



SCORPIO:

A 36-Core Research Chip
Demonstrating Snoopy Coherence
on a Scalable Mesh NoC
with In-Network Ordering



Bhavya Daya

Collaborators:

Chia-Hsin Chen

Suvinay Subramanian

Woo-Cheol Kwon

Sunghyun Park

Tushar Krishna

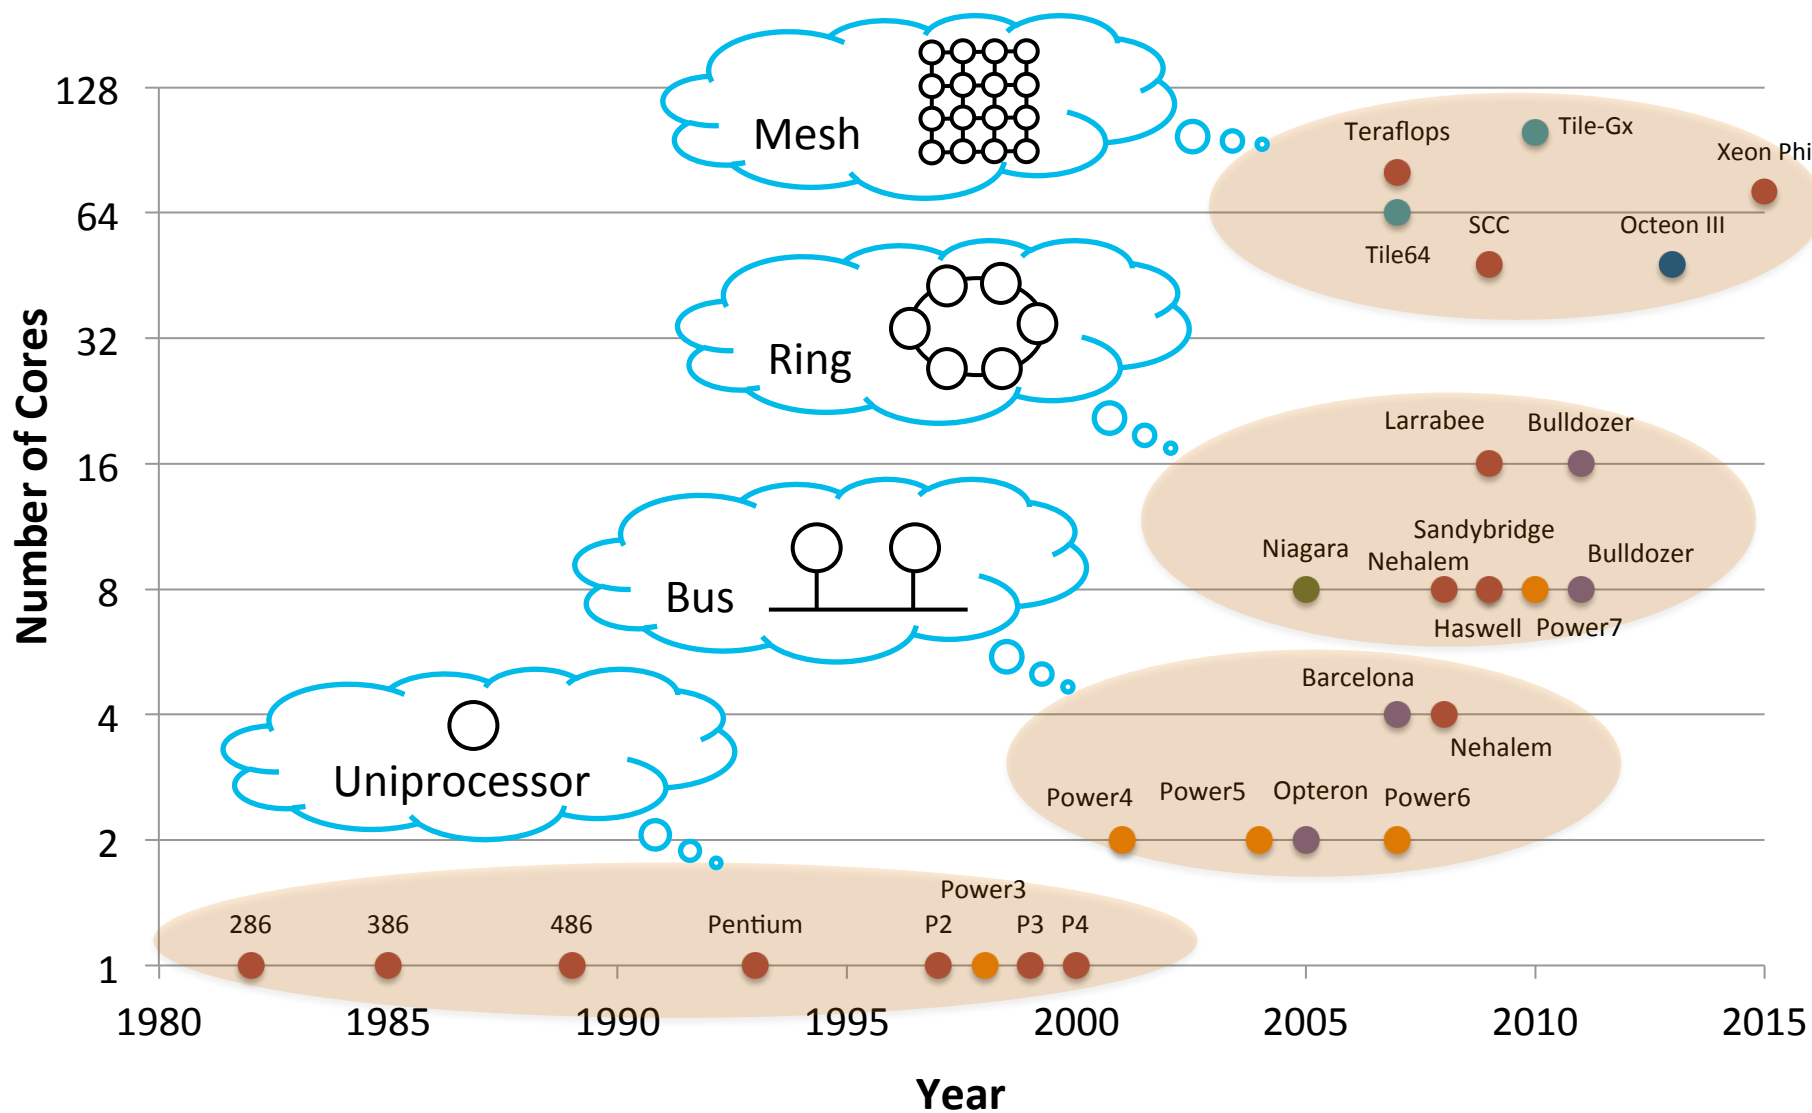
Jim Holt

Anantha Chandrakasan

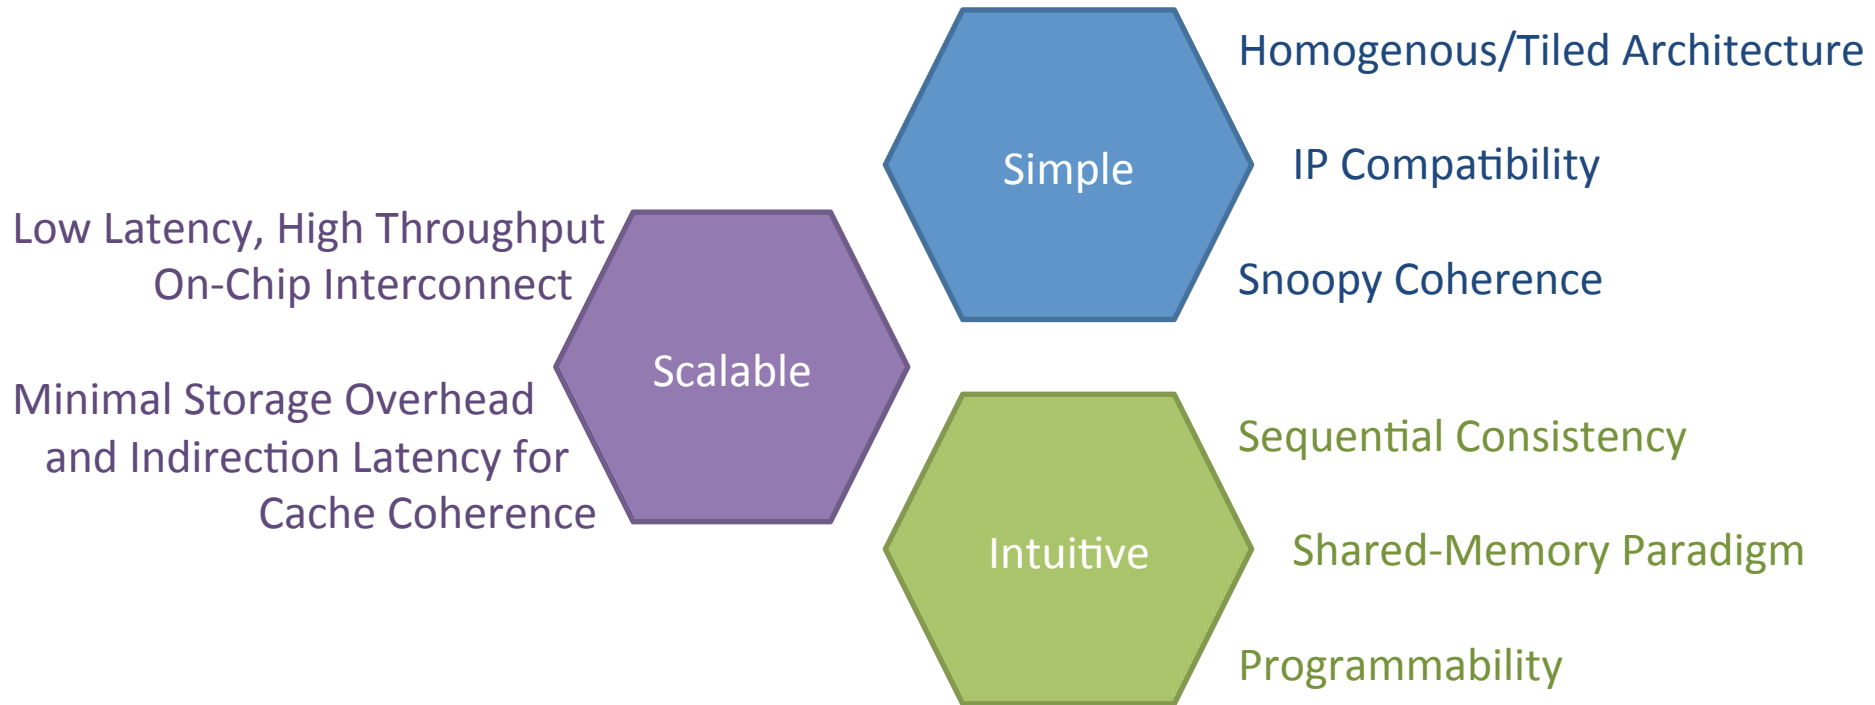
Li-Shiuan Peh



Evolution of On-Chip Networks



Objectives




Develop a simple, scalable, and intuitive multicore processor

