

Run-time Strategies for Energy-efficient Operation of Silicon-Photonic NoCs

Ajay Joshi

joshi@bu.edu

Integrated Circuits and Systems Group (ICSG)

Department of ECE, Boston University, Boston MA

Acknowledgement

□ Graduate Students/Postdocs


- Chao Chen, Jose Abellan, Tiansheng Zhang

□ Faculty

- Ayse Coskun, Jonathan Klamkin, Milos Popovic


- This research has been funded by the NSF grant CCF-1149549

Historic Trends in VLSI Systems




Data collected by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, C. Batten


Silicon-Photonic NoC Research




[Kirman'06]




[Shacham'08]




[Vantrease'08]




[Gu'09]




[Joshi'09]




[Batten'09]




[Pan'09]




[Cianchetti'09]




[Koka'10]




[Pan'10]




[Beamer'10]




[Xue'10]




[Hendary'11]




[Kao'11]



[Li'13]



[Zulfiqar'13]



[Chen'14]

Silicon-Photonic NoC Challenges

□ Bandwidth utilization

- Current applications/architectures do not need Tbps on-chip bandwidth

□ Packaging/Integration

- Interface electrical and photonic devices
- Coupling tens of off-chip laser sources to photonic NoC is challenging

□ Power consumption

- Large laser power and thermal tuning power could negate bandwidth benefits

Silicon-Photonic NoC Challenges

□ Bandwidth utilization

- Current applications/architectures do not need Tbps on-chip bandwidth

□ Packaging/Integration

- Interface electrical and photonic devices
- Coupling tens of off-chip laser sources to photonic NoC is challenging

