Keio University

1858

CALAMVS GLADIO FORTIOR

Interconnect Enhances Architecture:
Evolution of Wireless NoC from Planar to 3D

3D WiNoC Architectures

Hiroki Matsutani
Keio University, Japan

Outline: Wireless 3D NoC architectures

- Wireless 3D NoCs (3D WiNoCs) [8min]
  - Wireless 3D IC technology [Matsutani,NOCS'11]
  - 3D WiNoC design example (65nm) [Miura,IEEE Micro 13]
- Routing & topology exploration [8min]
  - Spanning tree optimization for irregular WiNoCs
  - Power management via vertical ON/OFF links [Matsutani,ASPDAC'13]
- Adding random NoC chip for 3D WiNoCs [5min]
  - Adding randomness induces small-world effects
  - Adding random NoC chip to NoC-less 3D ICs [Matsutani,DATE'14]
- Summary [1min]

Design cost of LSI is increasing

- System-on-Chip (SoC)
  - Required components are integrated on a single chip
  - Different LSI must be developed for each application
- System-in-Package (SiP) or 3D IC
  - Required components are stacked for each application

By changing the chips in a package, we can provide a wider range of chip family with modest design cost

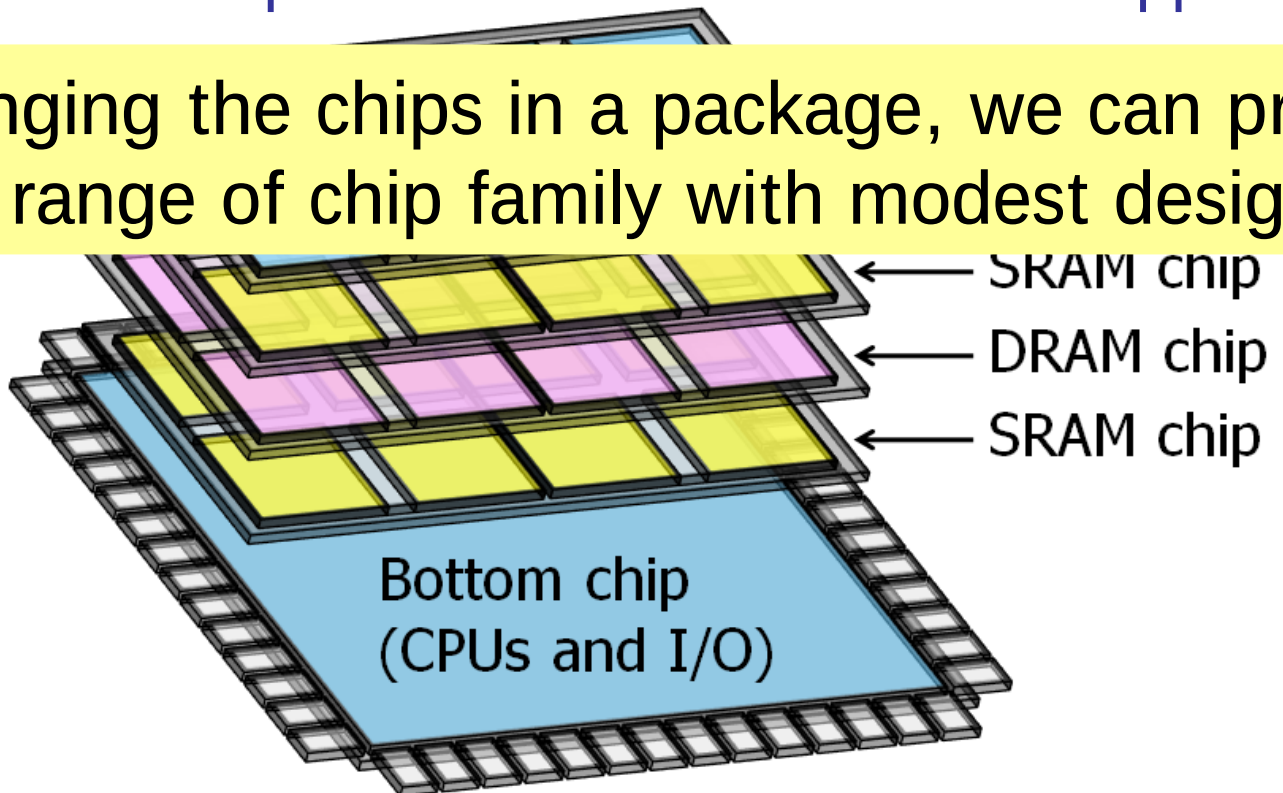

3D IC technology for going vertical

	Wired	Wireless
Two chips (face-to-face)	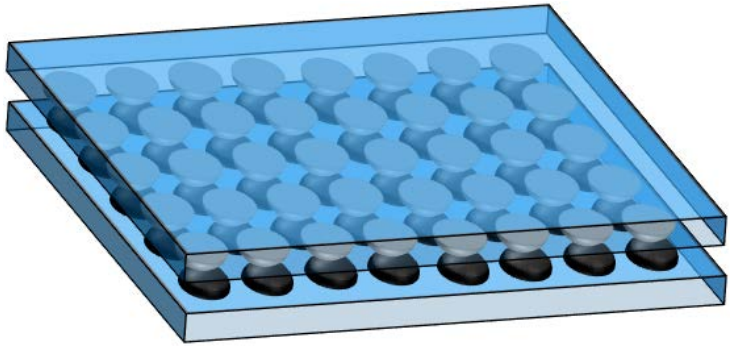 Microbump	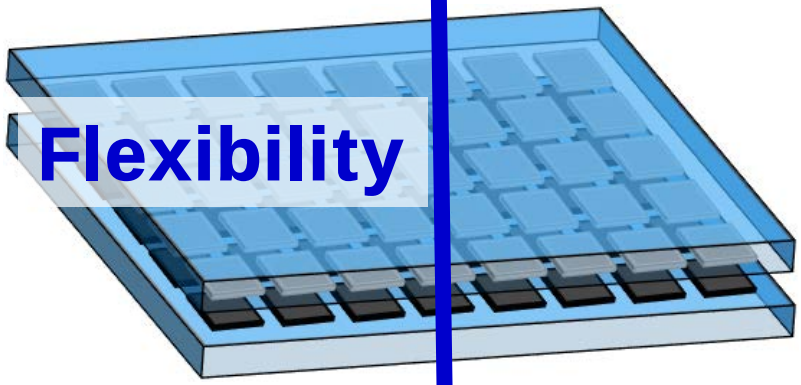 Flexibility Capacitive coupling
More than three chips	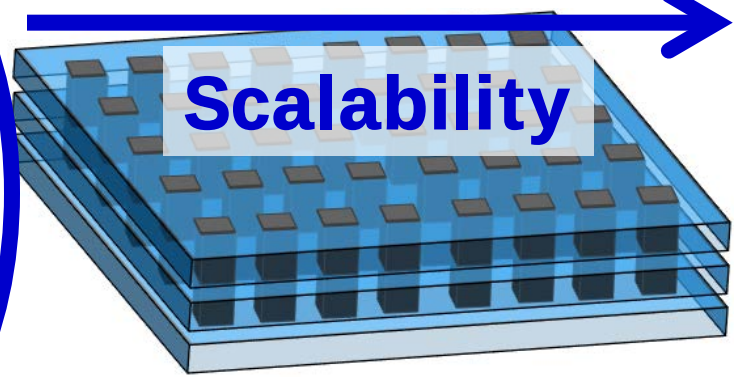 Scalability Through silicon via	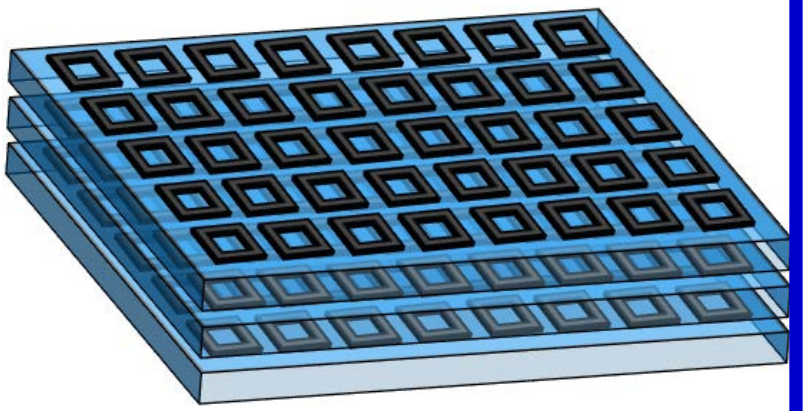 Inductive coupling

Wireless 3D NoC (3D WiNoC)

Three chips are stacked wirelessly

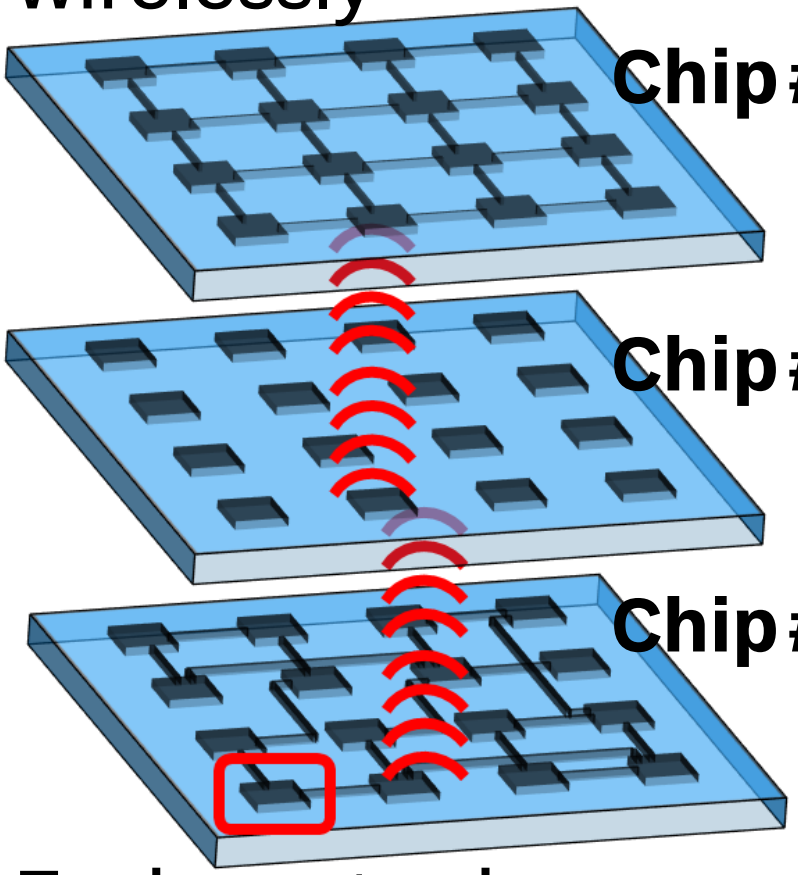

Each router has a wireless transceiver

Each router has TX/RX coils (inductors)

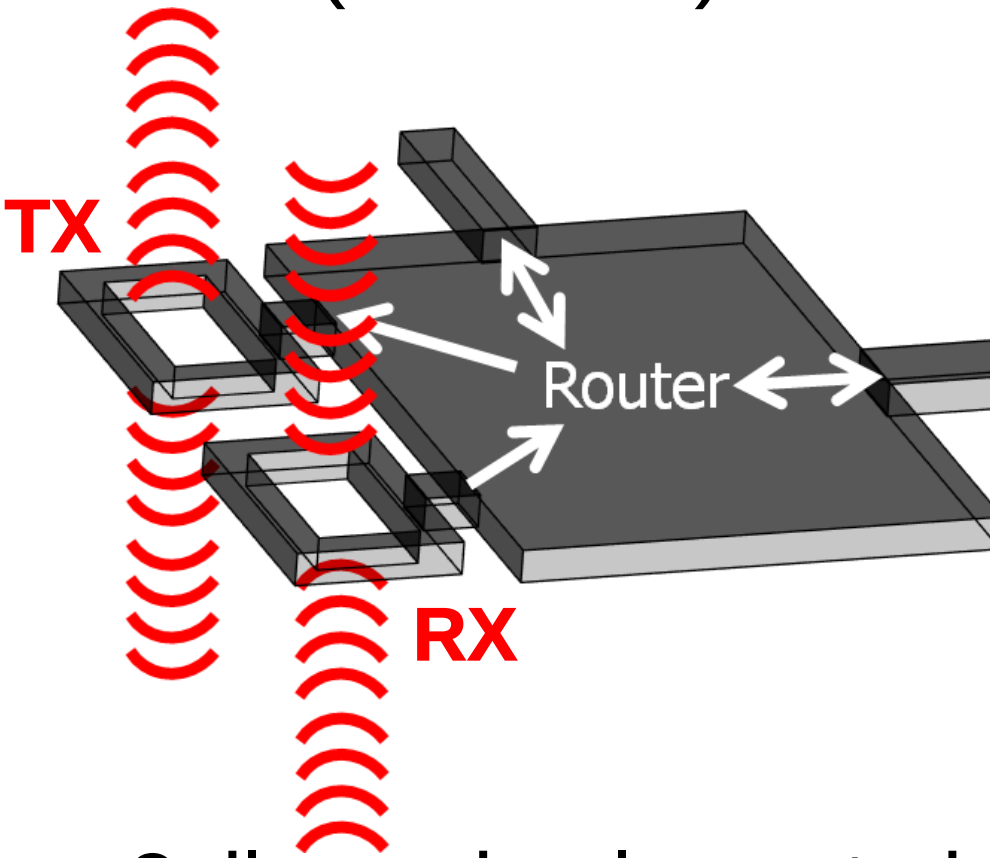

Coils are implemented with metal layers

An example: Cube-1 (2012)

- Test chips for building-block 3D systems
  - Two chip types: Host CPU chip & Accelerator chip
  - We can customize number & types of chips in SiP

[Miura, IEEE Micro 13]

- Cube-1 Host CPU chip
  - Two 3D wireless routers
  - MIPS-like CPU

- Cube-1 Accelerator chip
  - Two 3D wireless routers
  - Processing element array

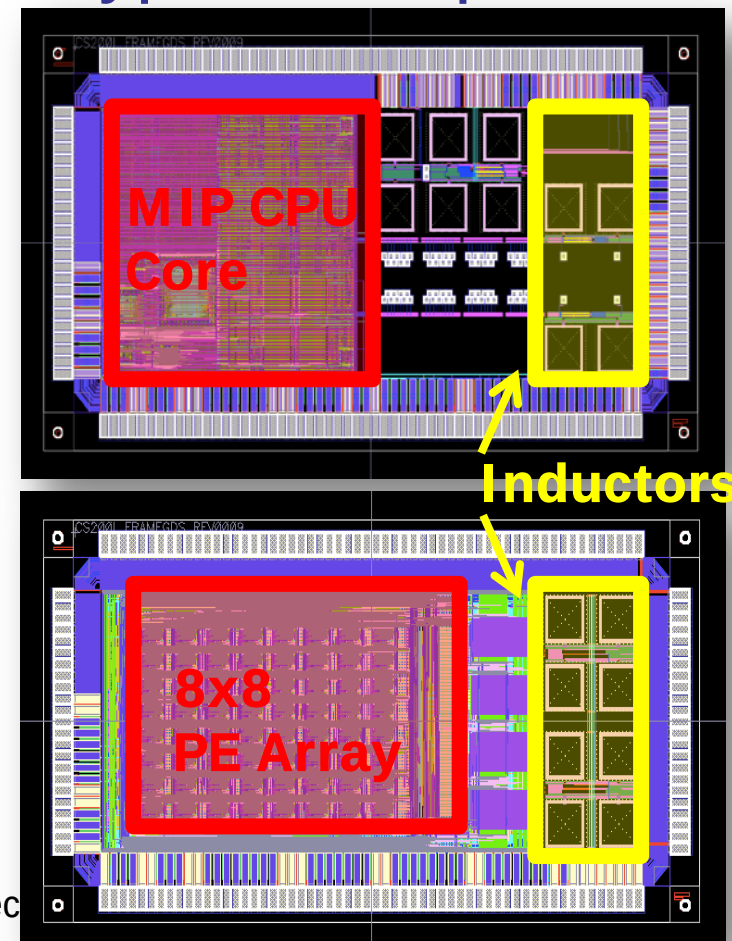

We can change the number and types of chips in a package according to application requirements

