Moneta:

A High-Performance Storage Architecture for Next-generation, Non-volatile Memories

Adrian M. Caulfield

Arup De, Joel Coburn, Todor I. Mollov, Rajesh K. Gupta, Steven Swanson

Non-Volatile Systems Laboratory,
Department of Computer Science and Engineering
University of California, San Diego

NVSL
Non-volatile Systems Laboratory

UCSDCSE
Computer Science and Engineering

The Future of Storage

	Hard Drives	PCIe-Flash 2007	PCIe-NVM 2013?	
				
Lat.:	7.1ms	68us	12us	
BW:	2.6MB/s	250MB/s	1.7GB/s	
	1x	104x	591x	= 2.89x/yr
	1x	96x	669x	= 2.95x/yr

*Random 4KB Reads from user space

Software Overheads

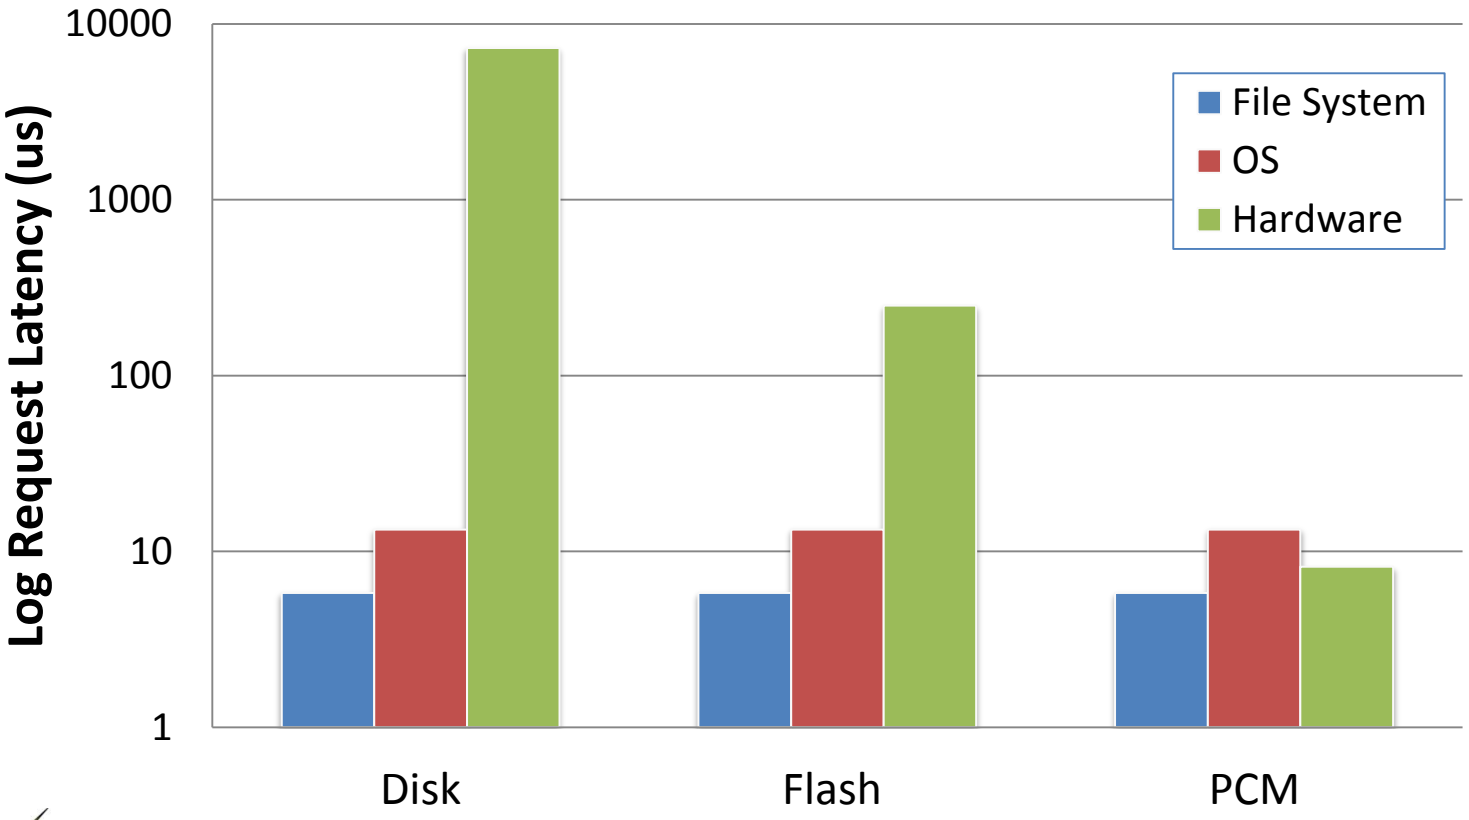

Architecting a High Performance SSD

- Hardware architecture and software layers limit performance
- HW/SW interface critical to good performance
- Careful co-design provides significant benefits
  - Increased bandwidth
  - Decreased latencies
  - Increased concurrency

Overview

- Motivation
- **Moneta Architecture**
- Optimizing Moneta
- Performance Analysis
- Conclusion