Focus on enabling **snoopy coherence on a **mesh interconnect****



Processor Overview



**IBM 45nm SOI, 143mm²
600M transistors**

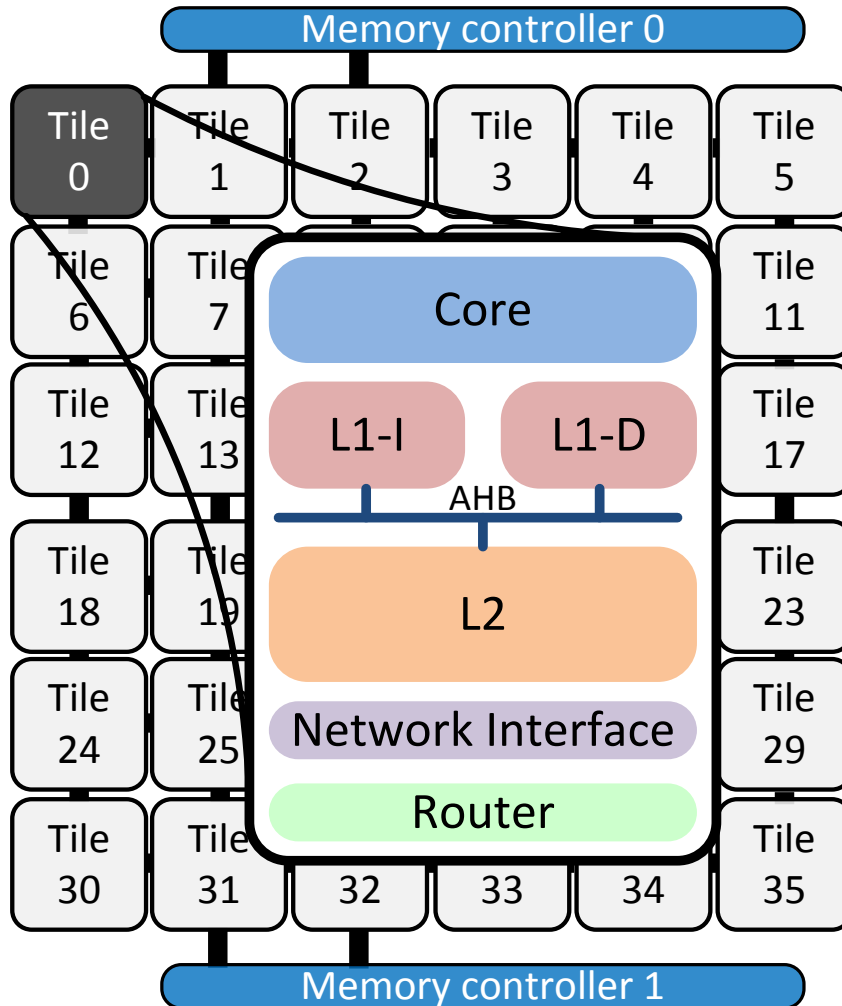
36 cores with total 4.5MB L2

**6x6 mesh on-chip network
supporting snoop coherence**

**Dual channel DDR2 memory
controllers**



Tile Architecture



Core

- Freescale e200 z760n3
- In-order
- Dual-issue

Private L1 cache

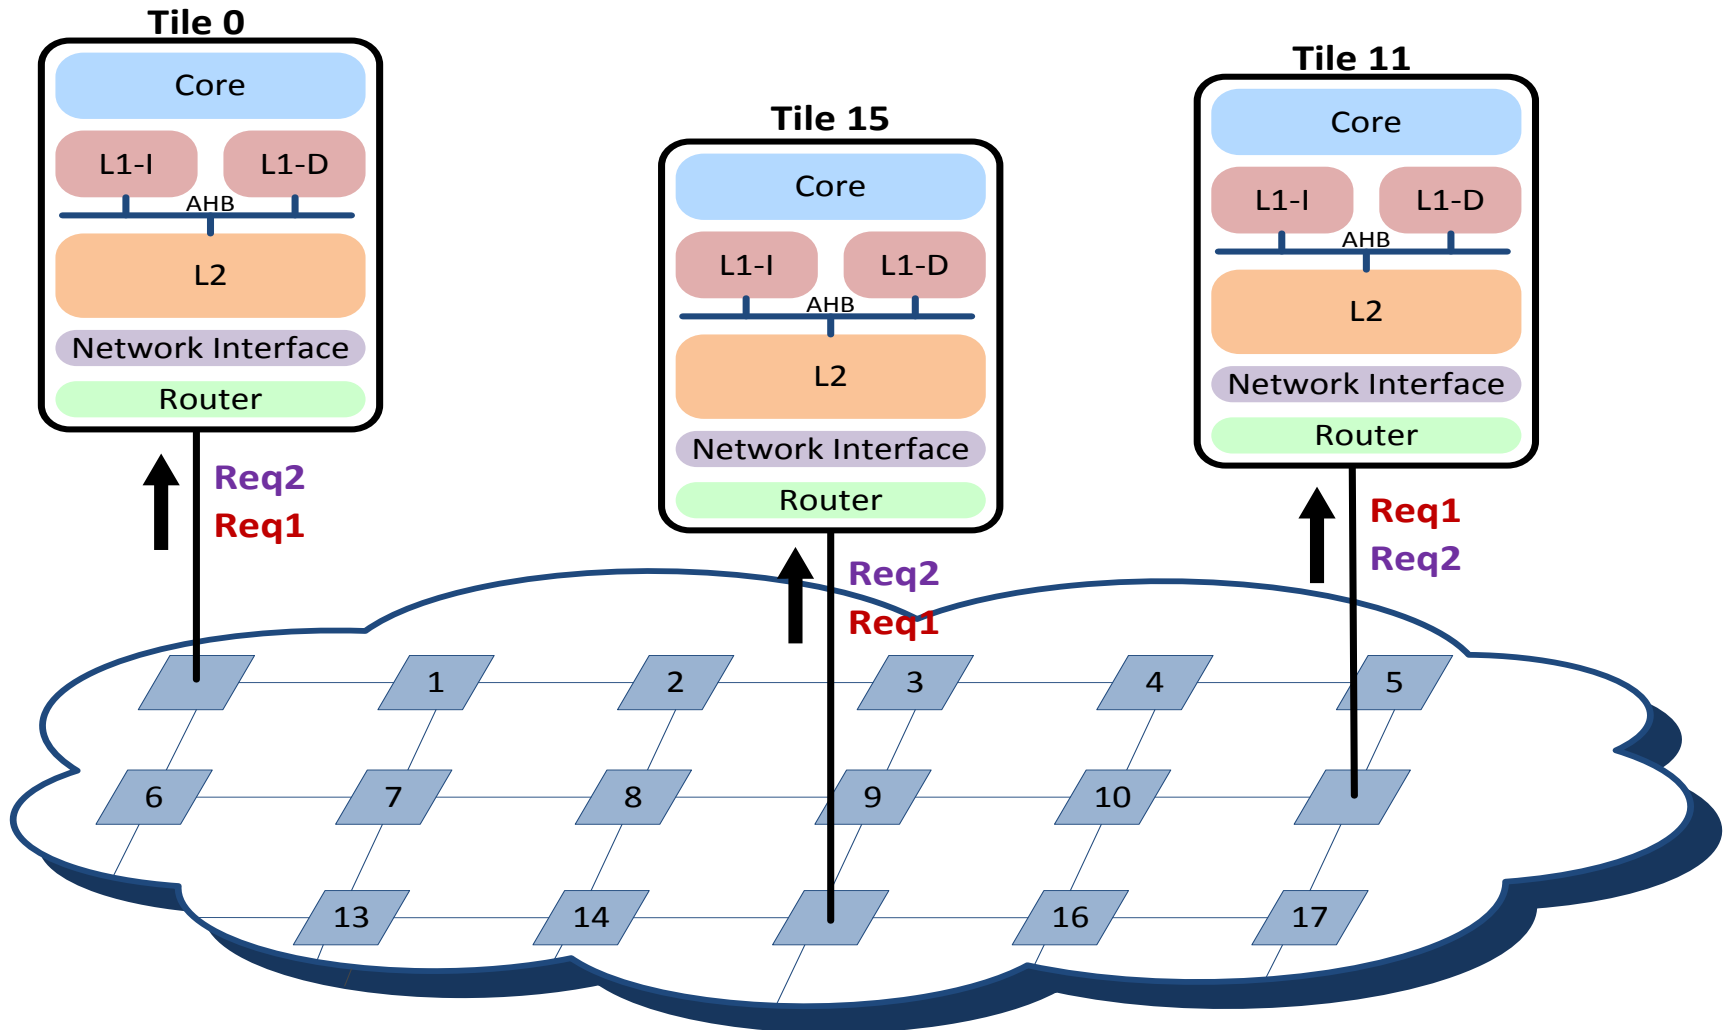
- Split 16KB for Inst / Data
- 4-way set associative

