2.5D/3D Systems with Silicon Photonic NoCs: Efficient Thermal Management, Opportunities, and Challenges

Ayse K. Coskun

Boston University, ECE Department

In collaboration with:

Ajay Joshi[1], Andrew B. Kahng[2,3], Jonathan Klamkin[4],

Tiansheng Zhang[1], Yenai Ma[1]

[1]Boston University ECE Dept.;

UCSD [2]ECE and [3]CSE Dept.;

[4]UCSB ECE Dept.

This research has been partially funded by the NSF grants CNS-1149703 and CCF-1149549.
Work at UCSD has been supported by NSF, Samsung and the IMPACT+ Center.

Today's Multi-/Many-core Computing Systems

- **Due to technology scaling & high computation needs, more resources are integrated on-chip**

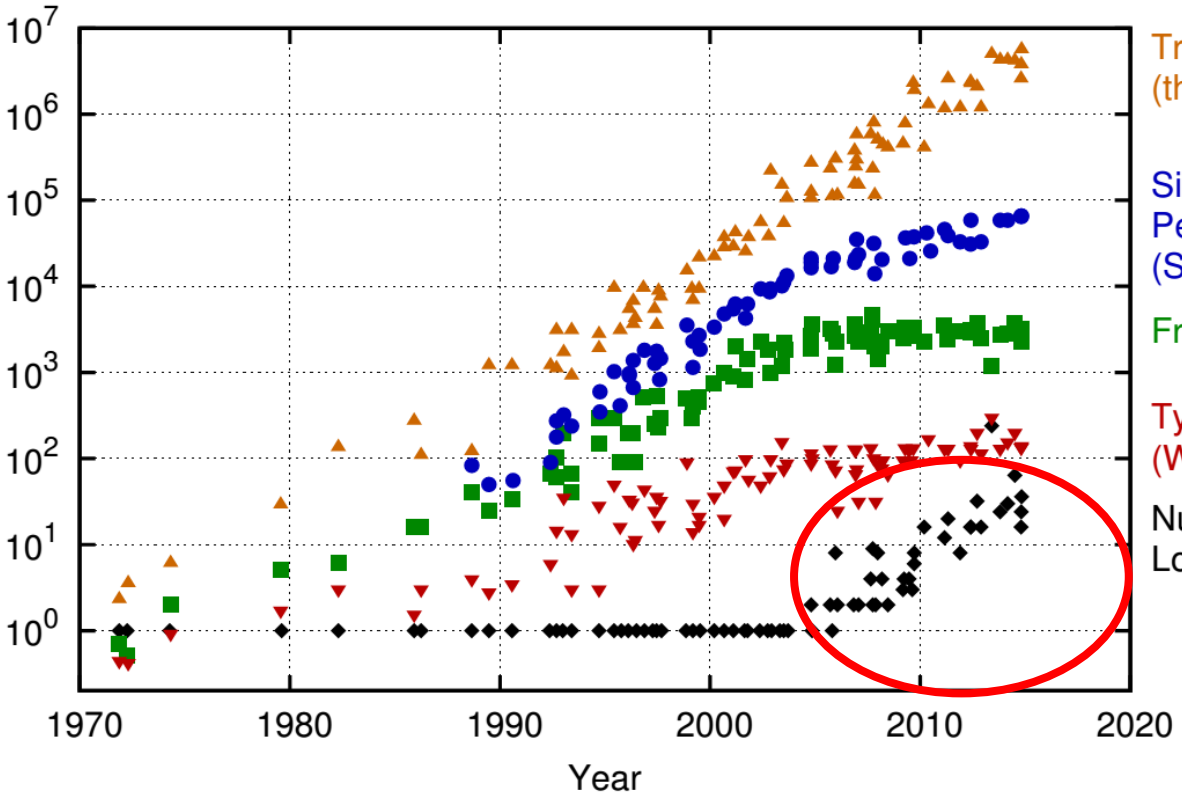

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2015 by K. Rupp

[Rupp, 40 years of microprocessor trend data, 2015]

Intel SCC
(48 cores, 2010)

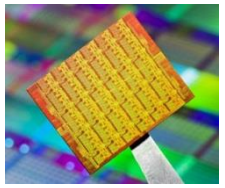

Tilera Tile Gx
(72 cores, 2012)

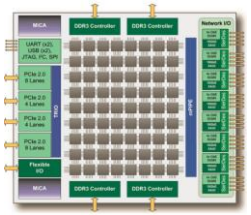

Intel Xeon Phi
(72 cores, 2015)

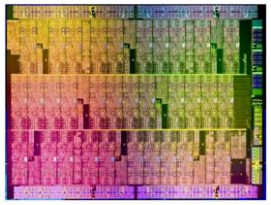

3D Stacking Technology & Its Benefits

- **3D stacking technology is a promising integration technology for future computation system design**
  - More on-chip resources compared to 2D designs

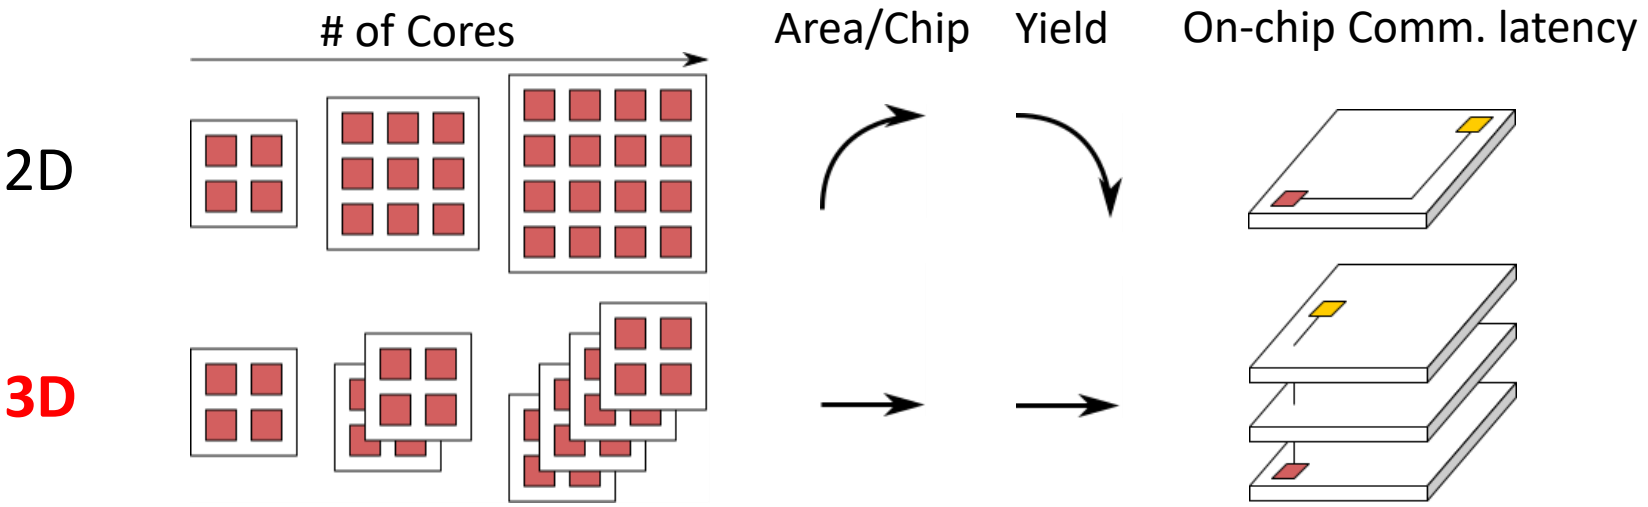

  - Various technologies integrated on a single chip
    - *On-chip stacking DRAM*
    - ***Silicon-photonic Network-on-Chip (PNoC)***

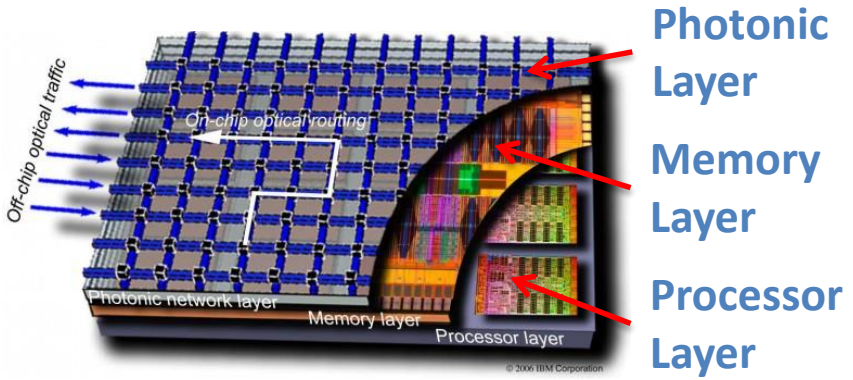

http://researcher.watson.ibm.com/researcher/view_group.php?id=2757

Challenges of 3D Stacking Technology

- **On-chip Resource Management**
  - *Under utilized resources* → **Performance and energy efficiency benefits left on the table**
  - *Increased power density* → **Potential thermal violations**

