

Silicon-Photonic Clos Networks for Global On-Chip Communication

Ajay Joshi^{*}, Christopher Batten[†], Yong-Jin Kwon[‡], Scott Beamer[‡],
Imran Shamim[†], Krste Asanović[‡], Vladimir Stojanović[†]

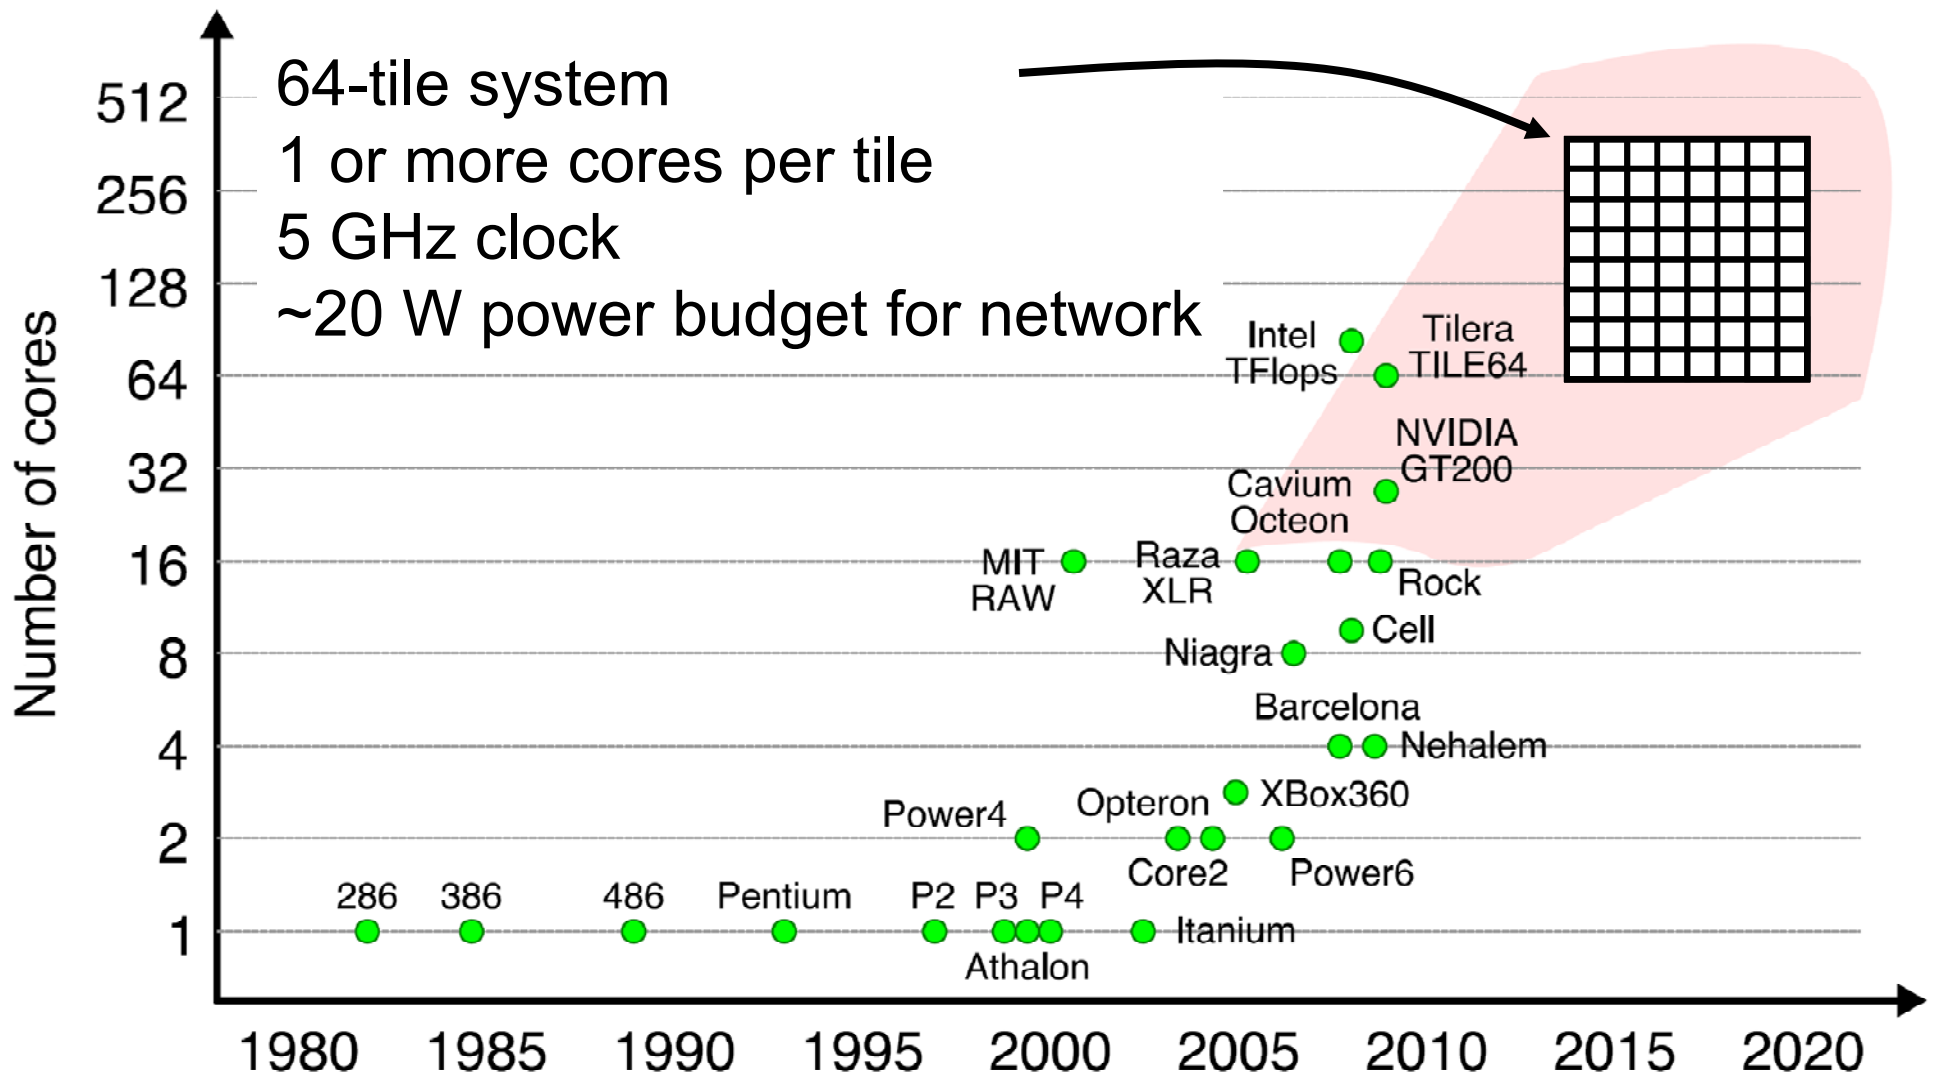
^{*}Boston University, Boston MA

[†]Massachusetts Institute of Technology, Cambridge MA


[‡]University of California, Berkeley, CA

HPEC 2009


Our target manycore system




Landscape of on-chip photonic networks



Mesh




CMesh



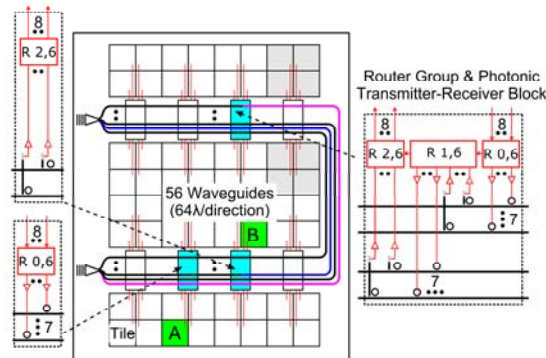
Clos




Crossbar



[Shacham'07]
[Petracca'08]

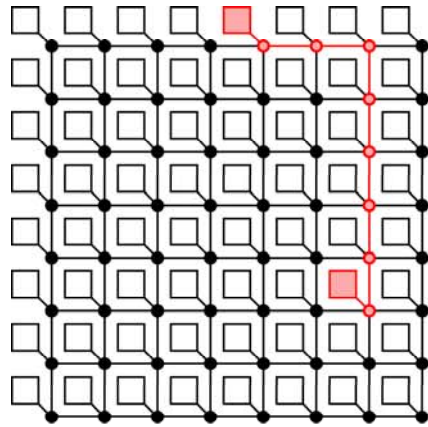


[This work]
[Pan'09]




[Vantrease'08]
[Psota'07]
[Kirman'06]


Outline




Mesh



CMesh



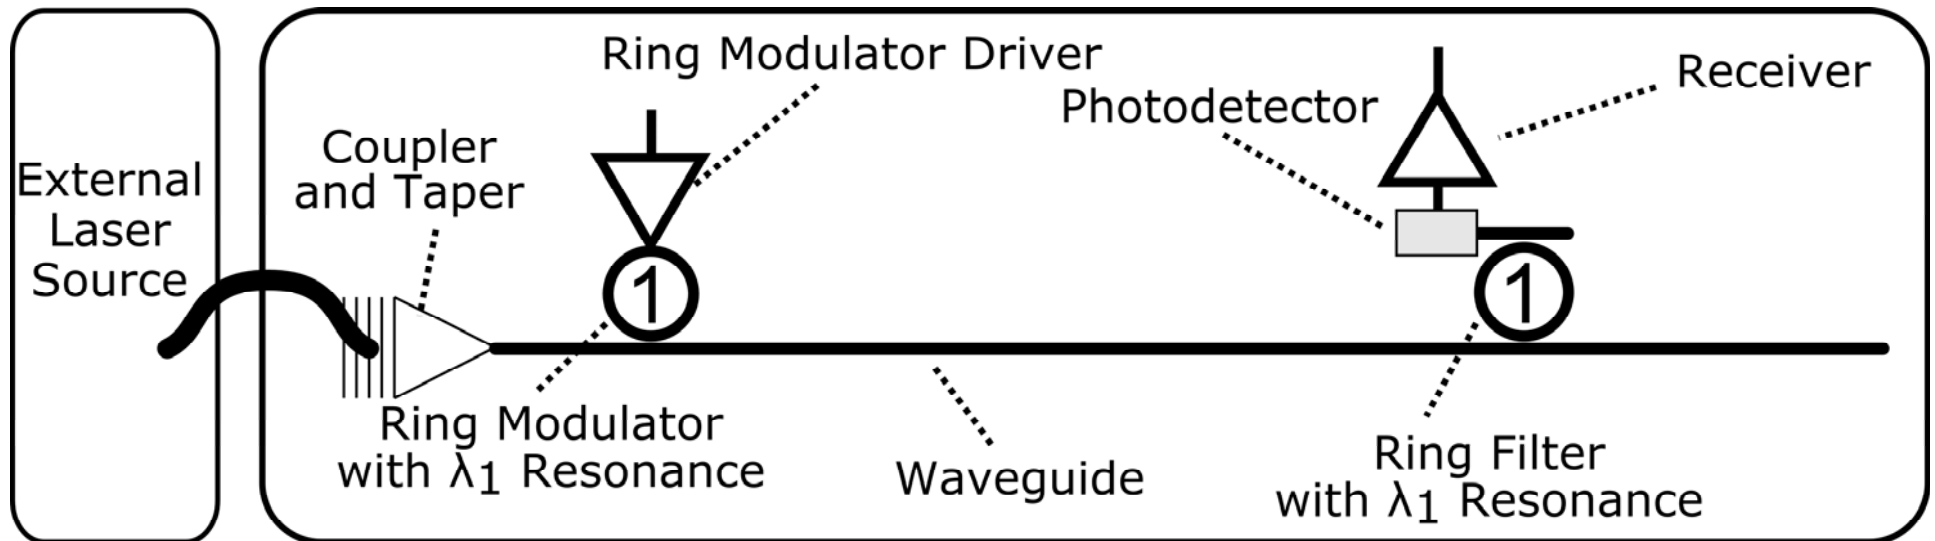
Clos



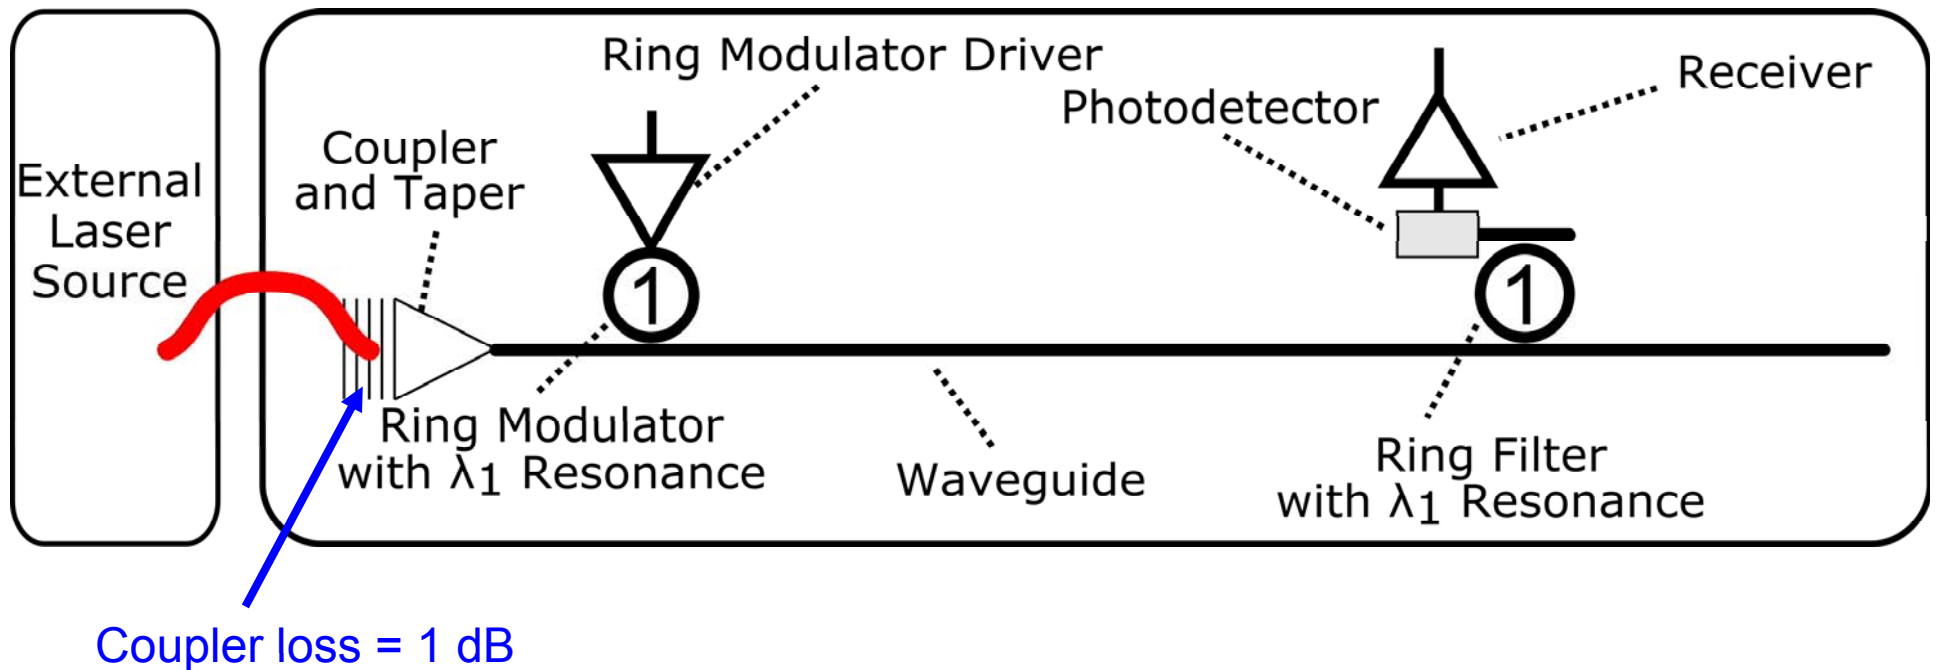
Crossbar

- ❑ Photonic interconnect technology
- ❑ Photonic networks
- ❑ Electrical vs Photonic networks

