MODELING AND EVALUATION OF CHIP-TO-CHIP SCALE SILICON PHOTONIC NETWORKS

Robert Hendry, Dessislava Nikolova,
Sébastien Rumley, Keren Bergman

Columbia University

HOTI 2014

Chip-to-chip optical networks

- Projected chip I/O bandwidth: tens of Tb/s
- Chip I/O bandwidth limited by pin count, data rate
- Promising solution: silicon photonics
  - Dense bandwidth via WDM
  - High data rates
  - Energy-distance independence in fiber

Outline

- **Silicon photonic chip-to-chip networks**
- Characterizing loss and WDM capacity
- Modeling power
- Determining network performance
- Conclusions

Microring-based silicon photonic links

- Microrings
  - Modulation
  - Switching
  - Filtering

Microring modulator (Cornell)

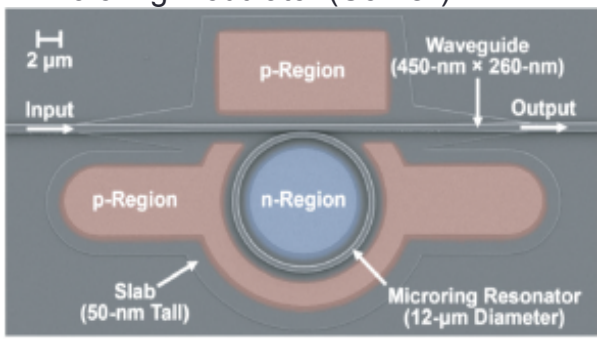

Demultiplexing filter (Kotura/ Oracle)

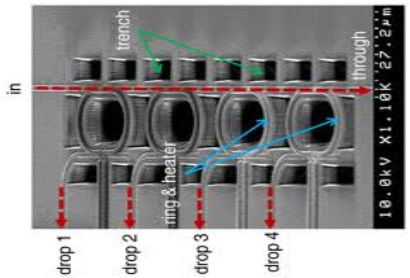

- Other optical devices:
  - Lasers, couplers, integrated photodetectors

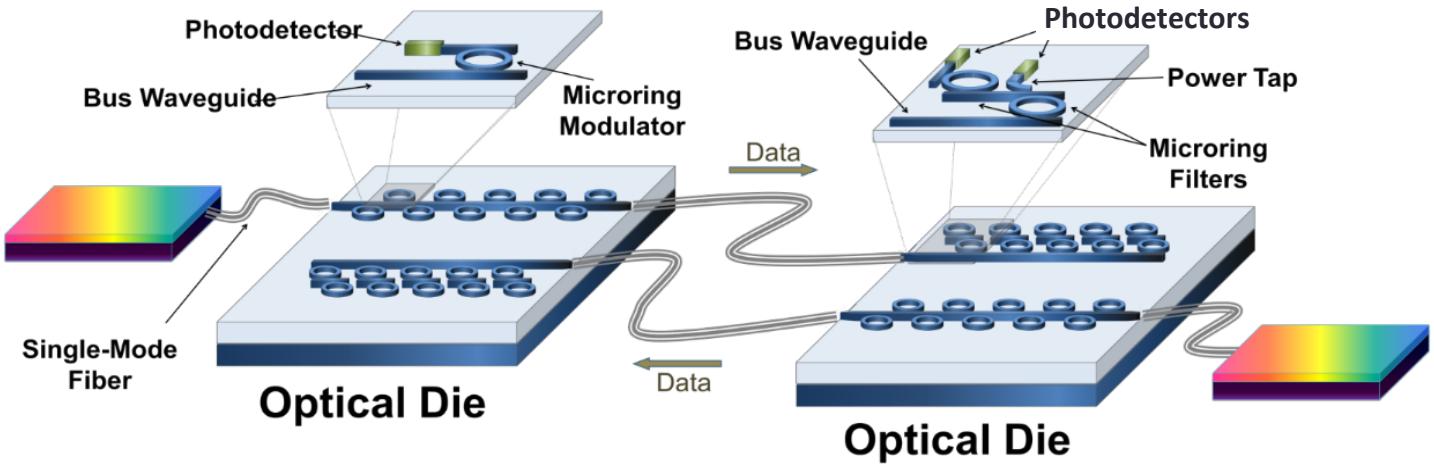

Chip-to-chip optical networks

- "Chip-to-chip" low radix, high bandwidth
- Chose two architectures to represent extremes of design space

