

# Architecting Phase Change Memory as a Scalable DRAM Alternative

Benjamin Lee<sup>†</sup>, Engin Ipek<sup>†</sup>, Onur Mutlu<sup>‡</sup>, Doug Burger<sup>†</sup>

<sup>†</sup> Computer Architecture Group  
Microsoft Research

<sup>‡</sup> Computer Architecture Lab  
Carnegie Mellon University

International Symposium on Computer Architecture  
22 June 2009

# Memory in Transition

## ► Charge Memory



- ▷ Write data by capturing charge  $Q$
- ▷ Read data by detecting voltage  $V$
- ▷ Examples: Flash, DRAM

## ► Resistive Memory

- ▷ Write data by driving current  $dQ/dt$
- ▷ Read data by detecting resistance  $R$
- ▷ Examples: PCM, MRAM, memristor

# Limits of Charge Memory

- ▷ Unscalable charge placement and control
- ▷ Flash: floating gate charge
- ▷ DRAM: capacitor charge, transistor leakage



# Towards Resistive Memory

## ▶ Scalable

- ▷ Program with current  $\propto$  cell size
- ▷ Map resistance to logical state

## ▶ Non-Volatile


- ▷ Set atomic structure in cell
- ▷ Incur activation cost

## ▶ Competitive

- ▷ Achieve viable delay, energy, endurance
- ▷ Scale to further improve metrics

## PCM Deployment

- ▶ Deploy PCM on the memory bus
- ▶ Begin by co-locating PCM, DRAM
- ▶ Begin by deploying in low-power platforms



# Outline

## ▶ Motivation

- ▷ Memory Scaling
- ▷ Charge Memory
- ▷ Resistive Memory

## ▶ Technology



- ▷ Phase Change Memory
- ▷ Technology Parameters
- ▷ Price of Scalability

## ▶ Architecture

- ▷ Design Objectives
- ▷ Buffer Organization
- ▷ Partial Writes



# Phase Change Memory

- ▶ Store data within phase change material [Ovshinsky68]
- ▶ Set phase via current pulse
- ▶ Detect phase via resistance (amorphous/crystalline)



# PCM Scalability

- ▶ Program with current pulses, which scale linearly
- ▶ PCM roadmap to 30nm [Raoux+08]
- ▶ Flash/DRAM roadmap to 40nm [ITRS07]





# PCM Non-Volatility

## ▶ Atomic Structure

- ▷ Program with current pulses
- ▷ Melt material at 650 °C
- ▷ Cool material to desired phase

## ▶ Activation Cost

- ▷ Crystallize with high activation energy
- ▷ Isolate thermal effects to target cell
- ▷ Retain data for >10 years at 85 °C

# Technology Parameters

- ▶ Survey prototypes from 2003-2008 [ISSCC][VLSI][IEDM][ITRS]
- ▶ Derive parameters for  $F=90\text{nm}$

## Density

- ▶ 9 -  $12F^2$  using BJT
- ▶  $1.5\times$  DRAM

## Latency

- ▶ 50ns Rd, 150ns Wr
- ▶  $4\times$ ,  $12\times$  DRAM

## Endurance

- ▶  $1\text{E}+08$  writes
- ▶  $1\text{E}-08\times$  DRAM

## Energy

- ▶  $40\mu\text{A}$  Rd,  $150\mu\text{A}$  Wr
- ▶  $2\times$ ,  $43\times$  DRAM

# Technology Parameters

- ▶ Survey prototypes from 2003-2008 [ISSCC][VLSI][IEDM][ITRS]
- ▶ Derive parameters for  $F=90\text{nm}$

## Density

- ▶ 9 -  $12F^2$  using BJT
- ▶  $1.5\times$  DRAM

## Latency

- ▶ 50ns Rd, 150ns Wr
- ▶  $4\times$ ,  $12\times$  DRAM

## Endurance

- ▶  $1\text{E}+08$  writes
- ▶  $1\text{E}-08\times$  DRAM

## Energy

- ▶  $40\mu\text{A}$  Rd,  $150\mu\text{A}$  Wr
- ▶  $2\times$ ,  $43\times$  DRAM