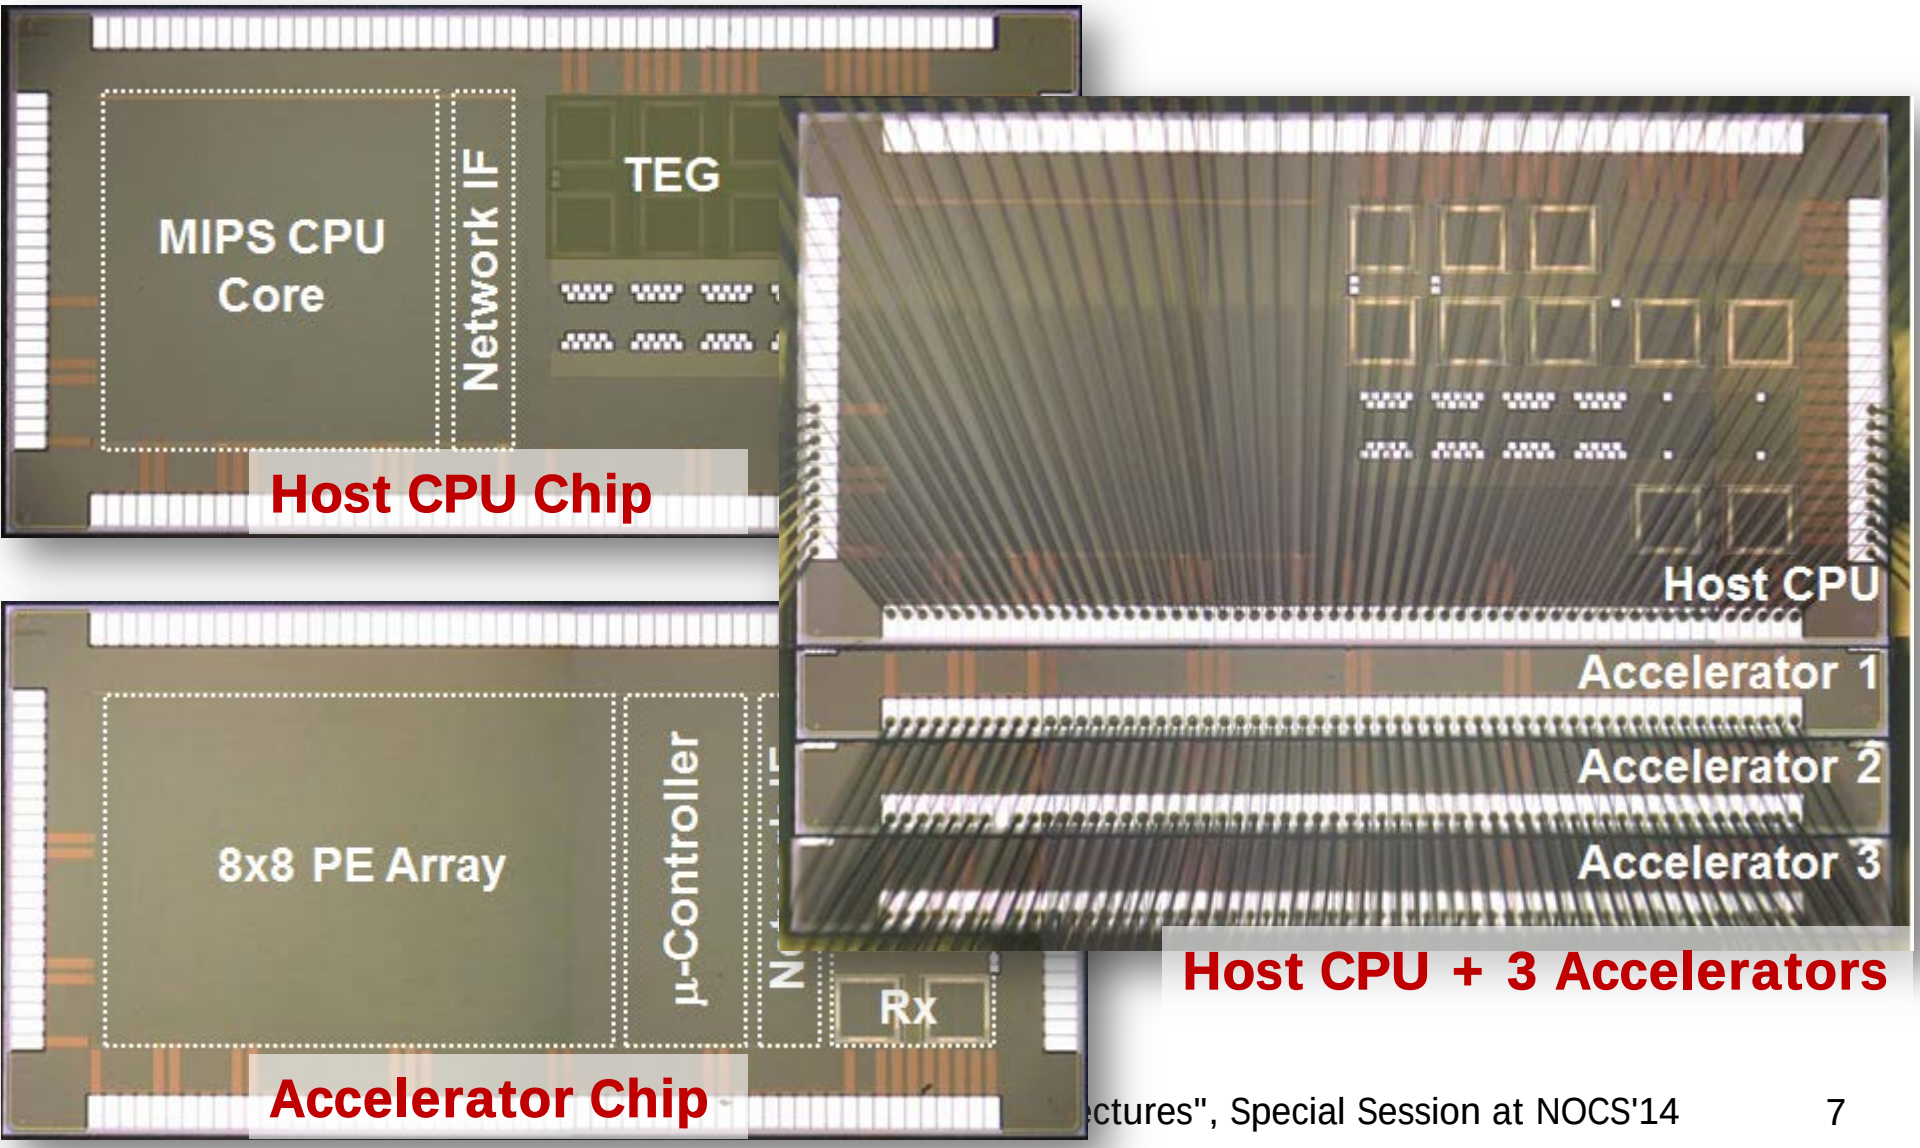

Either vertical P2P links or broadcast bus can be formed for 3D WiNoC

- T
  - Vertical point-to-point link between adjacent chips
  - Vertical shared bus (broadcast)

[Matsutani, NOCS'11]

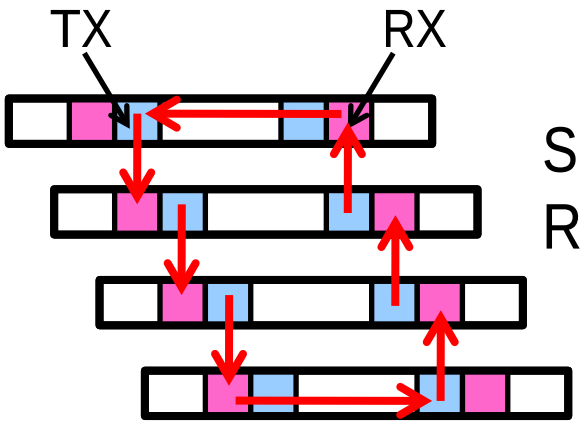

Stacking for Ring network

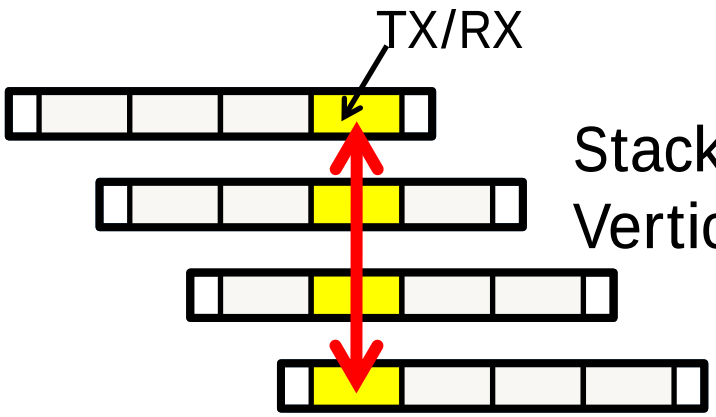

Stacking for Vertical bus

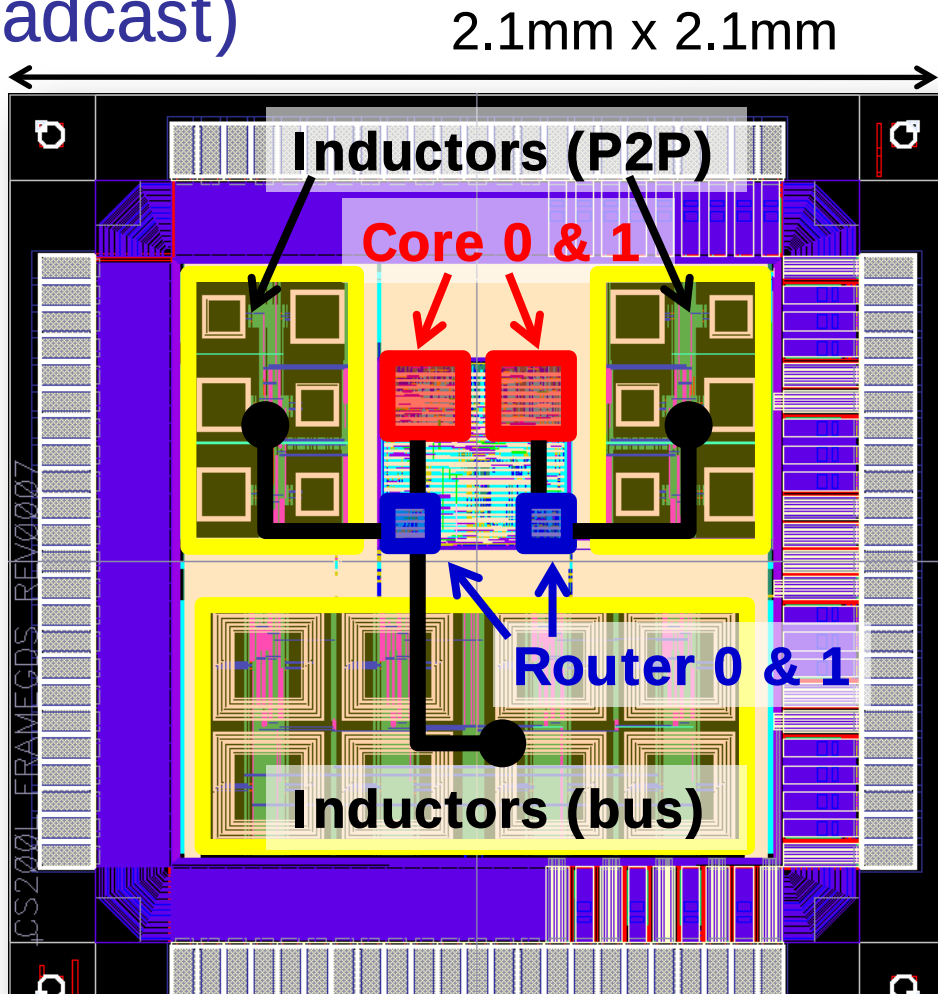

Outline: Wireless 3D NoC architectures

- Wireless 3D NoCs (3D WiNoCs) [8min]
  - Wireless 3D IC technology
  - 3D WiNoC design example (65nm)

[Matsutani,NOCS'11]
[Miura,IEEE Micro 13]

- Routing & topology exploration [8min]
  - Spanning tree optimization for irregular WiNoCs
  - Power management via vertical ON/OFF links

[Matsutani,ASPDAC'13]

- Adding random NoC chip for 3D WiNoCs [5min]
  - Adding randomness induces small-world effects
  - Adding random NoC chip to NoC-less 3D ICs

[Matsutani,DATE'14]

- Summary [1min]

Big picture: Wireless 3D NoC

- Arbitrary chips are stacked after fabrication
  - Each chip has vertical links at pre-specified locations, but we do not know internal topology of each chip
  - Some chips may not have horizontal NoC (vertical link only)

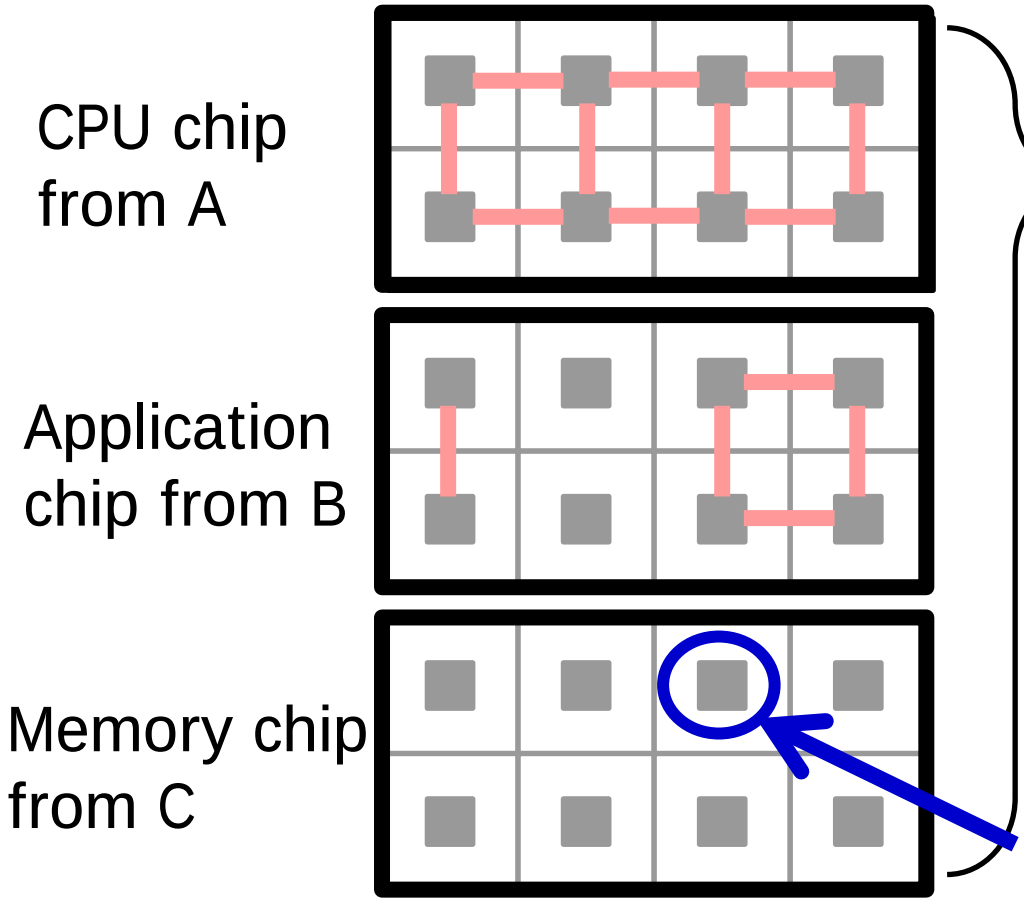

Required chips are stacked for each application

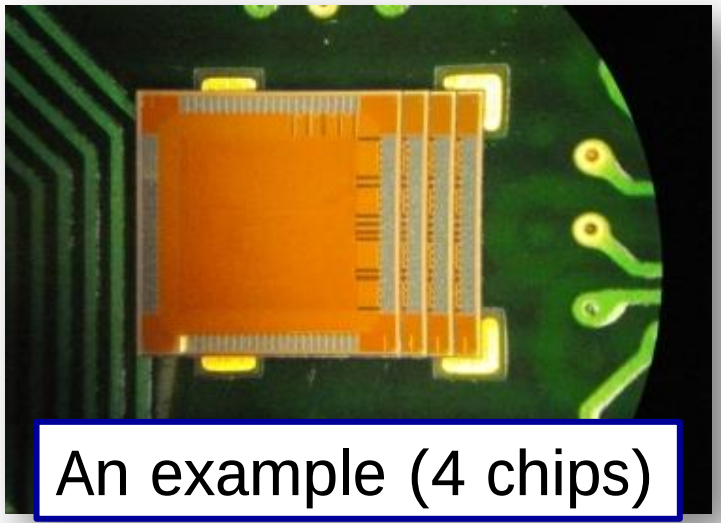

Inductor (vertical link)

Wireless 3D NoC: Up*/down* routing

- Up*/down* (UD) routing [Schroeder, JSAC'91]
  - Irregular network routing
  - A root node is selected
  - Direction (up or down) is assigned
  - Packets go up and then go down
- Example: 4 chips

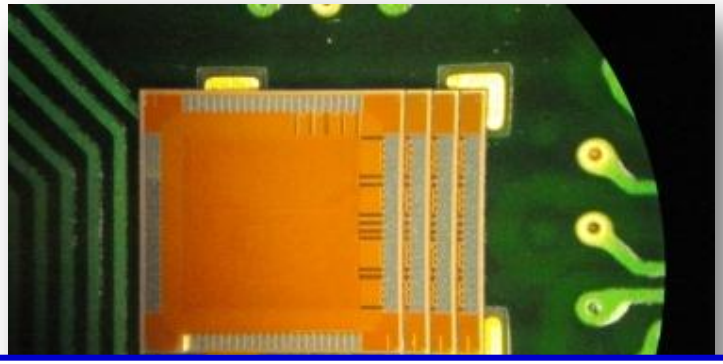

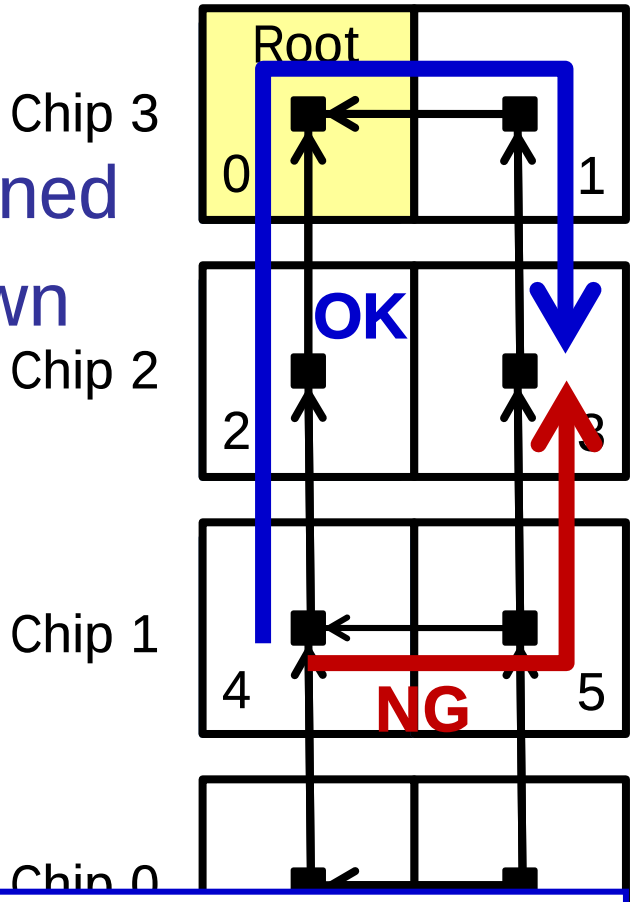

The best spanning tree root is selected by exhaustive or heuristic using communication traces (9sec for 64-tile)

Wireless 3D NoC: UD routing w/ VCs

- UD routing with multiple VCs
  - Each layer (VC) has its own spanning tree
  - Packets can transit multiple layers in descent order

[Koibuchi,ICPP'03]
[Lysne,TPDS'06]