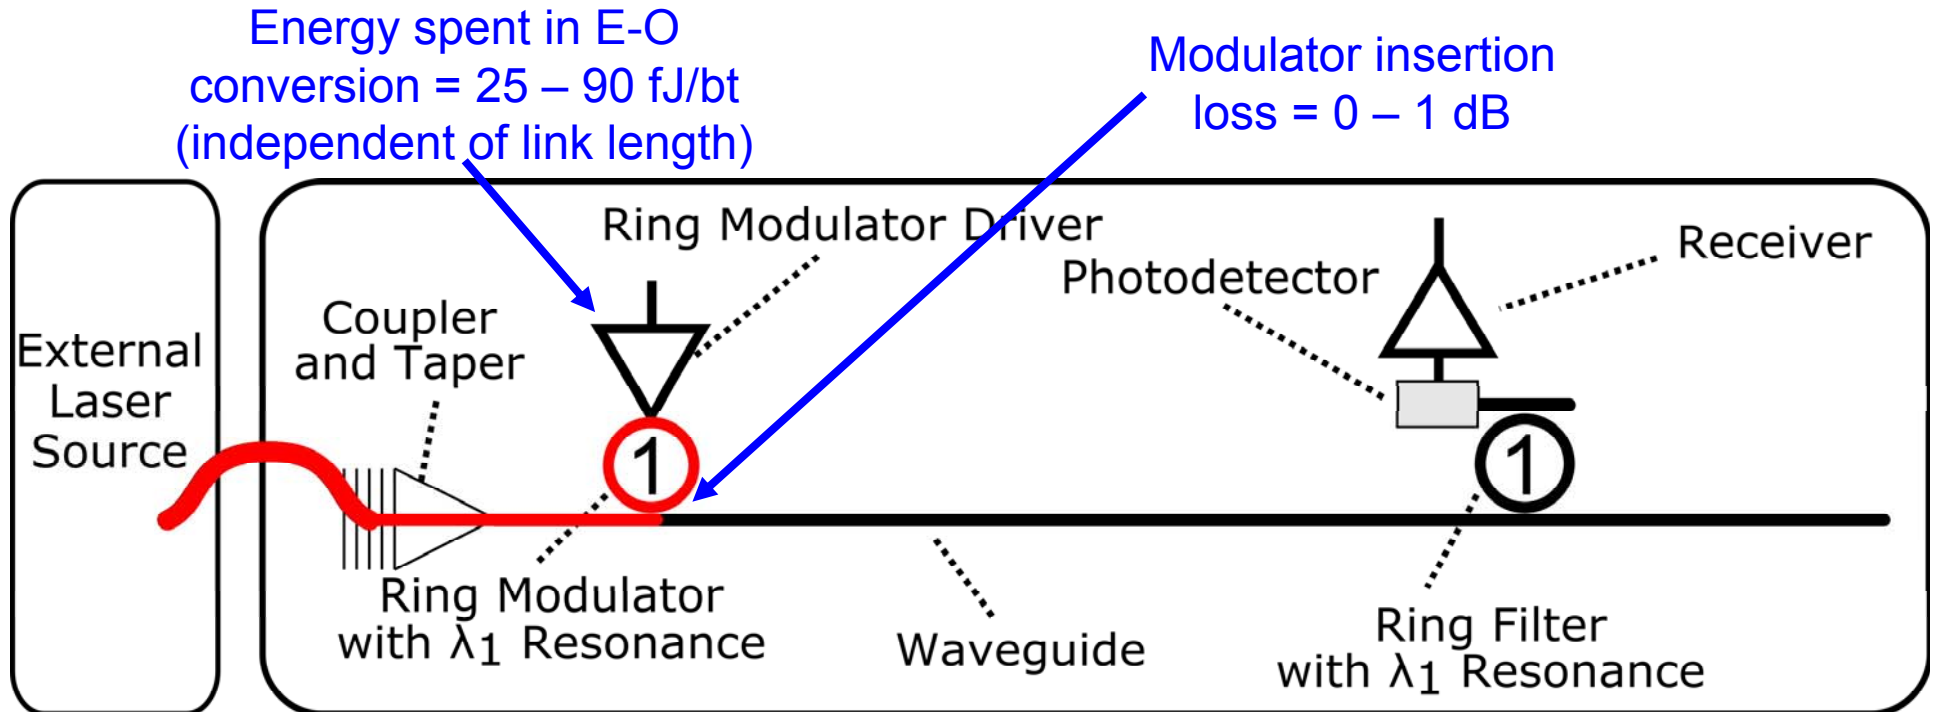
Photonic technology – Silicon photonic link