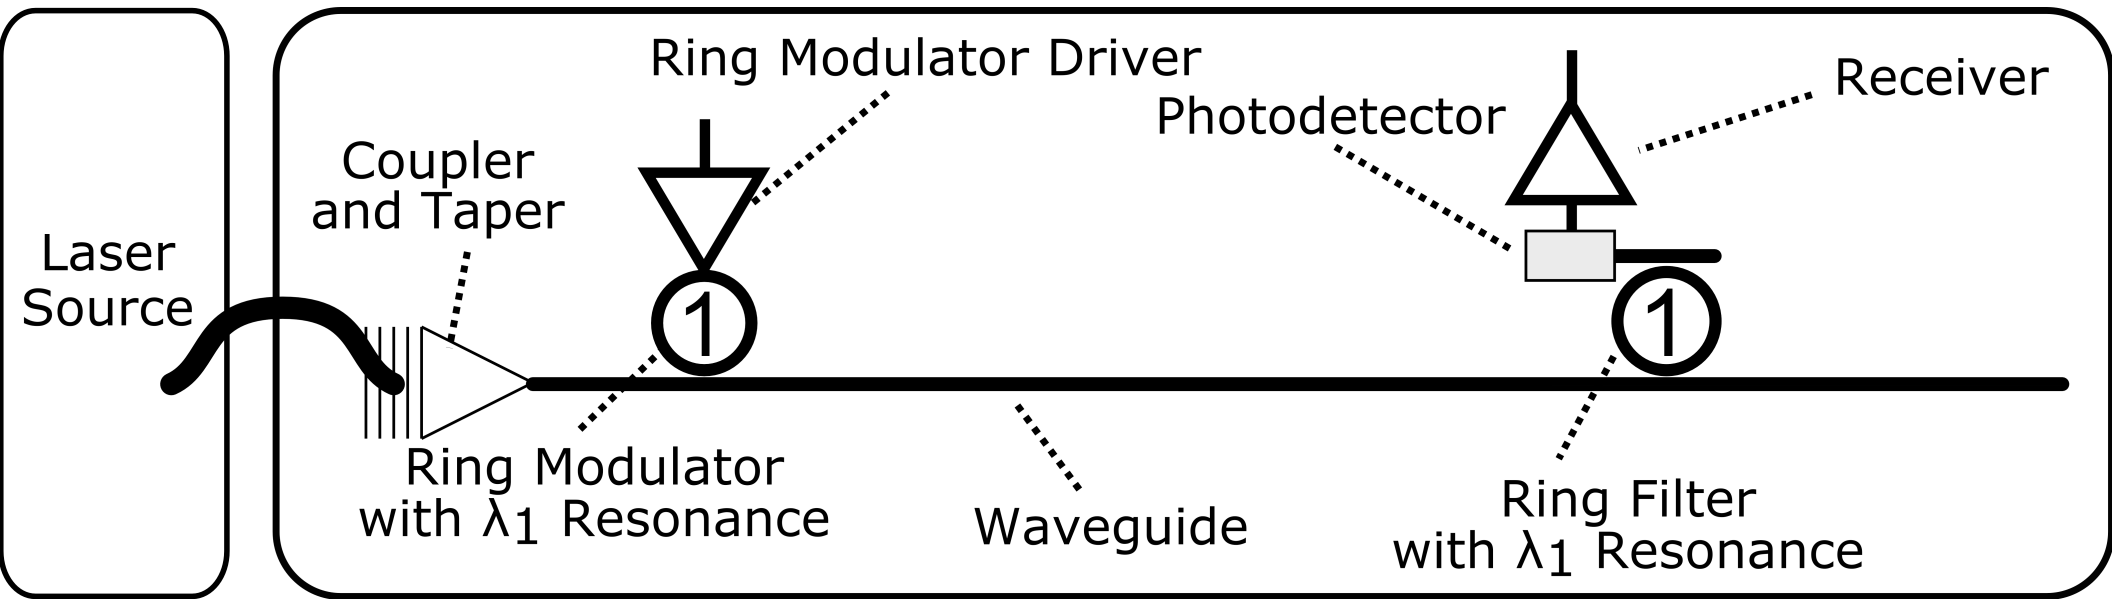
□ Power consumption

- Large laser power and thermal tuning power could negate bandwidth benefits

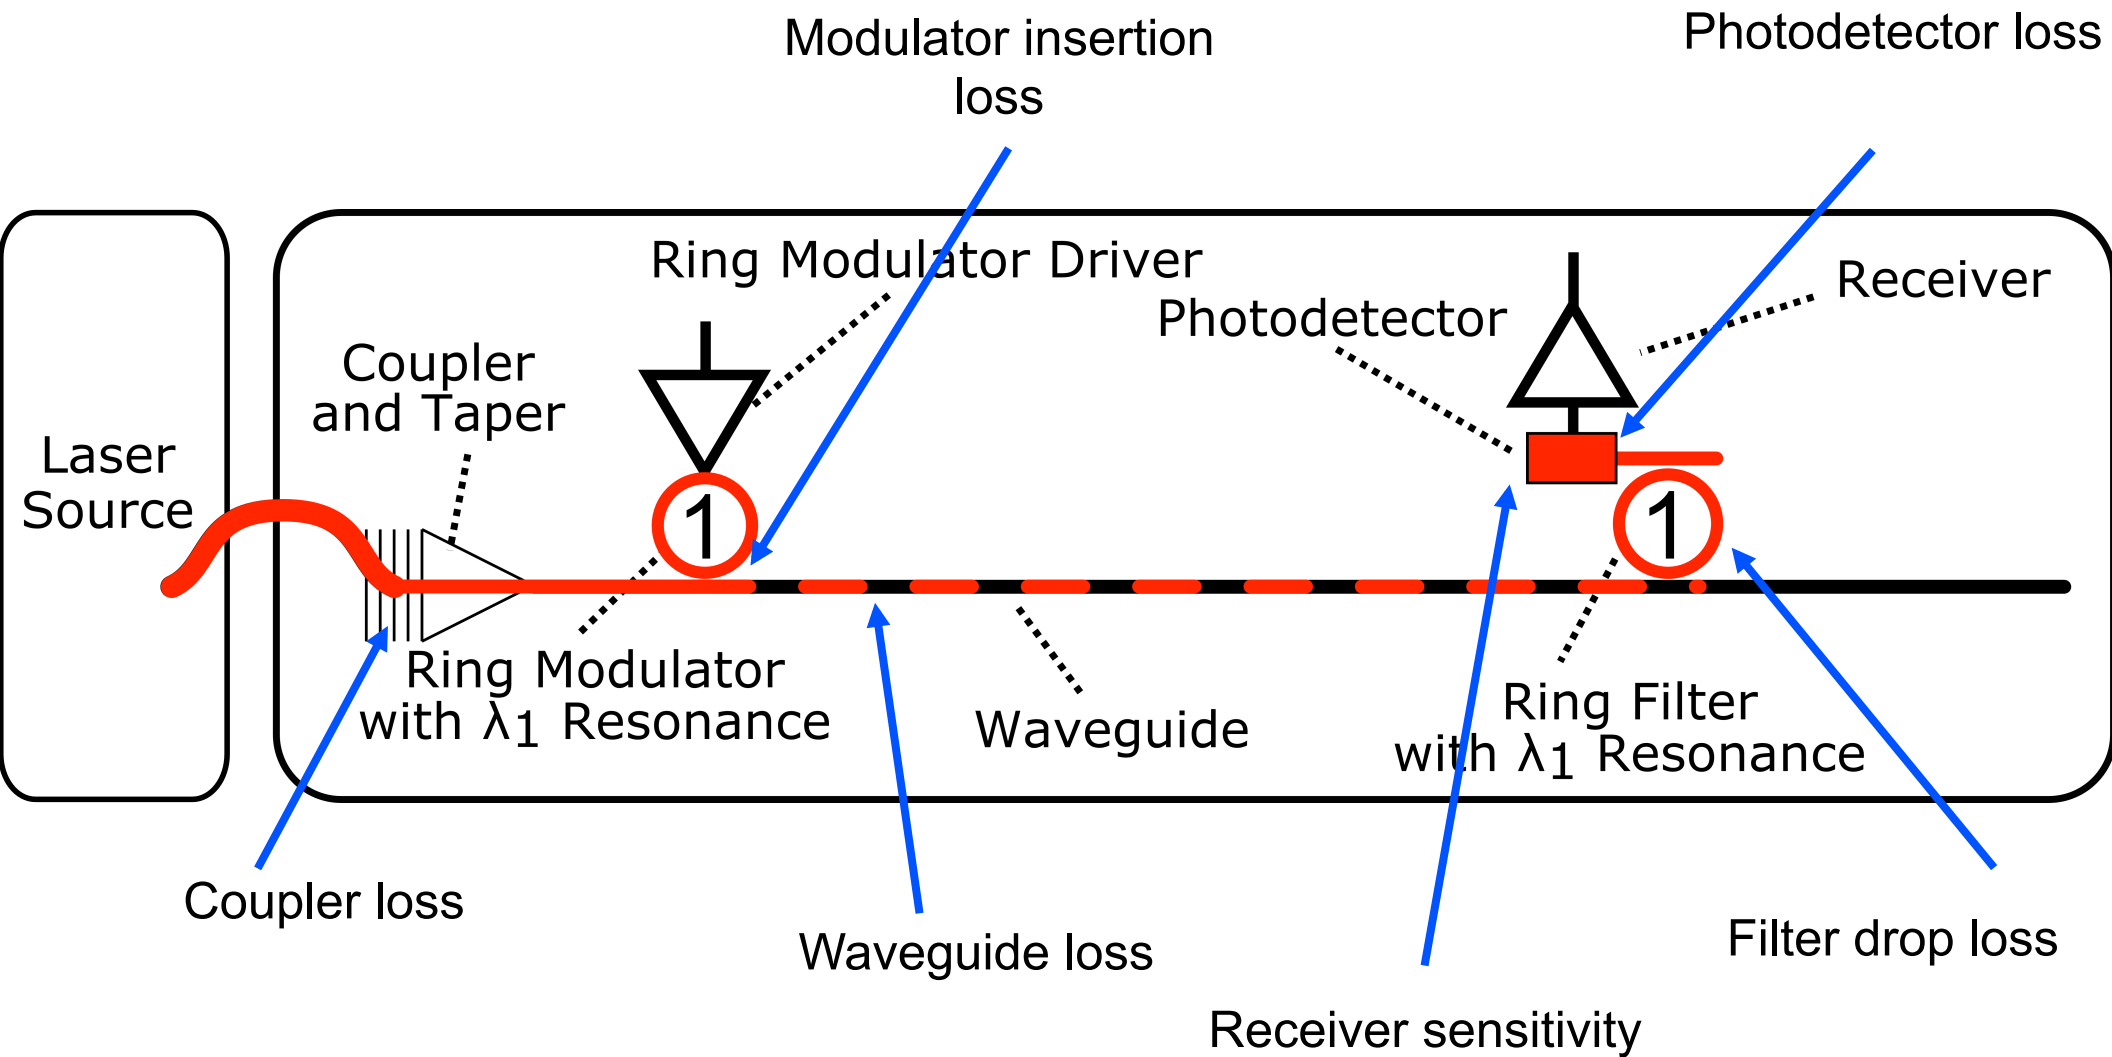
Outline

- Background
- Laser Power Management using NoC and Cache Reconfiguration
- Thermal management using job allocation
- Summary


Silicon photonic link



Silicon photonic link



Silicon photonic link



Laser source could consume 10-15 watts of power

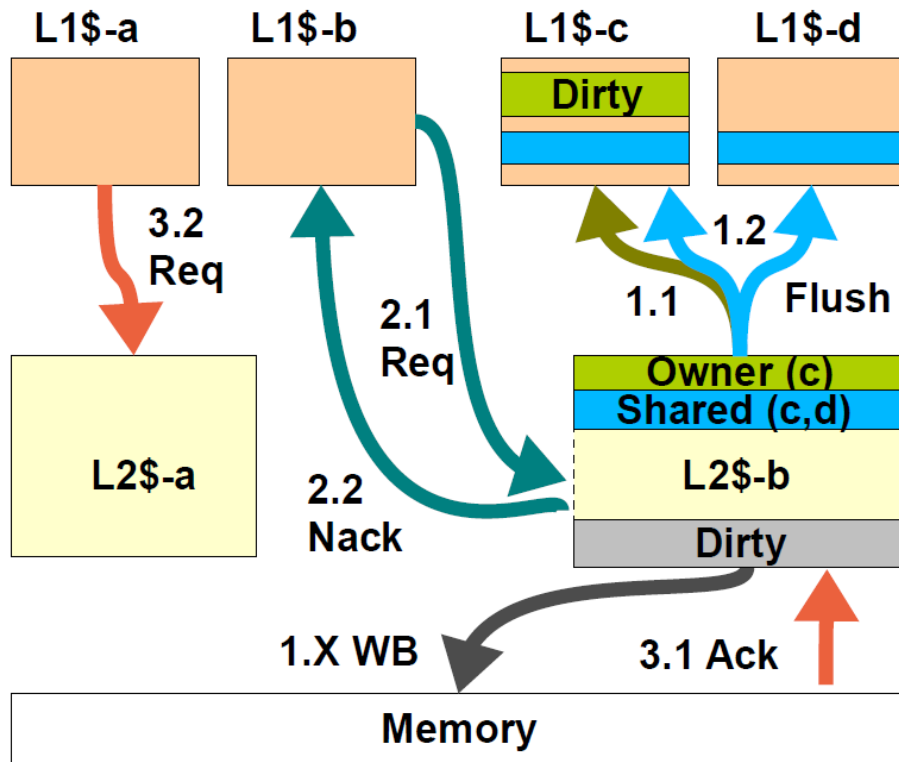
Related Work

- Low loss photonic devices [plenty of efforts in place]
- Channel sharing [Pan 2010] [Li 2013]
- NoC bandwidth scaling [Zhou 2013] [Chen 2013] [Demir 2014]
- Sharing/Placement of laser sources [Chen 2014]

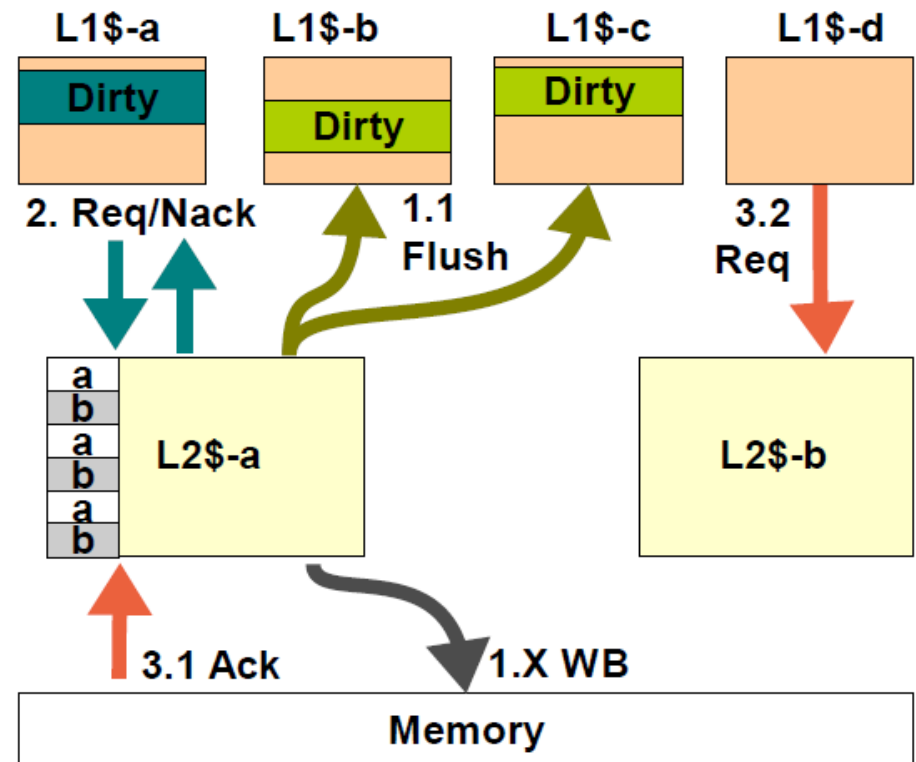
We use a combination of NoC and Cache reconfiguration to save laser power