# Technology Parameters

- ▶ Survey prototypes from 2003-2008 [ISSCC][VLSI][IEDM][ITRS]
- ▶ Derive parameters for  $F=90\text{nm}$

## Density

- ▶ 9 -  $12F^2$  using BJT
- ▶  $1.5\times$  DRAM

## Latency

- ▶ 50ns Rd, 150ns Wr
- ▶  $4\times$ ,  $12\times$  DRAM

## Endurance

- ▶  $1\text{E}+08$  writes
- ▶  $1\text{E}-08\times$  DRAM

## Energy

- ▶  $40\mu\text{A}$  Rd,  $150\mu\text{A}$  Wr
- ▶  $2\times$ ,  $43\times$  DRAM

# Technology Parameters

- ▶ Survey prototypes from 2003-2008 [ISSCC][VLSI][IEDM][ITRS]
- ▶ Derive parameters for  $F=90\text{nm}$

## Density

- ▶ 9 -  $12F^2$  using BJT
- ▶  $1.5\times$  DRAM

## Latency

- ▶ 50ns Rd, 150ns Wr
- ▶  $4\times$ ,  $12\times$  DRAM

## Endurance



- ▶  $1\text{E}+08$  writes
- ▶  $1\text{E}-08\times$  DRAM

## Energy

- ▶  $40\mu\text{A}$  Rd,  $150\mu\text{A}$  Wr
- ▶  $2\times$ ,  $43\times$  DRAM

# Price of Scalability

- ▷ 1.6× delay, 2.2× energy, 500-hour lifetime
- ▷ Implement PCM in typical DRAM architecture



# Outline

## ▶ Motivation

- ▷ Memory Scaling
- ▷ Charge Memory
- ▷ Resistive Memory

## ▶ Technology

- ▷ Phase Change Memory
- ▷ Technology Parameters
- ▷ Price of Scalability

## ▶ Architecture

- ▷ Design Objectives
- ▷ Buffer Organization
- ▷ Partial Writes

# Design Objectives

## ▶ **DRAM-Competitive**

- ▷ Reorganize row buffer to mitigate delay, energy
- ▷ Implement partial writes to mitigate wear mechanism

## ▶ **Area-Efficient**

- ▷ Minimize disruption to density trends
- ▷ Impacts row buffer organization

## ▶ **Complexity-Effective**

- ▷ Encourage adoption with modest mechanisms
- ▷ Impacts partial writes



# Buffer Organization

## ▶ On-Chip Buffers

- ▷ Use DRAM-like buffer and interface
- ▷ Evict modified rows into array

## ▶ Narrow Rows

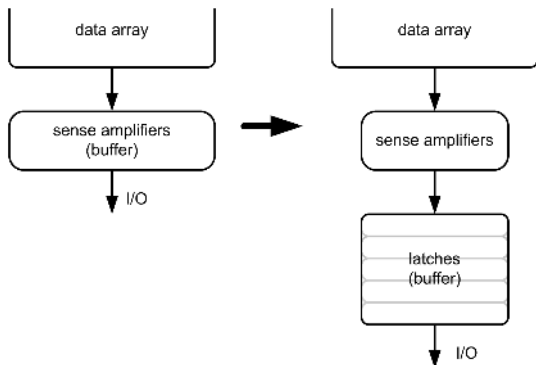
- ▷ Reduce write energy  $\propto$  buffer width
- ▷ Reduce peripheral circuitry, associated area