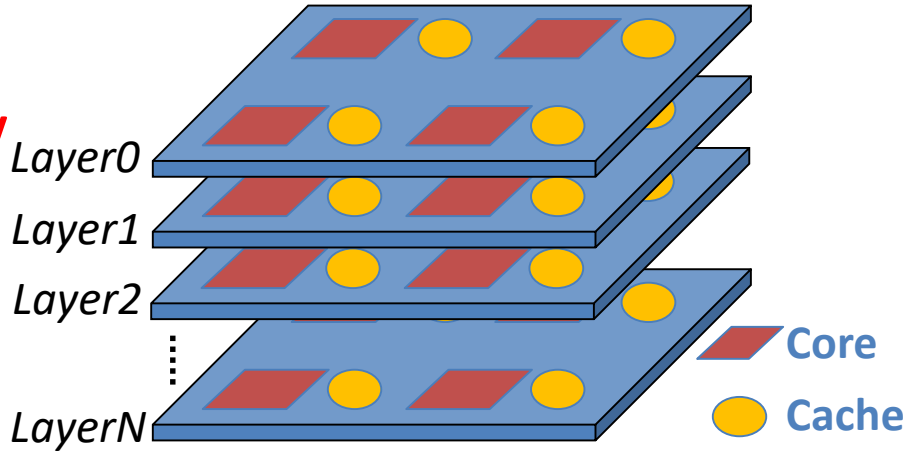

- **On-chip Thermal Management**
  - *Thermal and process sensitivity of devices in other technologies* → **Resilience problems or high power consumption**

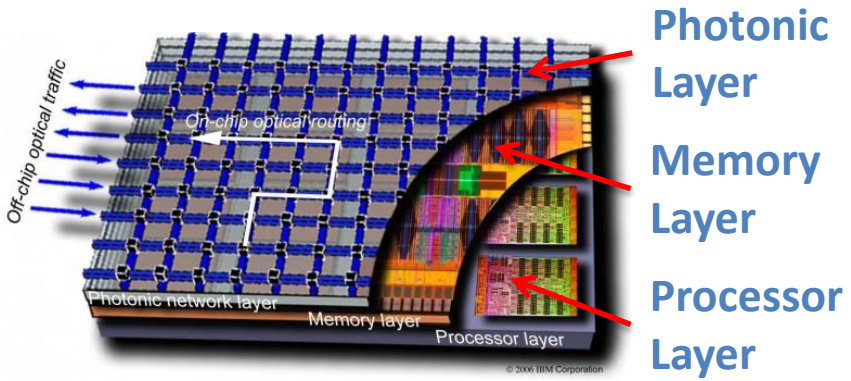

http://researcher.watson.ibm.com/researcher/view_group.php?id=2757

Silicon-Photonics Network-on-Chip

- **Silicon-Photonic Link**

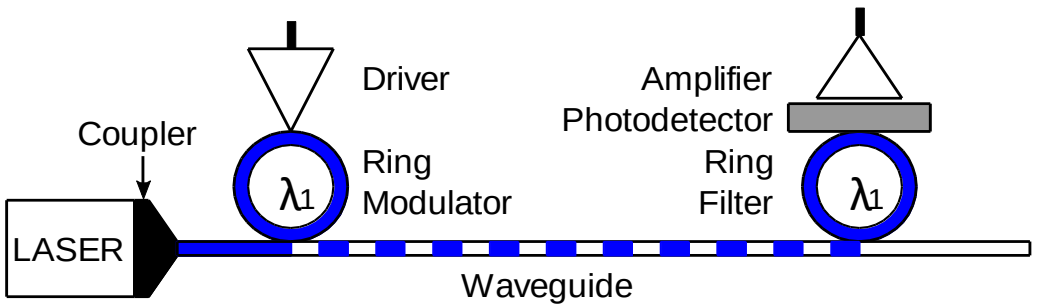

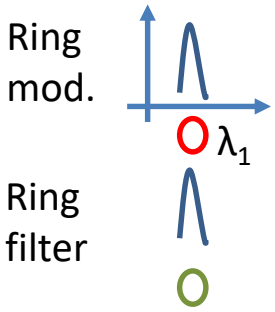

Integration methods:

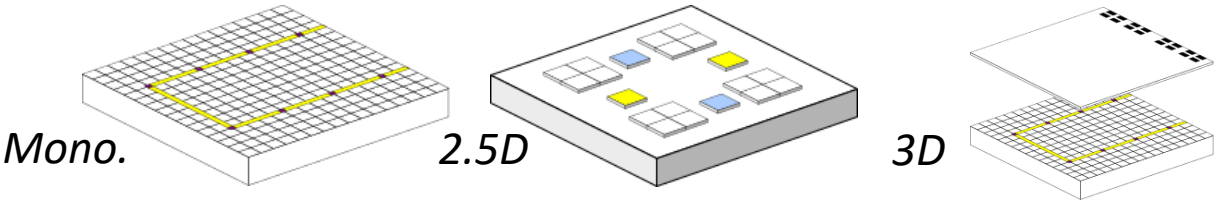

- **Silicon-Photonic Links vs. Electrical Links**

- Higher bandwidth density
- Lower long-distance communication latency
- Lower data-dependent energy consumption

- More sensitive to thermal variations
- More sensitive to process variations

Silicon-Photonics Network-on-Chip

- **Silicon-Photonic Link**

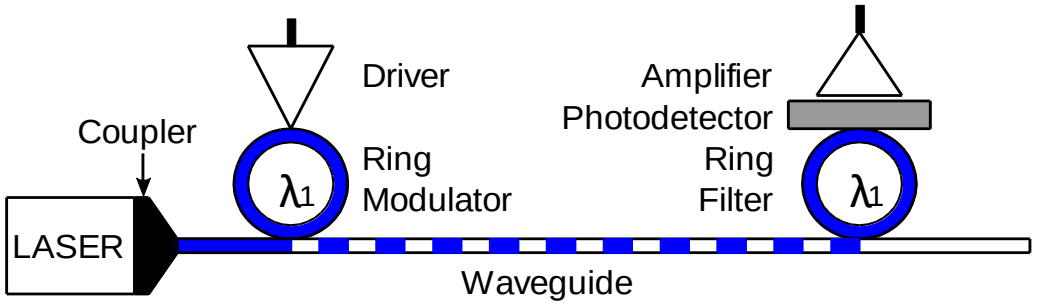

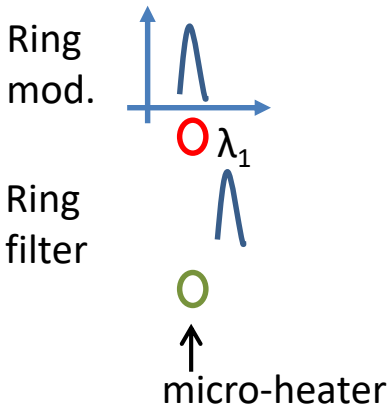

- **Silicon-Photonic Links vs. Electrical Links**

- Higher bandwidth density
- Lower long-distance communication latency
- Lower data-dependent energy consumption

- More sensitive to thermal variations
- More sensitive to process variations

} High thermal tuning power

- High optical loss
- Low laser source efficiency (due to high temp.)

} High laser source power

On-chip energy efficiency is a limiting factor for PNoC integration!

Manycore Systems with Silicon-Photonics NoCs

Integration methods:

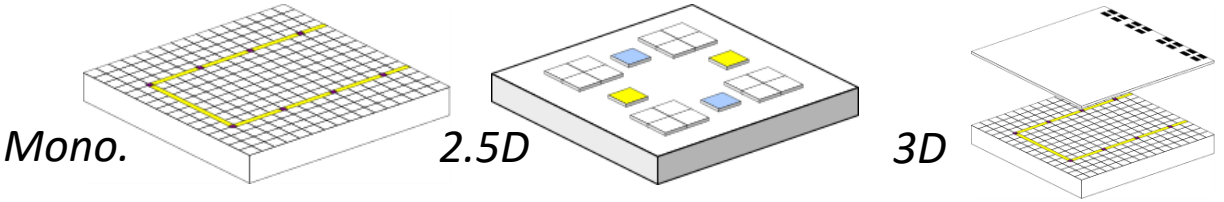

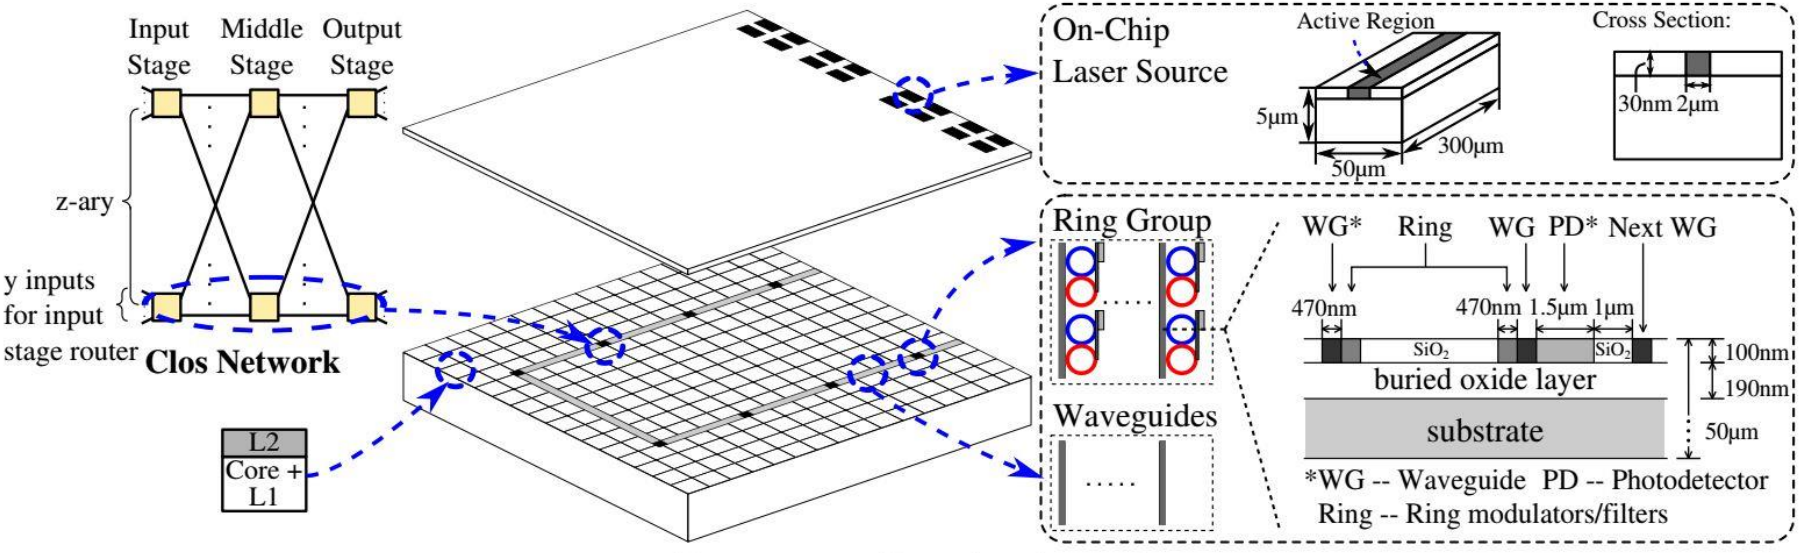

How do we address thermal sensitivity today?

Device-Level Techniques

- Cladding [***Djordjevic, Opt.Exp.'13***]
- Heaters [***Zhou, TACO'10] [Li, TVLSI'12***]

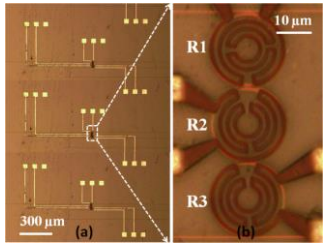

- Mach-Zehnder interferometers [***Biswajeet, Opt.Exp.'10***]

Runtime Management Techniques

- Aurora ***[Li, TCAD'15]***
  - Thermal Tuning
  - DVFS
  - Routing Algorithm

Design-Time Techniques

- Studies on PNoC placement's impact on signal to noise ratio ***[Li, DATE'15]***

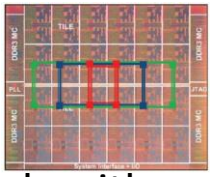

- Optical waveguide routing algorithms *to r*educe optical loss under a fixed netlist ***[Condrat, SLIP'08] [Ding, DAC'09] [Ramini, DATE'13]***

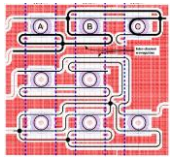

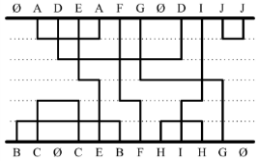

- *P & R Solutions for PNoC*
  - PROTON: An automatic tool for PNoC P & R ***[Boo, ICCAD'13]***
  - GLOW: A ILP based global router for PNoC ***[Ding, DAC'12]***

Our work aims at ***reducing the thermal tuning power & laser source power*** for PNoC via workload allocation, thermal tuning policies, and design-time techniques.

Cross-Layer Design Automation

Tooling for Design Space Exploration

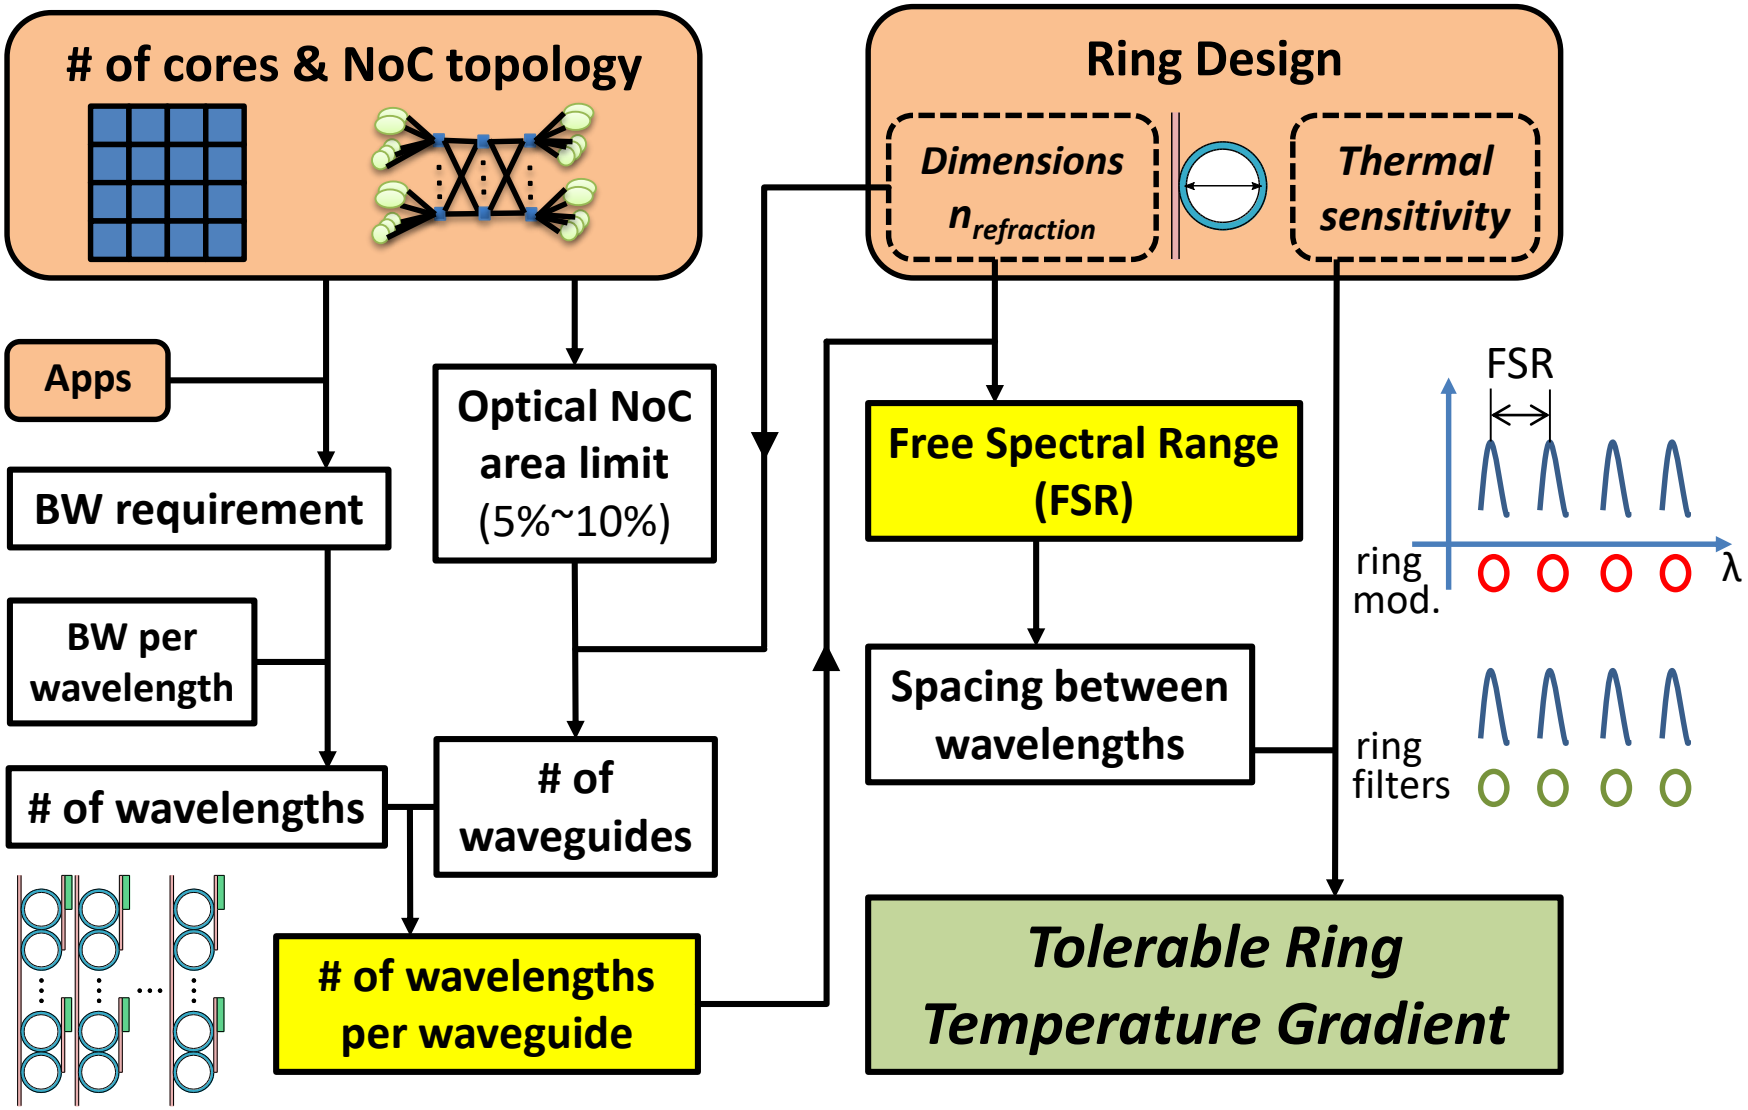

Target Many-core System w/ PNoC

[DATE'14, TCAD'17]

