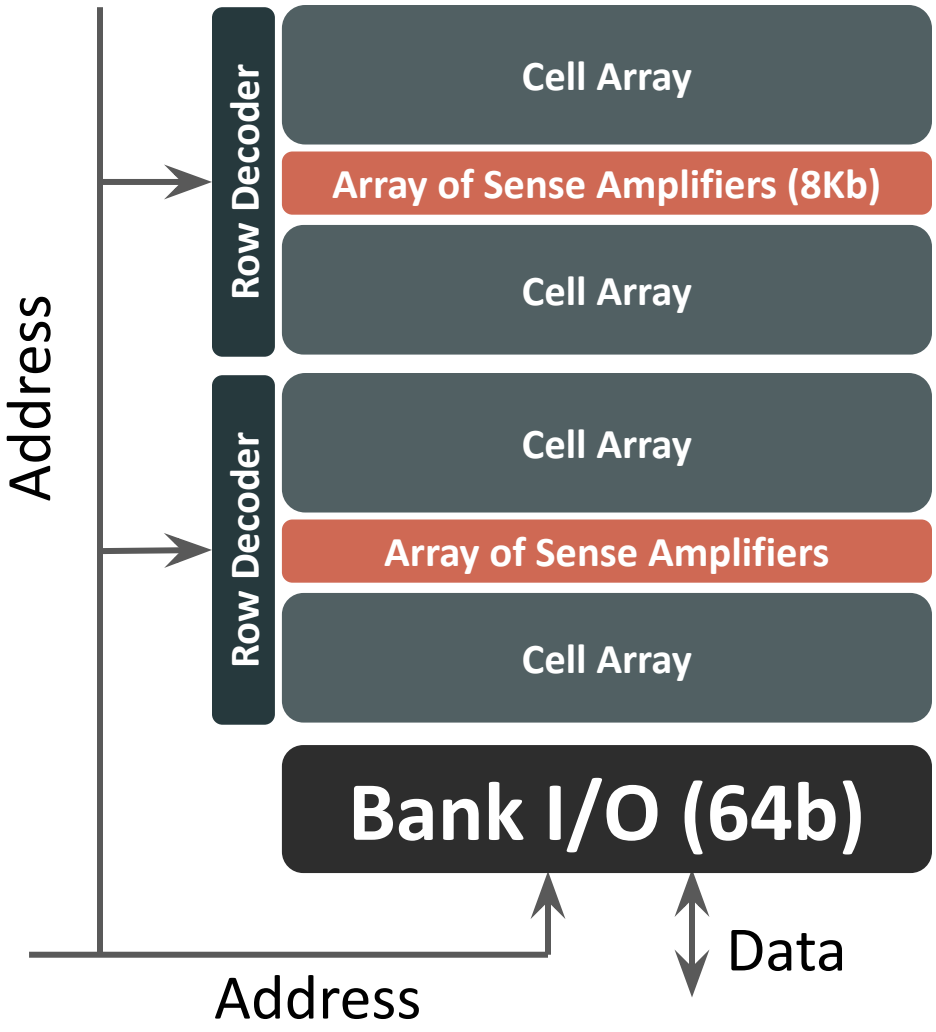
Address
Row Decoder
Cell Array
Array of Sense Amplifiers (8Kb)
Cell Array
Row Decoder
Cell Array
Array of Sense Amplifiers
Cell Array
Bank I/O (64b)
Data
Address

