Electronic-Photonic Integrated Circuit and System

Jiang Xu

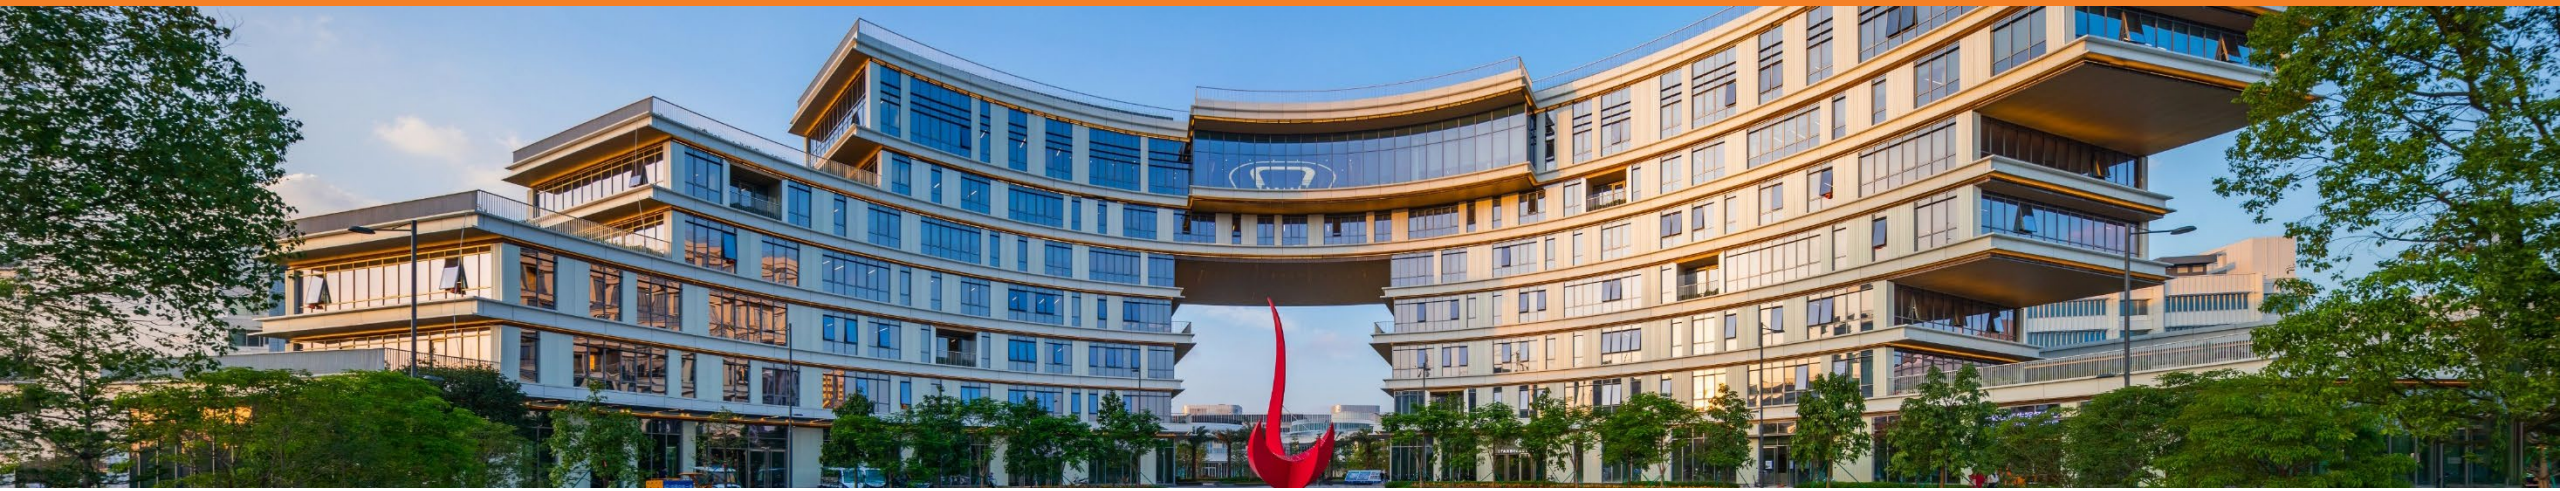

香港科技大学(广州)
THE HONG KONG UNIVERSITY OF SCIENCE AND TECHNOLOGY (GUANGZHOU)

微电子学域
MICROELECTRONICS THRUST

AI Outpaces Moore's Law

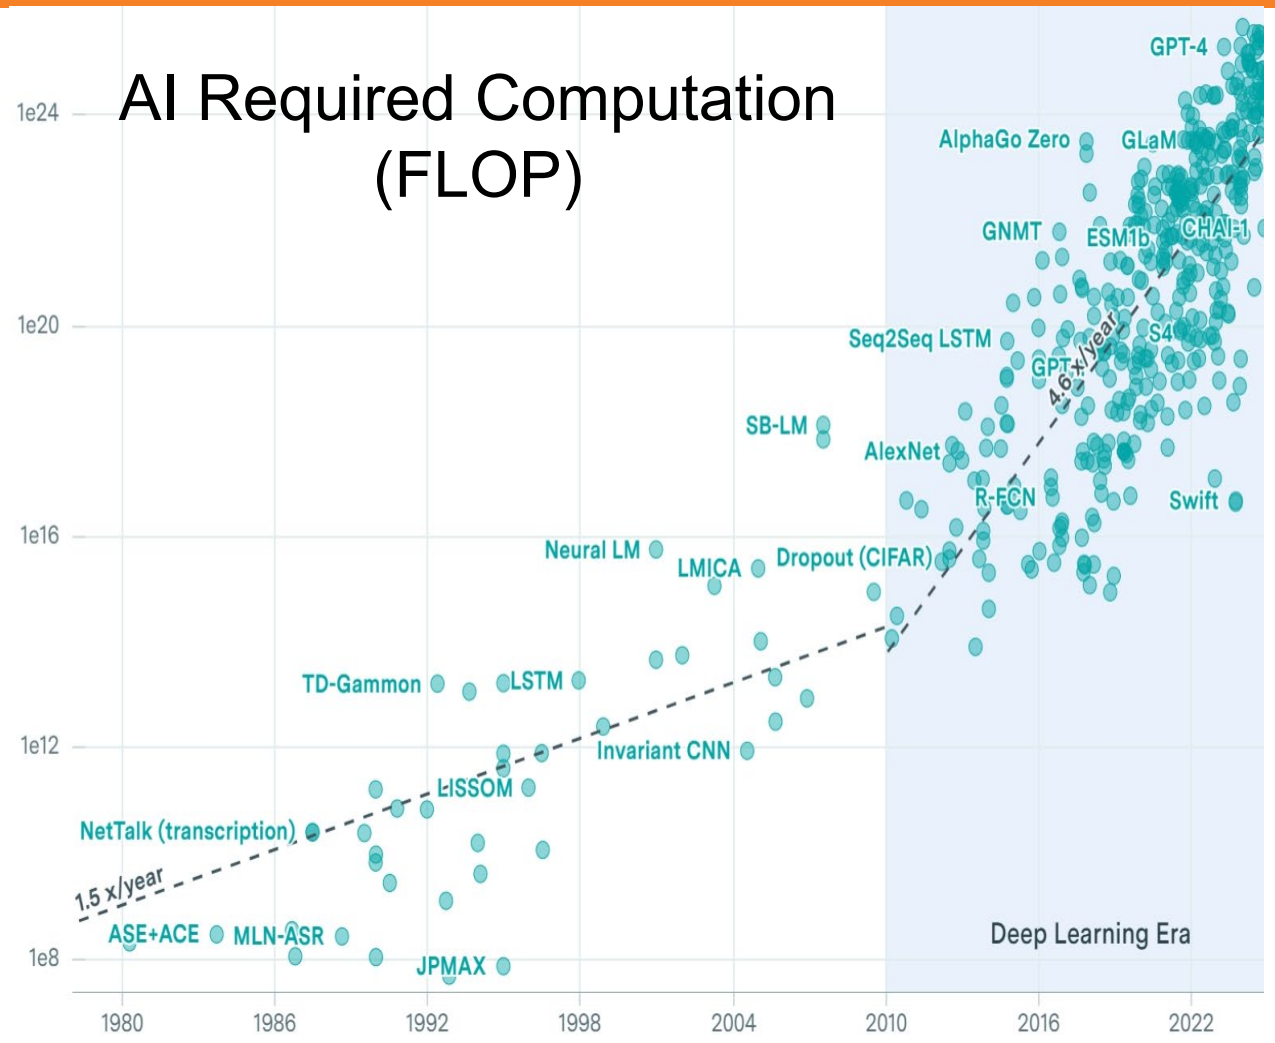

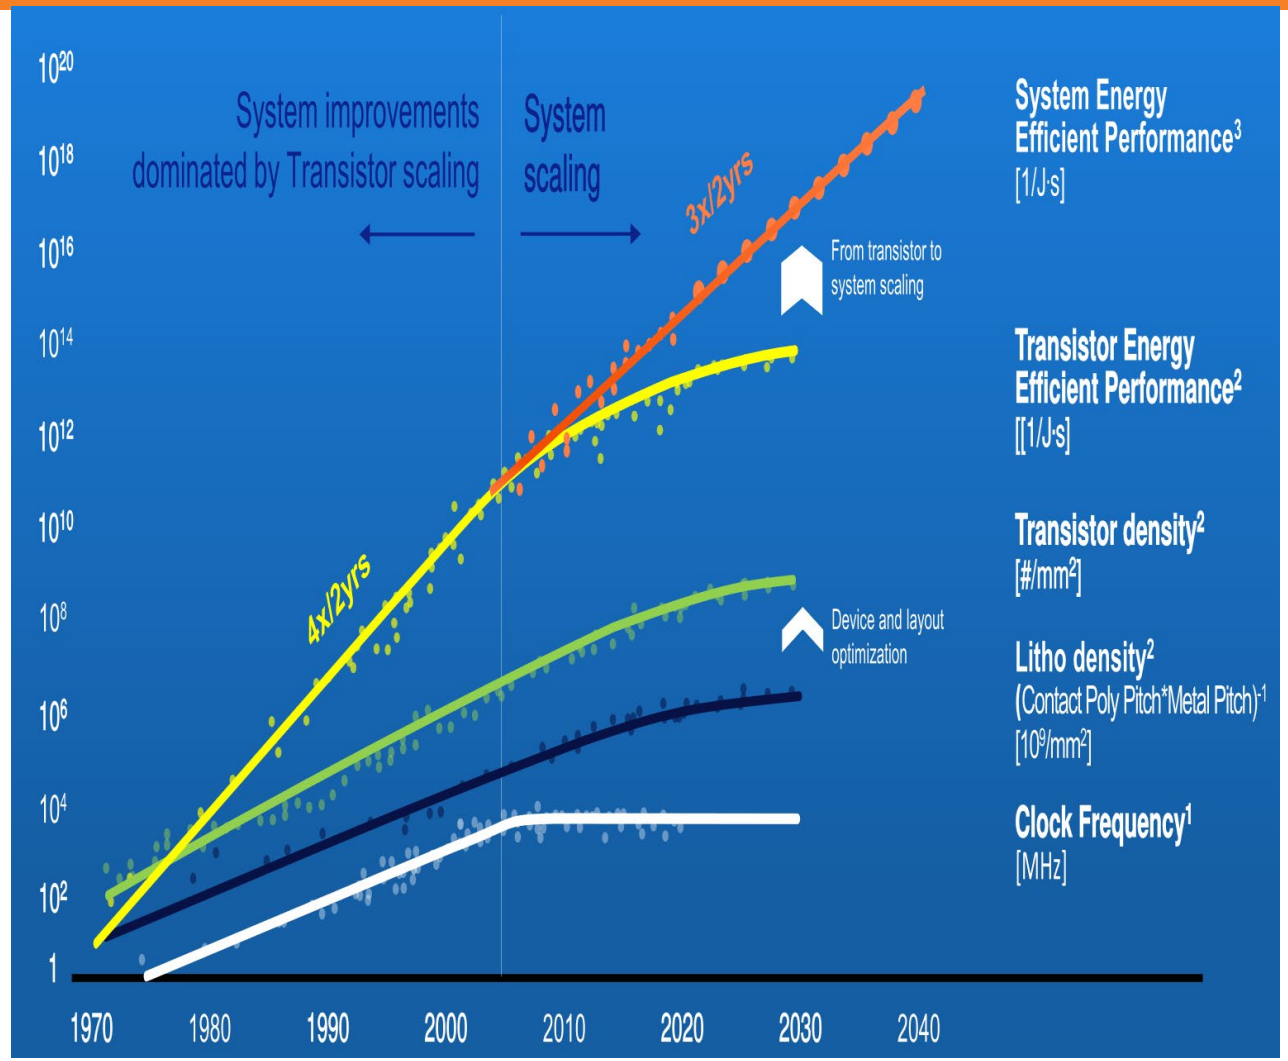

Source: Epoch AI, ASML, TSMC

Electronic-Photonic Integration is Coming

- Based on matured semiconductor technologies
  - Micron-scale nanosecond-level devices are widely demonstrated
- A choice across the industry for post-Moore's Law era
  - R&D: Intel, IBM, Acacia, Luxtera/Cisco, HPE, Oracle, Juniper, Finisar, Broadcom, Mellanox/Nvidia, ST, NTT, NEC, Fujitsu, Hamamatsu, Oclaro/Lumentum, NeoPhotonics, Ciena, Inphi, Infinera, Huawei …
  - Fab: TSMC、Intel、STMicroelectronics、GlobalFoundries、TowerJazz、AMS、SMIC、Samsung (2027)…
  - PEDA (photonic-electronic design automation): Cadence-PhoeniX-Lumerical, Siemens EDA-Lumerical-Luceda, Siemens EDA-VPIphotonics, PhoeniX/Synopsys, RSoft/Synopsys, Lumerical/Ansys …
  - Startup: Lightwire/Cisco, Kotura/Mellanox, Caliopa/Huawei, Aurrion/Juniper, Rockley, OneChip, Skorpios, Ayar, Sicoya, Elenion, Dust Photonics, AIO core, Lightelligence, Optalysys, Lightmatter, Fathom Computing, LightOn, Photoncounts …

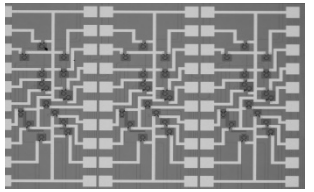

Integrated Optical Switches

R. Ji, *et al.* "Five-Port Optical Router Based on Microring Switches for Photonic Networks-on-Chip", IEEE Photonics Technology Letters 2013

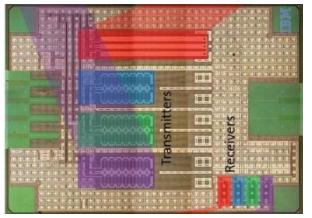

Integrated OE Interfaces

D.M. Gill, *et al.*, "Demonstration of Error Free Operation Up To 32 Gb/s From a CMOS Integrated Monolithic Nano-Photonic Transmitter", CLEO 2015

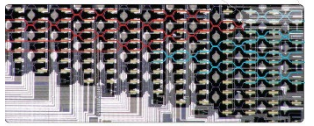

Optical Neural Network

Y. Shen *et al.*, "Deep learning with coherent nanophotonic circuits", Nature 2017

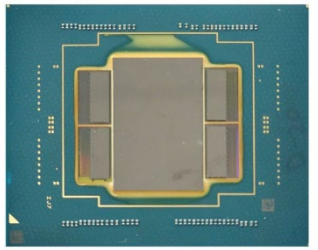

Intel 528 Thread Processor with Photonic Network

Hot Chips 2023

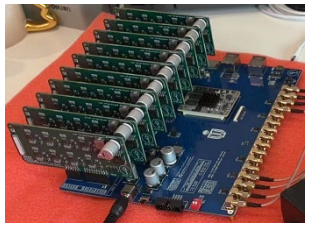

Optical Neural Network Full HW/SW System

Our Studies

- Silicon photonic devices
- Optical/electrical interconnects and OE interfaces
- Integrated optical switches
- High-radix optical switching fabrics for data center
- Optical network for memory system
- Optical neural computing
- Optical network for disaggregated computing
- Challenges
  - Crosstalk noise analysis
  - Thermal effects and analysis
  - Process variation
  - Nonlinear effects and optical power delivery
- Photonic-electronic design automation tools

Basic Optical Switching Element (BOSE)

- An abstraction for topology/architecture designs
  - Same for both Mach-Zehnder interferometer (MZI) and microresonator (MR)
  - Omit many details and keep only necessary characteristics
- Optical devices are still in the range of microns

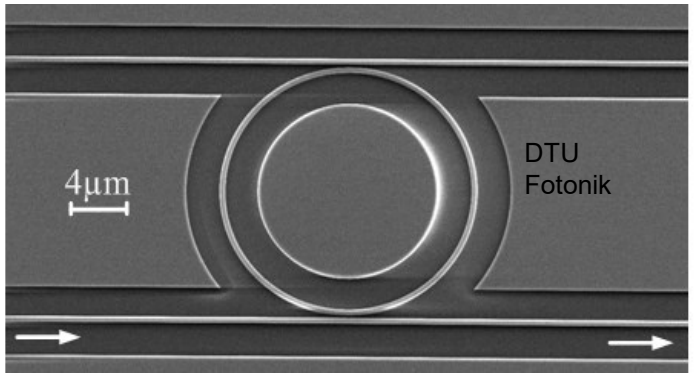

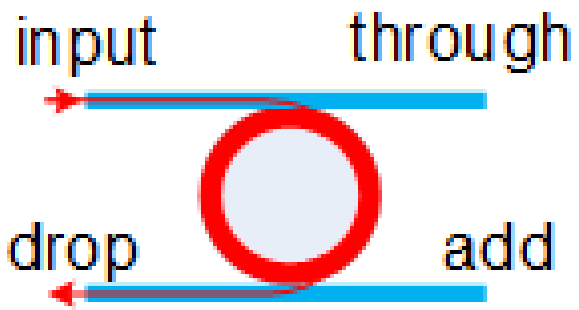

on-state

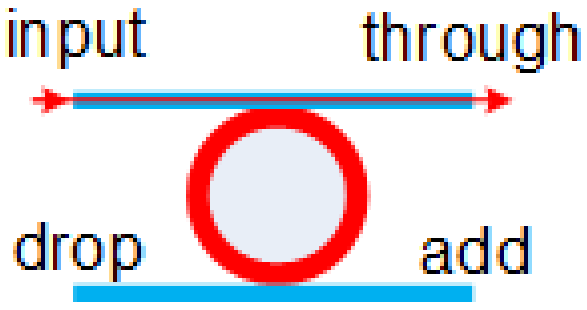

off-state

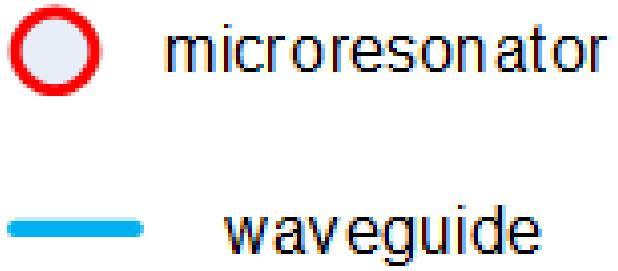

Optimized 5x5 Optical Crossbar

- For mesh and torus based optical networks
  - Ports are aligned to intended directions
  - No U-turns
- 20 BOSEs and 10 terminators
- Many waveguide crossings
  - Cause large loss in a network, on average 1.08dB loss per switch
- Could we do better?
  - Reduce optical power loss
  - Use less BOSEs and terminators