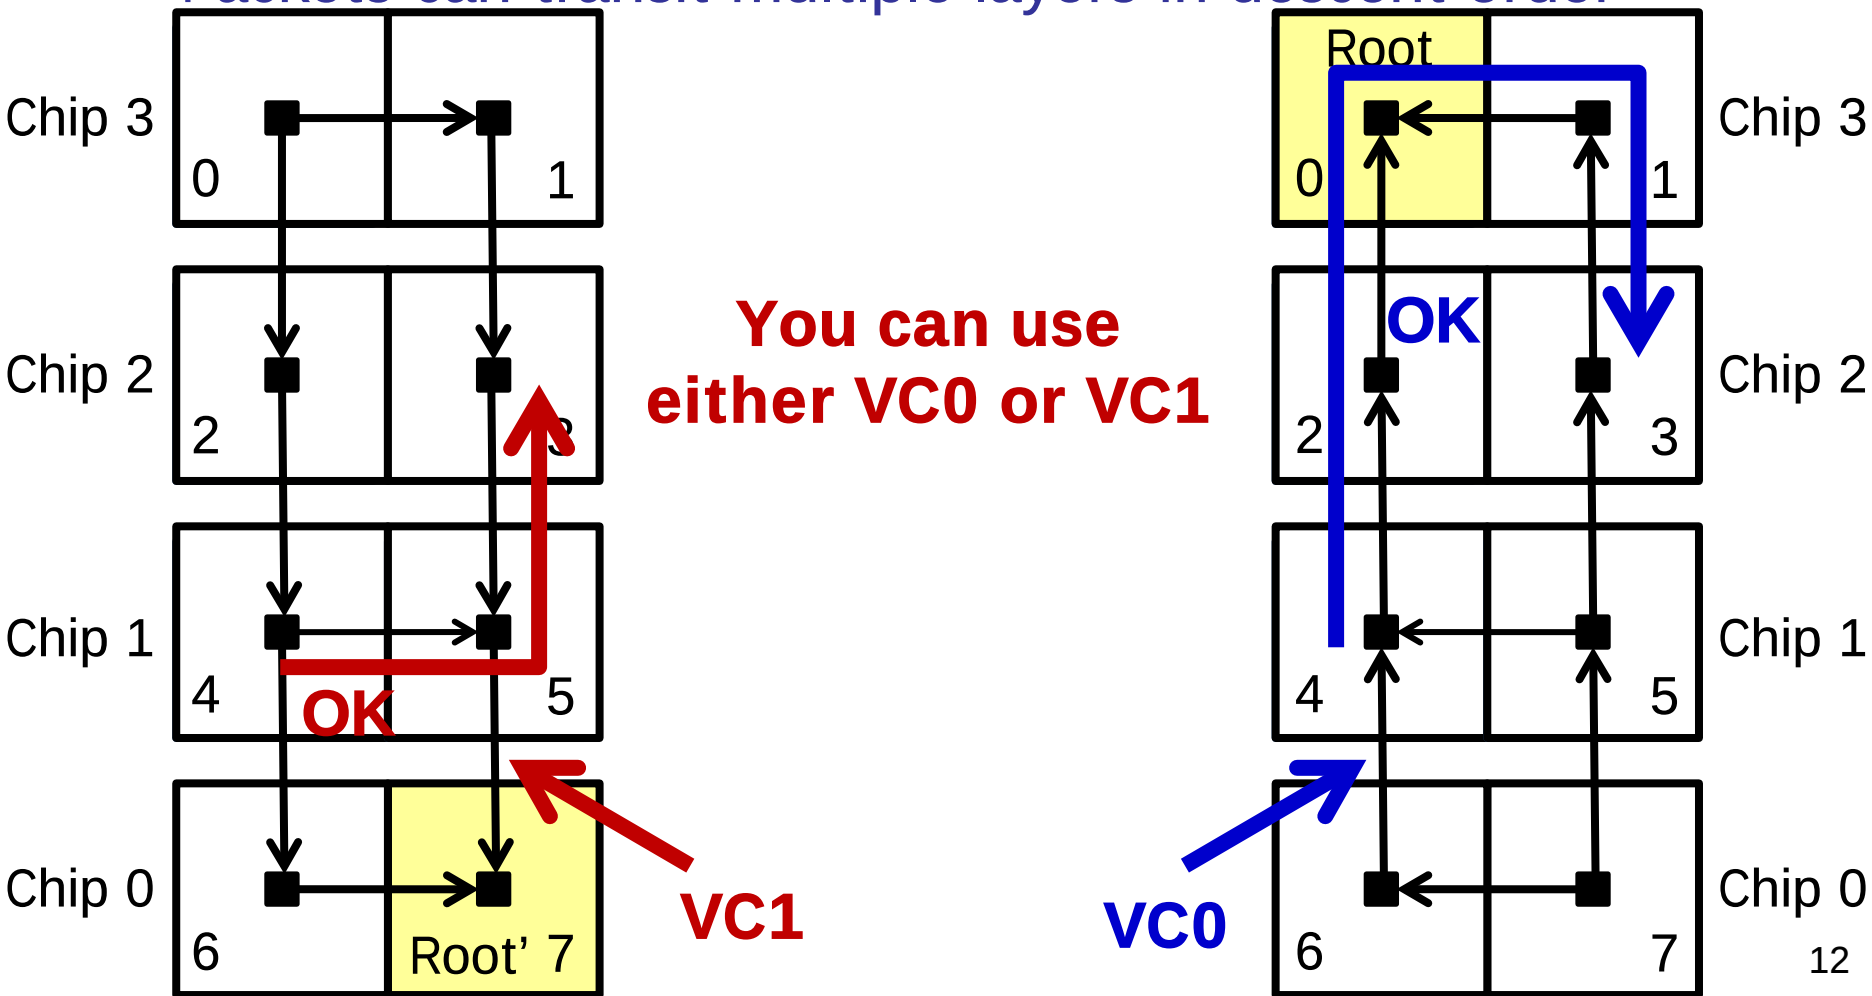

Full-system simulations: Topology

Among 1,000 random topologies, one with the most typical hop count value is selected for the full-system evaluation

- Each tile has router and core (e.g., processor or caches)
- Each horizontal link appears with 50% **(H50)**
- Each vertical link appears with 100% **(V100)**

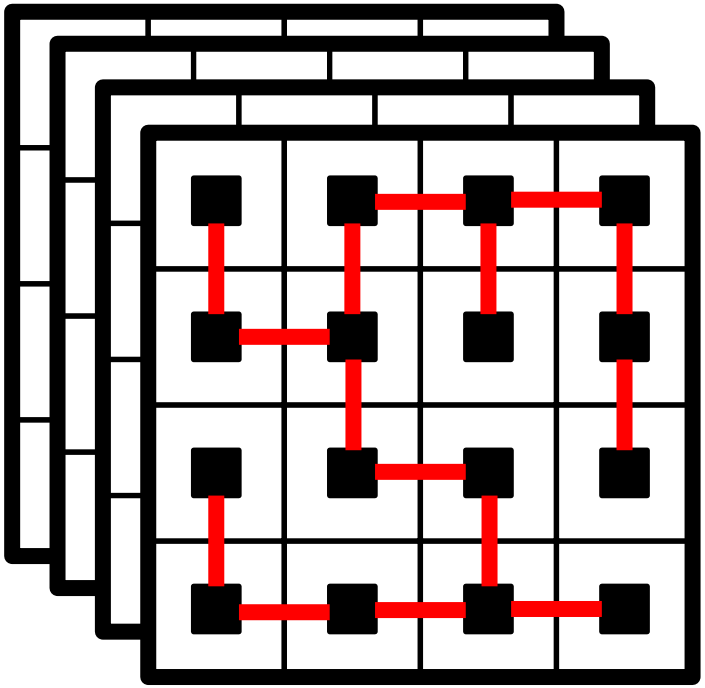

Hiroki Matsutani, "3D V

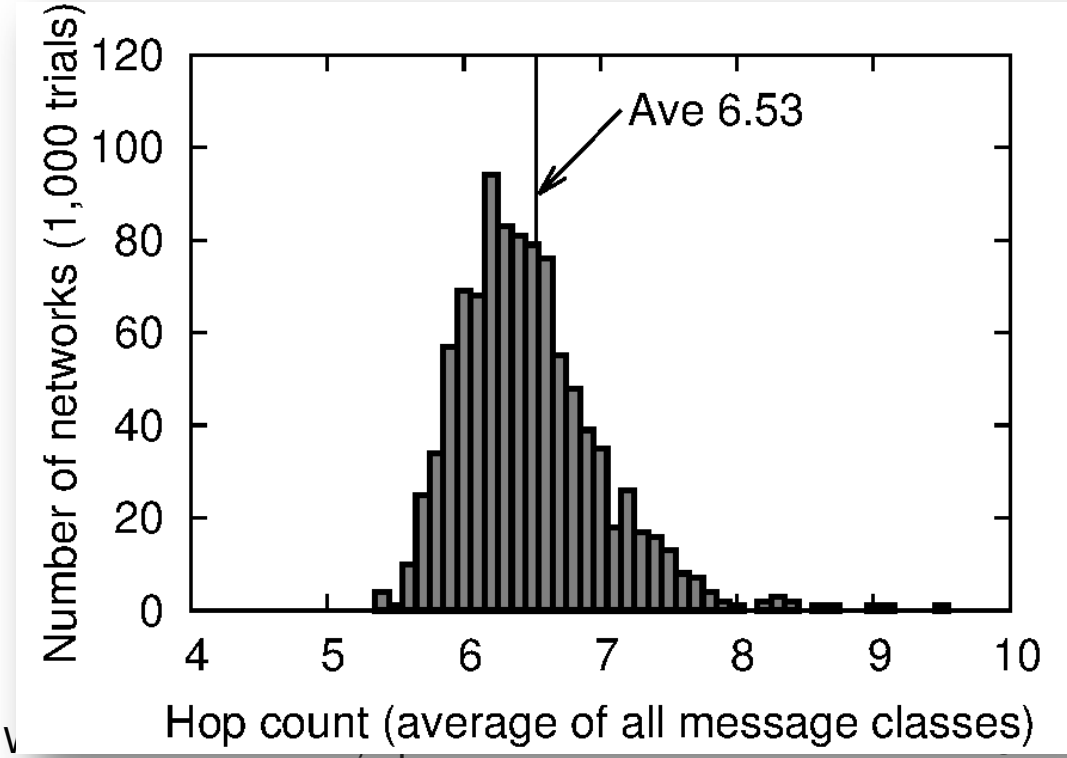

Full-system simulations: Gem5

- **H50-V100** topology (most typical one among 1,000)
  - Each horizontal link appears with 50% **(H50)**
  - Each vertical link appears with 100% **(V100)**

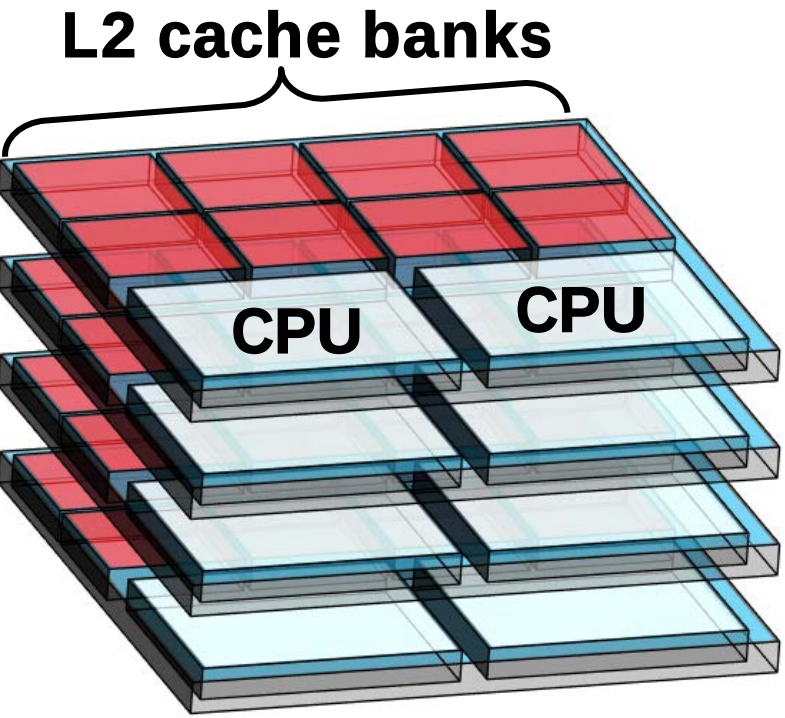

8 CPUs, 32 L2$ banks, and 4 memory controllers

Table 1: Simulation parameters

Processor architecture	X86-64
L1$ size & latency	32K / 1cycle
L2$ size & latency	256K / 6cycle
Memory size & latency	4G / 160cycle
Router latency	[BW] [VSA] [ST] [LT]
Router buffer size	5-flit per VC
Protocol	MOESI directory (3VC)

Table 2: Application programs

NPB (BT, CG, EP, FT, IS, LU, MG, SP, UA)

Full-system simulations: Gem5

- **H50-V100** topology (most typical one among 1,000)
  - Each horizontal link appears with 50% **(H50)**
  - Each vertical link appears with 100% **(V100)**
- The best spanning tree root is selected for each message class (i.e., 3 roots for 3 message classes)

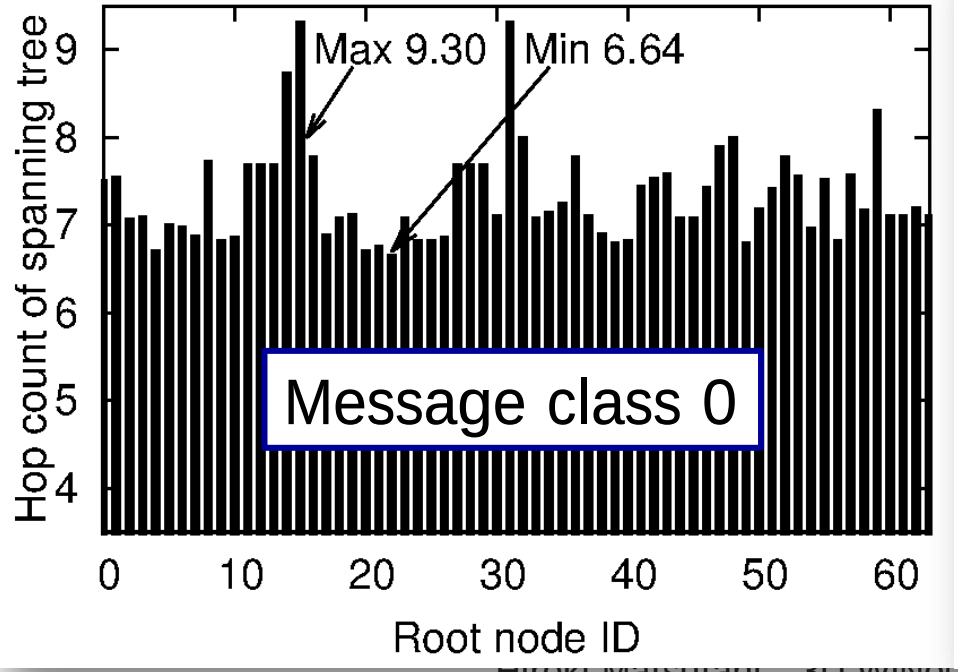

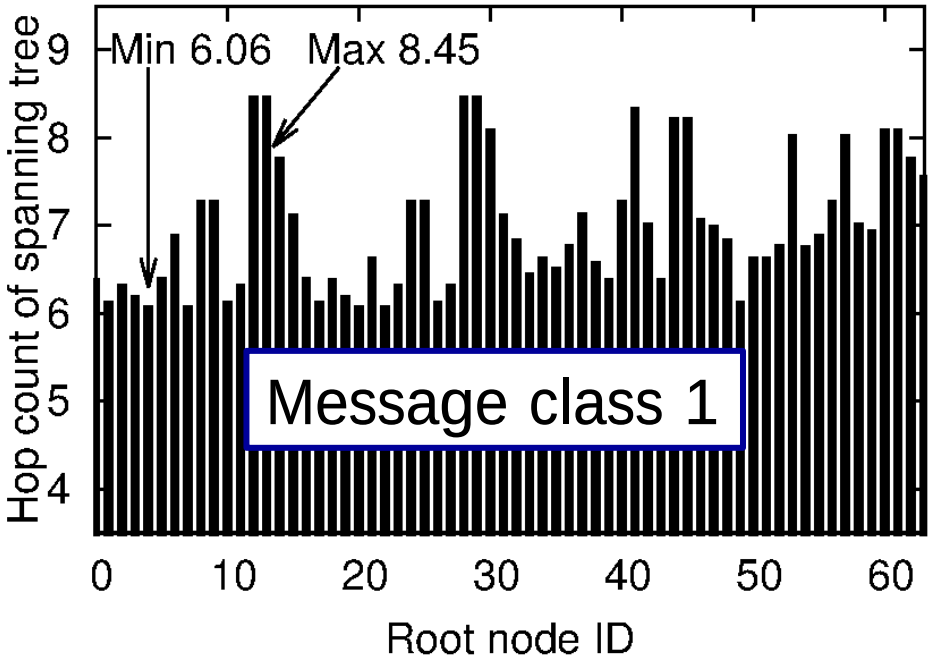

Hop count: H50-V100 topology

- Hop count w/ **the worst case** spanning tree root
- Hop count w/ **the best case** spanning tree root
- Hop count w/ two spanning tree roots (2 VCs)
- Ideal hop count (not deadlock-free)

