


Circuit-Switched Coherence

Natalie Enright Jerger*,
Li-Shiuan Peh+, Mikko Lipasti*


*University of Wisconsin - Madison
+Princeton University

2nd IEEE International Symposium on
Networks-on-Chip




Motivation

- Network on Chip for general purpose multi-core
 - Replacing dedicated global wires
 - Efficient/scalable communication on-chip
- Router latency overhead can be significant
 - Exploit application characteristics to lower latency
- Co-design coherence protocol to match network functionality



Executive Summary

- Hybrid Network
 - Interleaves circuit-switched and packet-switched flits
 - Optimize setup latency
 - Improve throughput over traditional circuit-switching
 - Reduce interconnect delay by up to 22%
- Co-design cache coherence protocol
 - Improves performance by up to 17%




Switching Techniques

- Packet Switching
 - Efficient bandwidth utilization

Best of both worlds?

Efficient bandwidth utilization + low latency

- Avoids router overhead after circuit is established

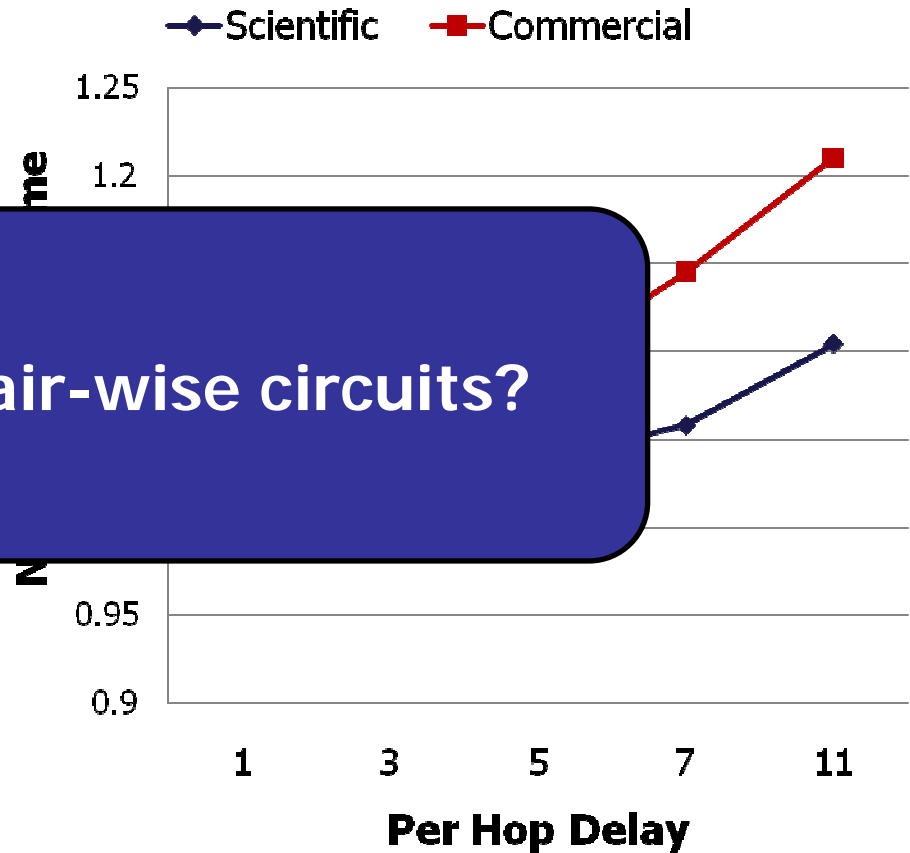


Circuit-Switched Coherence


■ Two key observations

- Coherence are correlated
- Significant sharing


Construct fast pair-wise circuits?



Commercial Workloads: SpecJBB, SpecWeb, TPC-H, TPC-W
Scientific Workloads: Barnes-Hut, Ocean, Radiosity, Raytrace




Traditional Circuit Switching




- Traditional circuit-switching hurts performance by up to ~7%

*Data collected for 16 in-order core chip multiprocessor




Circuit Switching Redesigned

- Latency is critical
- Utilize Circuit Switching for lower latency
 - A circuit connects resources across multiple hops to avoid router overhead
- Traditional circuit-switching performs poorly
- My contributions
 - Novel setup mechanism
 - Bandwidth stealing