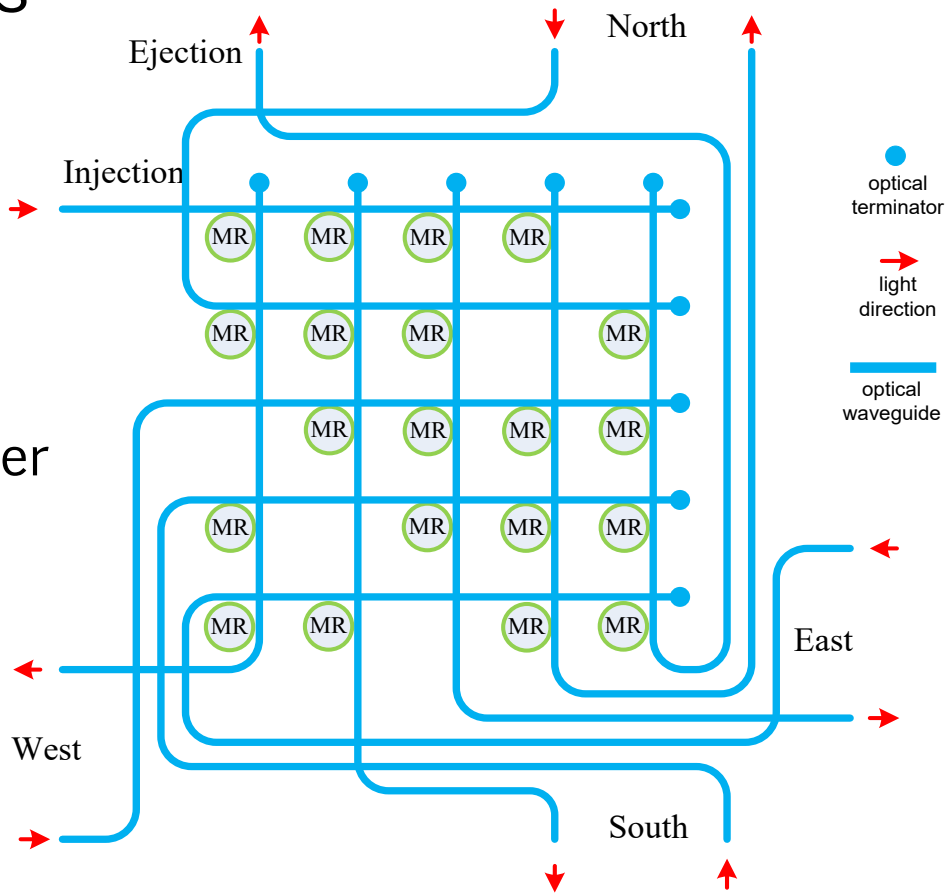

Cygnus Optical Switch

- 5x5 switching function
- Strictly non-blocking
  - Support any routing algorithm
- Ports are aligned to intended directions
- 16 BOSEs and 2 terminators
- Special feature
  - Passively switching between ports in the same dimension

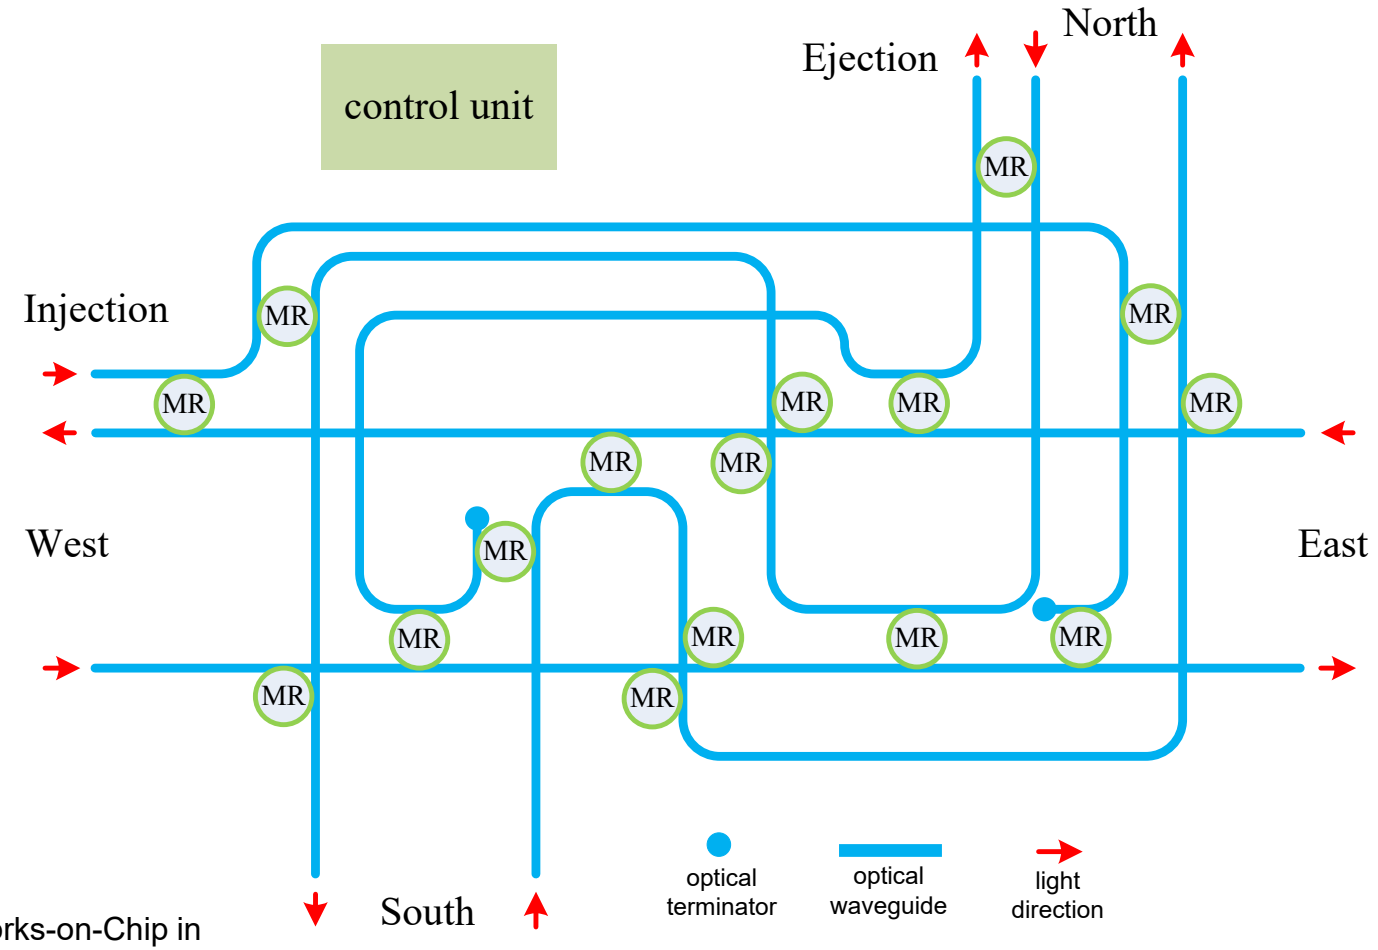

*Huaxi Gu, et al, "A Low-power Low-cost Optical Router for Optical Networks-on-Chip in Multiprocessor Systems-on-Chip", ISVLSI 2009

Implementation and Testing

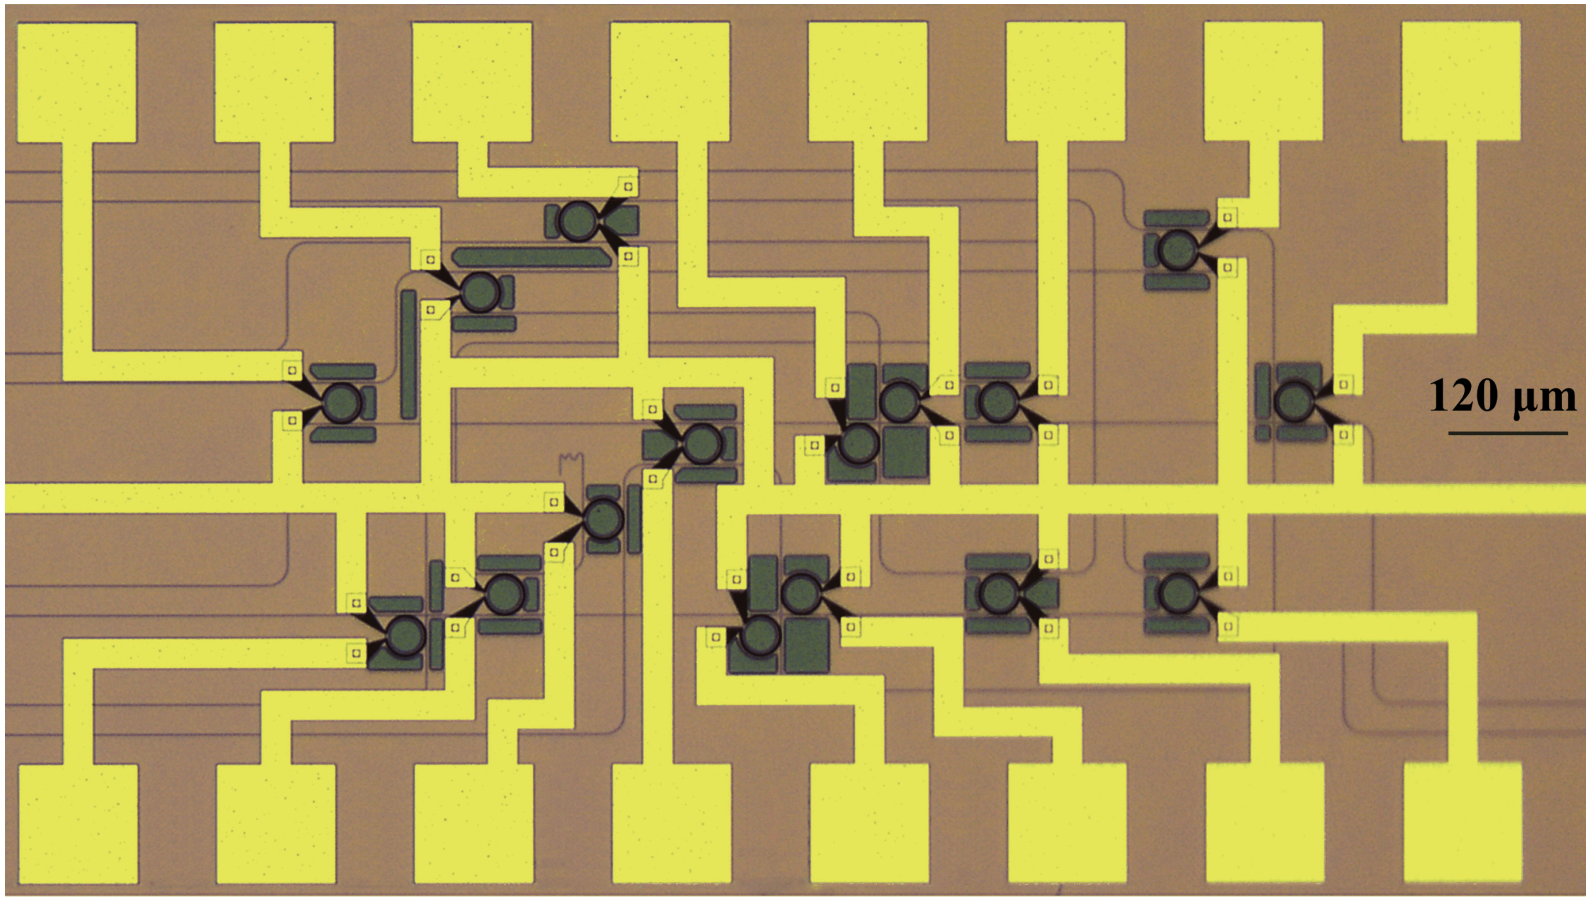

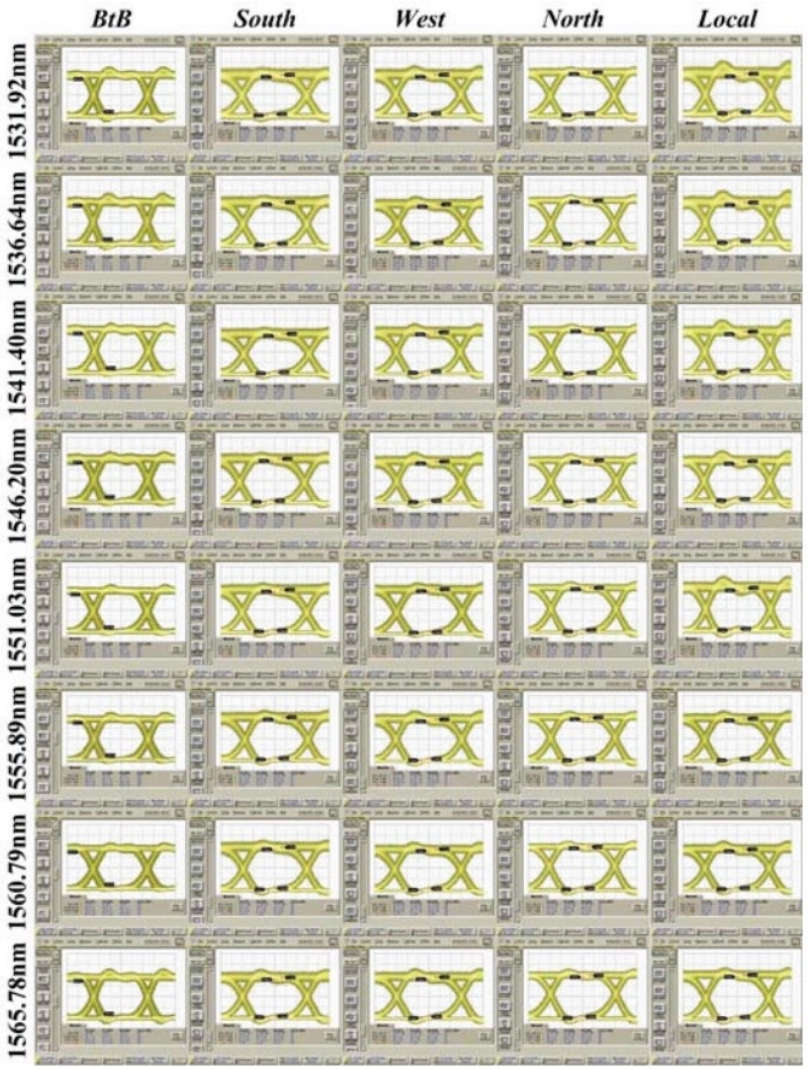

*R. Ji, *et al.* "Five-Port Optical Router Based on Microring Switches for Photonic Networks-on-Chip", IEEE Photonics Technology Letters, March 2013