Write-through ↓ ↑ Back-invalidate

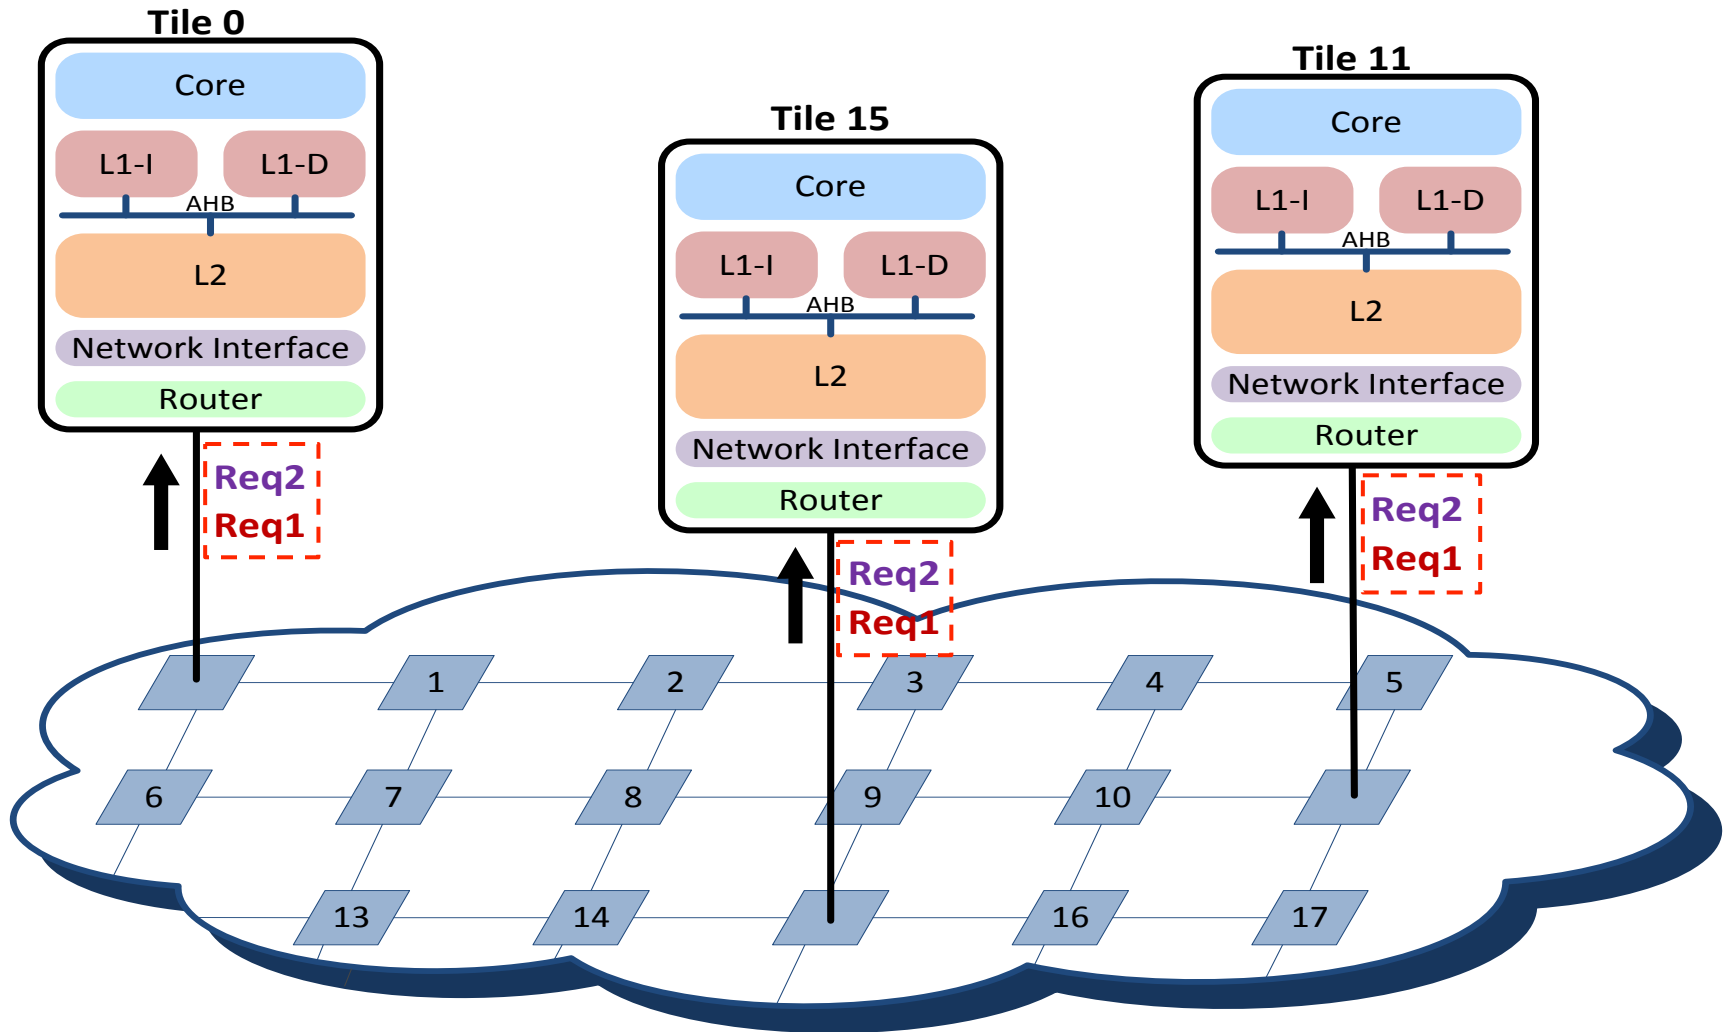
Private L2 cache

- 128KB
- 4-way set associative
- Inclusive
- MOSI Protocol

Globally Ordered Mesh Network

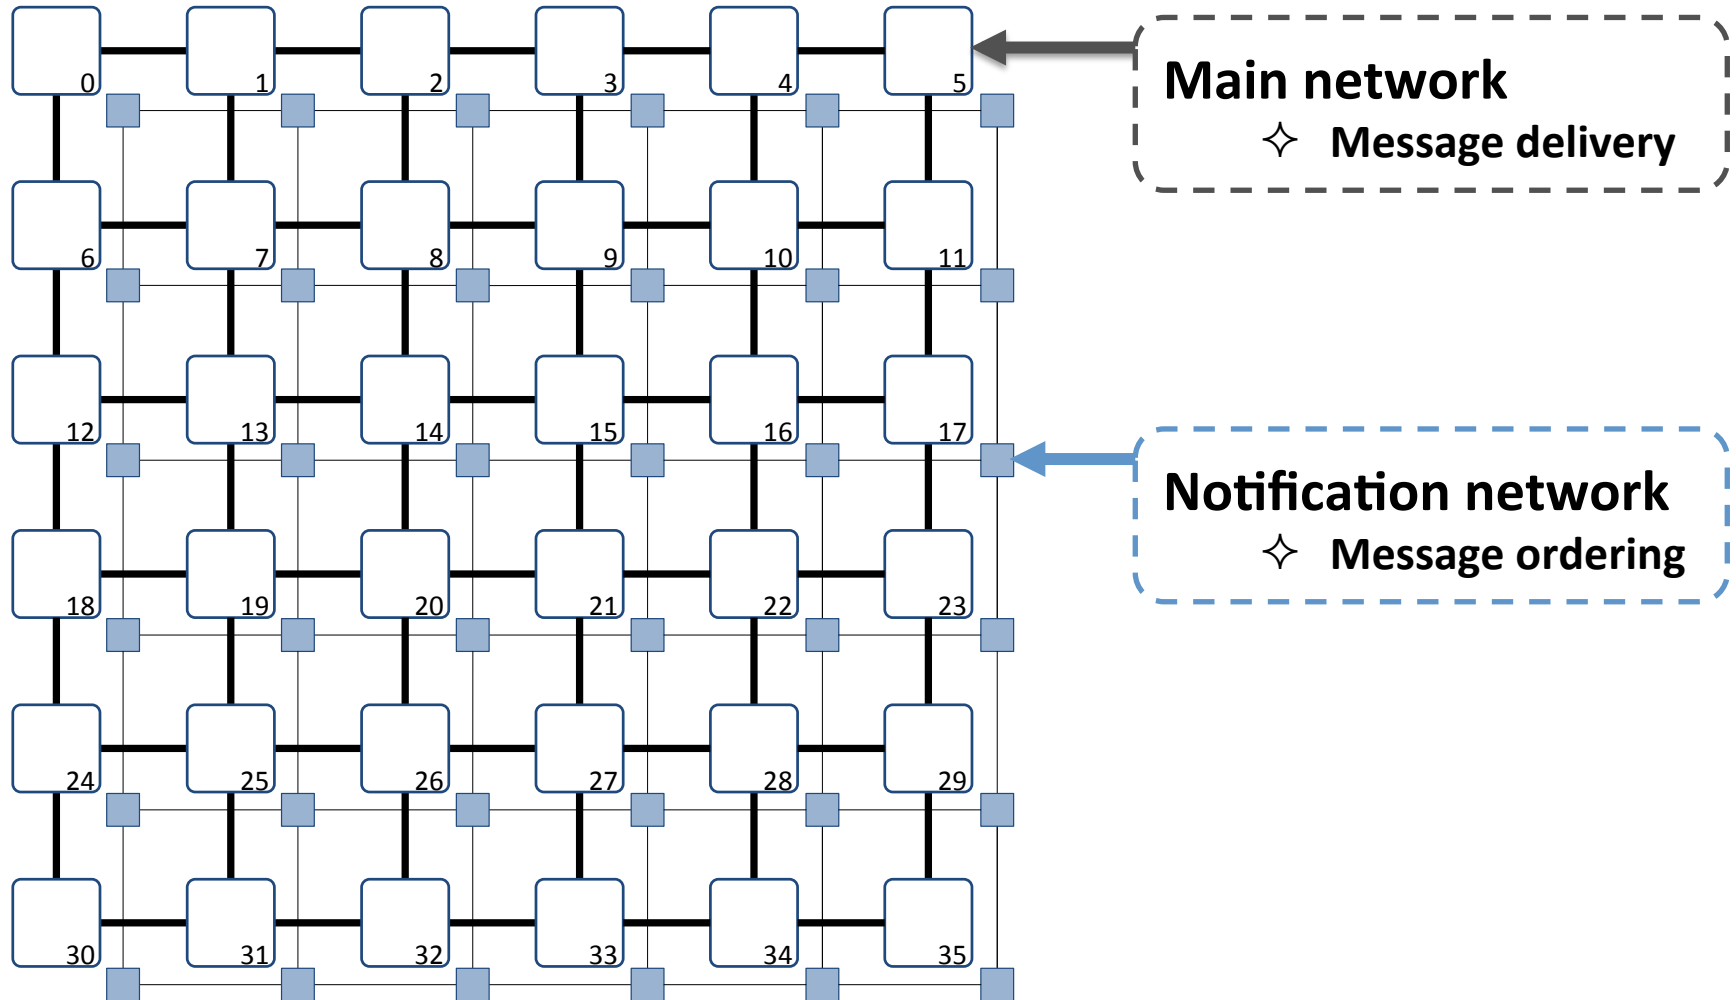


Globally Ordered Mesh Network






Decouple Ordering from Message Delivery





Notifications




Bounded latency (≤ 12 cycle)

- Non-blocking
- 1 cycle / hop broadcast mesh

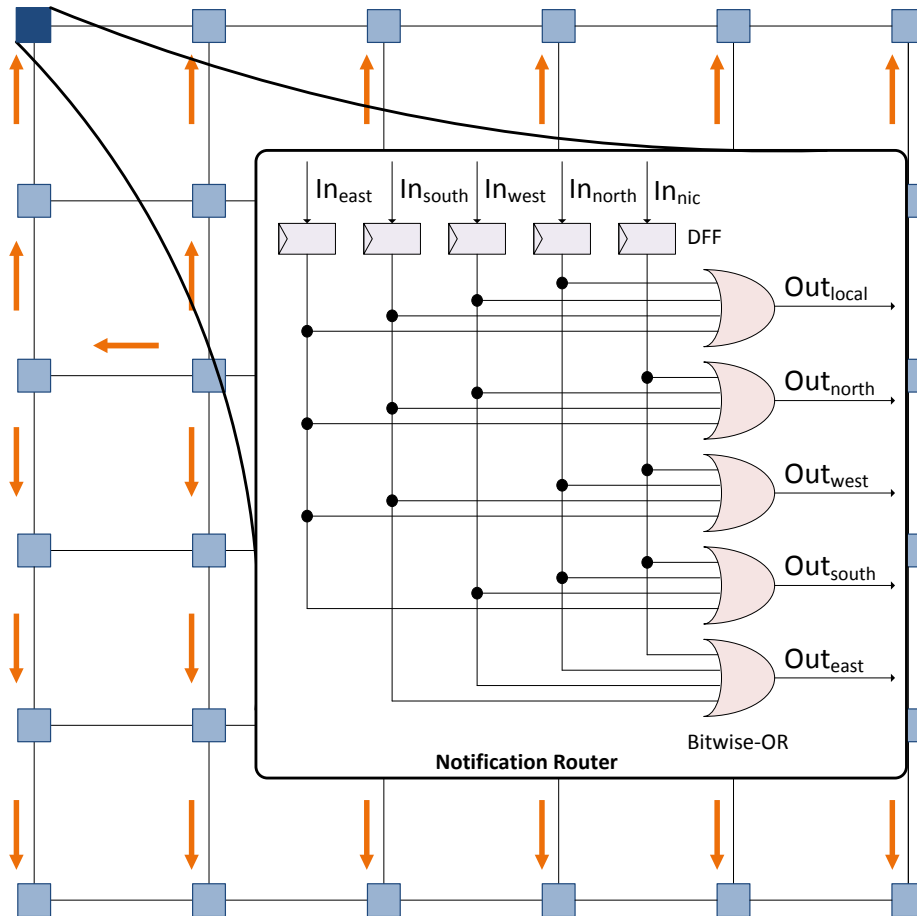
Notification bit-vector

- Dedicated 1 bit / tile
- Inject notification for each coherent request





Notifications




Bounded latency (≤ 12 cycle)