DRAM Chip

Shared internal bus

Memory channel - 8bits

DRAM Operation

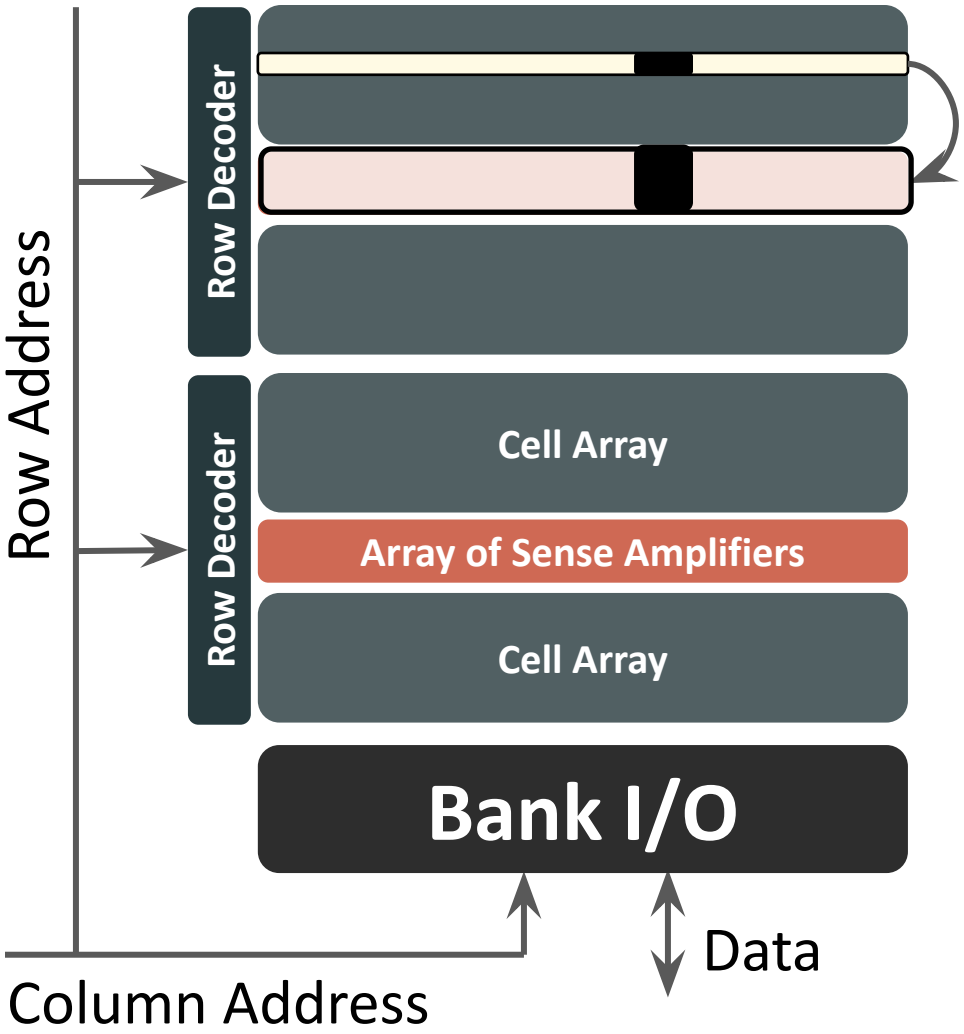

1. ACTIVATE Row
2. READ/WRITE Column
3. PRECHARGE

RowClone

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

Vivek Seshadri

Y. Kim, C. Fallin, D. Lee, R. Ausavarungnirun,
G. Pekhimenko, Y. Luo, O. Mutlu,
P. B. Gibbons, M. A. Kozuch, T. C. Mowry

Memory Channel – Bottleneck