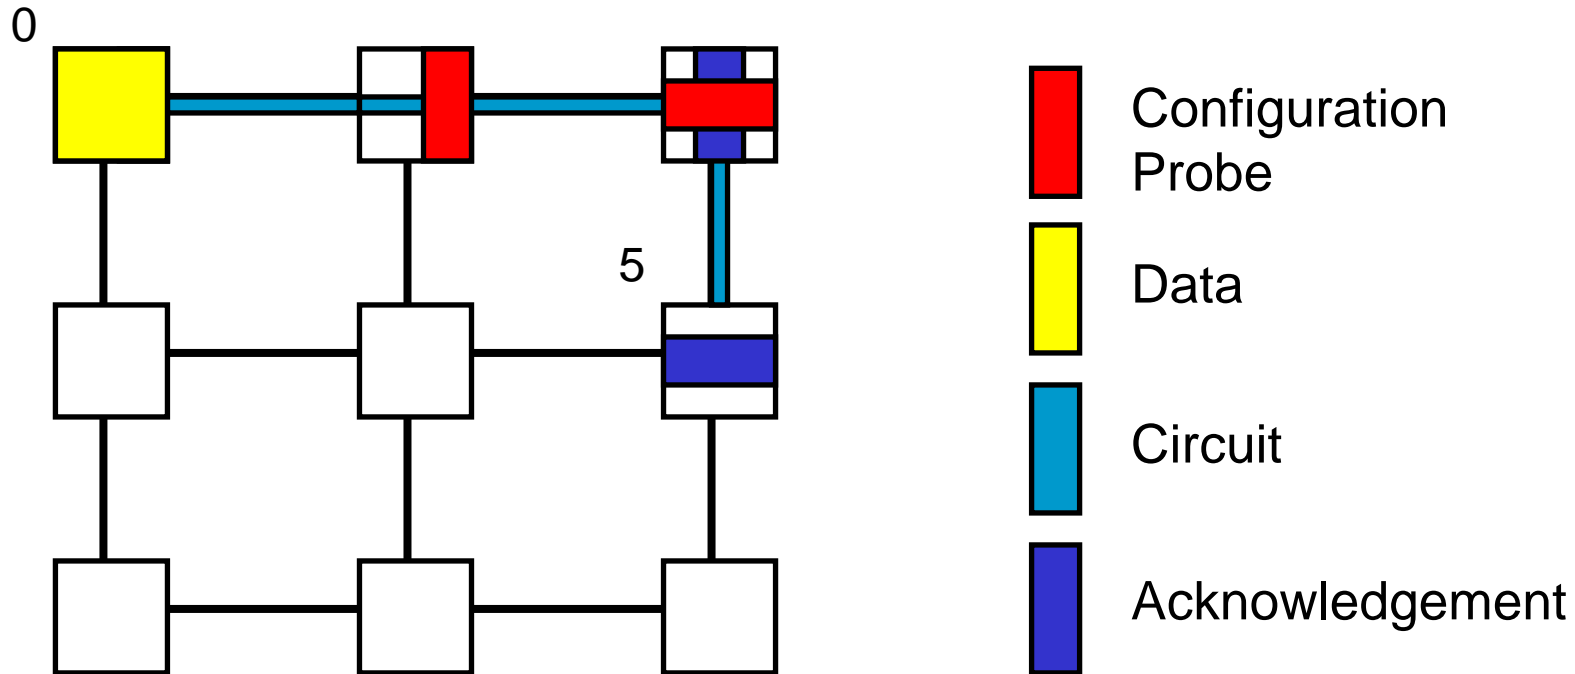


Outline


- Motivation
- Router Design
 - Setup Mechanism
 - Bandwidth Stealing
- Coherence Protocol Co-design
 - Pair-wise sharing
 - 3-hop optimization
 - Region prediction
- Results
- Conclusions



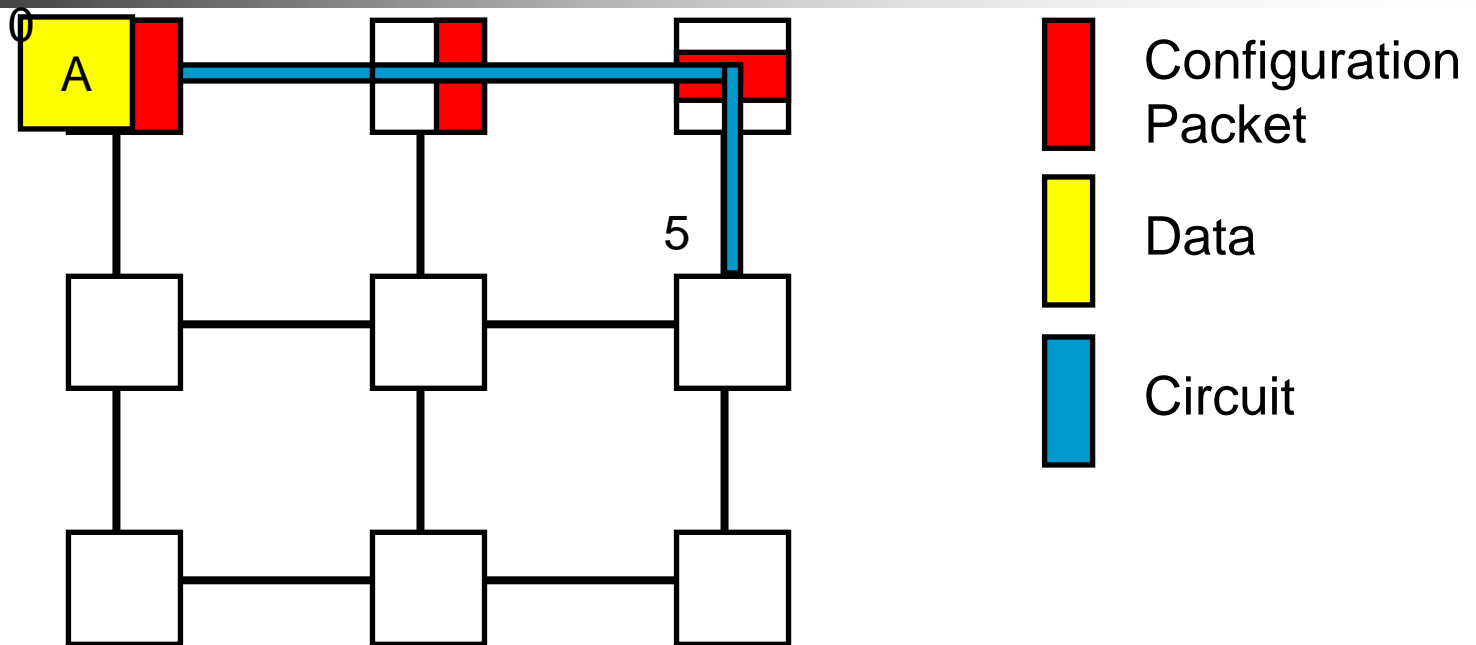
Traditional Circuit Switching Path Setup (with Acknowledgement)




- Significant latency overhead prior to data transfer
- Other requests forced to wait for resources