- Non-blocking
- 1 cycle / hop broadcast mesh

Notification bit-vector

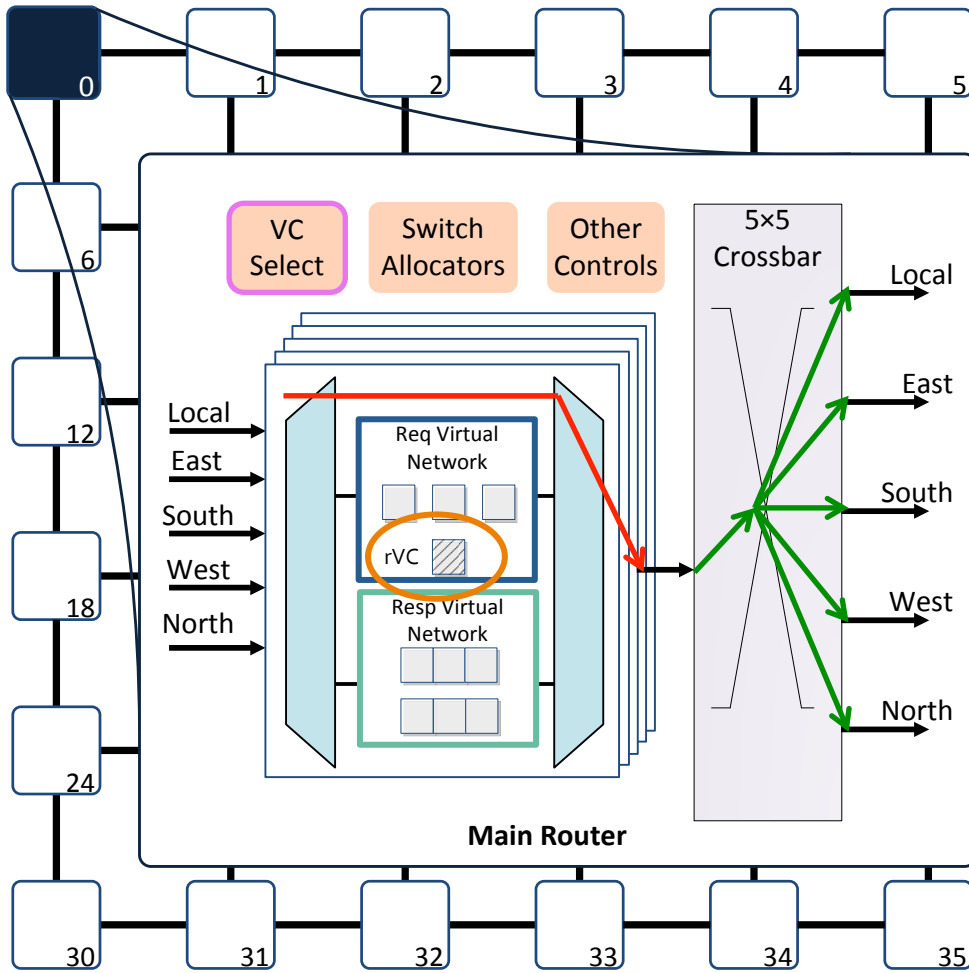
- Dedicated 1 bit / tile
- Inject notification for each coherent request



Low cost

- Only DFF + ORs

Main Network Router



Two Virtual Networks

- Globally Ordered Request (GO-REQ)
- Unordered Response (UO-RESP)

Optimizations

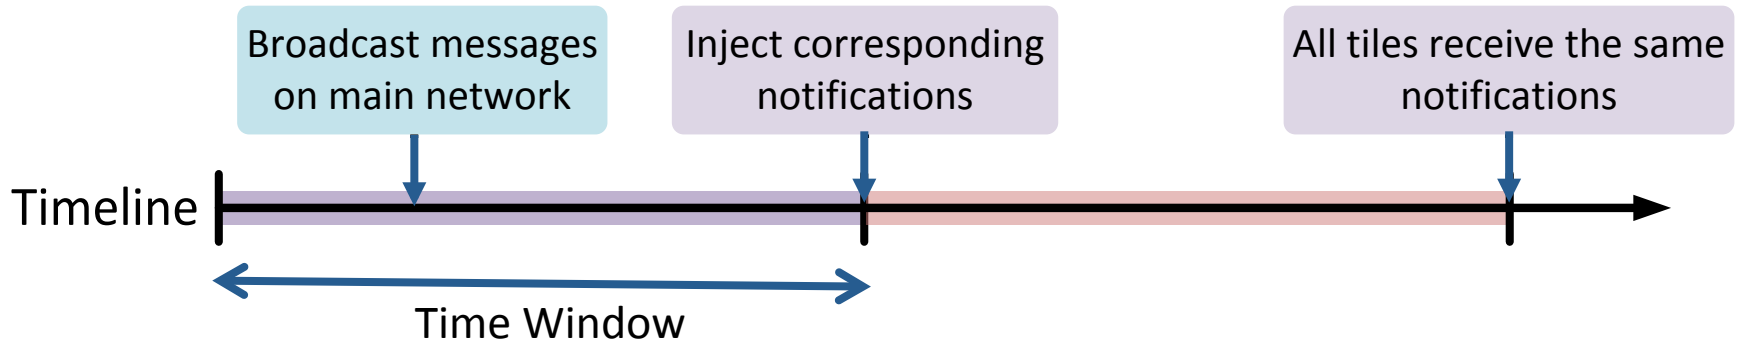
- Virtual Bypassing
- Broadcast Support
- VC Select Mechanism

Deadlock Avoidance

- XY Routing
- Separate virtual networks for requests and responses
- Reserved VC (rVC) for highest priority flit



Time Window



Length

At least equal to the fixed notification network latency bound

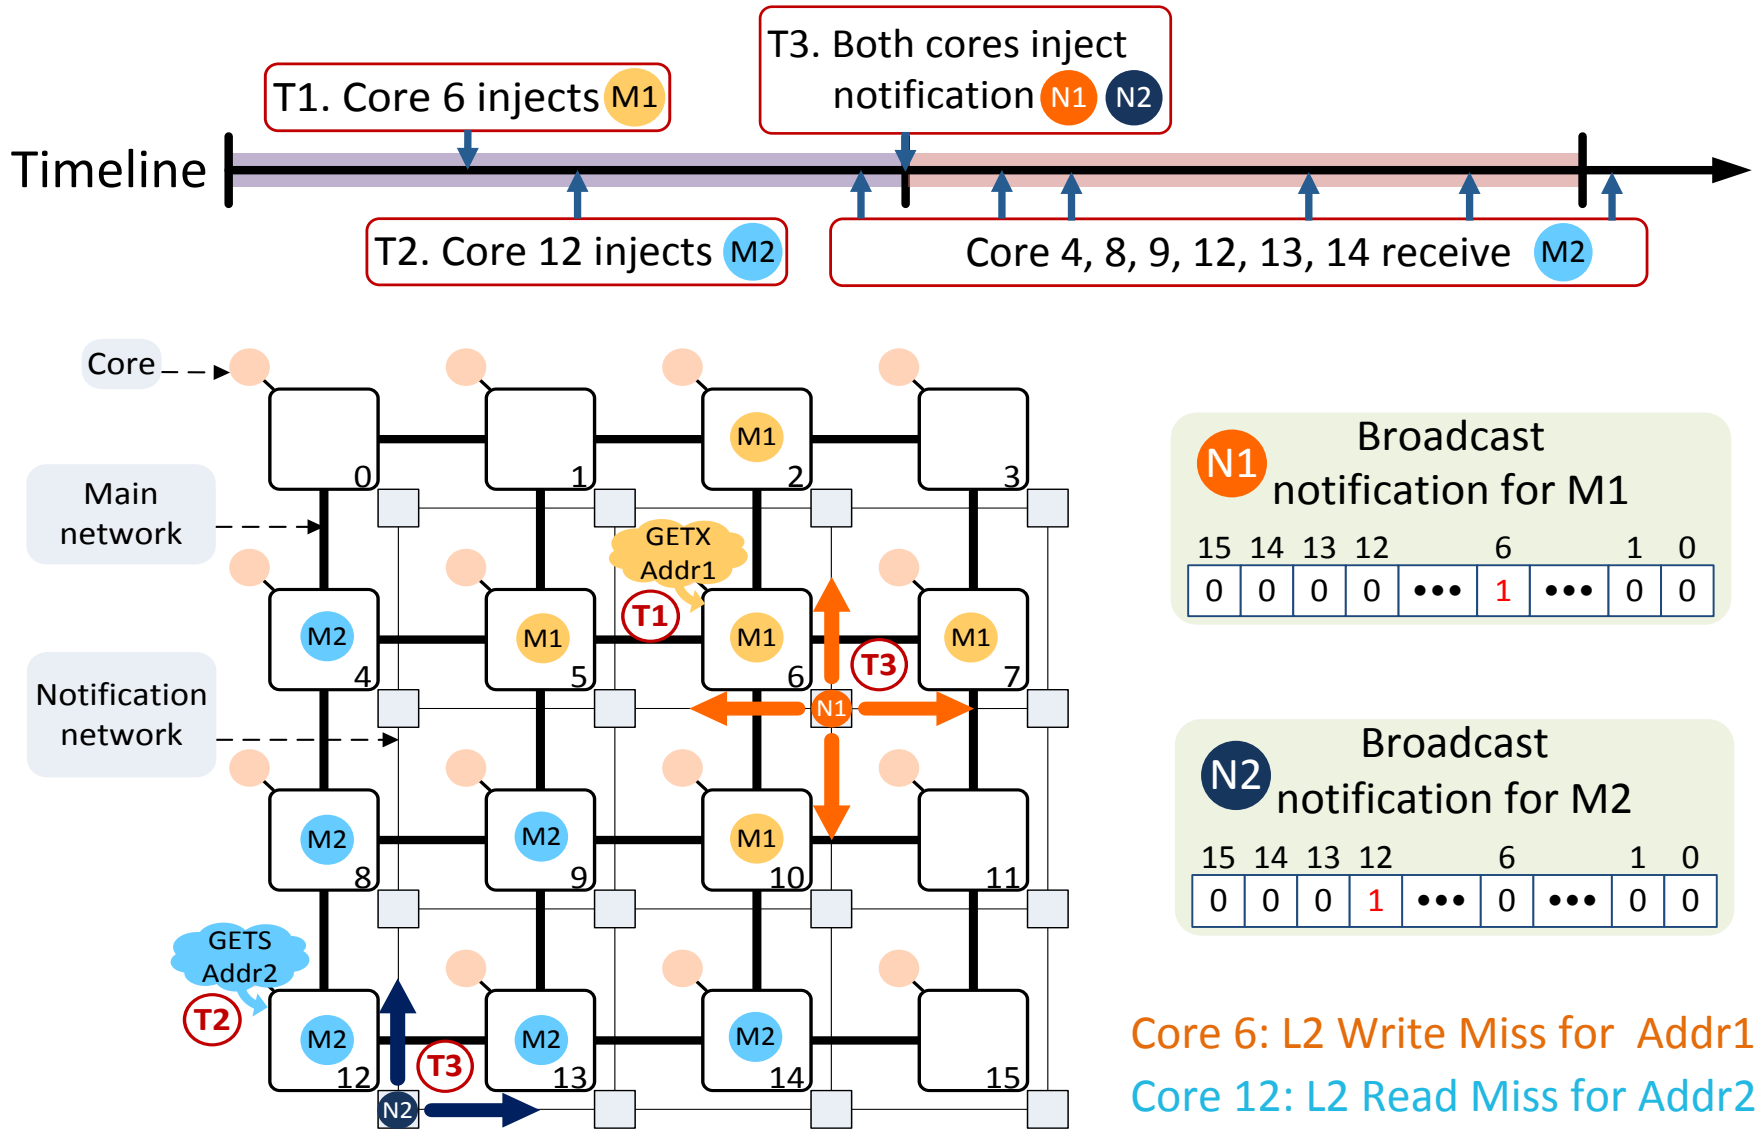
Fairness

Rotating global order priorities in each time window

Determine *locally* the *global order* according to consistent ordering rule at all nodes