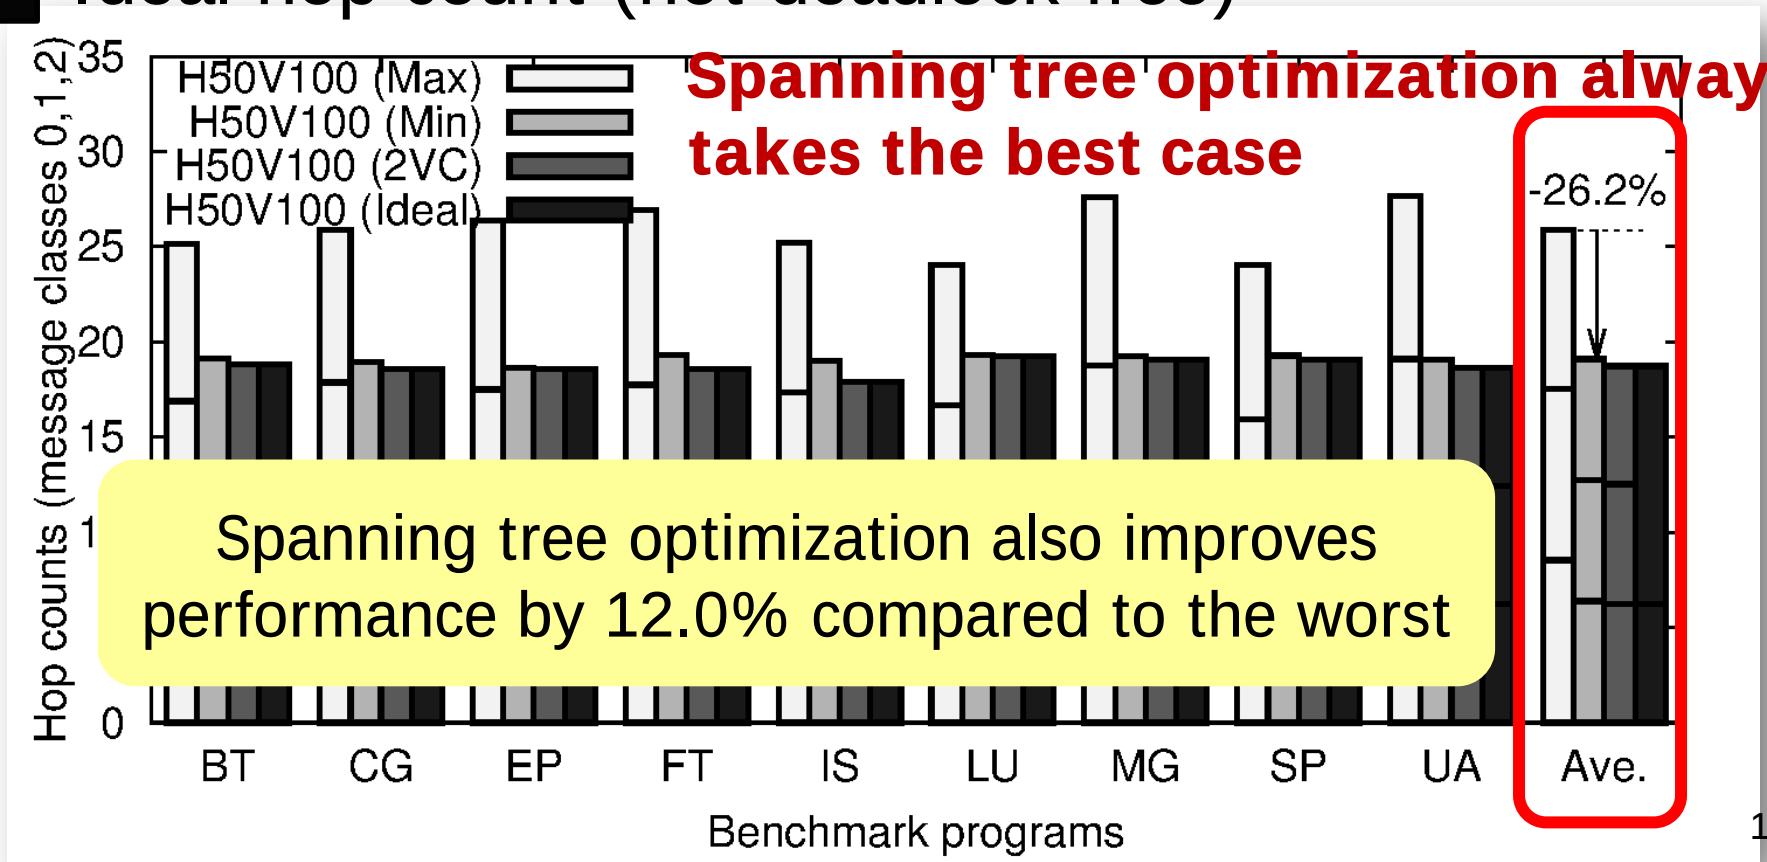

Spanning tree optimization also improves performance by 12.0% compared to the worst

App exec time: H50-V100 topology

H50-V100 w/ **the worst case**

H50-V100 w/ **the best case**

(50% horizontal links)

H100-V100 (all horizontal links available = 3D Mesh)

Spanning tree optimization always takes the best case

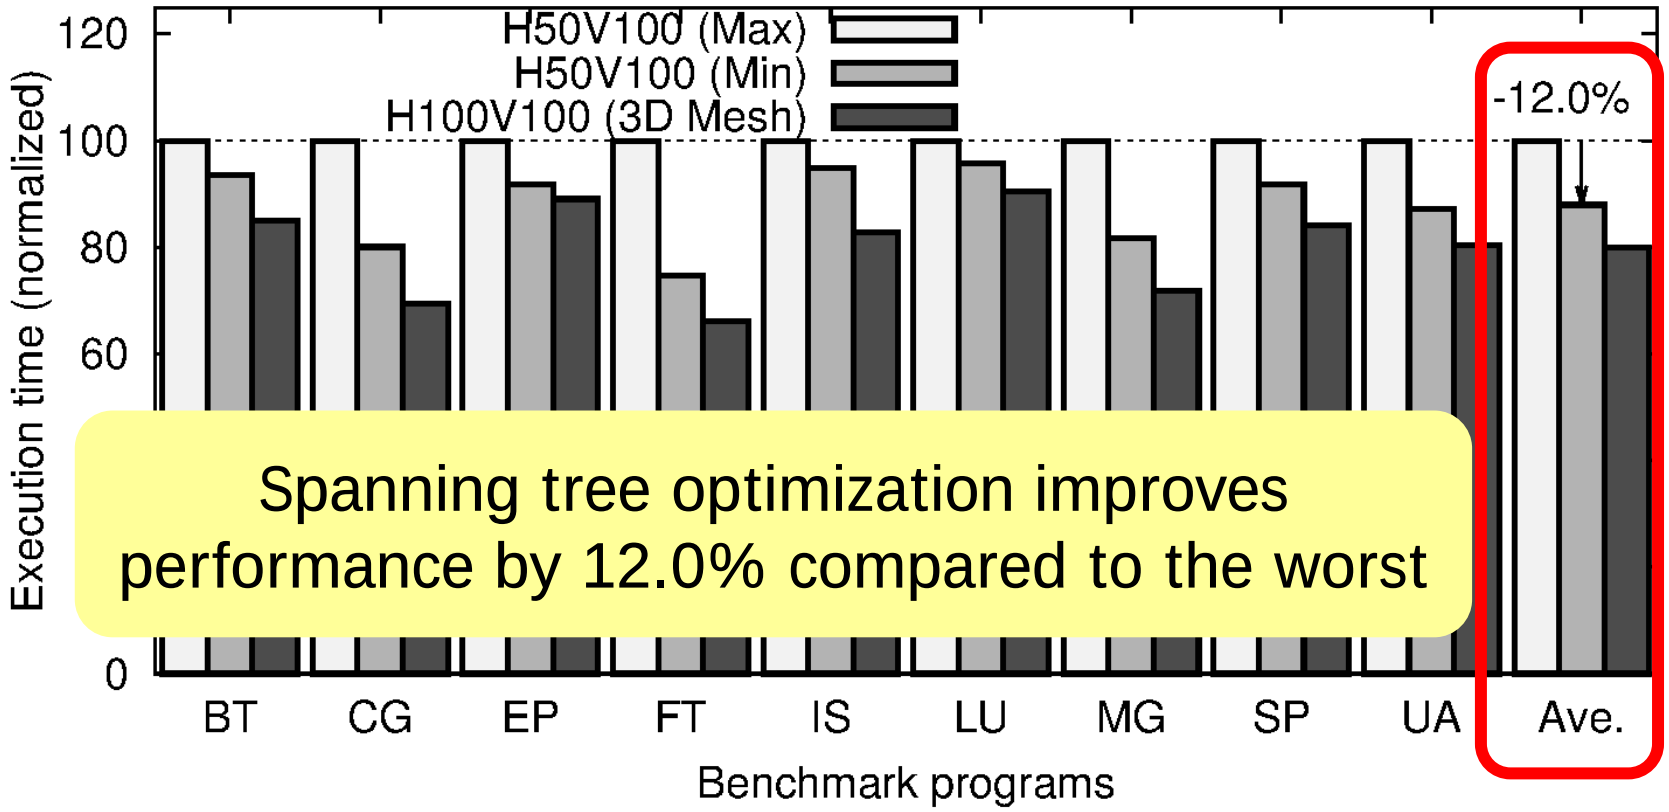

Spanning tree optimization improves performance by 12.0% compared to the worst

Energy per flit: H50-V100 topology

H50-V100 w/ **the worst case**

H50-V100 w/ **the best case**

(50% horizontal links)

H100-V100 (all horizontal links available = 3D Mesh)

Energy for routers, horizontal links, and vertical links

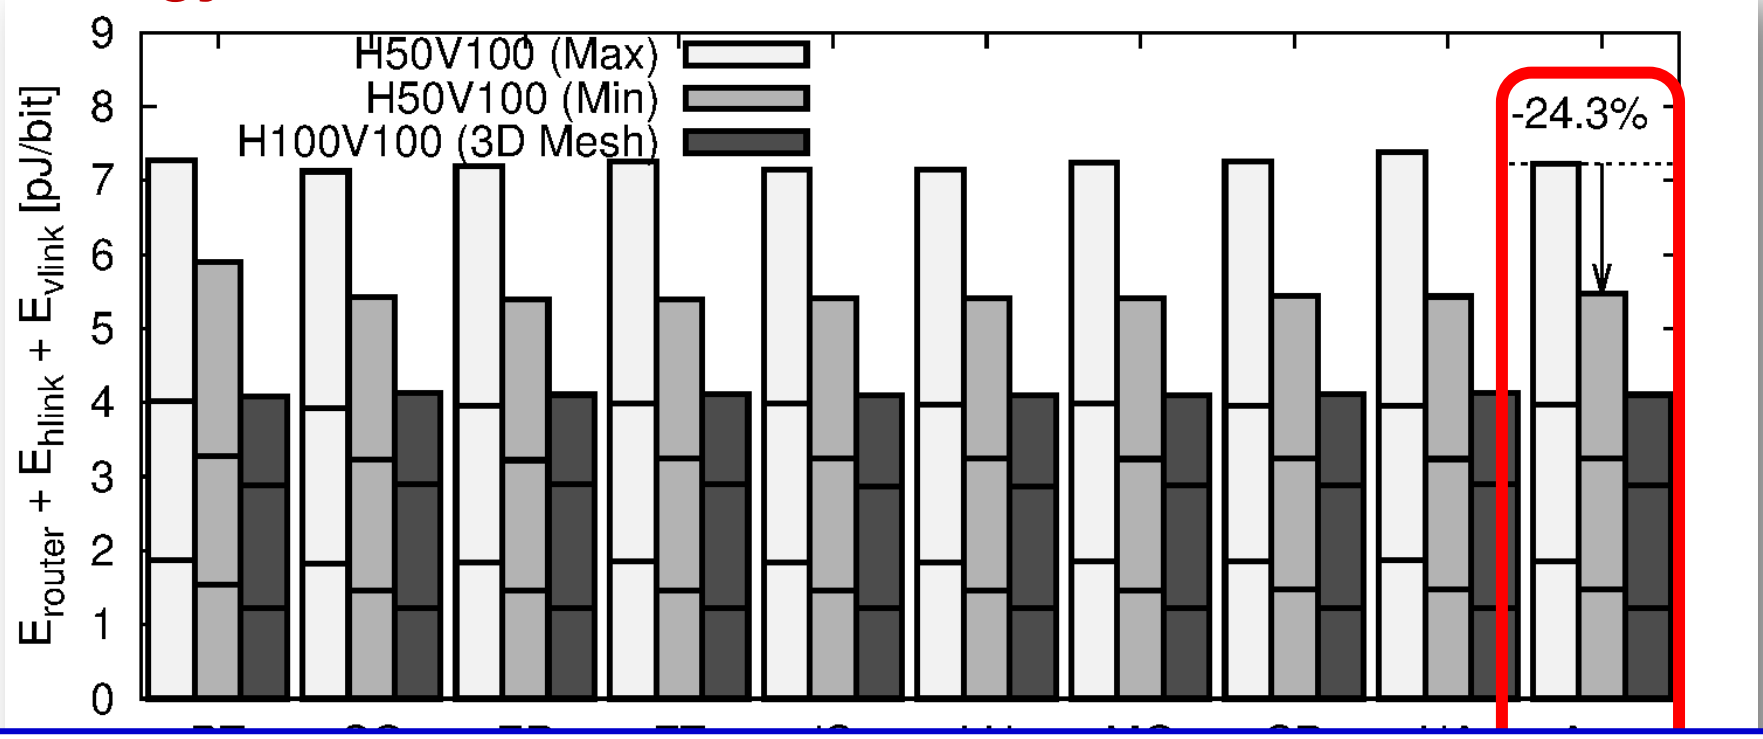

The next slides will explore vertical ON/OFF links (e.g., H100V50)

Power management: Vertical links

- 5.8mW per 2Gbps channel
- Stop power supply to inductors for low-power
  - **"H100-Vn"** topology (**n** =100, 50, 25, 15)
  - Performance/power is improved by spanning tree optimization

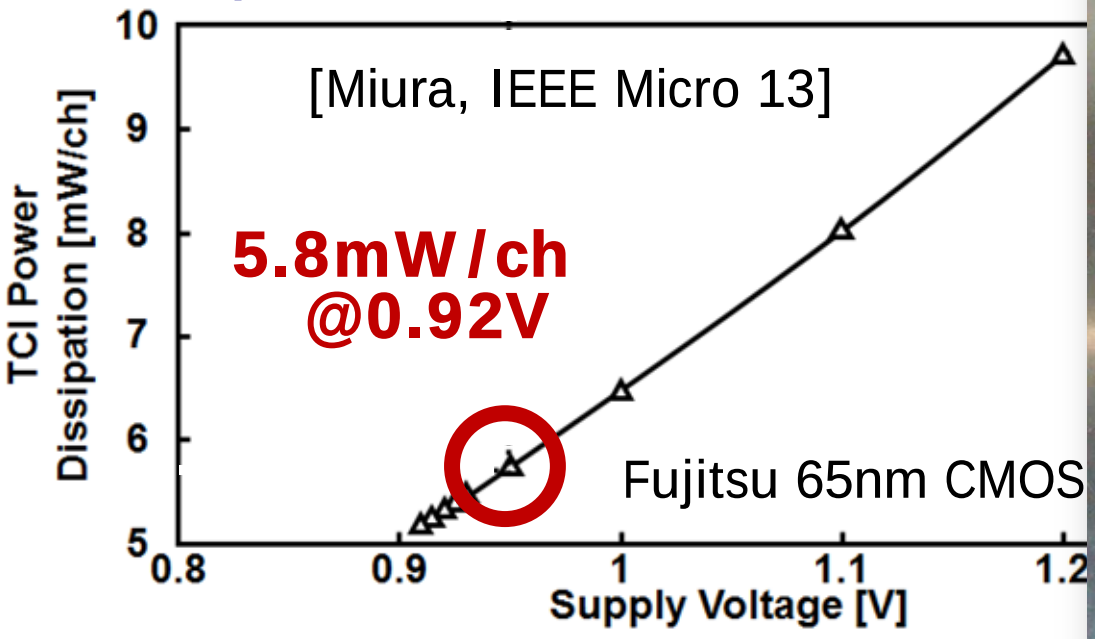

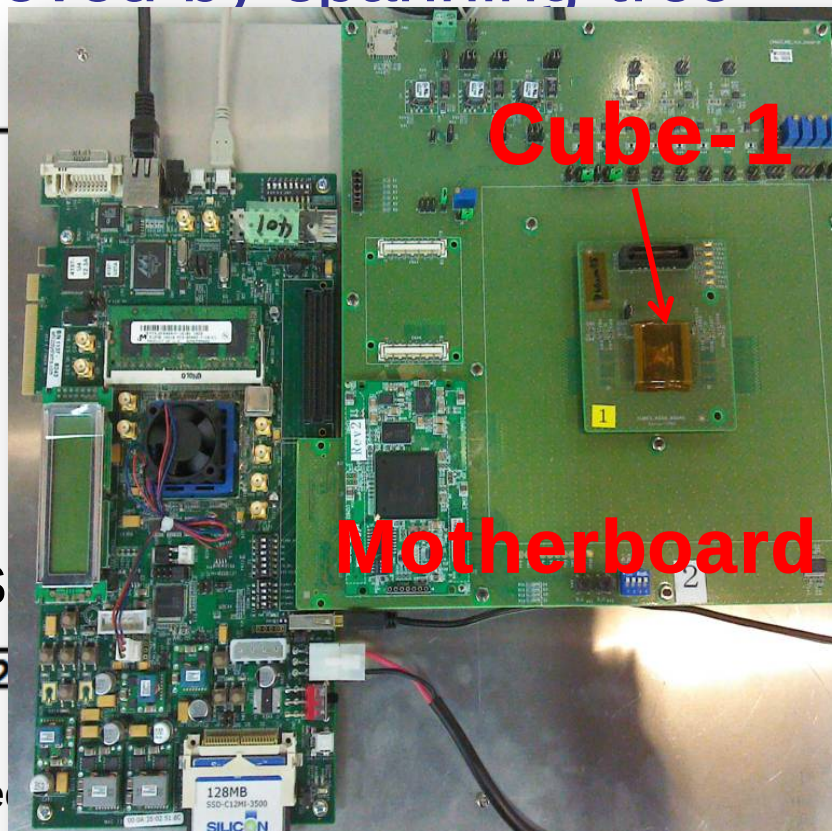

Hiroki Matsutani, "3D WiNoC Architec

Power mgmt: Vertical ON/OFF links

- **H100-Vn** topology (**n** =100, 50, 25, 15)
  - All horizontal links are available **(H100)**
  - **(100-n)**% of vertical links are power-gated **(Vn)**
  - Pick up most typical **H100-Vn** among 1,000 trials