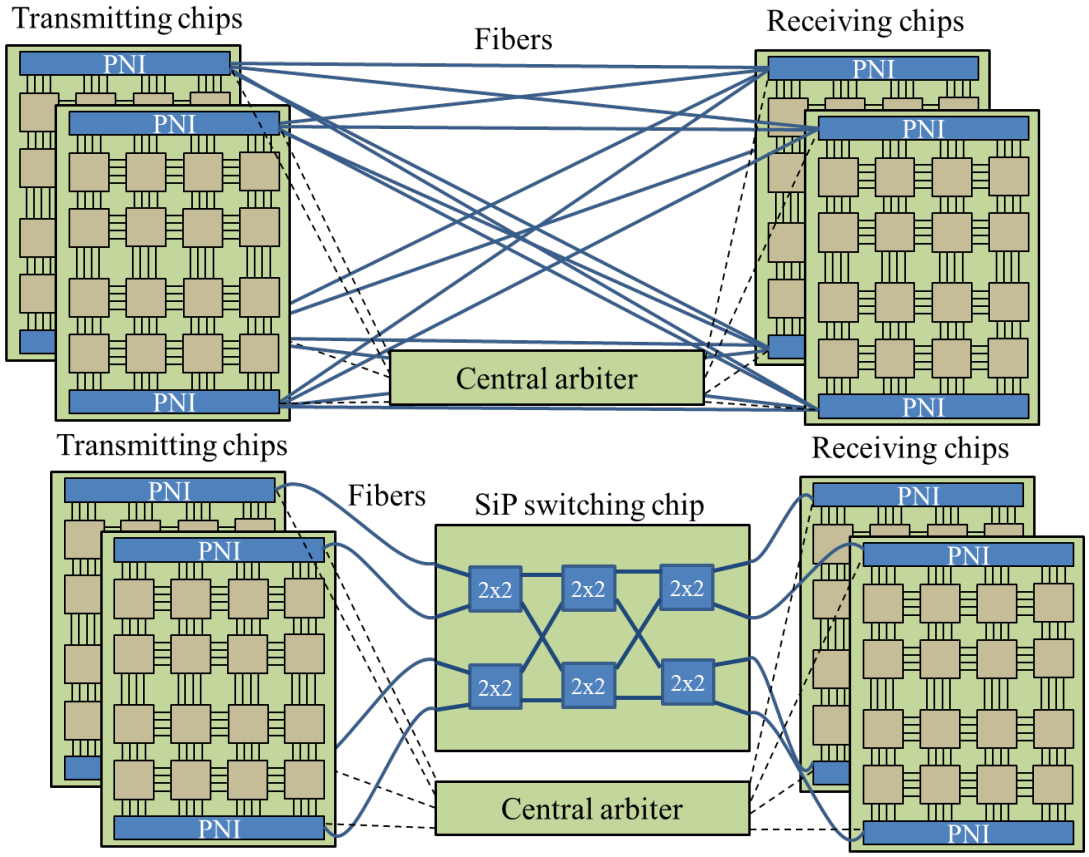

Full mesh architecture

Switched architecture

Full mesh

- One link at each source for each destination

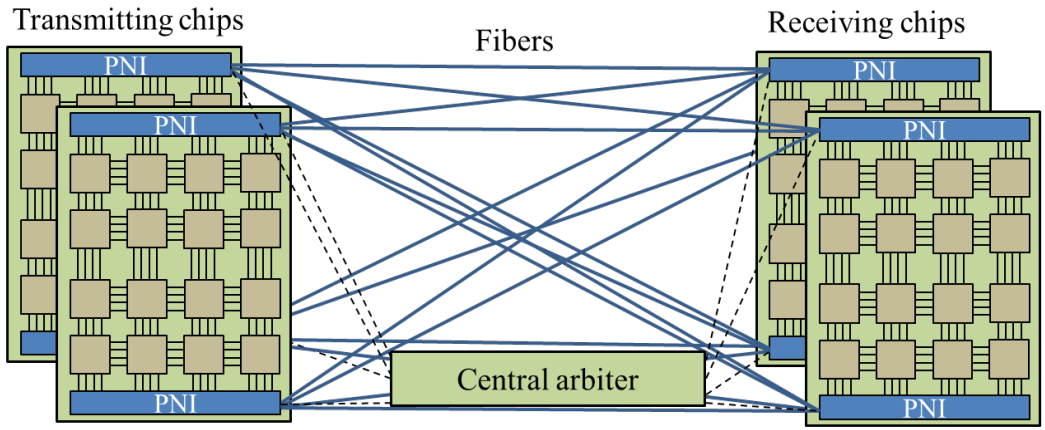

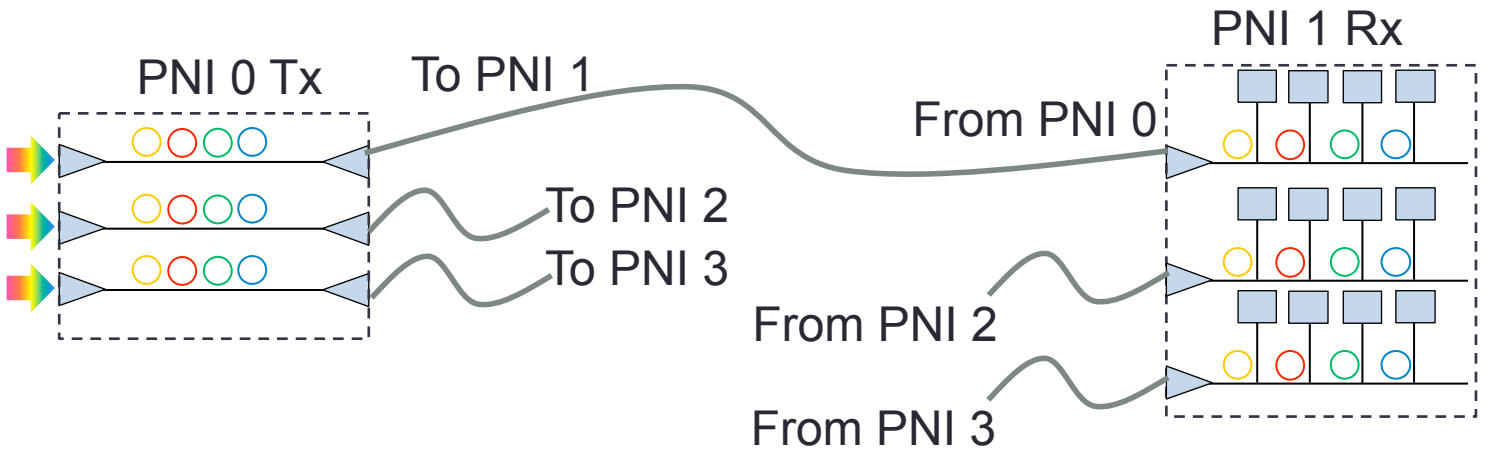

Switched architecture

- One input and output link per PNI

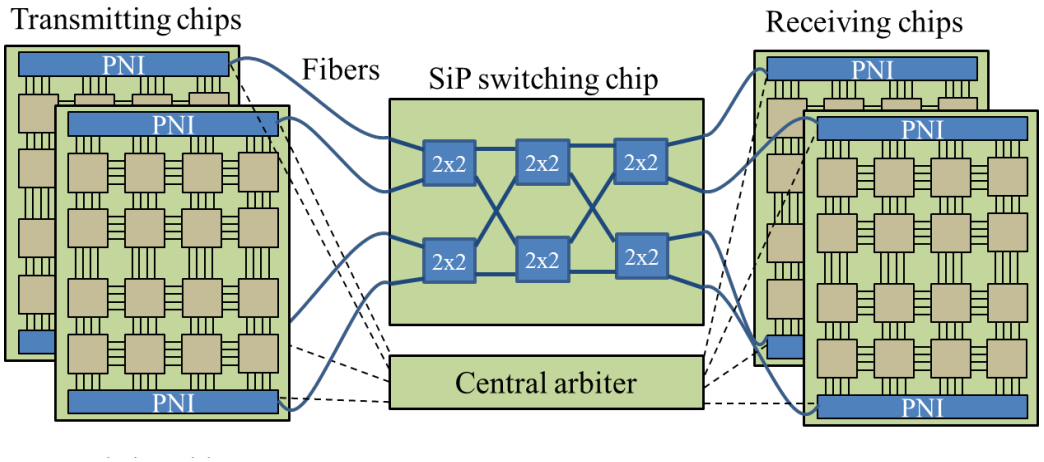

2x2 switch

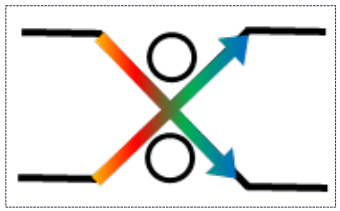

"Through" state

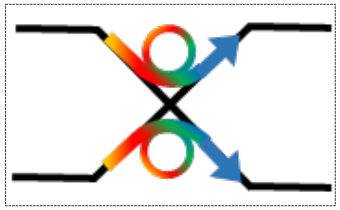

"Drop" state

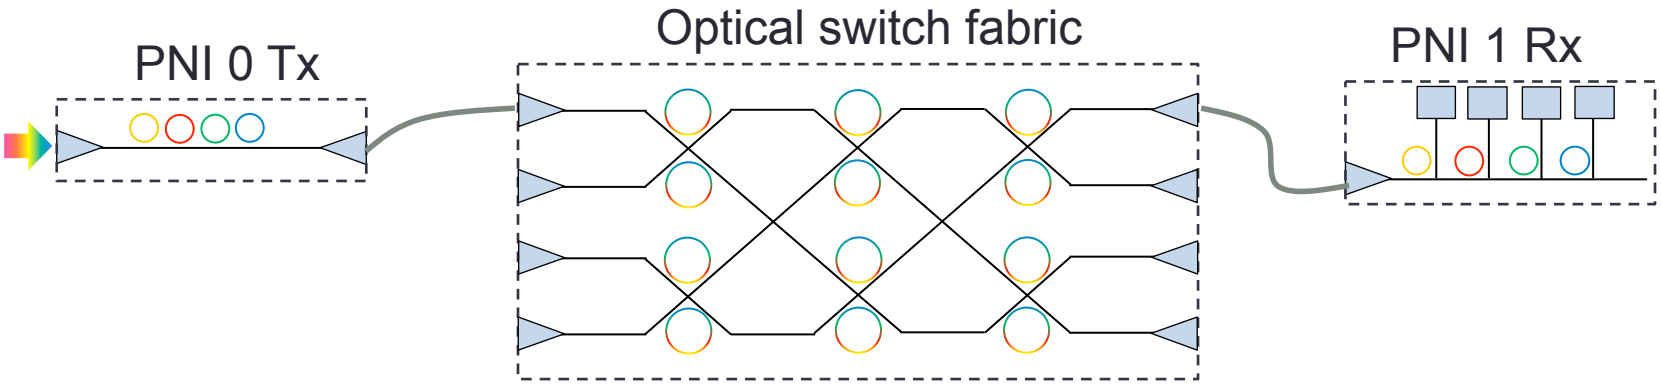

Comparing topologies

- **Laser power is the largest contributor to overall power in the network**
  - 4.5%, X. Zheng, et al. “Efficient WDM laser sources towards terabytes/s silicon photonic interconnects.” *Journal of Lightwave Technology,* vol. 31, no. 15, 2013.
- **Assume lasers are always on**
  - **Laser stabilization time on the order of microseconds**
  - **Context: small packets, short inter-arrivals**
- **Energy efficiency closely related to utilization of laser sources**
- Full mesh expectation:
  - No contention, lower queuing latency
  - More lasers, higher power, poor efficiency with load is low
- Switched architecture expectation:
  - Contention, higher queuing latency
  - Resource sharing improves utilization and therefore efficiency

Shared input/output waveguides

- Another way to share laser sources

...