Cache Reconfiguration Process

L2 cache bank deactivation




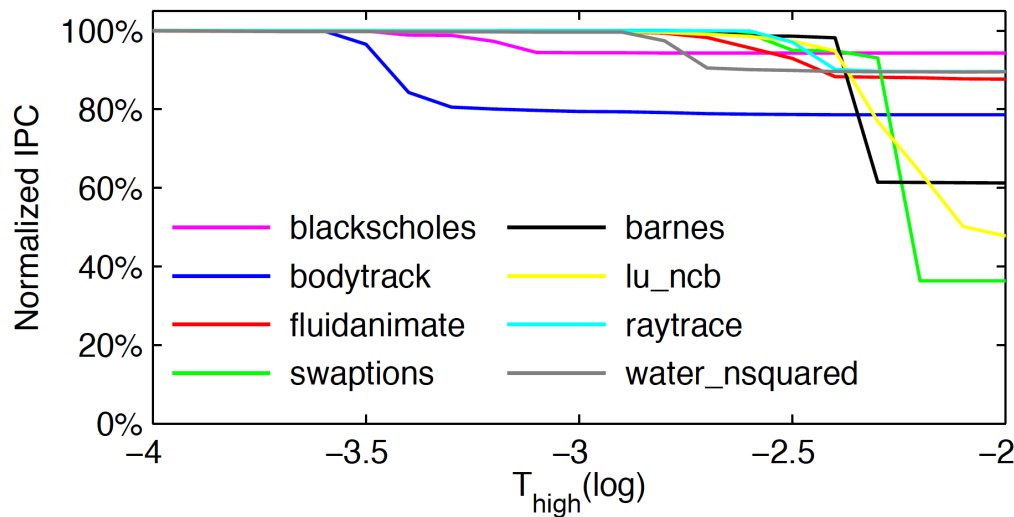
L2 cache bank activation



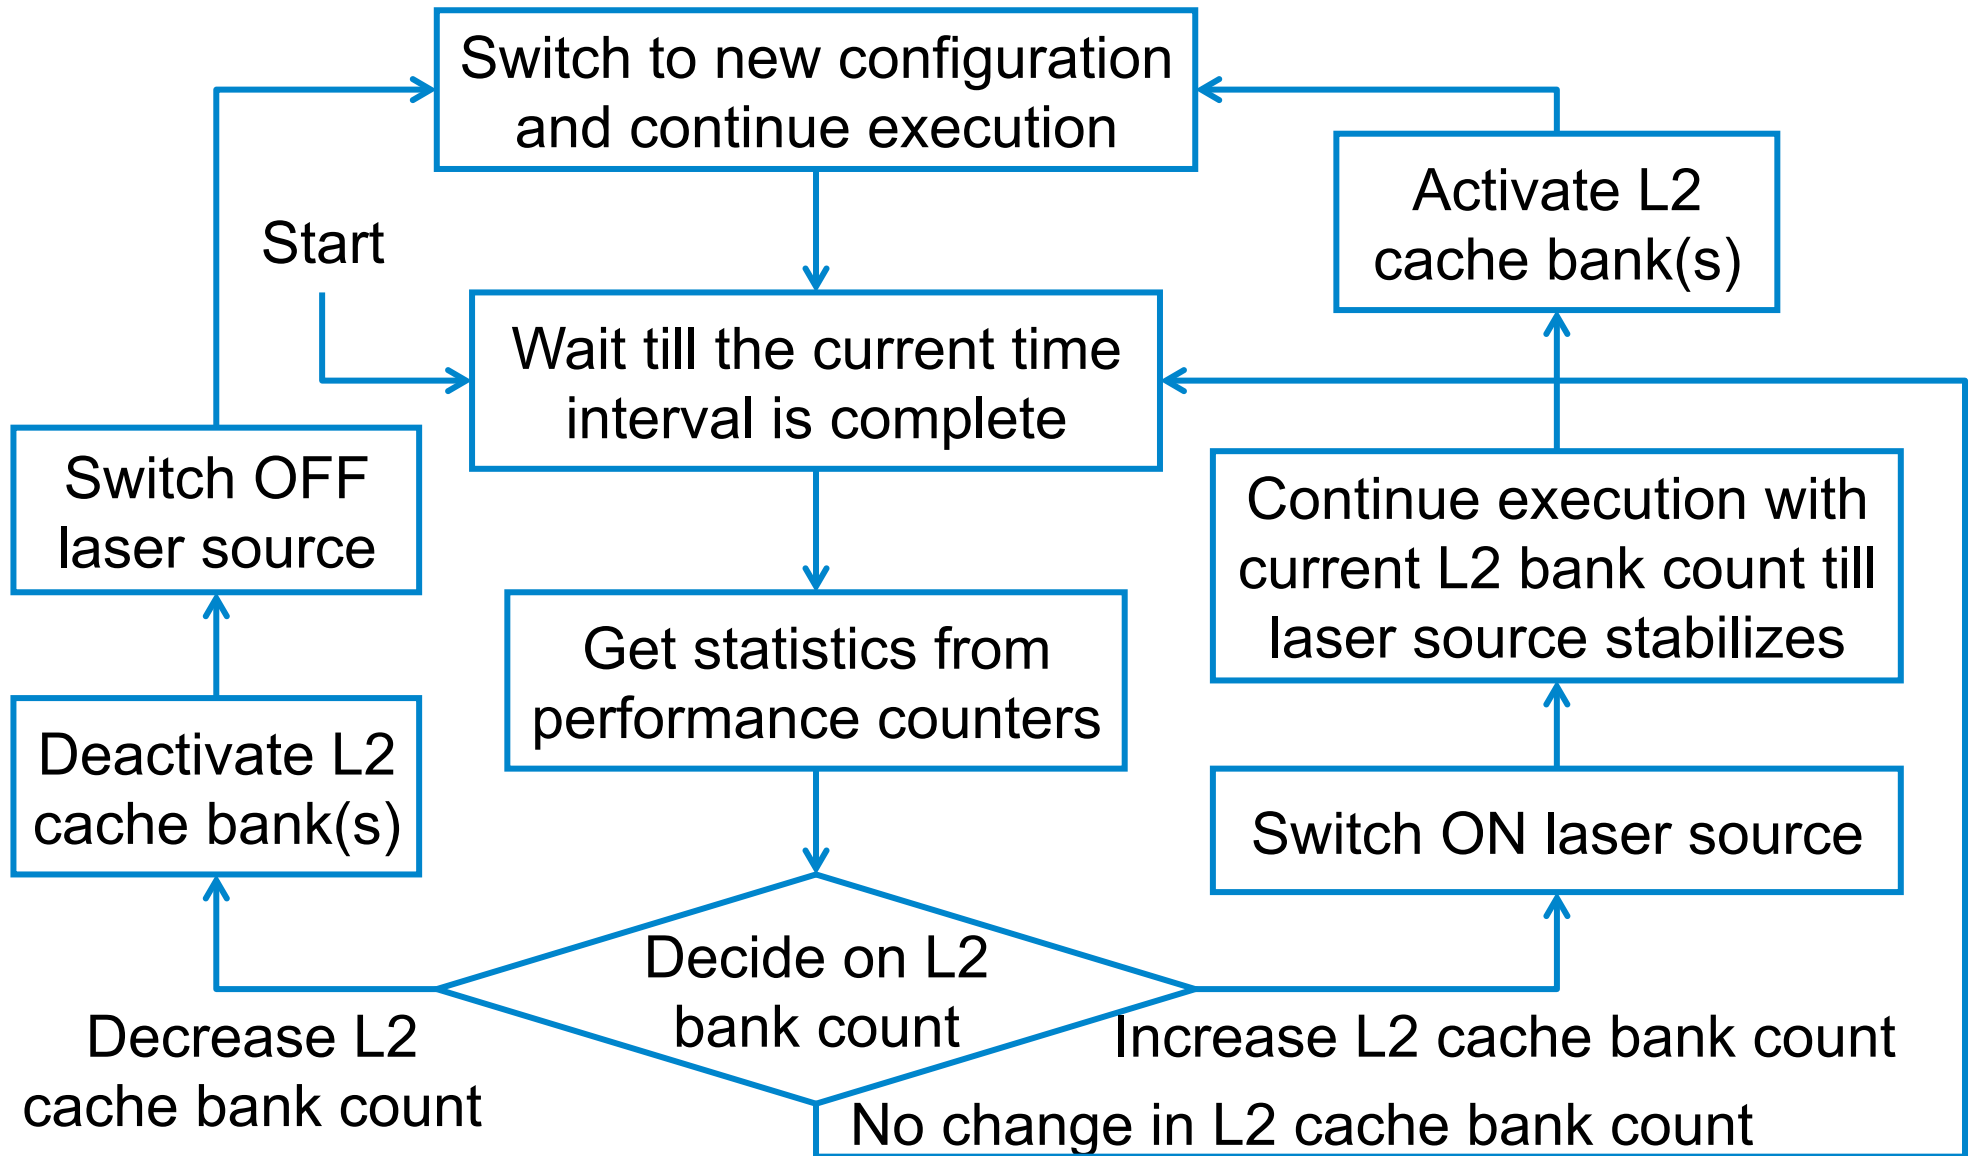
[Musalappa'06] [Sim'12] [Qureshi'07] [Wu'11]

Decision on Cache Reconfiguration

- L2 cache replacement rate is used for making decision on increasing or decreasing L2 cache bank count
- Dual threshold approach is adopted
 - We use $T_{\text{high}}(\log) = -3$ to keep the average performance degradation across all benchmarks less than 10%
 - We use $T_{\text{low}}(\log) = -4.5$ to minimize fluctuations in L2 cache bank count