- 256-core system with Clos network

Core Architecture: IA-32 core in Intel SCC [***Howard,ISSCC'11***], 16KB I/D L1 cache & 256KB L2 cache;

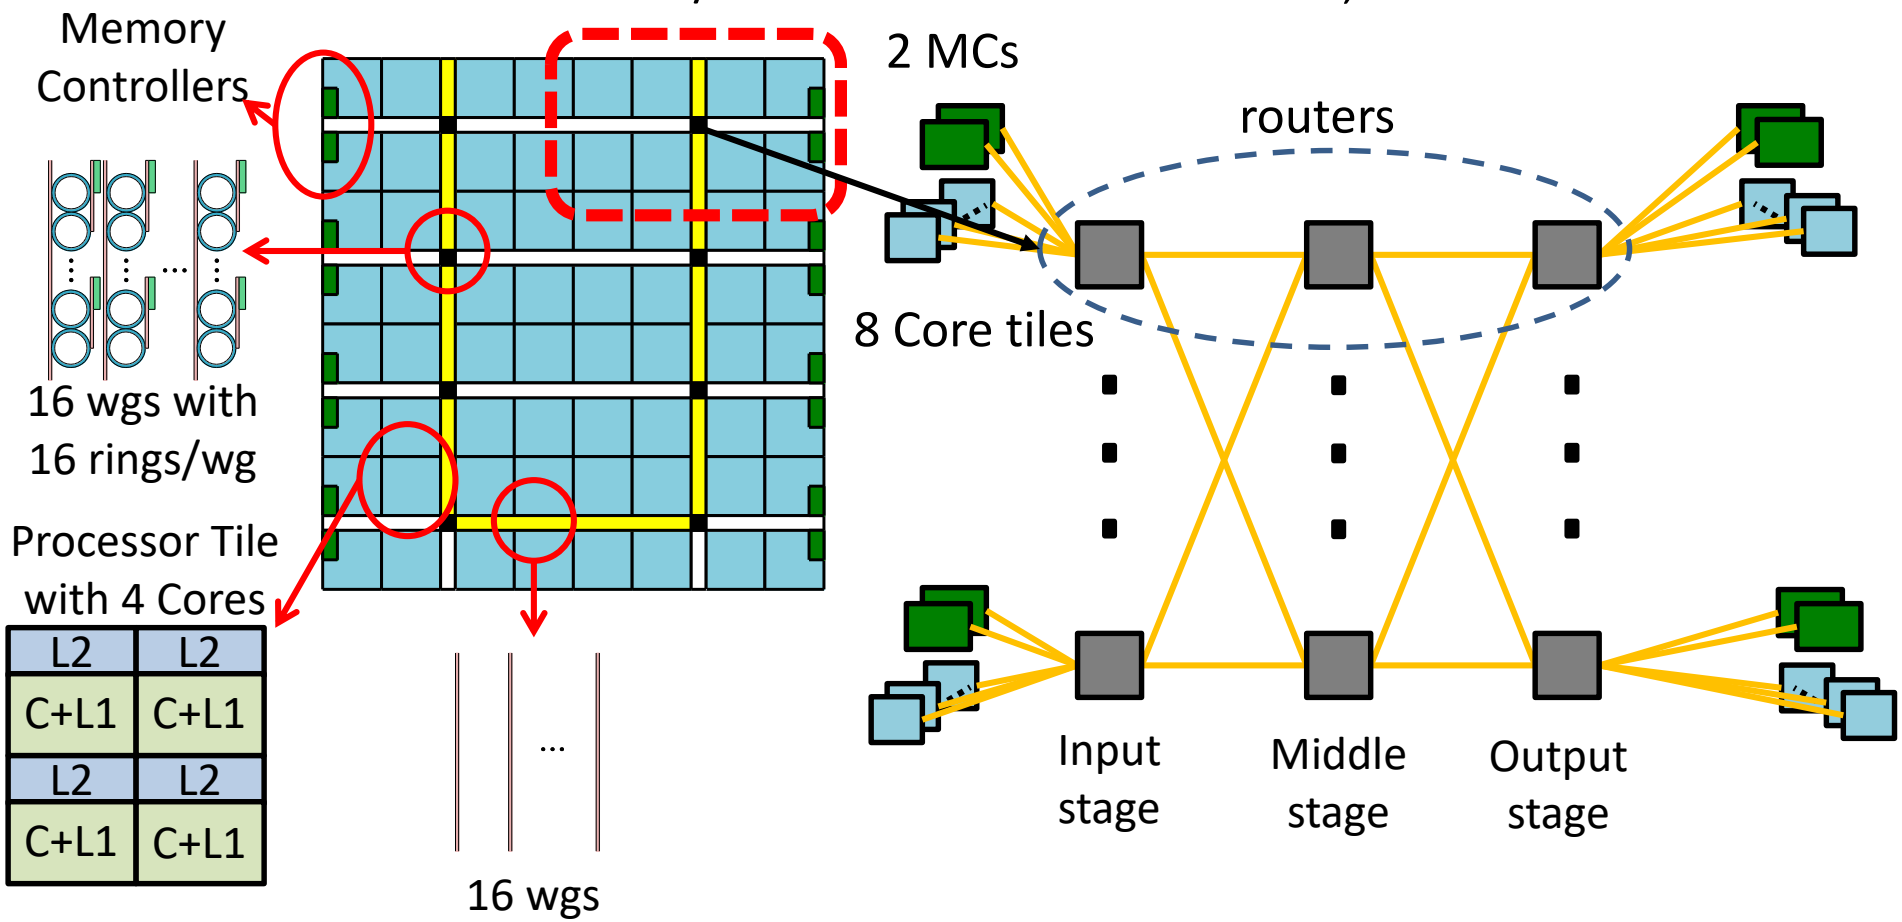

Floorplan Optimization Flow

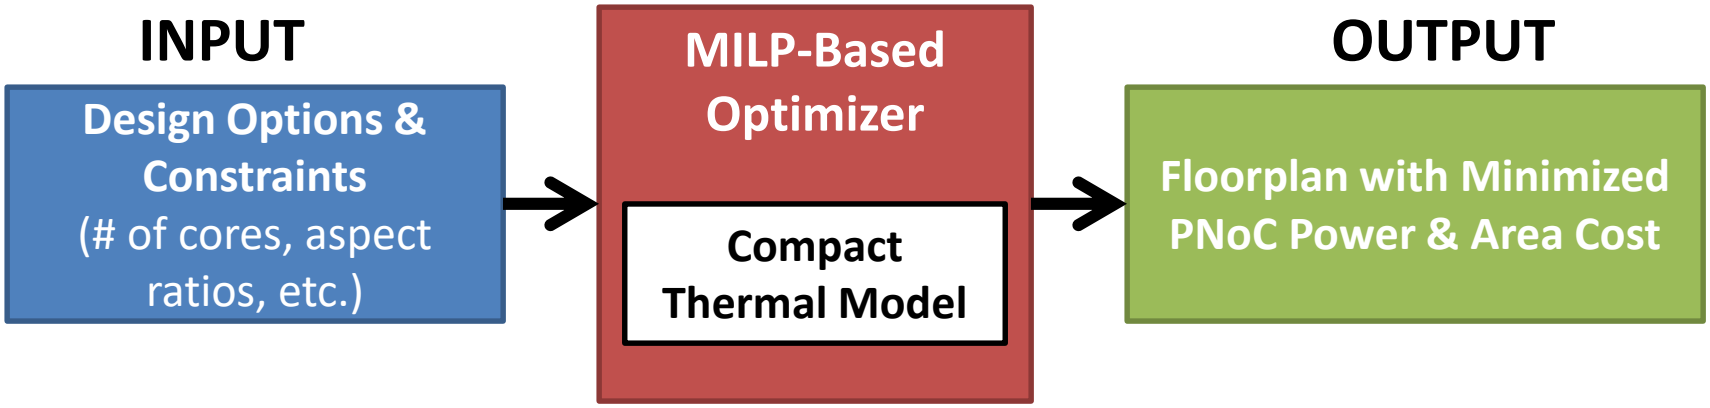

- **Optimization Goal:**

$$\textbf{Minimize:} \quad \alpha \cdot P_{PNoC} + \beta \cdot AREA_{PNoC}$$

$$P_{PNoC} = P_{laser} + P_{tuning} + P_{electrical}$$

  - **PNoC Power:**
    - P & R's impact on waveguide length, crossing and bending
    - Laser source efficiency
    - PNoC placement's impact on thermal tuning power
  - **PNoC Area:**
    - Area cost of router groups and waveguides