Moneta Architecture

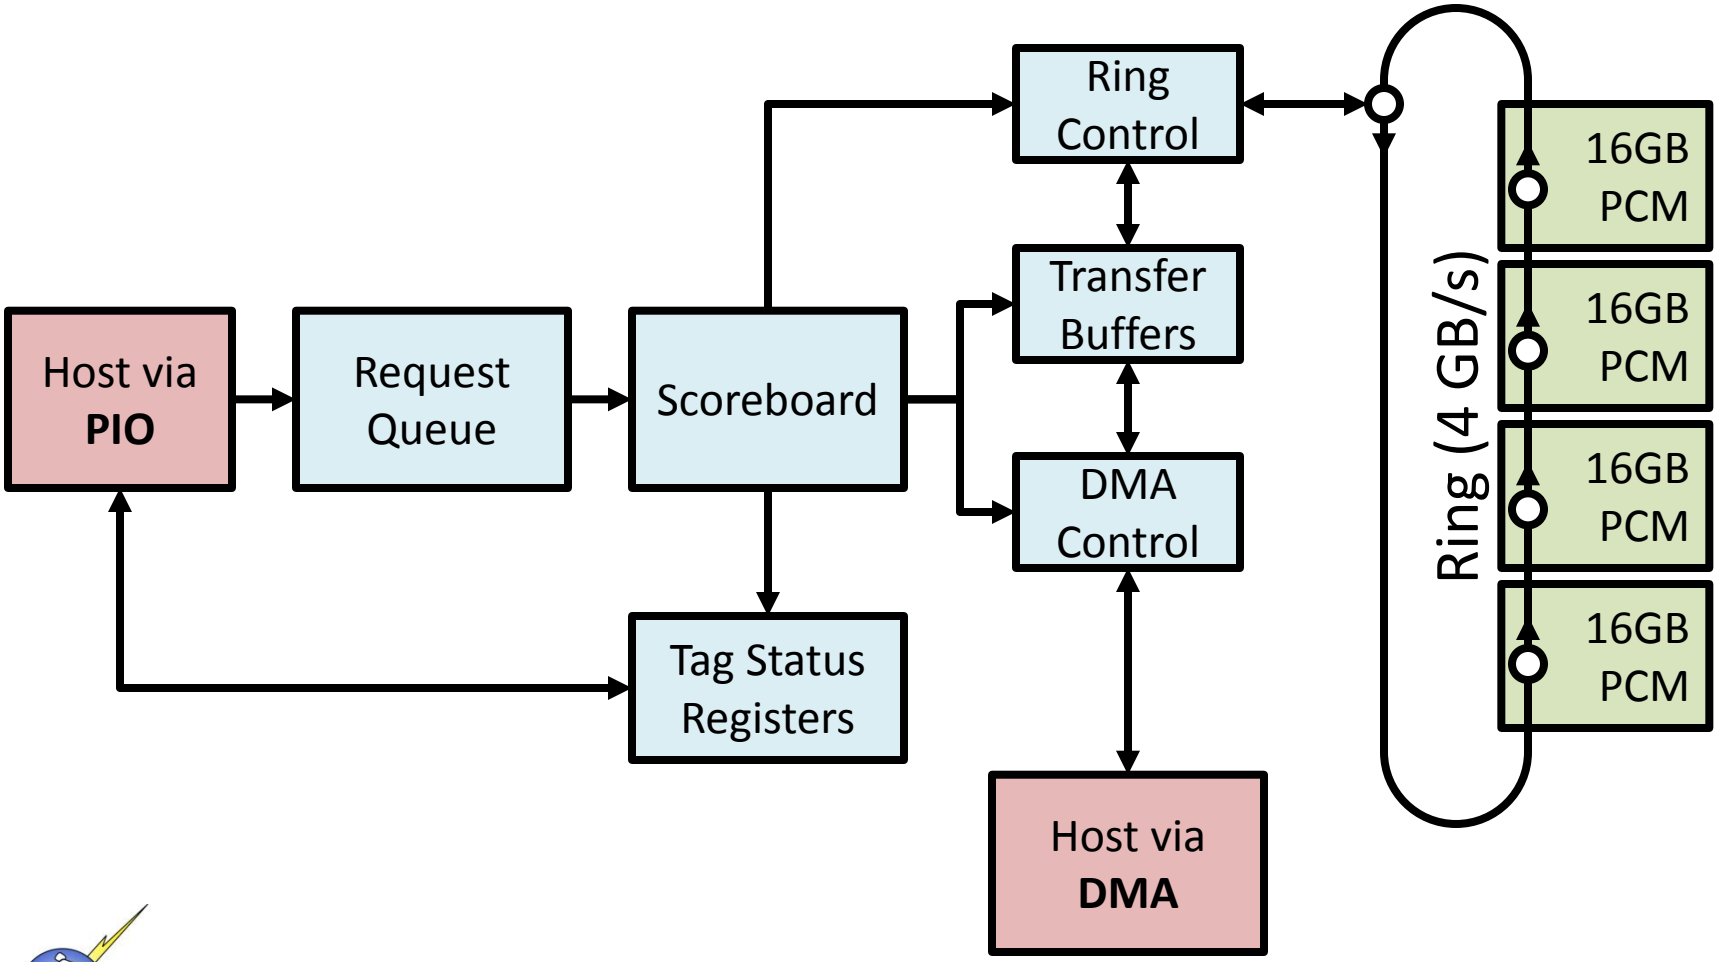

Moneta Architecture Overview

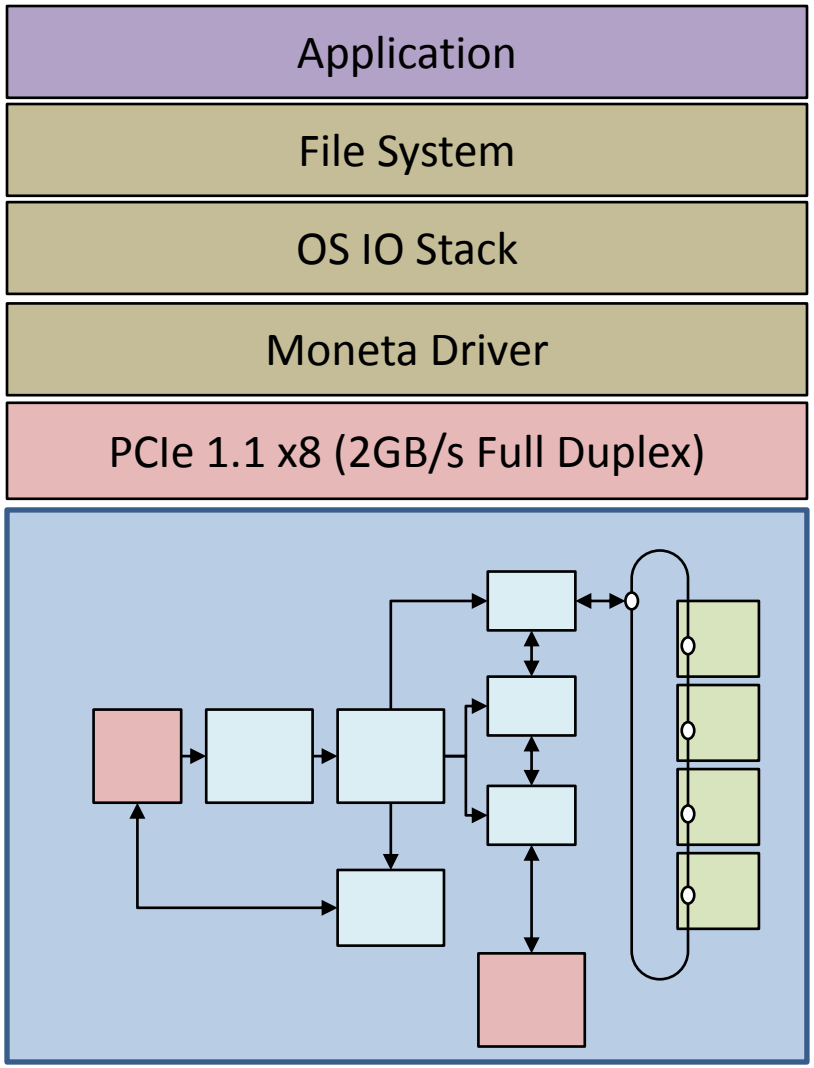

Advanced Memory Technology Characteristics

- Work with any fast NVM:
  - DRAM-like speed
  - DRAM-like interface
- Phase Change Memory
  - Coming soon
  - Simple wear leveling: Start Gap [Micro 2009]