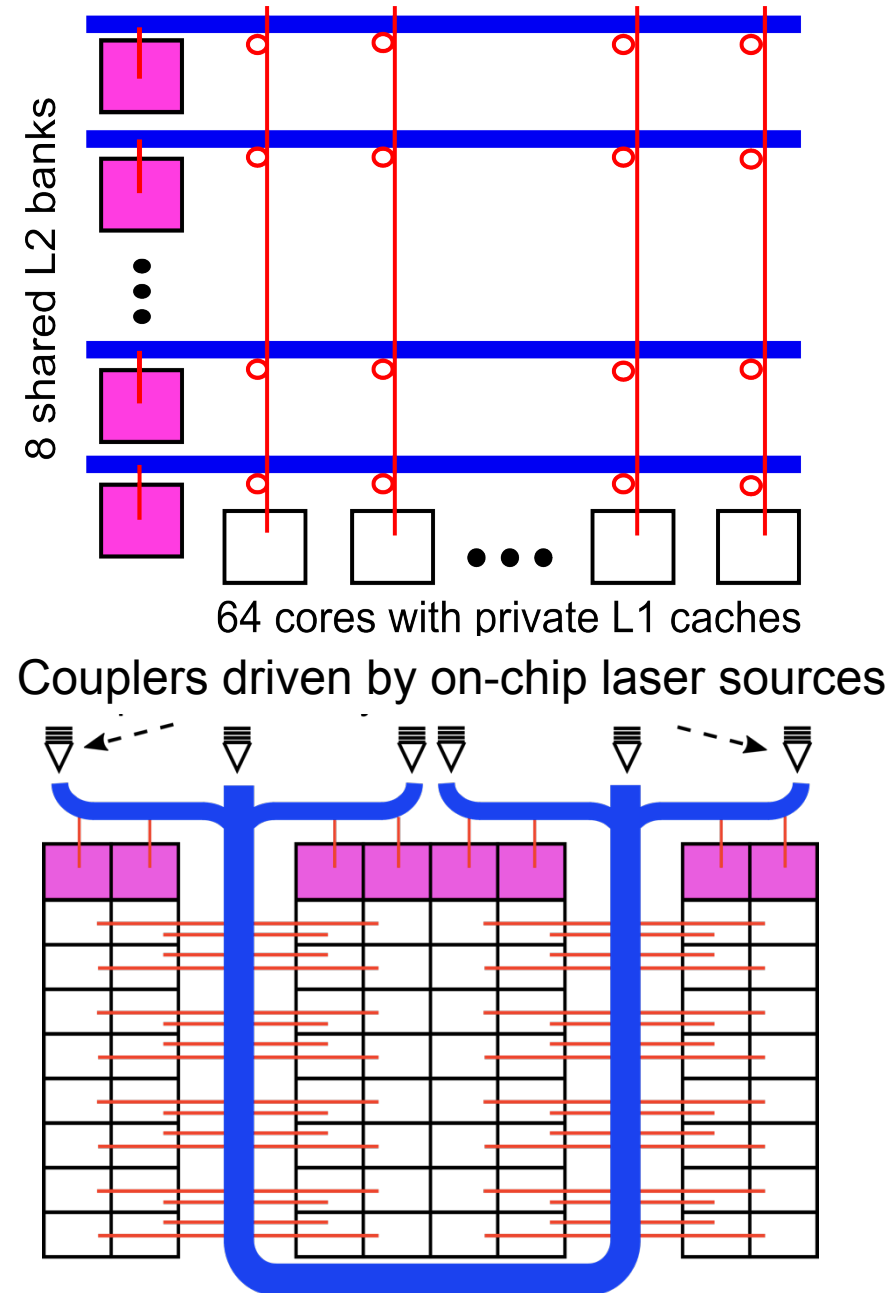


Reconfiguration Policy – Flow Chart

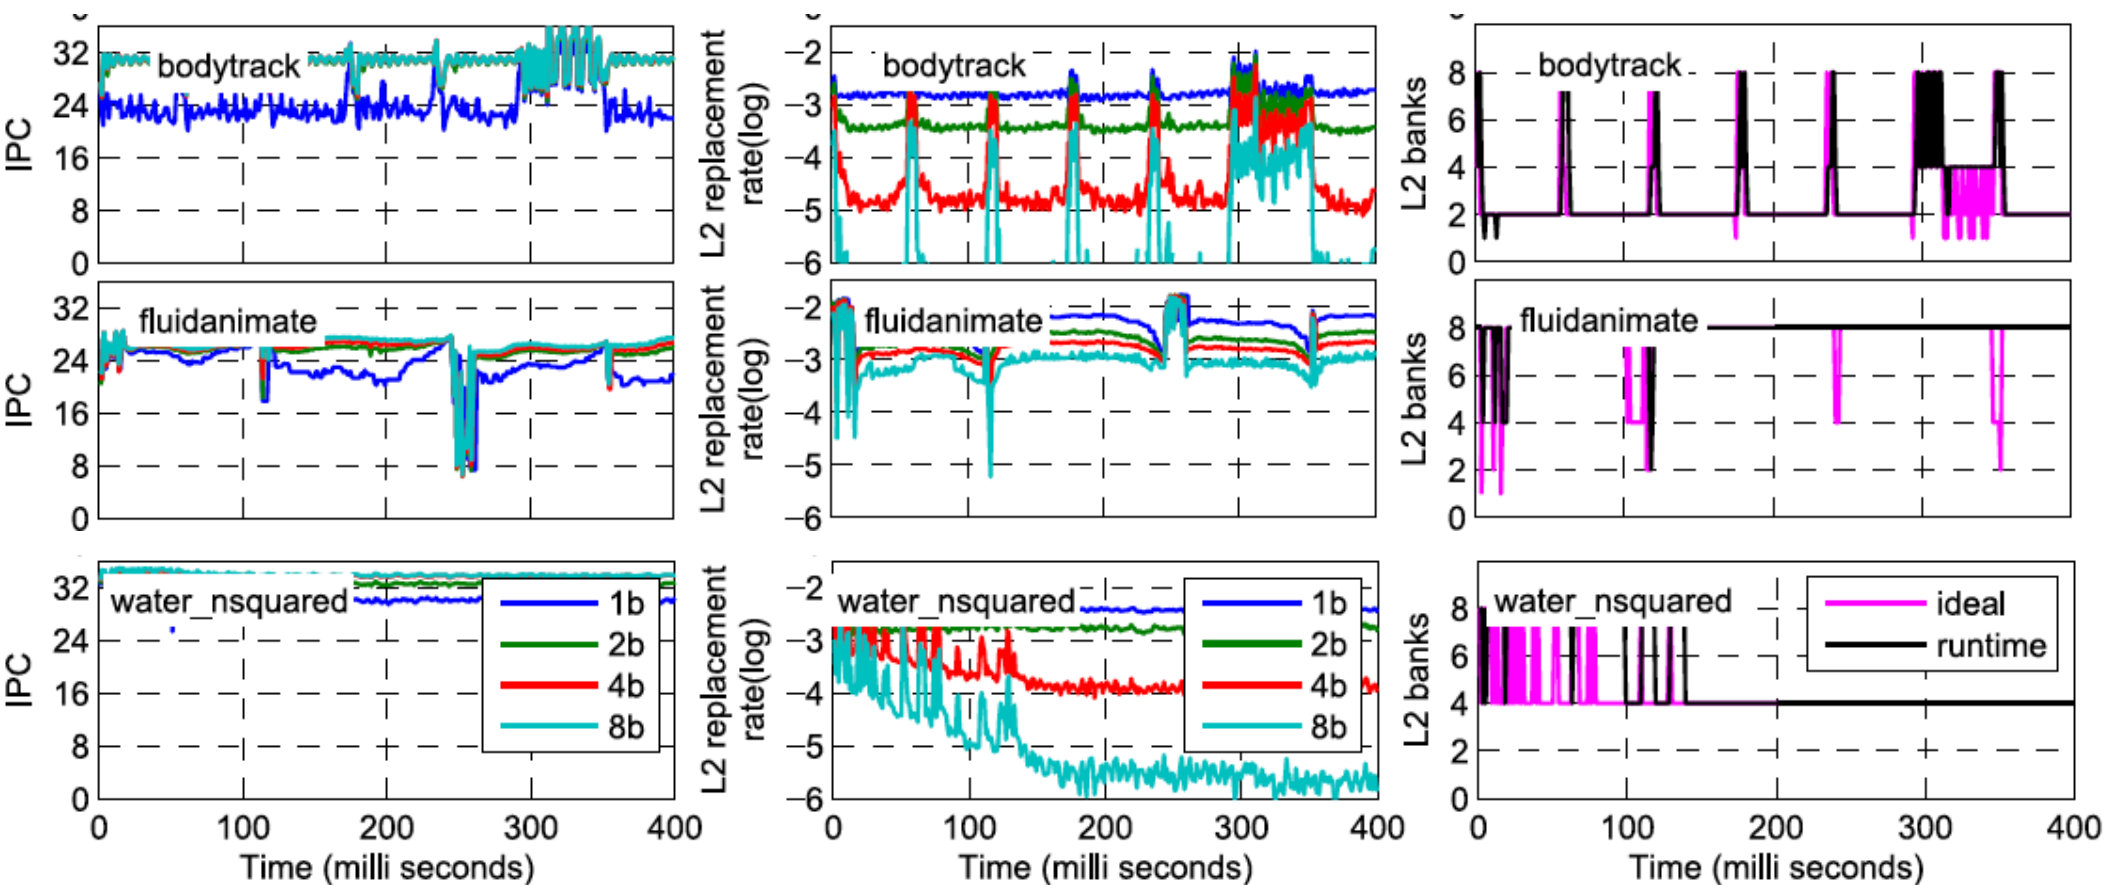


Target System

- 64-core in 22 nm @ 1.25 GHz
- 16 KB Private L1, 4 MB Shared L2 (8 banks)
- Crossbar NoC topology with 512-bit channels
 - L1-to-L2 communication uses Multiple Write Single Read (MWSR) arbitration
 - L2-to-L1 communication uses Single Write Multiple Read (SWMR) arbitration
- Evaluation
 - GEM5
 - McPat + Cacti + Inhouse setup for NoC