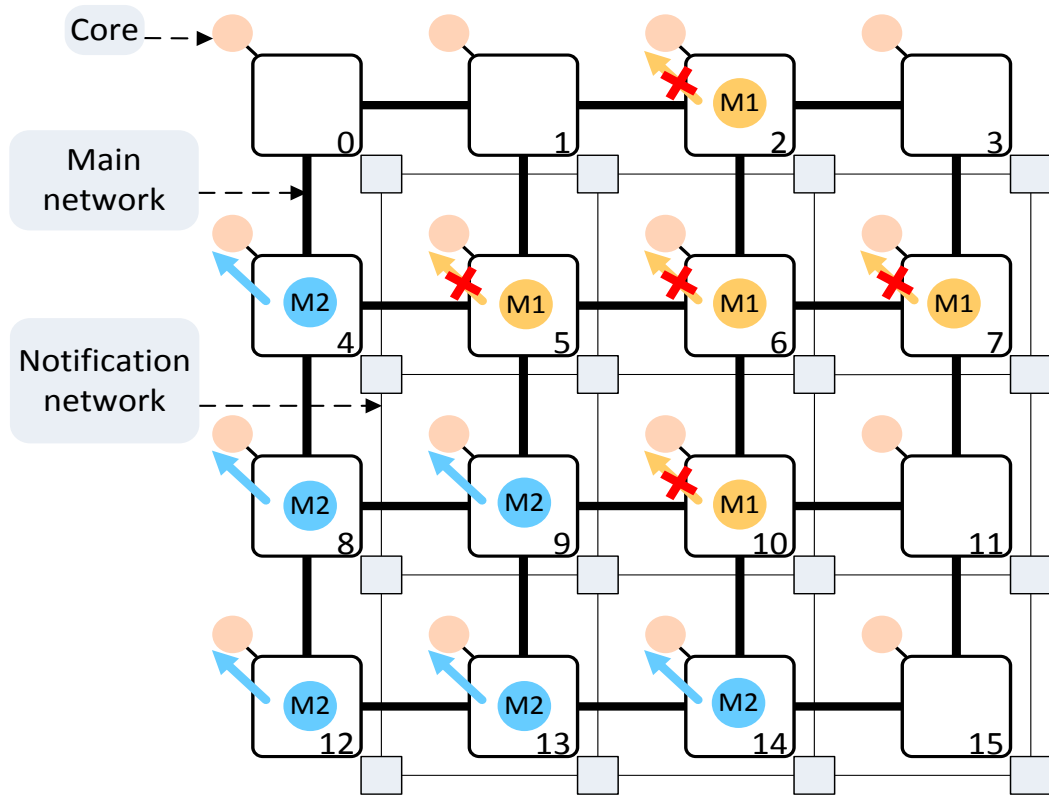
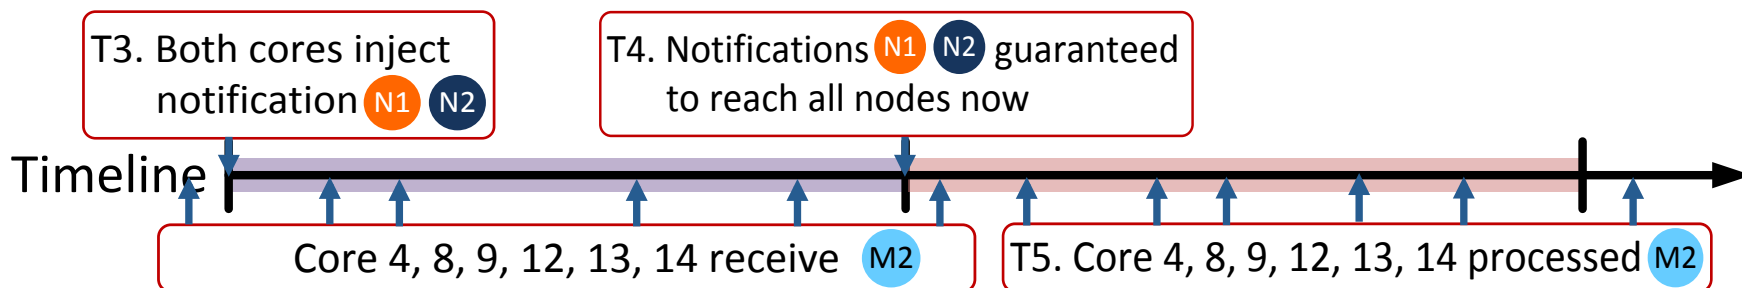


Journey of a Coherent Request



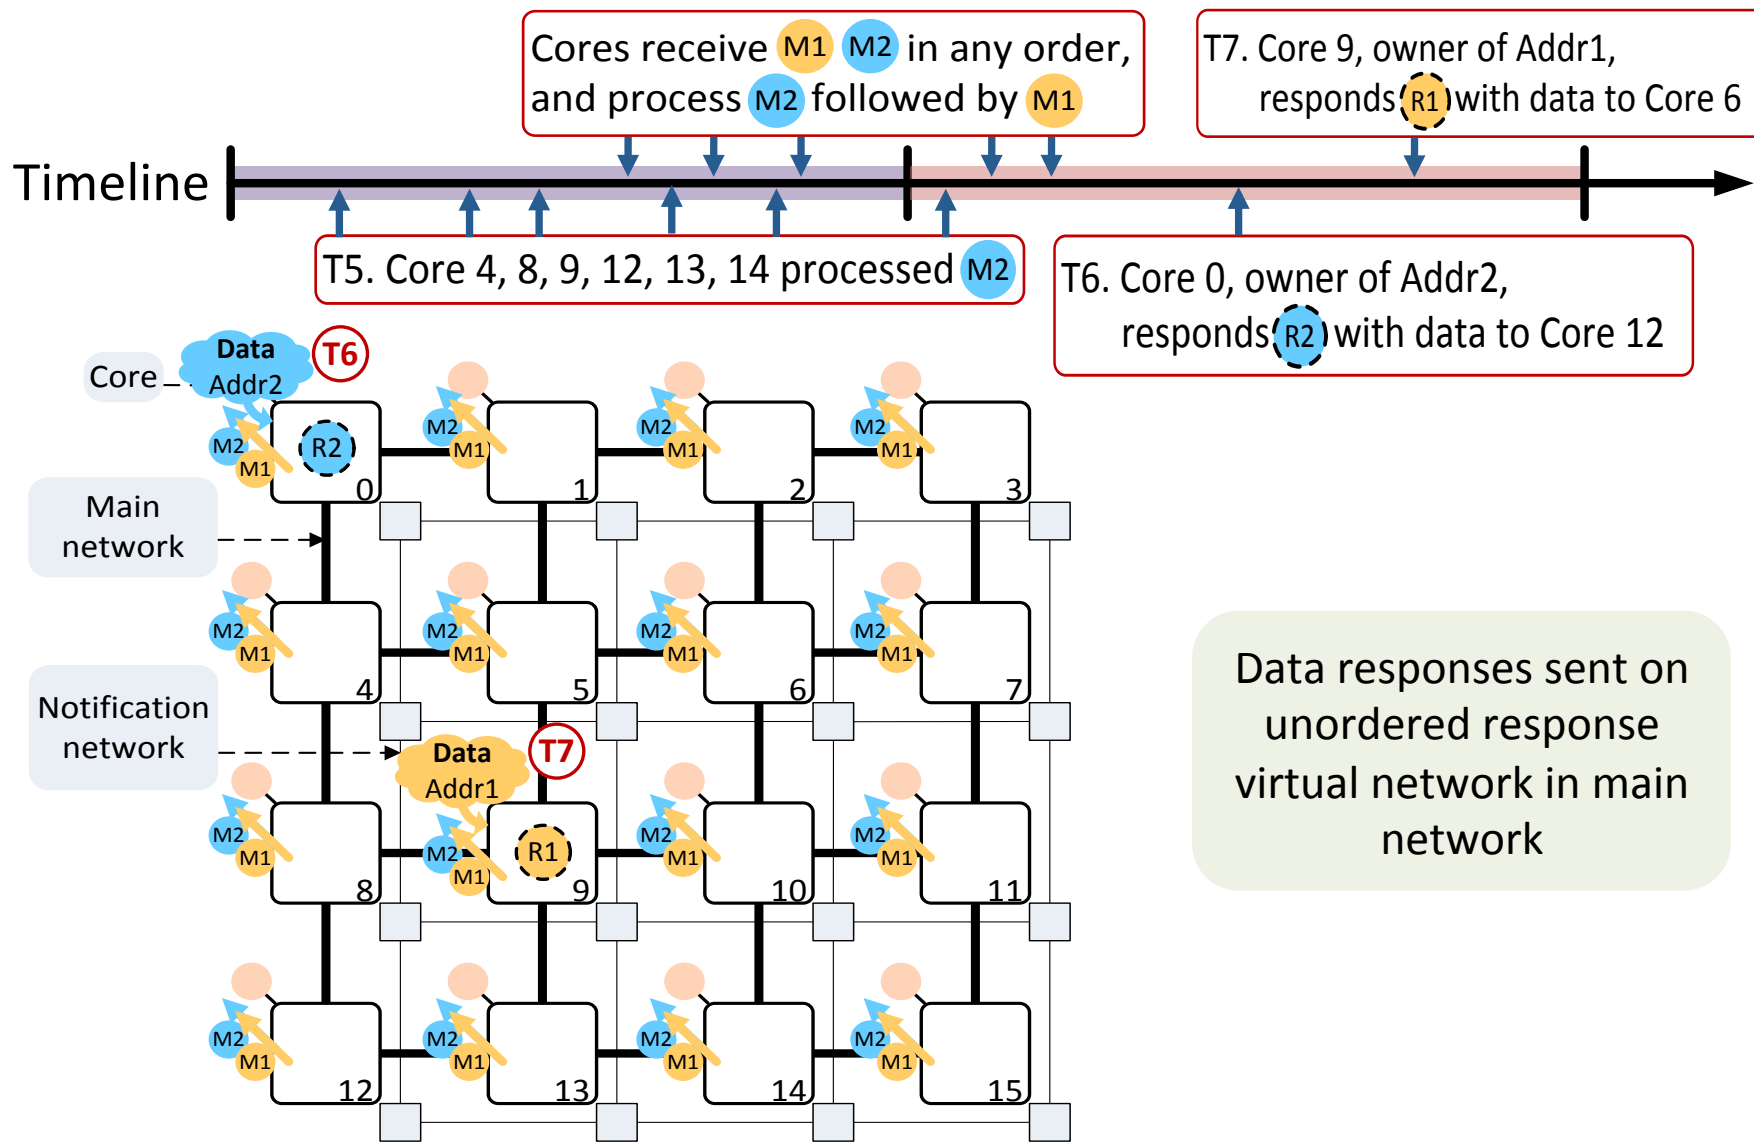


Journey of a Coherent Request





Journey of a Coherent Request





Architecture Analysis

Simulator	GEMS + GARNET
Access times	L1 – 1 cycle; L2 – 10 cycles; DRAM 90 cycles
LPD-D	Limited Pointer Distributed Directory
HT-D	AMD HyperTransport (Distributed)
SCORPIO-D	Snoopy Coherence: MOSI (Distributed)