Floorplan Optimization Flow

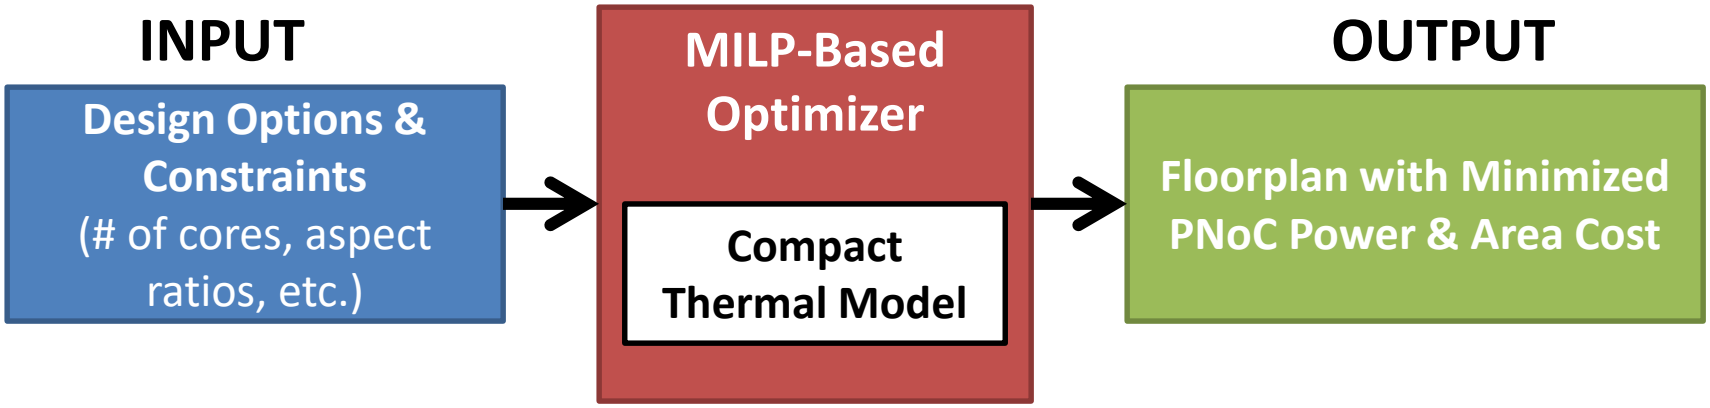

- ***Compact thermal model***

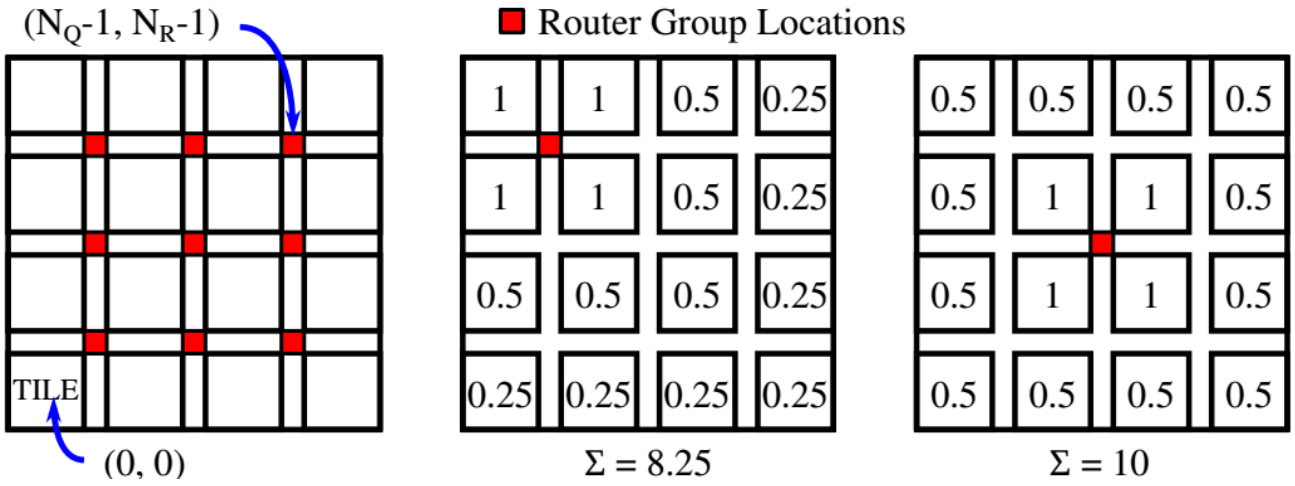

Floorplan Optimization Flow

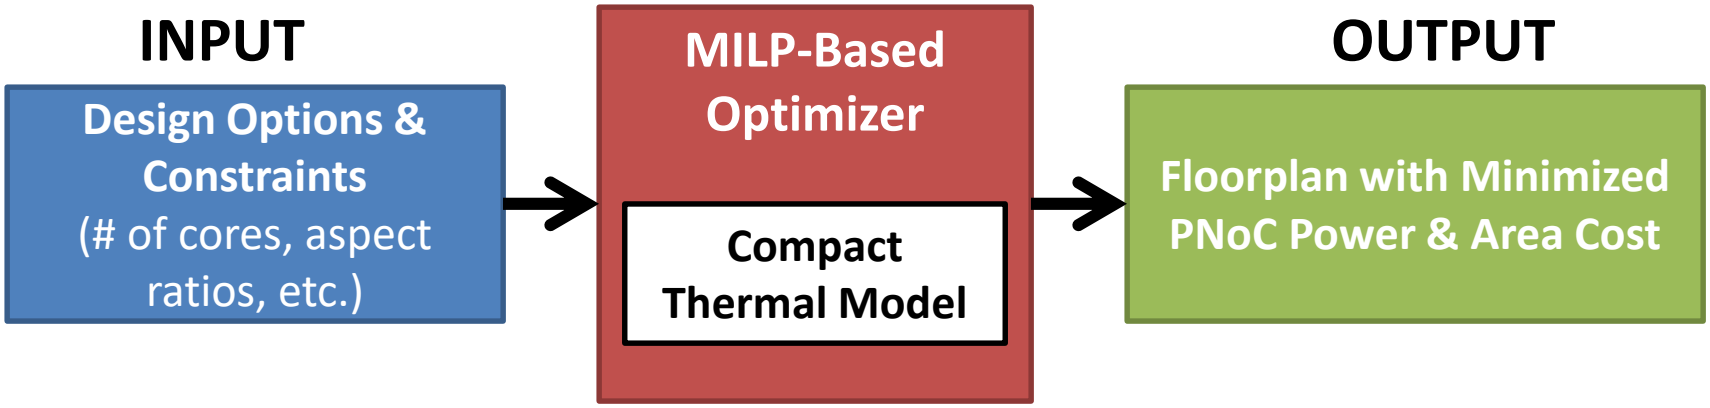

- ***Compact thermal model***

Power profile:

Accumulated thermal weight profiles

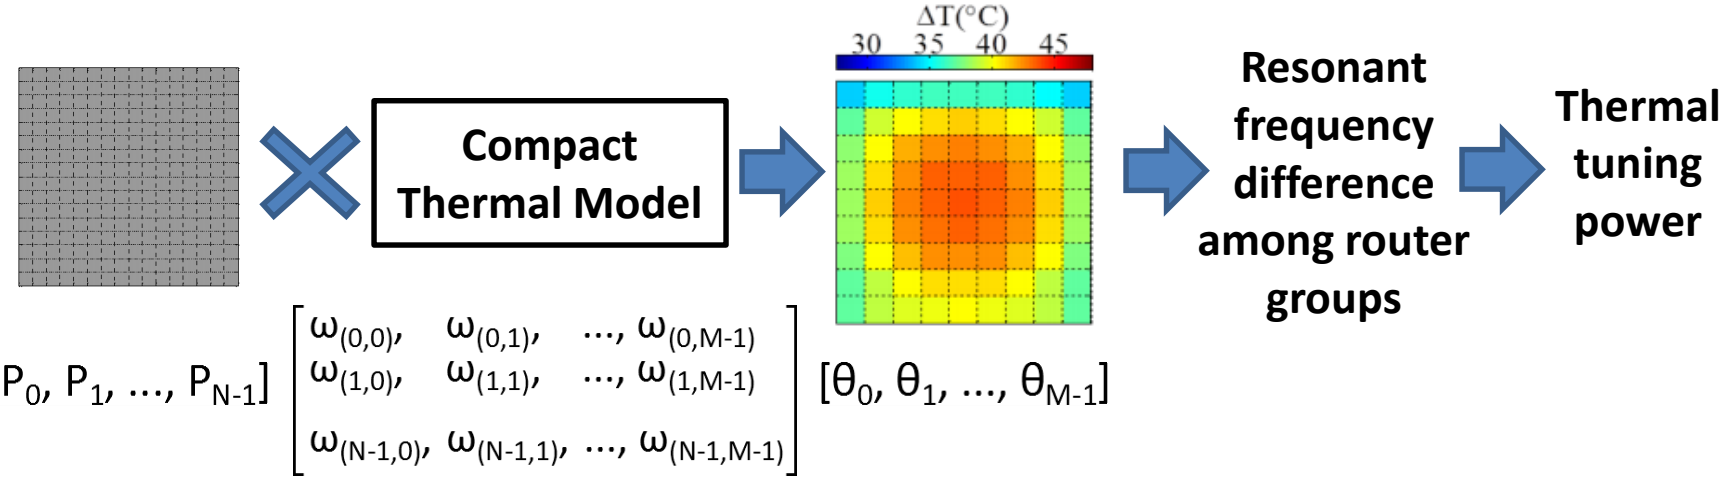

Size: 1×N N×M 1×M

Cross-layer PNoC P&R Optimization

Power Profiles

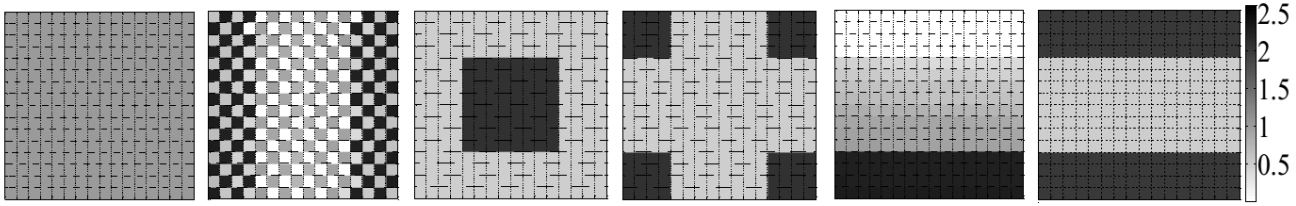

Thermal Conditions of Potential Ring Group Locations

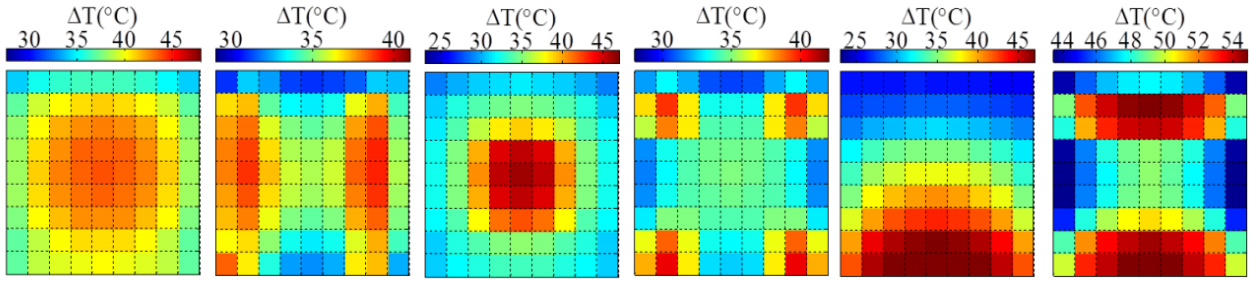

PNoC Layouts w/ Minimum PNoC Power

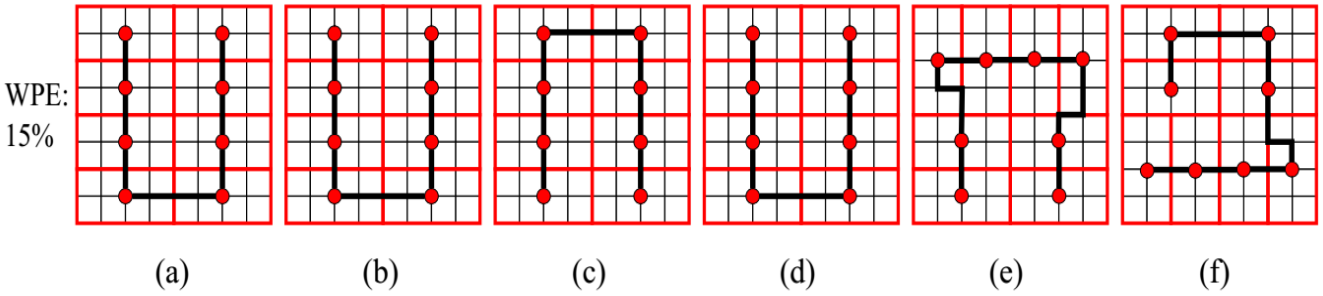

RingAware Workload Allocation Policy

- **Goals:**
  - Minimize the difference among ring temperatures
  - Reduce the overall chip temperature

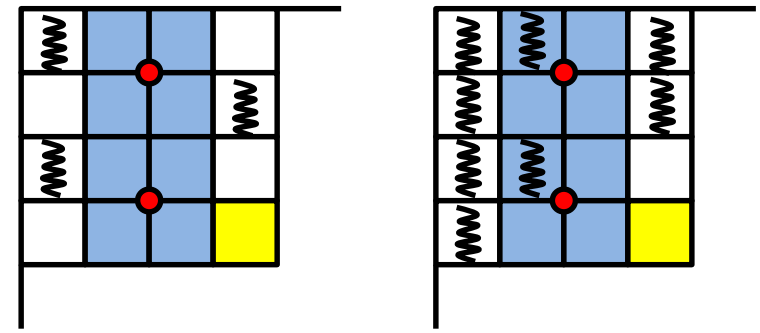

Center core RD0 cores

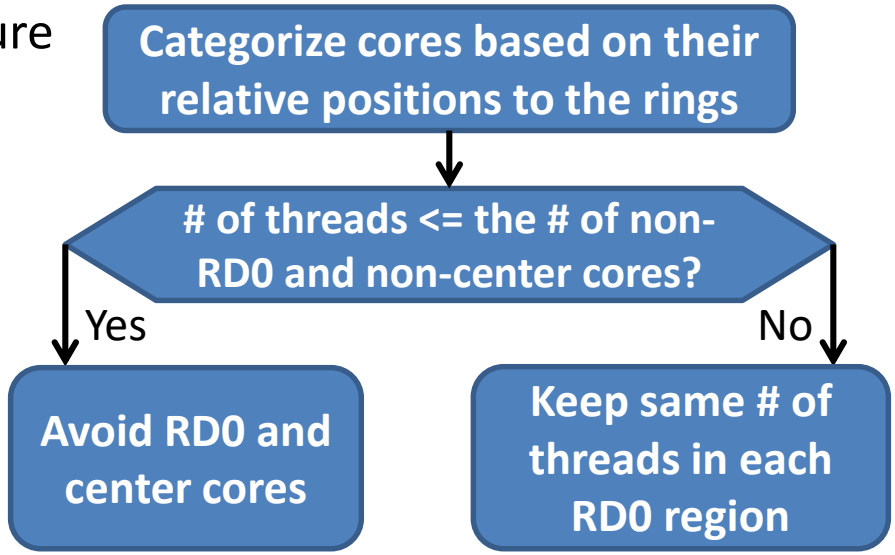

- ***Multi-program support***
  - Sort the threads based on their power dissipation & allocate high-power application first

FreqAlign Workload Allocation Policy

[TCAD'17]

- Process variation introduces resonant frequency shift after the system is manufactured
- Only balancing the temperature of ring groups is not enough to compensate the frequency mismatch
- On-chip laser sources' optical frequencies also need to match with corresponding rings' resonant frequency

FreqAlign + Adaptive Frequency Tuning

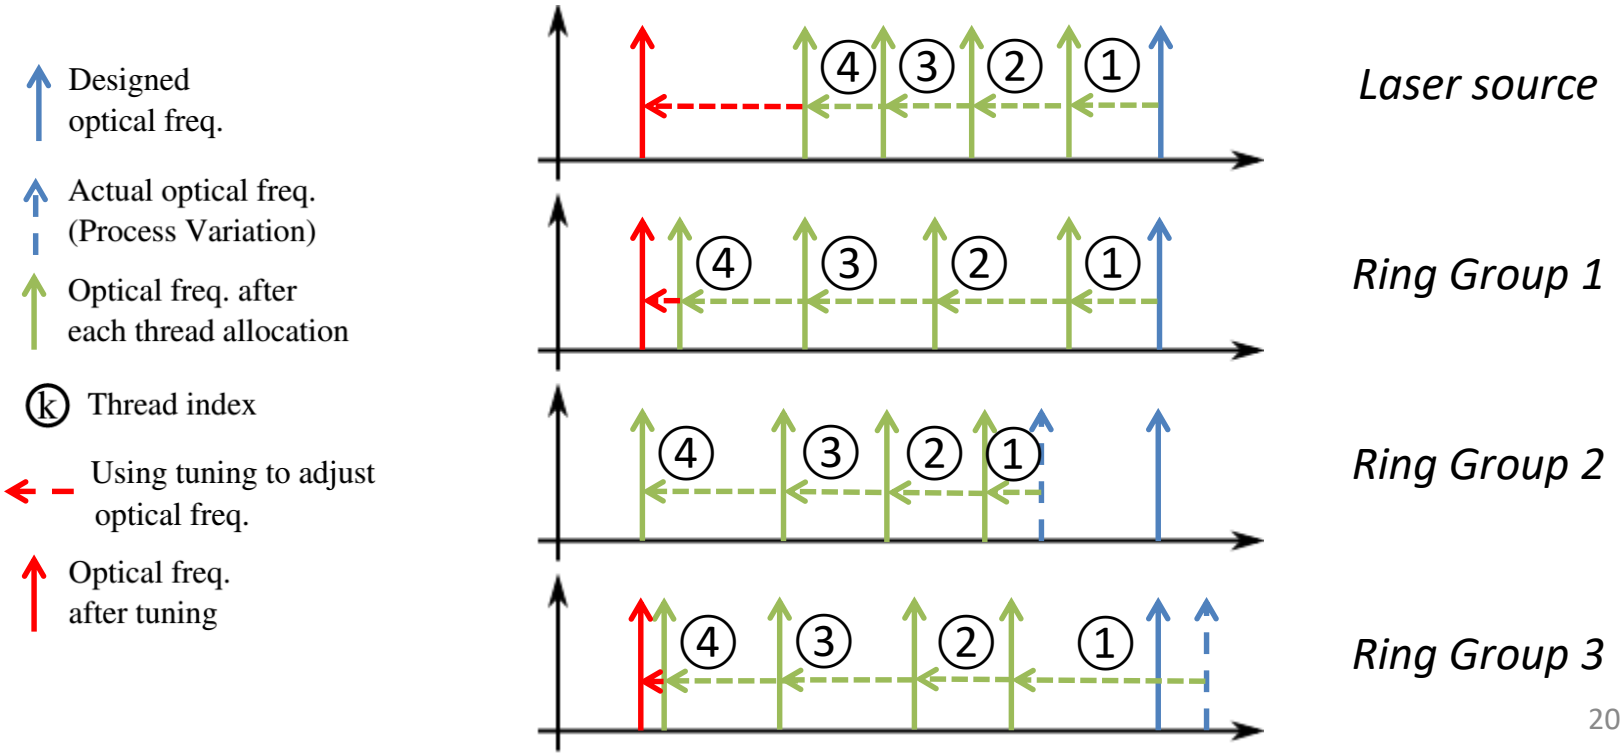

FreqAlign Workload Allocation Policy

[TCAD'17]