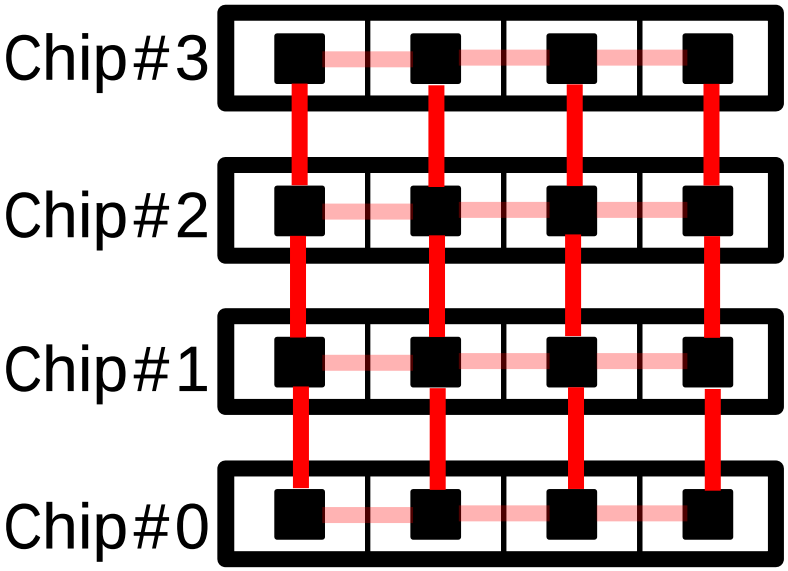

H100-**V100**

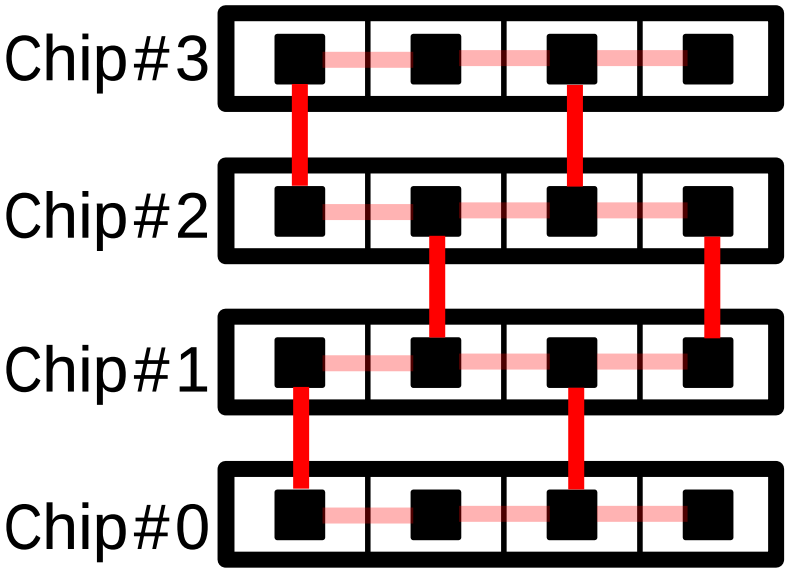

H100-**V50**

Power mgmt: Vertical off link selection

- Random+SPopt (computation cost: low)
  - Stop **(100-n)**% of vertical links **randomly**
  - Then, spanning tree optimization is performed to mitigate the performance penalty
- Exhaustive search (computation cost: **HUGE**)
  - Find the least important vertical links and stop them until **(100-n)**% of vertical links are removed

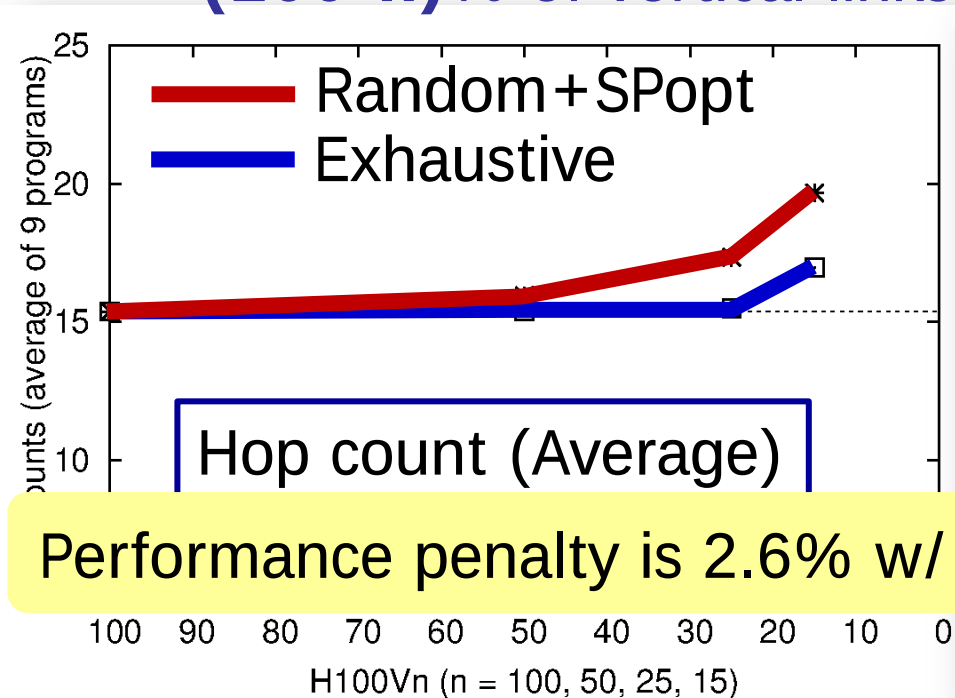

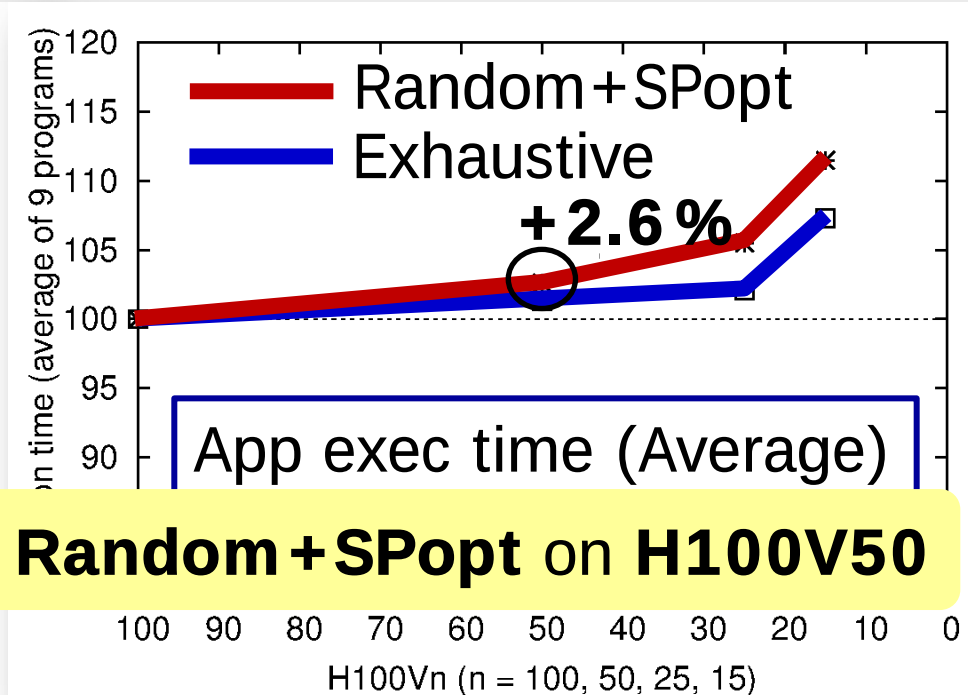

Performance penalty is 2.6% w/ **Random+SPopt** on **H100V50**

Outline: Wireless 3D NoC architectures

- Wireless 3D NoCs (3D WiNoCs) [8min]
  - Wireless 3D IC technology [Matsutani,NOCS'11]
  - 3D WiNoC design example (65nm) [Miura,IEEE Micro 13]
- Routing & topology exploration [8min]
  - Spanning tree optimization for irregular WiNoCs
  - Power management via vertical ON/OFF links [Matsutani,ASPDAC'13]
- Adding random NoC chip for 3D WiNoCs [5min]
  - Adding randomness induces small-world effects
  - Adding random NoC chip to NoC-less 3D ICs [Matsutani,DATE'14]
- Summary [1min]

Big picture: 3D WiNoC w/ Random

- Router IP macros have the same # of ports
  - Unused ports can be used for long-range links
- In addition, redundant links can be implemented
  - Statically multiplexed by FPGA-like switch boxes
  - By reconfiguring the switch boxes, an unique random topology is generated

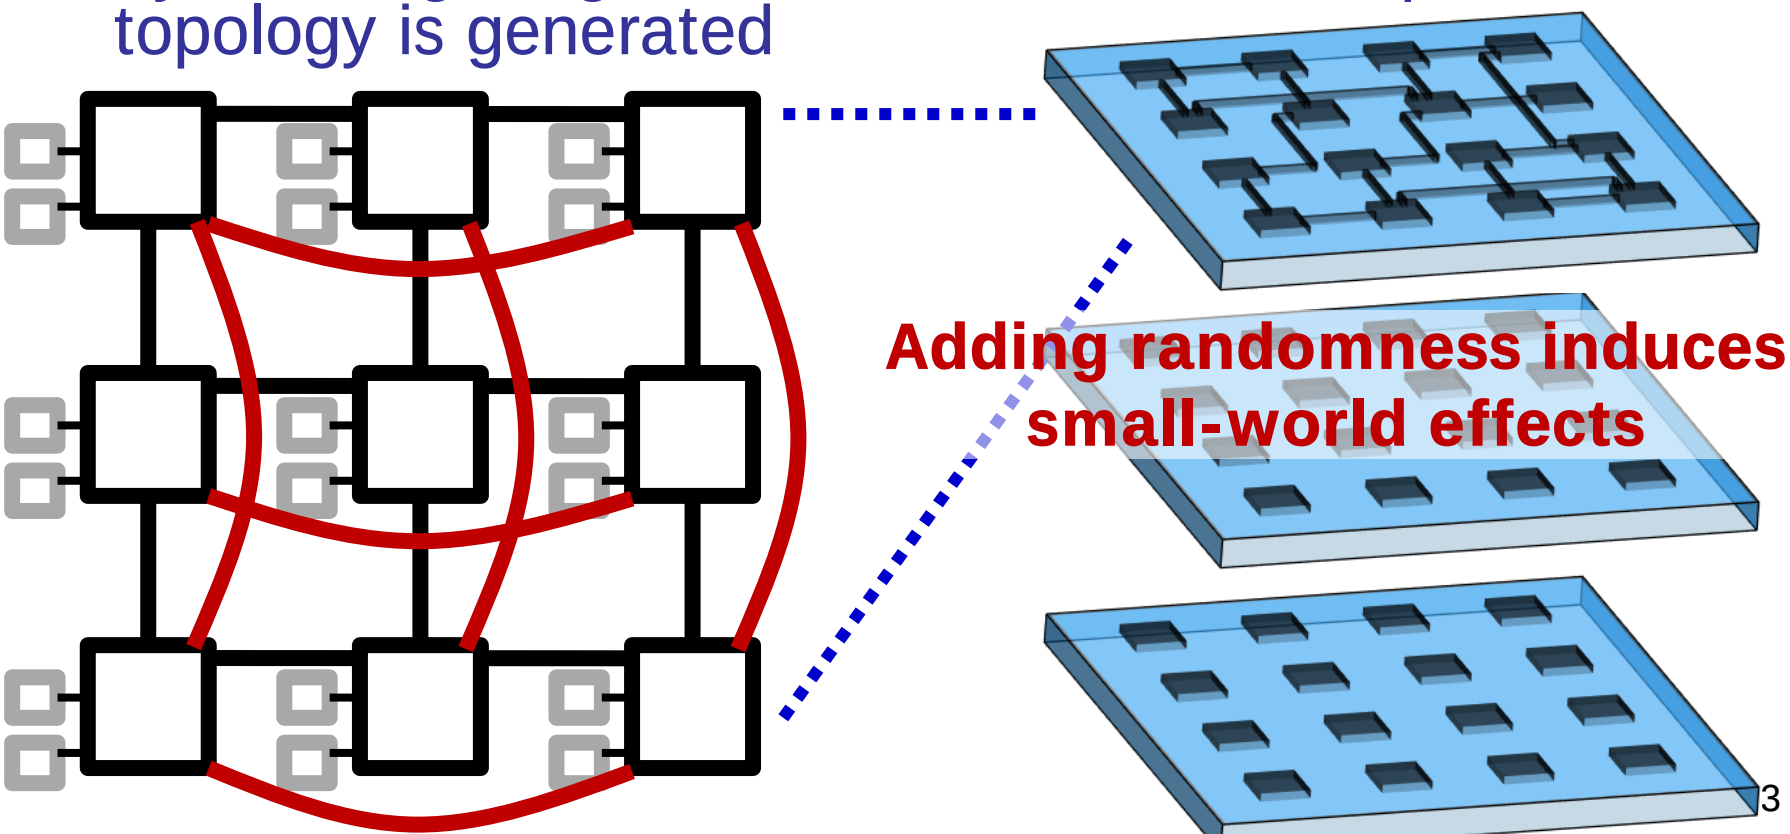

Case 1: Random NoC to NoC-less 3D ICs

- Each chip has inductors but does not have 2D NoC
  - Inductors in the same pillar form a **vertical broadcast bus**
  - Horizontal connectivity is not provided at all (i.e., NoC-less)

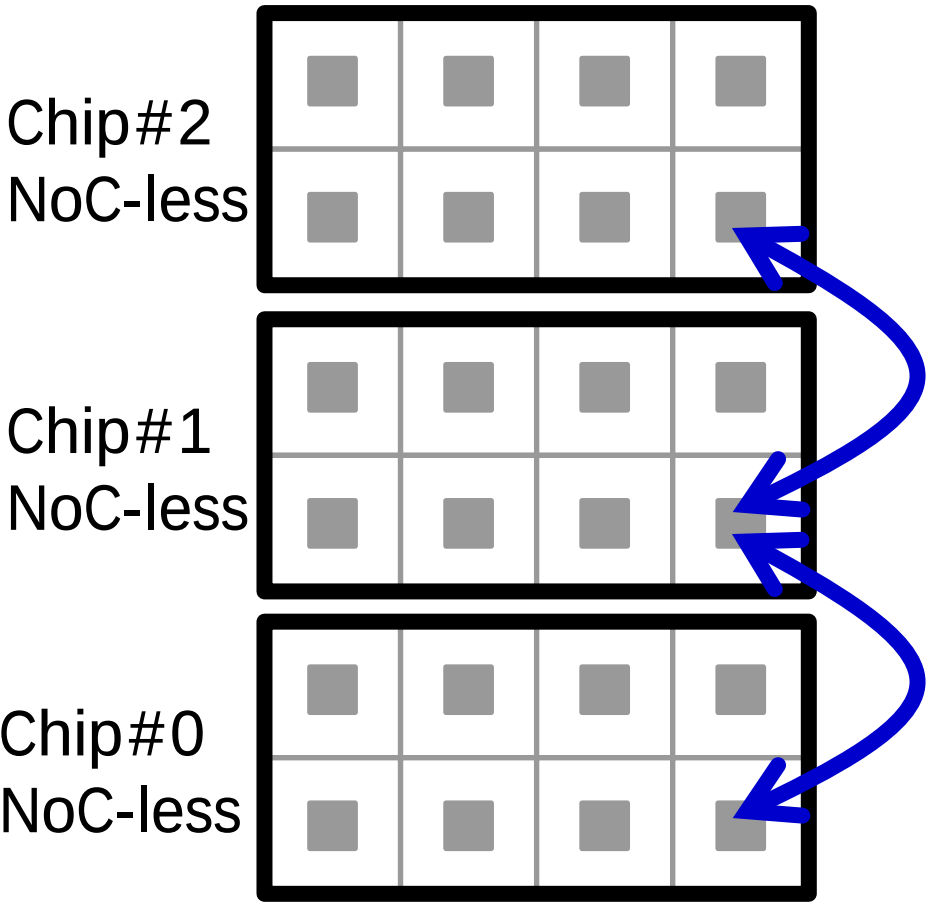

"– – –" Configuration

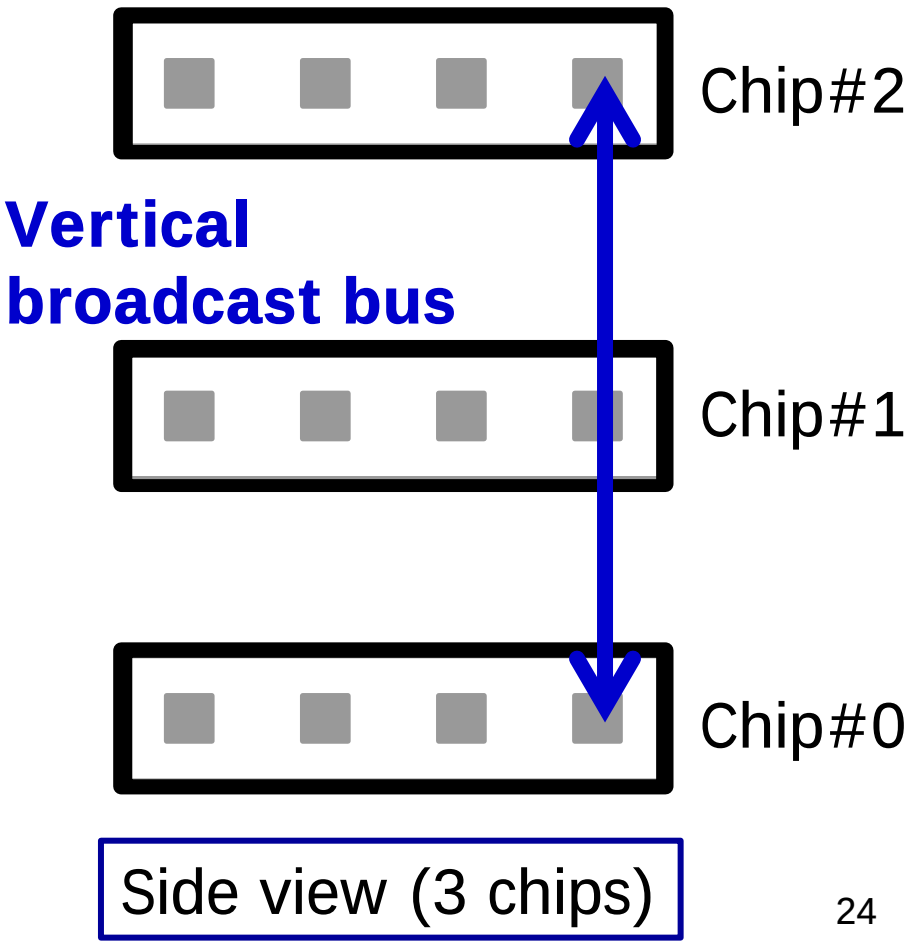

Side view (3 chips)