Strictly Non-blocking 6x6 and 7x7 Switches

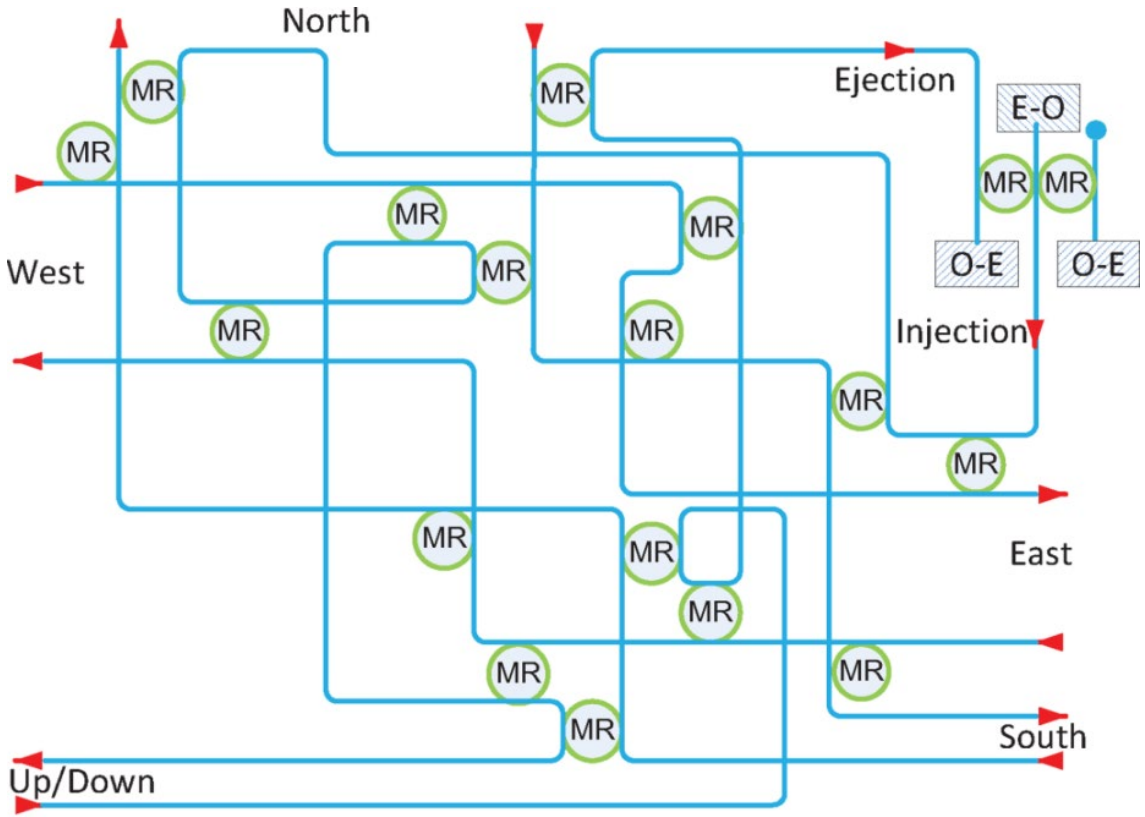

6x6 optical switch

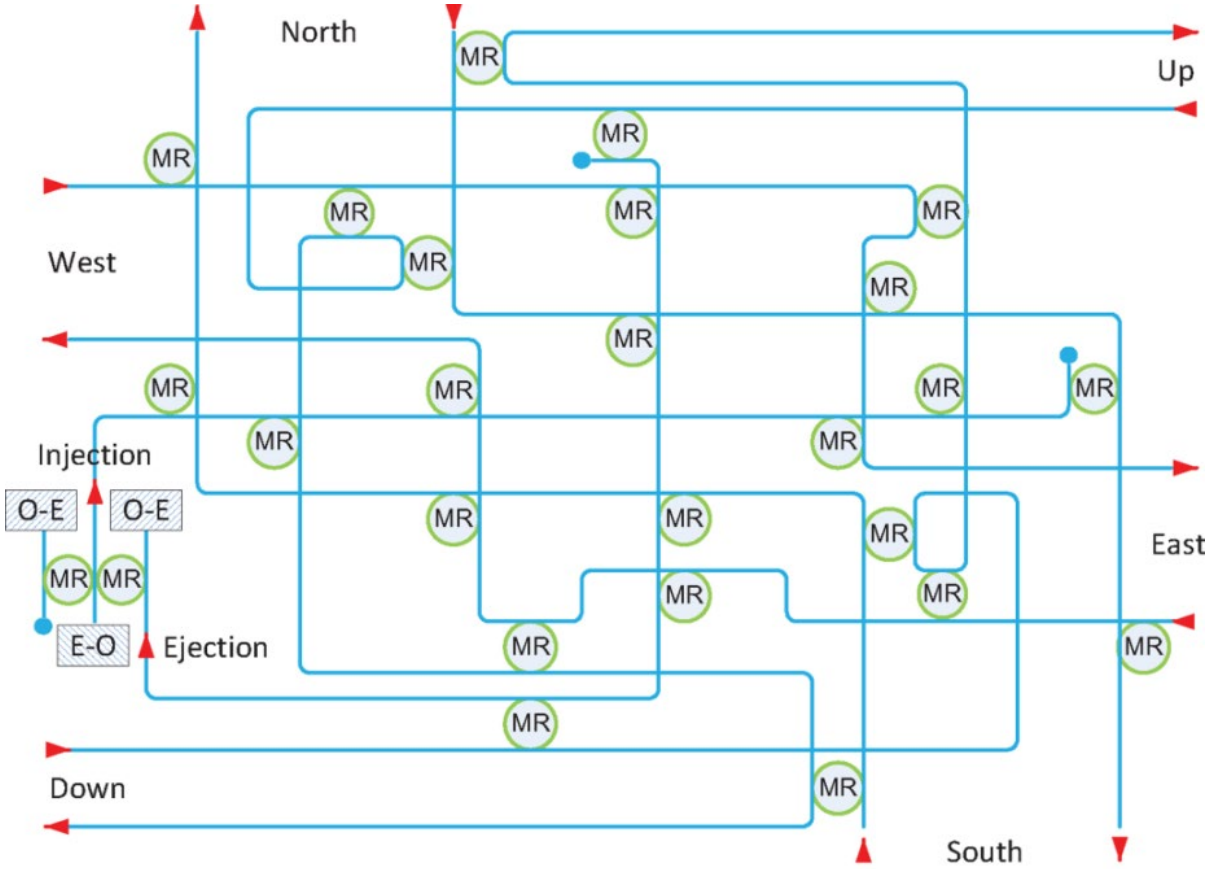

7x7 optical switch

Low-Radix Optical Switch Comparison

	Cygnus	Crux	OXY	ODOR	[14]	[15]	[16]	[17]
Routing algorithm	Arbitrary	XY routing	XY routing	XY routing	Arbitrary	Arbitrary	Arbitrary	XY routing
Nonblocking	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
Passive routing	Yes	Yes	Yes	Yes	No	No	No	Yes
N_{MR}	16	12	12	12	20	16	15	12
$N_{Terminator}$	2	2	2	5	10	2	0	2
$N_{Crossing}$	13	9	11	19	26	14	15	12
$L_{average}$	0.78dB	0.64dB	0.73dB	0.87dB	1.15dB	0.87dB	0.77dB	0.96dB

[14] A. Poon, *et al.* "Cascaded microresonator-based matrix switch for silicon on-chip optical interconnection," *IEEE Proceedings, 2009.*
[15] R. Ji, *et al.* "Five-port optical router for photonic networks-on-chip," *Optical Express,* Oct 2011.
[16] R. Min, *et al.* "Scalable non-blocking optical routers for photonic networks-on-chip," *IEEE Optical Interconnects Conference*, 2012.
[17] J. Chan, *et al.* "Physical-layer modeling and system-level design of chip-scale photonic interconnection networks," *IEEE TCAD,* 2011.

High-Radix Optical Switching Fabric

- Data center demand
  - High radix
  - High bandwidth
  - Low latency
  - Low power
  - Low cost
- Very challenging for electrical switches
- Optical switches are promising

- Many have been proposed
  - Benes
  - Fat-tree
  - Butterfly
  - Baseline
  - Clos
  - Crossbar
  - AWG
  - …
- High insertion loss and cost

DRAGON: High Radix Optical Switch Fabric

- Dual Radial and Angular Grating Optical Network
- Dual radial and angular topology
- Formally approved strictly non-blocking property
- Scalable to over 100 ports

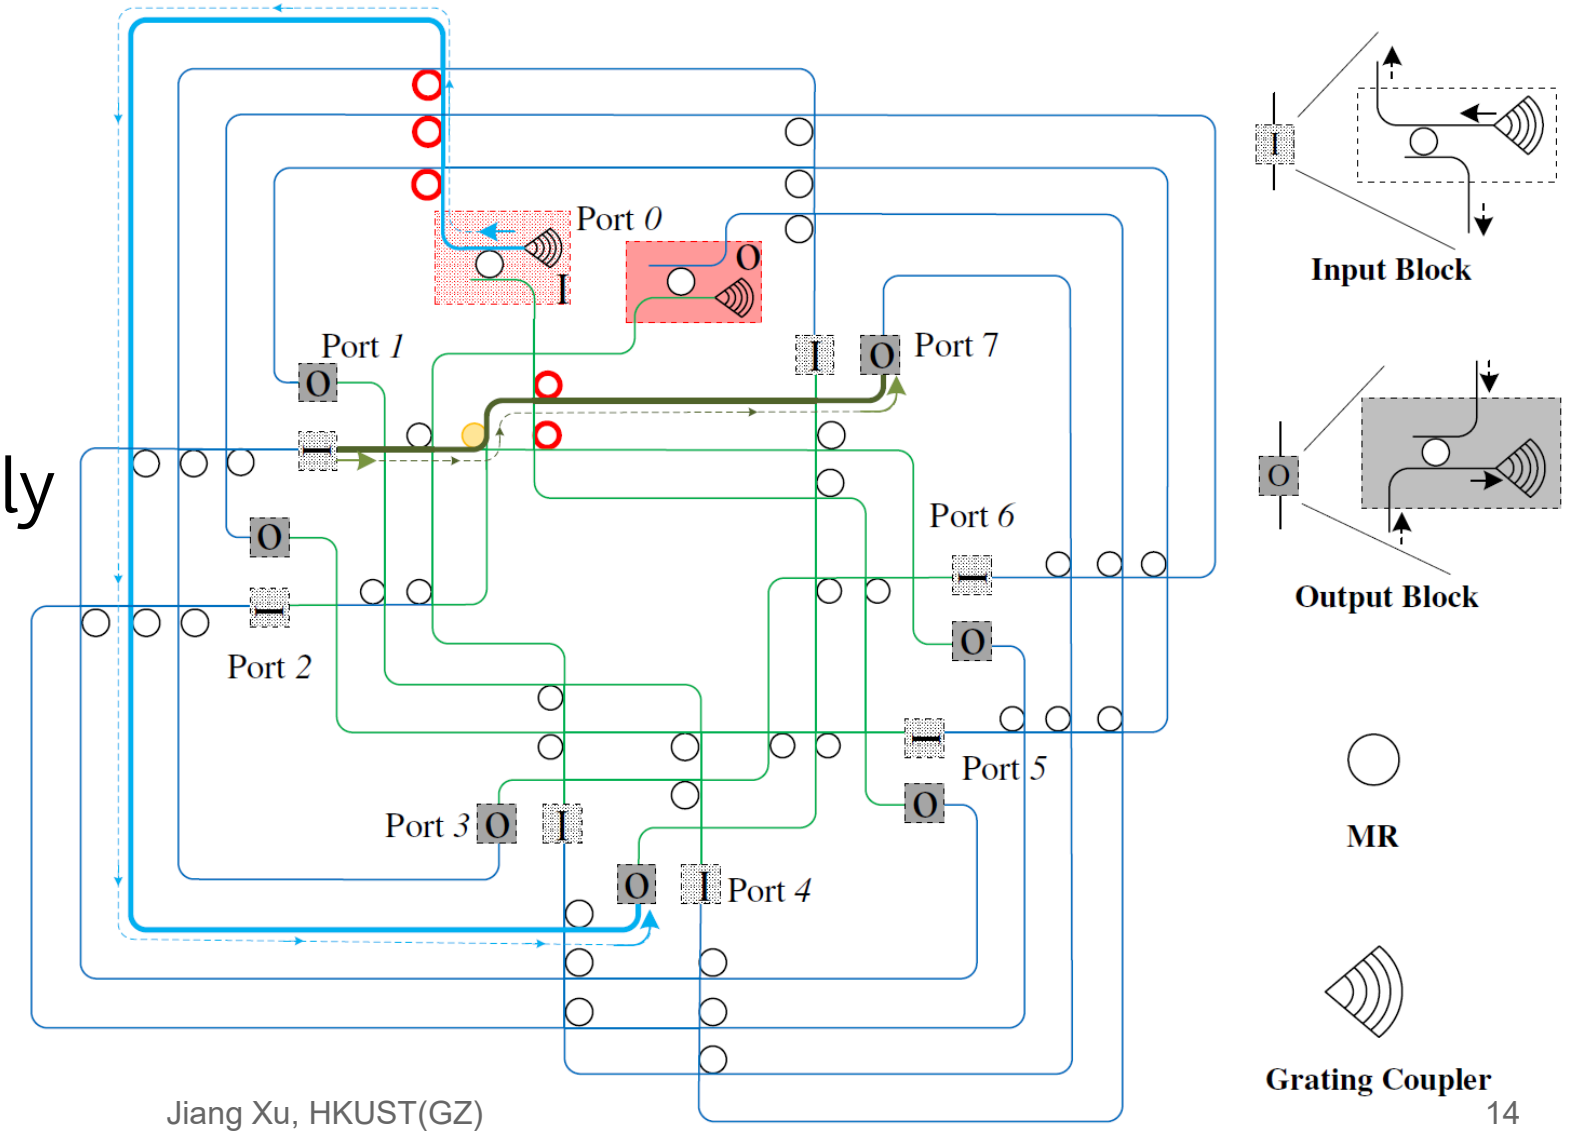

*Z. Wang, *et al.* "High-Radix Non-blocking Integrated Optical Switching Fabric for Data Center," *IEEE/OSA Journal of Lightwave Technology*, 2017.

Worst-Case and Average Loss

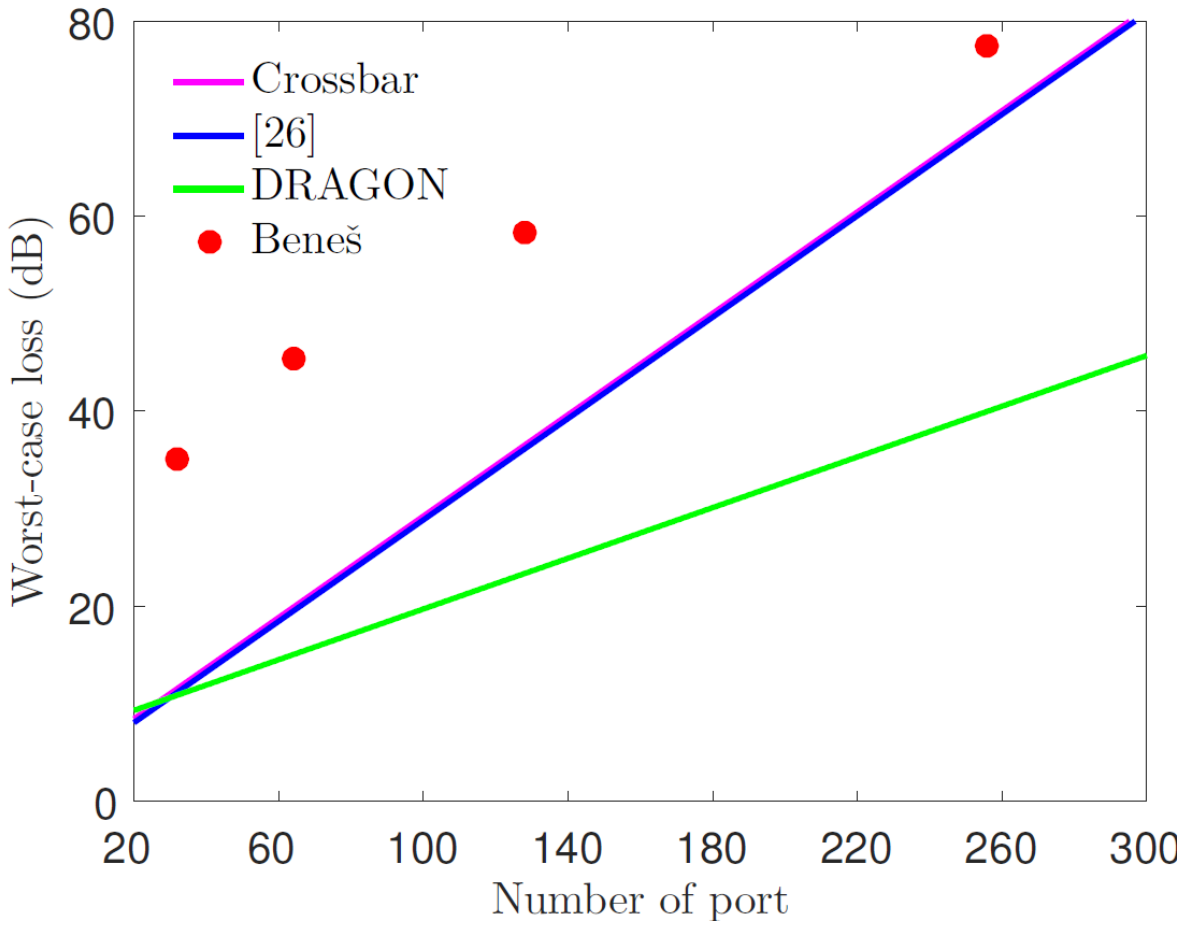

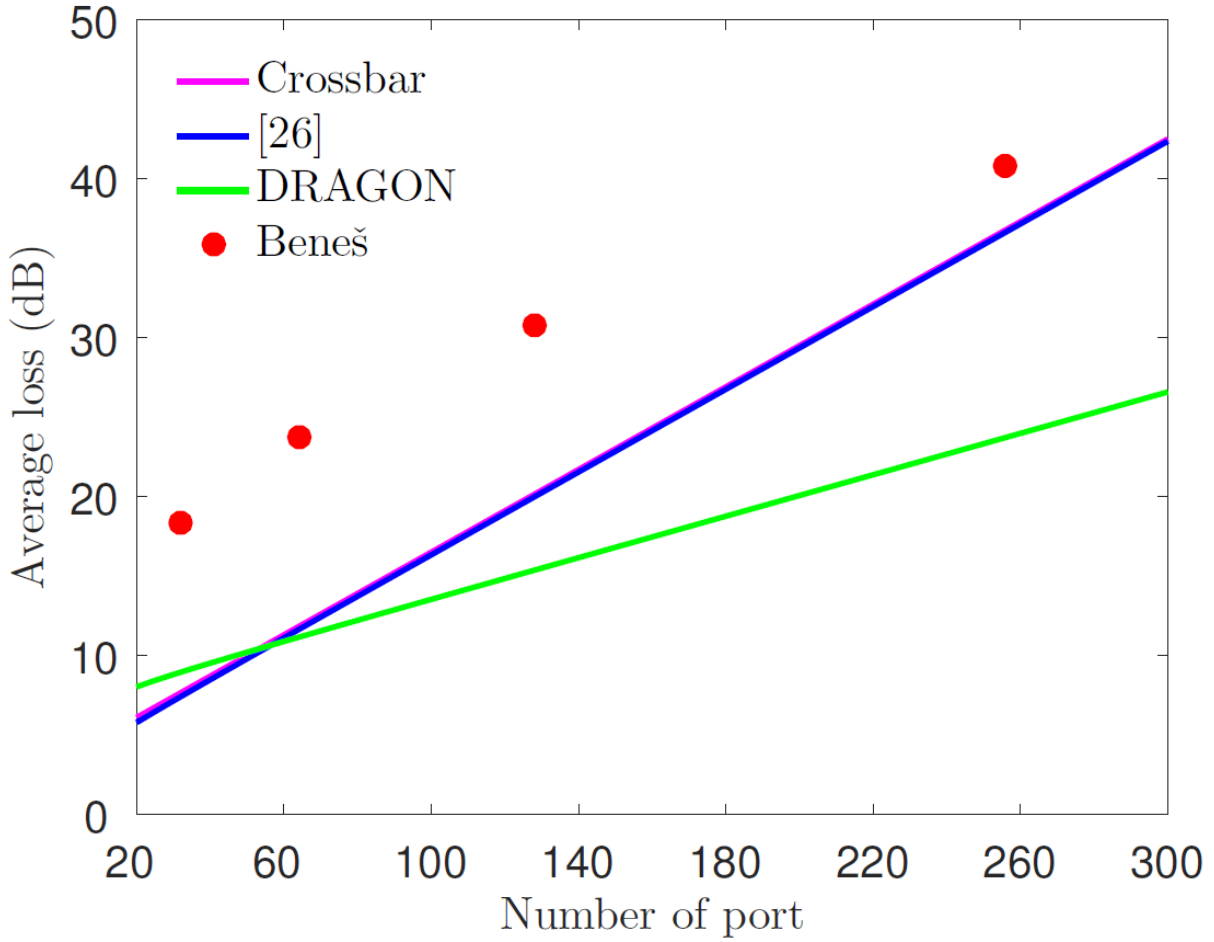

180-port DRAGON has 30dB and 19dB worst-case and average loss

Optical Crosstalk

- Average OSNR of 180x180-port DRAGON is 18dB
- Worst-case and average OSNR of 180x180-port DRAGON are 23dB better than [26]