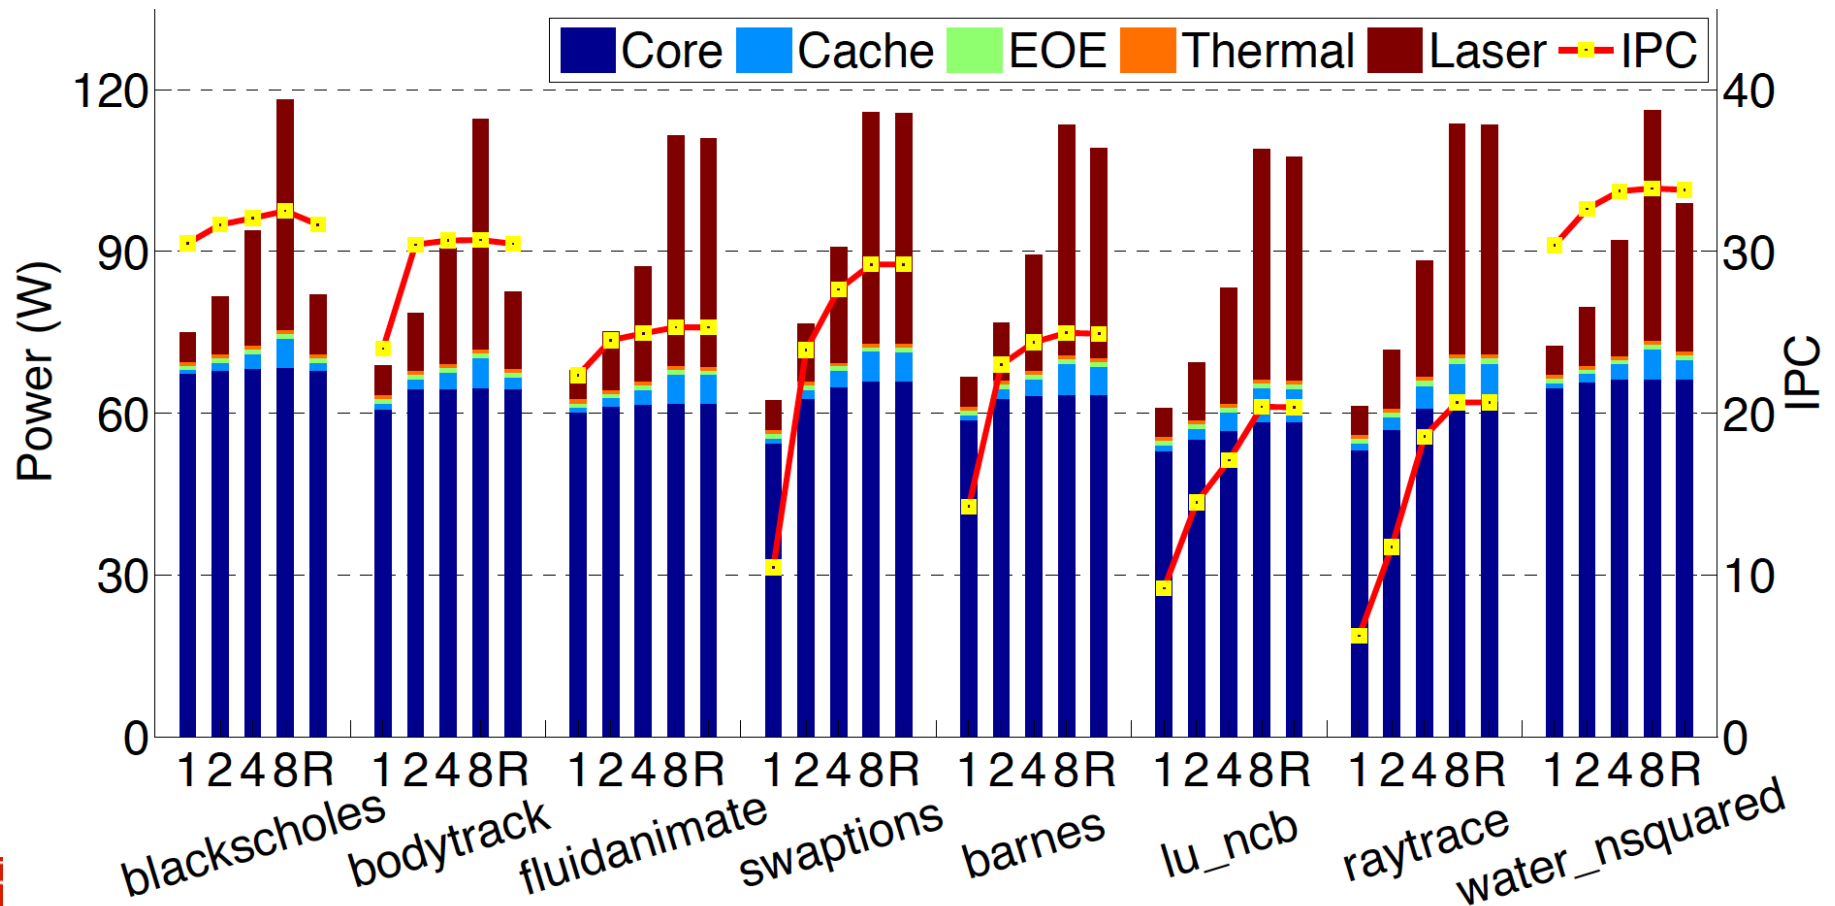
Evaluation – IPC, Replacement rate, Bank count



Evaluation – IPC, Power

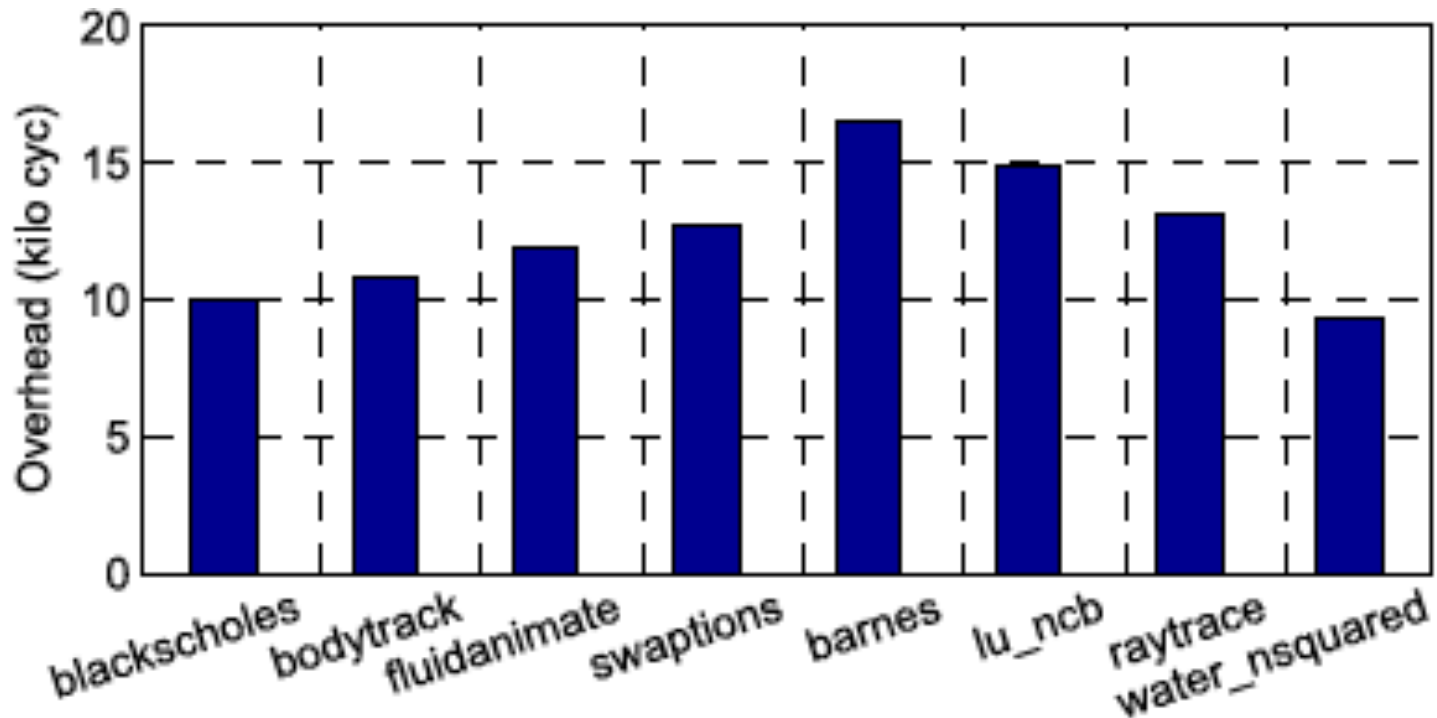
□ Target system: 64 cores and 8 L2 banks

- 23.8% saving in laser power on average (74.3% peak) with 0.65% IPC degradation on average (2.6% peak)
- 9.9% reduction in system power on average (30.6% peak) with 9.2% improvement in EDP on average (26.9% peak)



Evaluation – Reconfiguration Overhead

- Reconfiguration involves flushing L1 and L2 back to memory or fetching memory blocks to activated L2 banks
 - Upto 18,000 cycles required for reconfiguration
 - 150 uJ energy overhead for DRAM accesses