Novel Circuit Setup Policy




- Overlap circuit setup with 1st data transfer
- Reconfigure existing circuits if no unused links available
 - Allows piggy-backed request to *always* achieve low latency
 - Multiple circuit planes prevent frequent reconfiguration




Setup Network

- Light-weight setup network
 - Narrow
 - Circuit plane identifier (2 bits) +
 - Destination (4 bits)
- Low Load
 - No virtual channels → small area footprint
- Stores circuit configuration information
 - Multiple narrow circuit planes prevent frequent reconfiguration
- Reconfiguration
 - Buffered, traverses packet-switched pipeline

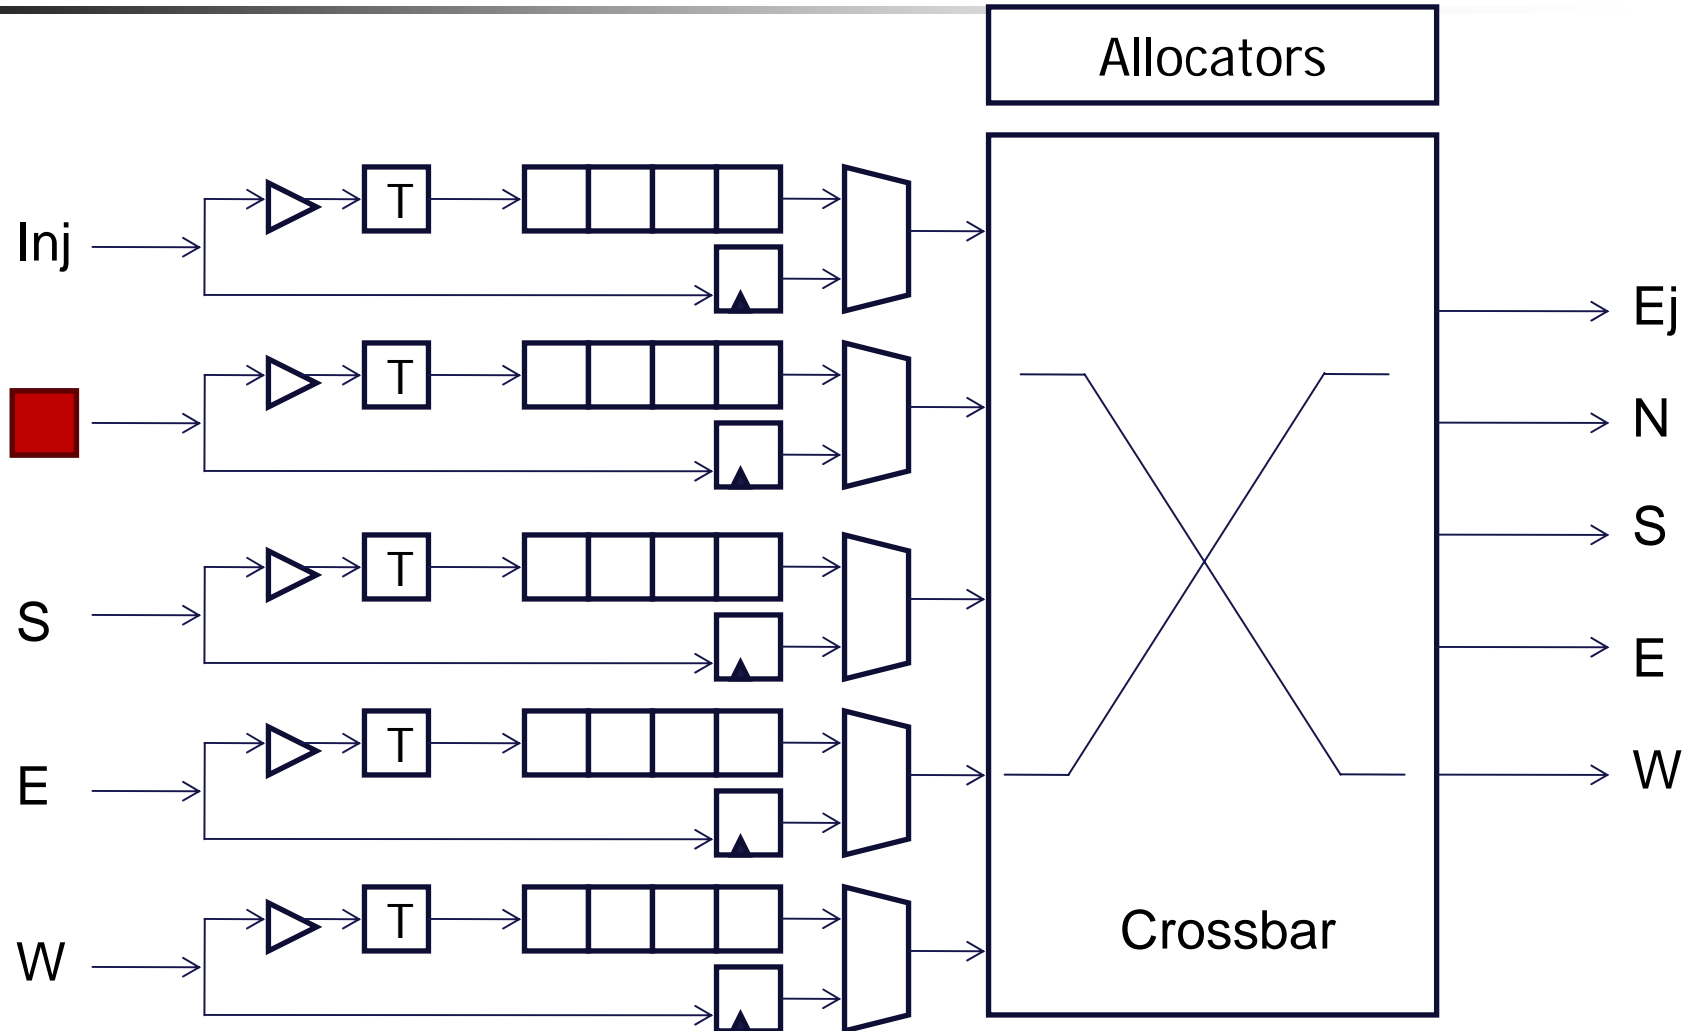



Packet-Switched Bandwidth Stealing

- Remember: problem with traditional Circuit-Switching is **poor** bandwidth