Moneta: Modeling Advanced NVMs

- Built on RAMP's BEE3 board
- PCIe 1.1 x8 host connection
- 250MHz design
- DDR2 DRAM emulates NVMs
  - Adjust timings to match PCM
  - RAS-CAS Delay for reads
  - Precharge latency for writes

	Read	Write
Projected PCM Latency	48 ns	150 ns

Latency projections from [B.C. Lee, ISCA'09]

Overview

- Motivation
- Moneta Architecture
- **Optimizing Moneta**
  - **Interface & Software**
  - Microarchitecture
- Performance Analysis
- Conclusion

Software Optimization is Critical

- Baseline Latency (4KB)
  - Hardware: 8.2 us
  - Software: 13.4 us
- Optimizations
  - Remove IO Scheduler
  - Atomic Operations
    - Lock Free Data Structures
    - Atomic HW/SW Interface
  - Spin-waits vs. sleeping
- Final 5us SW latency
  - 62% reduction vs. Base

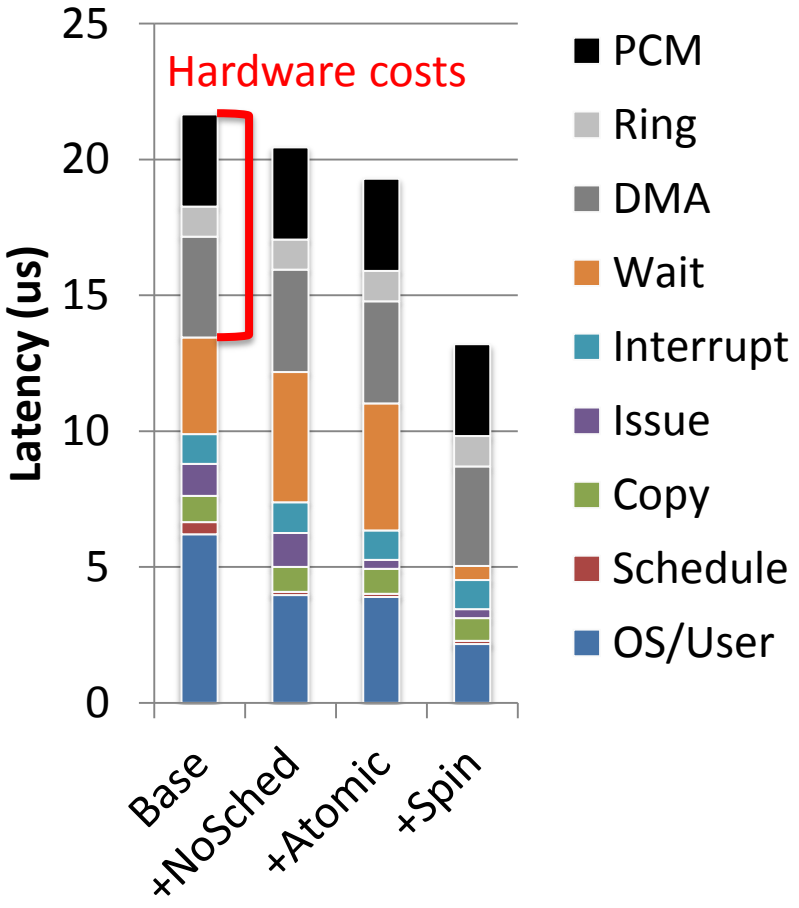

Moneta Bandwidth

Random Write Accesses

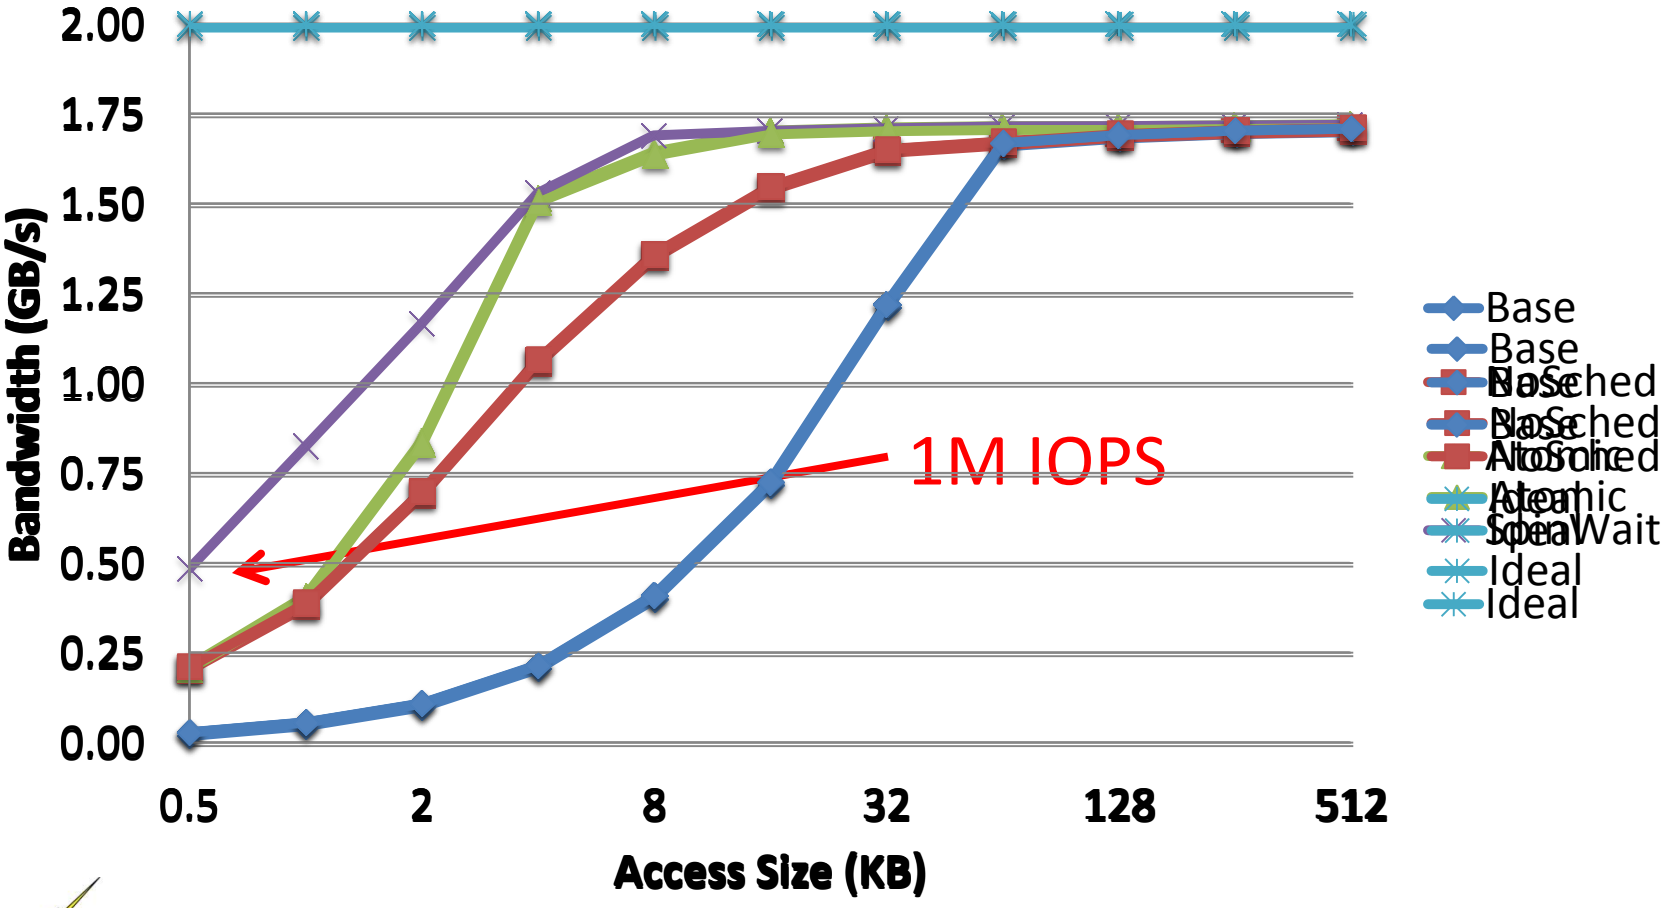

CPU Utilization

Random 4KB Read/Write Accesses

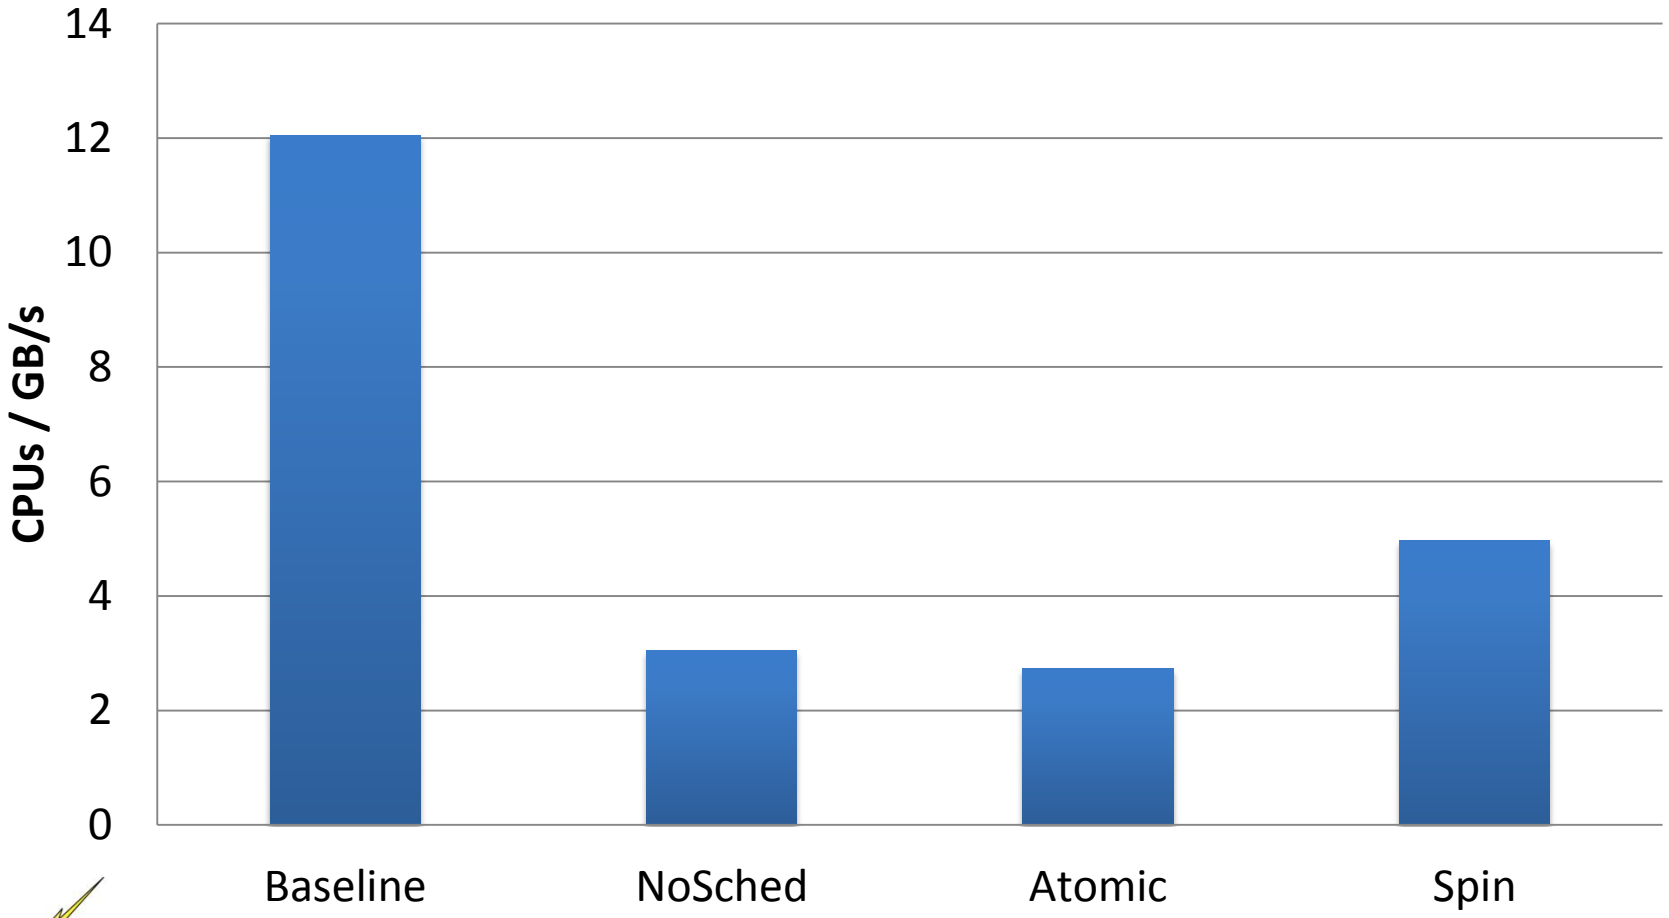

Overview

- Motivation
- Moneta Architecture
- Optimizing Moneta
  - Interface & Software
  - **Microarchitecture**
- Performance Analysis
- Conclusion