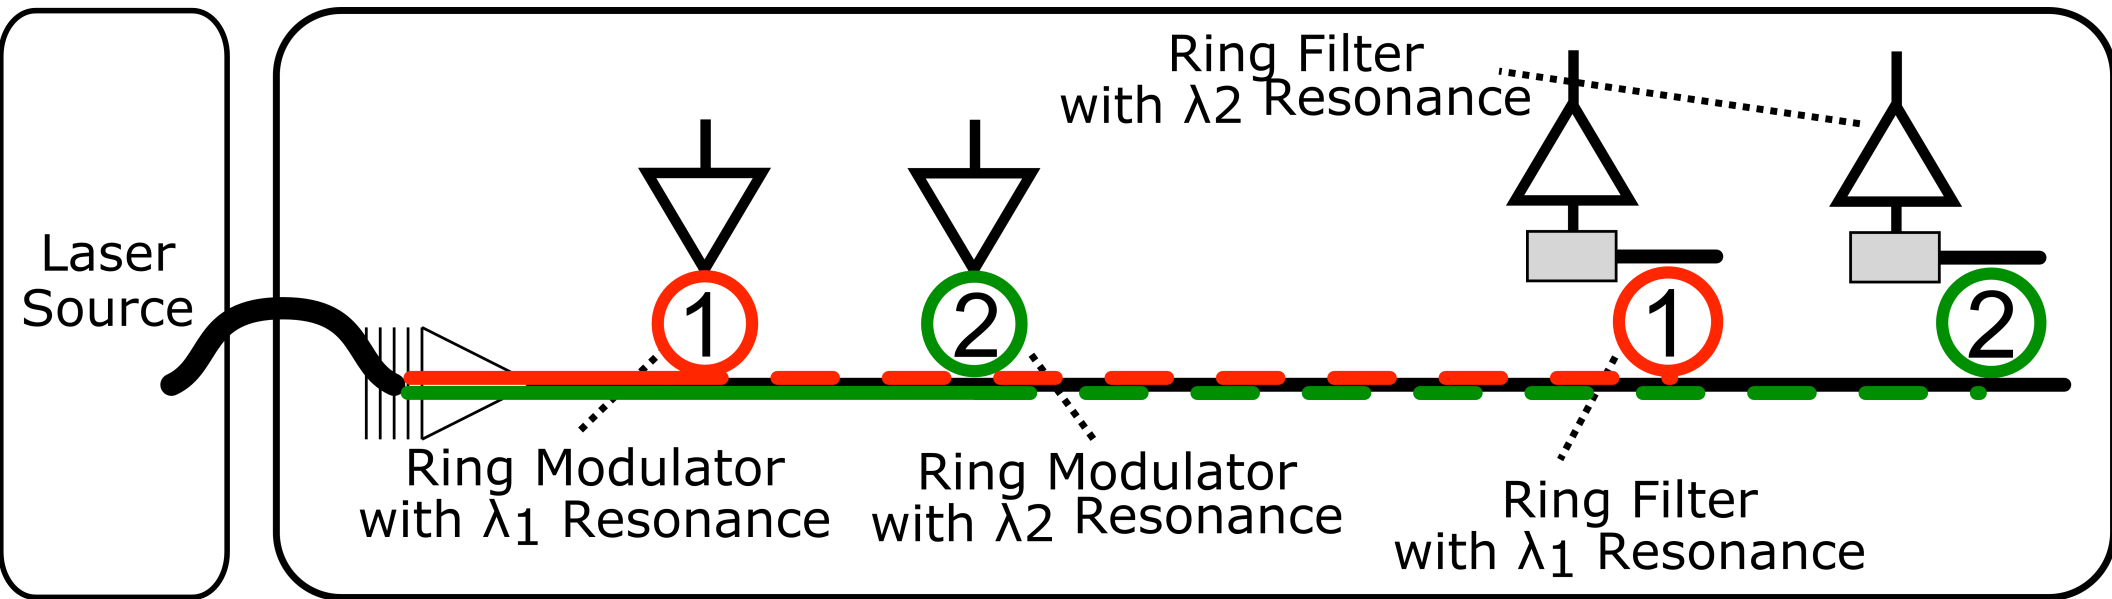
Summary I

- ❑ We proposed to manage the laser power by reconfiguring the NoC bandwidth based on the temporal and spatial variations in the cache size required by applications
- ❑ We adopted dual threshold approach to determine the L2 bank count at runtime
- ❑ On a 64-core target system, our proposed technique reduces laser power and system power by 23.8% and 9.9%, respectively, and improves EDP by 9.2% on average

Outline

- Background
- Laser Power Management using NoC and Cache Reconfiguration
- Thermal management using job allocation
- Summary

Silicon photonic link - DWDM



- Dense WDM (as much as 64 λ /wg, 10 Gbps/ λ) improves bandwidth density
- Thermal management becomes very challenging
 - Thermal tuning could cost more than 10 W of power

Related Work



- Thermal control through novel ring filters/modulators designs
 - Cladding [Djordjevic, Optical Exp.' 2013]
 - Heaters [Zhou, TACO'2010; Li, TVLSI'2012]
 - Mach-Zehnder interferometers [Biswajeet, Optical Exp.'2010]
- Techniques for thermal management in manycore systems
 - DVFS [Quan, DAC'2001]
 - Workload migration [Zhou, TACO'2010]
 - Liquid cooling [Coskun, VLSI-SoC'2009]

We use workload allocation technique to minimize the thermal gradients among photonic devices

Target Manycore System Architecture

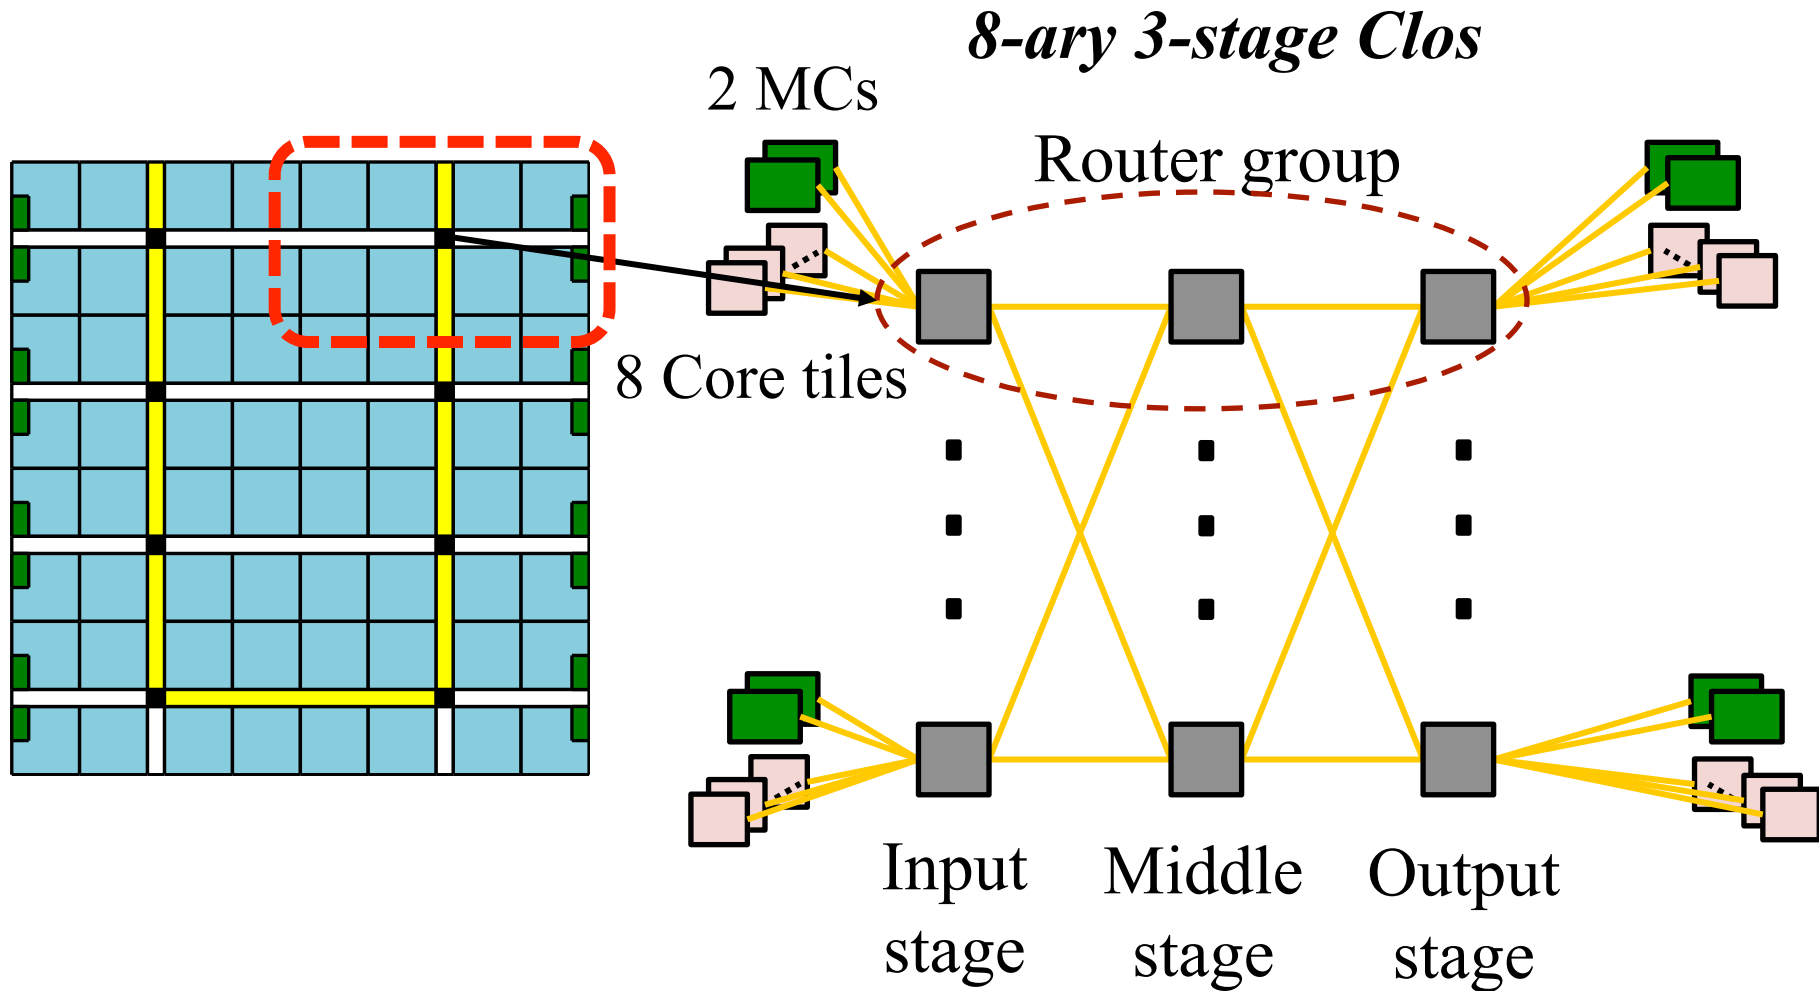
- 256-core system with Clos network

Core Architecture: IA-32 core in Intel SCC [*Howard, ISSCC2011*],
 16KB I/D L1 cache & 256KB L2 cache;
 Average power consumption: 1.166W



Target Manycore System Architecture

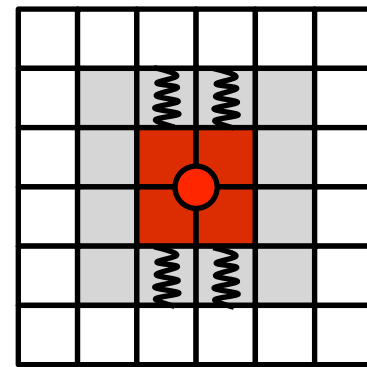
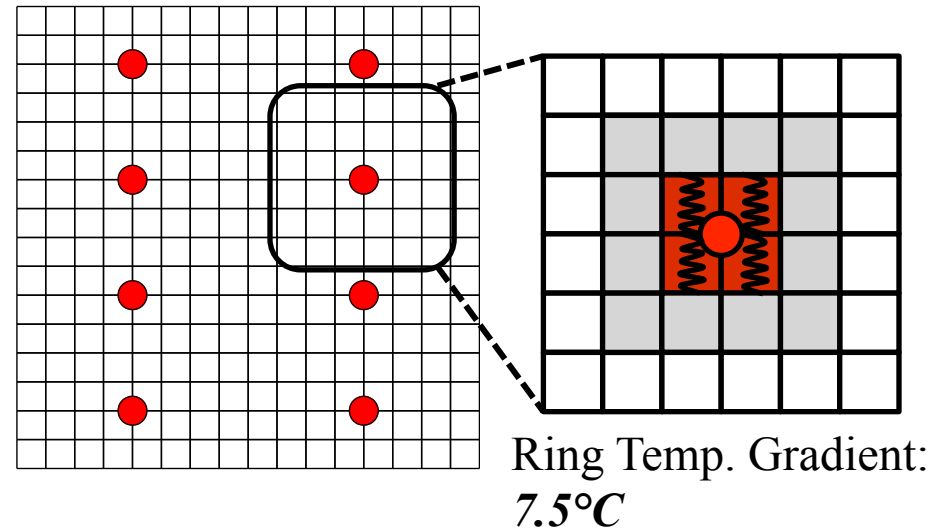
- 256-core system with Clos network



