Coherent Network Interfaces for Fine-Grain Communication

Shubhendu S. Mukherjee

Babak Falsafi, Mark D. Hill, and David A. Wood

Wisconsin Wind Tunnel Project

University of Wisconsin-Madison

The Problem

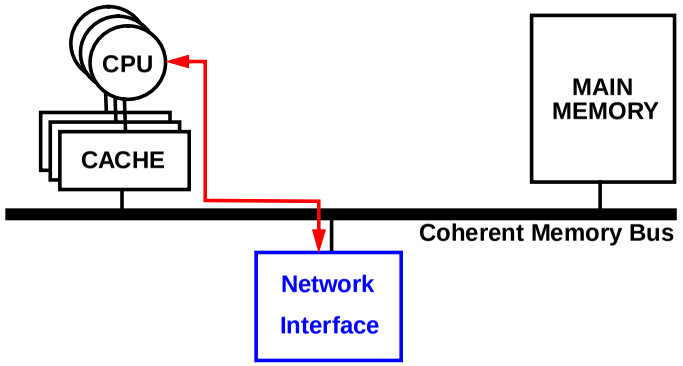

Uncached access to Network Interface is inefficient

Exploit Cache Coherence

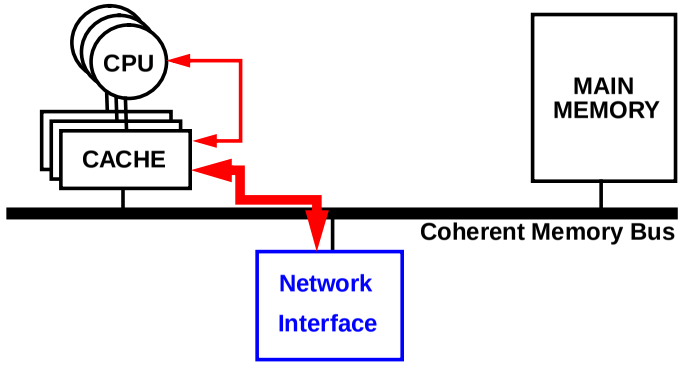

Cache status, control, & messages

Transfer message in cache block units (e.g., 32 - 128 bytes)

Coherent Network Interfaces

+ Efficient polling
+ Efficient message send/receive
+ Plentiful buffering in main memory

Significant performance improvement

- 17-53% on a coherent memory bus
- 30-88% on a coherent I/O bus

Outline

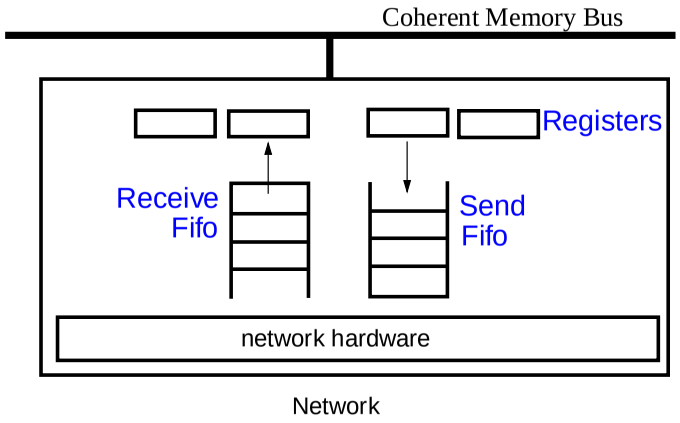

Outline

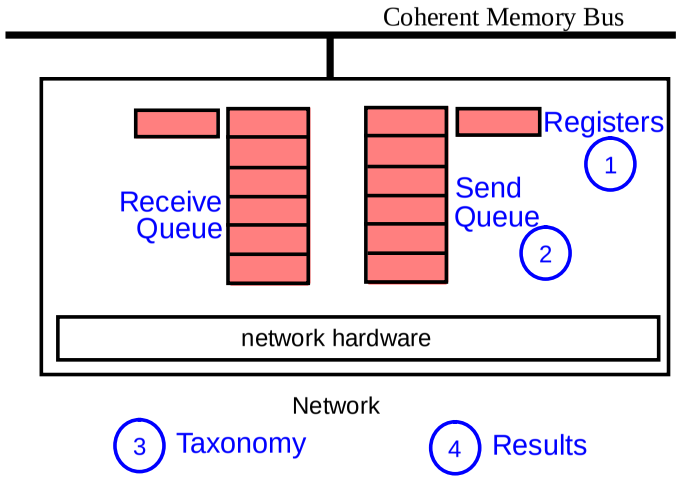

Cachable Device Registers

Generalization of CCRs [Reinhardt et al., 1996]

