Ordered Mesh Network Interconnect (OMNI) Design & Implementation of In-Network Coherence

Suvinay Subramanian, Advisor: Li-Shiuan Peh

Collaborators: Chia-Hsin Owen Chen, Bhavya Daya, Tushar Krishna, Woo-Cheol Kwon, Sunghyun Park

MIT EECS — ELECTRICAL ENGINEERING AND COMPUTER SCIENCE

Introduction

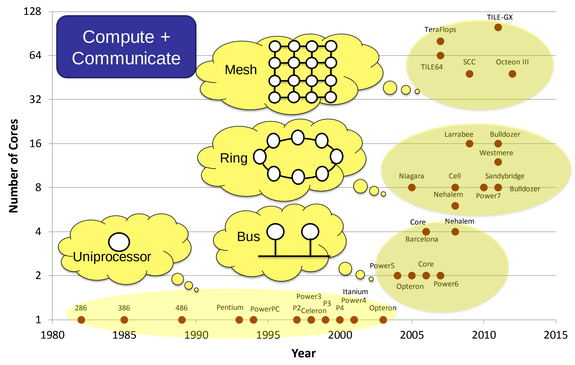

- Moore's Law "The number of transistors incorporated in a chip will approximately double every 24 months" → Increased compute density
- Shift to multi-core processors → better performance vis-à-vis power efficiency
- **WE NEED** A scalable coherence mechanism to maintain "consistent" view of "shared memory"
- **WE NEED** A scalable communication fabric

The SCORPIO Processor

Snoopy **CO**herent **R**esearch **P**rocessor with **I**nterconnect **O**rdering

Specification
- 36 Power-ISA cores
- 32 KB L1; 128 KB L2
- 6x6 Mesh network: OMNI

Why Snoopy Coherence?
- Efficient cache-to-cache transfers ☺
- No directory overhead ☺
- No indirection latency ☺

Challenge
Achieve global ordering for messages

OMNI
- Enable in-network snoopy coherence
- Scalable communication fabric

Challenge
Low latency, low power broadcast network

OMNI Design

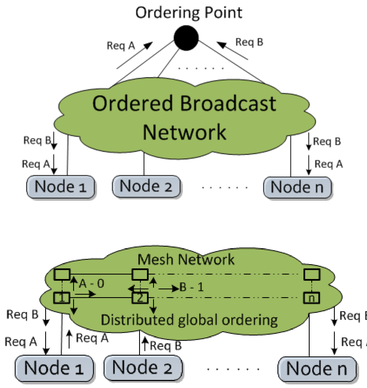

Traditional interconnects (eg: buses) route all messages to central ordering point