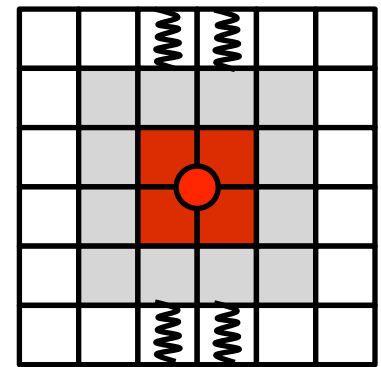
Ring-Aware Thermal Management

- Core power impacts on ring temperature
- Goals:
 - Minimize the difference among ring temperatures
 - Reduce the overall chip temperature
- Approach:
 - Classify the cores based on their distances to a ring group
 - Determine the impact of each region on ring temperature
 - Allocate threads so as to minimize thermal gradients

● Rings ■ RD0 cores □ RD1 cores
□ RD2 cores ⚡ Threads



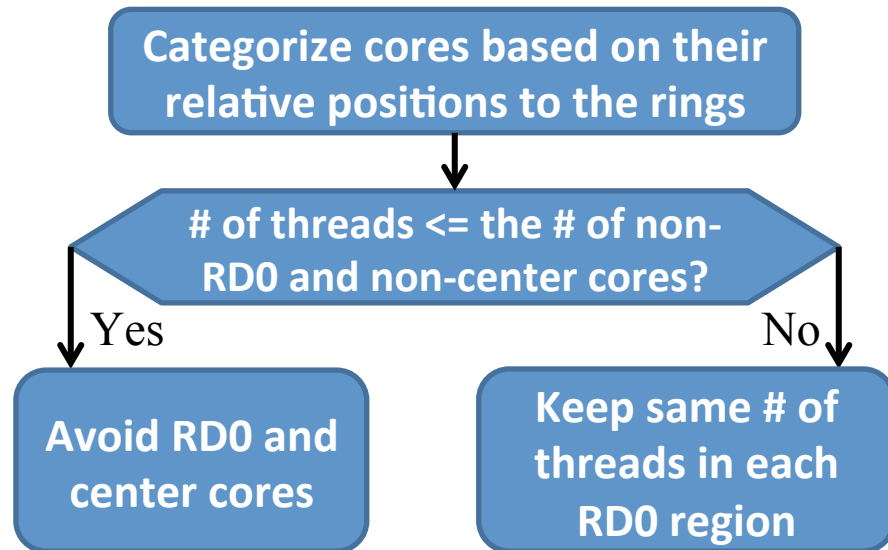

$<1^{\circ}\text{C}$



$<1^{\circ}\text{C}$

Ring-Aware Thermal Management

- **Ring-aware workload allocation**



- **Multi-program support**

- Sort the applications and threads based on their power dissipation & allocate high-power application first

Experimental Methodology



□ Simulation Platform

- Performance: Sniper simulator + SPLASH-2 & PARSEC
- Power: McPAT 0.8 + Temperature dependent leakage power model
- Thermal: HotSpot 5.02

□ Simulated Systems:

- 256-core system with silicon-phonic Clos NoC
- Tech. Node: 22nm; Area: 340 mm²
- Single-application and multi-program workloads
- Utilization scenarios: 32, 64, 96, 128, 156, 180, 206, 230 and 256 threads