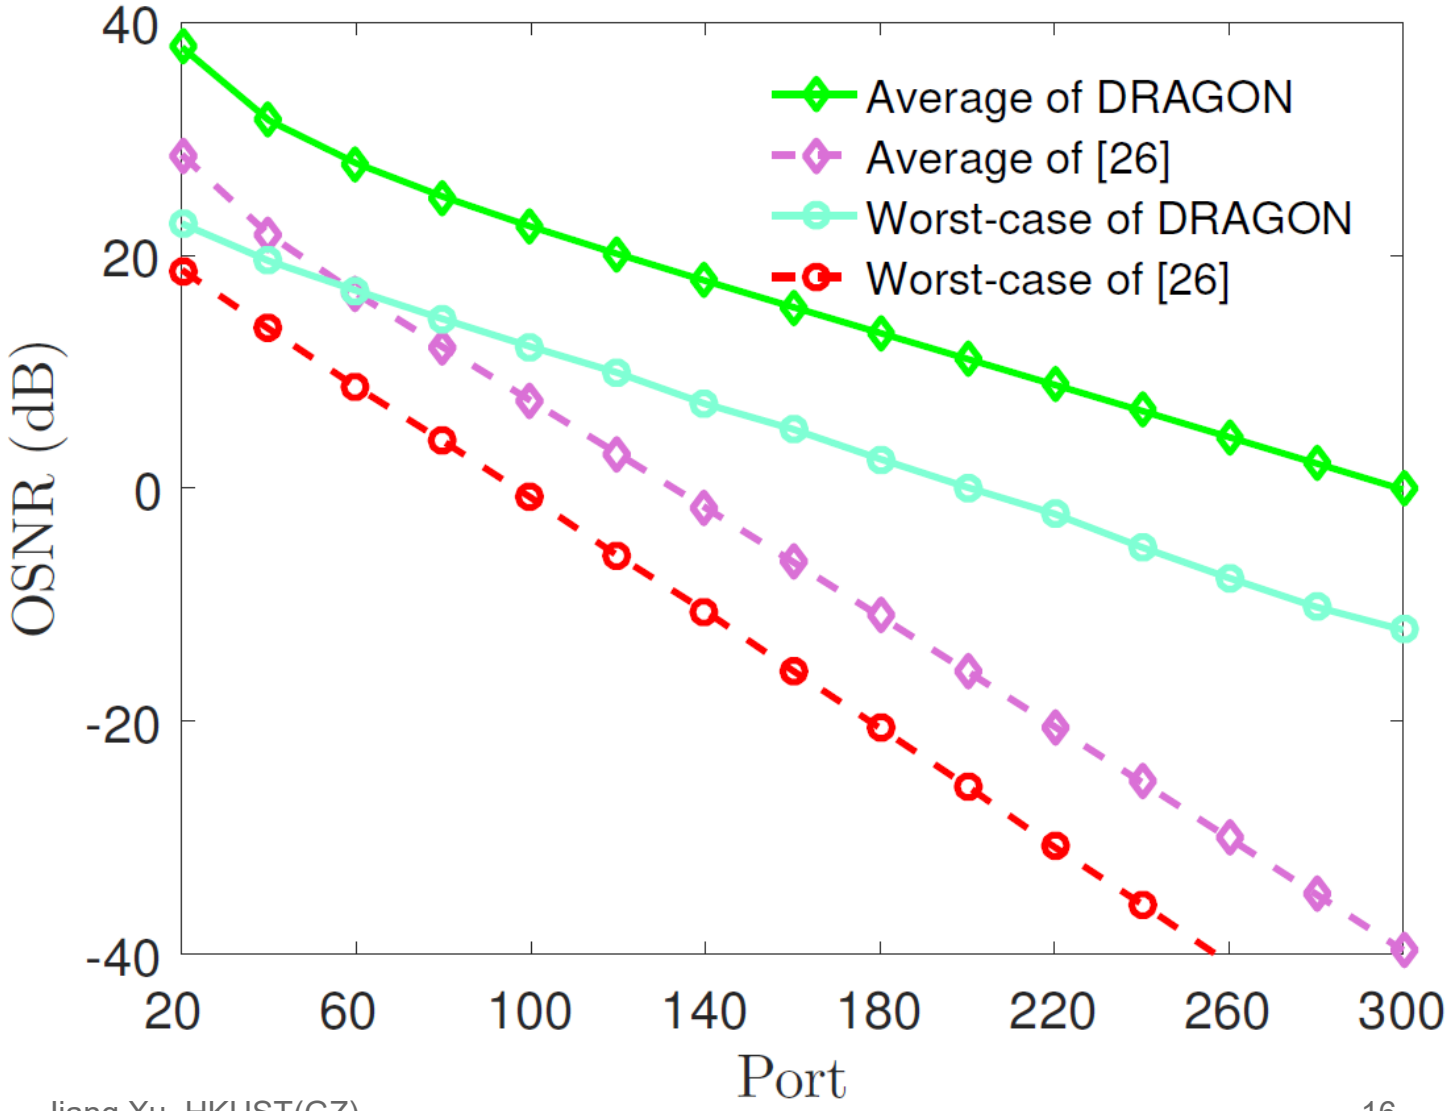

Fabrication Sensitivity Analysis

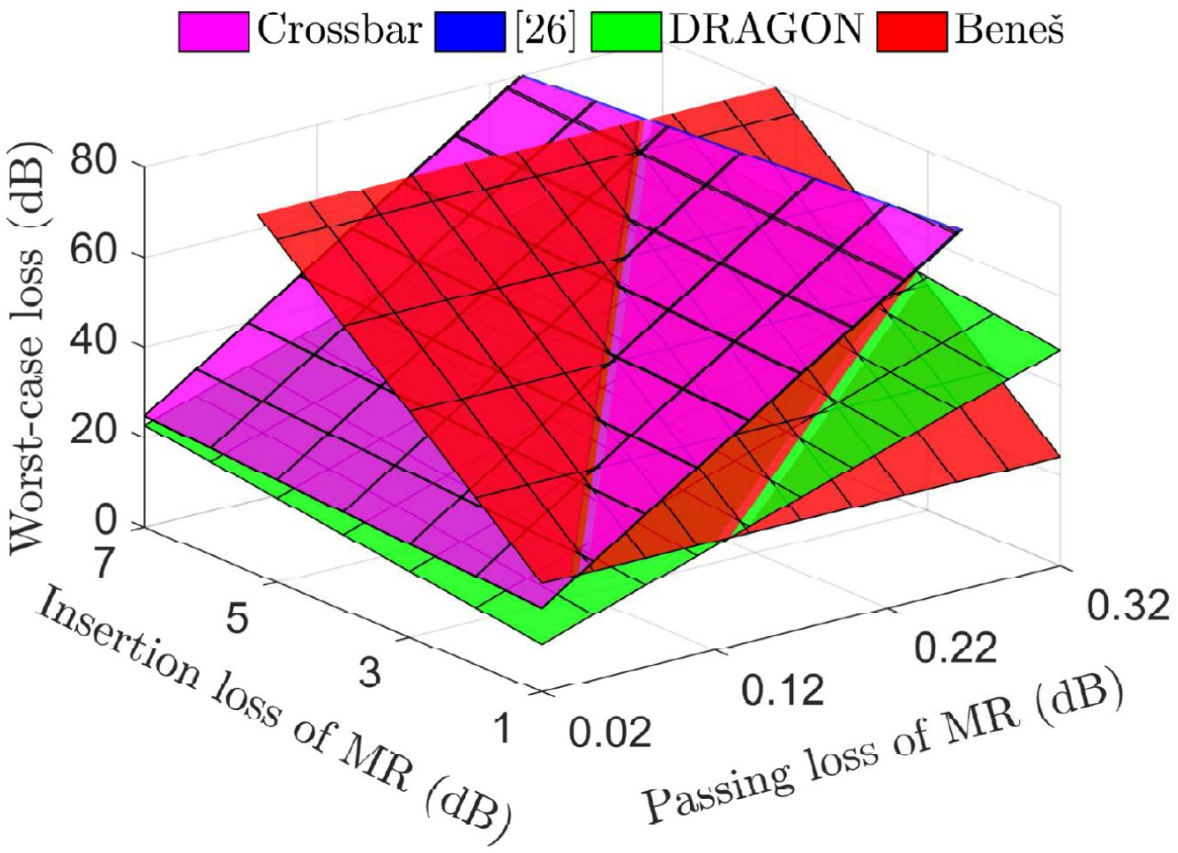

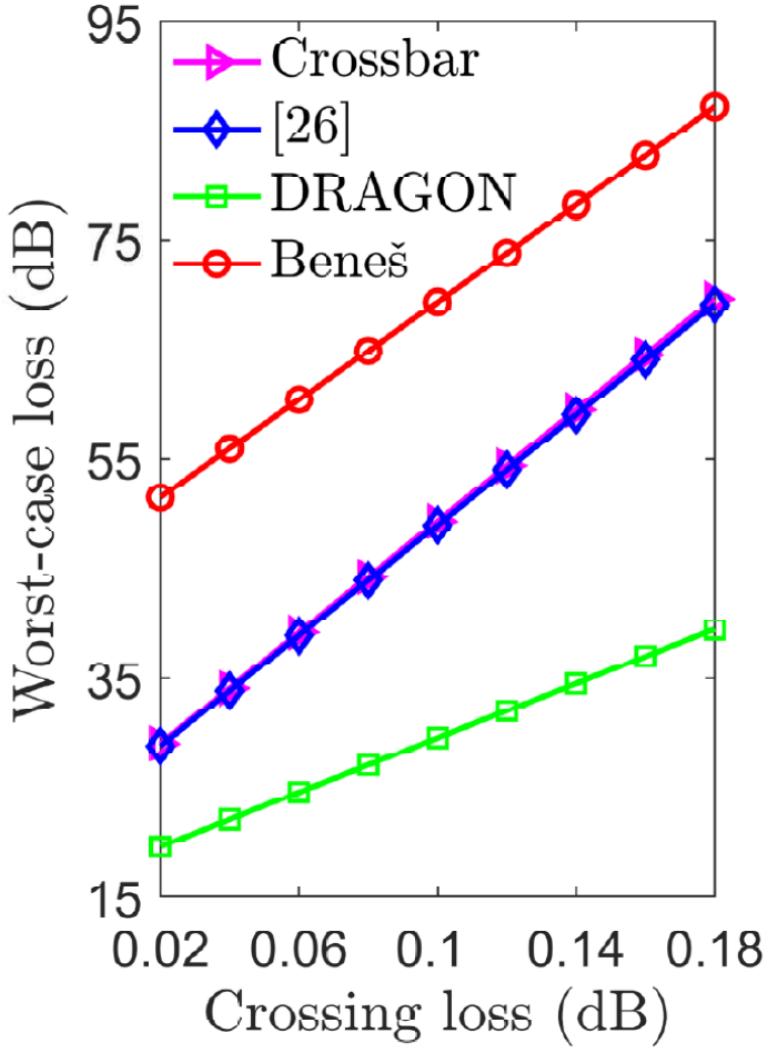

Throughput and Latency Comparison

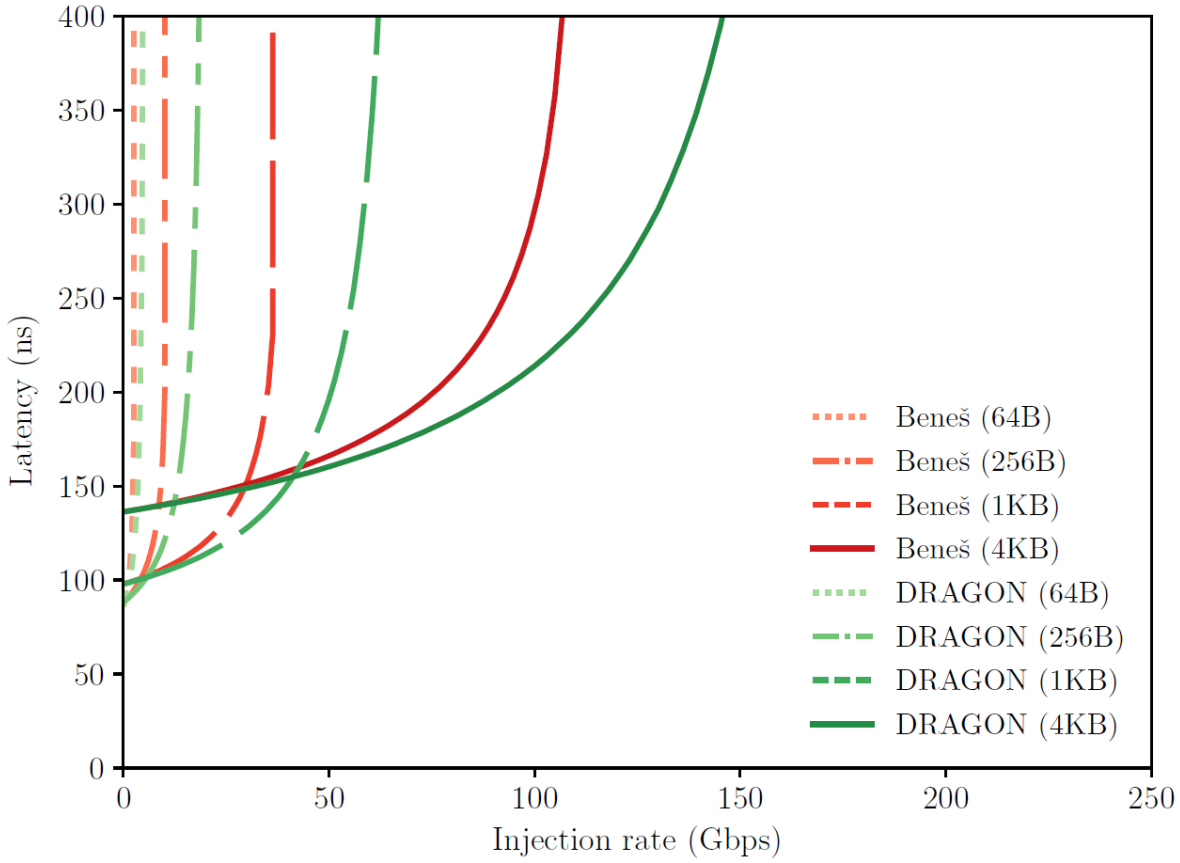

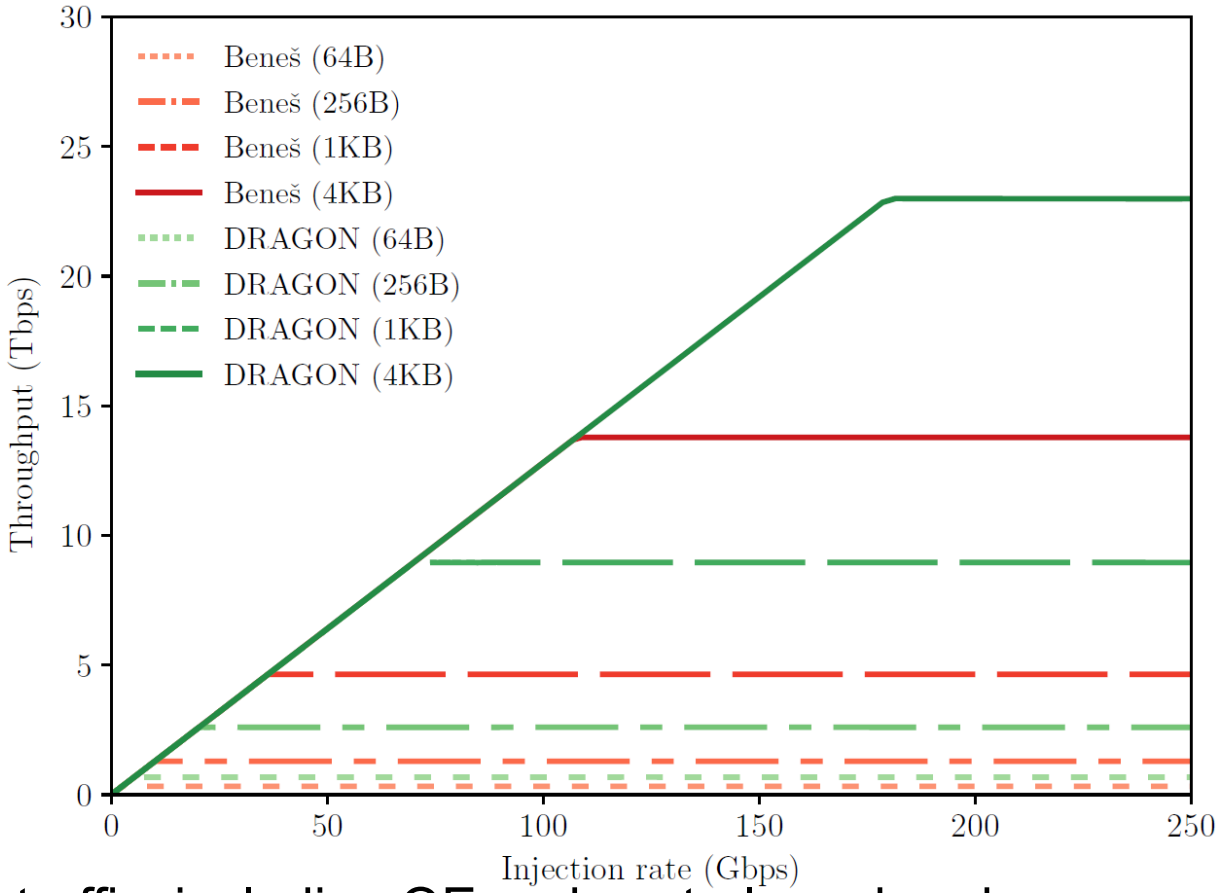

*3m fiber (15ns) on each port, 128-port, uniform traffic, including OE and control overheads

Energy Efficiency and Cost

- DRAGON has the highest energy efficiency
  - 128-port, Including OE, control, and turning overheads
- Low cost
  - 4N more MRs than [26]

ENERGY CONSUMPTION (UNIT: pJ/bit)

Packet size	64 B	256 B	1 KB	4 KB
Beneš	133	127	126	126
Crossbar	12.9	10.4	9.9	10.3
DRAGON	8.2	5.6	5.1	5.5

HARDWARE COST COMPARISON

Switching fabric	Optical pins	3-serial-coupled MRs	Crossings
DRAGON	$2N$	$4N(N-1)$	$N(N-2)$
Crossbar	$2N$	$4N^2$	N^2
[26]	$2N$	$4N(N-2)$	$N(N-2)$
Beneš	$2N$	$2^{n+1}(2n-1)^\star$	$2^{n-1}(2^n+n-2)^\star$

* $n = \lceil \log_2 N \rceil$

FODON: Ultra-High Radix Optical Switch Fabric

- Floorplan Optimized Delta Optical Network
- Optimized stage switch radix
- Optimized stage switch design
- Optimized floorplan to minimize waveguide crossing