All nodes see all messages in the same order

OMNI: Network provides mechanism for global ordering

No central ordering point

Idea: Nodes agree on global order at synchronized time intervals

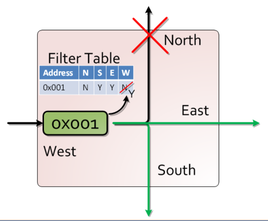

Filter redundant broadcasts to save bandwidth and power

Routers maintain & propagate sharing information

Walkthrough Example

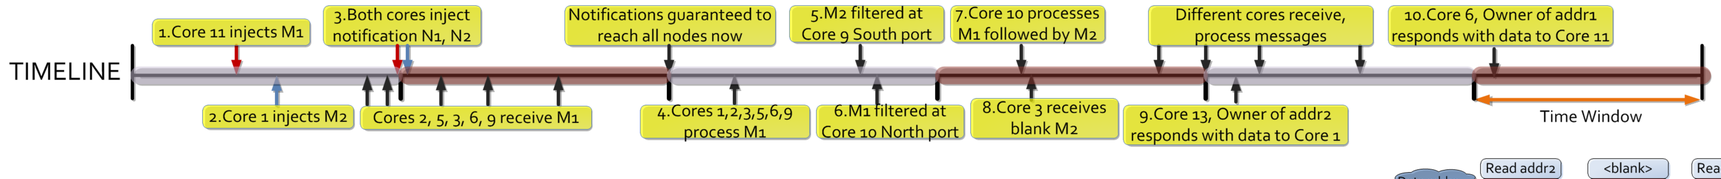

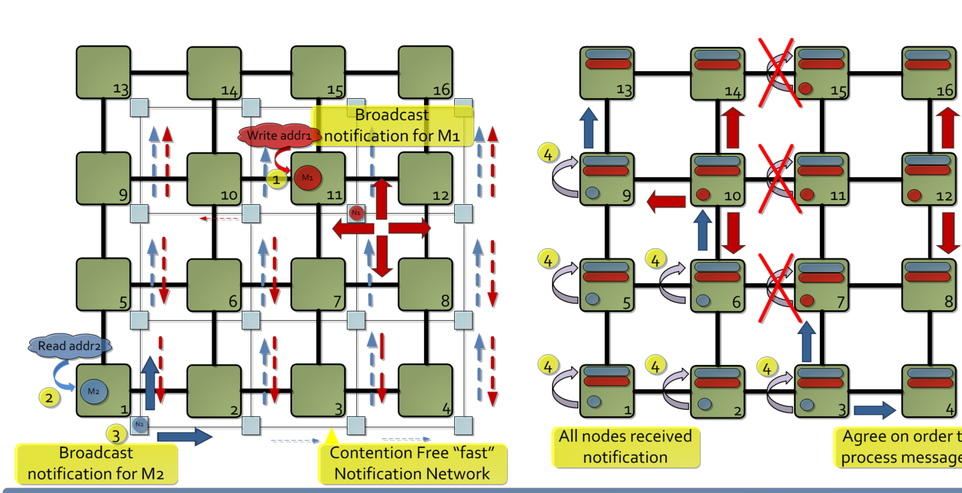

Router Microarchitecture

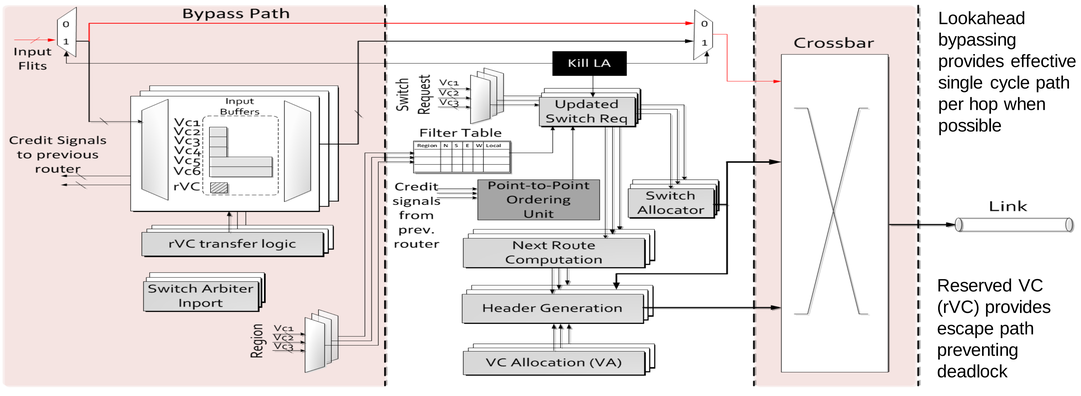

Lookahead bypassing provides effective single cycle path per hop when possible

Reserved VC (rVC) provides escape path preventing deadlock

Network Interface Block Diagram

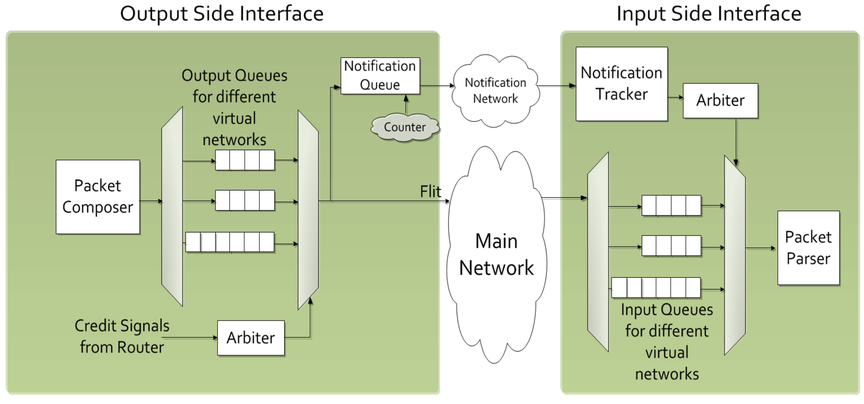

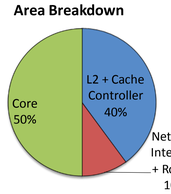

Evaluations

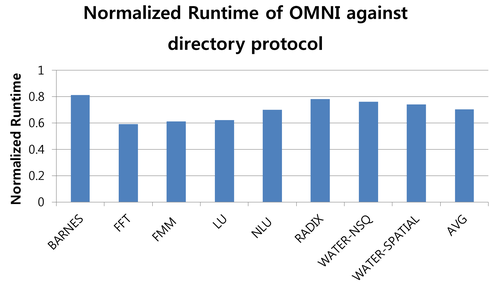

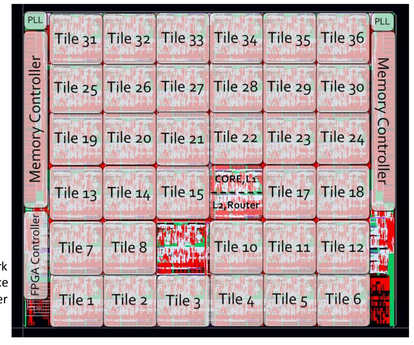

Average Power		Technology Node
Critical Path	1.2 ns	45 nm commercial

Channel Width	Number of Virtual Networks	Number of Virtual Channels
16 bytes	3	4+2+2 = 8