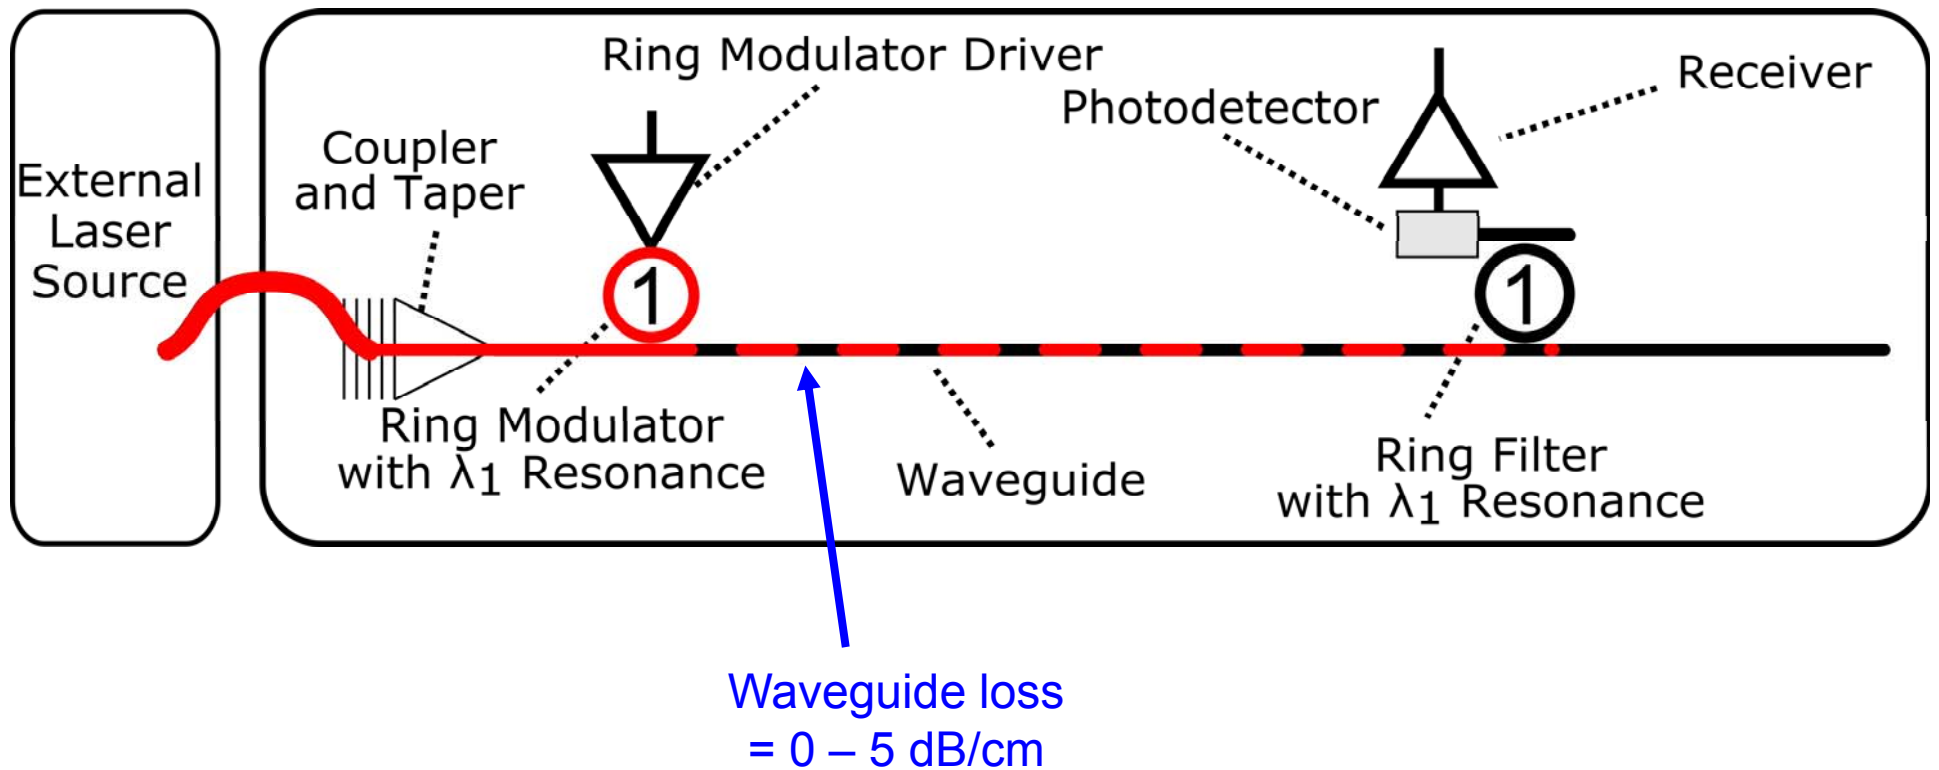
Silicon photonic link – Coupler



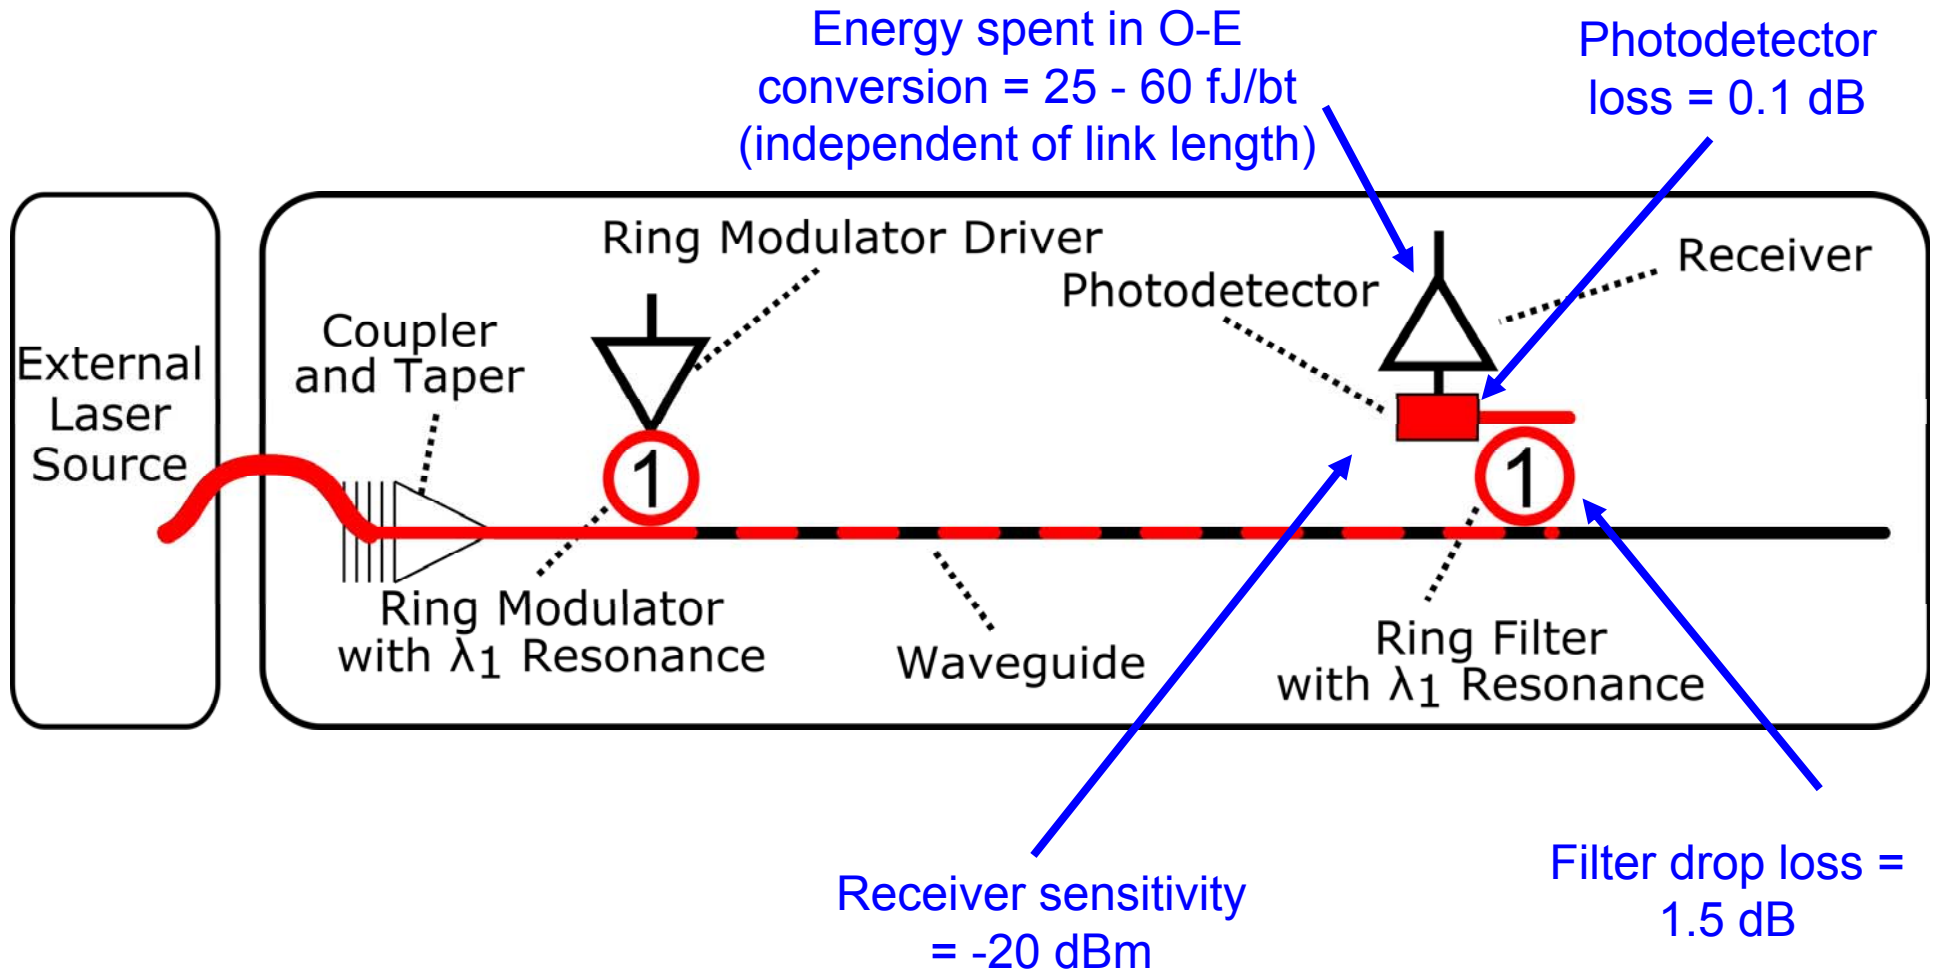
Silicon photonic link – Ring modulator



Silicon photonic link – Waveguide

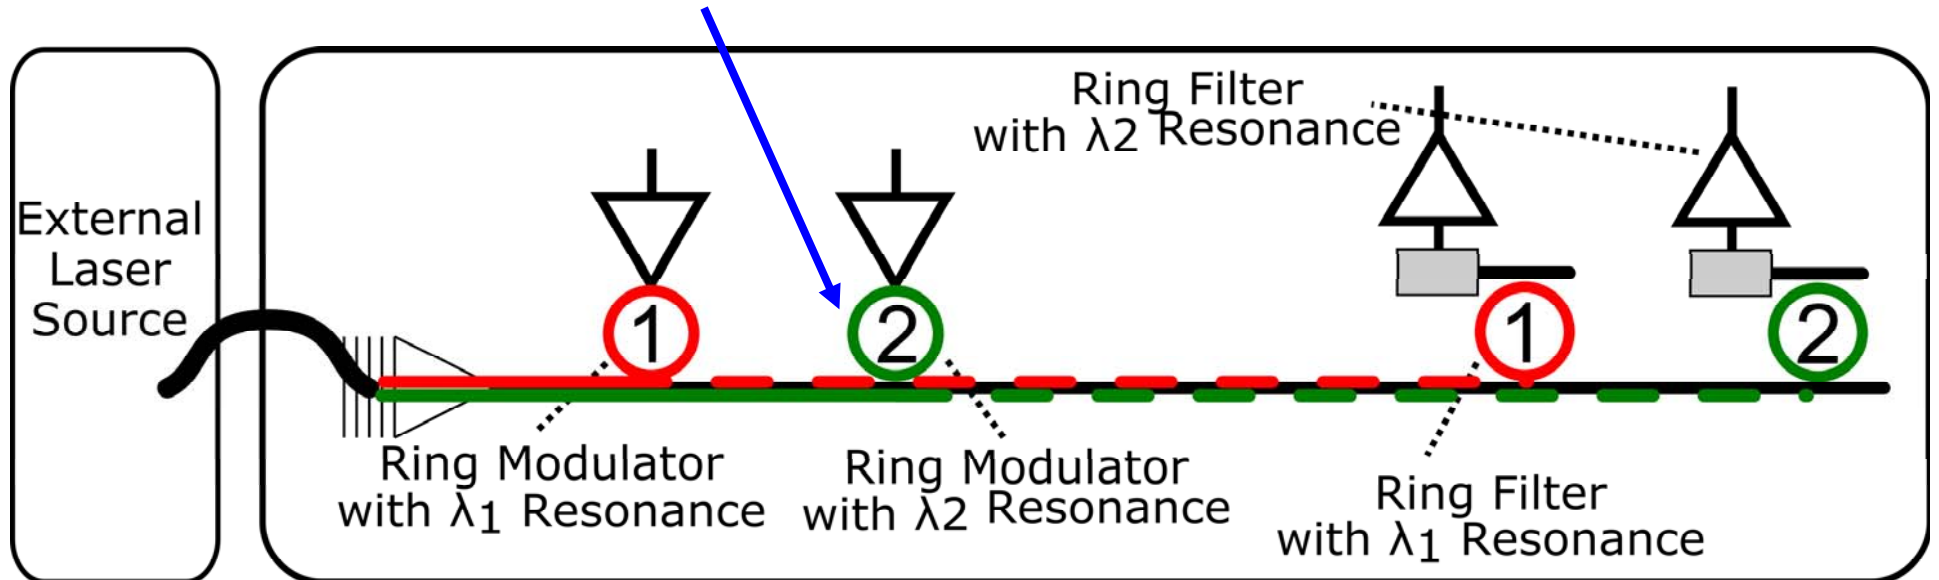


Silicon photonic link – Ring filter, photodetector



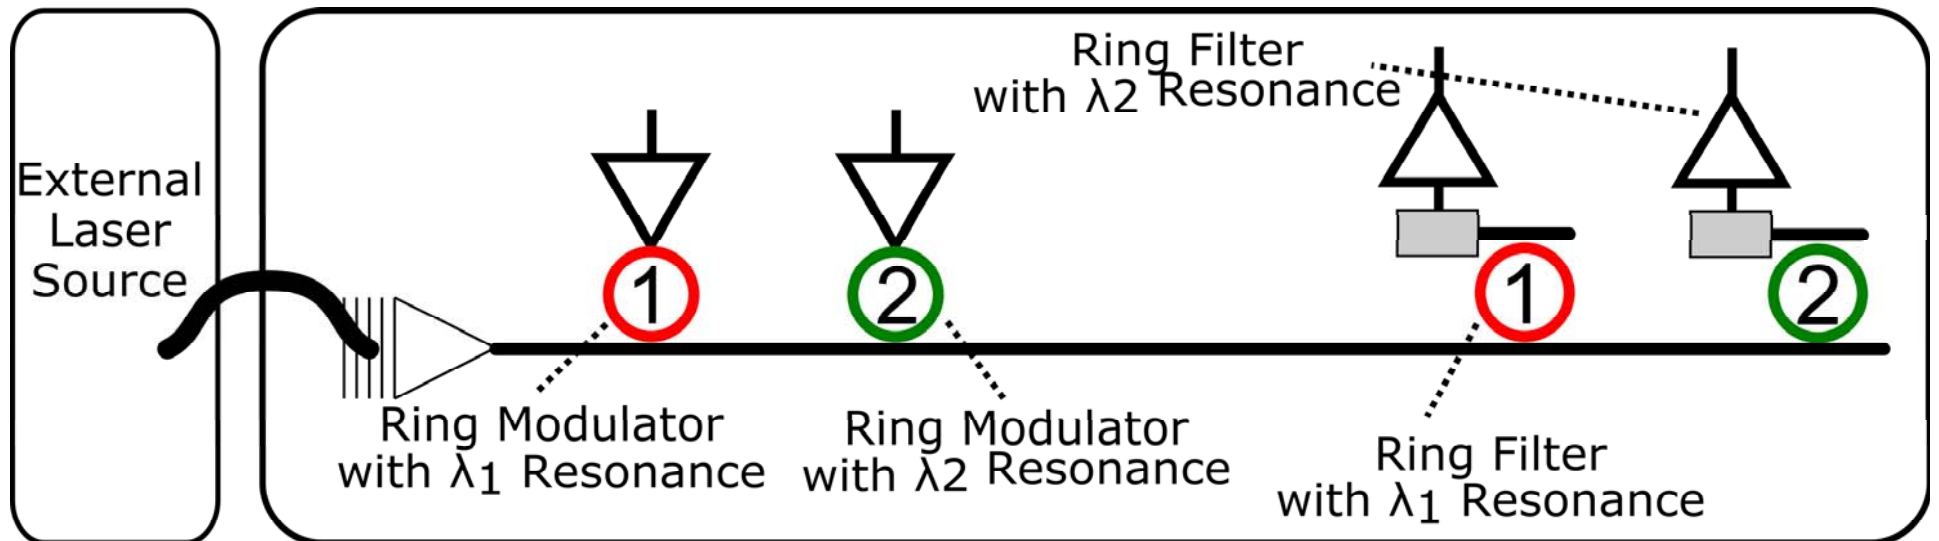
Silicon photonic link – WDM

Through ring loss =
 $1e-4 - 1e-2$ dB/ring




- ❑ Dense WDM ($128 \lambda/\text{wg}$, $10 \text{ Gbps}/\lambda$) improves bandwidth density

Silicon photonic link – Energy cost



- ❑ E-O-E conversion cost – 50-150 fJ/bt (independent of length)
- ❑ Thermal tuning energy (increases with ring count)
- ❑ External laser power (dependent on losses in photonic devices)