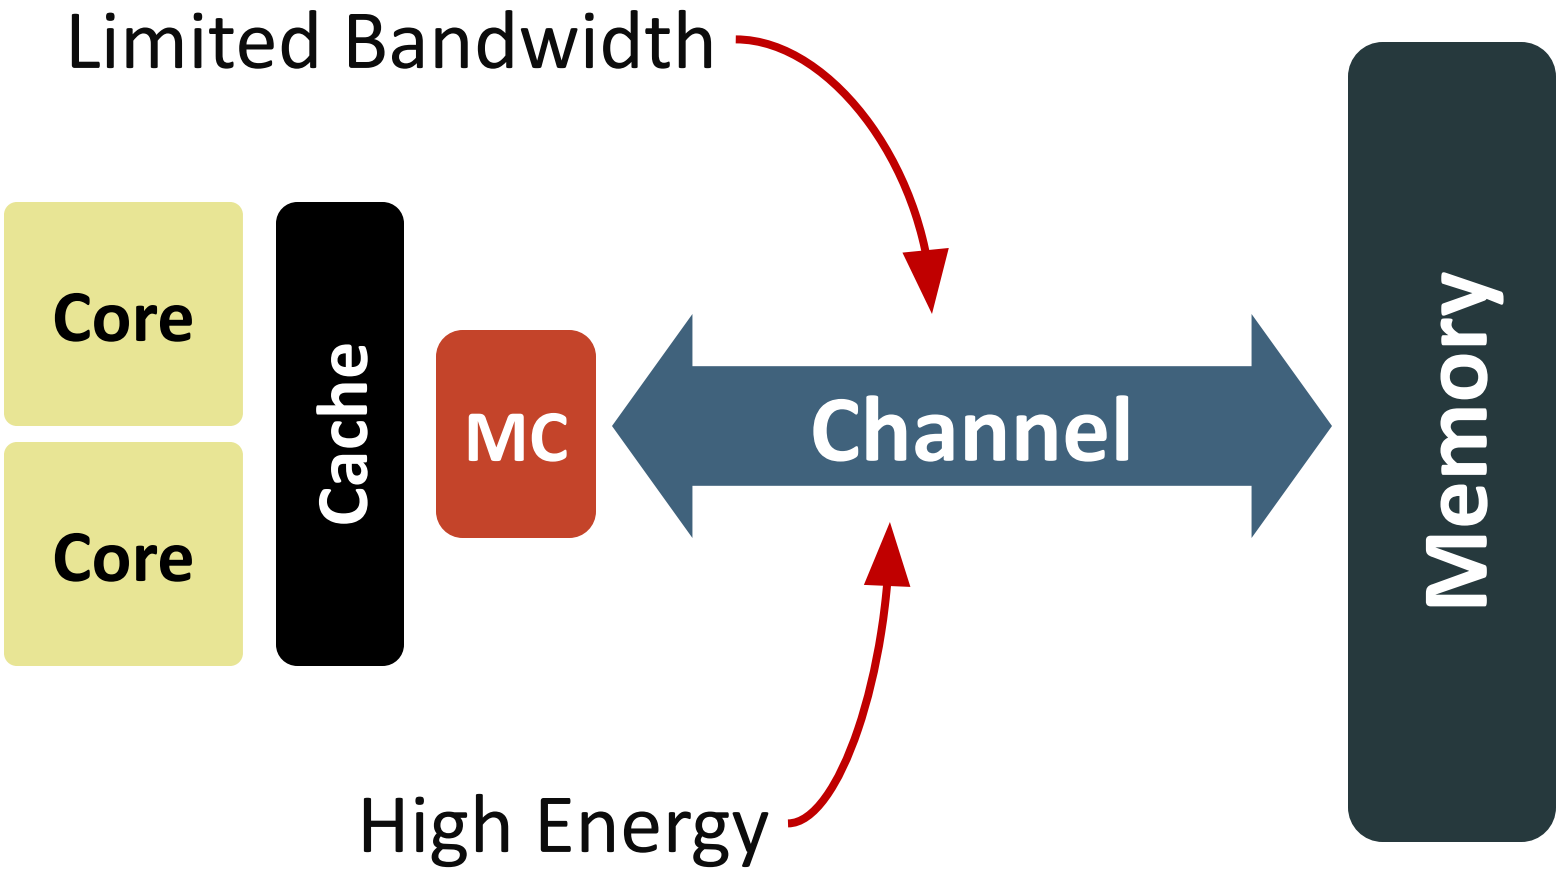

Goal: Reduce Memory Bandwidth Demand

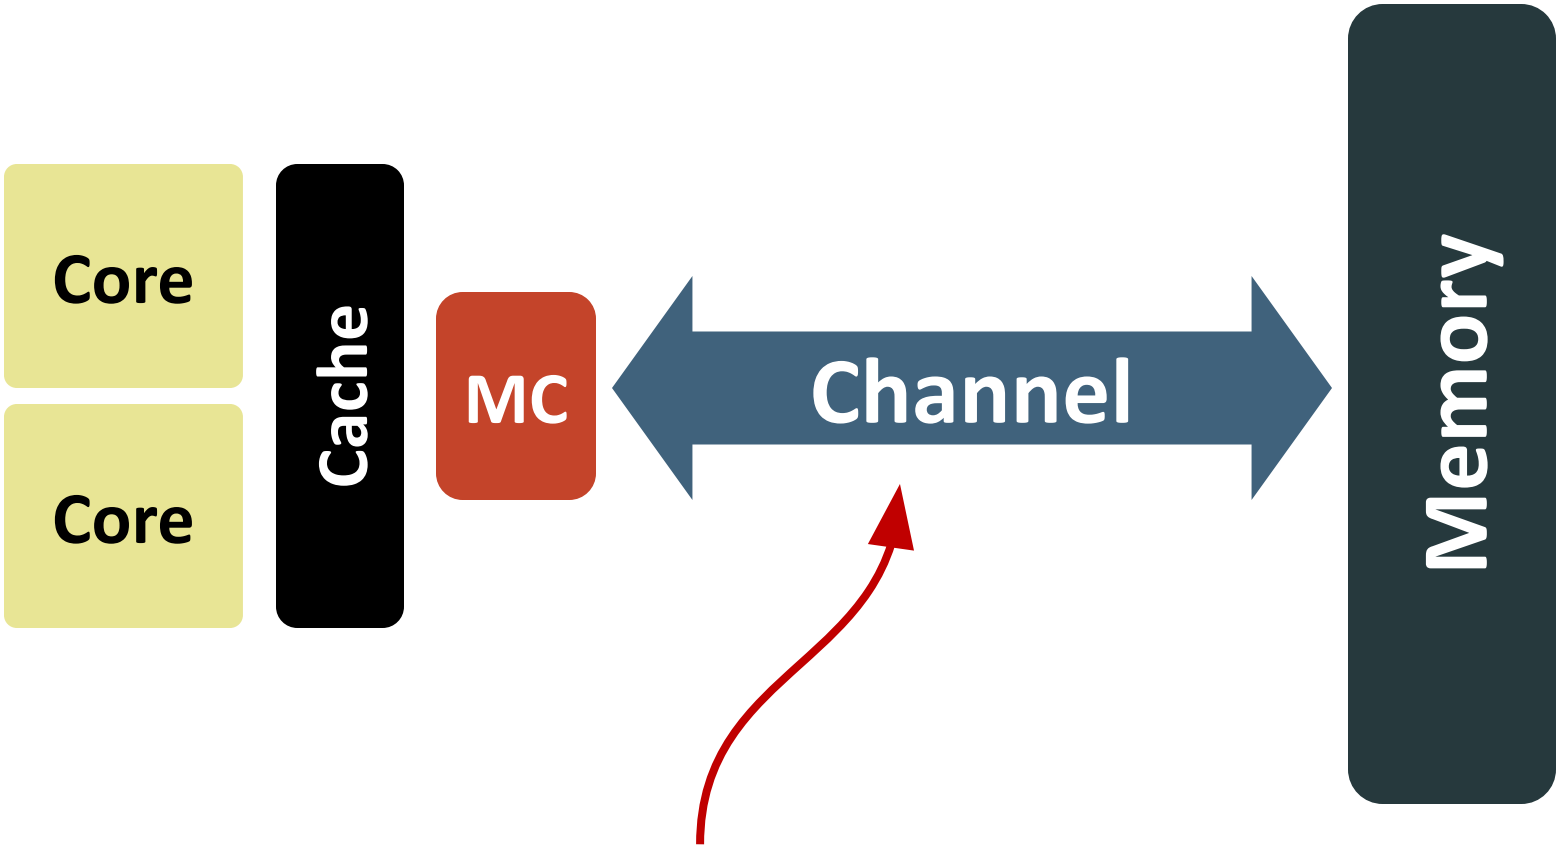

Reduce unnecessary data movement

Bulk Data Copy and Initialization

Bulk Data Copy

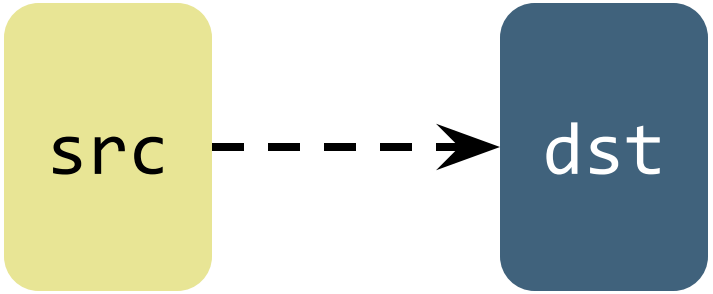

Bulk Data Initialization

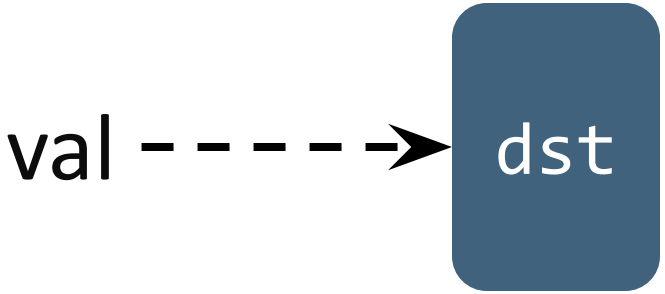

Bulk Data Copy and Initialization

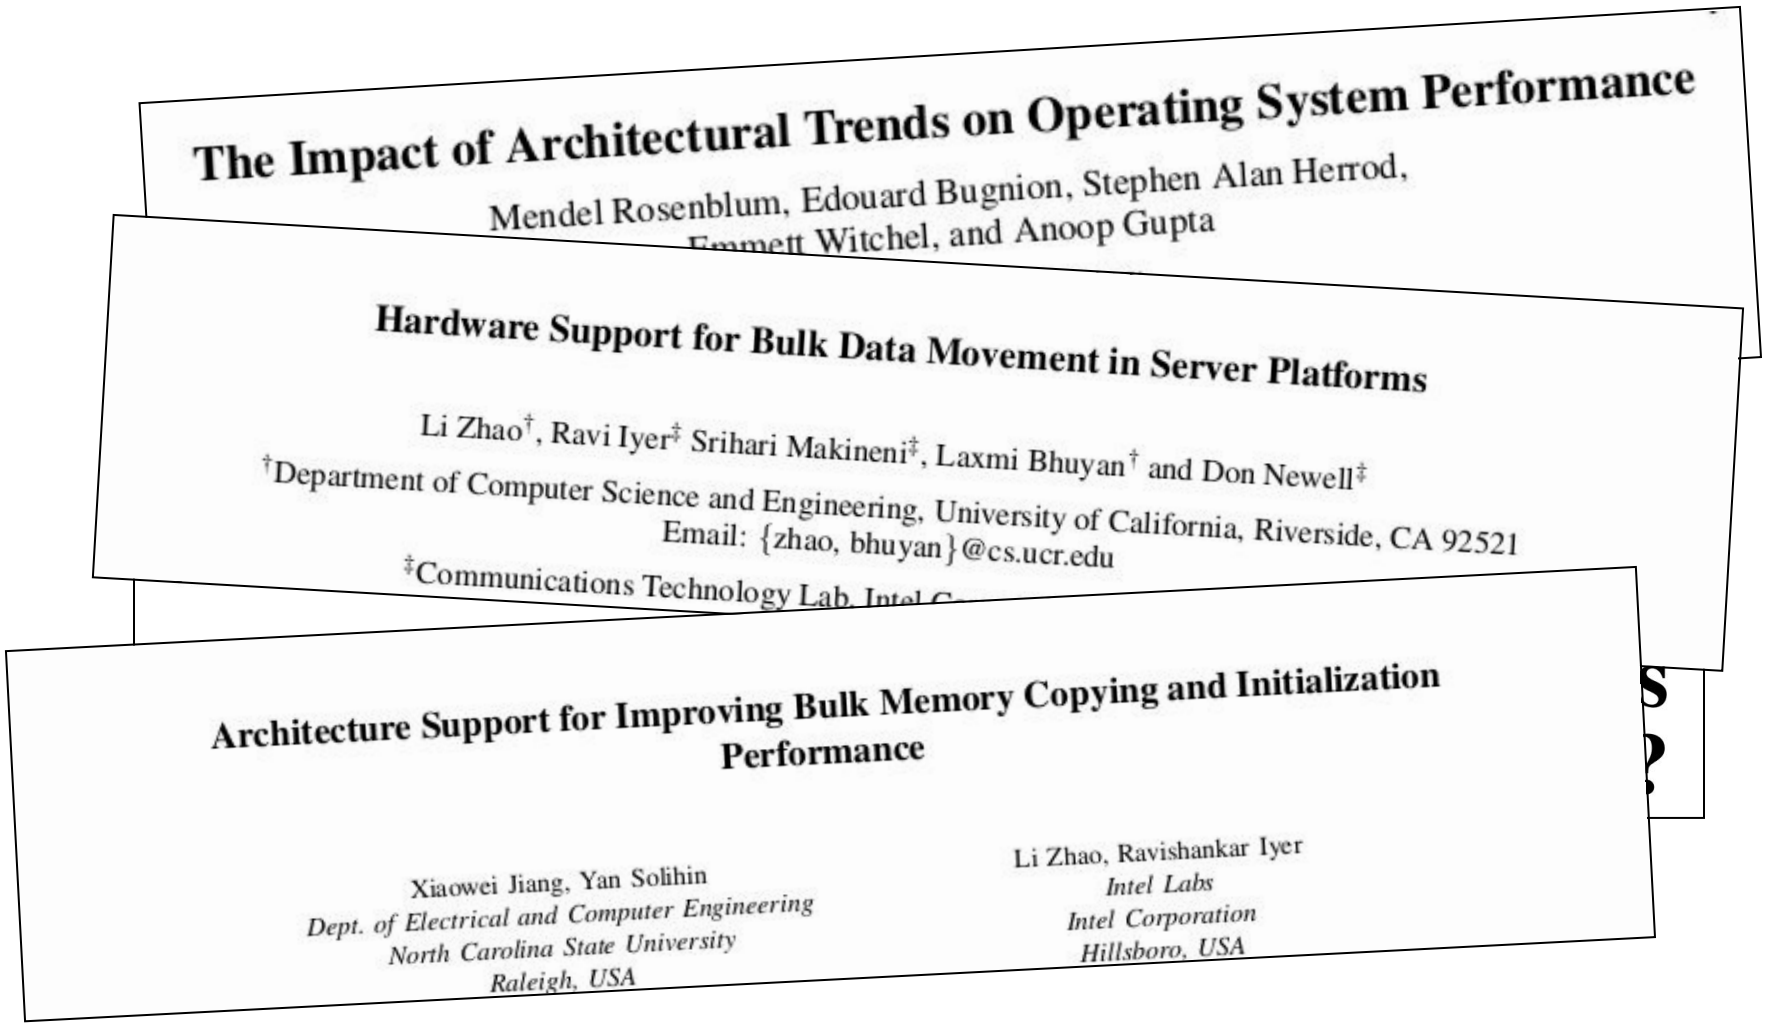

The Impact of Architectural Trends on Operating System Performance

Mendel Rosenblum, Edouard Bugnion, Stephen Alan Herrod,
Emmett Witchel, and Anoop Gupta

Hardware Support for Bulk Data Movement in Server Platforms

Li Zhao†, Ravi Iyer‡ Srihari Makineni‡, Laxmi Bhuyan† and Don Newell‡
†Department of Computer Science and Engineering, University of California, Riverside, CA 92521
Email: {zhao, bhuyan}@cs.ucr.edu
‡Communications Technology Lab, Intel C...