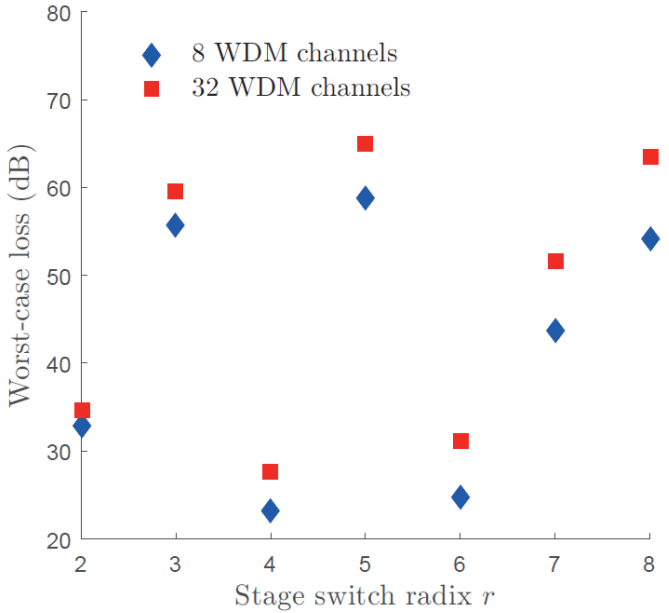

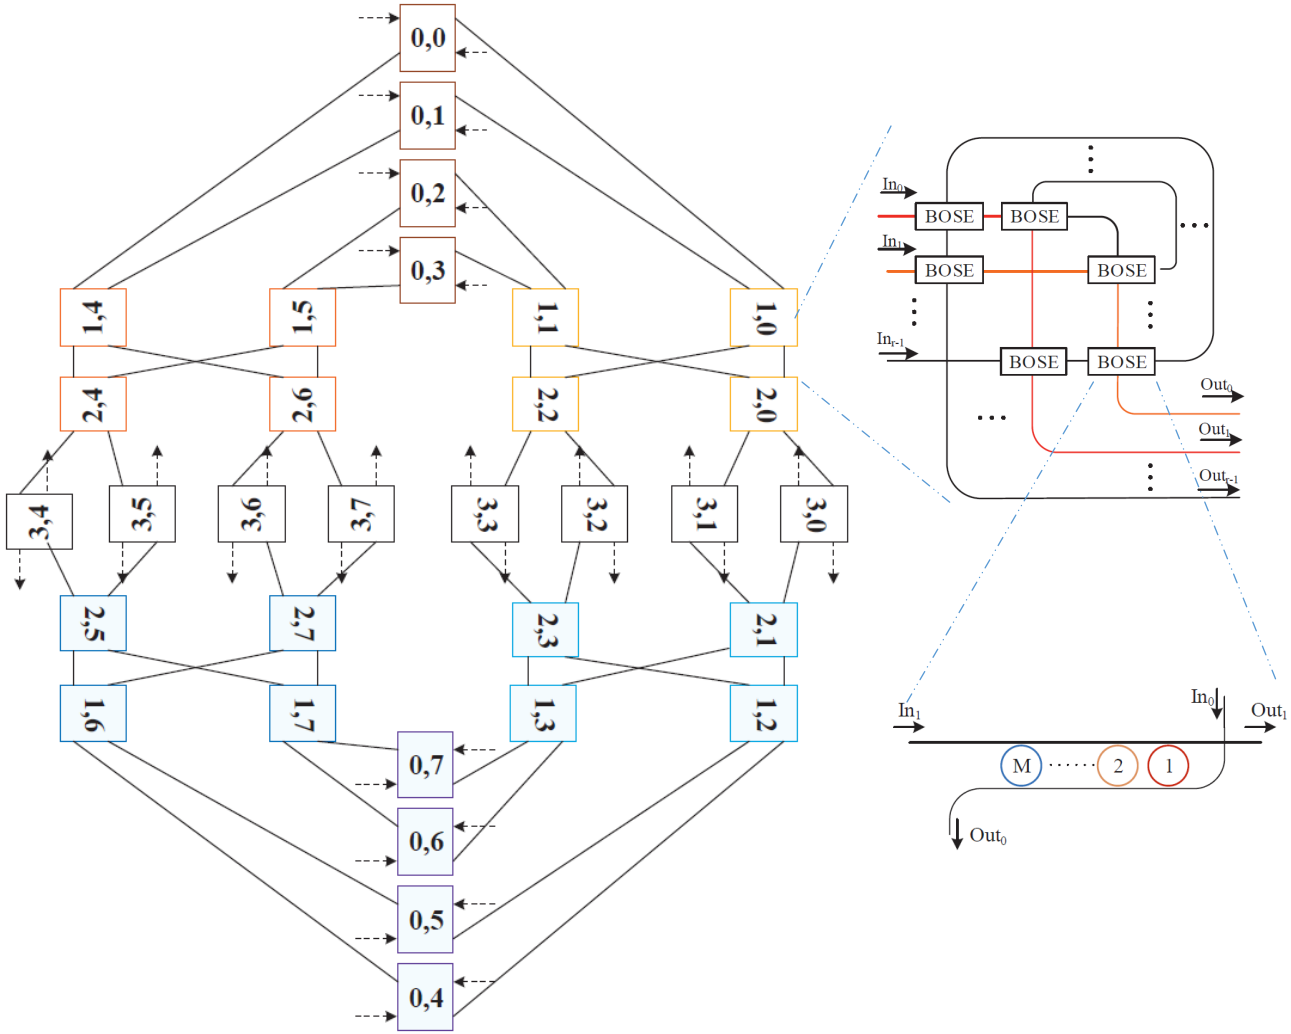

*Z. Wang, *et al.* "Low-Loss High-Radix Integrated Optical Switch Networks for Software-Defined Servers," *IEEE/OSA Journal of Lightwave Technology*, 2016.

Worst-Case and Average Loss

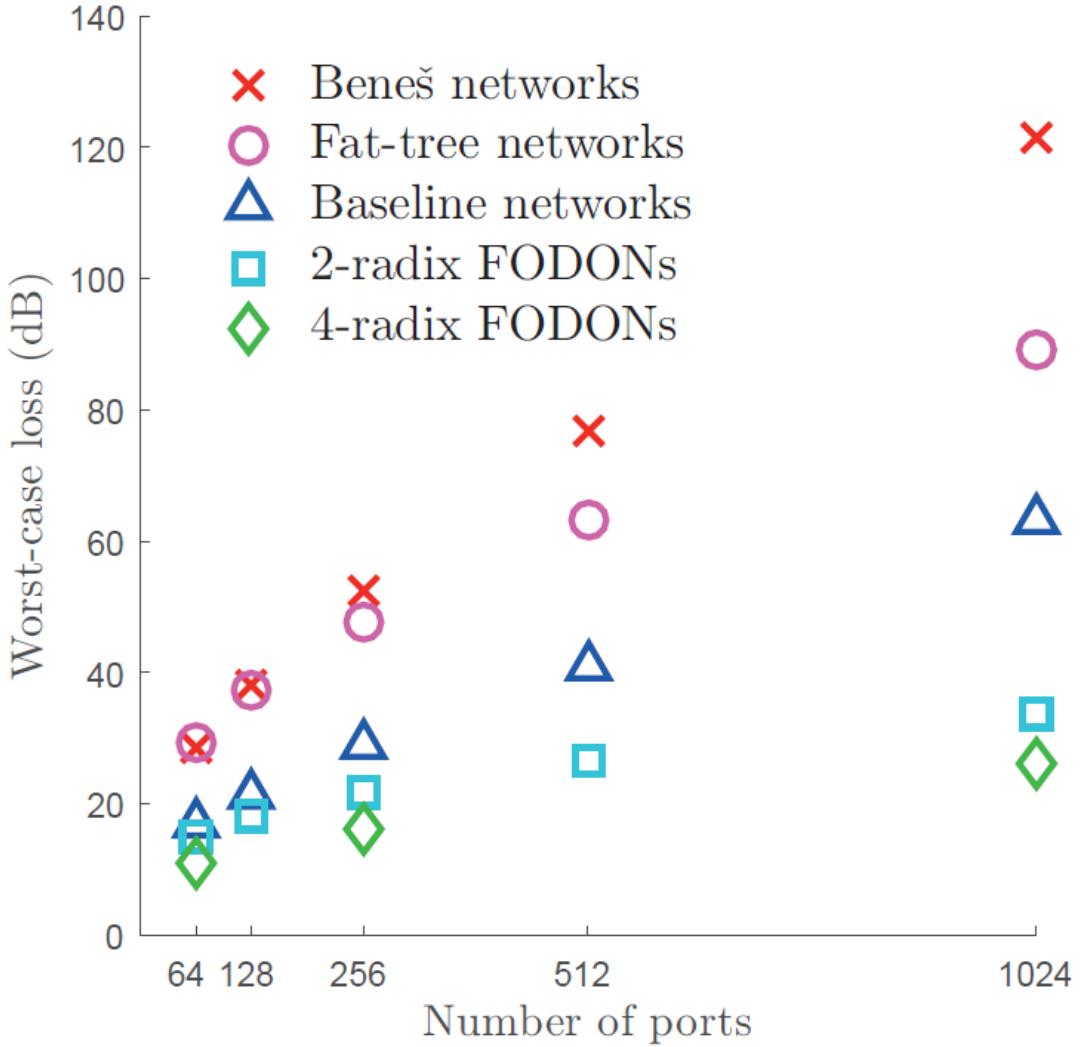

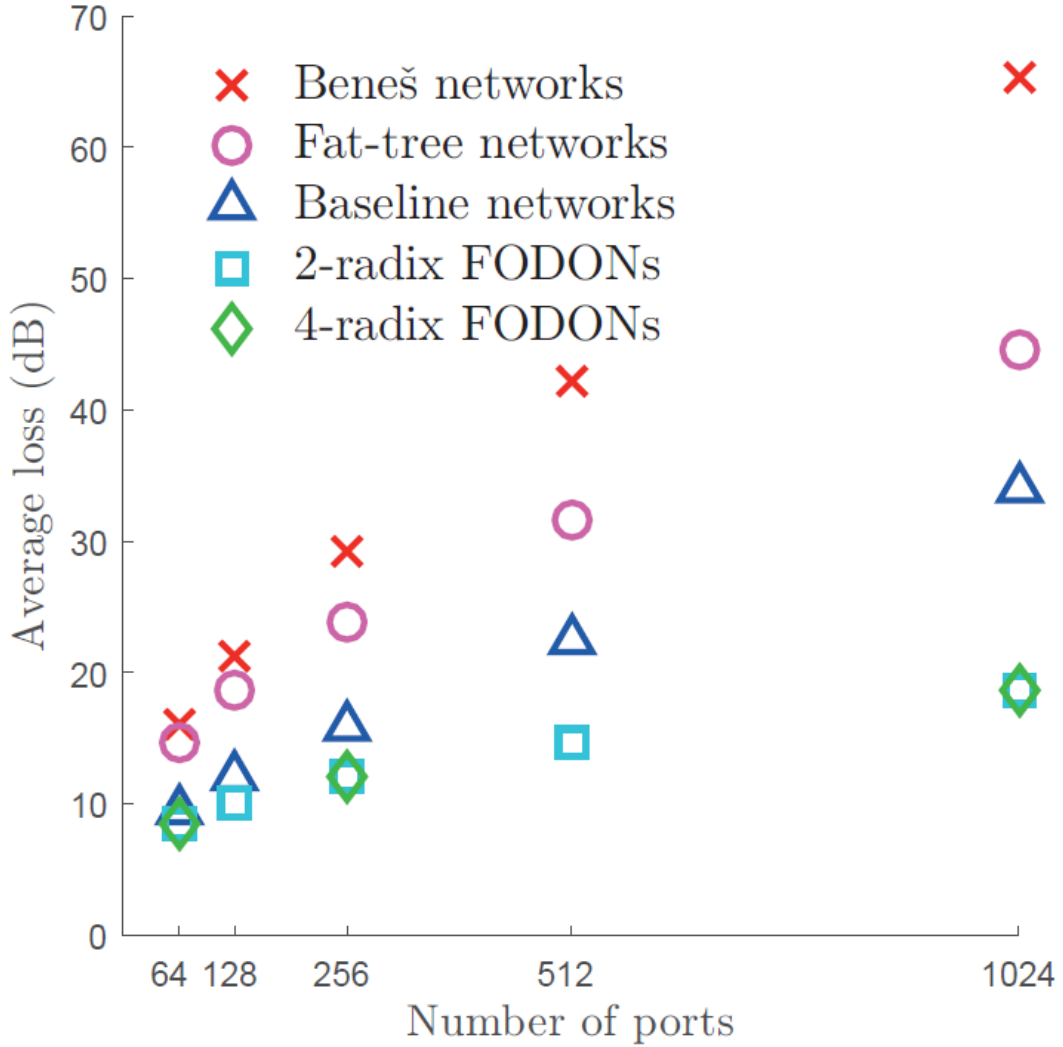

Cost

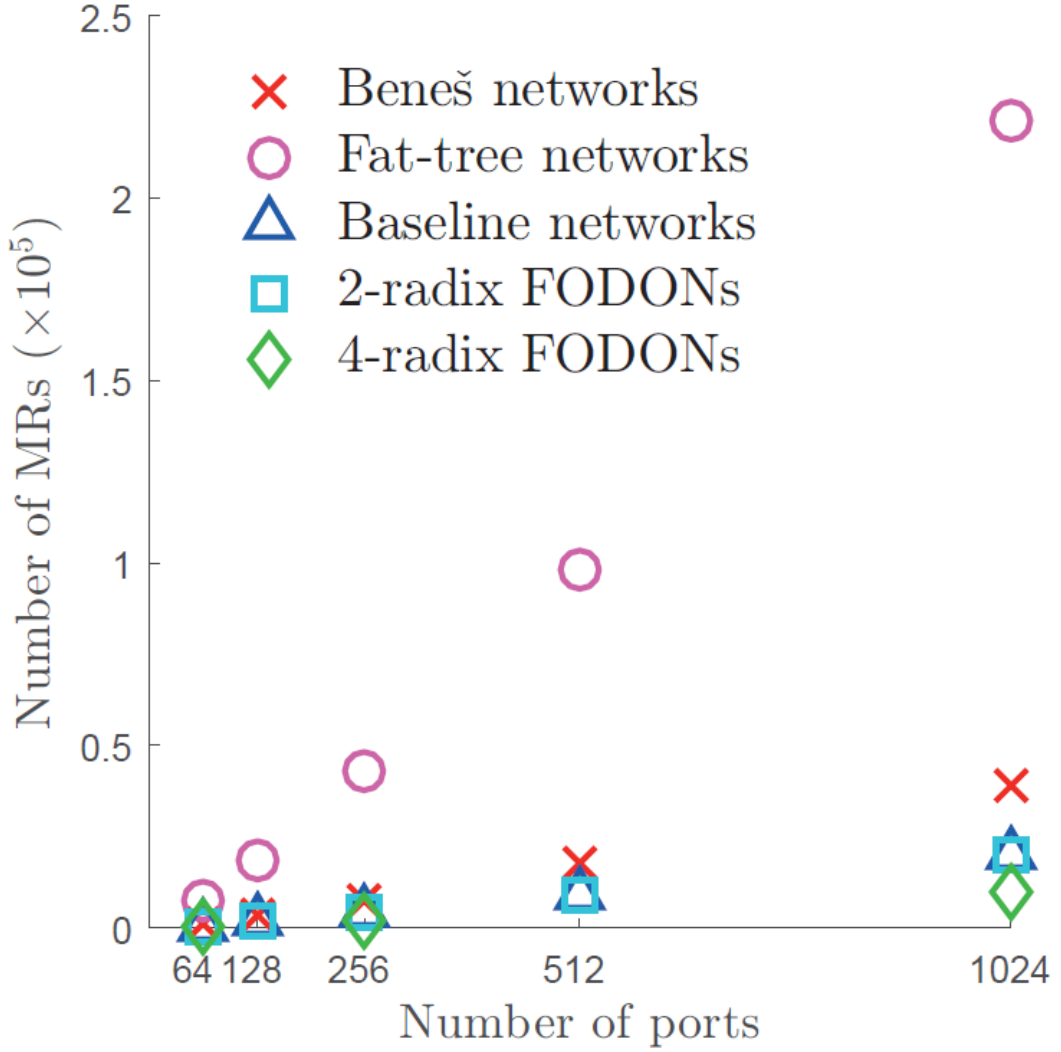
Beneš networks
Fat-tree networks
Baseline networks
2-radix FODONs
4-radix FODONs
Number of MRs ($\times 10^5$)
2.5
2
1.5
1
0.5
0
64 128 256 512 1024
Number of ports

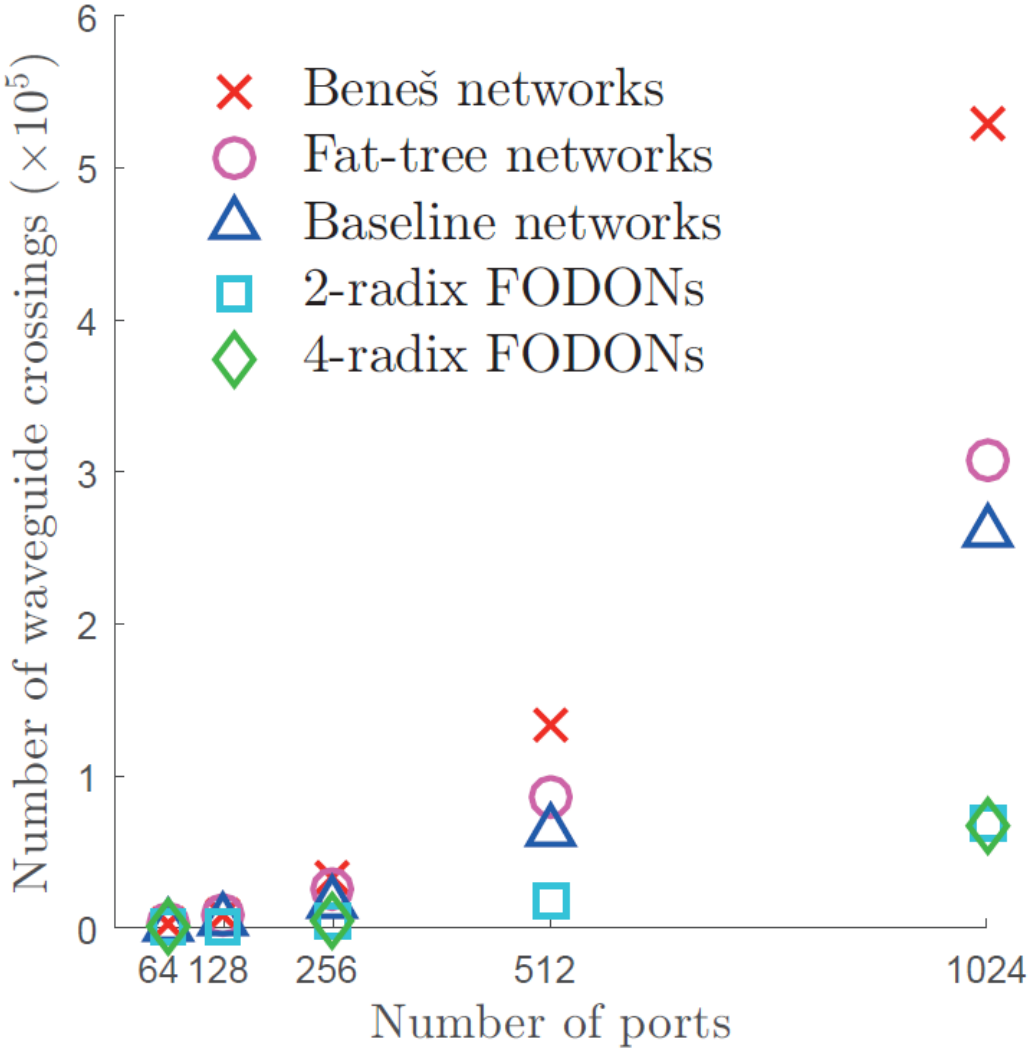
Beneš networks
Fat-tree networks
Baseline networks
2-radix FODONs
4-radix FODONs
Number of waveguide crossings ($\times 10^5$)
6
5
4
3
2
1
0
64 128 256 512 1024
Number of ports