- **Target many-core system:**

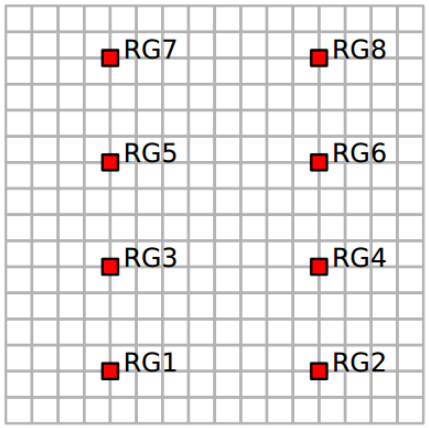

- **FreqAlign:**
  - Keep track of the optical frequency shifts of ring groups (in **RG weight array**)
  - Record every core's thermal impact on every ring group
  - Choose the core to minimize the frequency difference among all ring groups

- **Workflow:**

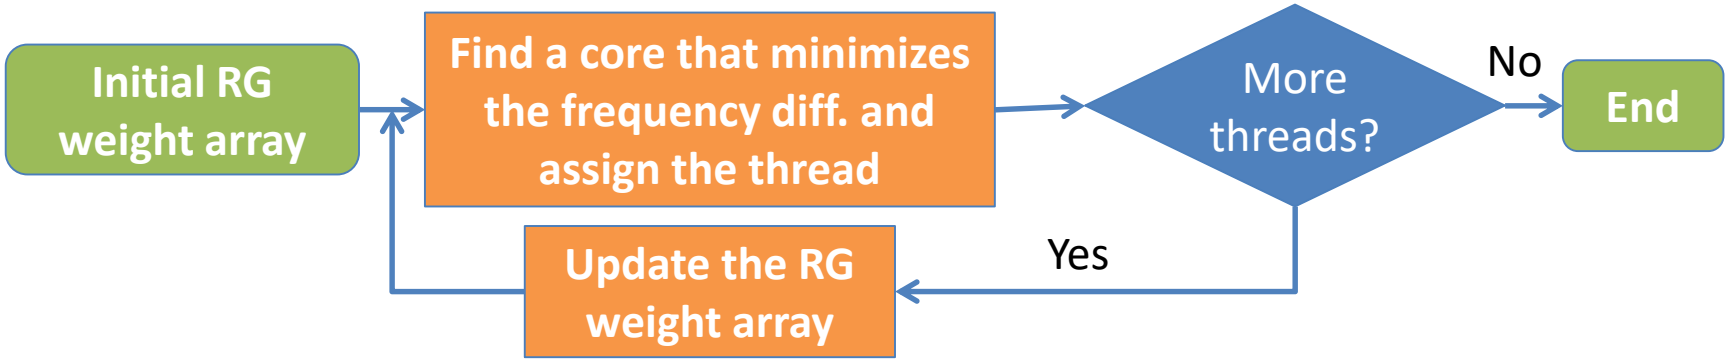

Experimental Methodology

- **Simulation Framework:**

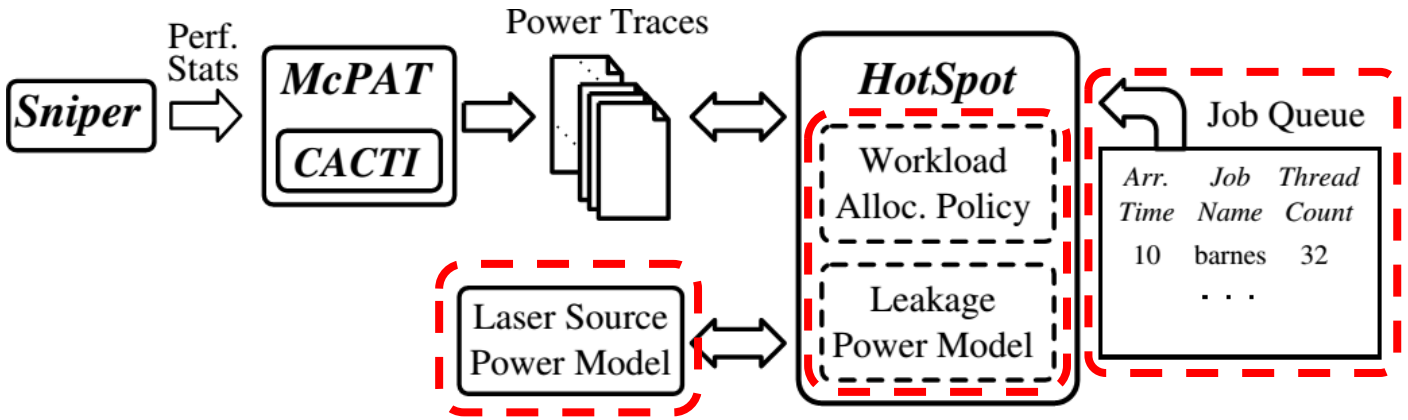

- **Workload Sets:** Selected benchmarks from SPLASH2, PARSEC and UHPC:

Workload Sets	Job 1	Job2
HP + HP	md	shock
HP + MP	md	blackscholes
HP + LP	shock	lu_cont
MP + MP	barnes	blackscholes
MP + LP	barnes	water_nsq
LP + LP	lu_cont	canneal

- **Tested Policies:**
  - Workload Allocation Policy: **Cluster, RingAware, FreqAlign**
  - Thermal Tuning Policy: **Target Frequency Tuning (TFT), Adaptive Frequency Tuning (AFT)**

Experimental Results for Many-core System w/o Process Variations

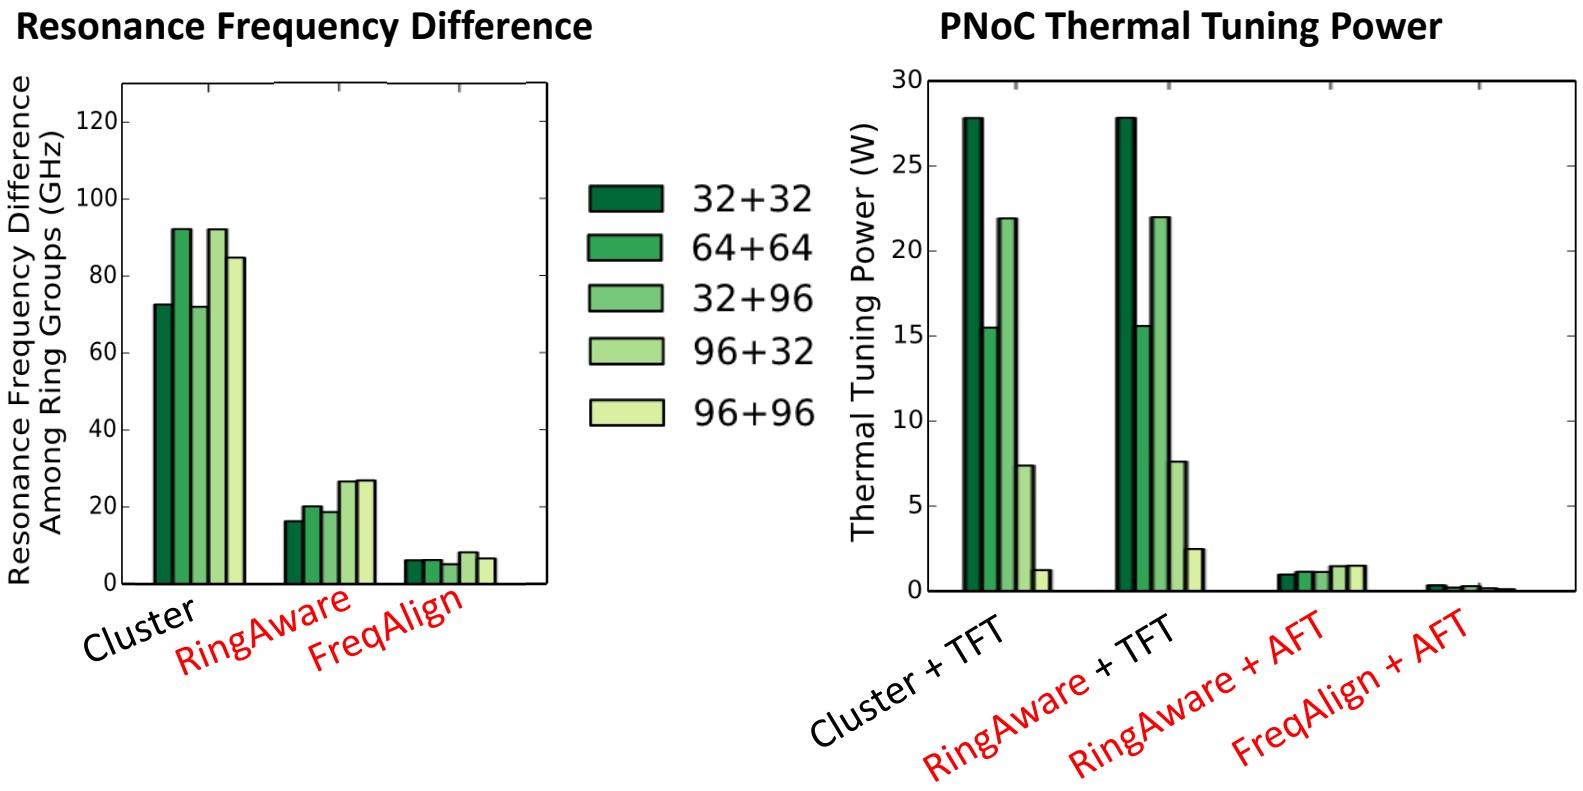

- Compared to *RingAware*, *FreqAlign* reduces the resonant frequency difference by **60.6%** on average;
- Compared to *RingAware + TFT*, *FreqAlign + AFT* reduces the tuning power by **14.93W** on average.