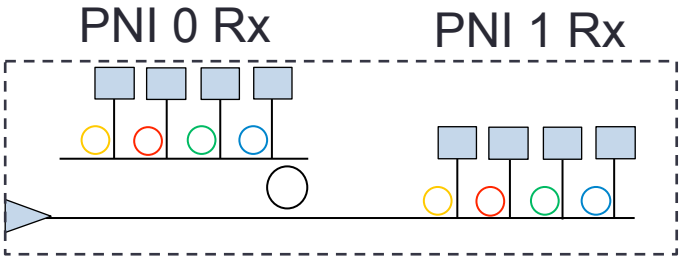

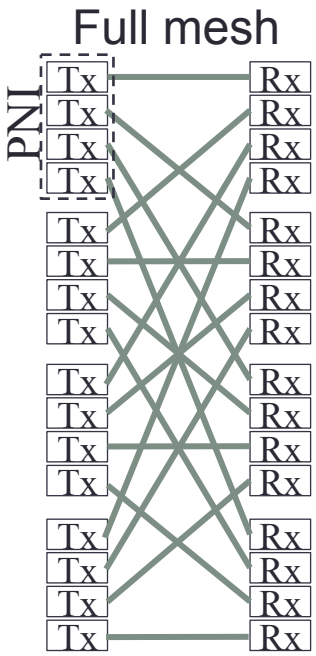

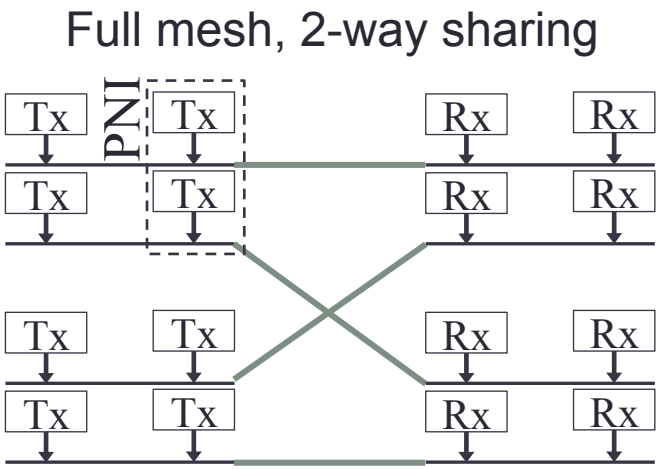

- Sacrificing performance for better utilization

Shared input/output waveguides

- Another way to share laser sources

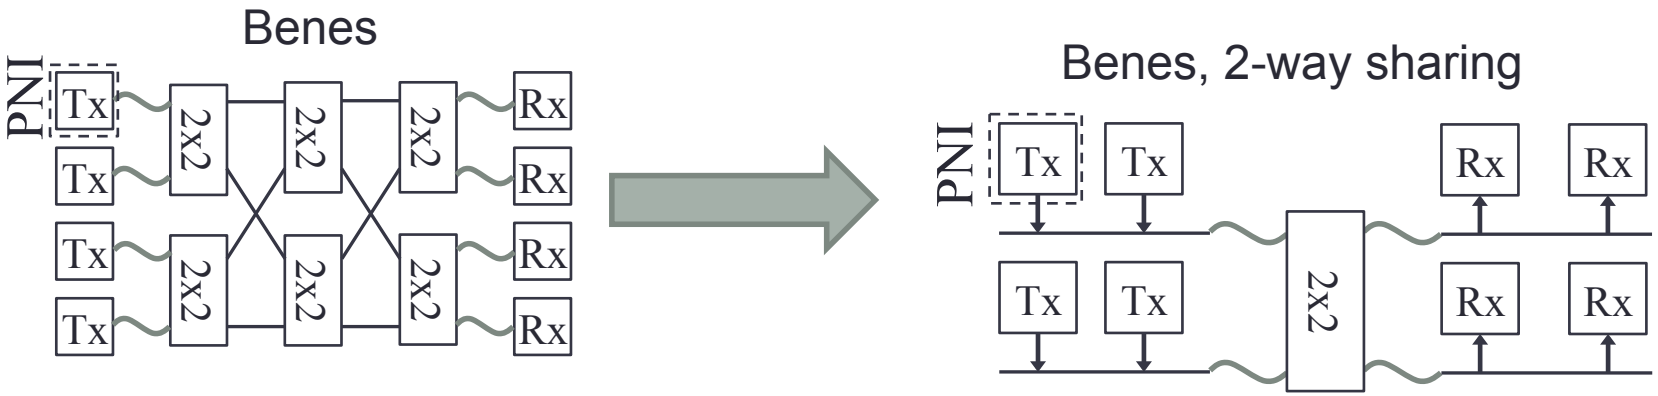

Design space

- Topology
  - Benes
  - Full mesh
- Sharing
  - No sharing, or two-way sharing
- Network radix
  - 4, 8, or 16
- **Goal: to find optimal topologies for given bisectional bandwidth requirements**
- Ex: “Benes-4T-2S”, “FM-16T-1S”