Fabrication Sensitivity Analysis

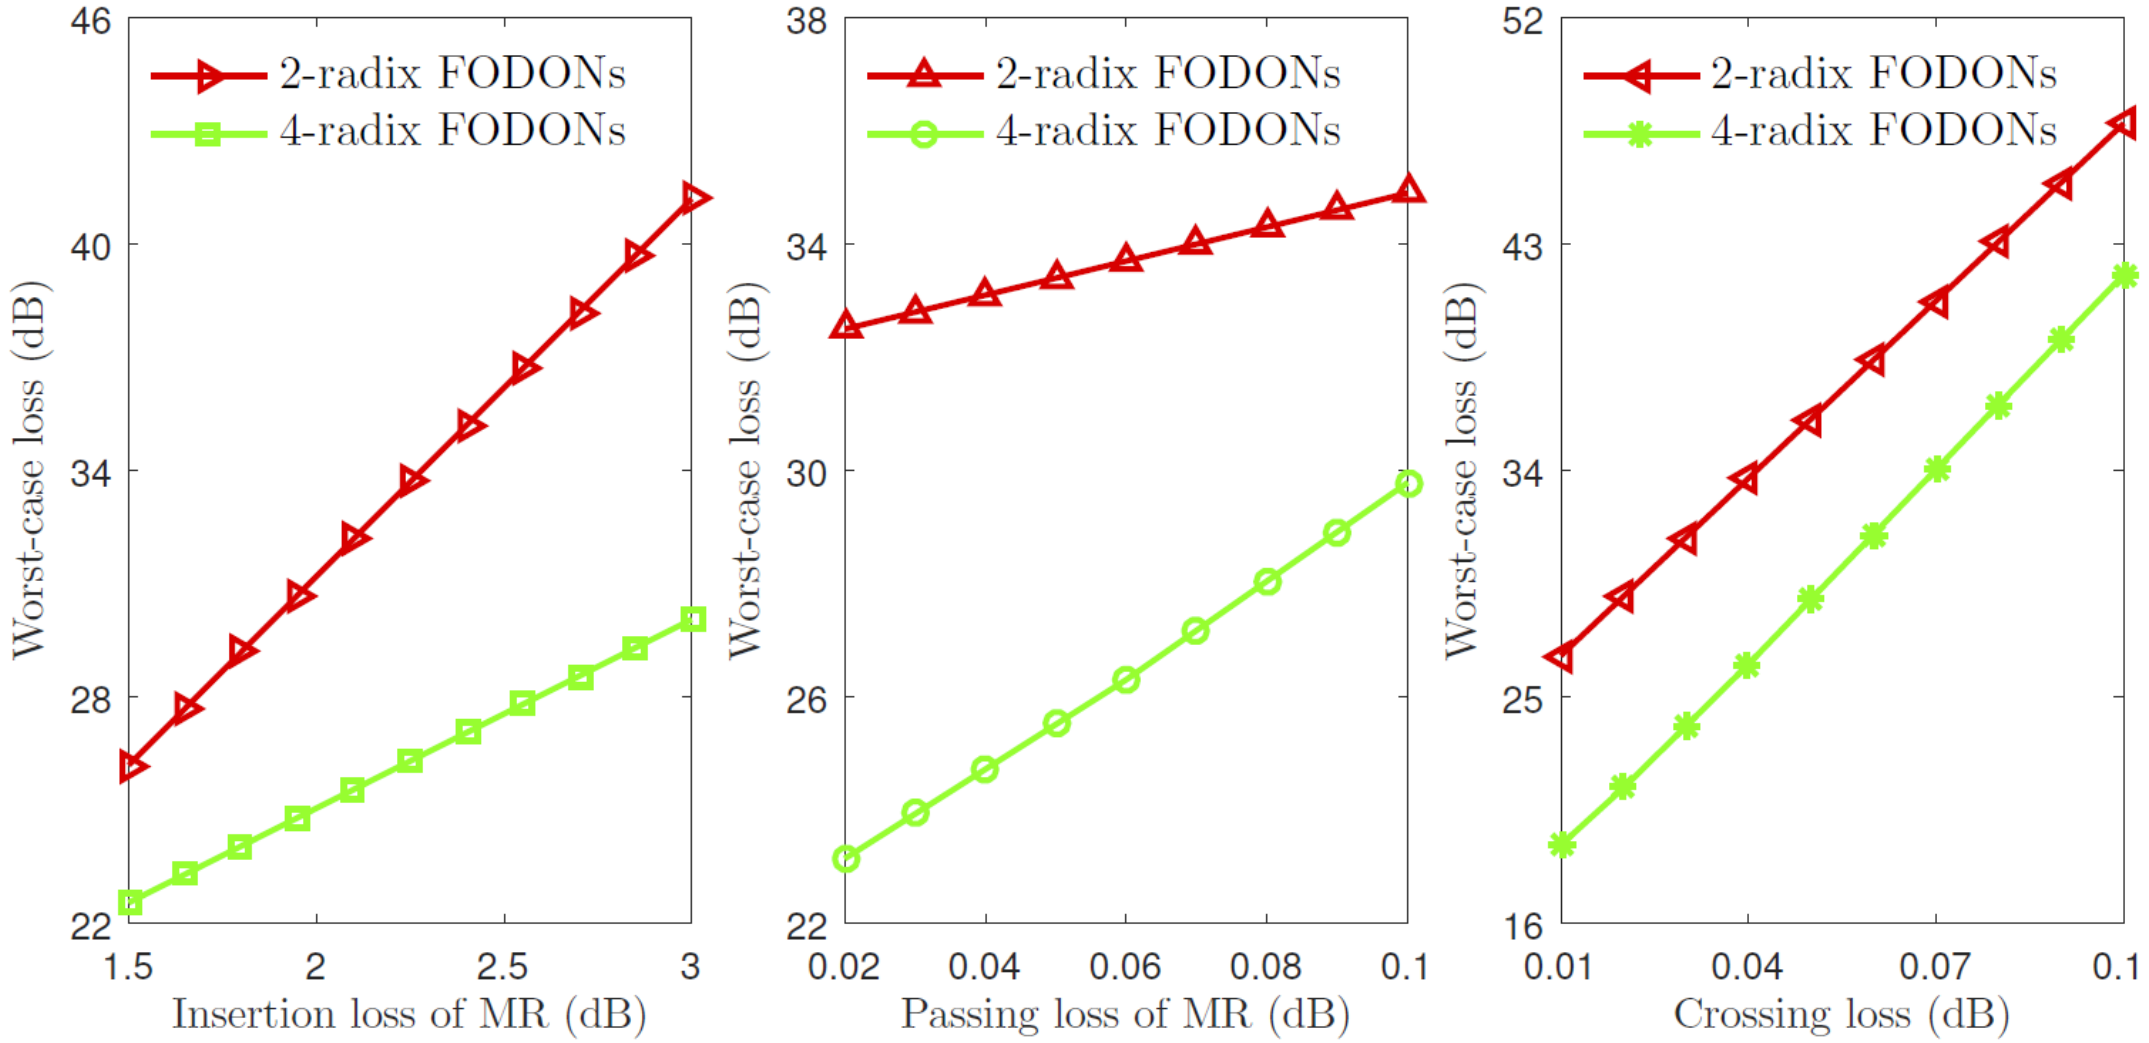

Switch Performance

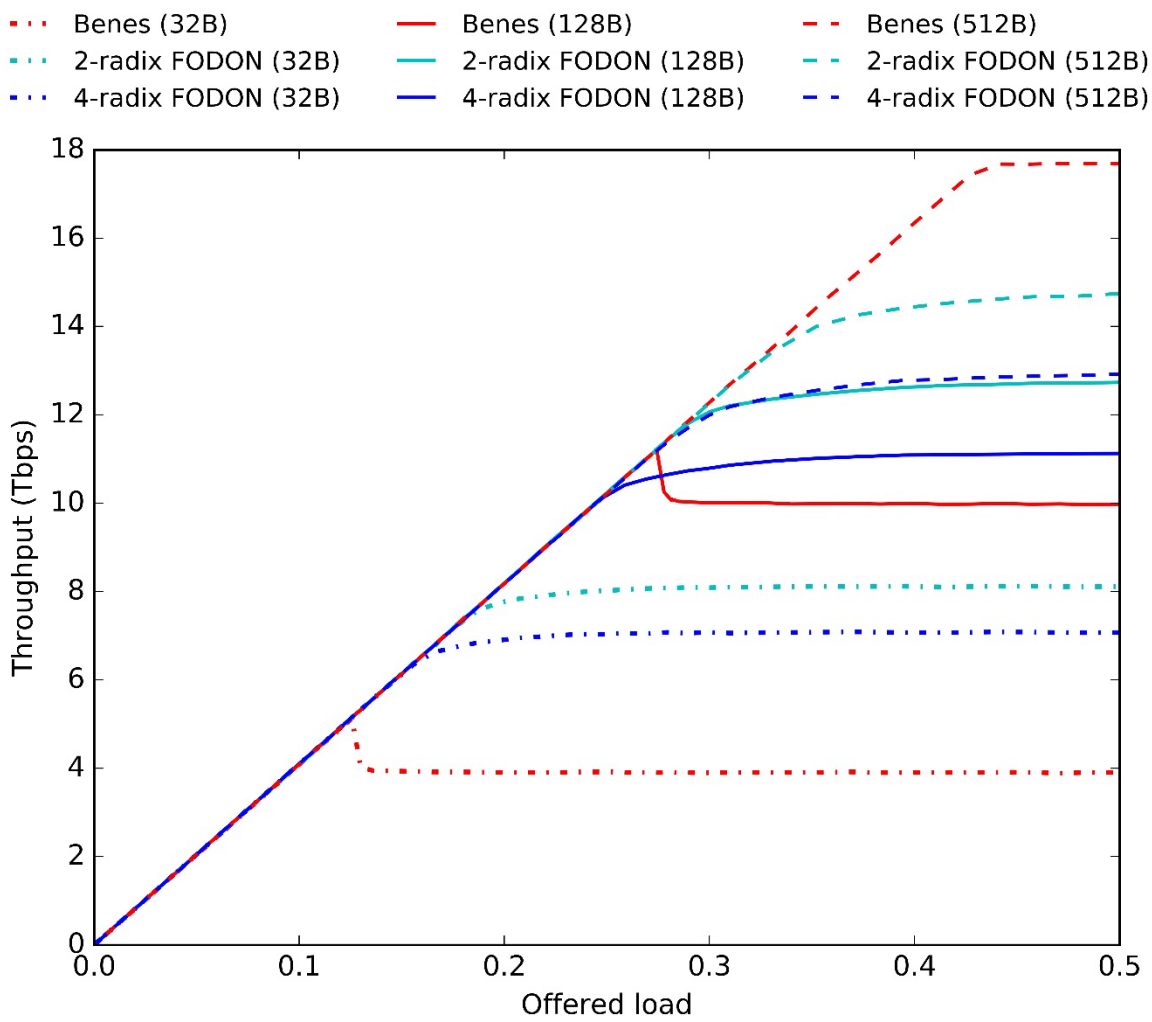

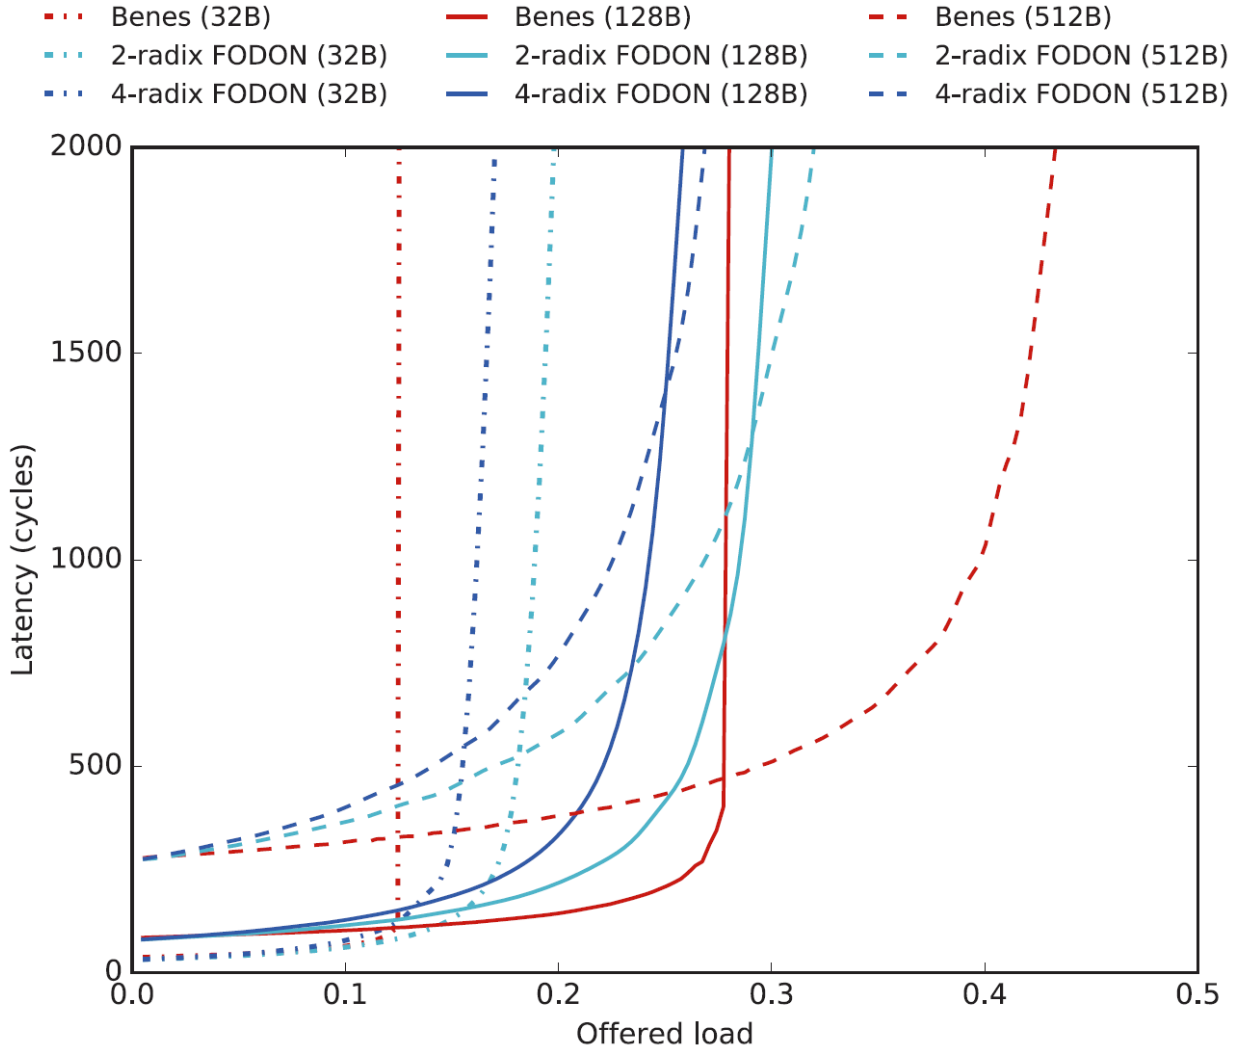

Manycore Processor with Photonic Chiplet

- Try to minimize changes to current processor designs
- Photonic devices are on separate chiplets
  - Different from logic and memory dies
- Lasers are off-chip
  - Easy to replace
  - Share among chips
  - Low power modes

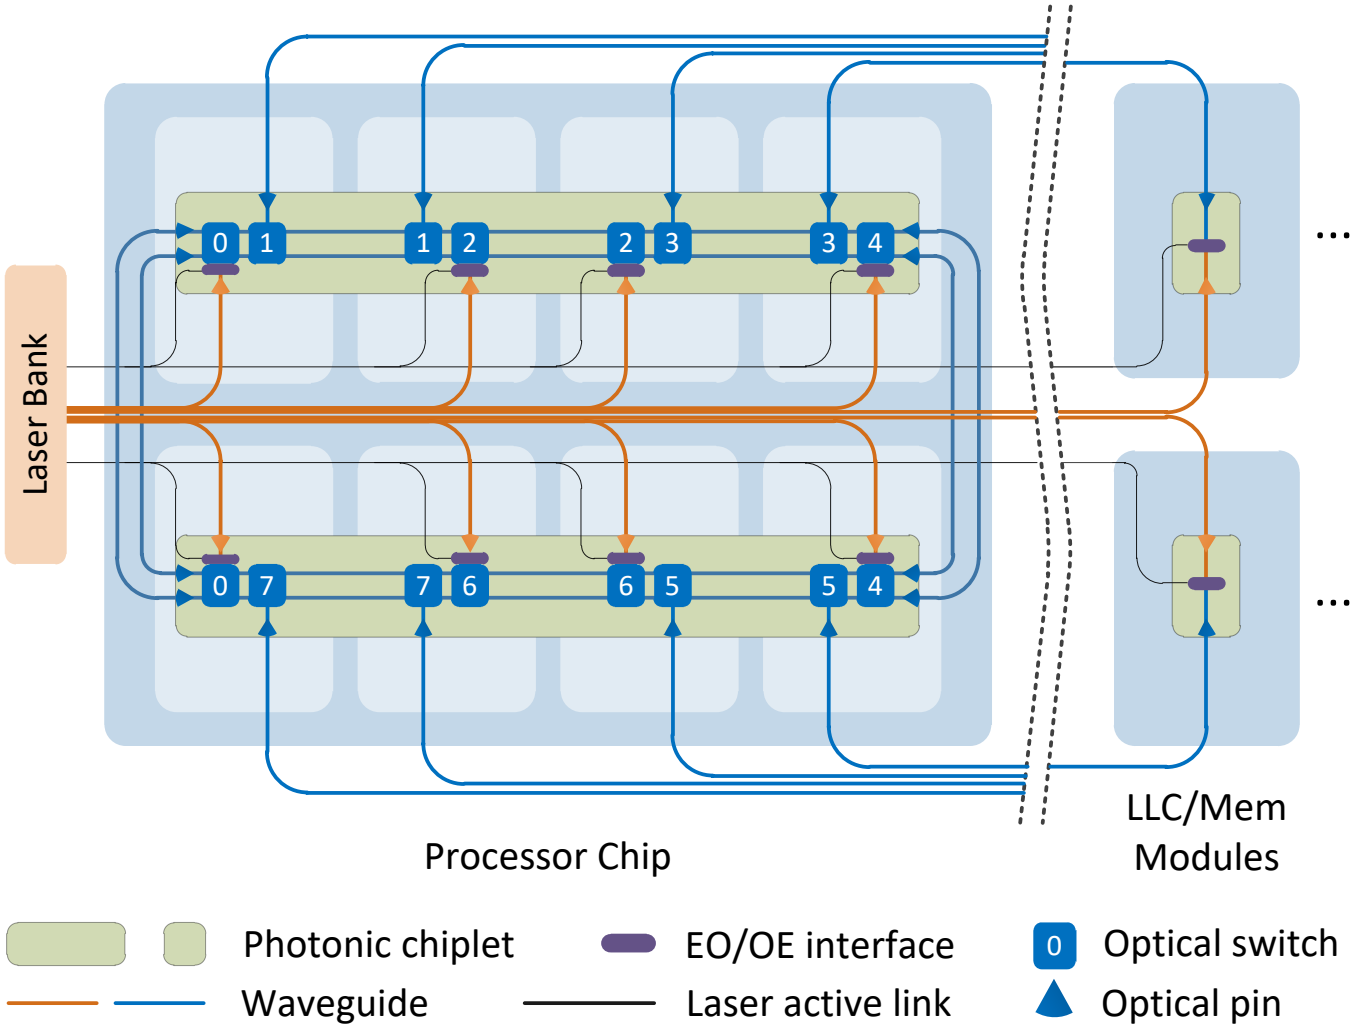

* Zhehui Wang, Zhifei Wang, Jiang Xu, *et al.*, "CAMON: Low-Cost Silicon Photonic Chiplet for Manycore Processors," IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems 2020