Electrical technology



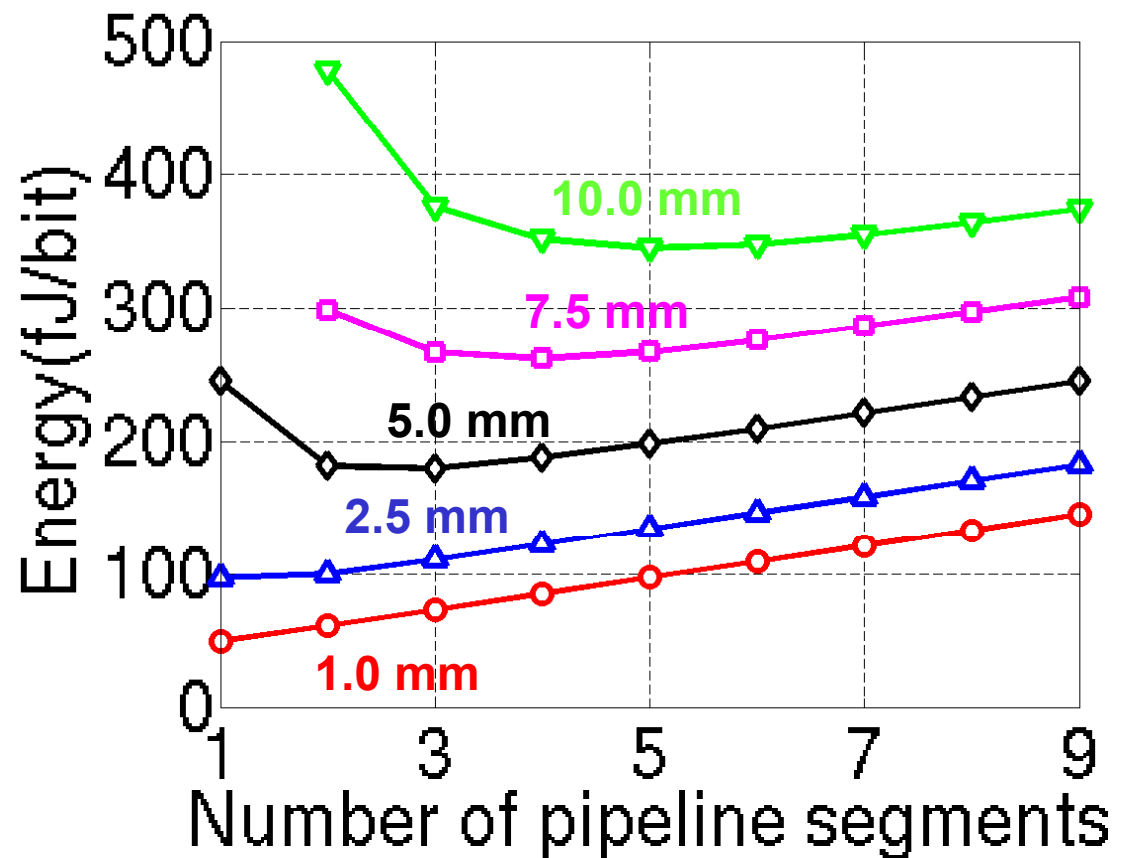
Repeater inserted pipelined wires

Design constraints


- 22 nm technology
- 500 nm pitch
- 5 GHz clock

Design parameters

- Wire width
- Repeater size
- Repeater spacing



Electrical technology



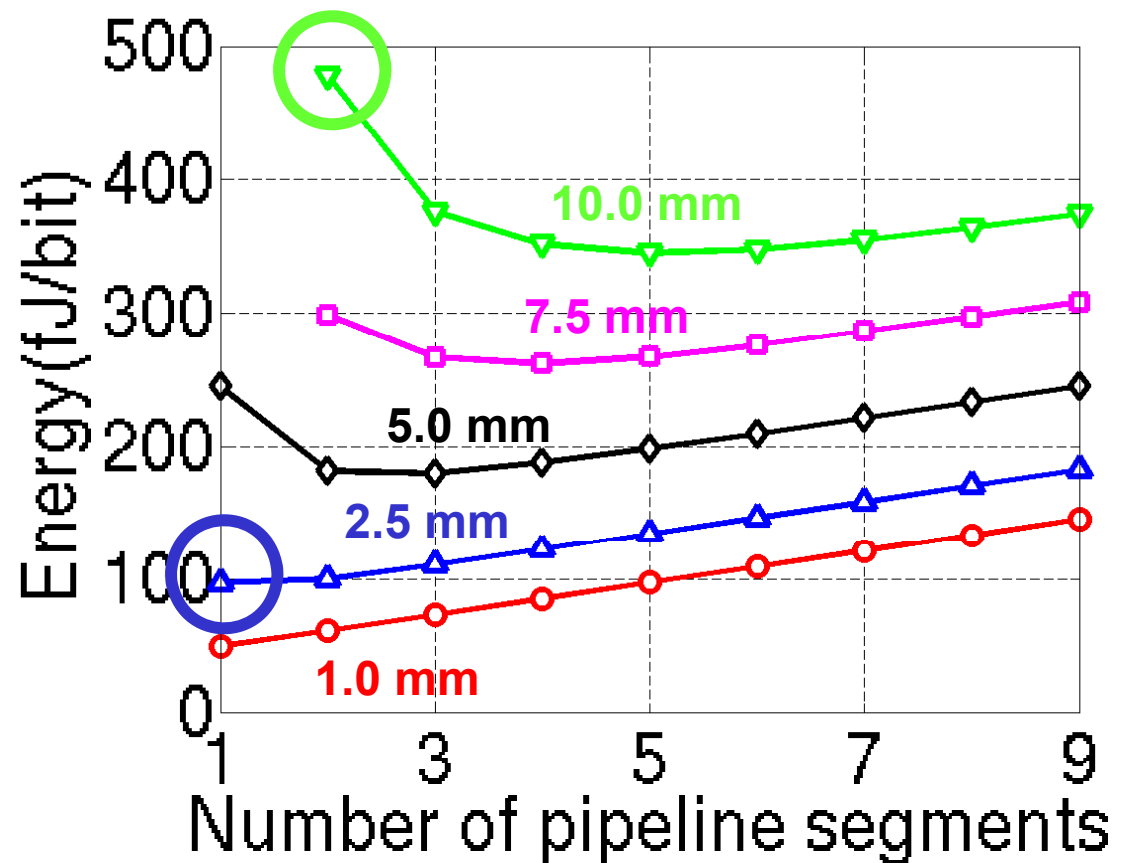
Repeater inserted pipelined wires

Design constraints


- 22 nm technology
- 500 nm pitch
- 5 GHz clock

Design parameters

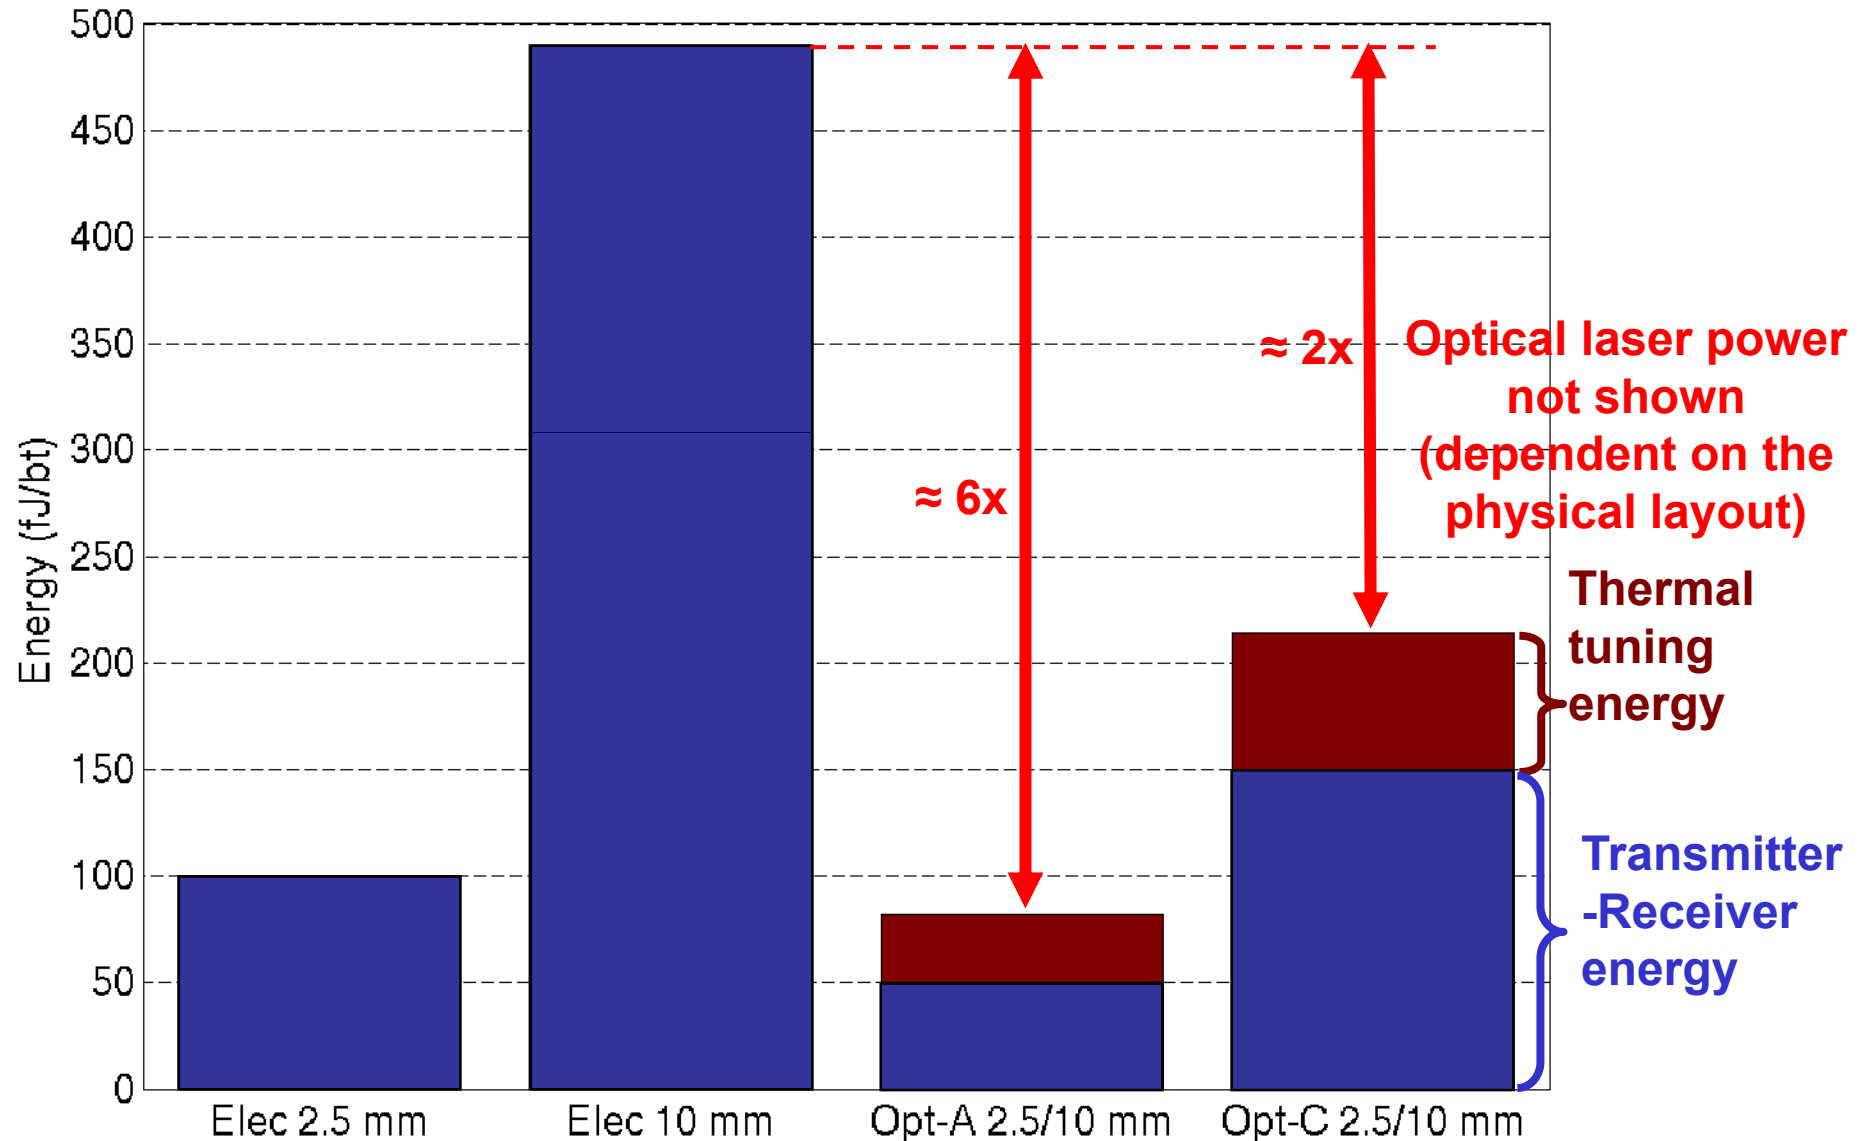
- Wire width
- Repeater size
- Repeater spacing




Electrical vs Optical links – Energy cost



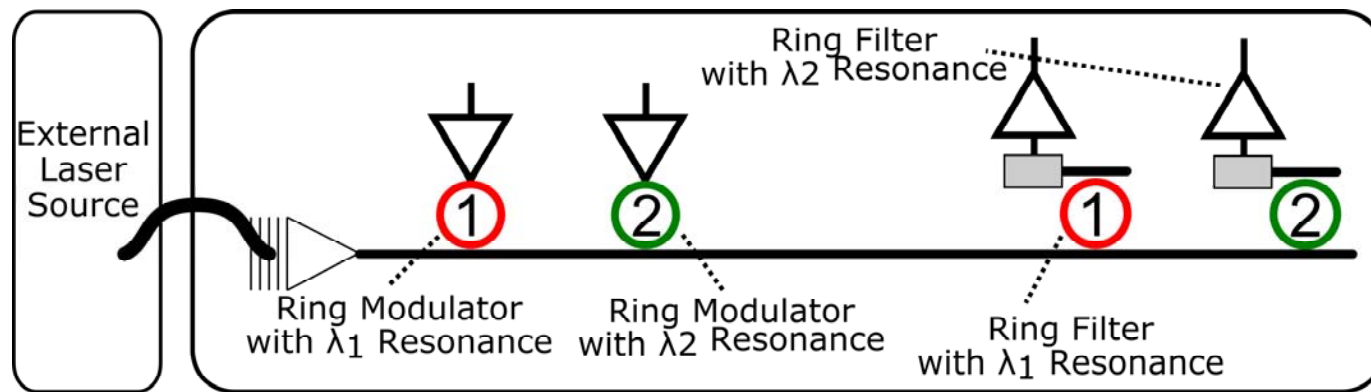
Electrical vs Optical links – Energy cost



Electrical vs Optical links – Bandwidth density




Repeater inserted pipelined wires – 10 Gbps/ μ




Wavelength-division multiplexed photonic link – 320 Gbps/ μ

30x bandwidth density advantage using optical links


Outline



Mesh



CMesh




Clos




Crossbar

- ❑ Photonic interconnect technology
- ❑ **Photonic networks**
- ❑ Electrical vs Photonic networks


Outline



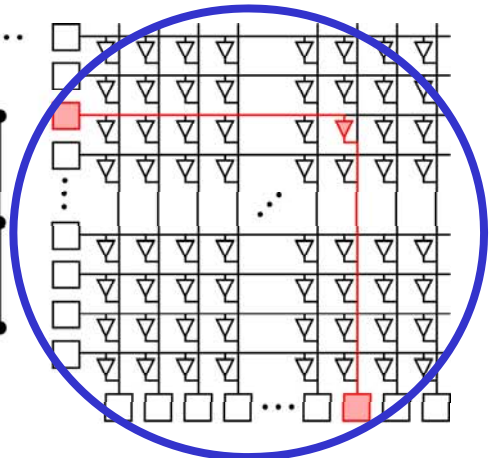
Mesh



CMesh



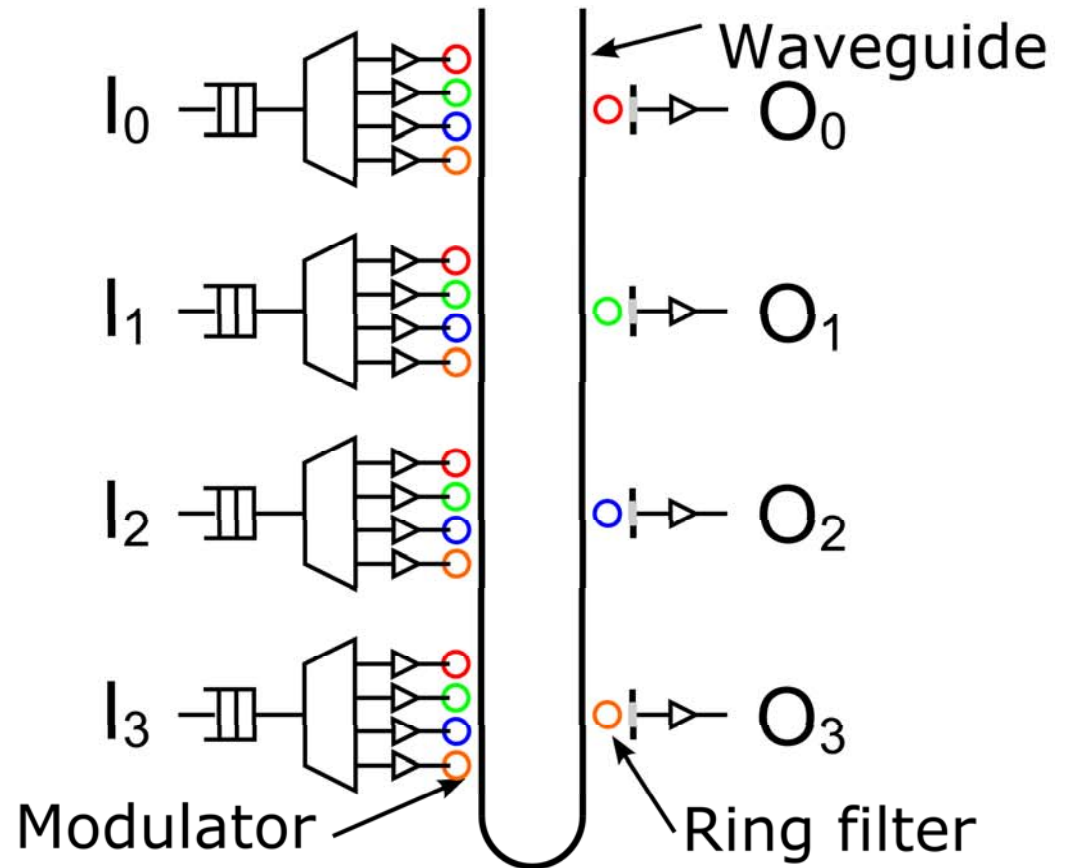

Clos



Crossbar

- ❑ Photonic interconnect technology
- ❑ **Photonic networks**
- ❑ Electrical vs Photonic networks