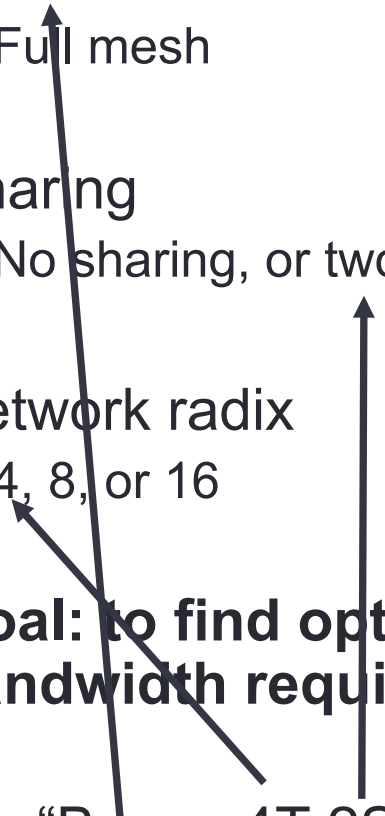

Outline

- Silicon photonic chip-to-chip networks
- **Characterizing loss and WDM capacity**
- Modeling power
- Determining network performance
- Conclusions

Determining worst-case loss

- Combine losses of all devices along worst-case path of light

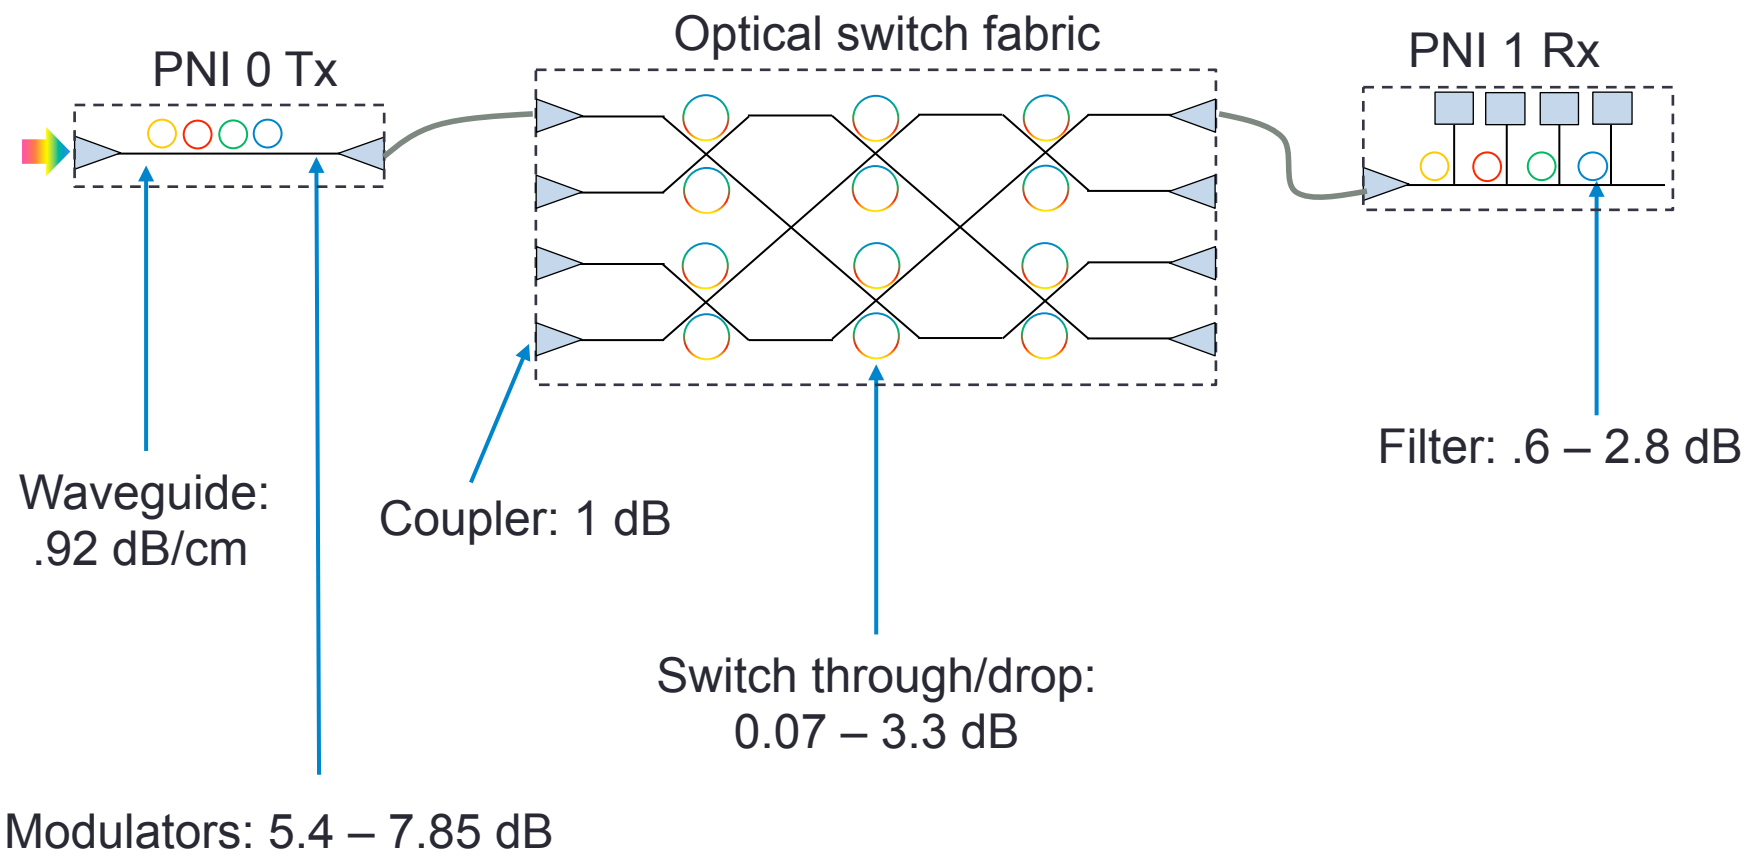

Complexity vs. link capacity

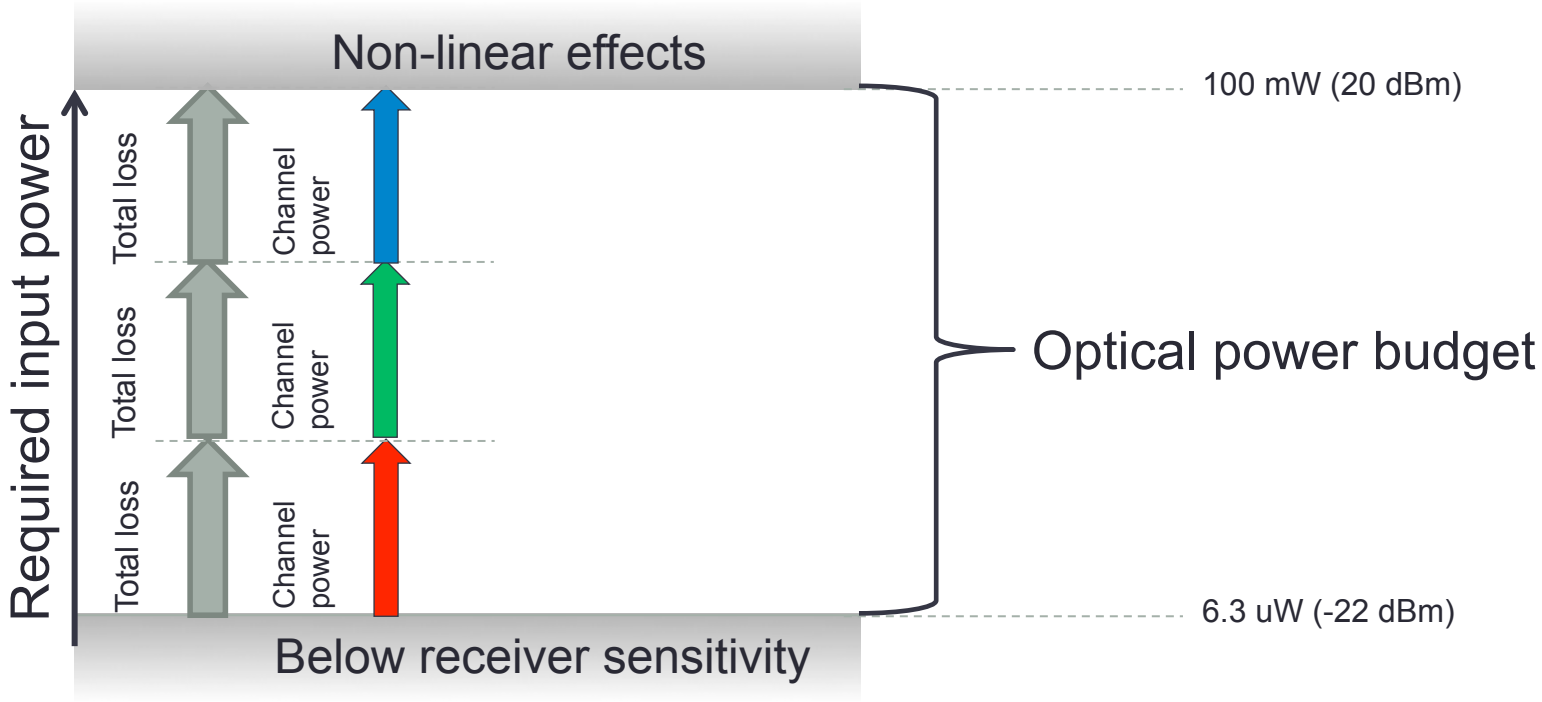

- 10 Gb/s per wavelength channel, OOK modulation
- Intermodulation crosstalk limits WDM capacity to 125 wavelengths
  - Assuming 50nm spectrum
  - K. Padmaraju, et al. “Intermodulation Crosstalk Characteristics of WDM Silicon Microring Modulators.” *IEEE Photonics Letters*, vol. 26, no. 14, 2014.

Peak Bisectional Bandwidth

- Loss → maximum wavelengths per link
- Maximum wavelengths *x* 10 Gb/s *x* number of links in bisection → peak bisectional bandwidth