Performance & Performance/Energy

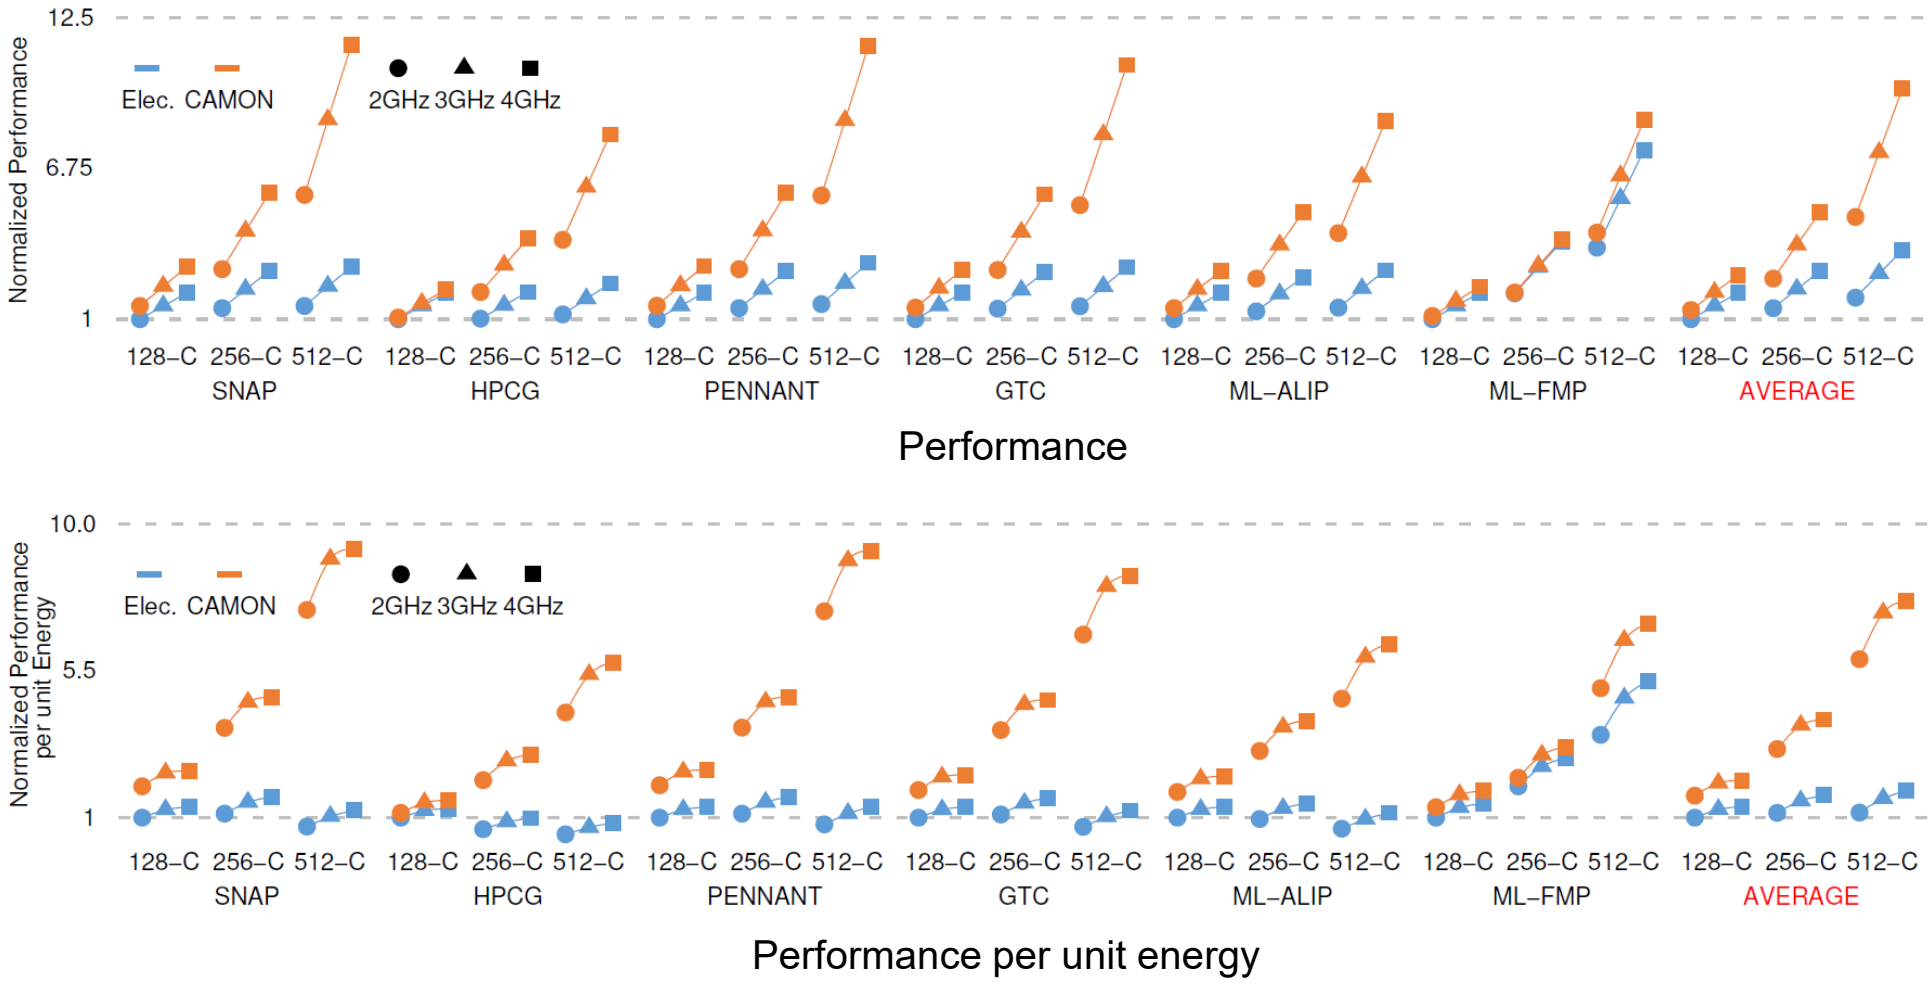

- Speedup the manycore processor by 2.7X
- Improve the full-system performance per unit energy by 4.6X

Bandwidth & Latency

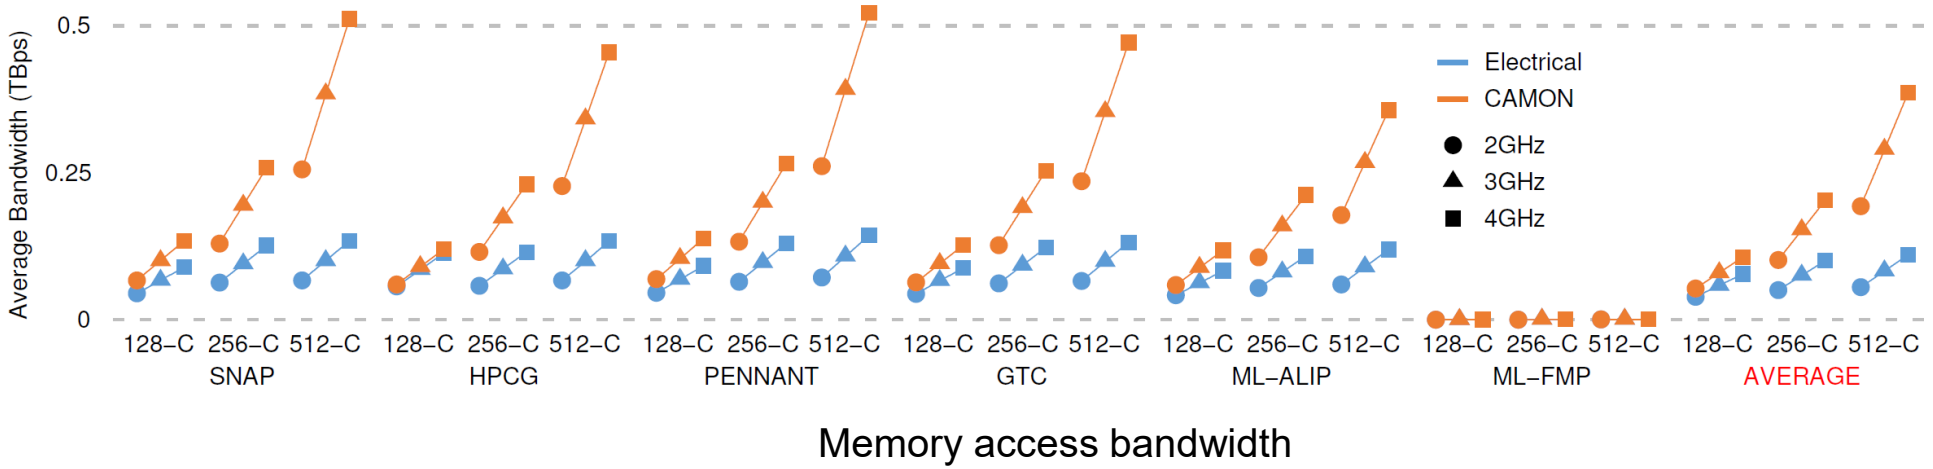

Memory access bandwidth

- Show 2.3X memory access bandwidth

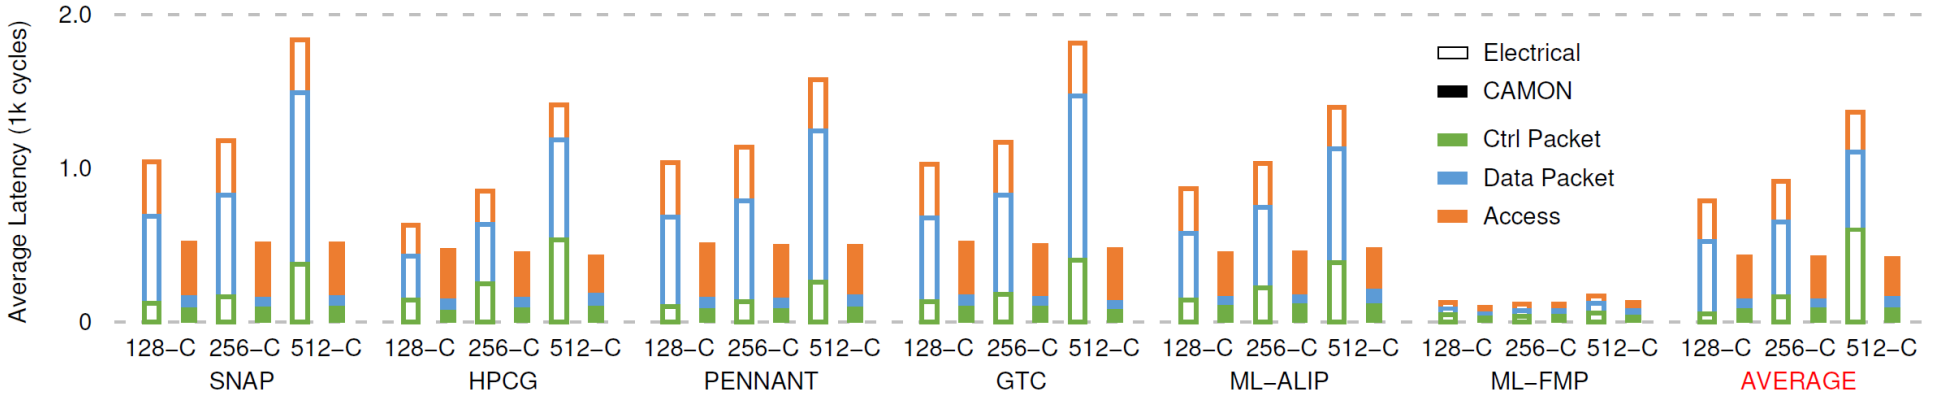

Memory access latency

- Reduce 48% memory access latency

Photonic Chiplet-based GPU

- Scaling GPU with photonic chiplet
- Electrical intra-chiplet network
- Optical inter-chiplet interconnect

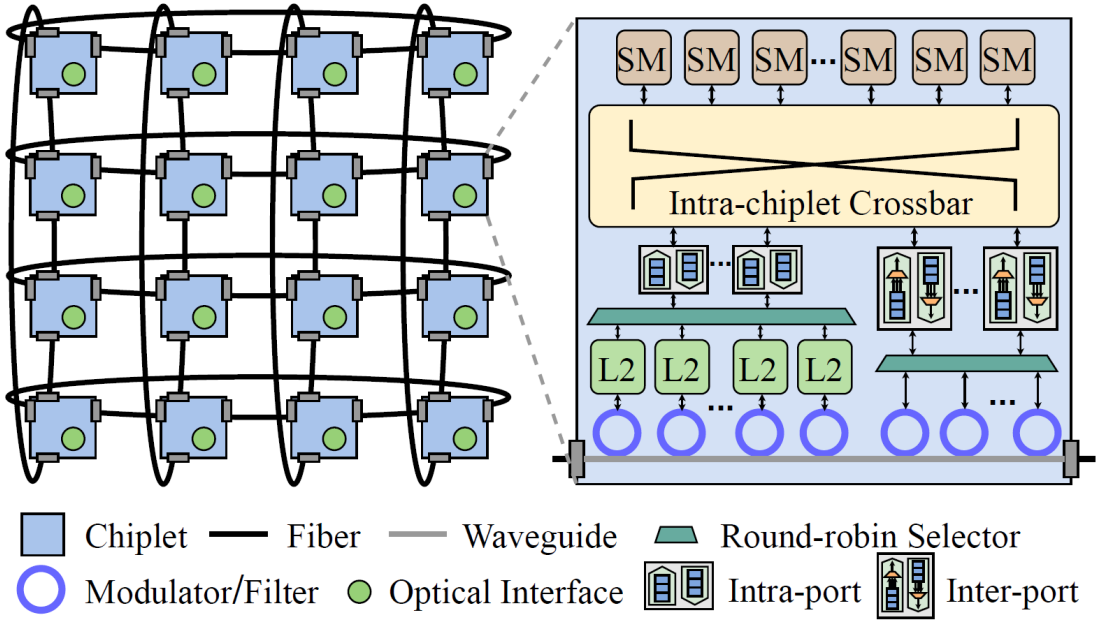

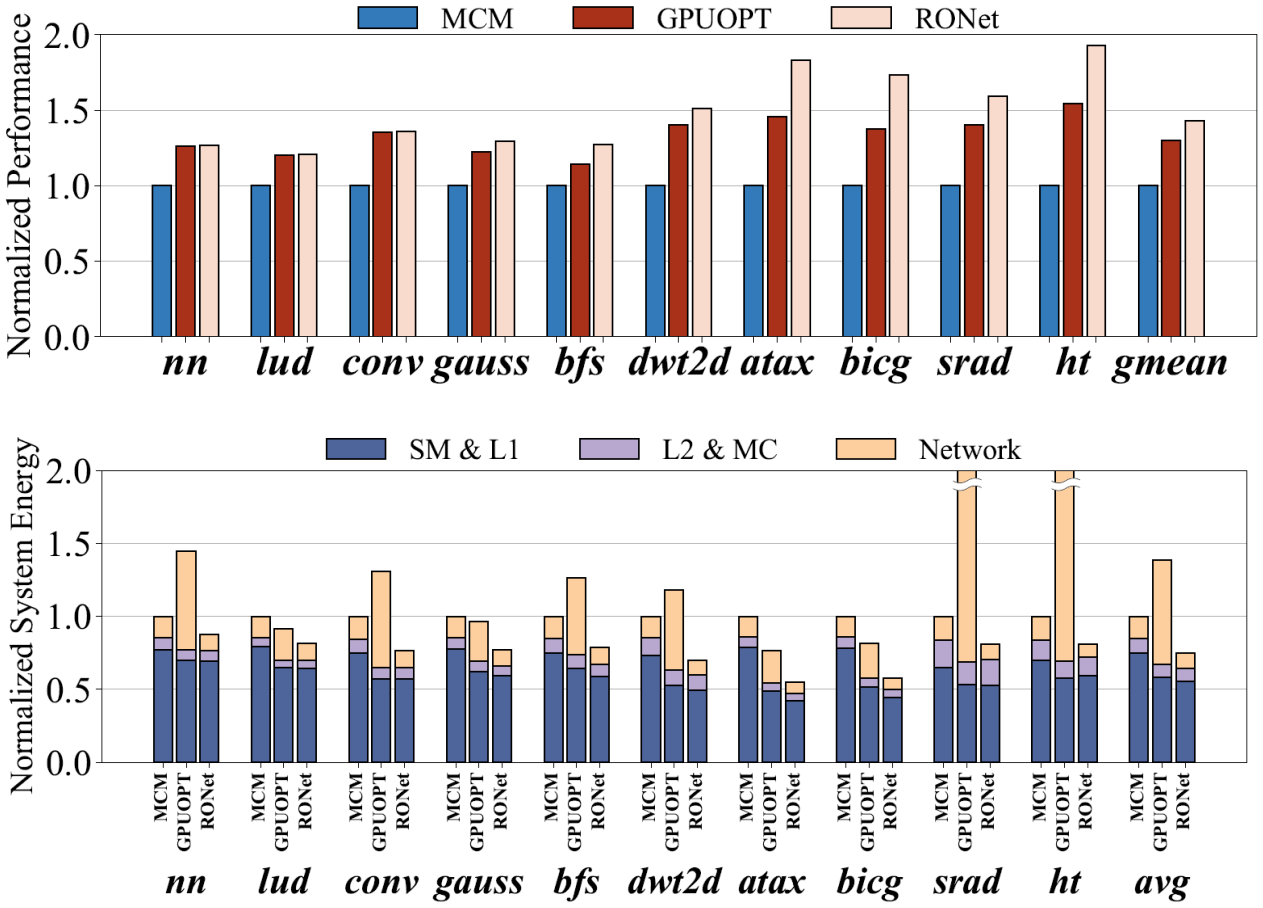

*Chengeng Li, *et al* "RONet: Scaling GPU System with Silicon Photonic Chiplet", ICCAD 2023