Case 1: Random NoC to NoC-less 3D ICs

- Each chip has inductors but does not have 2D NoC
  - Inductors in the same pillar form a **vertical broadcast bus**
  - Adding a **2D Mesh NoC** to such **NoC-less 3D IC**

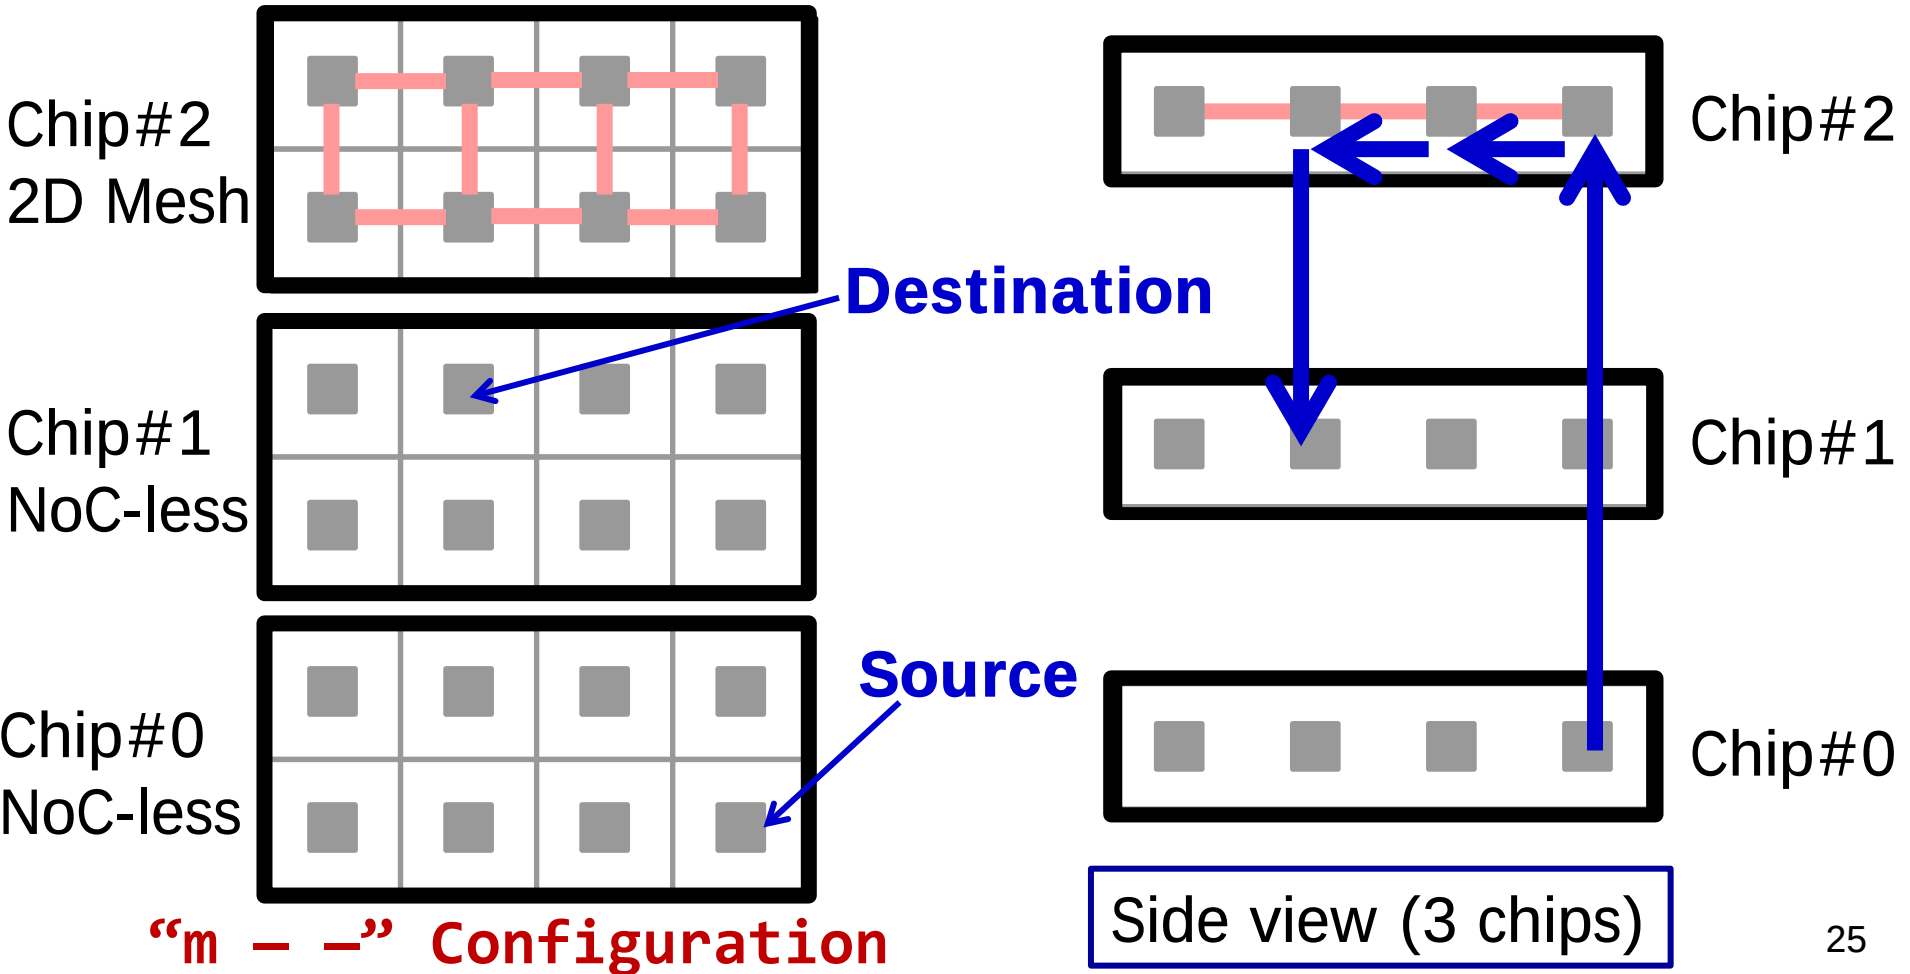

Case 1: Random NoC to NoC-less 3D ICs

- Each chip has inductors but does not have 2D NoC
  - Inductors in the same pillar form a **vertical broadcast bus**
  - Adding a **Random NoC** to such **NoC-less 3D IC**

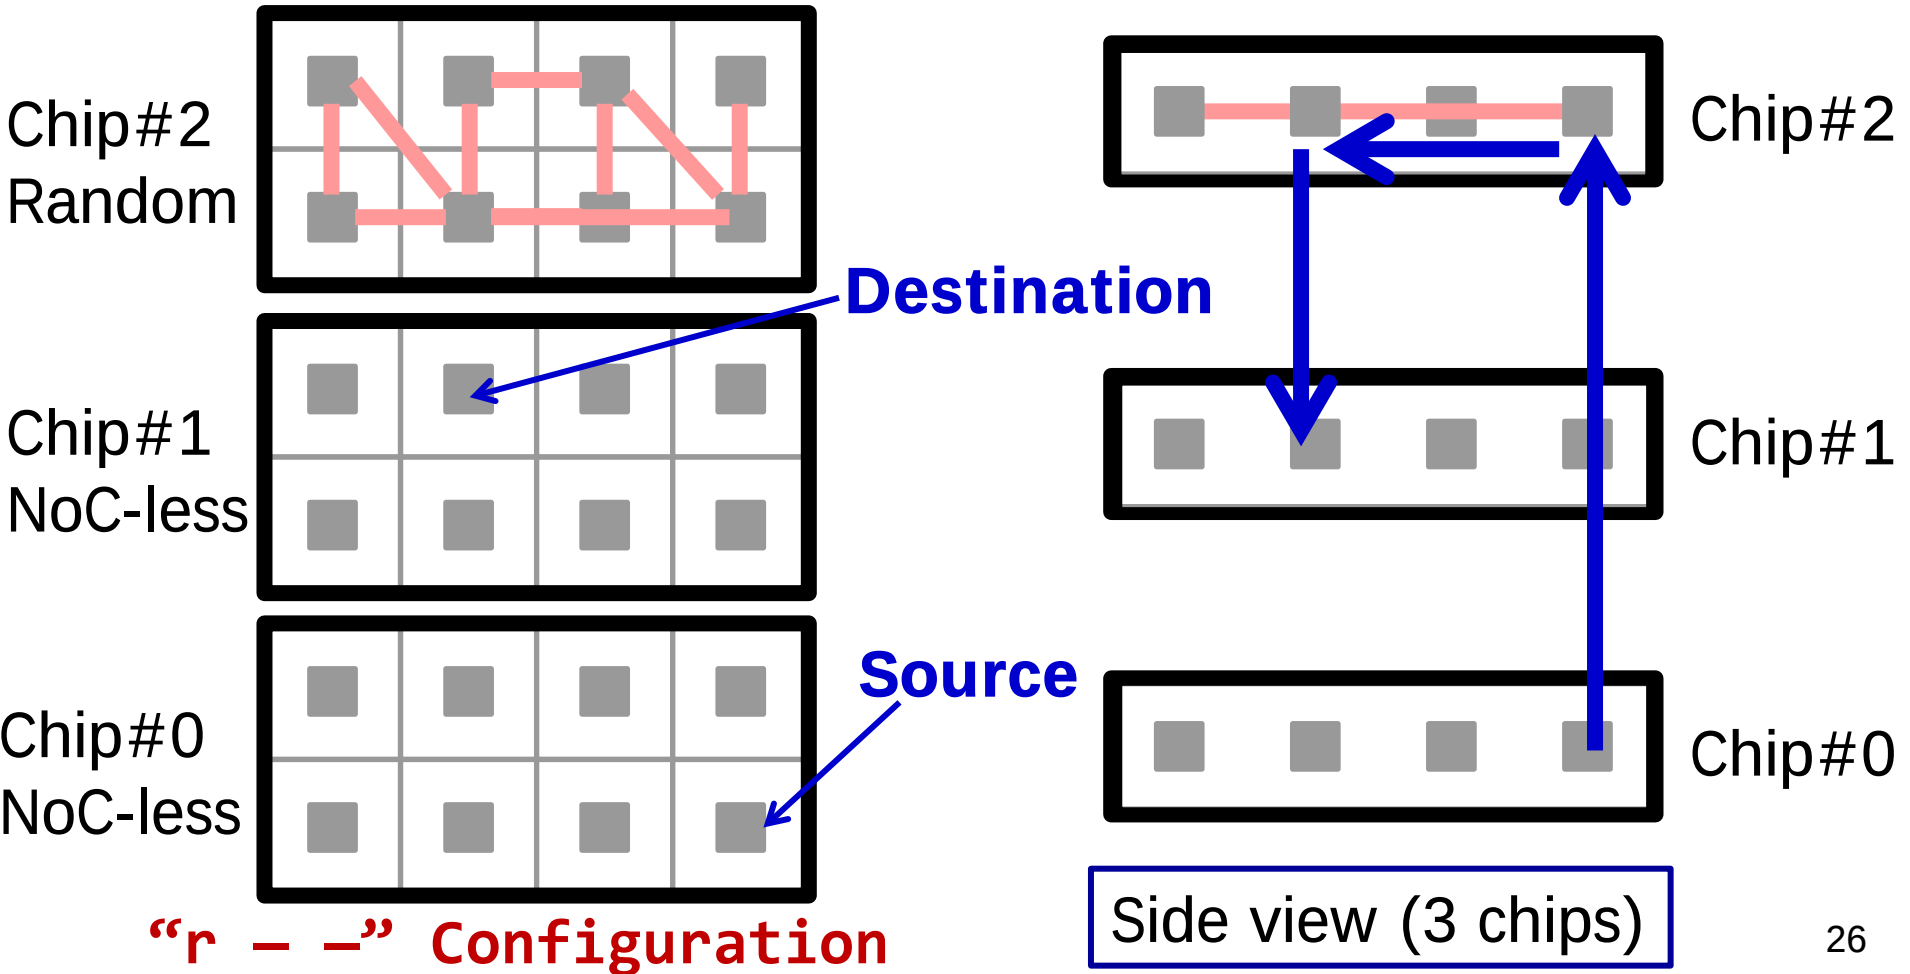

Case 2: Replacing regular NoC w/ random

- 3D WiNoC that consists of three 2D Mesh NoC layers
  - Inductors in neighboring chips form a **vertical P2P link**
  - E.g., regular 3D Mesh topology (i.e., regular 3D NoC)

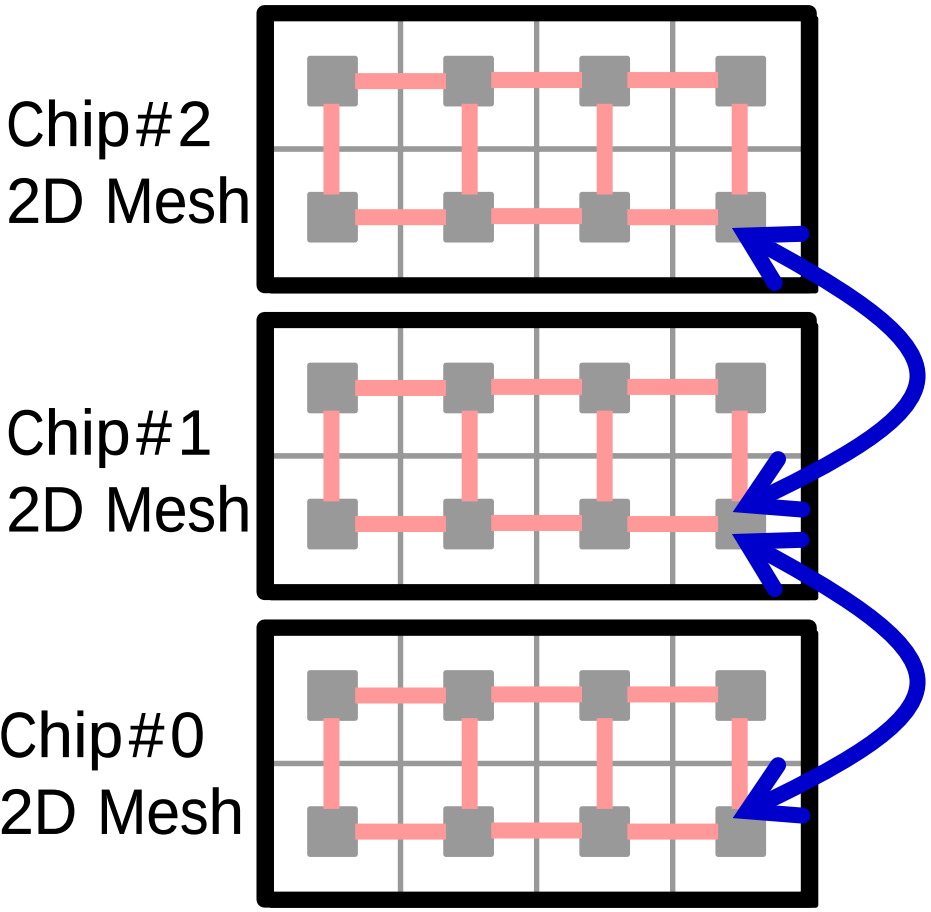

"m m m" Configuration

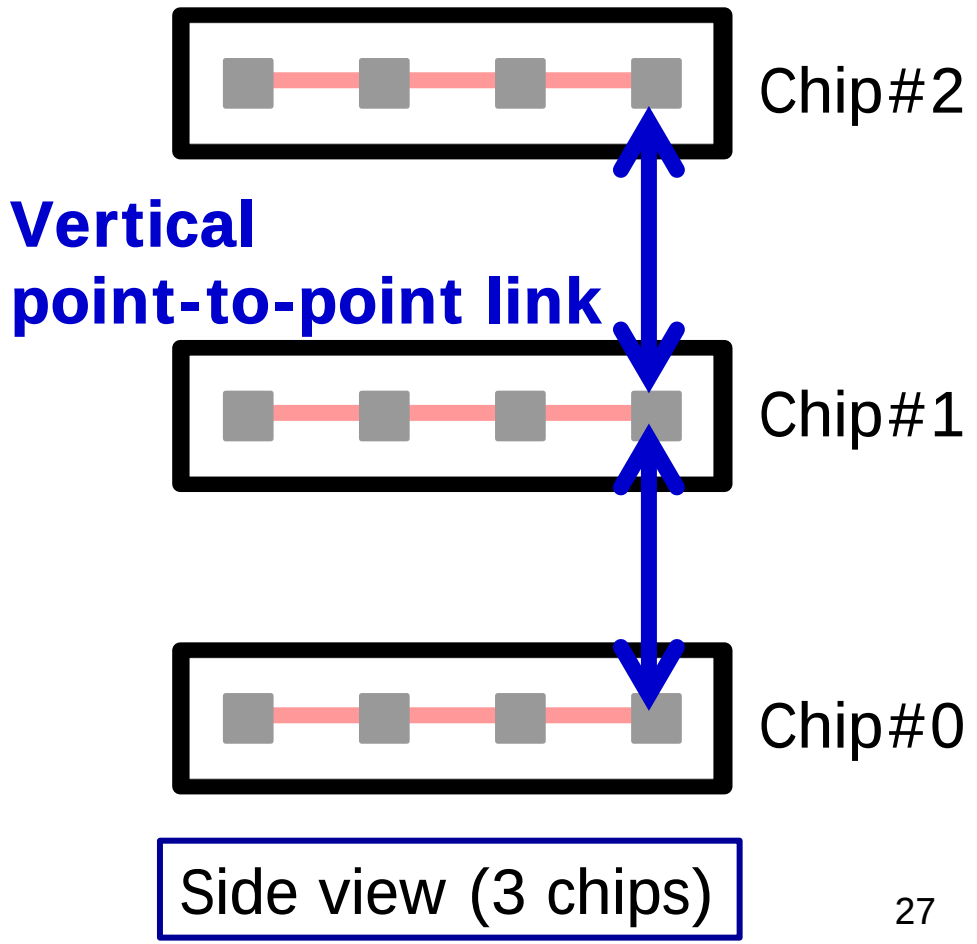

Side view (3 chips)