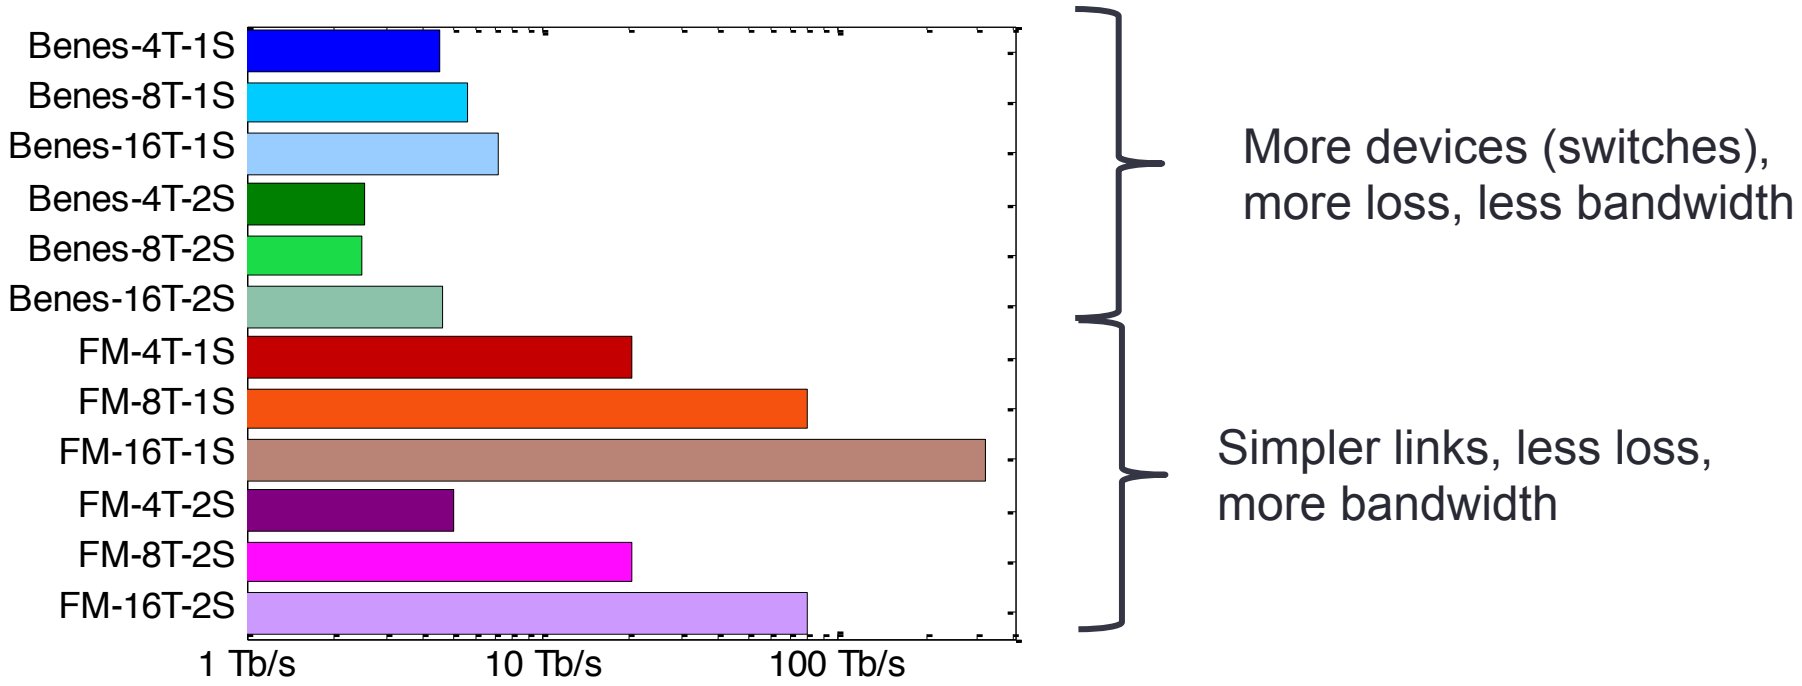

T = number of PNIs, S = number of Tx/Rx per link

Peak Bisectional Bandwidth

- Loss → maximum wavelengths per link
- Maximum wavelengths *x* 10 Gb/s *x* number of links in bisection → peak bisectional bandwidth