Laser Source Placement and Sharing Examples

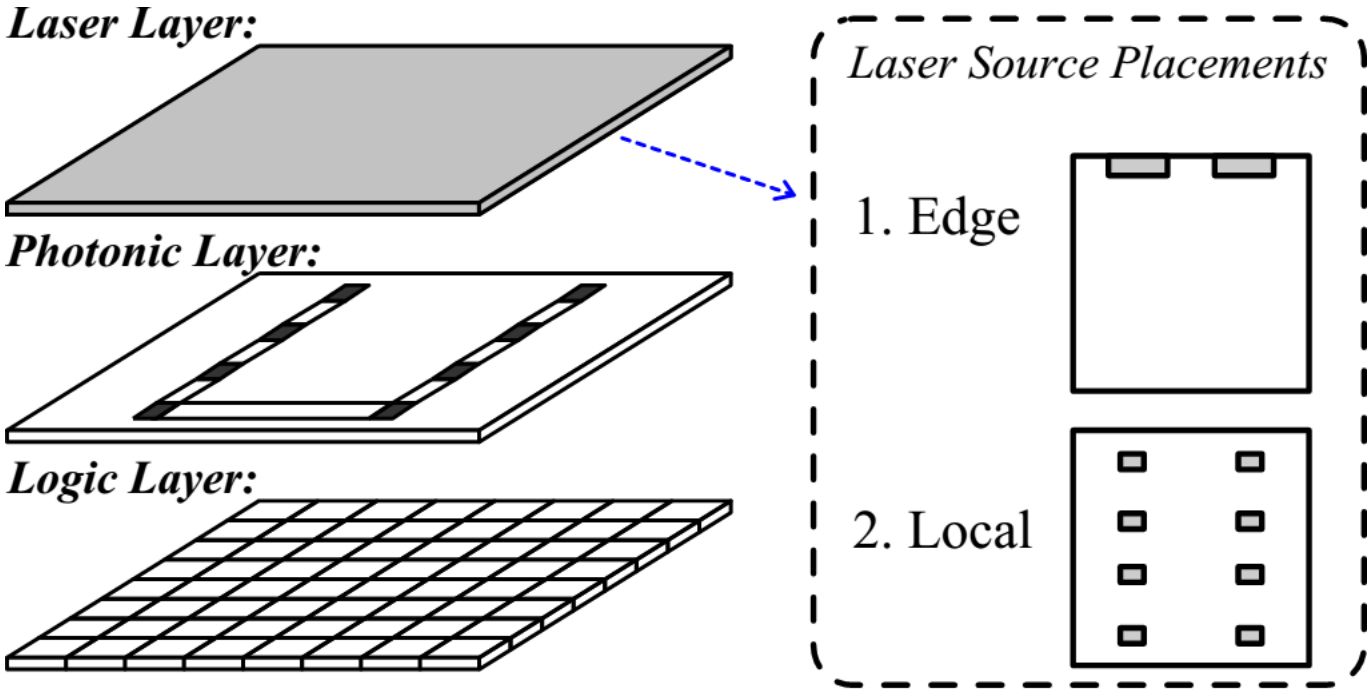

	Edge Placement	Local Placement
Sharing Degree	High	Low
Propagation loss	High	Low

Higher Sharing Degree
→ Higher η_{WPE}, Lower # of laser sources
→ Lower laser source power consumption

Higher Propagation loss
→ Higher required output optical power
→ Higher laser source power consumption

Take-Aways

- Cross-layer, thermally-aware optimizer for floorplanning of PNoCs
- Runtime workload allocation for thermal tuning power reduction
- ***Cross-layer simulation & optimization flow:*** an enabler to design energy-efficient systems with PNoCs

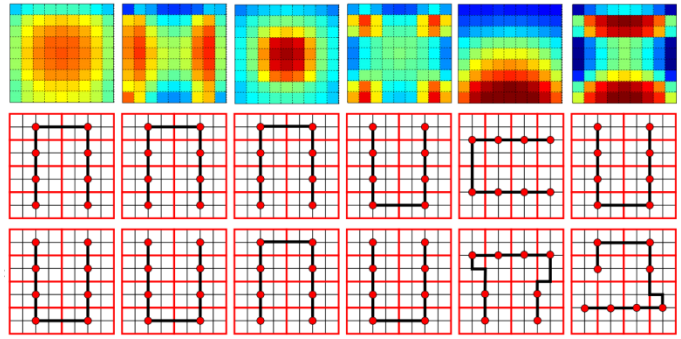

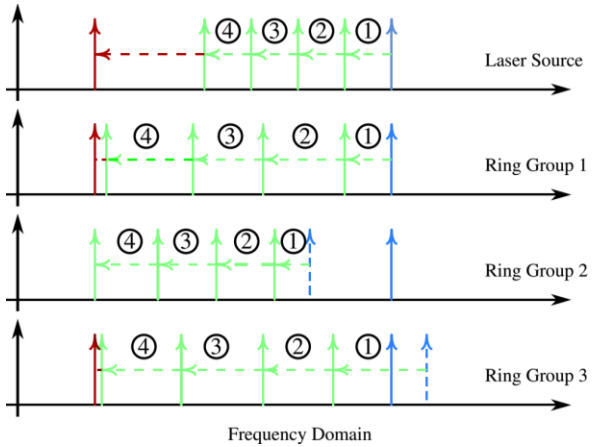

Performance and Energy Aware Computing Laboratory

http://www.bu.edu/peaclab

Current graduate students:

Ozan Tuncer, Onur Sahin, Emre Ates,

Yijia Zhang, Aditya Narayan, Prachi Shukla

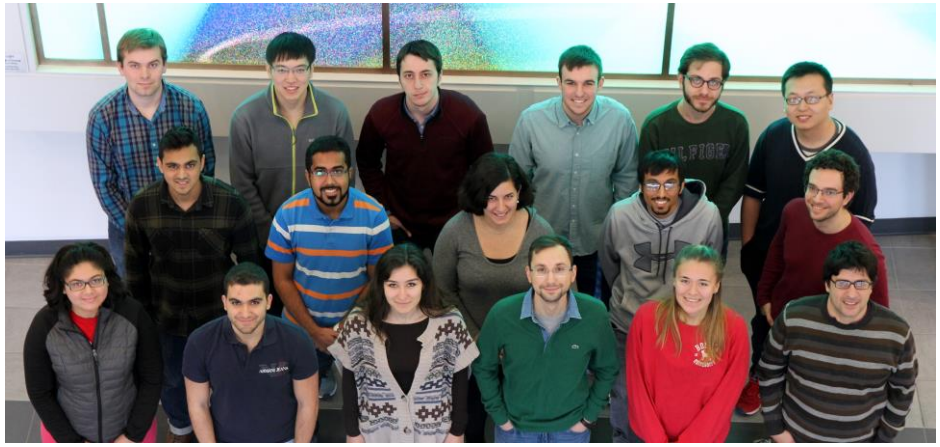

Research Scientist: Dr. Ata Turk

Alumni: Dr. Jie Meng, Dr. Can Hankendi, Dr. Hao Chen, Dr. Fulya Kaplan, Dr. Tiansheng Zhang,

Cyril Saade, John Knollmeyer, Nathaniel Michener, Ann Lane, Katsu Kawakami, John Furst, Samuel Howes, Jon Bell, Benjamin Havey, Ryan Mullen

Collaborators:

D. Atienza & Y. Leblebici @ EPFL,

J. Ayala @ UCM, J. M. Moya @ UPM,

C. Isci, S. Duri @ IBM TJ Watson,

T. Brunschwiler @IBM Zurich,

L. Benini @ ETHZ/U. of Bologna,

M. Caramanis and A. Joshi @ BU,

K. Gross & K. Vaidyanathan @ Oracle,

J. Klamkin @ UCSB,

V. Leung, and A. Rodrigues @ Sandia Labs,

S. Reda @ Brown University,

A. Kahng and D. Tullsen @ UCSD.

Many masters students, especially:

Dan Rossell, Charlie De Vivero

Visitors:

Dr. Marina Zapater, Dr. Andrea Bartolini

Recent Funding

IBM Sandia National Laboratories ORACLE NSF AMD