 - Need to overcome this limitation
- Hybrid Circuit-Switched Solution: Packet-switched messages snoop incoming links
 - When there are no circuit-switched messages on the link
 - A waiting packet-switched message can **steal idle** bandwidth



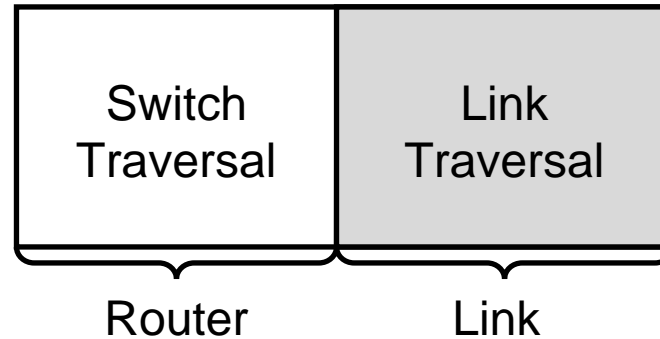
Hybrid Circuit-Switched Router Design







HCS Pipeline

- Circuit-switched messages: 1 stage




- Packet-switched messages: 3 stages
 - Aggressive Speculation reduces stages



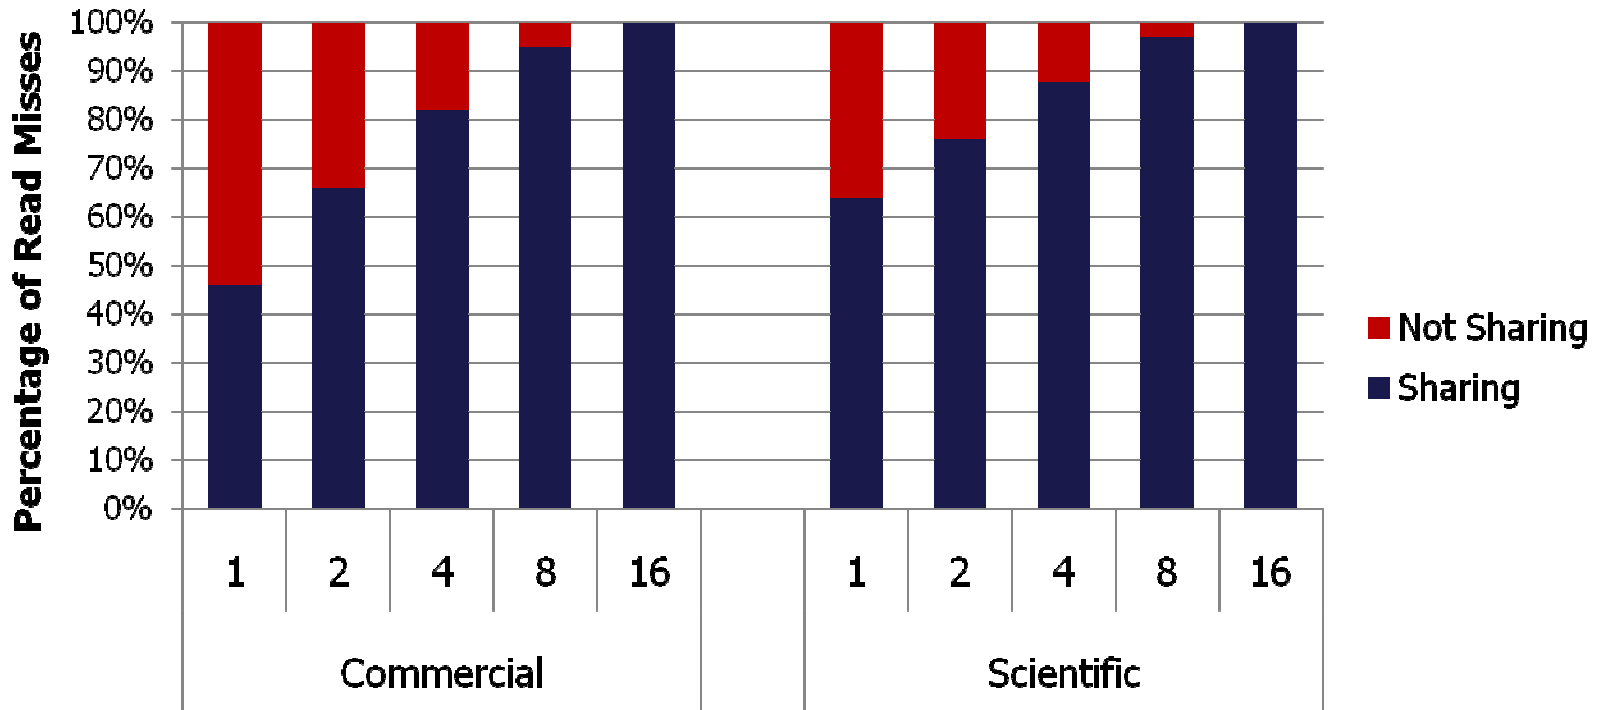


Outline

- Motivation
- Router Design
 - Setup Mechanism
 - Bandwidth Stealing
- Coherence Protocol Co-design
 - Pair-wise sharing
 - 3-hop optimization
 - Region prediction
- Results
- Conclusions