Balancing Bandwidth Usage

- Full duplex PCIe should see better R/W BW
- Smarter HW request scheduling = more BW
  - Two request queues: one for reads, one for writes
  - Alternate between Qs on each buffer allocation

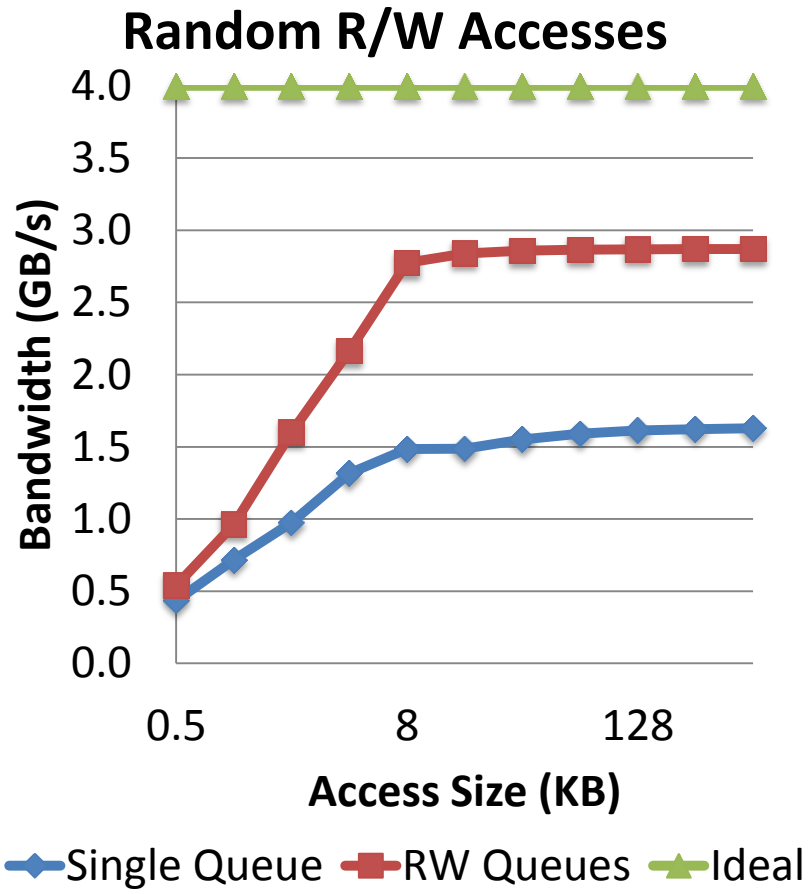

Round-Robin Scheduling

- Prevent requests from starving other requests
- Allocate a buffer and then put request at back of queue
- Attains much better small request throughput in the presence of large requests
- 12x improvement in small request BW

NVMs for Storage vs DRAM Replacement

- Write coalescing
  - Storage must guarantee durability by closing row
  - DRAM leaves row open to enable coalescing
- Row buffer size should match access size
  - Large accesses for storage
  - Cache-line sized for memory
- Peak memory activity limited by PCIe BW
- Storage and DRAM replacement are different and should be optimized differently

Overview

- Motivation
- Moneta Architecture
- Optimizing Moneta
- **Performance Analysis**
- Conclusion

System Overview

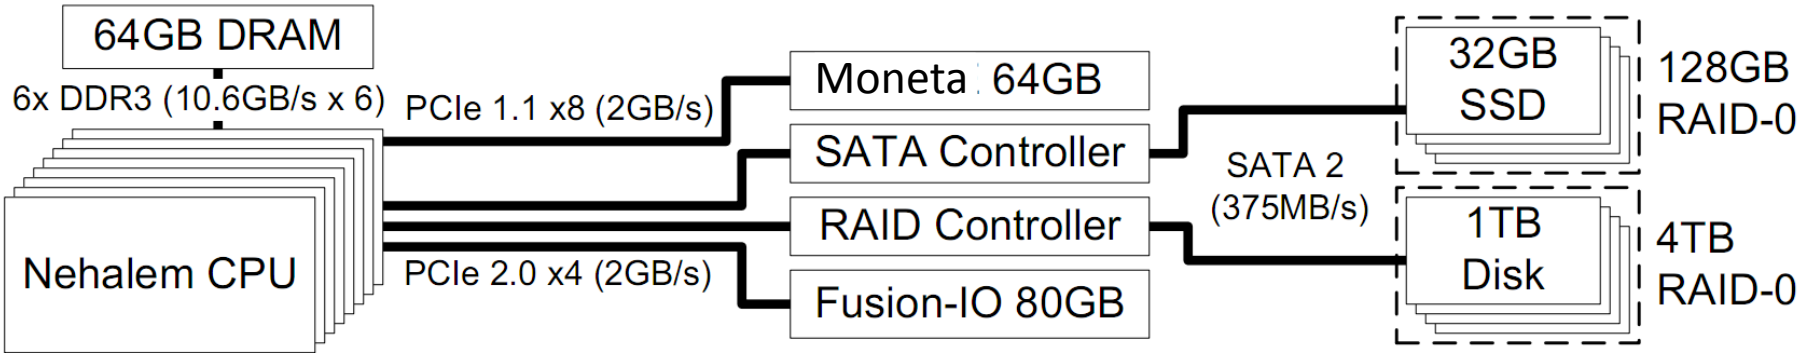

Memory and Device	Interconnect	Capacity
Fusion-IO IODrive	PCIe 4x	80GB
SLC NAND Flash SW RAID-0	PCIe 4x SATA 2 Controller	128GB
Disk HW RAID-0	PCIe 4x RAID Controller	4TB
Moneta	PCIe 8x	64GB
Moneta-4x	PCIe 4x	64GB