Distributed Multiplexer Crossbar



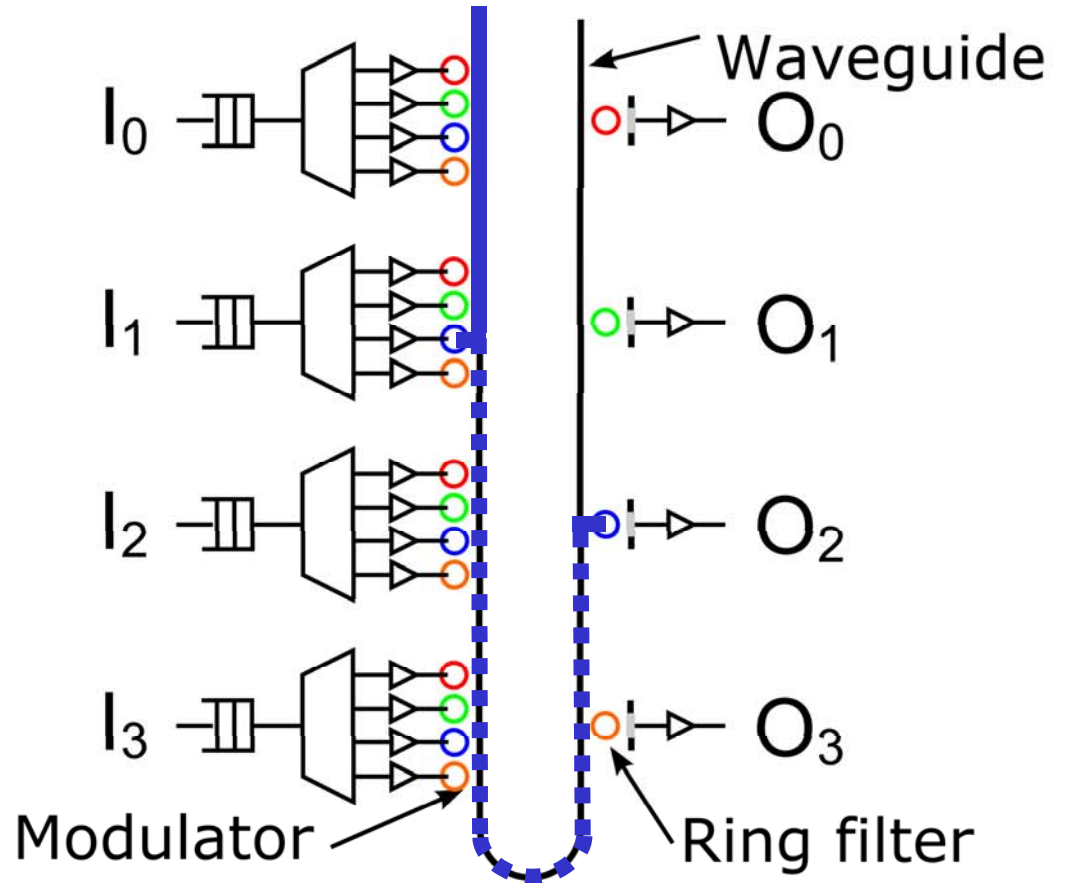
Electrical design

Photonic design

Distributed Multiplexer Crossbar

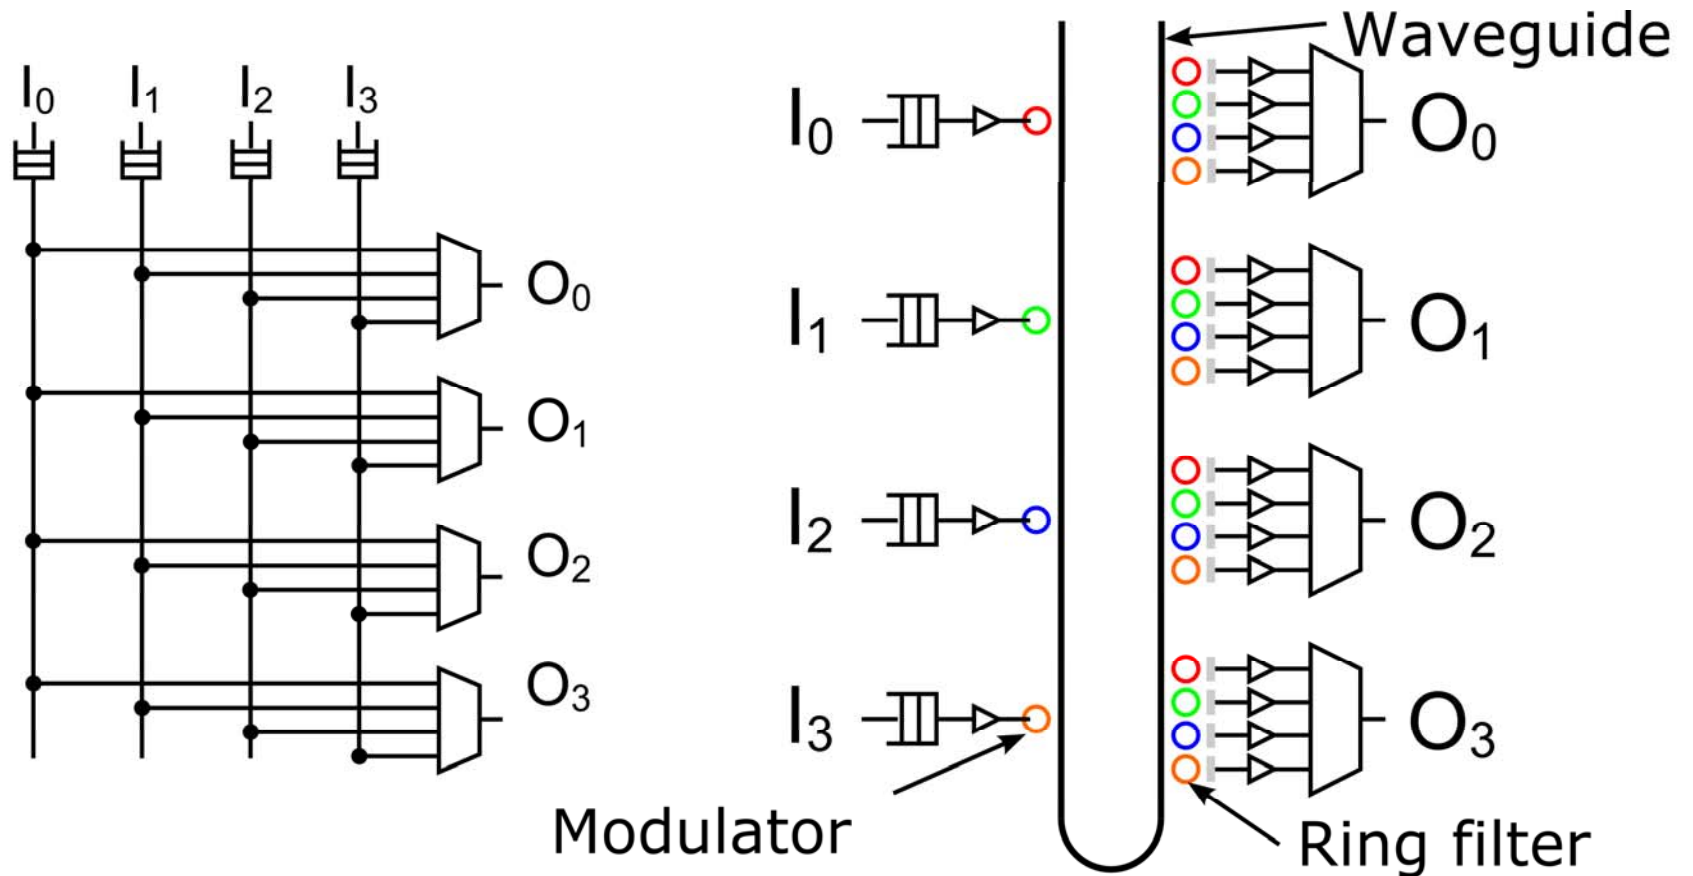


Electrical design



Photonic design

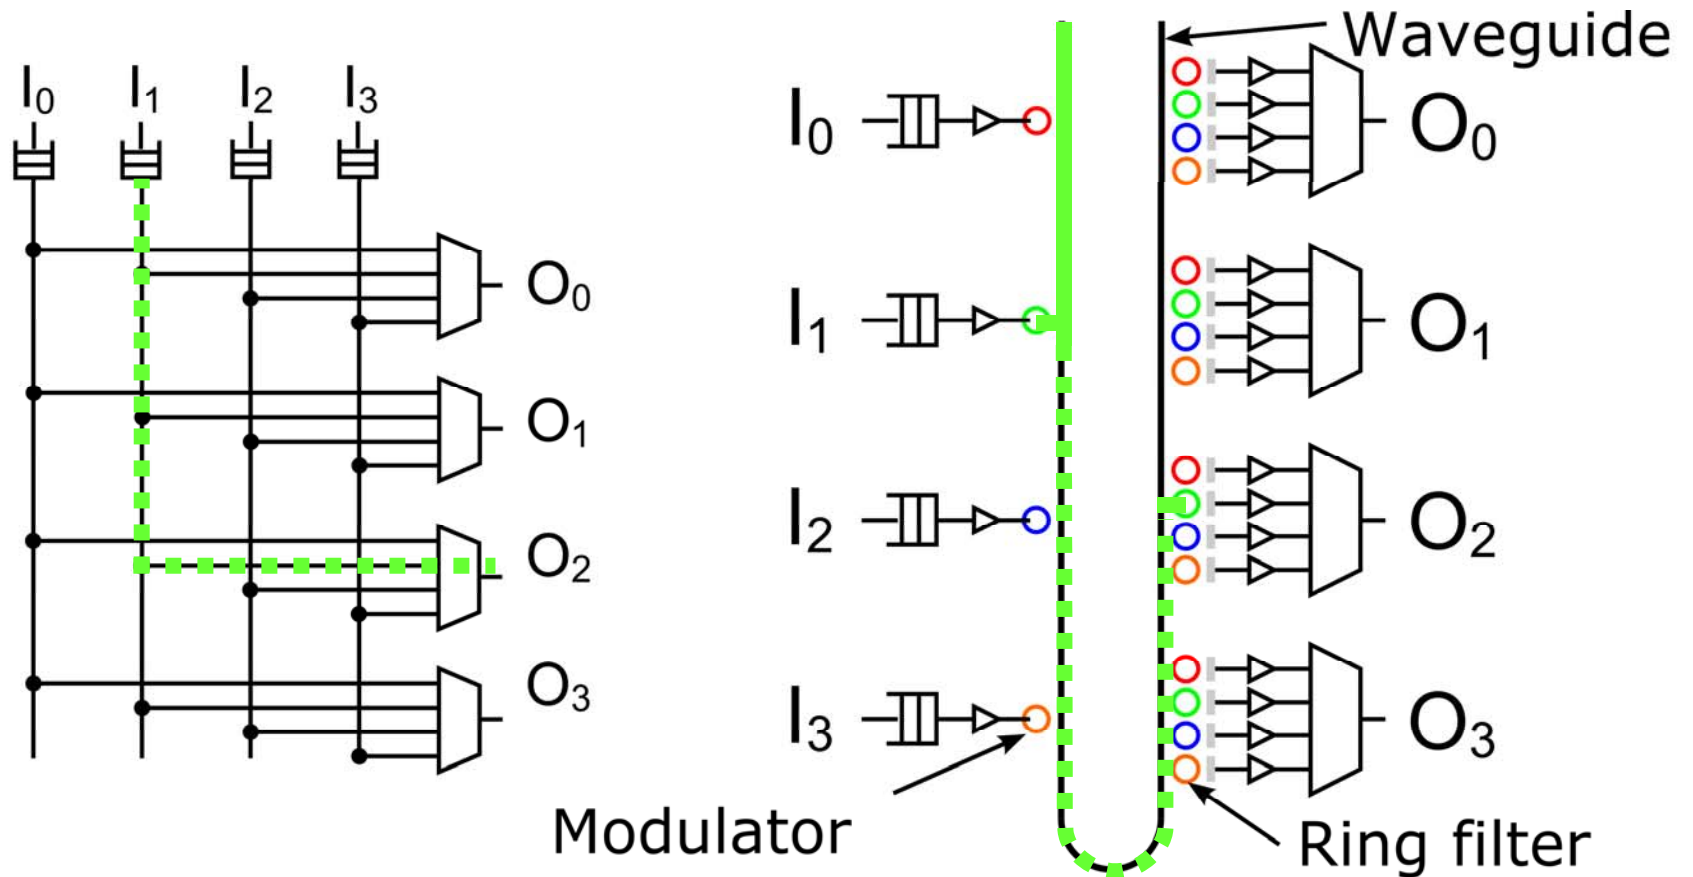
Centralized Multiplexer Crossbar



Electrical design

Photonic design


Centralized Multiplexer Crossbar



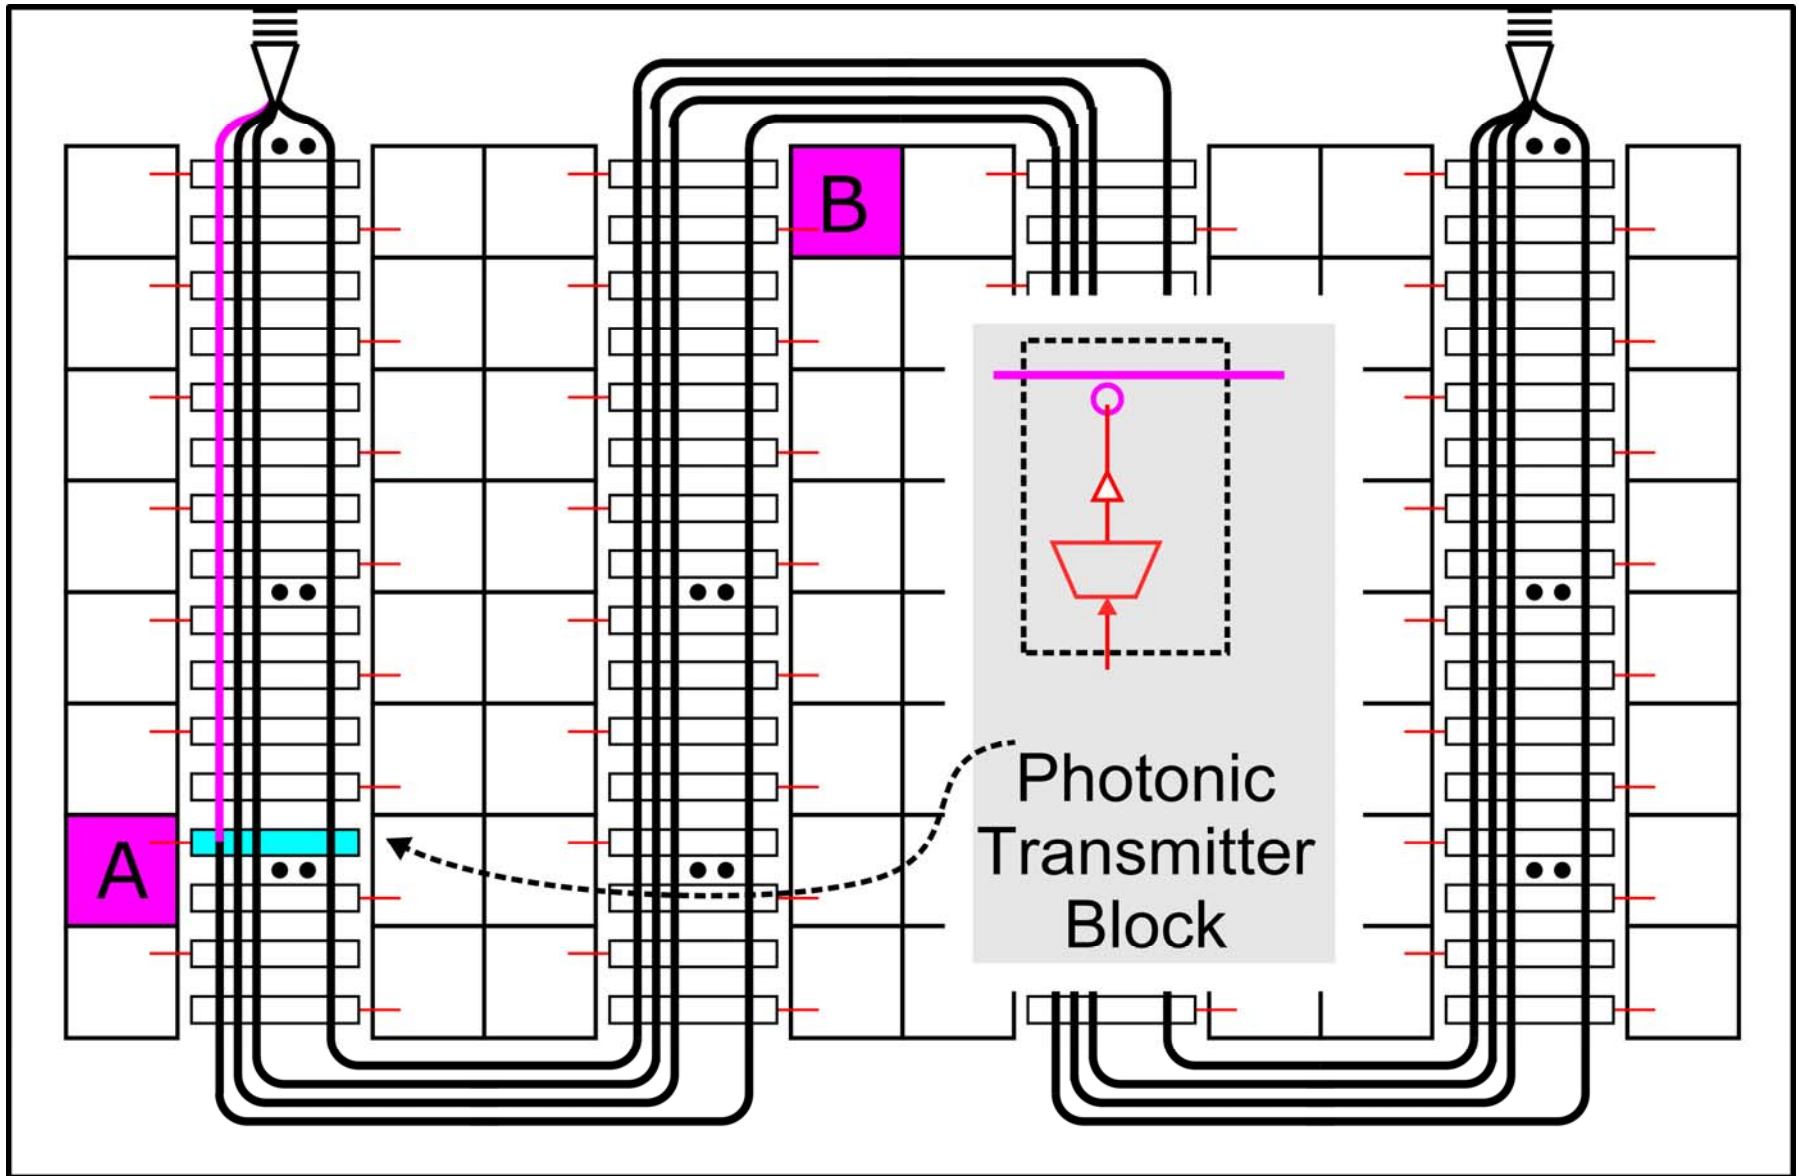
Electrical design