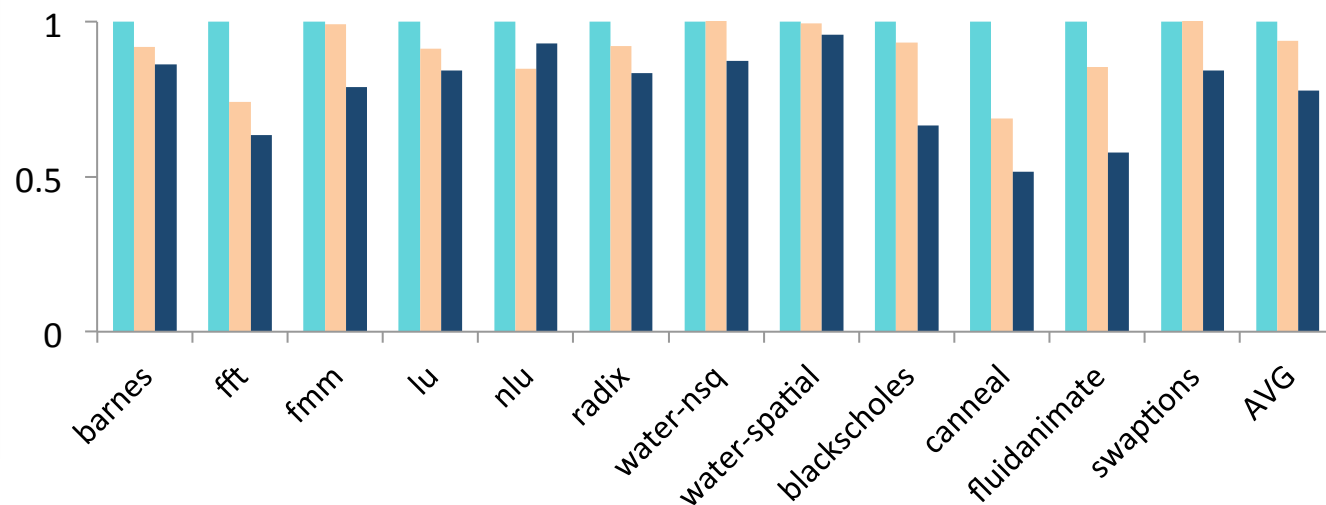
	LPD-D	HT-D	SCORPIO-D	
What is tracked?	Few sharers	Presence of owner	Presence of owner	Isolate
Ordering point	Directory	Directory	Network	Storage overhead
				Indirection latency