XDD Bandwidth Comparison

50/50 Read/Write Random Accesses

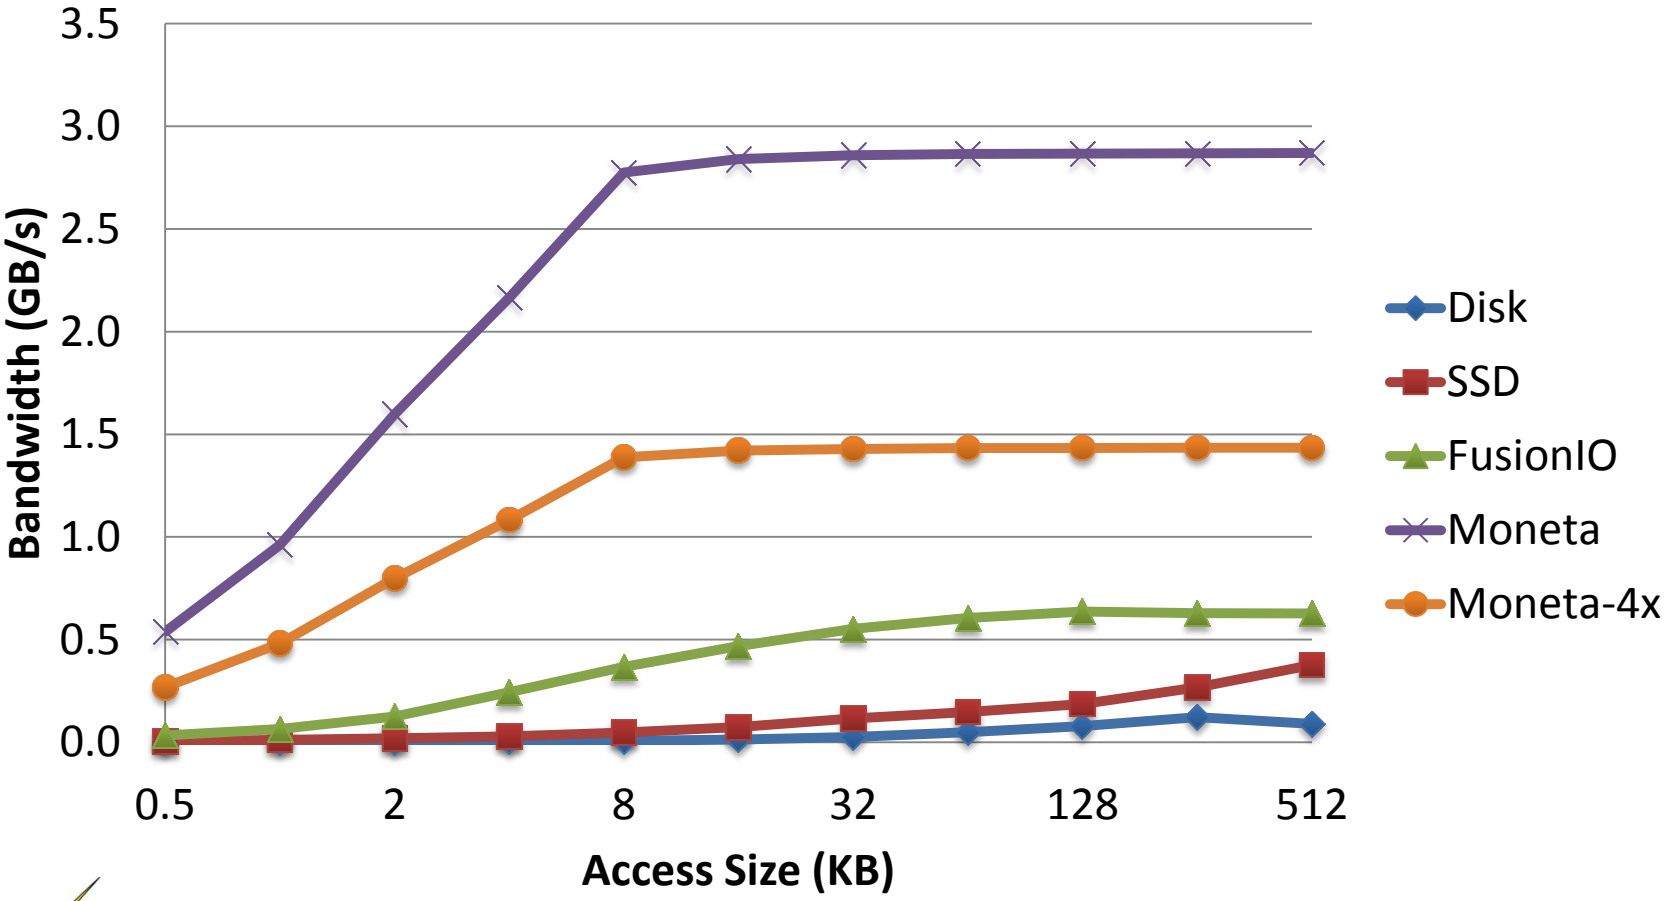

Bandwidth w/o File System

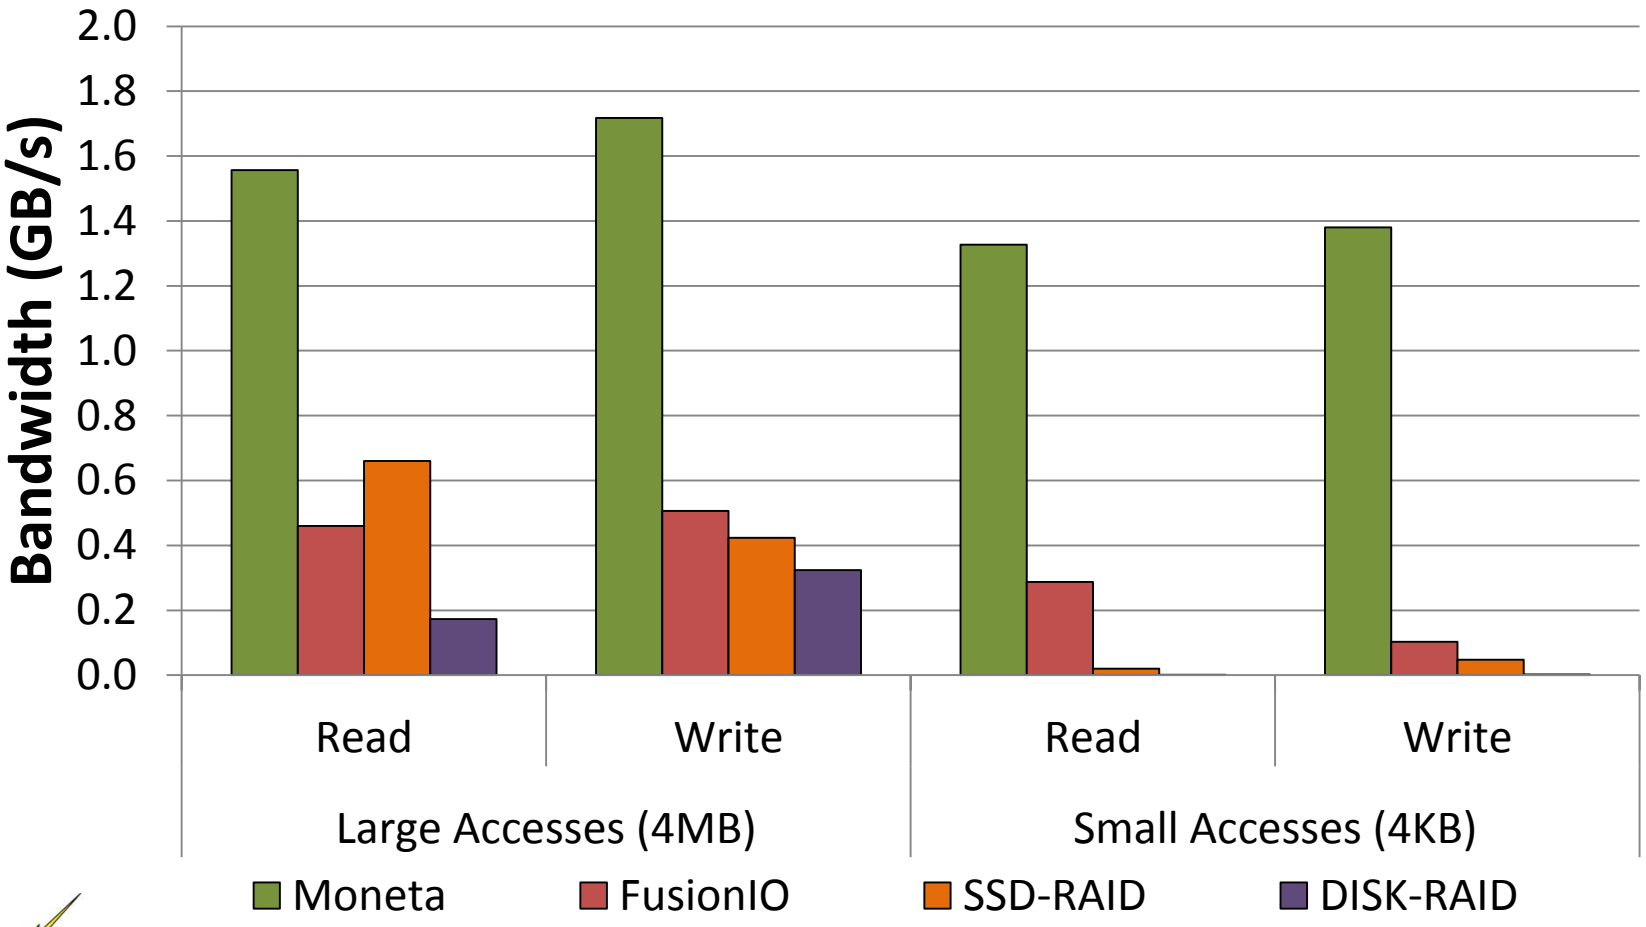

Bandwidth with XFS

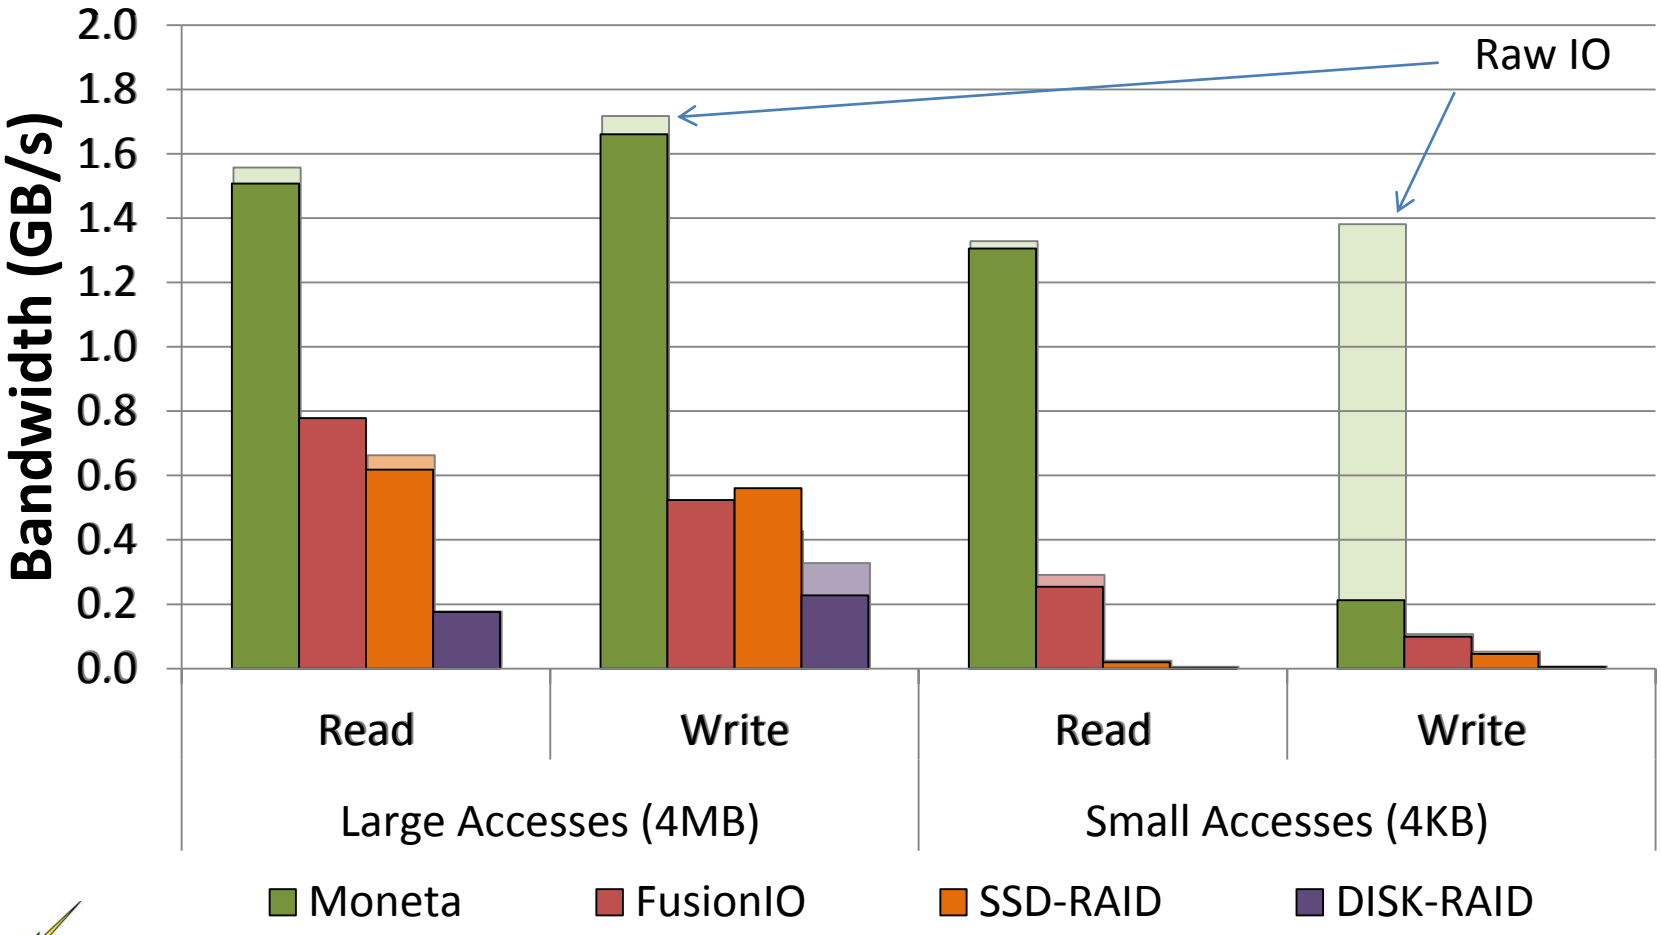

Database & App Performance