Photonic design

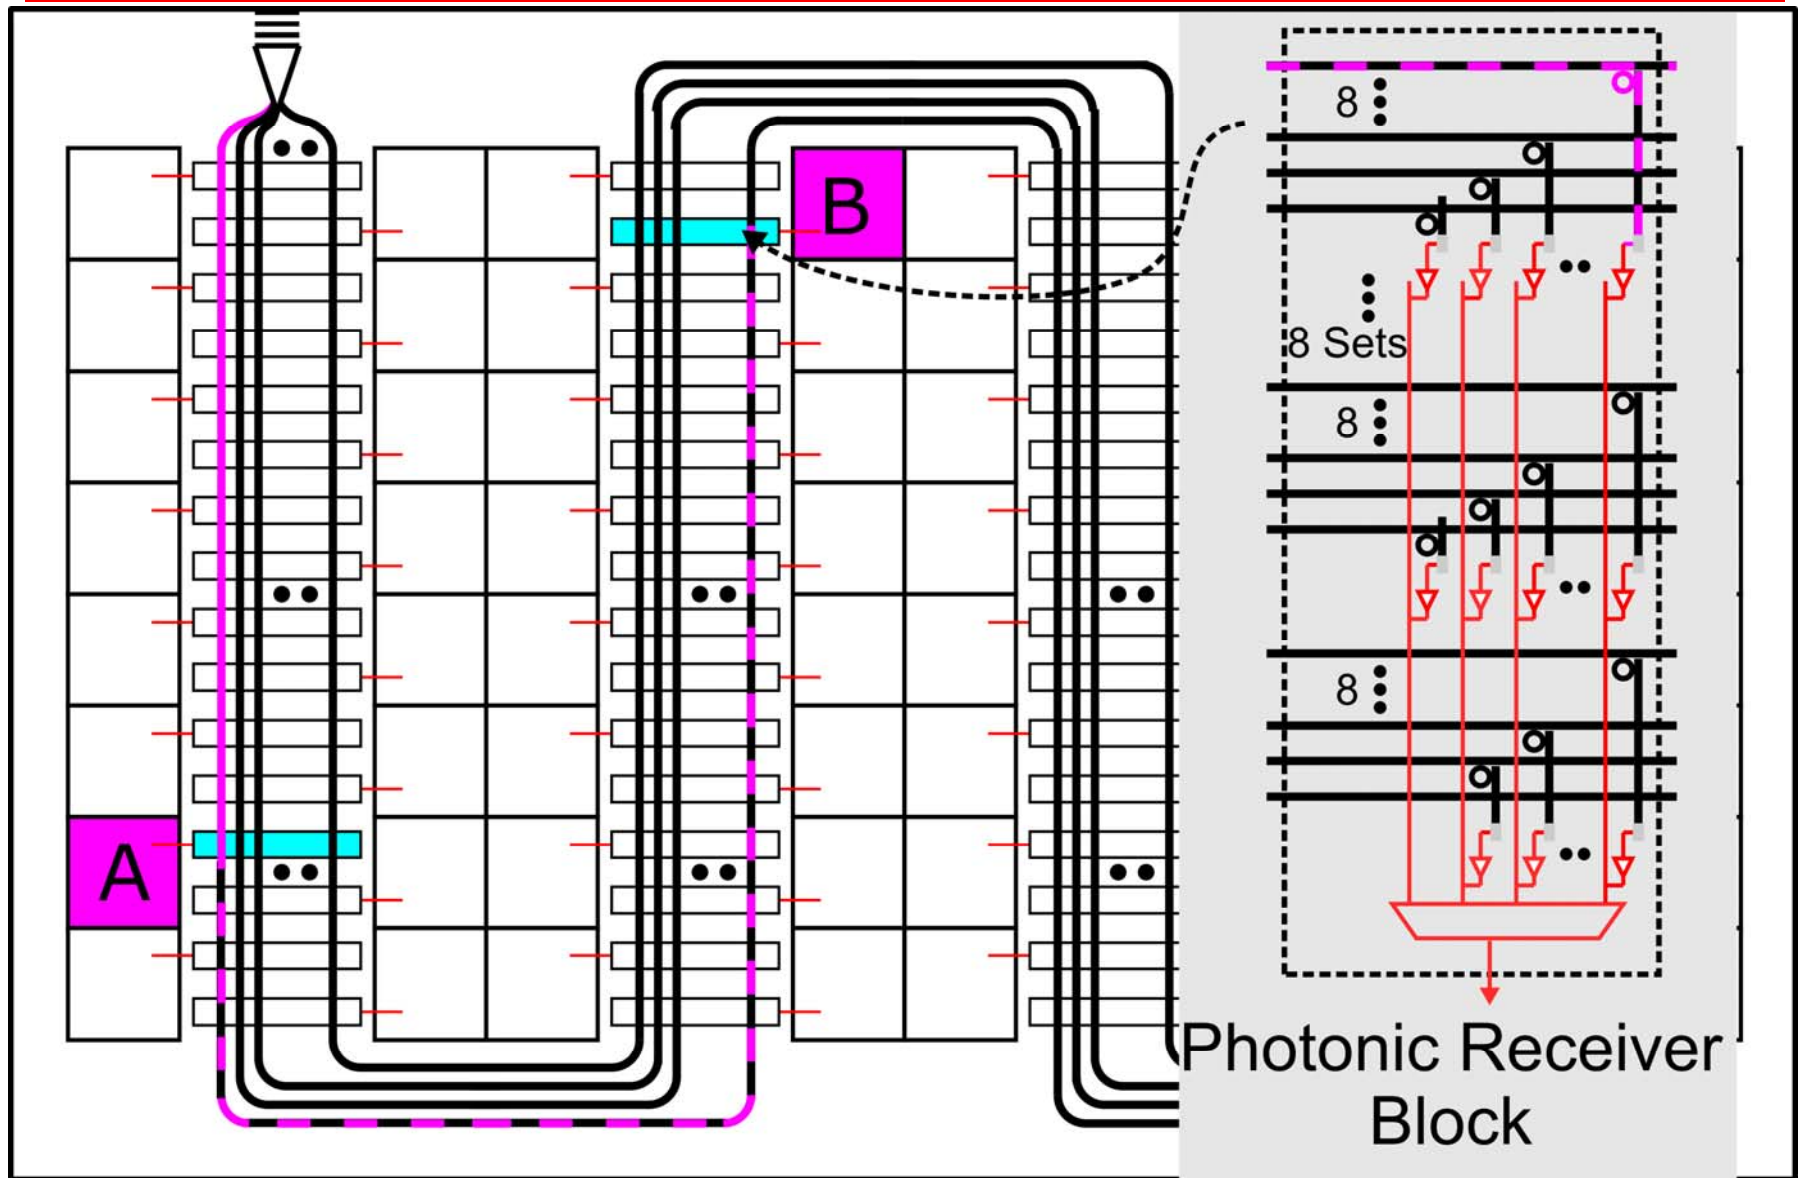
Photonic crossbar for a 64-tile system



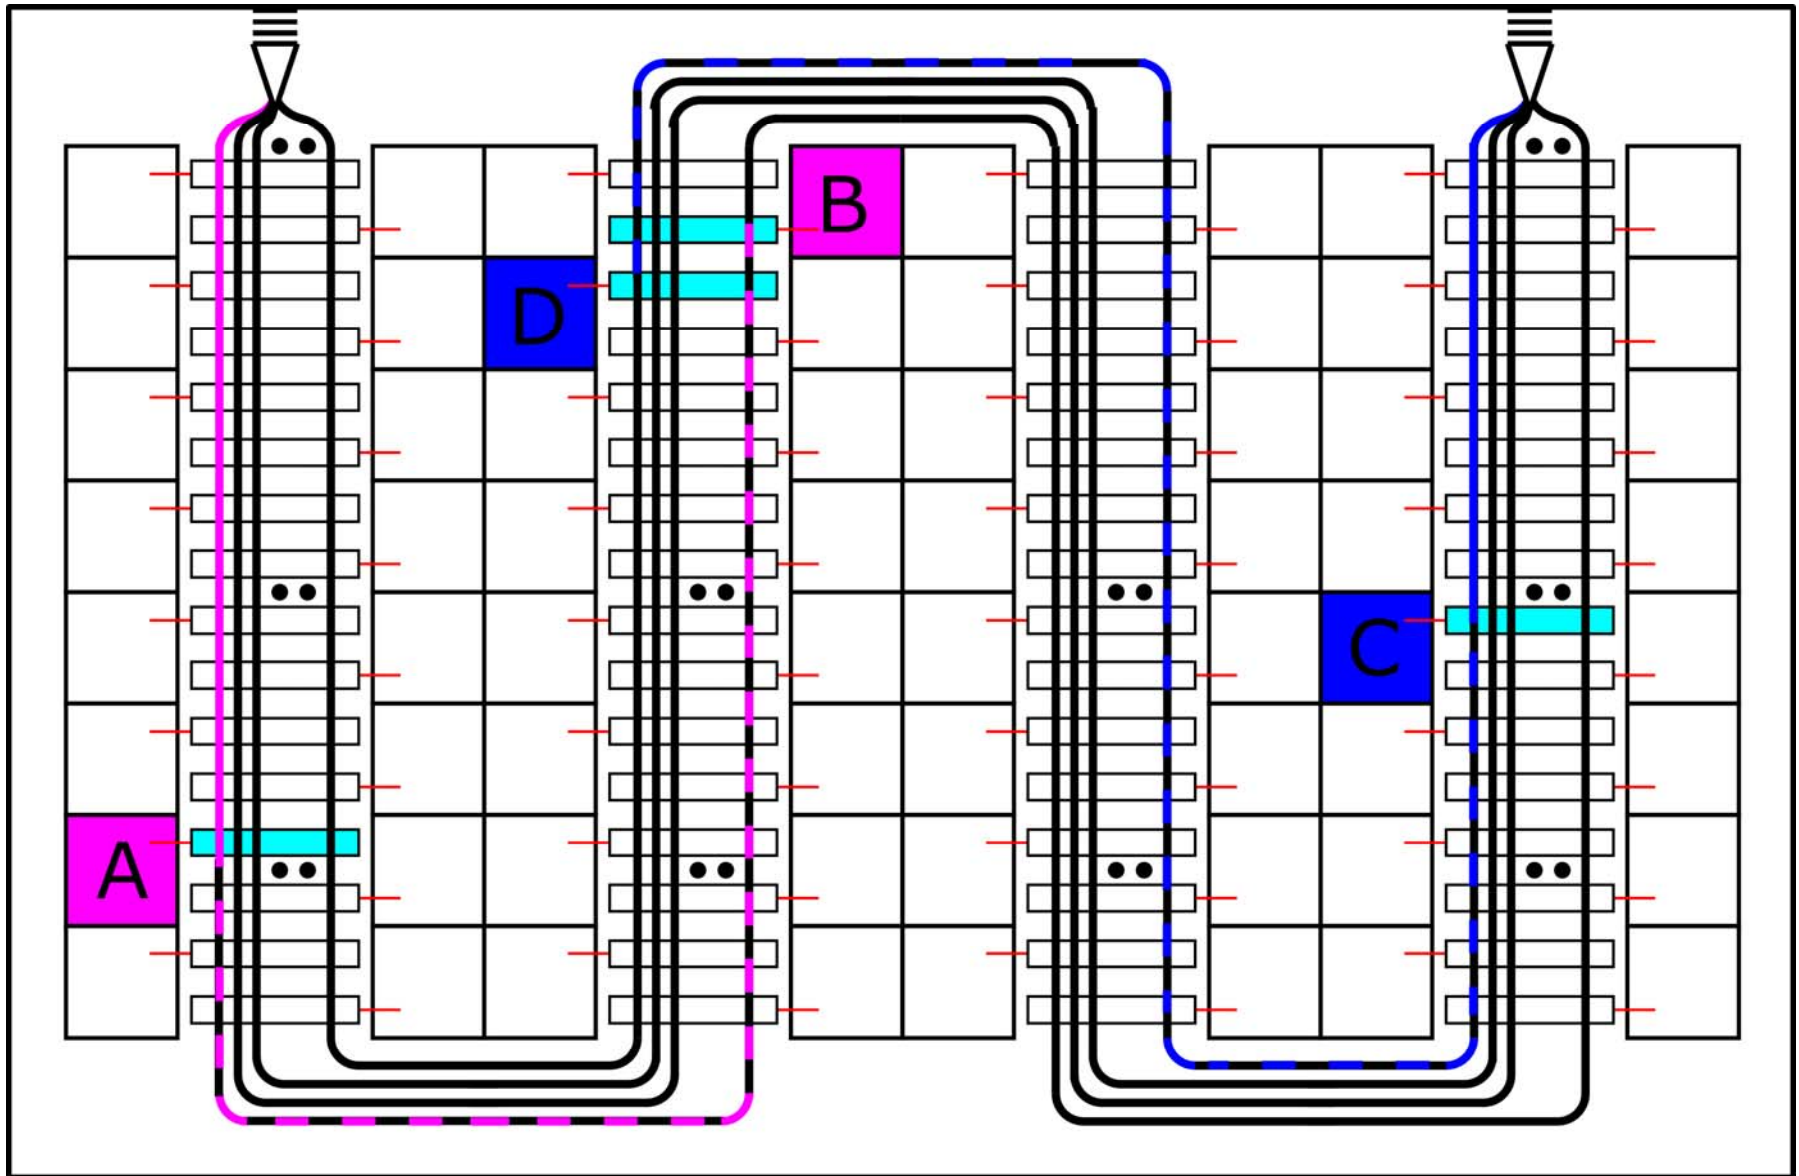
Photonic crossbar for a 64-tile system



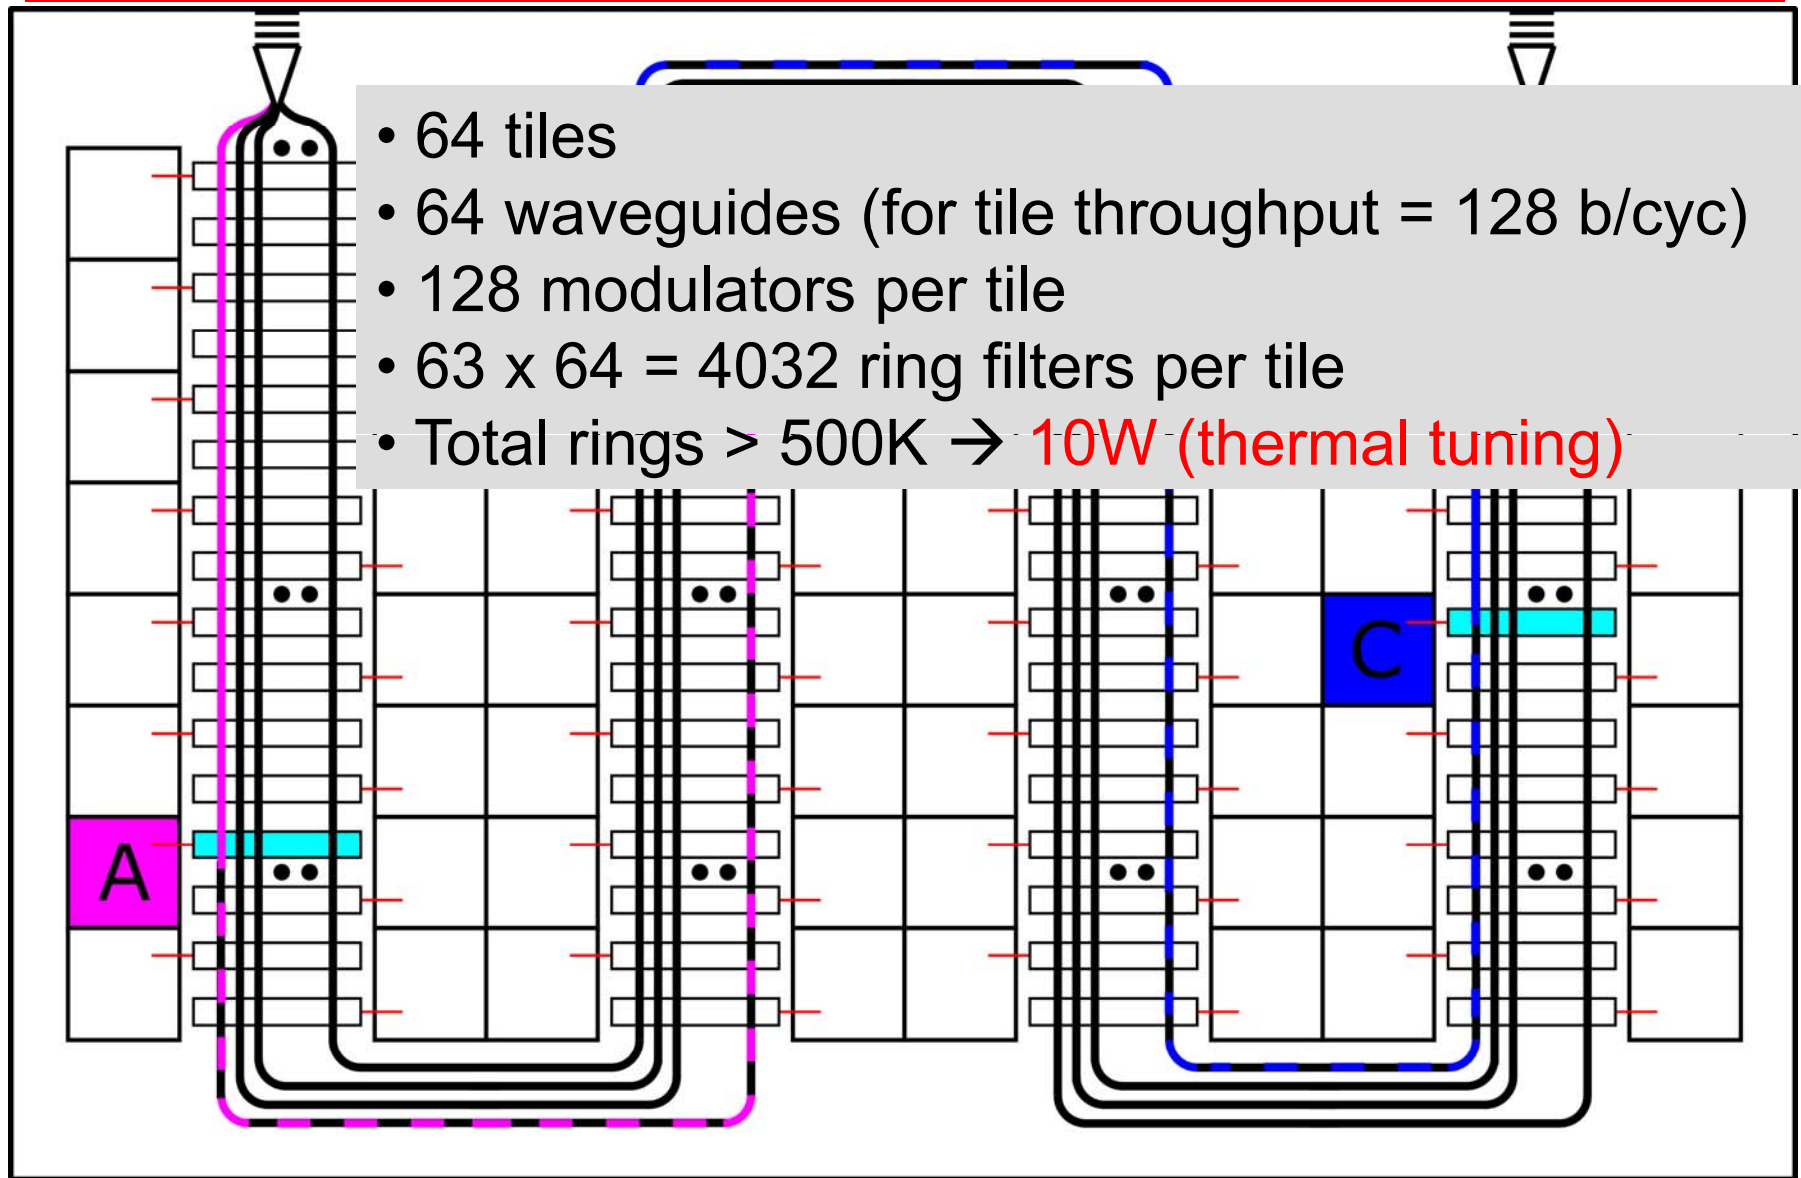
Photonic crossbar for a 64-tile system



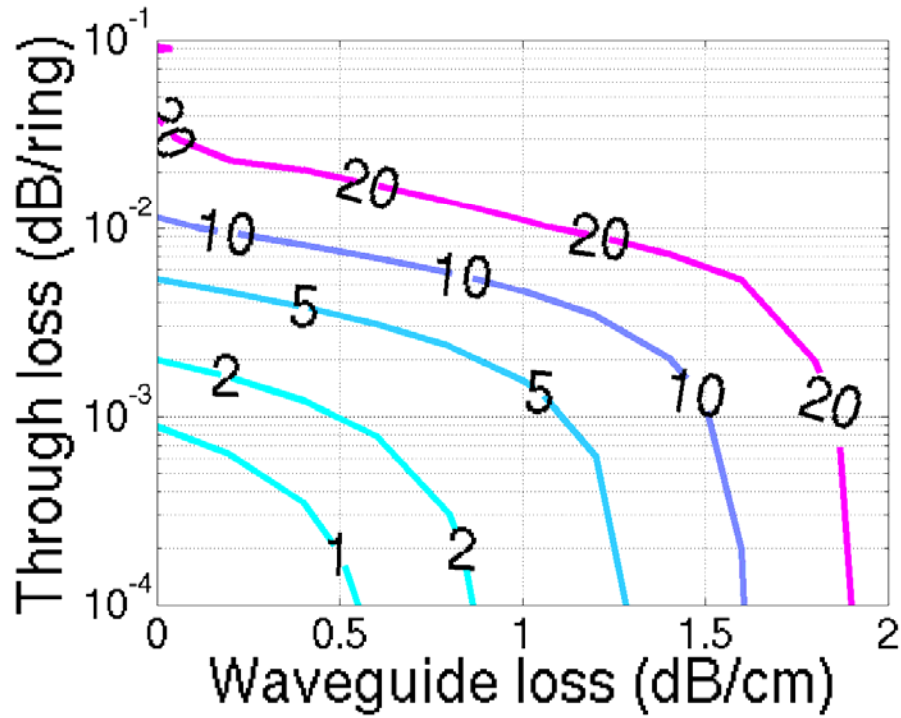
Photonic crossbar for a 64-tile system




Photonic crossbar for a 64-tile system



Photonic device requirements in a crossbar




Optical laser power (W) contour




Percent area of photonic devices contour

Photonic device requirements in a crossbar



Optical laser power (W) contour




Percent area of photonic devices contour




Waveguide loss and Through loss limits for 2 W optical laser power (30% laser efficiency) constraint

Outline




- ❑ Interconnect technologies
- ❑ **Photonic networks**
- ❑ Electrical vs Photonic networks