Normalized Runtime


3
6
C
O
R
E
S

LPD-D HT-D SCORPIO-D



- 22% lower than LPD-D
- 17% lower than HT-D

6
4
C
O
R
E
S

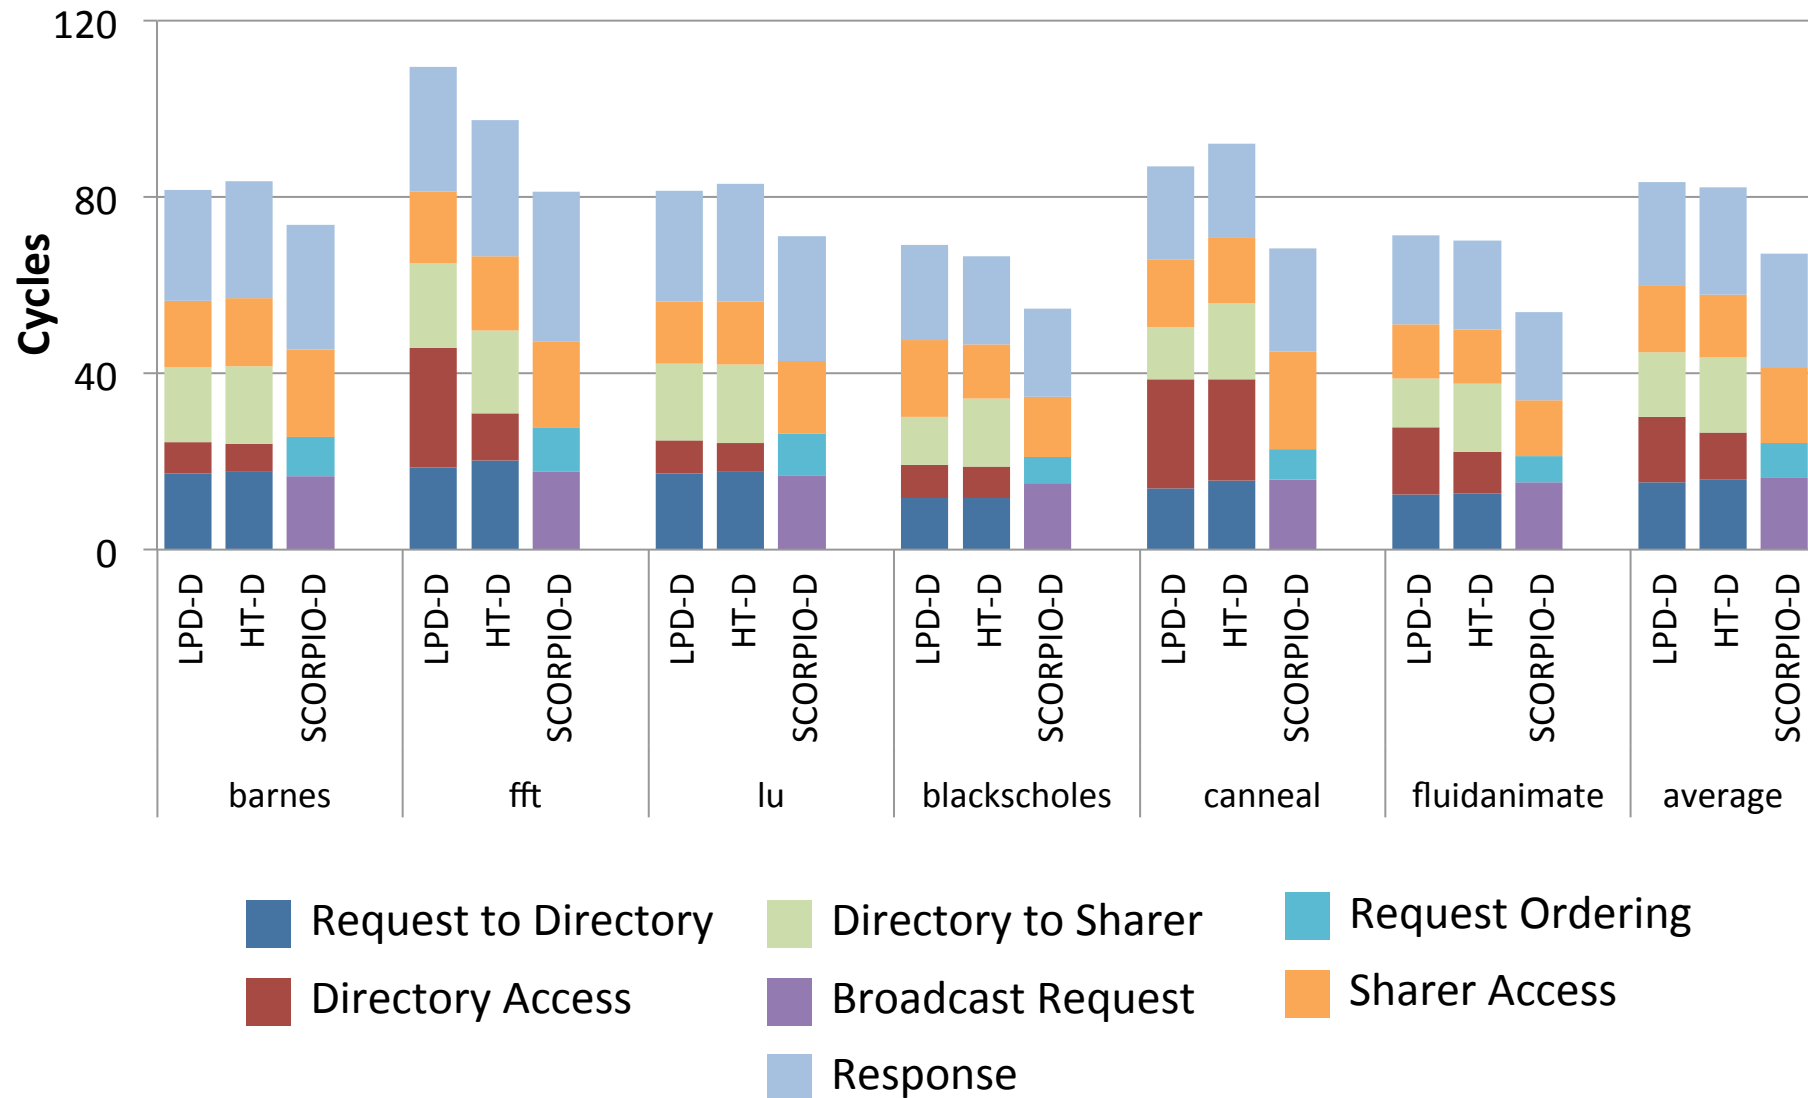


- 24% lower than LPD-D
- 13% lower than HT-D



L2 Service Latency Breakdown

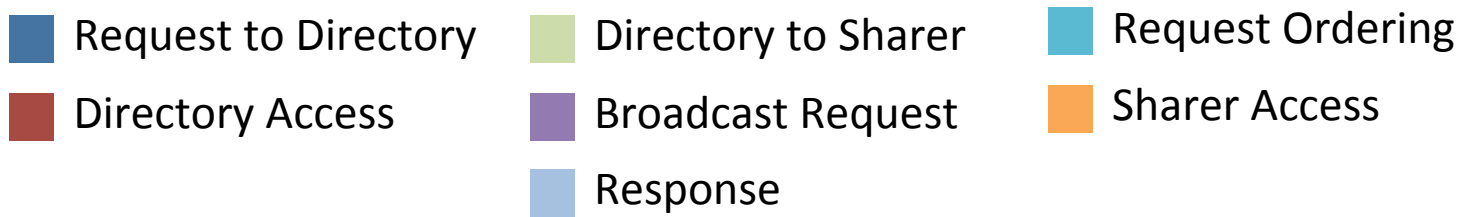
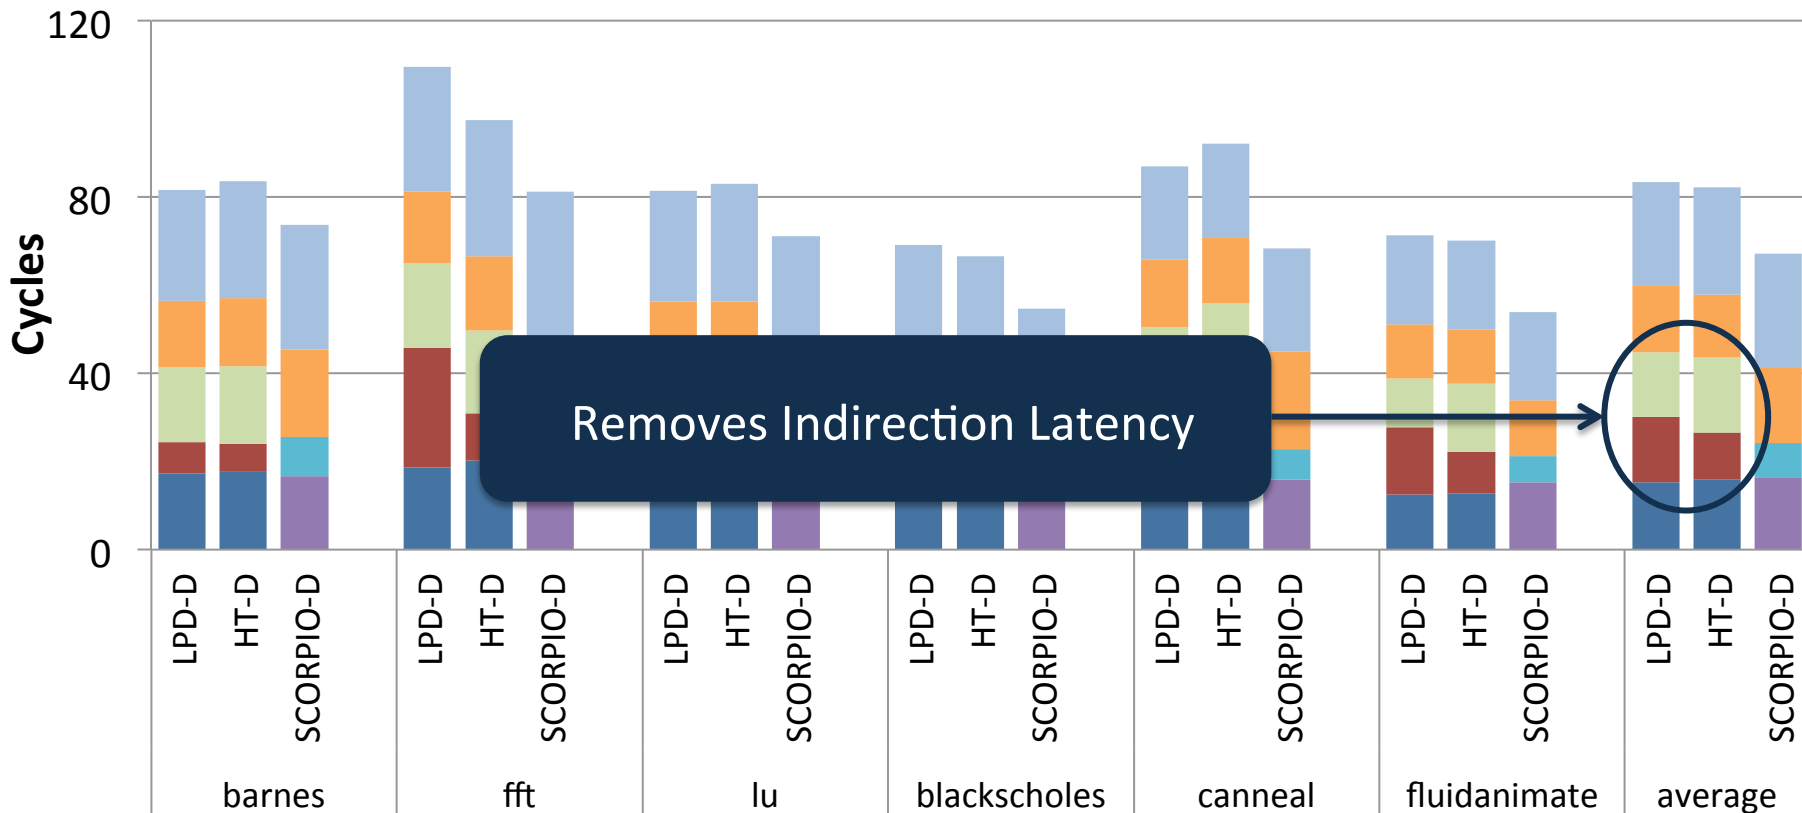
Served by Other Caches On-Chip