Architecture Support for Improving Bulk Memory Copying and Initialization Performance

Xiaowei Jiang, Yan Solihin
Dept. of Electrical and Computer Engineering
North Carolina State University
Raleigh, USA

Li Zhao, Ravishankar Iyer
Intel Labs
Intel Corporation
Hillsboro, USA

Bulk Copy and Initialization – Applications

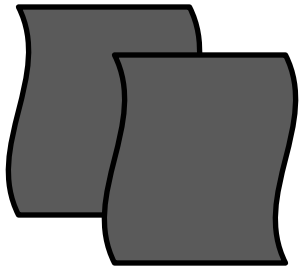

Forking

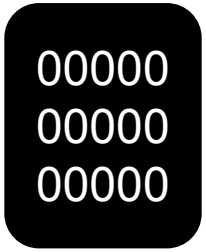

Zero initialization (e.g., security)

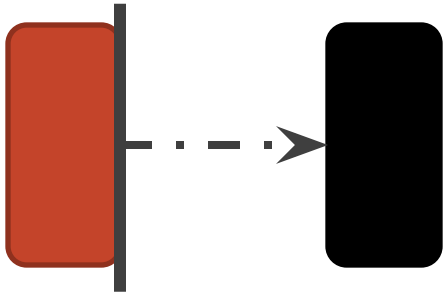

Checkpointing

VM Cloning
Deduplication

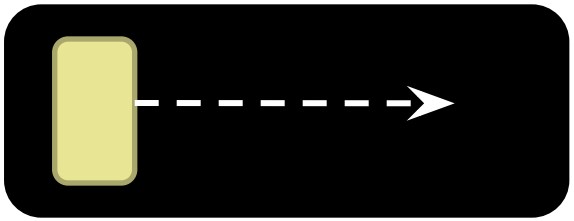

Page Migration

• • •
Many more

Shortcomings of Existing Approach

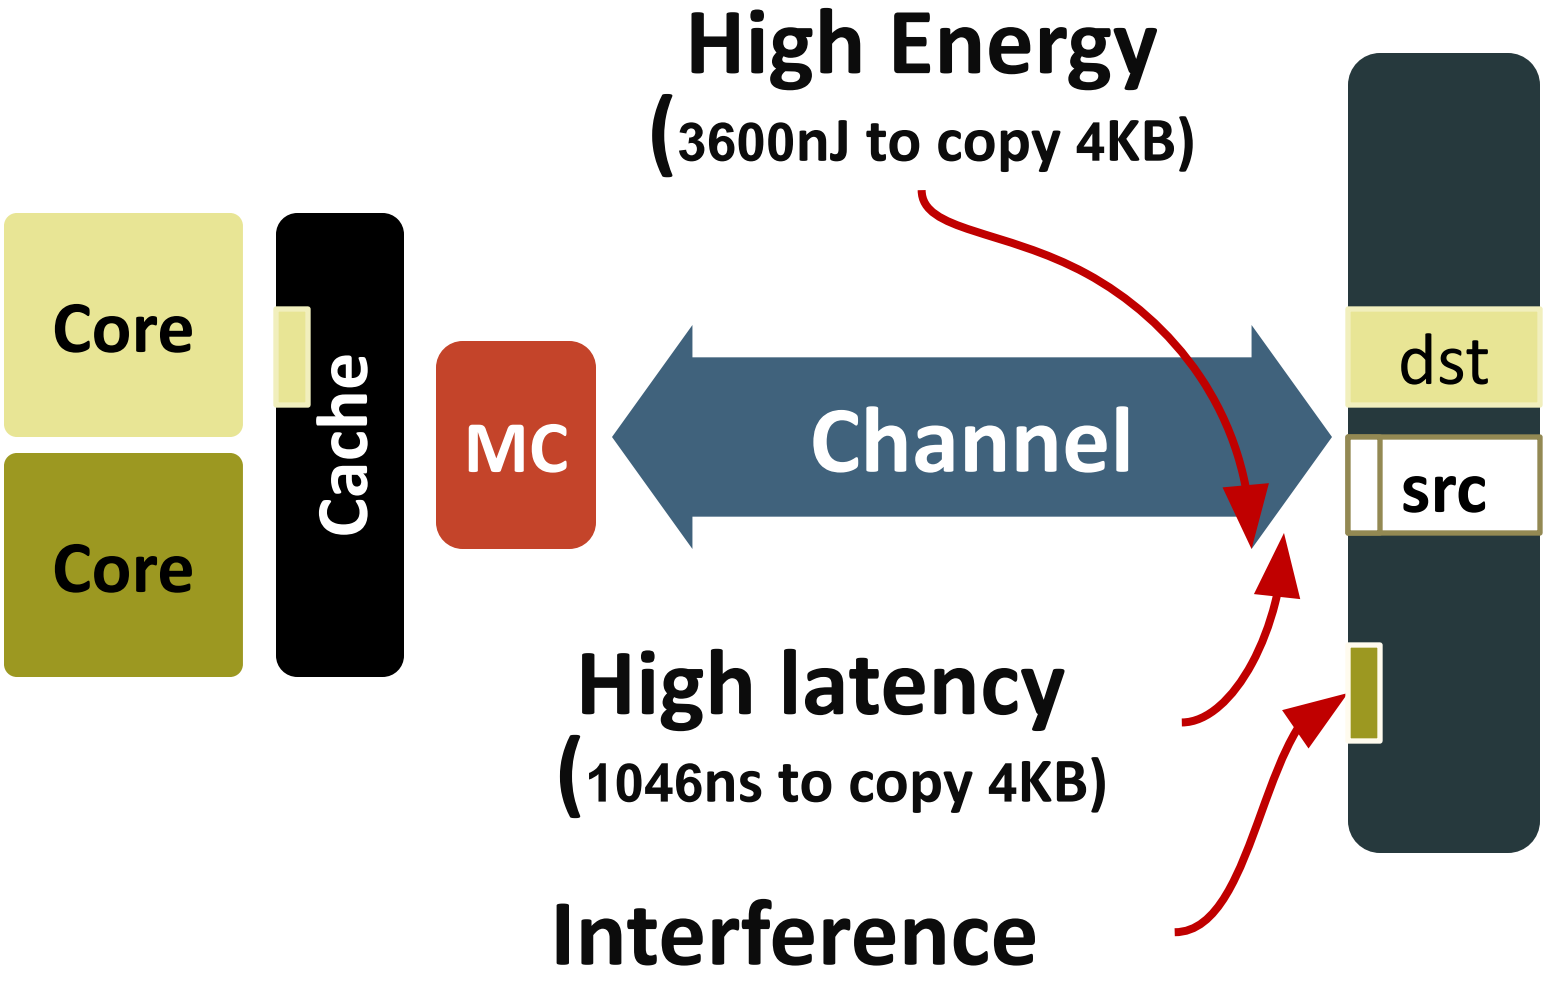

Our Approach: In-DRAM Copy with Low Cost