Case 2: Replacing regular NoC w/ random

- 3D WiNoC that consists of three 2D Mesh NoC layers
  - Inductors in neighboring chips form a **vertical P2P link**
  - Replacing **regular 2D Mesh** with **random NoC**

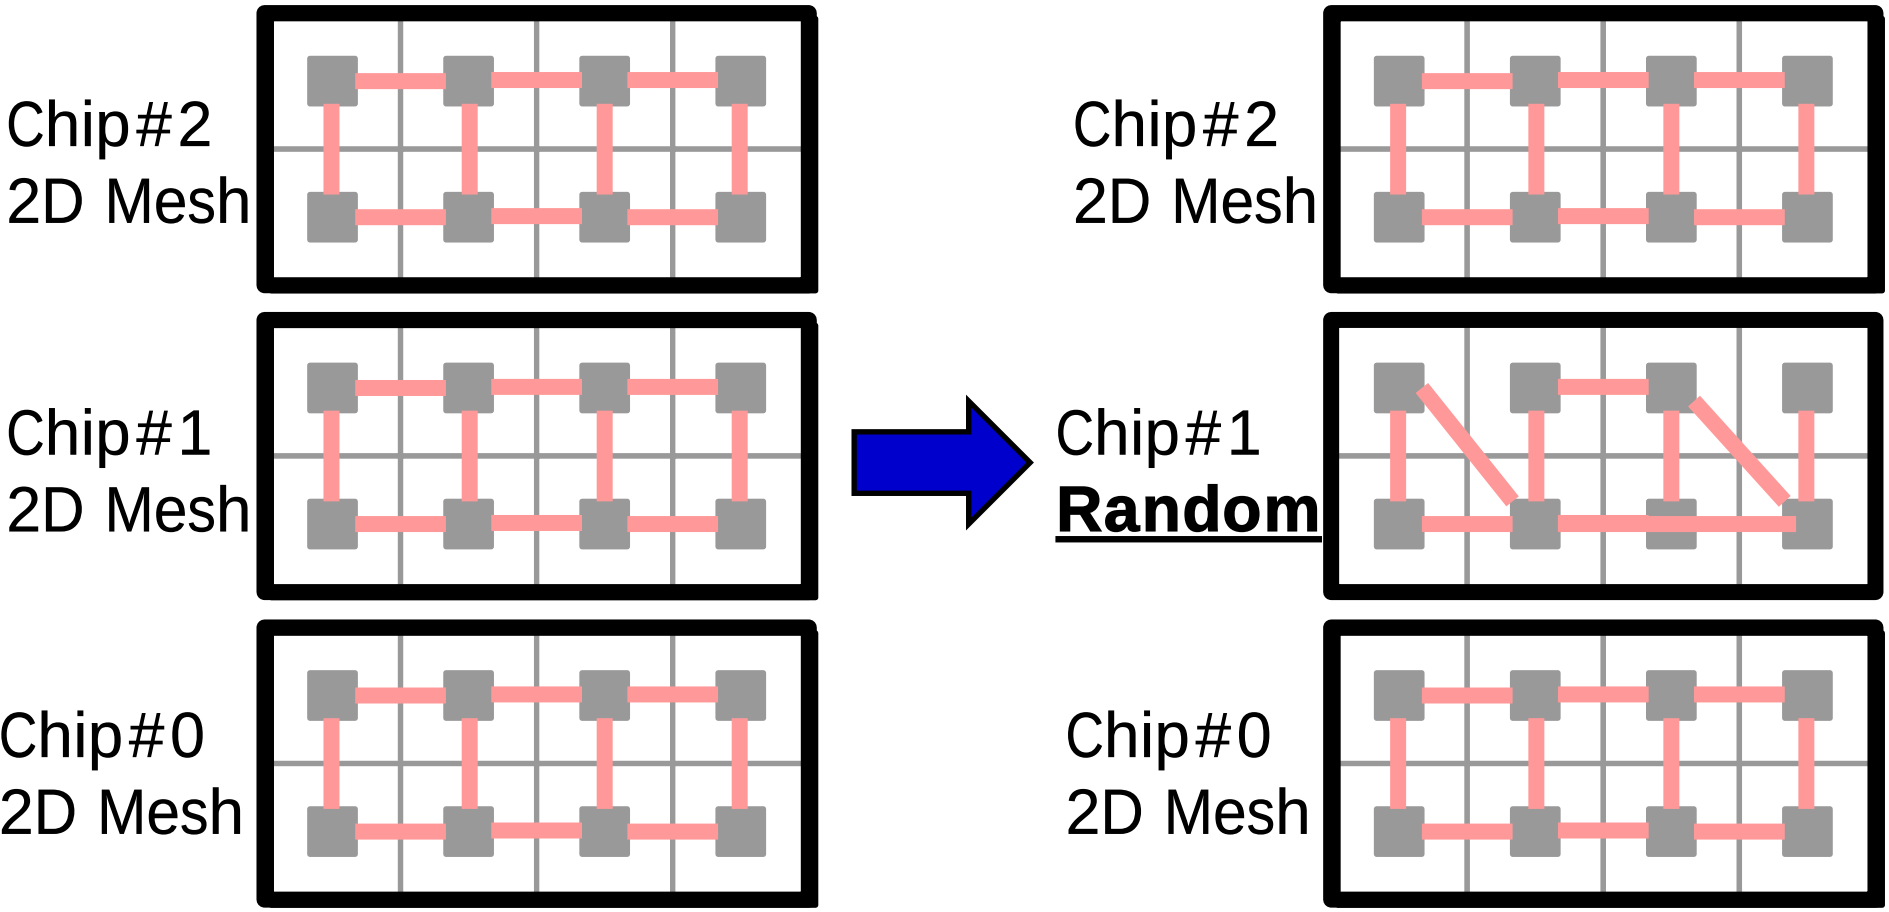

Q1: How many random chips do we need?

Number of random chips vs. Average latency

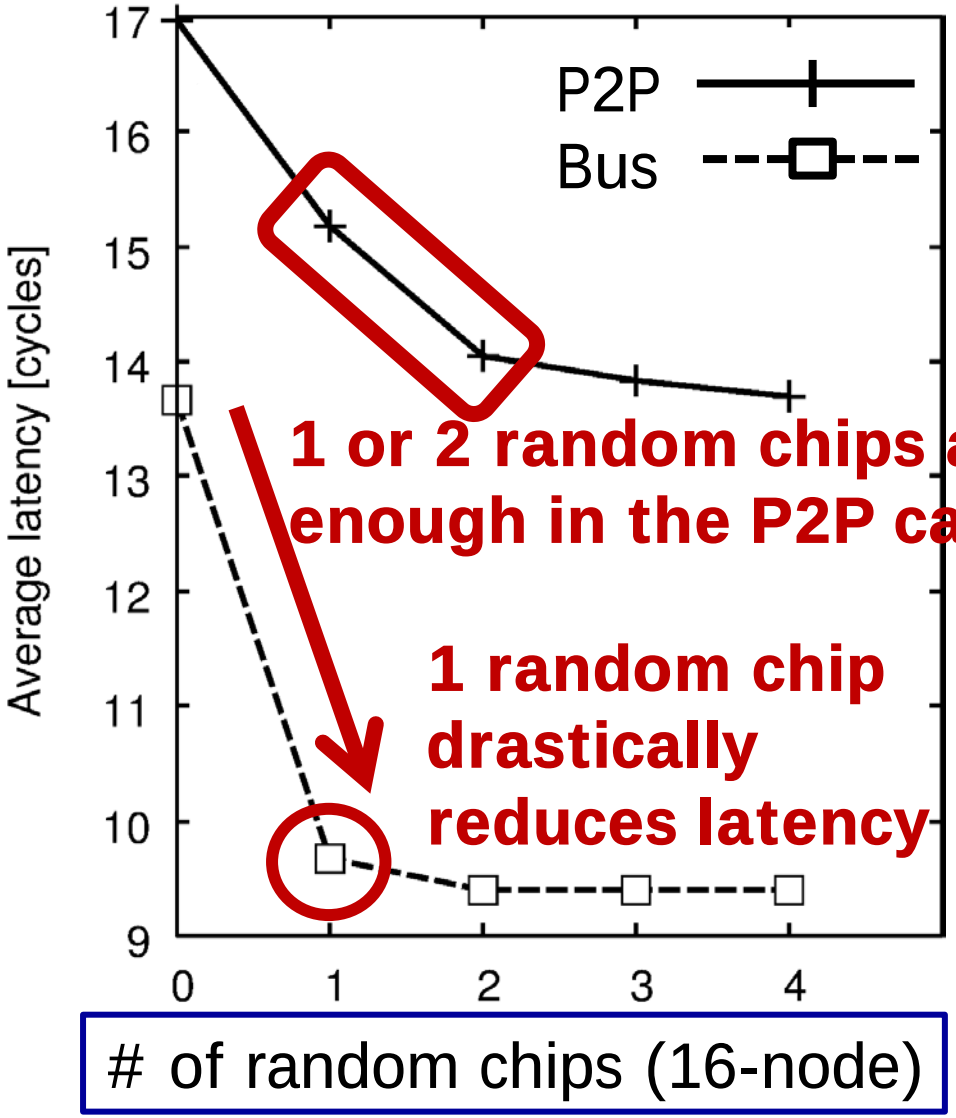

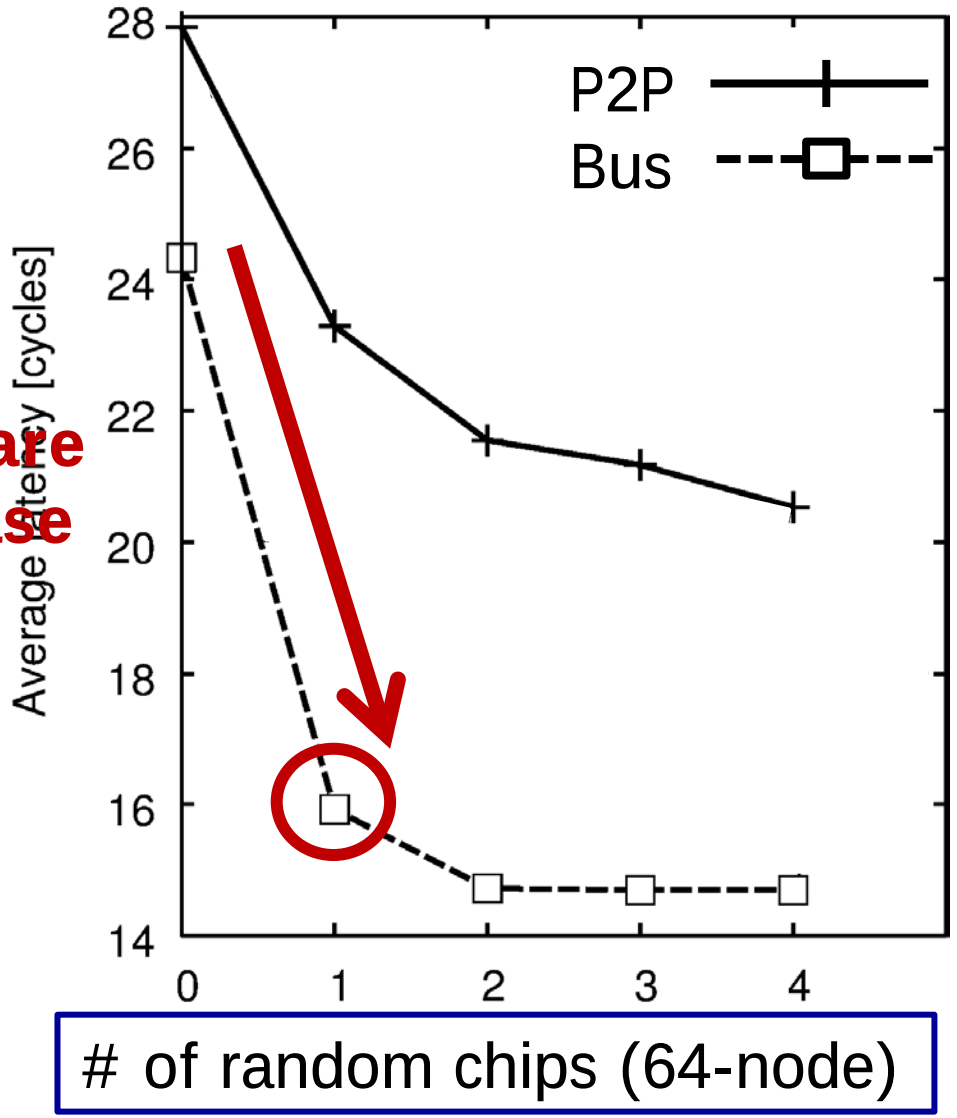

Q2: How should we design random chip?

Max. link length (left) & Horizontal degree (right)

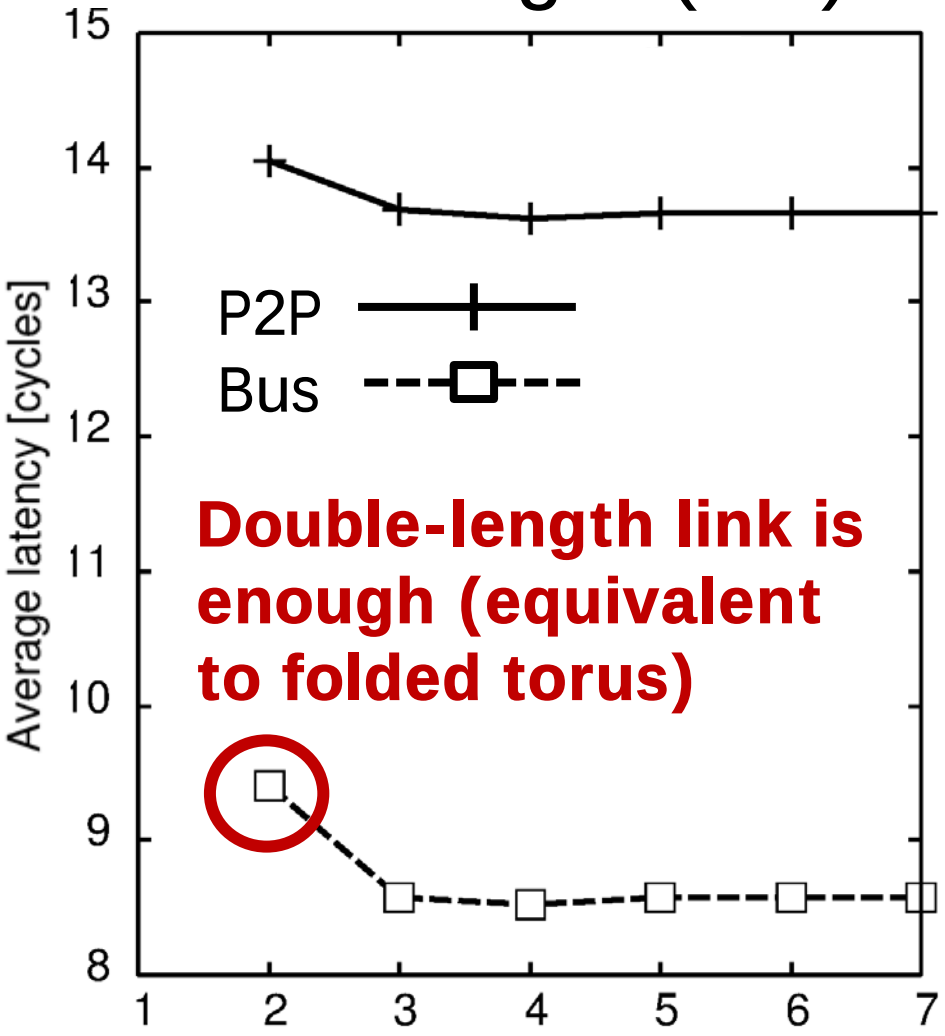

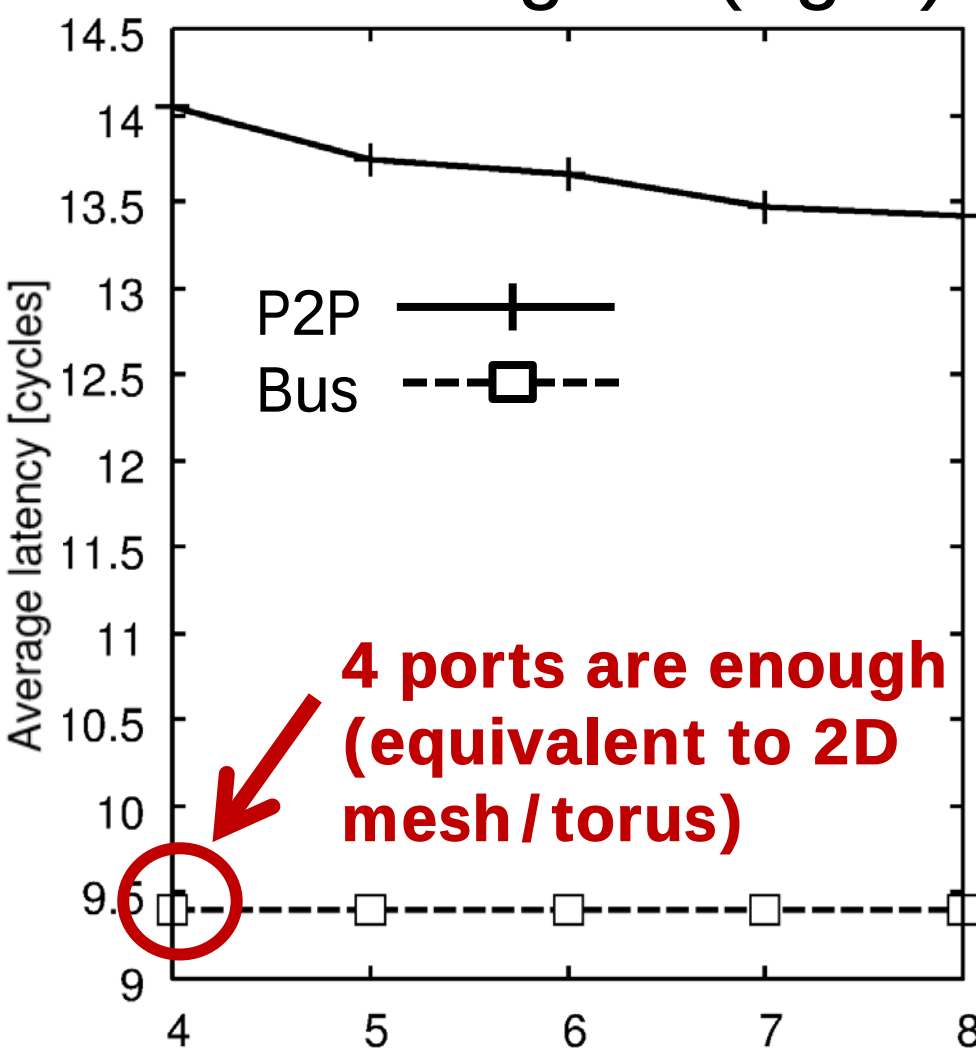

Packet latency: P2P (mmmm mrrm rrrr)

m m m m

m r r m

r r r r

m m m m	Mesh	Mesh	Mesh	Mesh	P2P
m r r m	Mesh	Rand	Rand	Mesh	P2P
r r r r	Rand	Rand	Rand	Rand	P2P
- - - m	None	None	None	Mesh	**Bus**
- - - r	None	None	None	Rand	**Bus**
m - - r	Mesh	None	None	Rand	**Bus**