Sharing Characterization

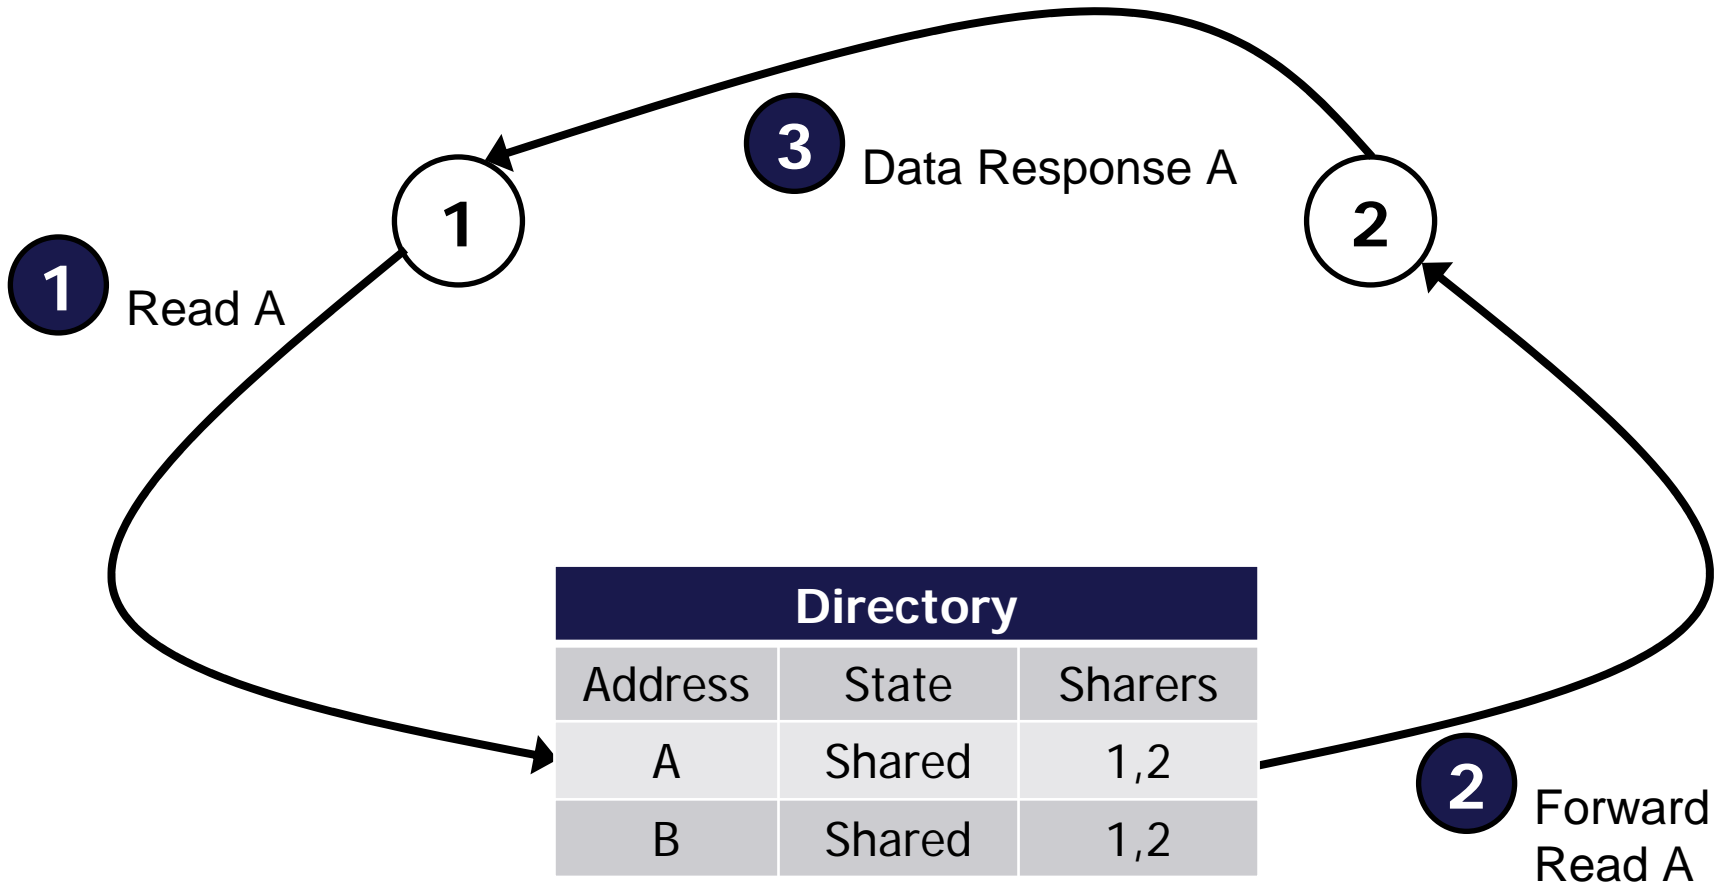



- Temporal sharing relationship: 67-76% of misses are serviced by 2 most recently shared with cores

Commercial Workloads: SpecJBB, SpecWeb, TPC-H, TPC-W
Scientific Workloads: Barnes-Hut, Ocean, Radosity, Raytrace