Single application workload Multi-program workload

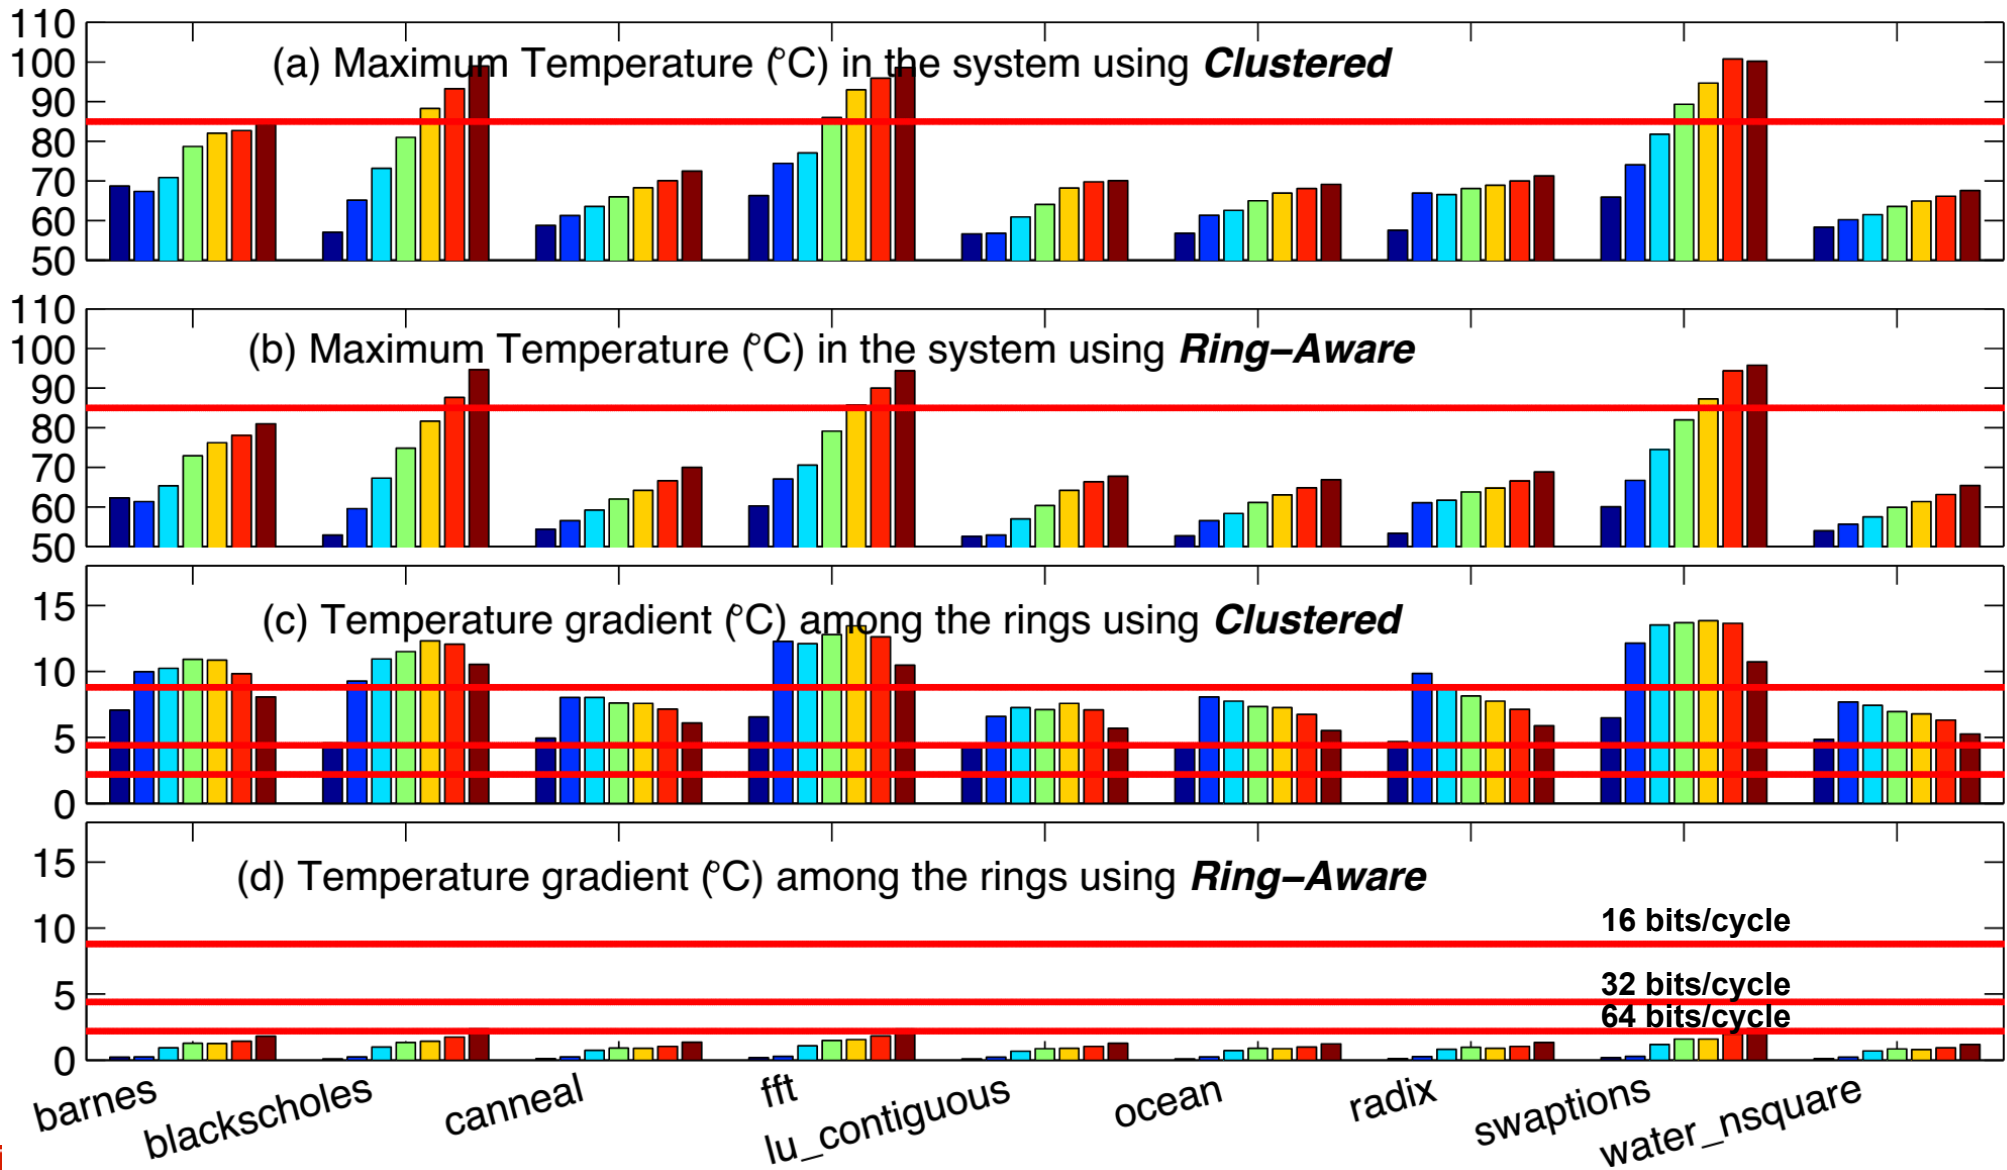


– *Allocation Policies:*

- *Clustered*
- *Chessboard*
- *Ring-Aware*

Evaluation of Single-Application Workloads

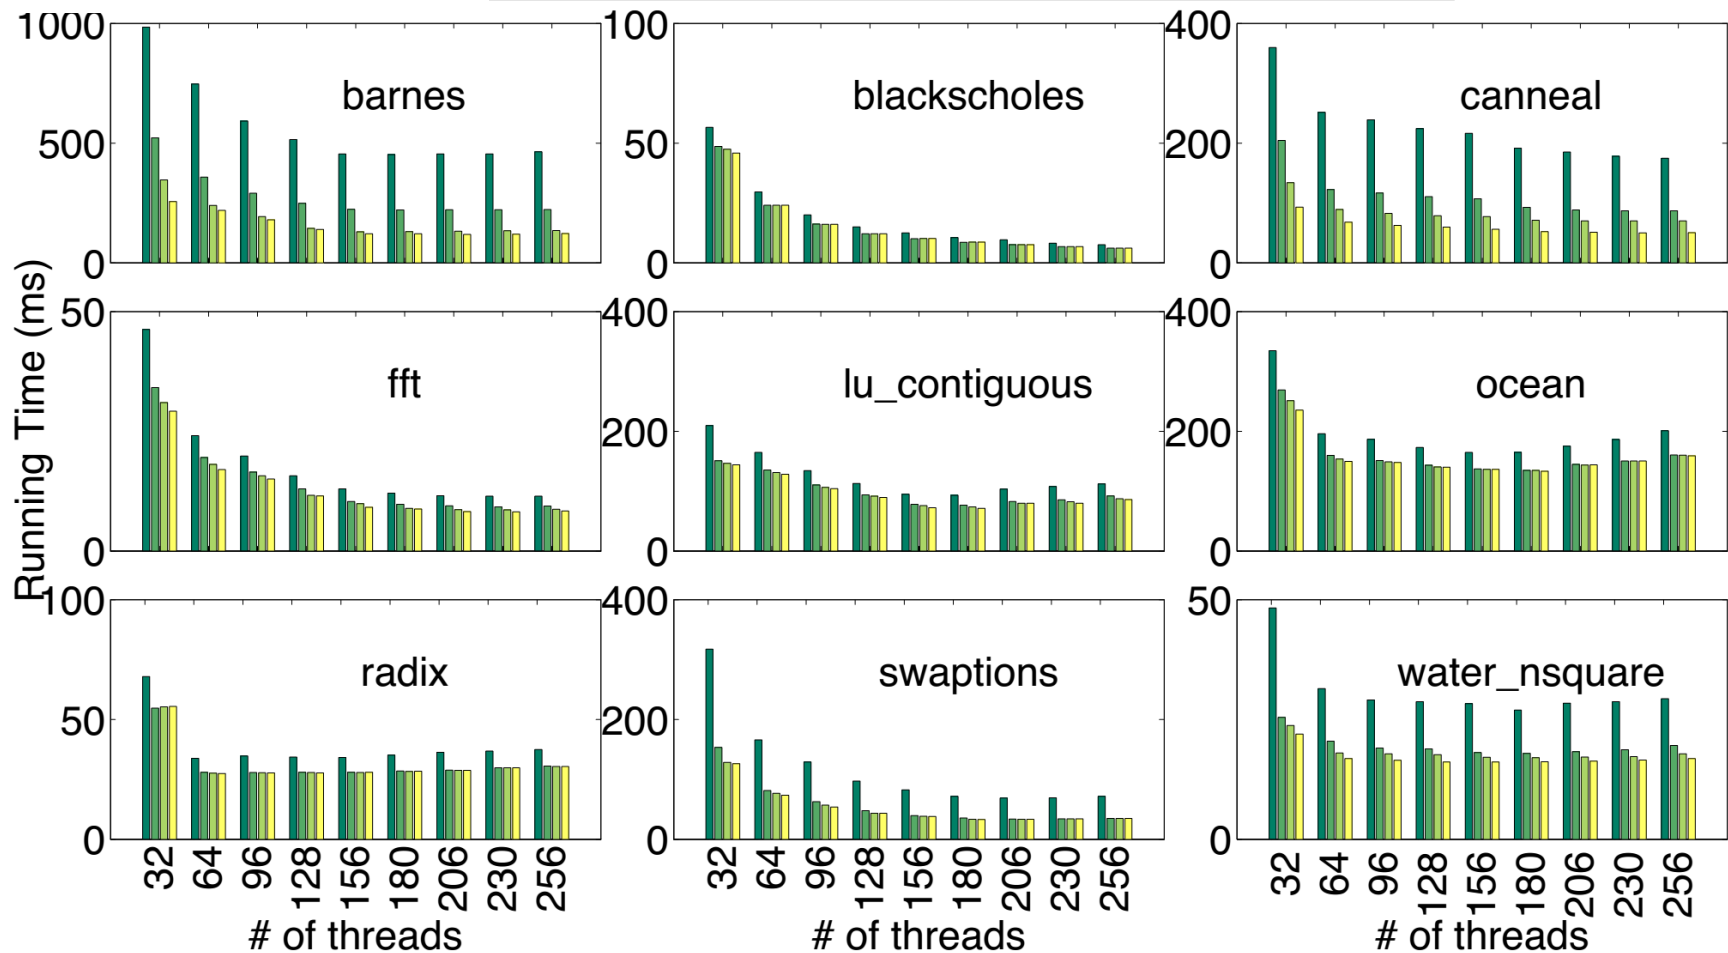
of threads: 32 64 96 128 156 180 206



Performance Results

Ring Temp. Gradient: $>8.8^{\circ}\text{C}$ $\leq 8.8^{\circ}\text{C}$ $\leq 4.4^{\circ}\text{C}$ $\leq 2.2^{\circ}\text{C}$




Channel bandwidth: 8bits/cyc 16bits/cyc 32bits/cyc 64bits/cyc




When using more than 50% of the cores, several applications have significantly better performance with the Ring-Aware approach

Evaluation of Multi-Program Workloads

- Mapping Policies: In-order left (*Inorder*), random (*Rand*), *Proposed*



Low-pow app
 High-pow app



Evaluation of Multi-Program Workloads


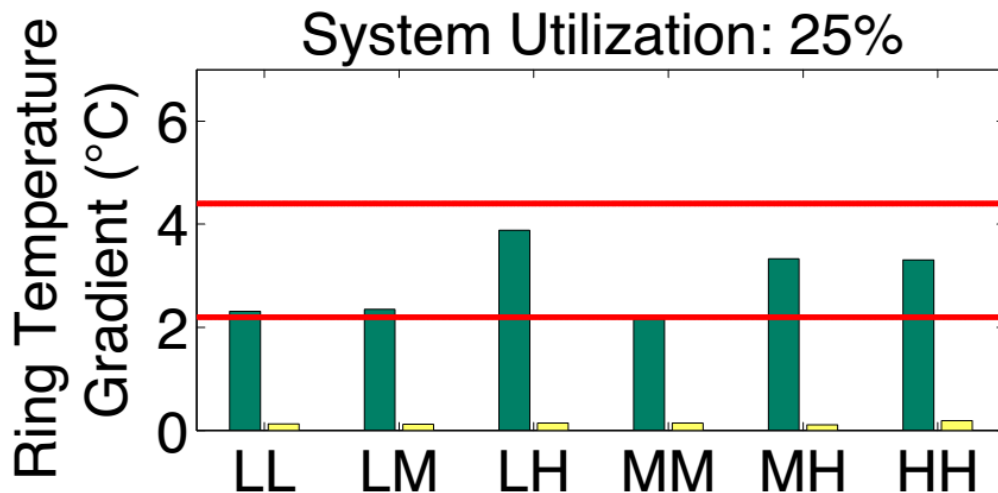
□ Diverse multi-program workloads

L: low-power

M: medium-power

H: high-power

LL		water_nsquare (L), lu_contiguous (L)		HH		barnes (H), fft (H)
LH		barnes (H), lu_contiguous (L)		LM		canneal (M), ocean (L)
MM		radix (M), blackscholes (M)		MH		radix (M), swaptions (H)



Chess_Inorder results in 6.1%, 13.8% and 8% lower performance compared to Ring_Proposed for LH, MH and HH, respectively.

Summary II

- We used a cross-layer approach for thermal analysis & design of silicon photonic NoC
- We proposed a Ring-Aware job allocation policy to reduce thermal gradients among photonic devices
- Our policy enables us to operate the photonic links at their maximum bandwidth and in turn maximize application performance

Summary

- Numerous challenges need to be overcome to make photonic NoC viable
 - Bandwidth utilization
 - Packaging/Integration
 - Power consumption
 - NoC and Cache reconfiguration can be used to lower laser power
 - Software-based workload allocation policy can be used to reduce thermal tuning power