An Opportunity for Leveraging Hardware Optimizations

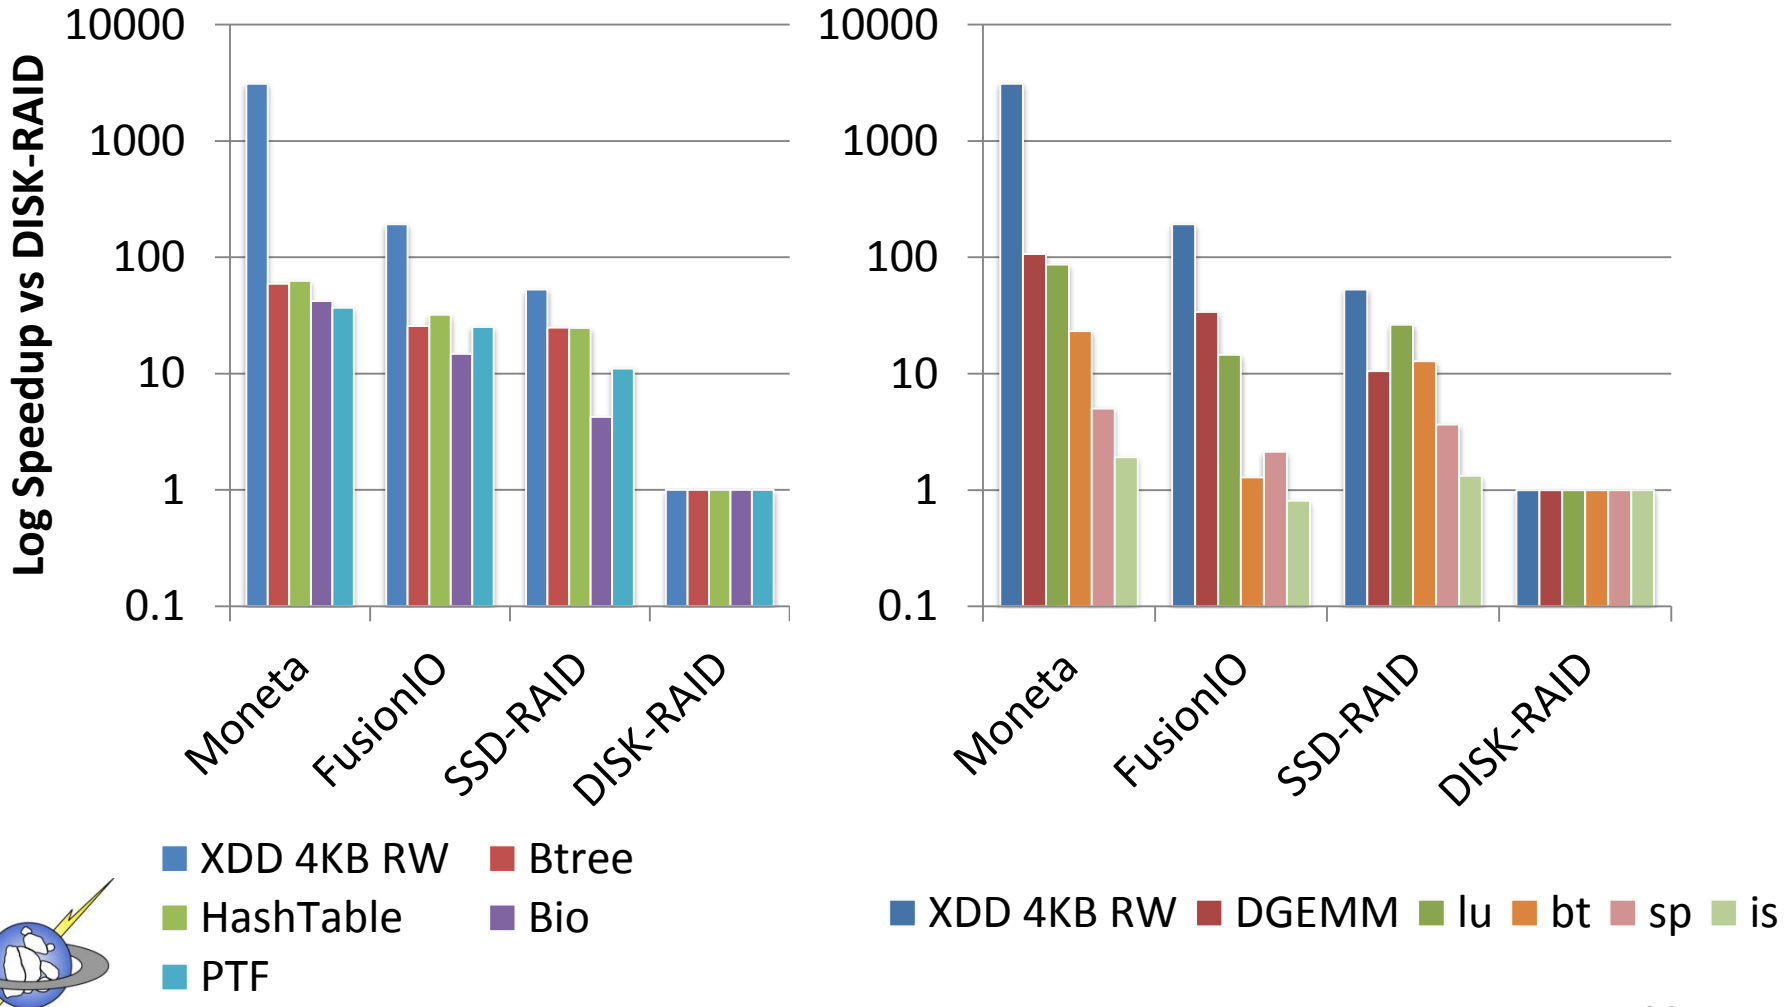

Conclusion

- Fast advanced NVMs are coming soon
- We built Moneta to understand the impact of NVMs
  - Found that the interface and microarchitecture are both critical to getting excellent performance
  - Many opportunities to move software optimizations into hardware
- Many open questions exist about the architecture of fast SSDs and the systems they interface with

Thank You!

Any Questions?

NVSL
Non-volatile Systems Laboratory

UCSDCSE
Computer Science and Engineering