Directory Coherence

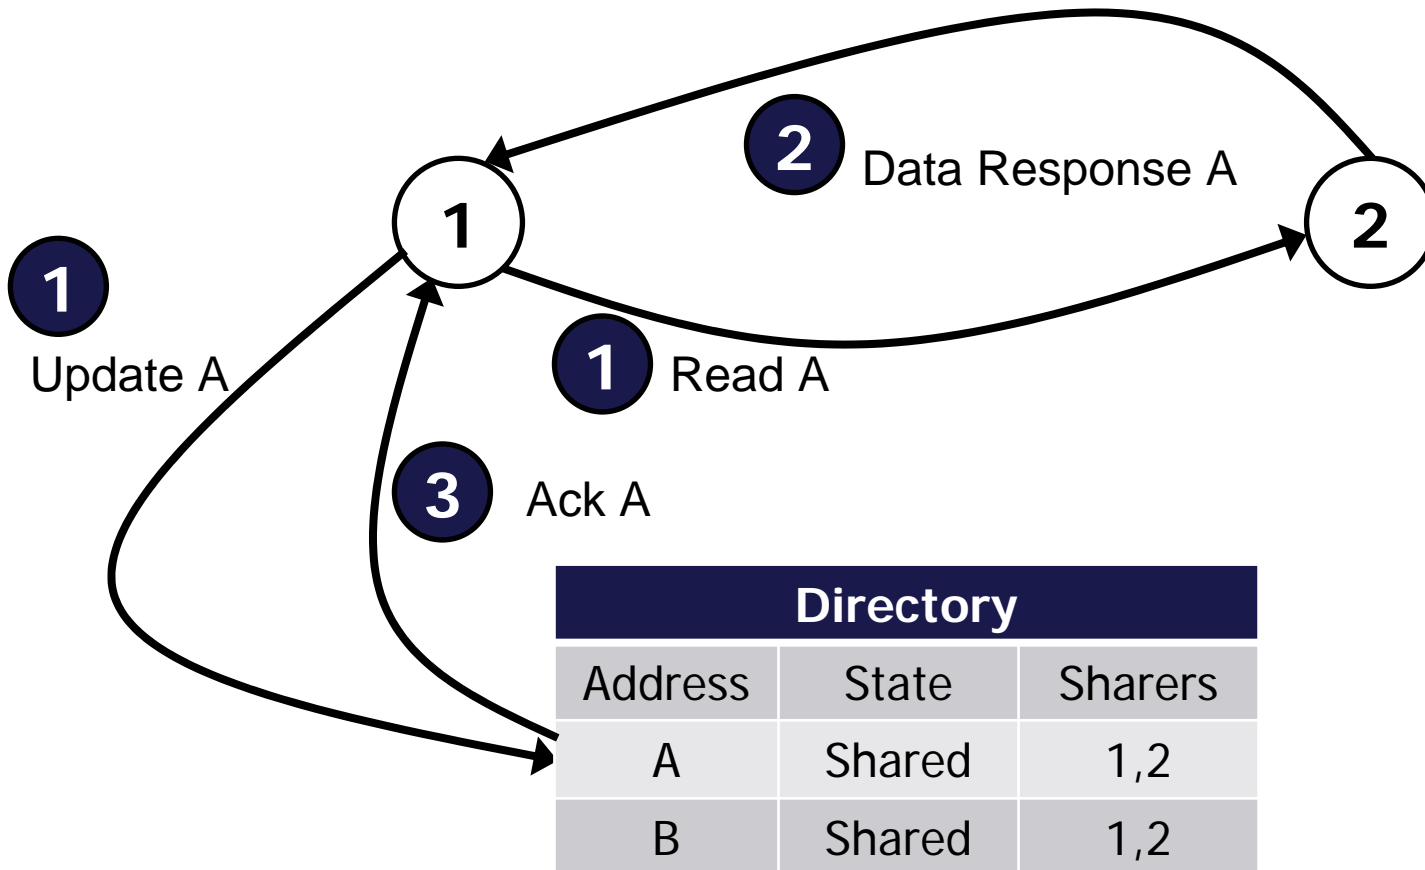






Coherence Protocol Co-Design

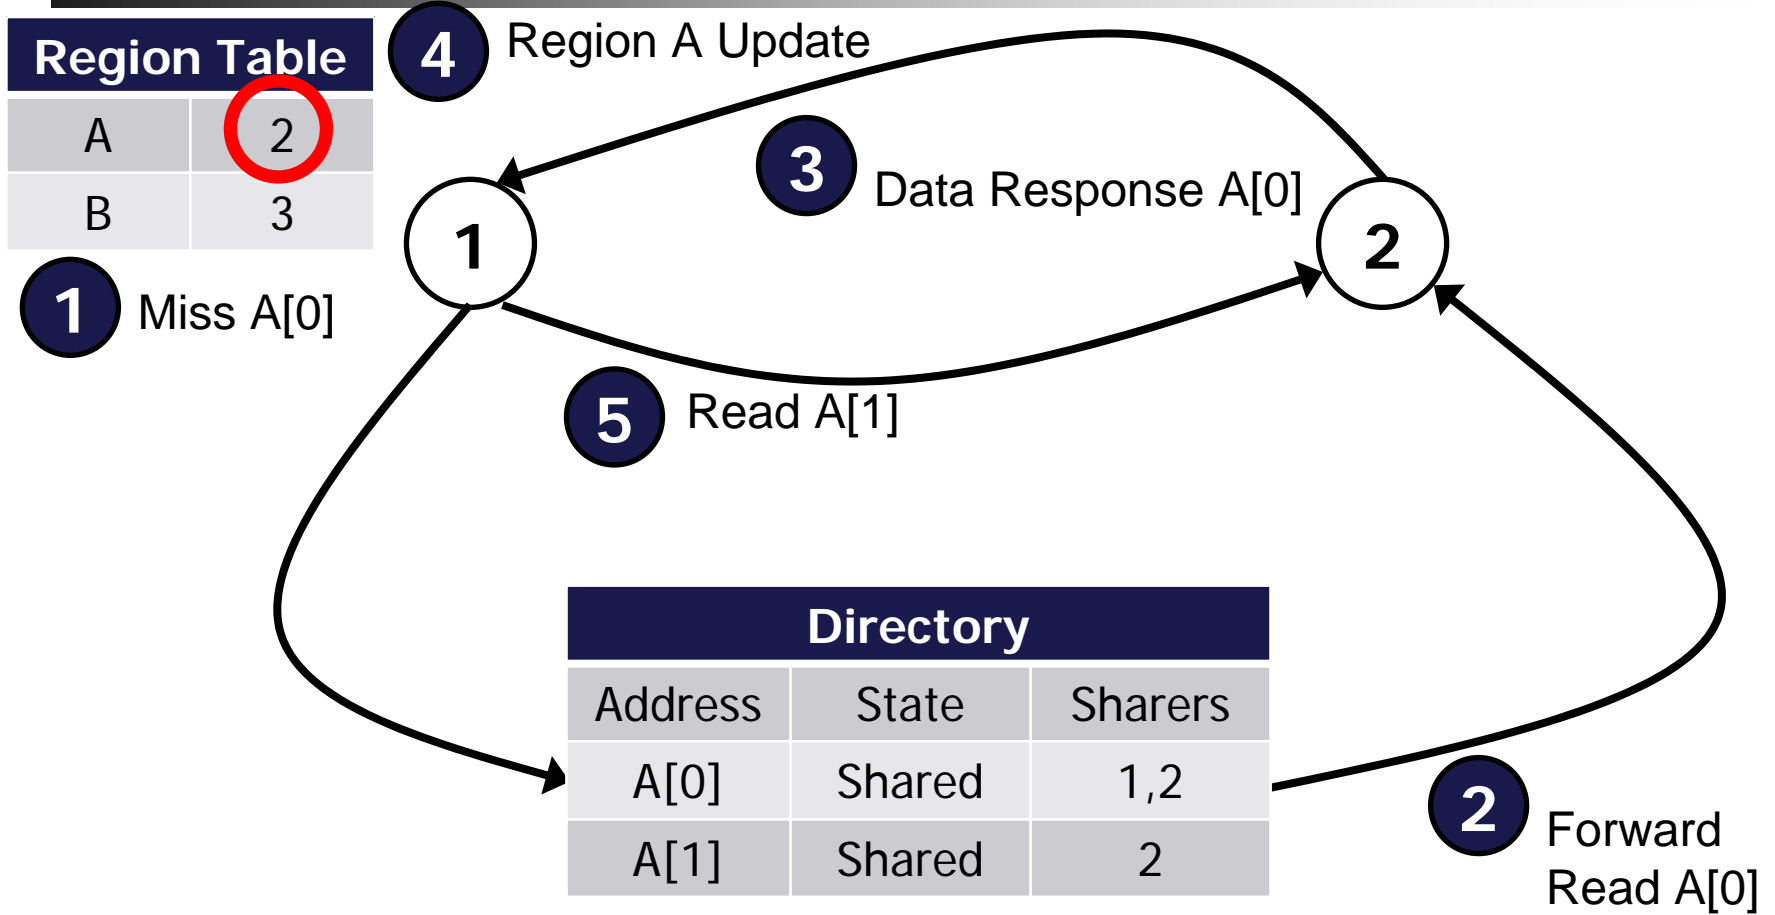
- Goal: Better exploit circuits through coherence protocol
- Modifications:
 - Allow a cache to send a request directly to another cache
 - Notify the directory in parallel
 - Prediction mechanism for pair-wise sharers
- Directory is sole ordering point

Circuit-Switched Coherence Optimization






Region Prediction




- Each memory region spans 1KB
 - Takes advantage of spatial and temporal sharing




Simulation Methodology

- PHARMSim
 - Full-system multi-core simulator
 - Detailed network level model
 - Cycle accurate router model
 - Flit-level contention modeled
- More results in paper




Simulation Workloads

Commercial	
SPECjbb	Java server workload 24 warehouse, 200 requests
SPECweb	Web server, 300 requests
TPC-W	Web e-commerce, 40 transactions
TPC-H	Decision support system
Scientific	
Barnes-Hut	8k particles, full run
Ocean	514x514, parallel phase
Radiosity	Parallel phase
Raytrace	Car input, parallel phase
Synthetic	
Uniform Random	Destination select with uniform random distribution
Permutation Traffic	Each node communicates with one other node (pair-wise)




Simulation Configuration


Processors	
Cores	16 in-order general purpose
Memory System	
L1 I/D Caches	32 KB 2-way set associative 1 cycle
Private L2 caches	512 KB 4-way set associative 6 cycles 64 Byte lines
Shared L3 Cache	16 MB (1MB bank/tile) 4-way set associative 12 cycles
Main Memory Latency	100 cycles
Interconnect: 4x4 2-D Mesh	
Packet-switched baseline	Optimized 1-3 router stages 4 Virtual channels with 4 Buffers each
Hybrid Circuit Switching	1 router stage 2 or 4 Circuit planes