+ Efficient polling

+ Block Transfer

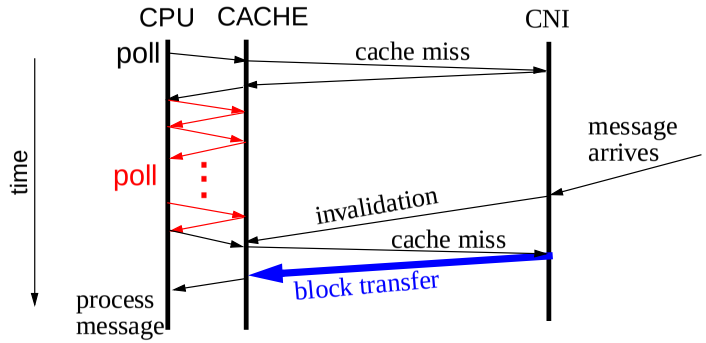

Queue

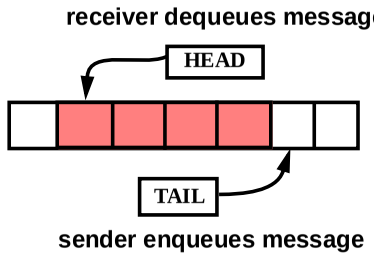

Replace hardware fifos with queues

Cachable Queue

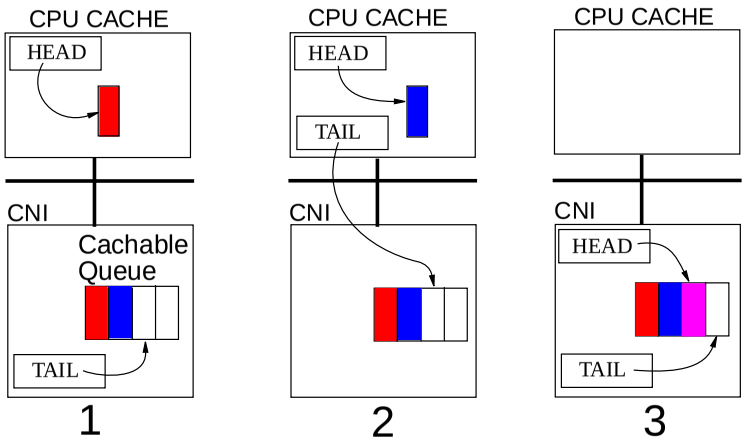

Reducing Coherence Traffic

HEAD: Lazy Pointer

sender infrequently reads HEAD

TAIL: Message Valid Bit [MIT *T-NG, Berkeley PAM]

no coherence traffic on TAIL

Queue Entries: Sense Reverse

stale message or valid message?