Clos network using point-to-point channels

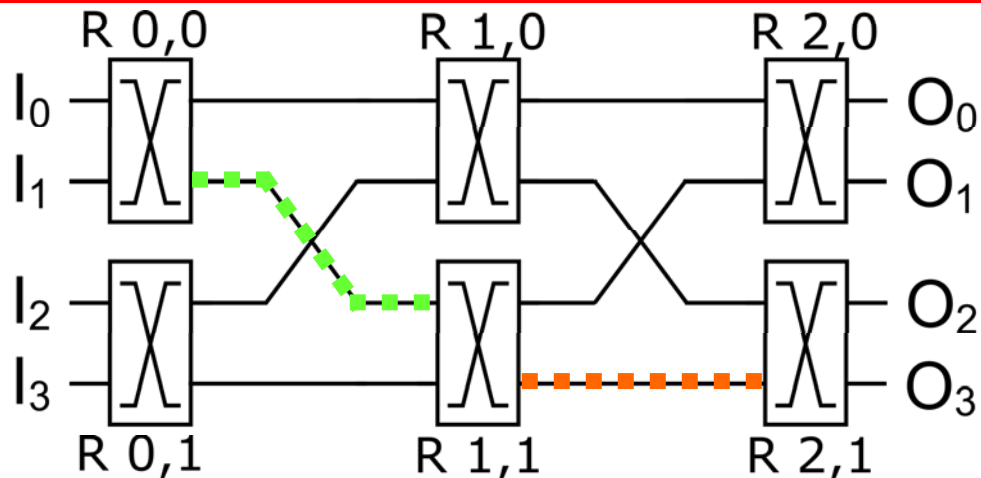


Electrical design




Photonic design

Clos network using point-to-point channels




Electrical design

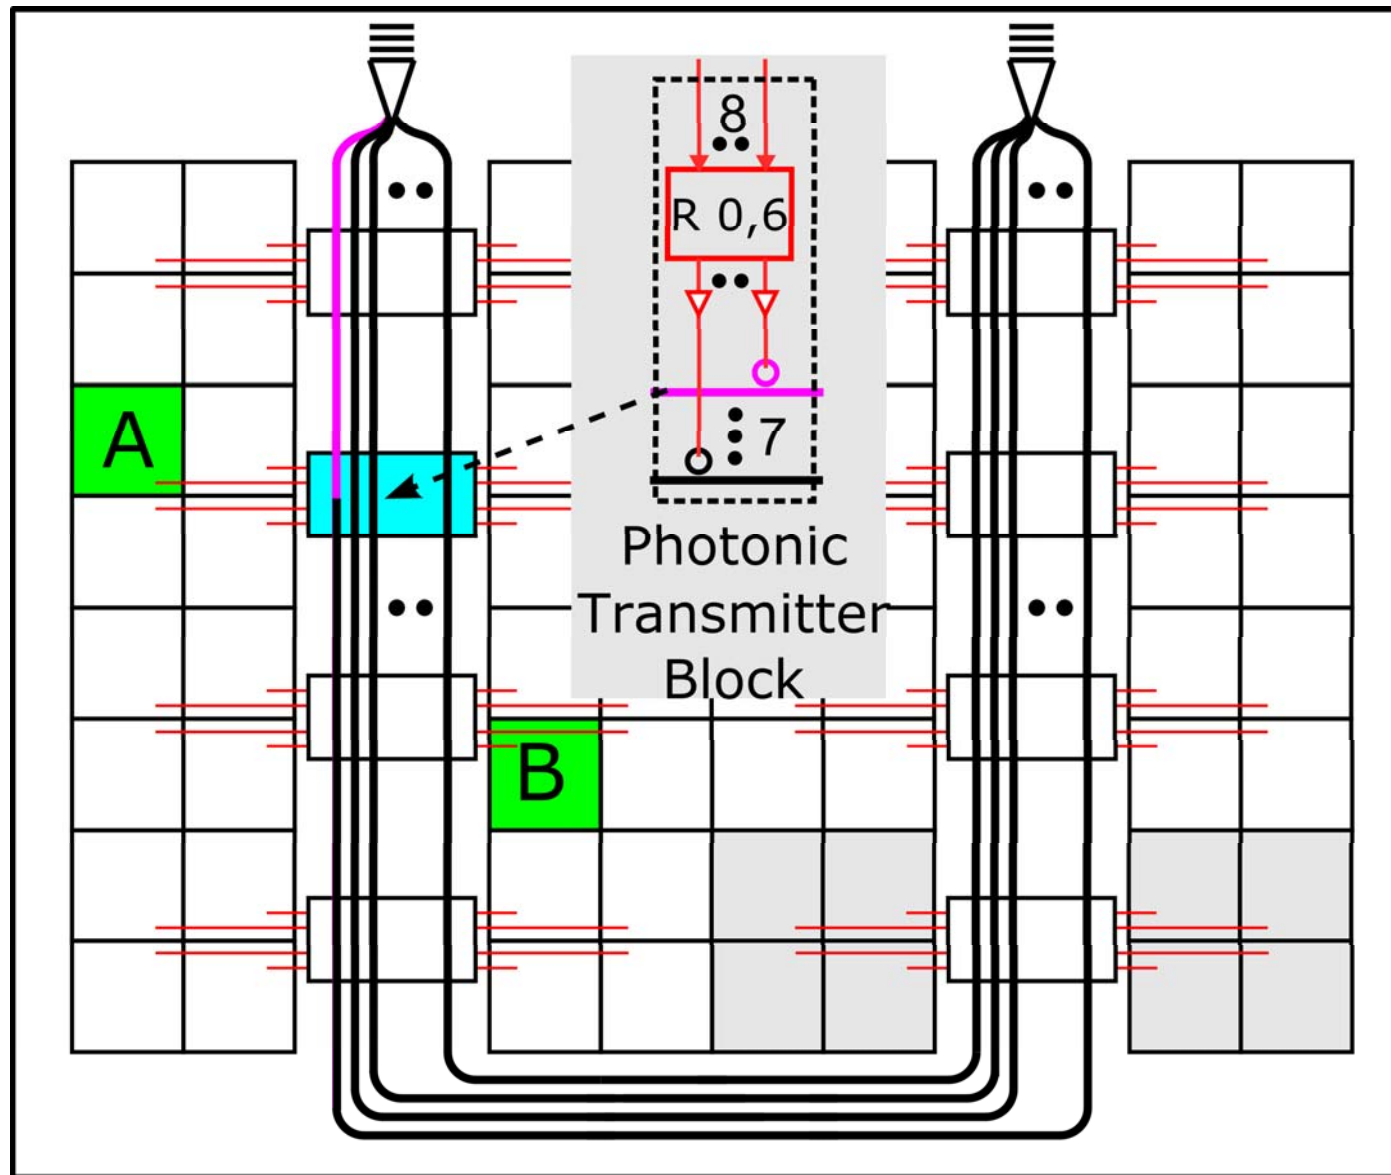


Photonic design

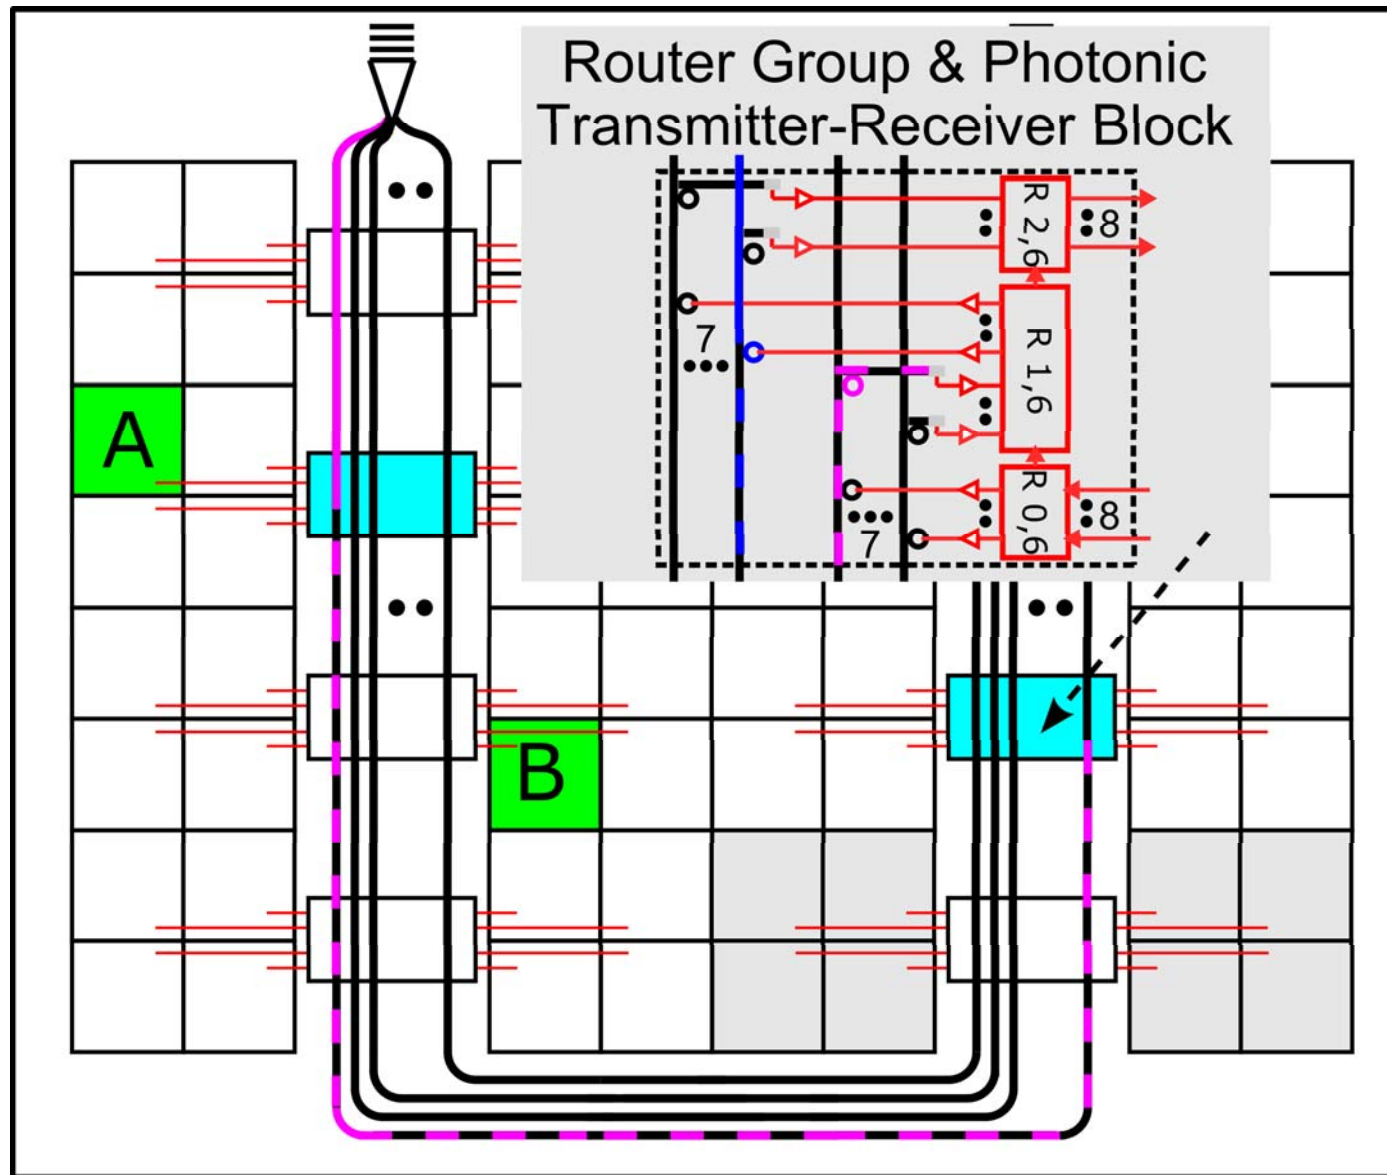
Photonic Clos for a 64-tile system



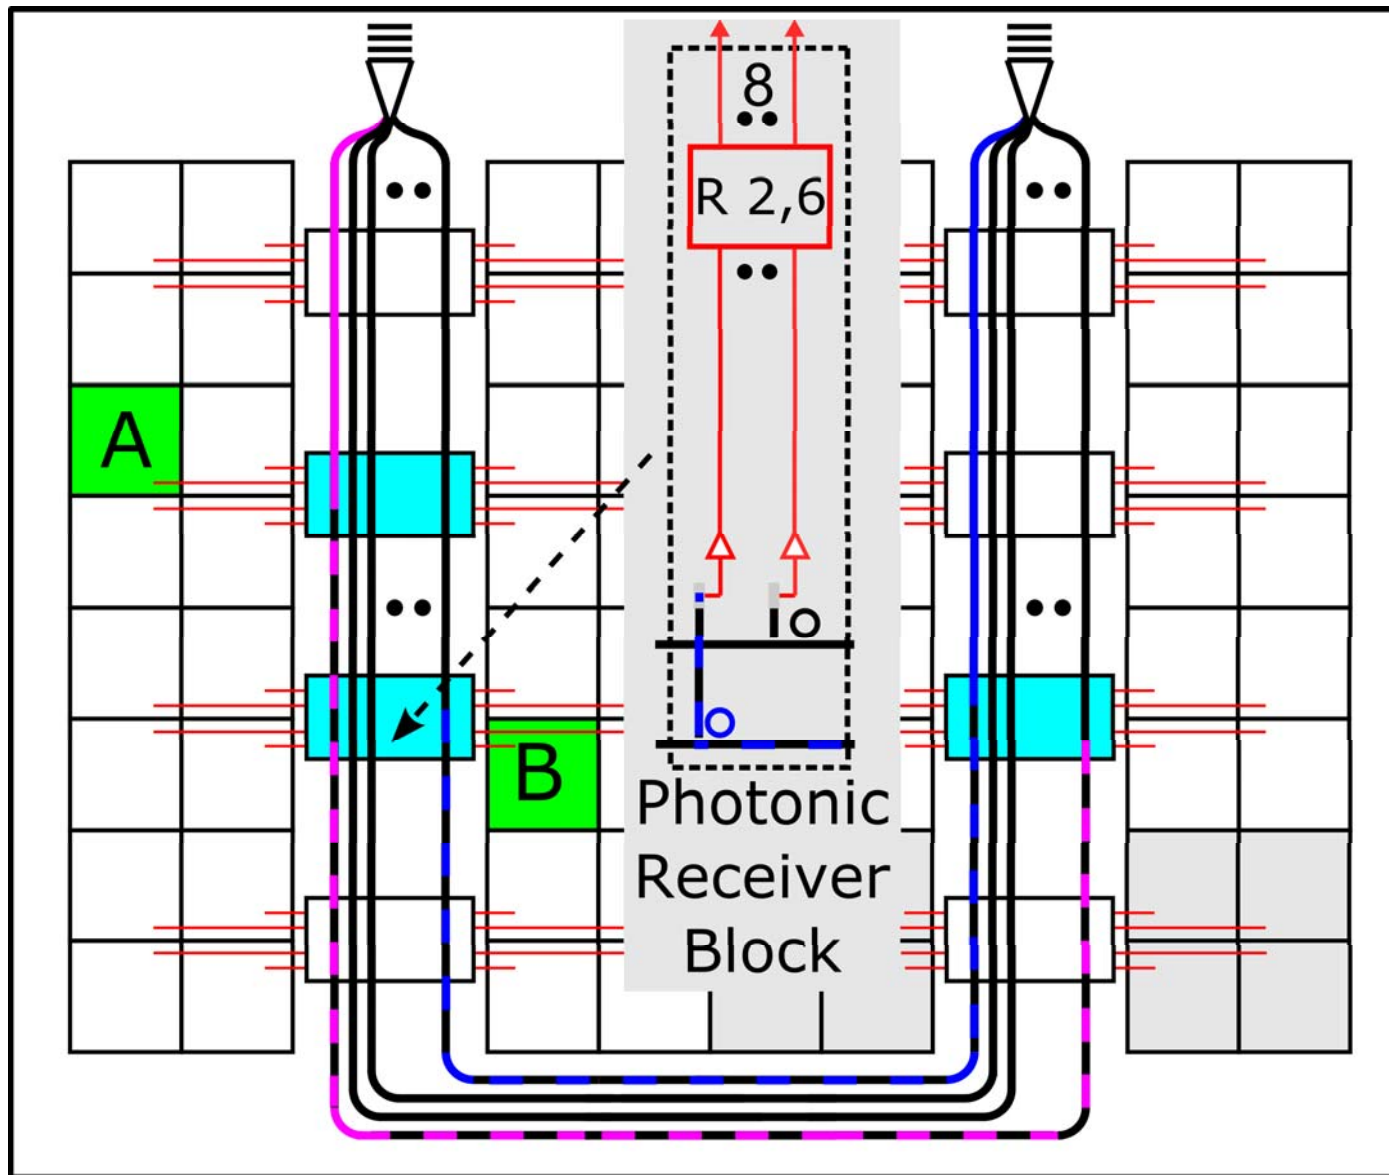
Photonic Clos for a 64-tile system




Photonic Clos for a 64-tile system



Photonic Clos for a 64-tile system




Photonic Clos for a 64-tile system




- 64 tiles
- 56 waveguides (for tile throughput = 128 b/cyc)
- 128 modulators per cluster
- 128 ring filters per cluster
- Total rings $\approx 28K \rightarrow 0.56W$ (Thermal tuning)

Photonic device requirements in a Clos



Optical laser power (W) contour




Percent area of photonic devices contour




Waveguide loss and Through loss limits for 2 W optical laser power (30% laser efficiency) constraint

Photonic device requirements in a Clos



Optical laser power (W) contour



Percent area of photonic devices contour



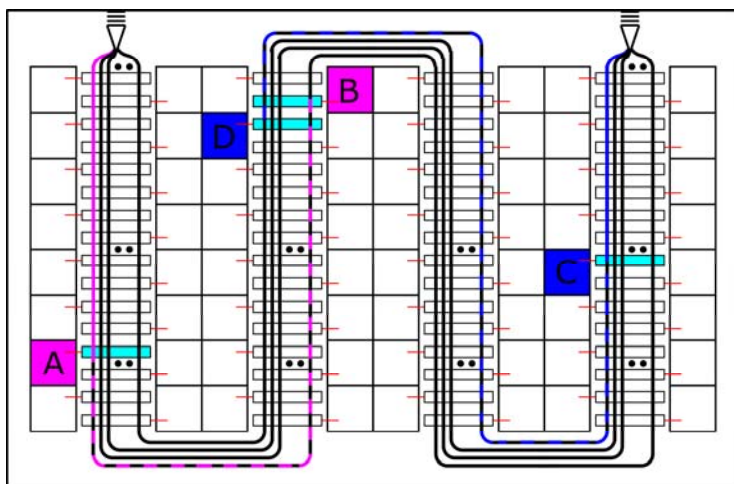
Optical loss tolerance for Crossbar




Optical loss tolerance for Clos

Photonic Crossbar vs Photonic Clos

Crossbar




Clos




- ❑ 10 W power for thermal tuning circuits ($1 \mu\text{W}/\text{ring}/\text{K}$)
- ❑ For 2 W optical laser power
 - Waveguide loss $< 1 \text{ dB}/\text{cm}$
 - Through loss $< 0.002 \text{ dB}/\text{ring}$

- ❑ 0.56 W power for thermal tuning circuits ($1 \mu\text{W}/\text{ring}/\text{K}$)
- ❑ For 2 W optical laser power
 - Waveguide loss $< 2 \text{ dB}/\text{cm}$
 - Through loss $< 0.05 \text{ dB}/\text{ring}$


Outline




Mesh



CMesh



Clos




Crossbar


- ❑ Photonic interconnect technology
- ❑ Photonic networks
- ❑ **Electrical vs Photonic networks**

Simulation setup

- ❑ Cycle-accurate microarchitectural simulator
- ❑ Traffic patterns based on partition application model
 - Global traffic – UR, P2D, P8D
 - Local traffic – P8C
- ❑ 64-tile system, 512-bit messages
- ❑ Events captured during simulations to calculate power

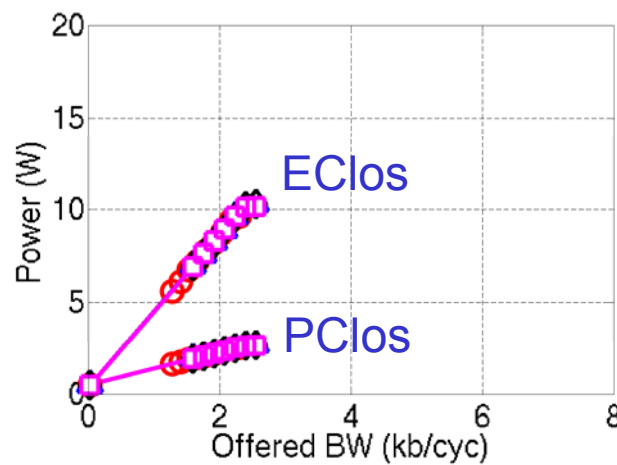


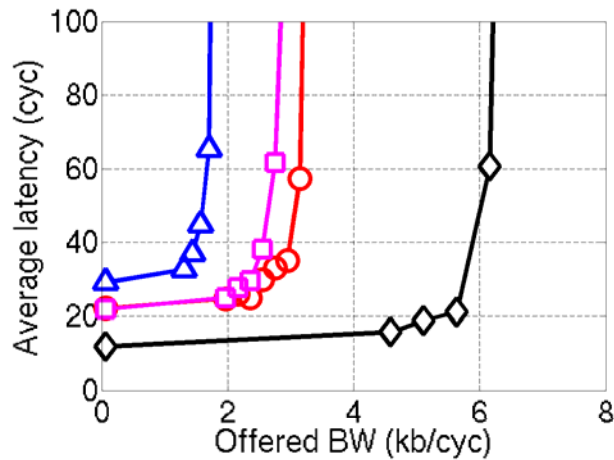


CMesh



Clos

Power-Bandwidth tradeoff





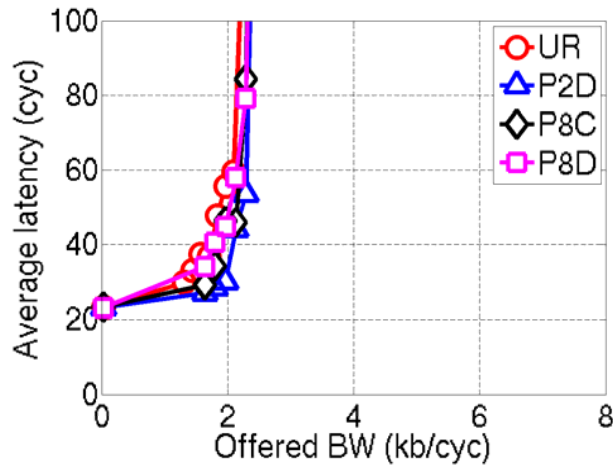

CMeshX2

Channel width = 128b

Clos

Channel width = 64b

Power-Bandwidth tradeoff








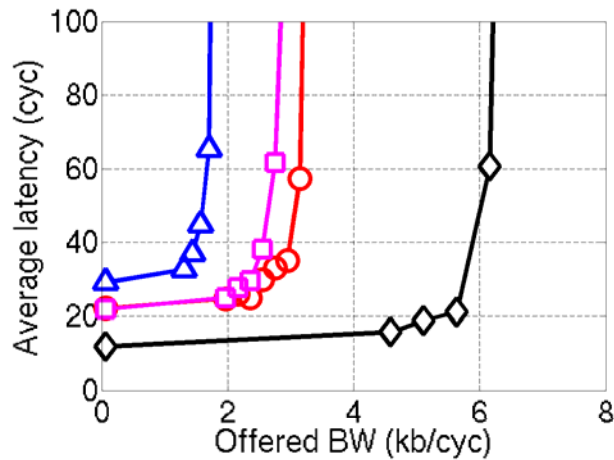
CMeshX2

Channel width = 128b

Clos

Channel width = 64b

Power-Bandwidth tradeoff



CMeshX2

Channel width = 128b

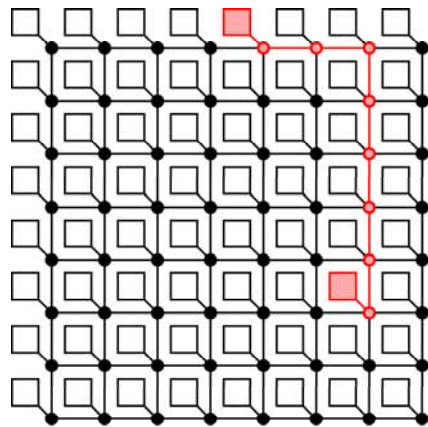
Clos

Channel width = 64b


Clos

Channel width = 128b


Conclusion




Mesh



CMesh



Clos



Crossbar


- ❑ Accurate baseline electrical design required
- ❑ Need to carefully account for the energy components in optical interconnects
 - E-O-E conversion, Thermal tuning power, Optical laser power
- ❑ Clos network provides comparable throughput at lower energy for global traffic patterns
- ❑ More work required on the photonic device design

Acknowledgement


- MIT photonic device team
 - Franz Kärtner, Rajeev Ram, Judy Hoyt, Henry Smith
 - Jason Orcutt, Anatoly Khilo, Benjamin Moss, Charles Holzwarth, Jonathan Leu, Michael Georgas, Jie Sun, Miloš Popović, Hanqing Li
- Funding sources
 - DARPA
 - Intel Corp.

Backup

Clos network using intermediate crossbar




Electrical design




Photonic design

Clos network using intermediate crossbar




Electrical design




Photonic design

Clos network using intermediate crossbar



Electrical design



Photonic design