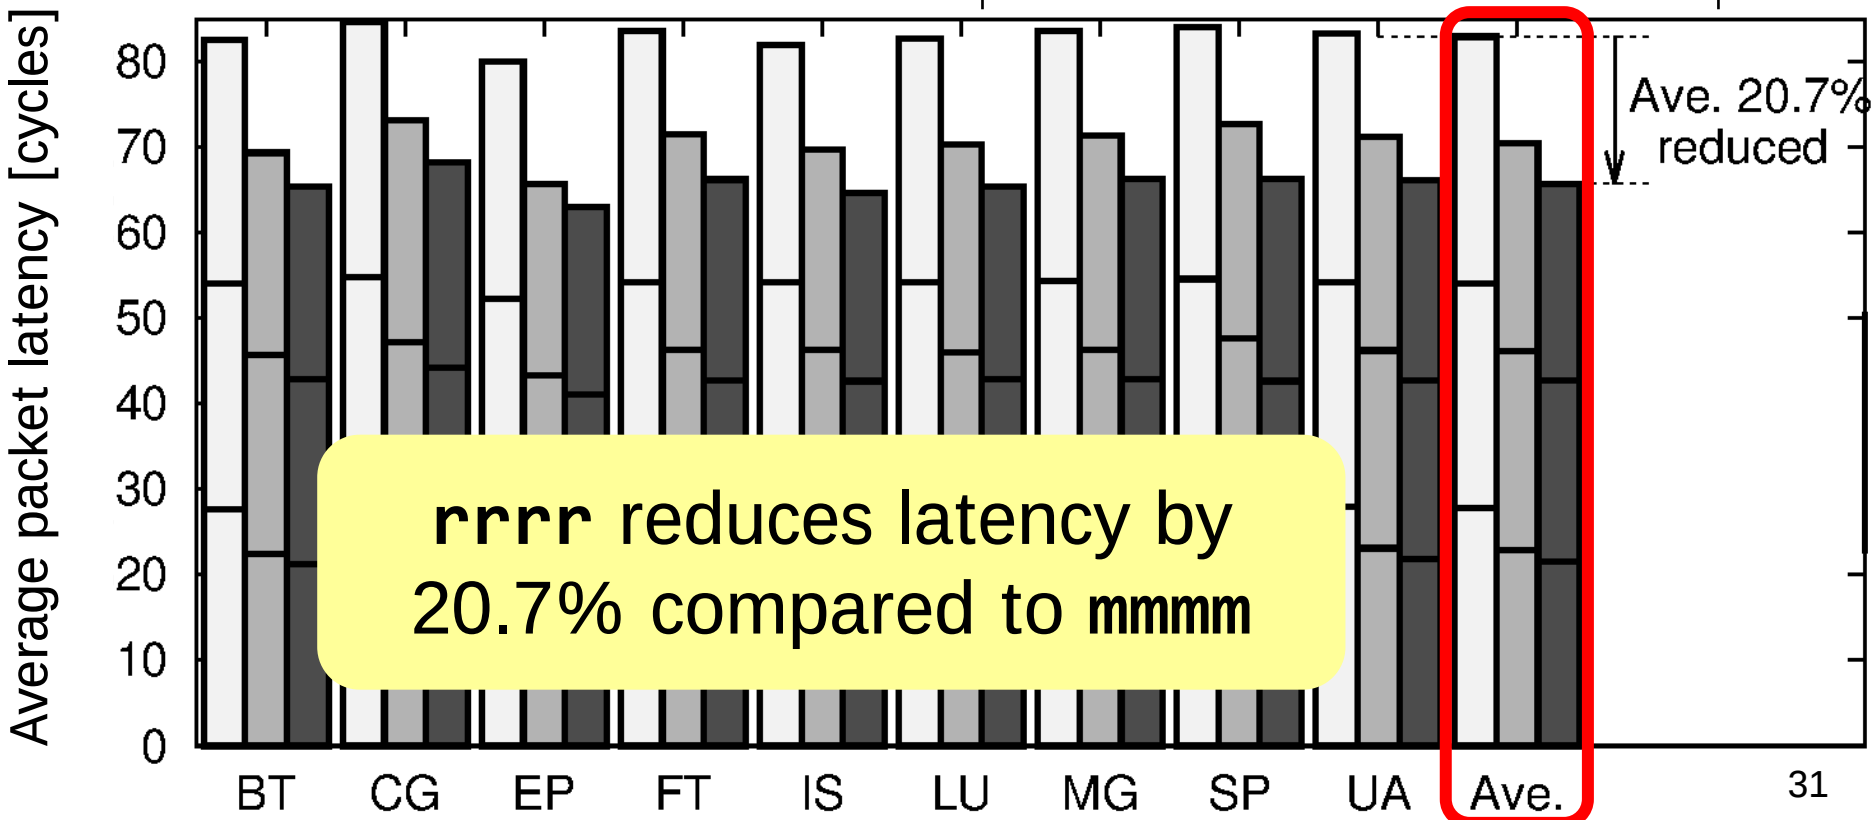

rrrr reduces latency by 20.7% compared to **mmmm**

Packet latency: Bus (---m ---r m--r)

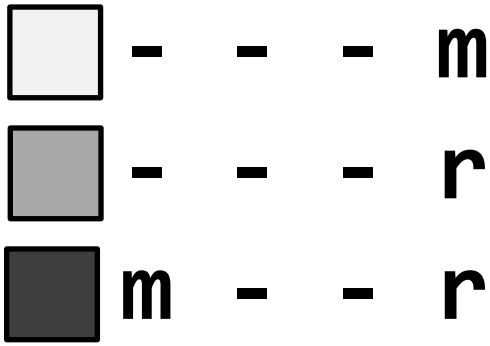

m m m m	Mesh	Mesh	Mesh	Mesh	P2P
m r r m	Mesh	Rand	Rand	Mesh	P2P
r r r r	Rand	Rand	Rand	Rand	P2P
- - - m	None	None	None	Mesh	**Bus**
- - - r	None	None	None	Rand	**Bus**
m - - r	Mesh	None	None	Rand	**Bus**

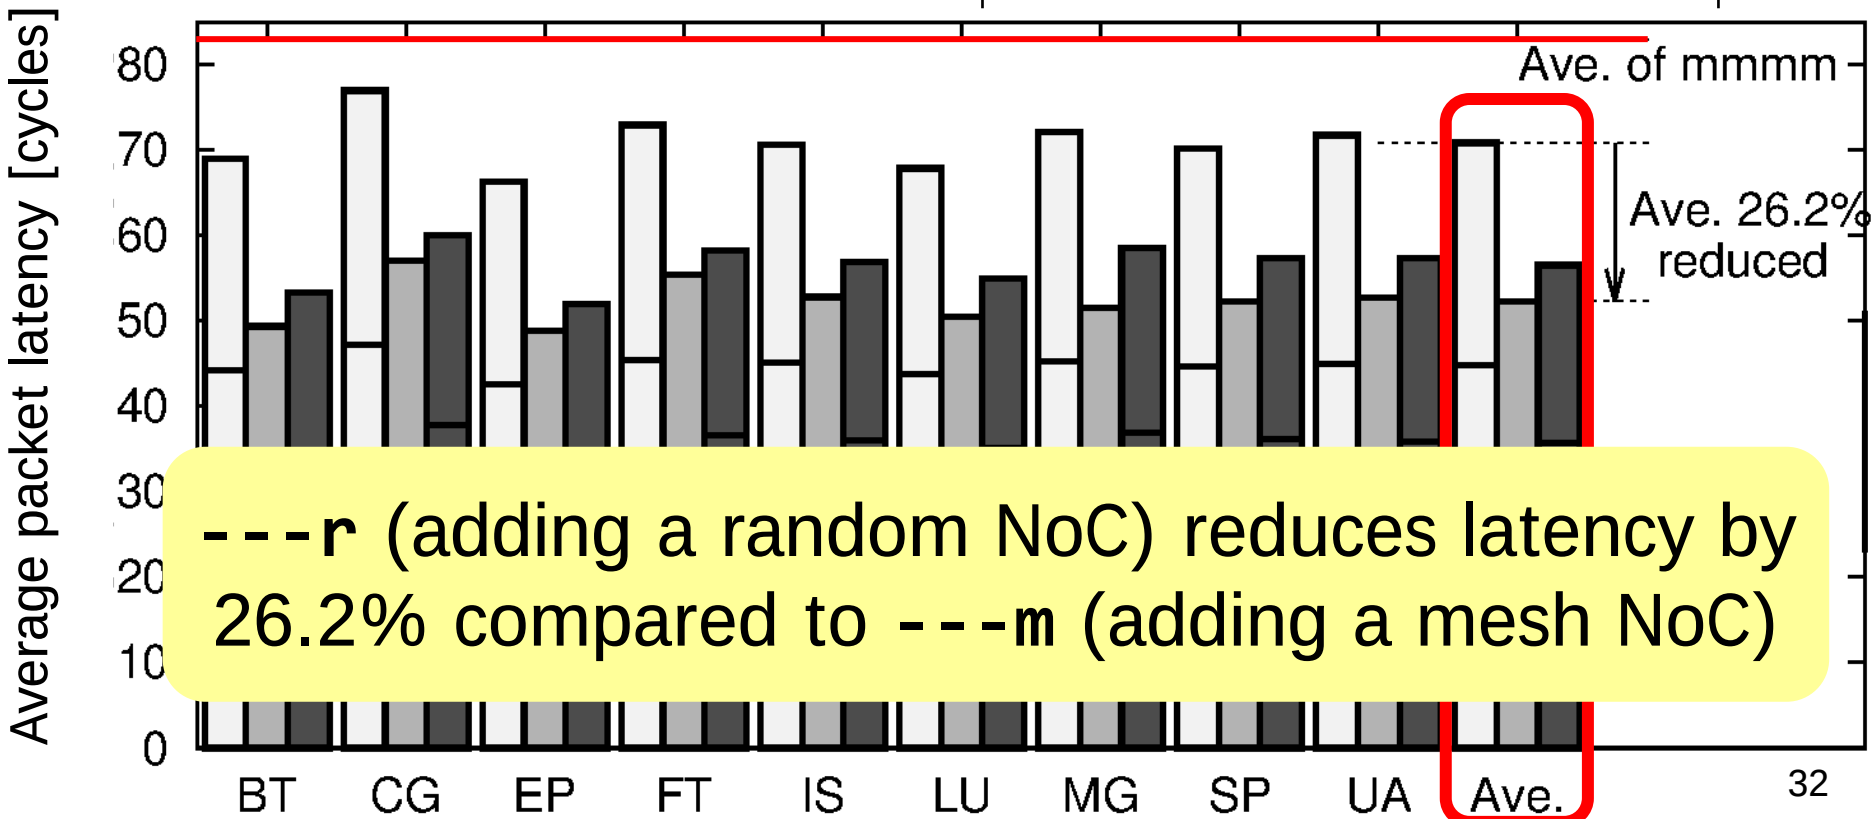

---r (adding a random NoC) reduces latency by 26.2% compared to ---m (adding a mesh NoC)

Outline: Wireless 3D NoC architectures

- Wireless 3D NoCs (3D WiNoCs) [8min]
  - Wireless 3D IC technology [Matsutani,NOCS'11]
  - 3D WiNoC design example (65nm) [Miura,IEEE Micro 13]
- Routing & topology exploration [8min]
  - Spanning tree optimization for irregular WiNoCs
  - Power management via vertical ON/OFF links [Matsutani,ASPDAC'13]
- Adding random NoC chip for 3D WiNoCs [5min]
  - Adding randomness induces small-world effects
  - Adding random NoC chip to NoC-less 3D ICs [Matsutani,DATE'14]
- Summary [1min]

Summary: 3D WiNoC Architectures

- Inductive-coupling 3D SiP
  - A low cost alternative to build low-volume custom systems by stacking off-the-shelf known-good-dies
  - No special process technology is required; inductors are implemented with metal layers
- Cube-1: A practical 3D WiNoC system
  - Two types: Host CPU chip & Accelerator chips
  - We can customize number & types of chips in SiP

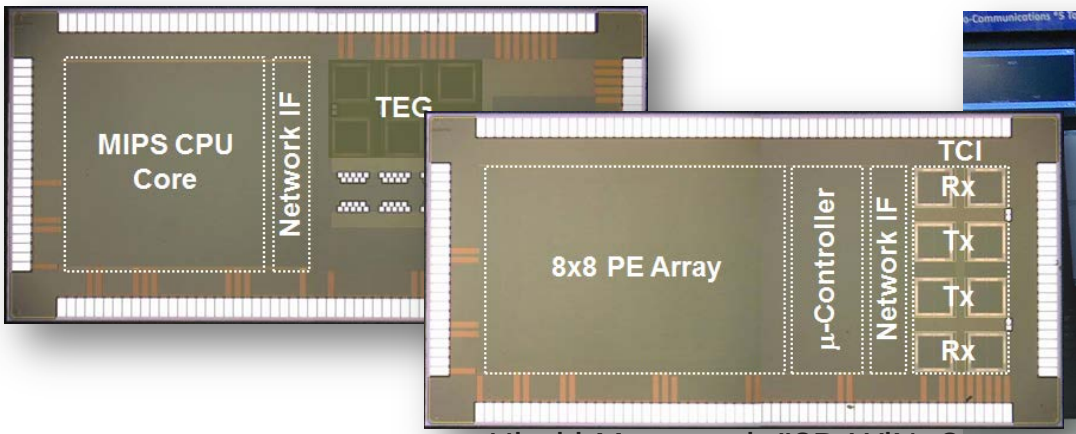

Hiroki Matsutani, "3D WiNoC

Summary: 3D WiNoC Architectures

- Routing & topology exploration
  - Spanning tree optimization for adhoc 3D WiNoCs
  - Power reduction by **randomly** stopping vertical links
  - Performance/power is improved by spanning tree optimization
- Adding randomness shortens comm. latency
  - Adding **random NoC chip** to **NoC-less 3D ICs**
  - Replacing **regular 2D NoC** with **random 2D NoC**

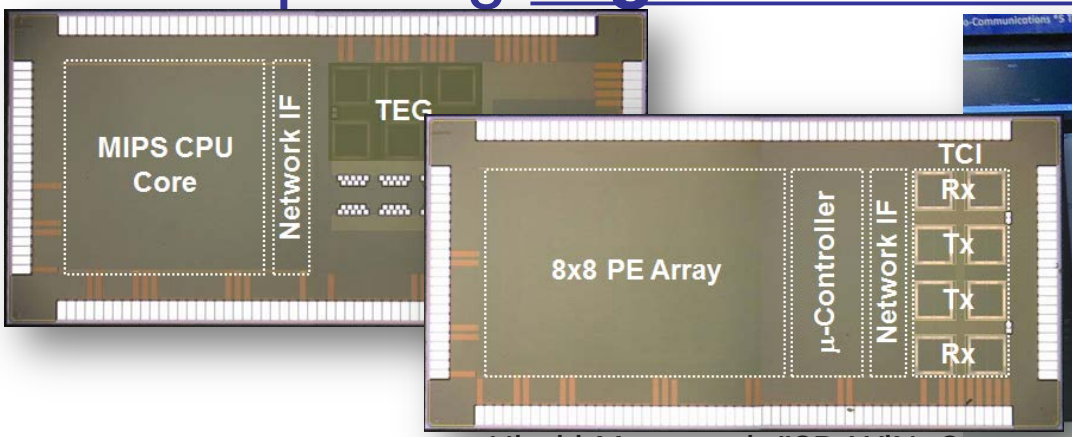

Hiroki Matsutani, "3D WiNoC

References (1/2)

- Cube-0: The first real 3D WiNoC
  - H. Matsutani, et.al., "A Vertical Bubble Flow Network using Inductive-Coupling for 3-D CMPs", NOCS 2011.
  - Y. Take, et.al., "3D NoC with Inductive-Coupling Links for Building-Block SiPs", IEEE Trans on Computers (2014).
- Cube-1: The heterogeneous 3D WiNoC
  - N. Miura, et.al., "A Scalable 3D Heterogeneous Multicore with an Inductive ThruChip Interface", IEEE Micro (2013).
- MuCCRA-Cube: Dynamically reconfigurable processor
  - S. Saito, et.al., "MuCCRA-Cube: a 3D Dynamically Reconfigurable Processor with Inductive-Coupling Link", FPL 2009.

References (2/2)

- Vertical bubble flow control on Cube-0
  - H. Matsutani, et.al., "A Vertical Bubble Flow Network using Inductive-Coupling for 3-D CMPs", NOCS 2011.
  - Y. Take, et.al., "3D NoC with Inductive-Coupling Links for Building-Block SiPs", IEEE Trans on Computers (2014).
- Spanning trees optimization for 3D WiNoCs
  - H. Matsutani, et.al., "A Case for Wireless 3D NoCs for CMPs", ASP-DAC 2013 (Best Paper Award).
- Small-world effects (randomness)
  - H. Matsutani, et.al., "Low-Latency Wireless 3D NoCs via Randomized Shortcut Chips", DATE 2014.
  - M. Koibuchi, et.al., "A Case for Random Shortcut Topologies for HPC Interconnects", ISCA 2012.

Backup slides