RSON: Rack-Scale Optical Network for Disaggregated Computing Systems

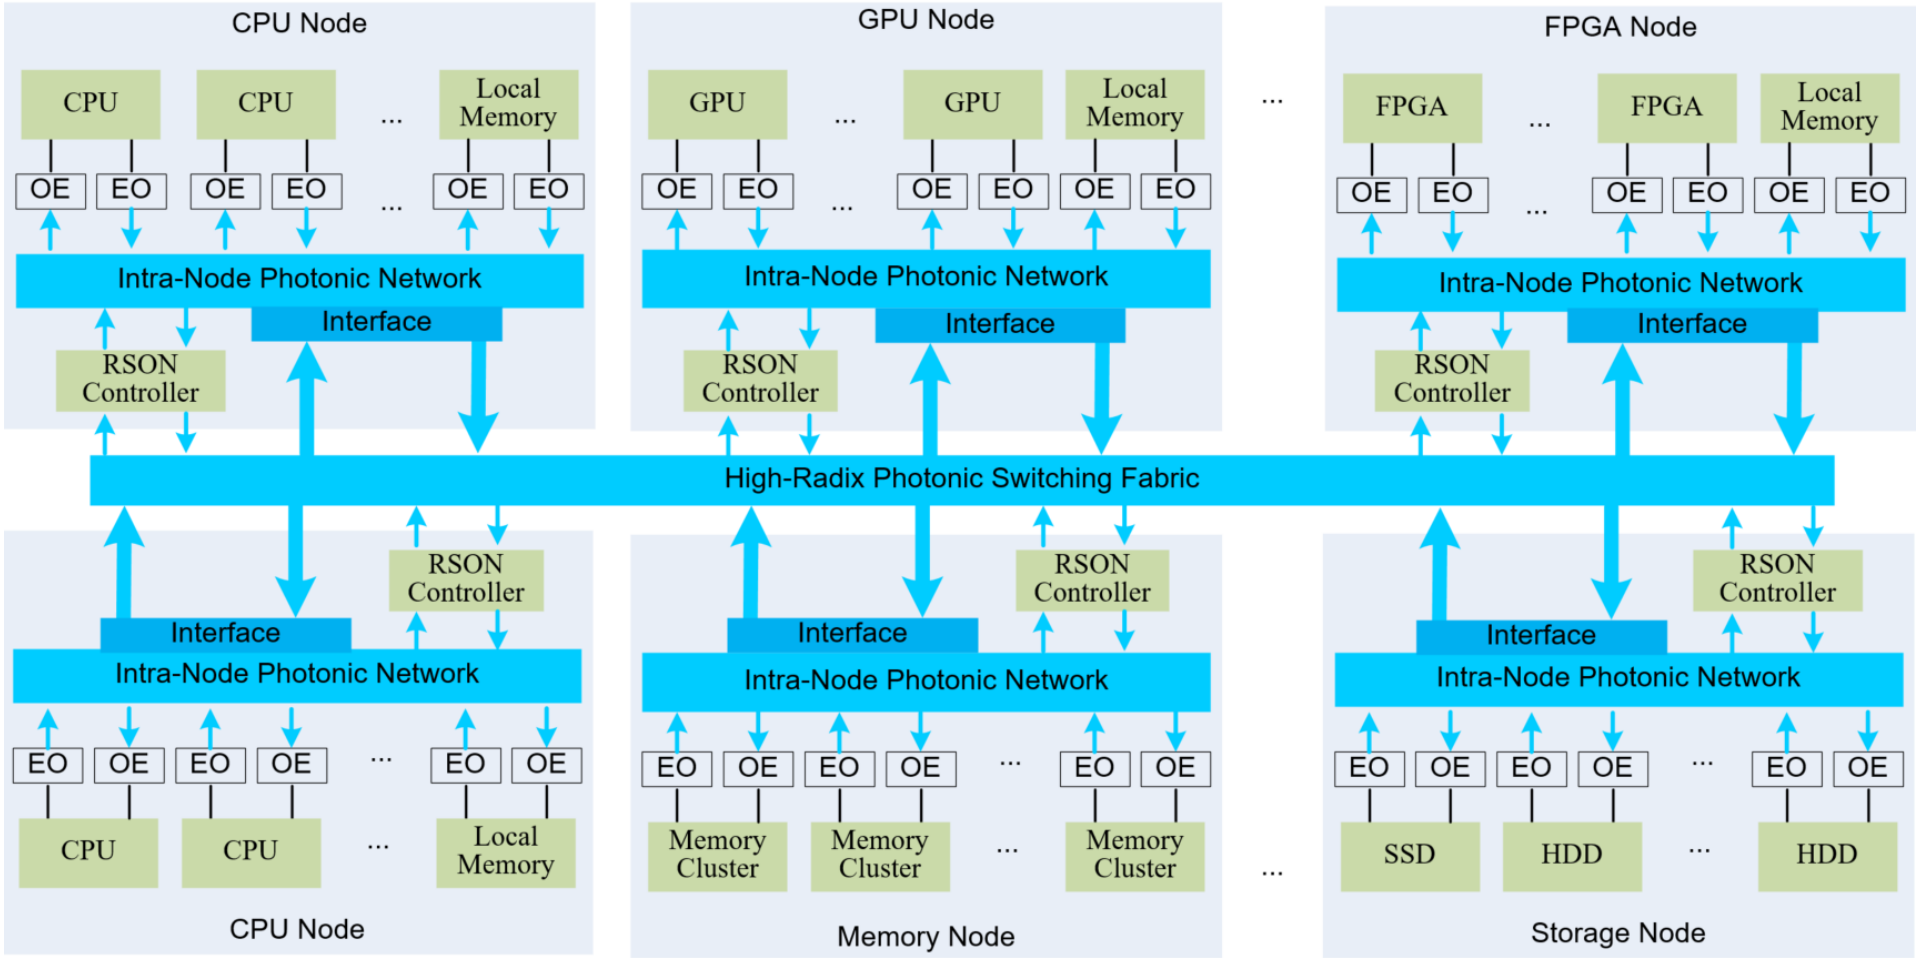

* P. Yang, *et al.* "Multi-Domain Inter/Intra-Chip Silicon Photonic Networks for Energy-Efficient Rack-Scale Computing Systems", IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 2020

Full-System Energy Efficiency

- 2048 cores and 4TB memory in 64 nodes
- 3X~8X higher full-system energy efficiency

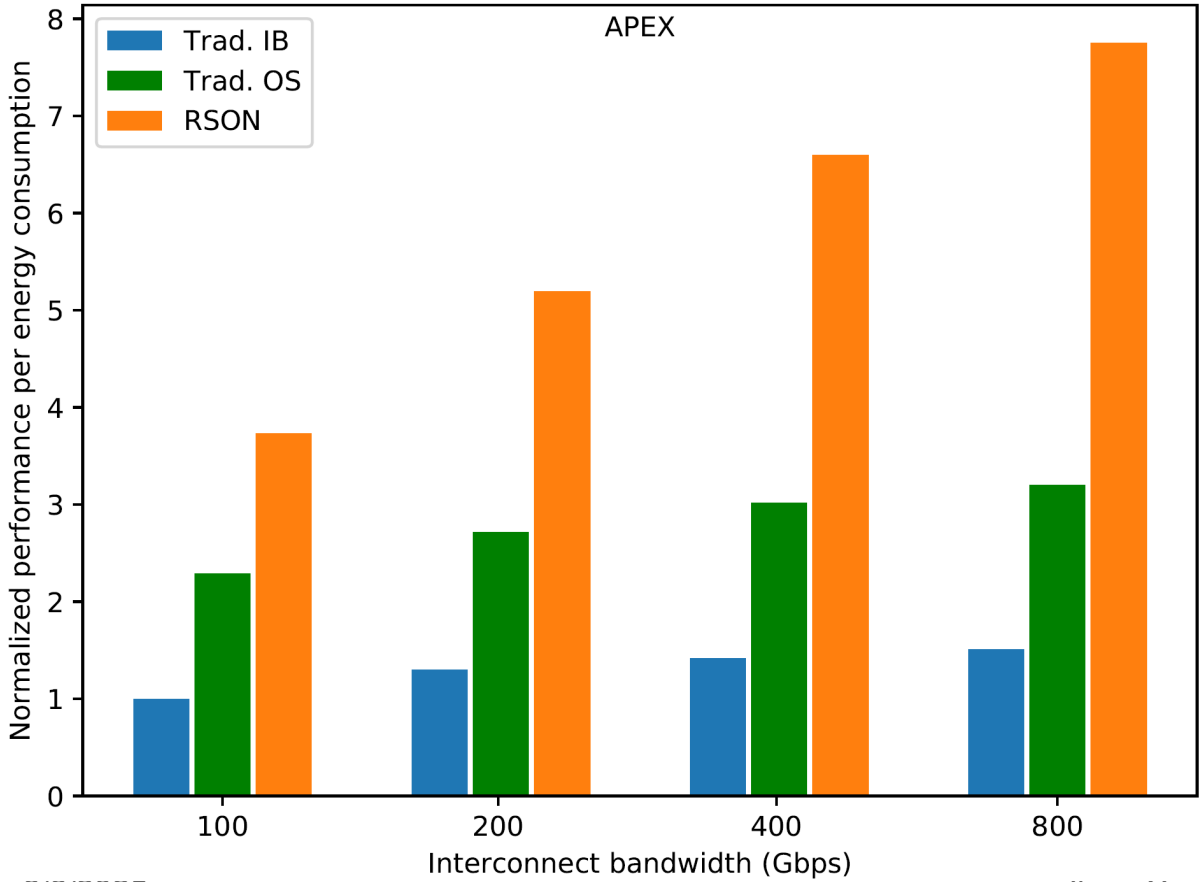

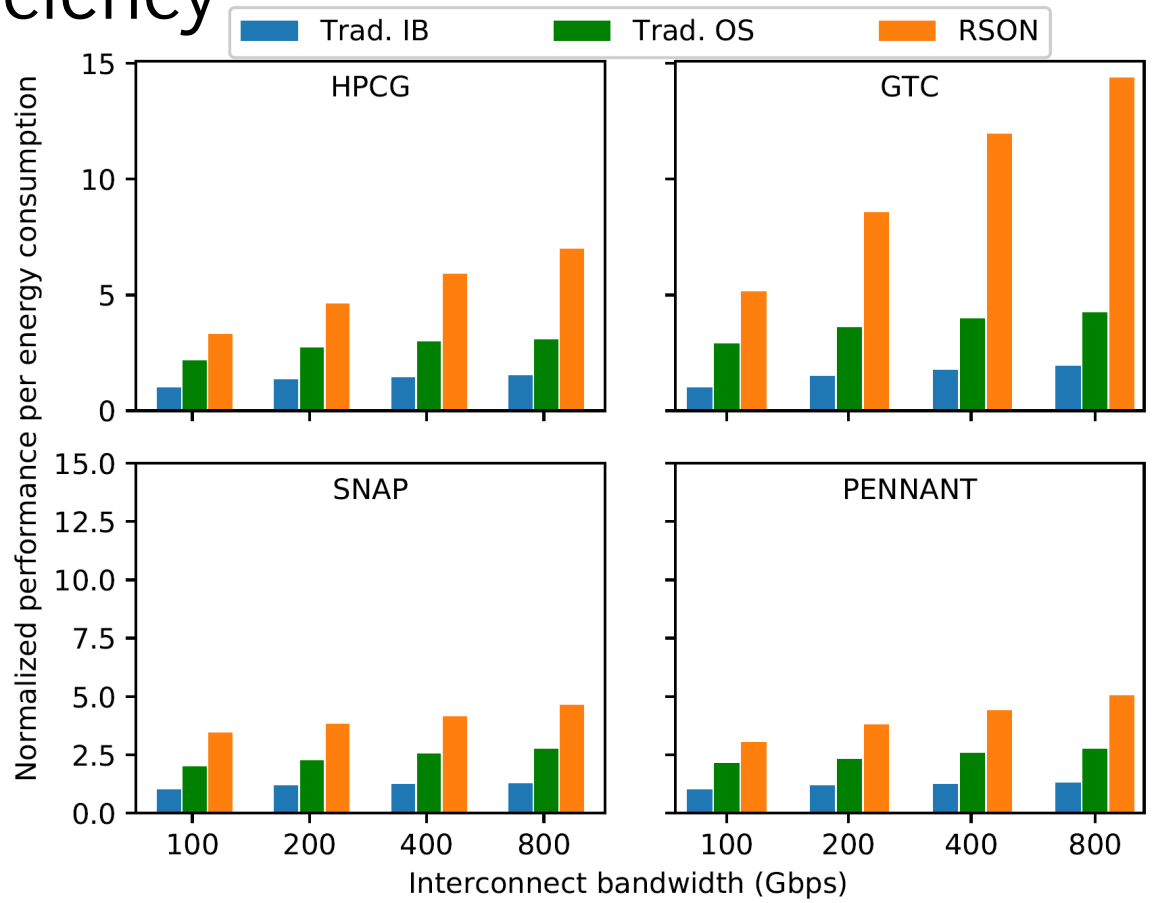

UONN Optical Neural Network Accelerator

- Support a wide range of neural networks
- Low-cost, energy-efficient, reconfigurable architecture

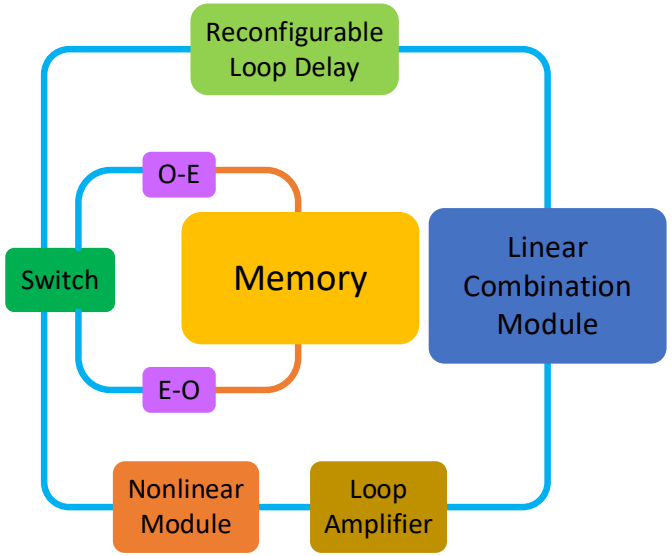

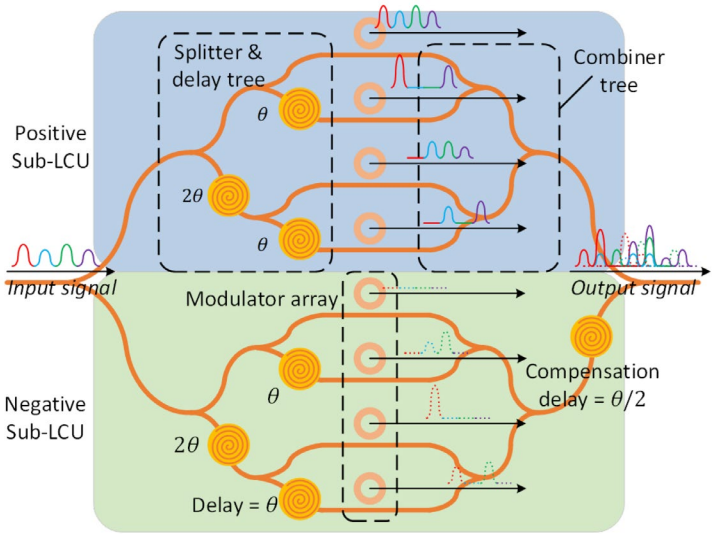

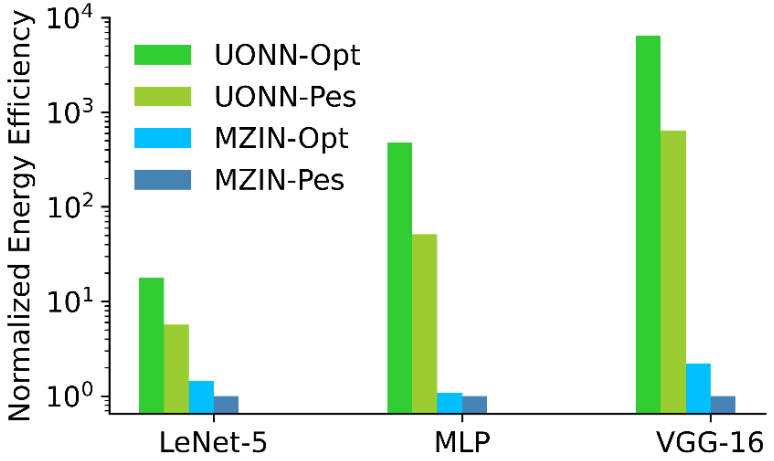

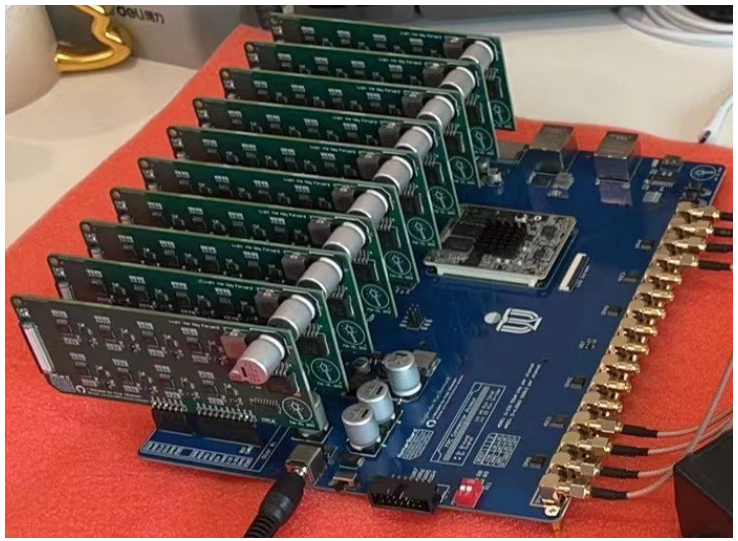

*Jiaxu Zhang, *et al*, "UONN: Energy-Efficient Optical Neural Network," Asia Communications and Photonics Conference, 2021
Jiaxu. Zhang, *et al*, "Energy-Efficient and Low-Cost Optical Neural Network," OPIC, 2022

Publicly Released Tools for Photonic-Electronic Integration

- JADE heterogeneous system design, modeling, and simulation environment
- COSMIC heterogeneous system benchmark suite
- BOSIM SPICE-compatible silicon photonic device design, modeling, and simulation platform
- CLAP optical crosstalk and loss modeling and analysis platform
- OTemp optical thermal effect modeling and analysis platform
- OEIL optical and electrical interface and link design, modeling and analysis platform
- MCSL realistic network-on-chip traffic patterns
- PowerSoC power delivery system design, modeling, and analysis platform

Bibliography for inter/intra-chip optical networks

eexu.home.ece.ust.hk

香港科技大学(广州)
THE HONG KONG
UNIVERSITY OF SCIENCE AND
TECHNOLOGY (GUANGZHOU)

微电子学域
MICROELECTRONICS THRUST
功能枢纽
FUNCTION HUB