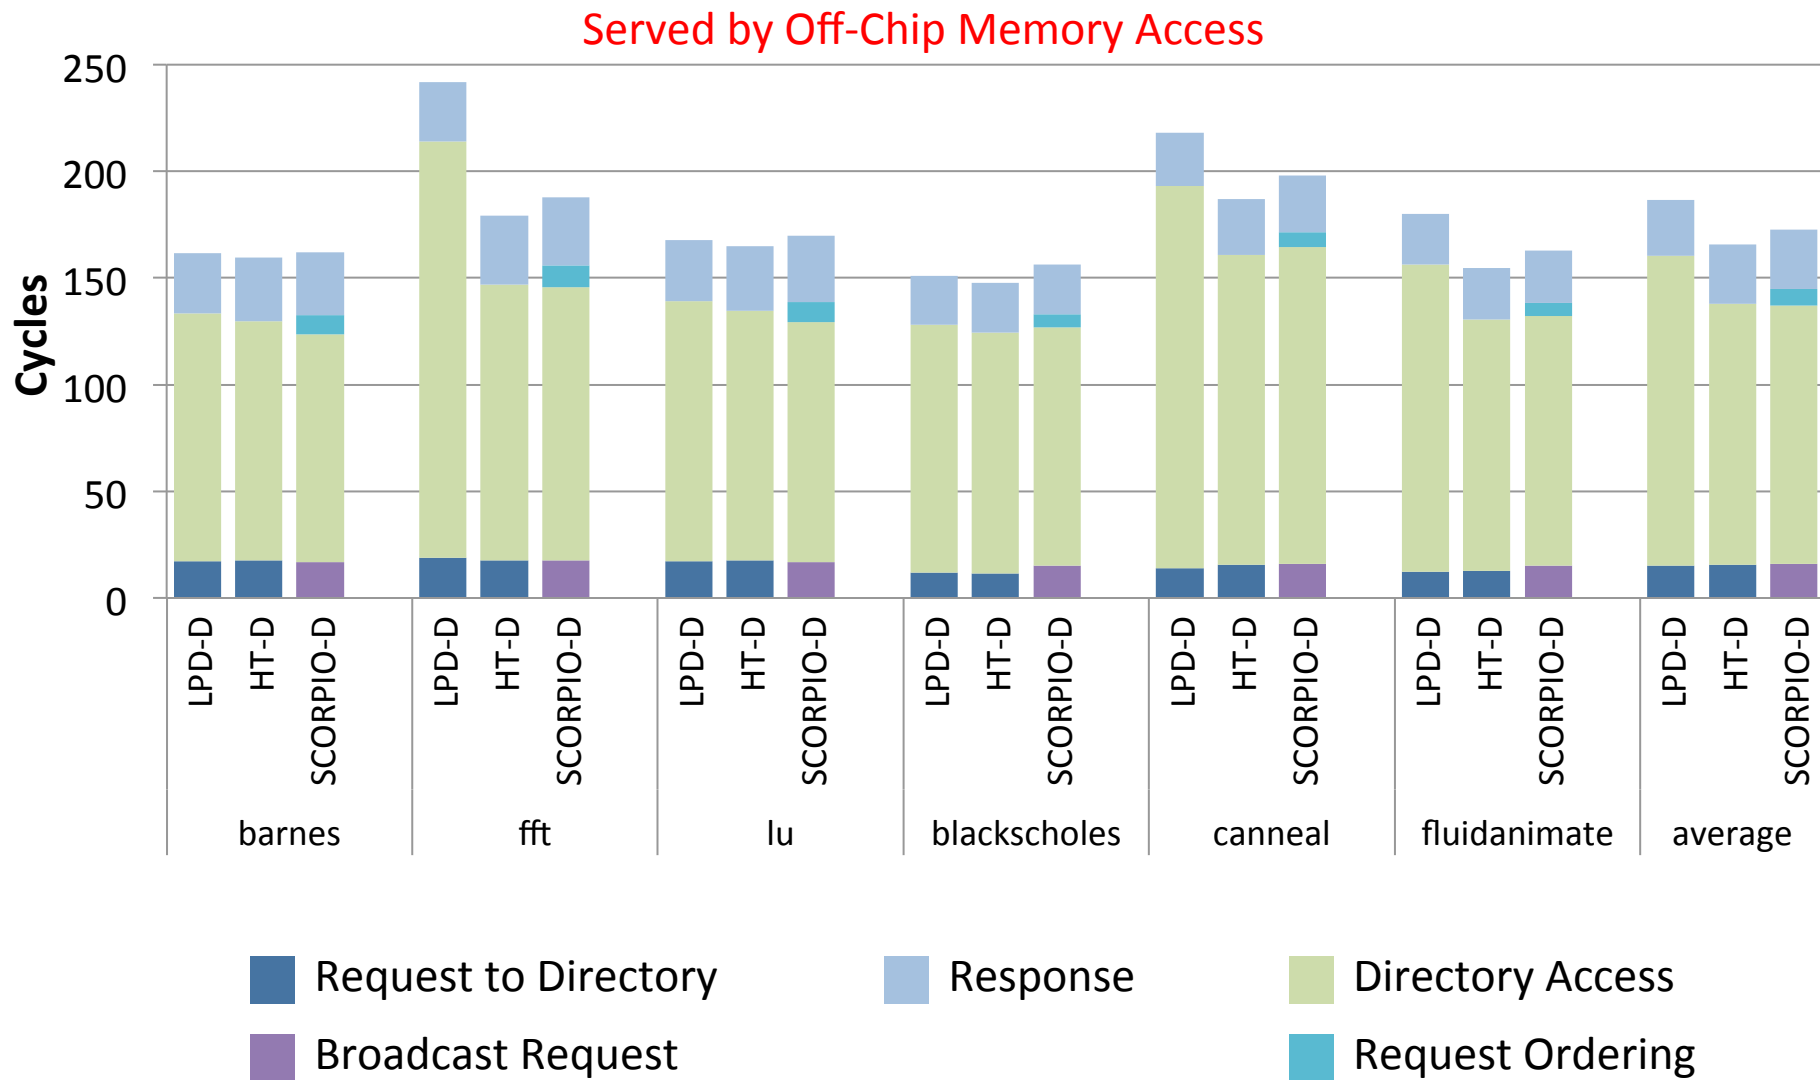
L2 Service Latency Breakdown

Served by Other Caches On-Chip



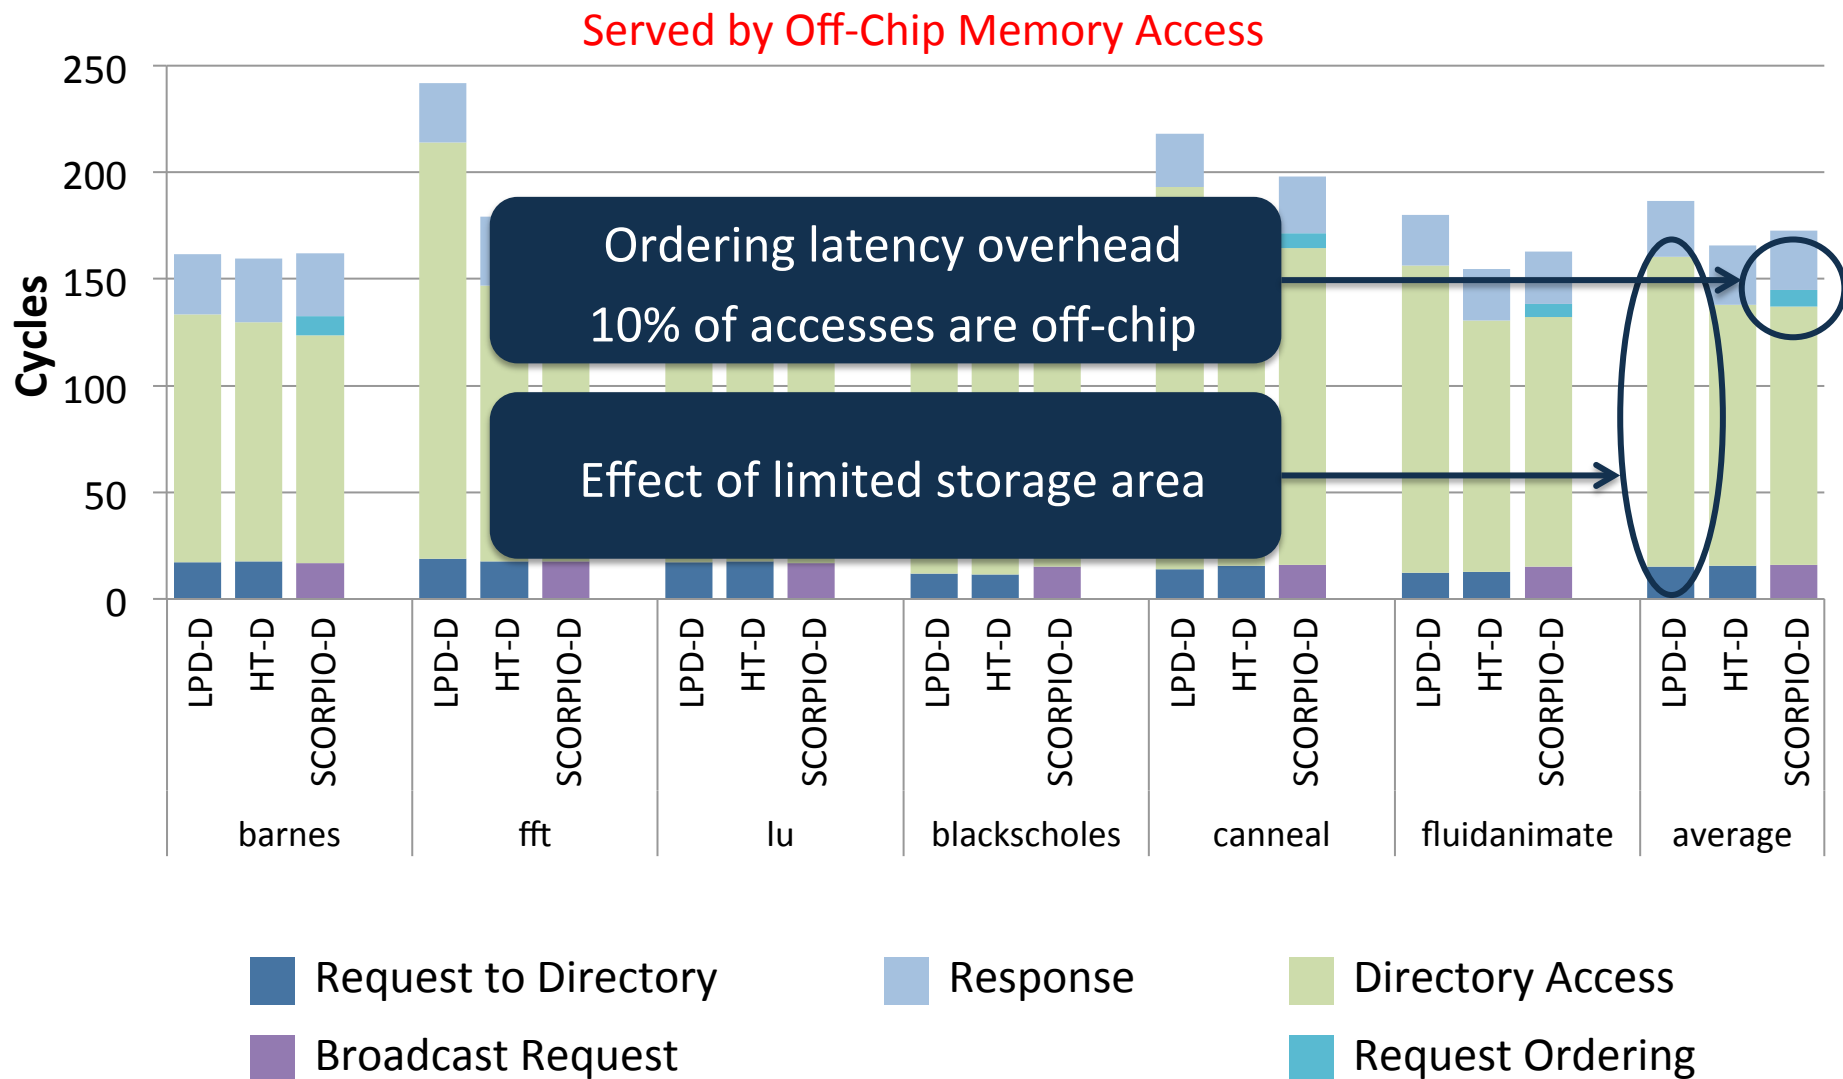


L2 Service Latency Breakdown





L2 Service Latency Breakdown



Practical Network Design



Token Coherence (TokenB)

- T tokens for each block of shared memory
- T = # of processors

$$\text{Area Overhead} \propto \text{Number Cachelines}$$


In-Network Snoop Ordering (INSO)

- Periodic expiration required
- Worsening ordering latency and network traffic with small expiration window

Timestamp Snooping (TS)

- Each request tagged with an ordering time (OT)

$$\begin{aligned} \text{Buffer Area} &\propto \text{\# of Cores} \\ \text{Per Core} &\propto \text{Msgs/Core} \end{aligned}$$





Practical Network Design

Token Coherence (TokenB)

- T tokens for each block of shared data
- $T = \# \text{Cores} \times \text{Cache Size}$

Area Overhead \propto Number of Cachelines

In-Network Snoop Ordering (INSO)

- Periodic expiration

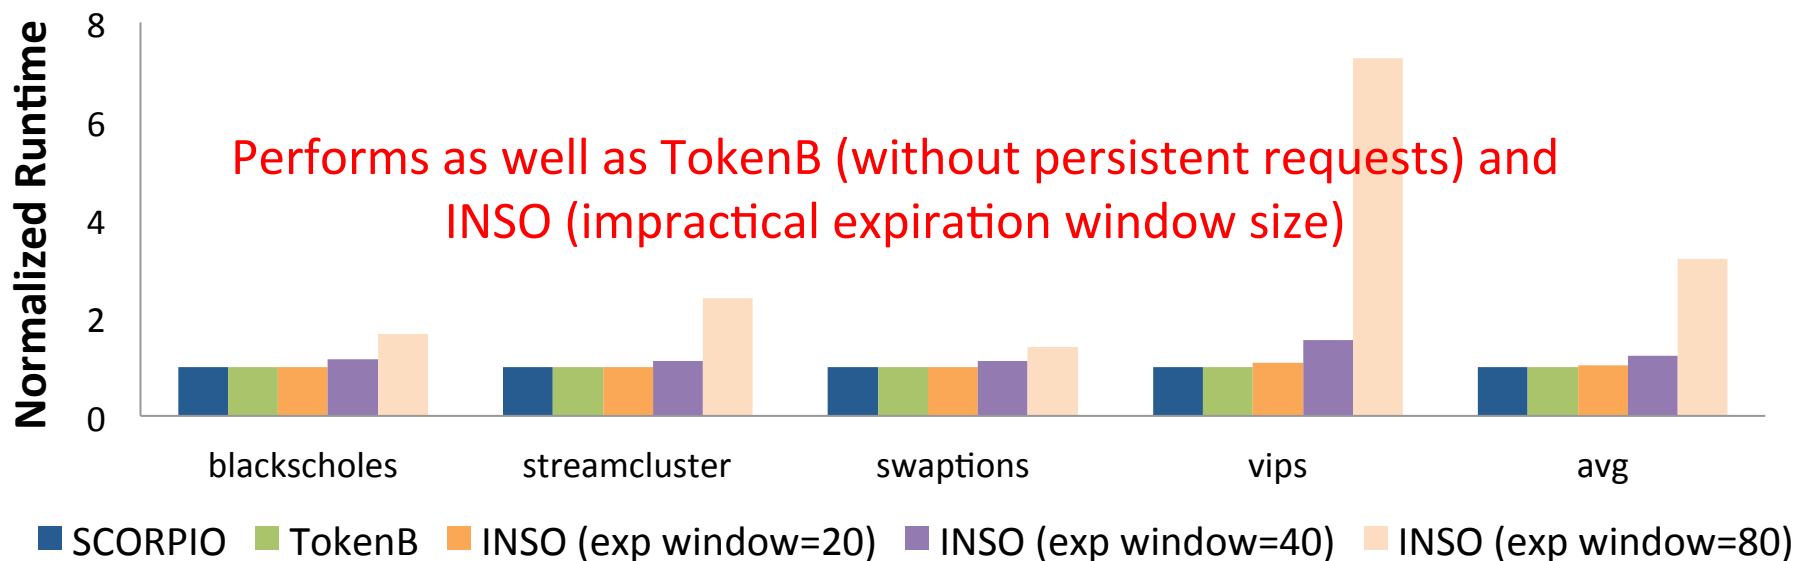
2 buffers at each endpoint
 traffic with small expiration window

Timestamp Snooping (TS)

- Each request tagged with time

Buffer Area Per Core \propto # of Cores \times Msgs/Core


**SCORPIO Minimum Area Requirement for Correctness:
 2 buffers at each endpoint**





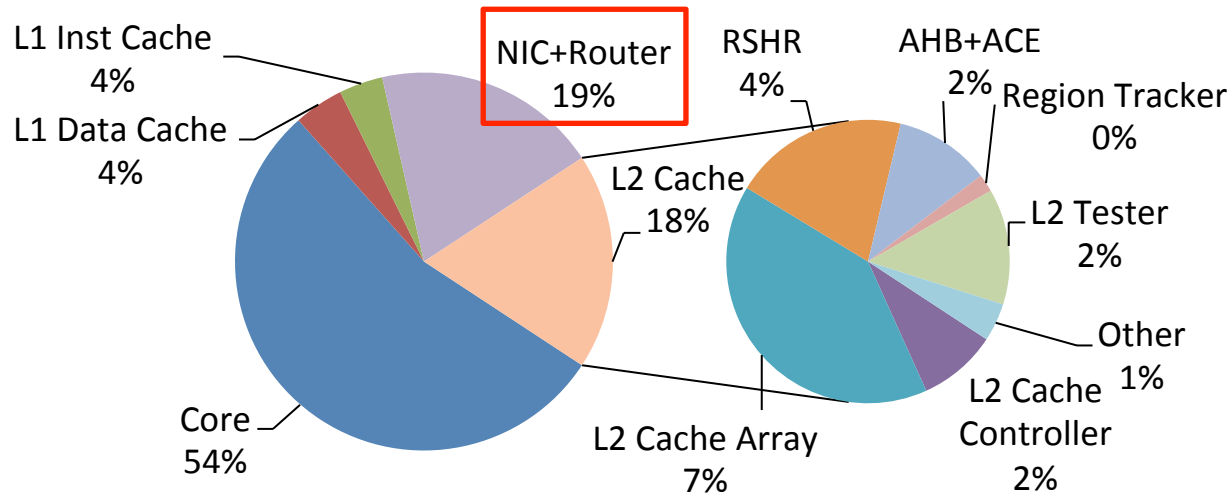
Network Cost

Post-synthesis frequency: 1 GHz
 Post-layout frequency: 833 MHz



Network occupies only 10% of the area


Area




Network consumes 19% of the power

Power

Conclusion



SCORPIO Die



SCORPIO Packaged & Assembled



Mounting SCORPIO Chip to Custom Designed PCB Board



Intend to Run Linux OS and Applications on the Chip



Releasing SCORPIO NoC RTL after Measurements Completed