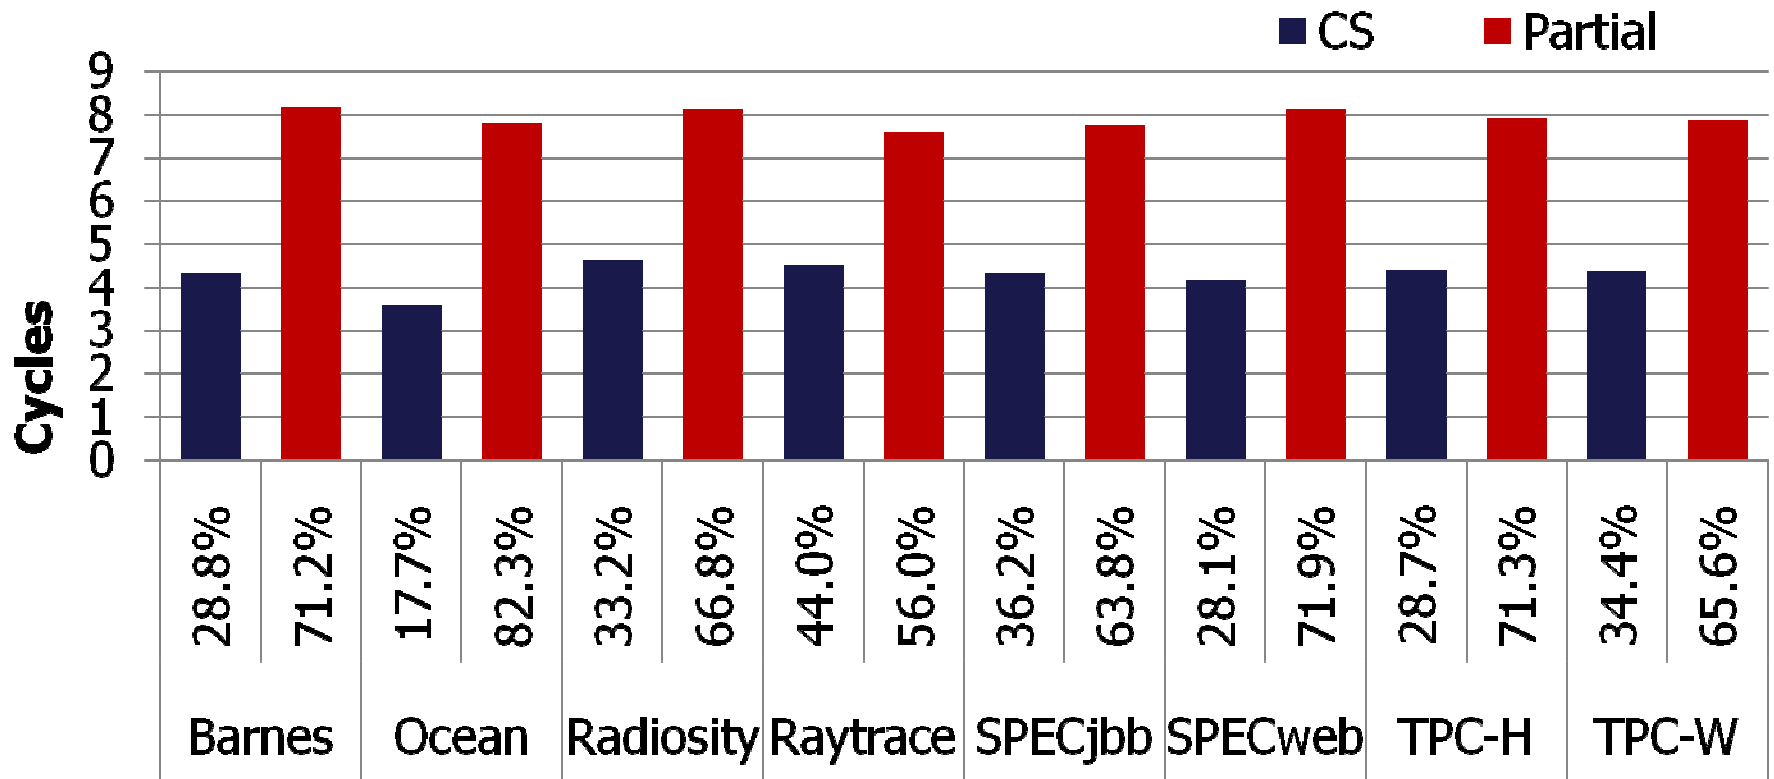
Network Results




- Communication latency is key: shave off precious cycles in network latency




Flit breakdown




- Reduce interconnect latency for a significant fraction of messages




HCS + Protocol Optimization




- Improvement of HCS + Protocol optimization is greater than the sum of HCS or Protocol Optimization alone.
 - Protocol Optimization drives up circuit reuse, better utilizing HCS



Uniform Random Traffic




- HCS successfully overcomes bandwidth limitations associated with Circuit Switching




Related Work


- Router optimizations
 - Express Virtual Channels [Kumar, ISCA 2007]
 - Single-cycle router [Mullins, ISCA 2004]
 - Many more...
- Hybrid Circuit-Switching
 - Wave-switching [Duato, ICPP 1996]
 - SoCBus [Wiklund, IPDPS 2003]
- Coherence Protocols
 - Significant research in removing overhead of indirection



Circuit-Switched Coherence Summary


- Replace packet-switched mesh with **hybrid** circuit-switched mesh
 - Interleave circuit and packet switched flits
- Reconfigurable circuits
- Dedicated bandwidth for frequent pair-wise sharers
- Low Latency and low power
 - Avoid switching/routing
- Devise novel coherence mechanisms to take advantage of benefits of circuit switching





Thank you

www.ece.wisc.edu/~pharm
enrightn@cae.wisc.edu



Circuit Setup

- Novel Setup Policy
 - Overlap circuit setup with first data transfer
 - Store circuit information at each router
 - Reconfigure existing circuits if no unused links available
 - Allows piggy-backed request to *always* achieve low latency
 - Multiple narrow circuit planes prevent frequent reconfiguration
- Reconfiguration
 - Buffered, traverses packet-switched pipeline