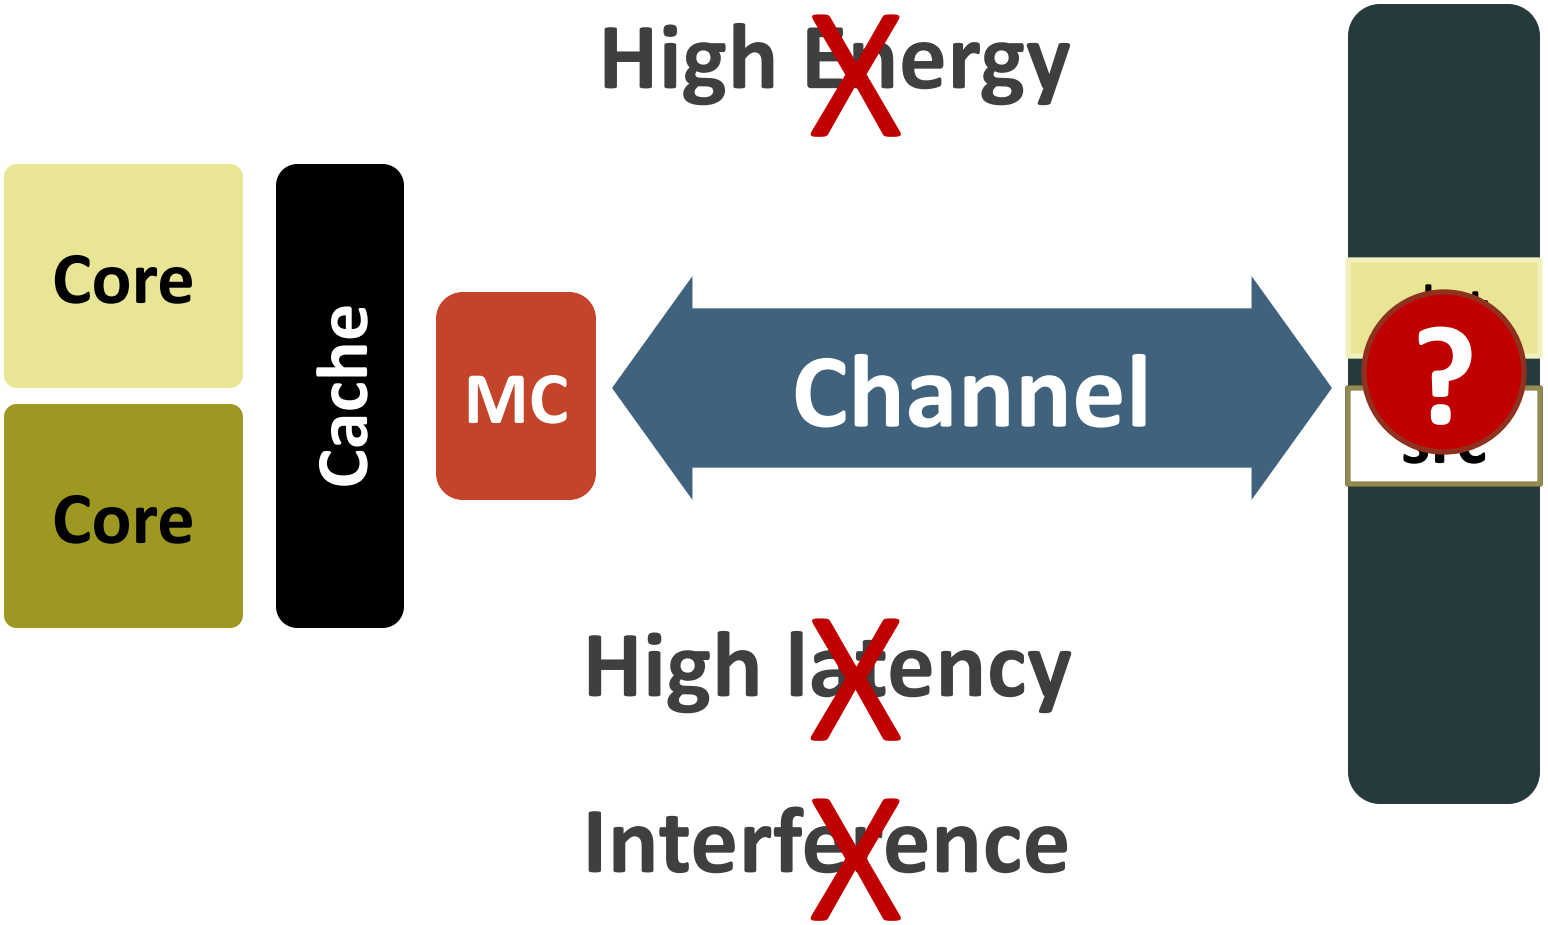

RowClone: In-DRAM Copy

Two Key Observations

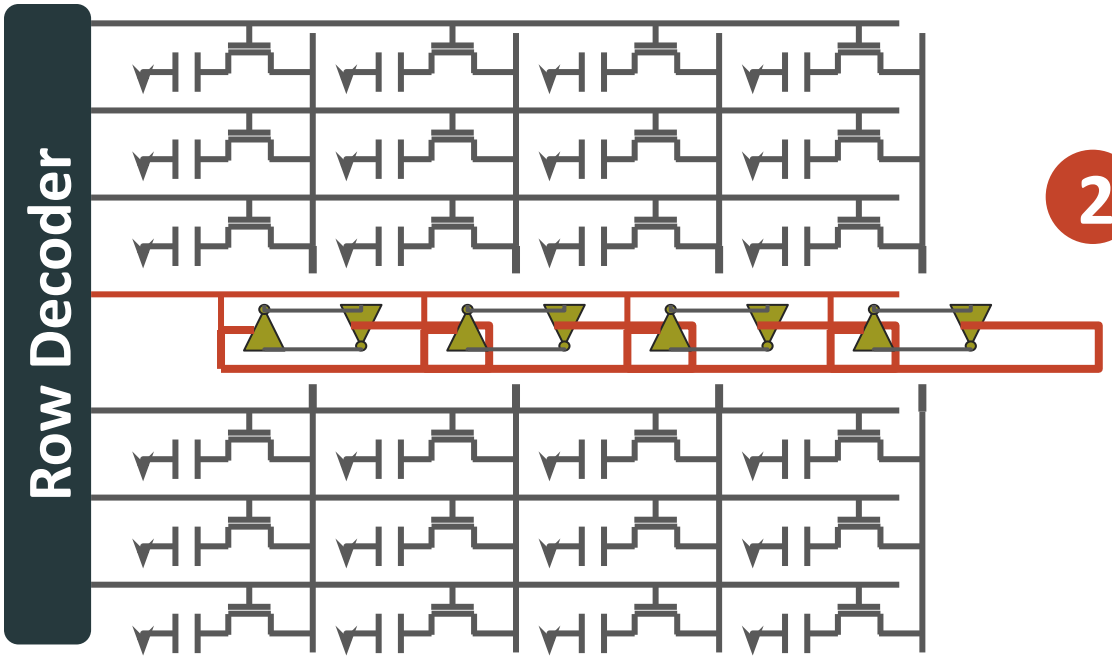

2 Many DRAM cells share the same sense amplifier

1 Any operation on one sense amplifier can be easily performed in bulk

Bulk Copy in DRAM – RowClone

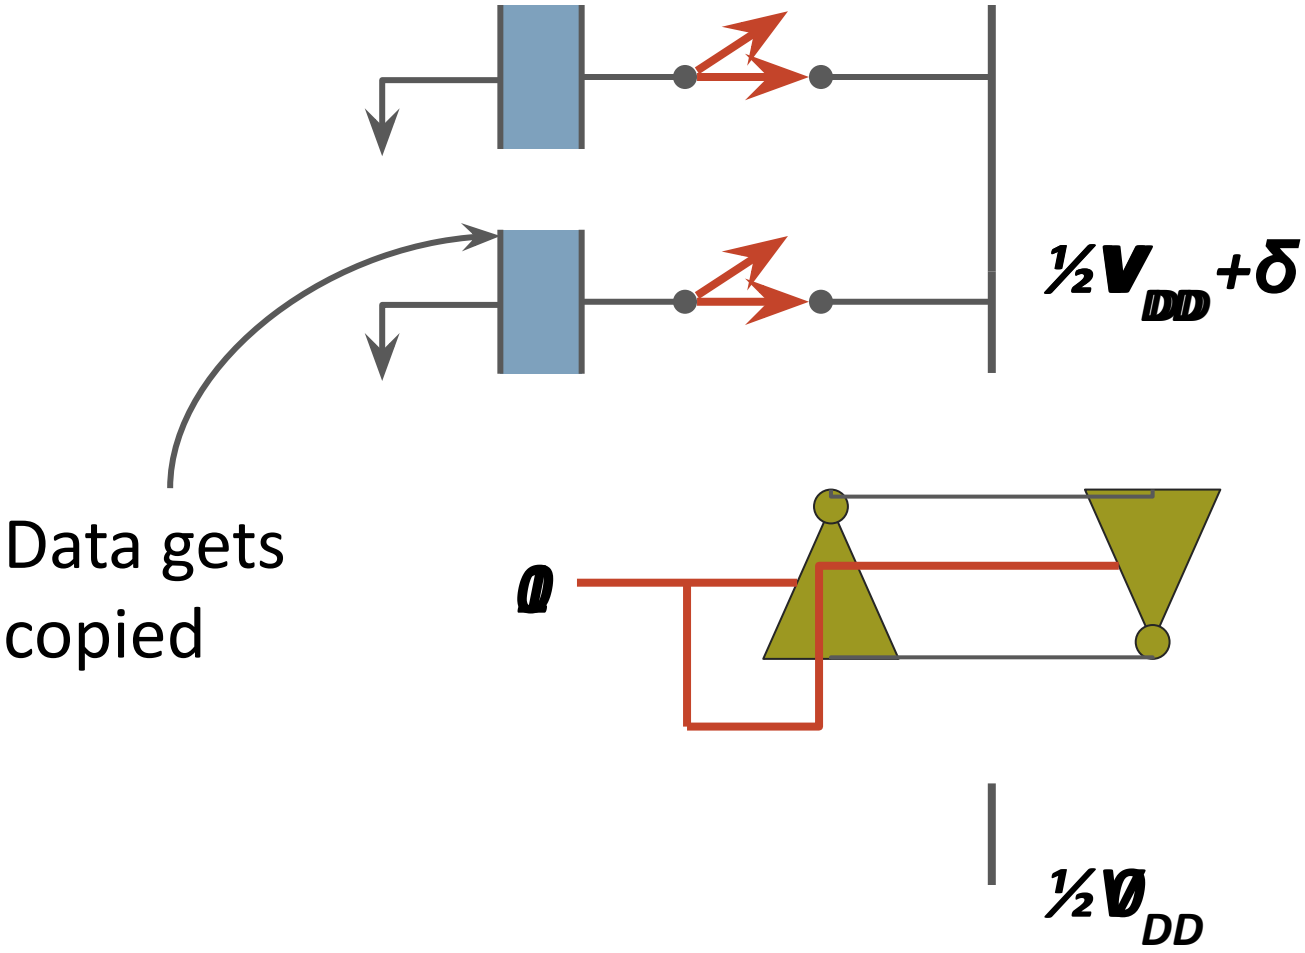

Fast Parallel Mode – Benefits

Bulk Data Copy (4KB across a module)