## ▶ Multiple Rows

- ▷ Reduce eviction frequency
- ▷ Improve locality, write coalescing

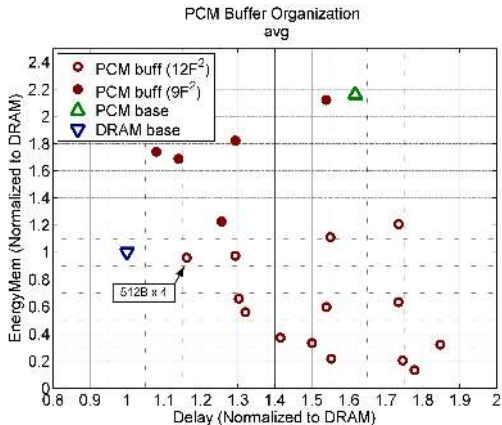
## Buffer Area Strategy

- ▷ Narrow rows :: fewer expensive S/A's (44T)
- ▷ Multiple rows :: more inexpensive latches (8T)



# Buffer Design Space

- ▶ Explore area-neutral buffer designs
- ▶ Identify DRAM-competitive buffer design



# Wear Reduction

## ▶ **Wear Mechanism**

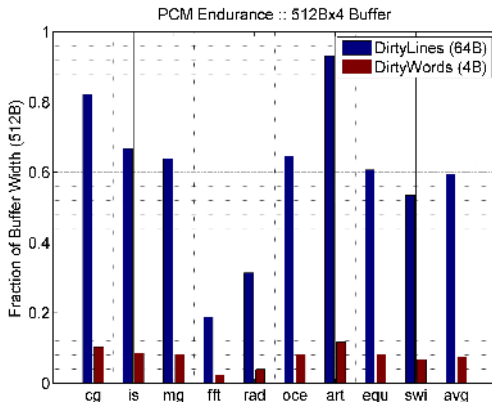
- ▶ Writes induce phase change at 650 °C
- ▶ Contacts degrade from thermal expansion/contraction
- ▶ Current injection is less reliable after 1E+08 writes

## ▶ **Partial Writes**

- ▶ Reduce writes to PCM array
- ▶ Write only stored lines (64B), words (4B)
- ▶ Add cache line state with 0.2%, 3.1% overhead



## Partial Writes

- ▷ Derive PCM lifetime model
- ▷ Quantify eliminated writes during buffer eviction



# Scalable Performance

- ▶ 1.2× delay, 1.0× energy, >5-year lifetime
- ▶ Scaling improves energy, endurance



## Also in the paper...

### ▶ **Technology Survey**

- ▷ Survey of circuit/device prototypes
- ▷ PCM architectural timing, energy models
- ▷ Scaling analysis, implications

### ▶ **Buffer Organization**

- ▷ Transistor-level area model
- ▷ Buffer sensitivity analysis

### ▶ **Partial Writes**

- ▷ Endurance model
- ▷ Bus activity analysis

# Conclusion & Future Directions

## ▶ **Memory Scaling**

- ▷ Fundamental limits in charge memory
- ▷ Transition towards resistive memory

## ▶ **Phase Change Memory**

- ▷ Scalability and non-volatility
- ▷ Competitive delay, energy, endurance
- ▷ DRAM alternative on the memory bus

## ▶ **Applied Non-Volatility**

- ▷ Instant start, hibernate
- ▷ Inexpensive checkpointing
- ▷ Safe file systems



# PCM File System (PFS)

J. Condit et al. “**Better I/O through byte-addressable, persistent memory.**” SOSP-22: Symposium on Operating System Principles, October 2009. (To Appear)

## ▶ File System Properties

- ▶ Consistency :: COW with atomicity, ordering
- ▶ Safety :: Reflect writes to PCM in O(ms), not O(s)
- ▶ Performance :: Outperform NTFS on RAM disk

## ▶ Architectural Support

- ▶ Atomic 8B writes with capacitive support
- ▶ Ordered writes with barrier-delimited epochs

# Architecting Phase Change Memory as a Scalable DRAM Alternative

Benjamin Lee<sup>†</sup>, Engin Ipek<sup>†</sup>, Onur Mutlu<sup>‡</sup>, Doug Burger<sup>†</sup>

<sup>†</sup> Computer Architecture Group  
Microsoft Research

<sup>‡</sup> Computer Architecture Lab  
Carnegie Mellon University

International Symposium on Computer Architecture  
22 June 2009