odd pass: 1 = valid message

even pass: 0 = valid message

Home for Cachable Queue

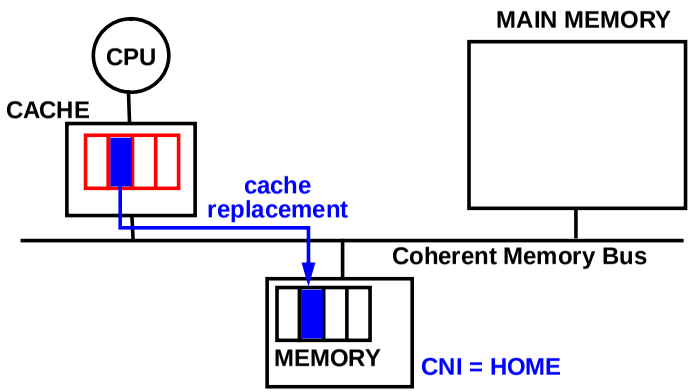

Home for Cachable Queue

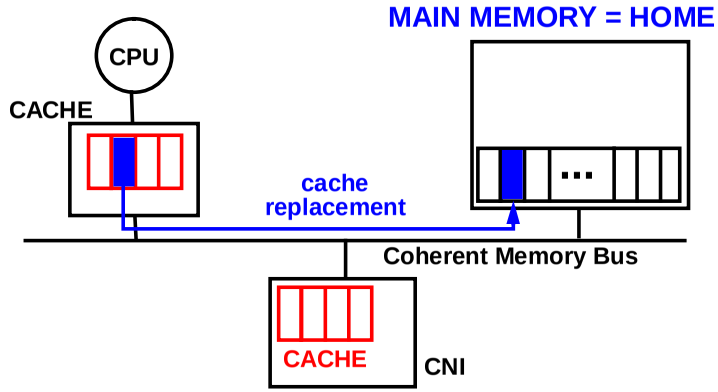

Plentiful Buffering

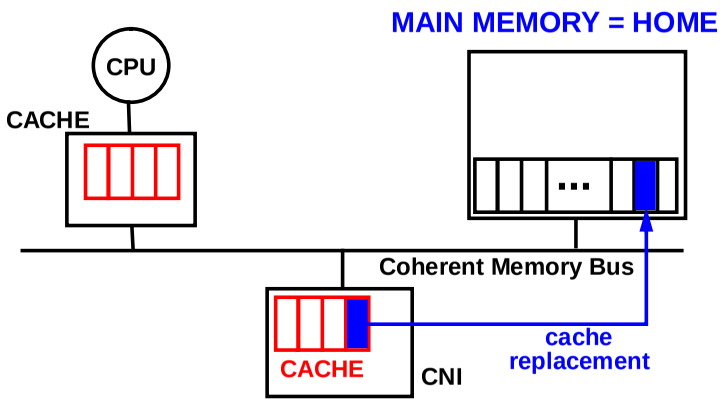

Fast Cache-To-Cache Transfer

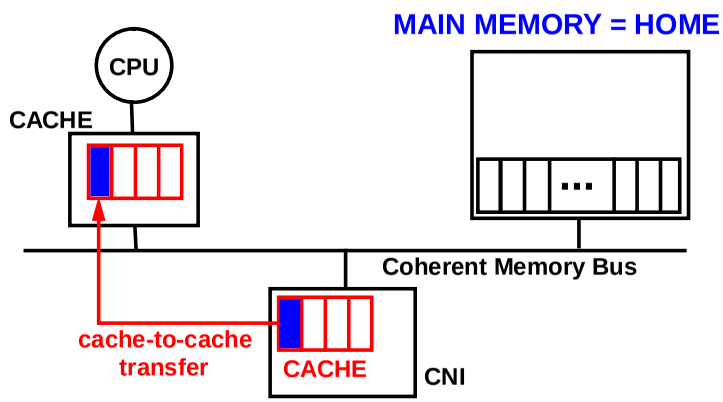

Taxonomy

NI_iX / CNI_iX

- i = exposed queue size
- X = empty, queue not exposed
  - = Q, queue explicit
  - = Q_m, queue explicit + home in main memory

NI_{2w} = ~ CM-5, uncached registers

$CNI_{16}Q$ = cachable queue

$CNI_{16}Q_m$ = cachable queue + cache in CNI

Results

Simulation parameters

- network interface on memory bus
- 100 cycle network latency
- end-to-end hardware flow control
- 64-byte cache blocks
- 256 byte network messages (4 cache blocks)
- 200 MHz dual issue hyperSPARC-like processor
- 100 MHz memory bus with MBus coherence protocol
- 16 node parallel machine

Microbenchmarks

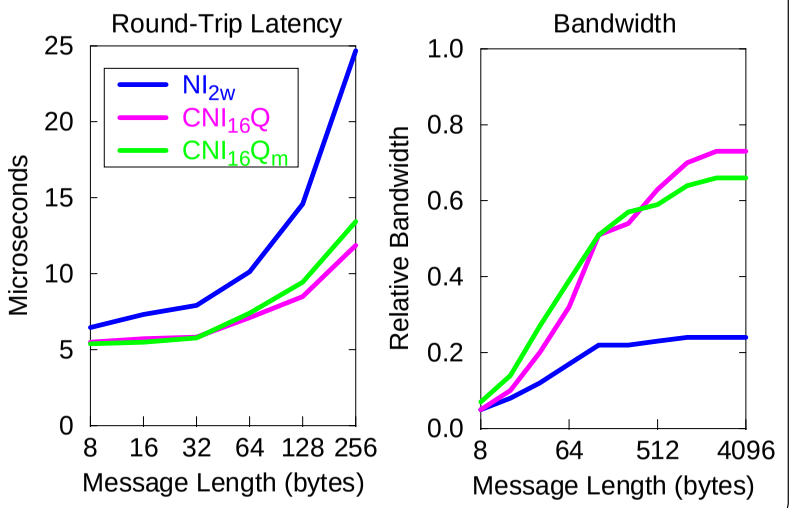

Macrobenchmarks

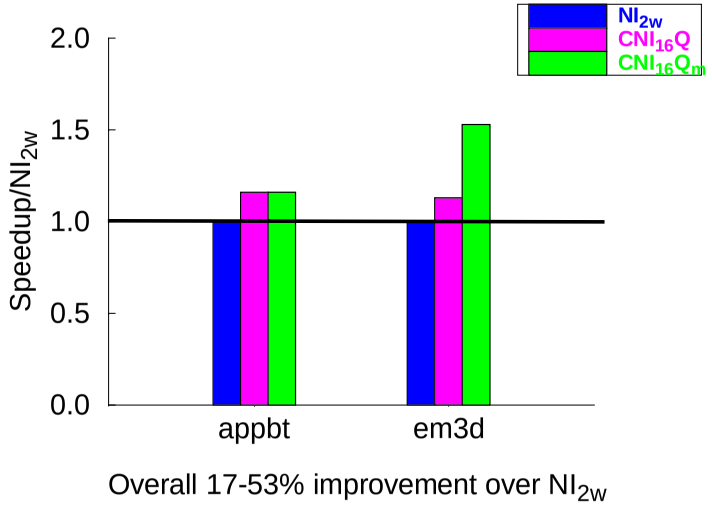

Overall 17-53% improvement over NI_{2w}

Summary

Exploit cache coherence

Use Cachable Queues to transfer messages

- \+ Efficient polling
- \+ Efficient message send/receive
  - block transfer
  - optimizations
- \+ Plentiful buffering in main memory

Wisconsin Wind Tunnel Project: http://www.cs.wisc.edu/~wwt/