Latency **11X** ↓ 1046ns to 90ns

Energy **74X** ↓ 3600nJ to 40nJ

No bandwidth consumption

Very little changes to the DRAM chip

Fast Parallel Mode – Constraints

- Location constraint
  - Source and destination in same subarray
- Size constraint
  - Entire row gets copied (no partial copy)

1. Can still accelerate many existing primitives (*copy-on-write, bulk zeroing*)
2. Alternate mechanism to copy data across banks (*pipelined serial mode – lower benefits than Fast Parallel*)

End-to-end System Design

- Software interface
  - `memcpy` and `meminit` instructions
- Managing cache coherence
  - Use existing DMA support!
- Maximizing use of Fast Parallel Mode
  - Smart OS page allocation

Applications Summary

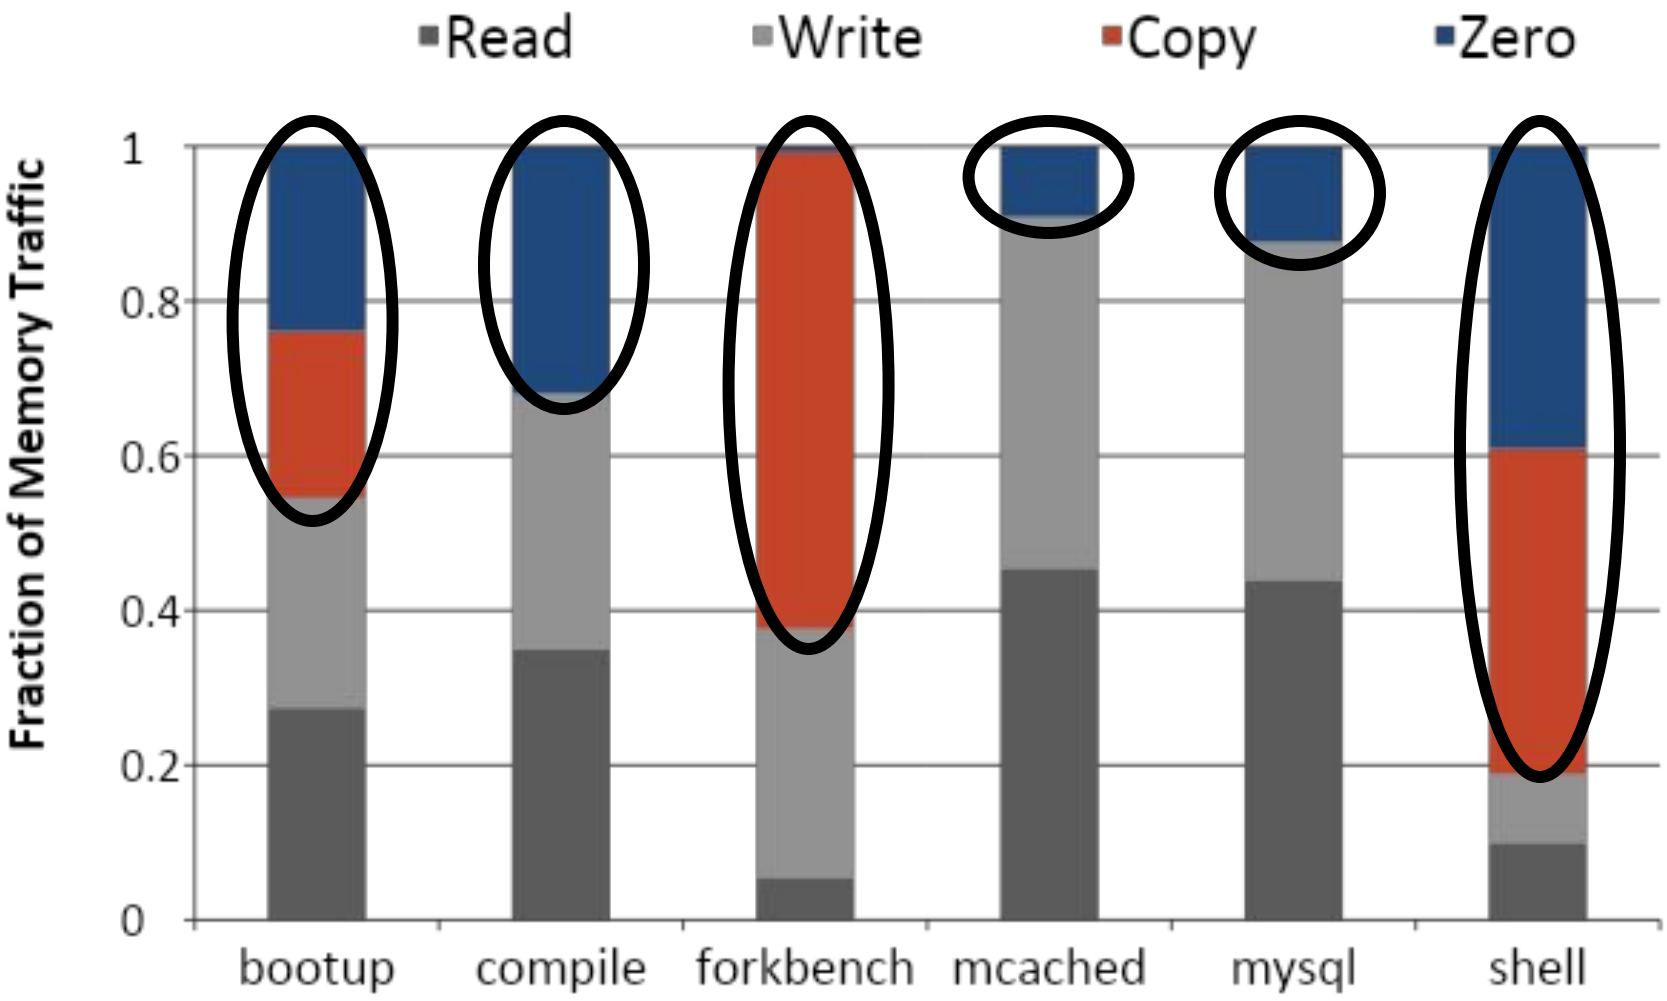

Results Summary