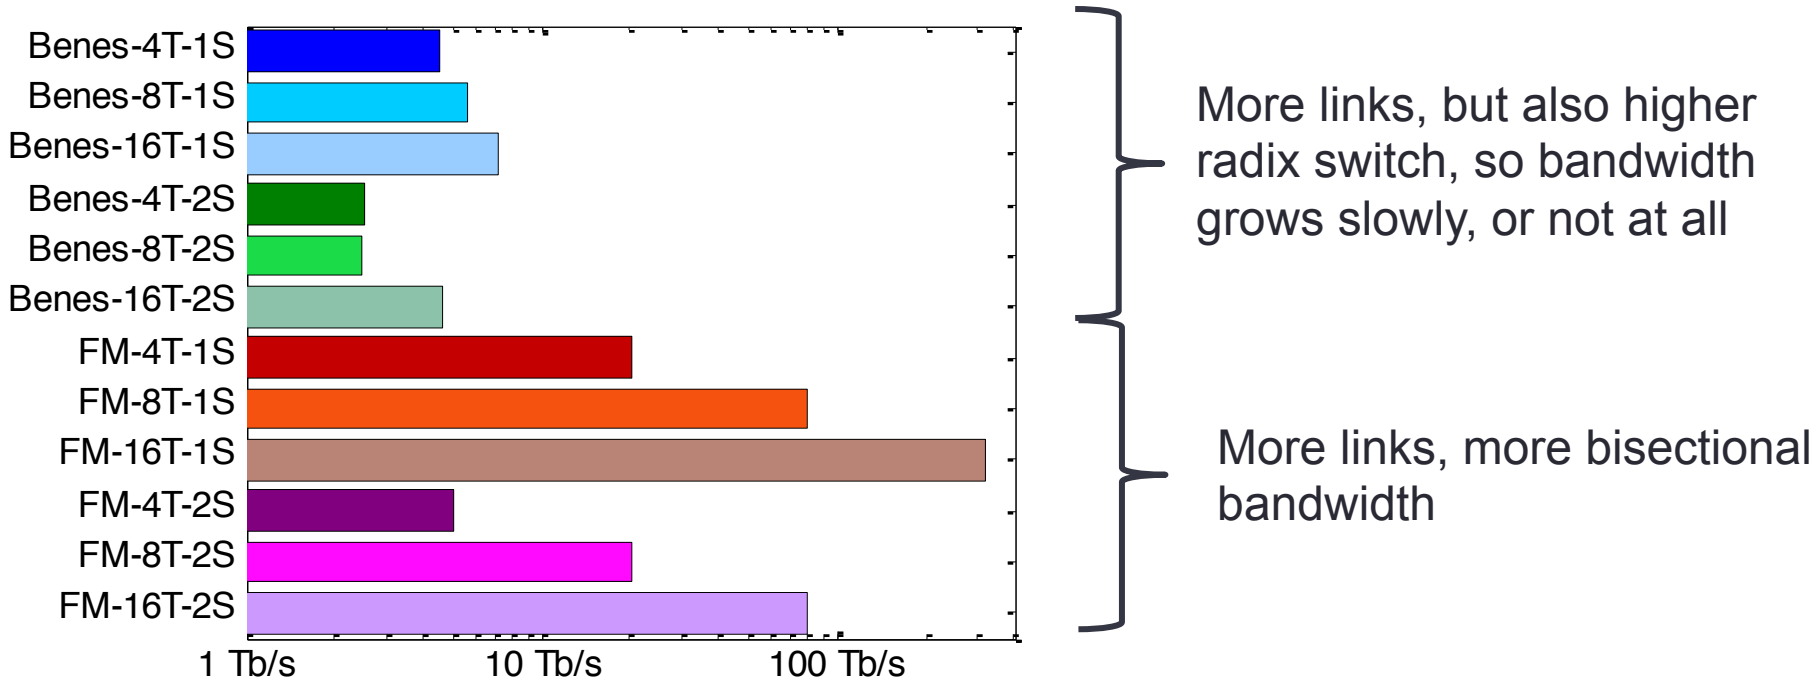

T = number of PNIs, S = number of Tx/Rx per link

Outline

- Silicon photonic chip-to-chip networks
- Characterizing loss and WDM capacity
- **Modeling power**
- Determining network performance
- Conclusions

Power modeling

- Microring tuning, trimming
  - Thermal fluctuations
  - Imperfect fabrication
- Laser power
  - A function of loss and number of wavelengths used
- Static dissipation in photodetectors
- Dynamic modulation, switching power
- Not modeling network interfaces

Device	Type/Origin	Power/Device (mW)
Modulator	Thermal	0.875
	Driver circuitry	1.35
	Dissipation in ring	0.1
Switch	Thermal	3.5
Filter	Thermal	0.875
Detector	Static	3.95
Laser	Static	1250

Outline

- Silicon photonic chip-to-chip networks
- Characterizing loss and WDM capacity
- Modeling power
- **Determining network performance**
- Conclusions

Impact of layout on network performance

- Poisson arrivals, uniform random destination
- Fixed message size (256B)
- Assume we have an arbitration scheme that can reach 100% utilization across the chip-scale network
  - Models indicate queuing and head-to-tail latency

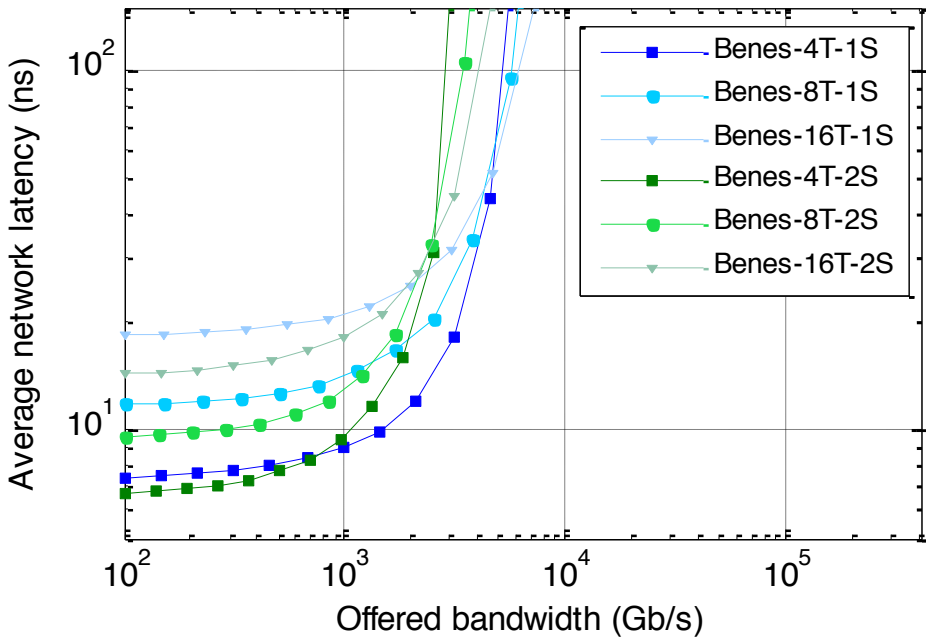

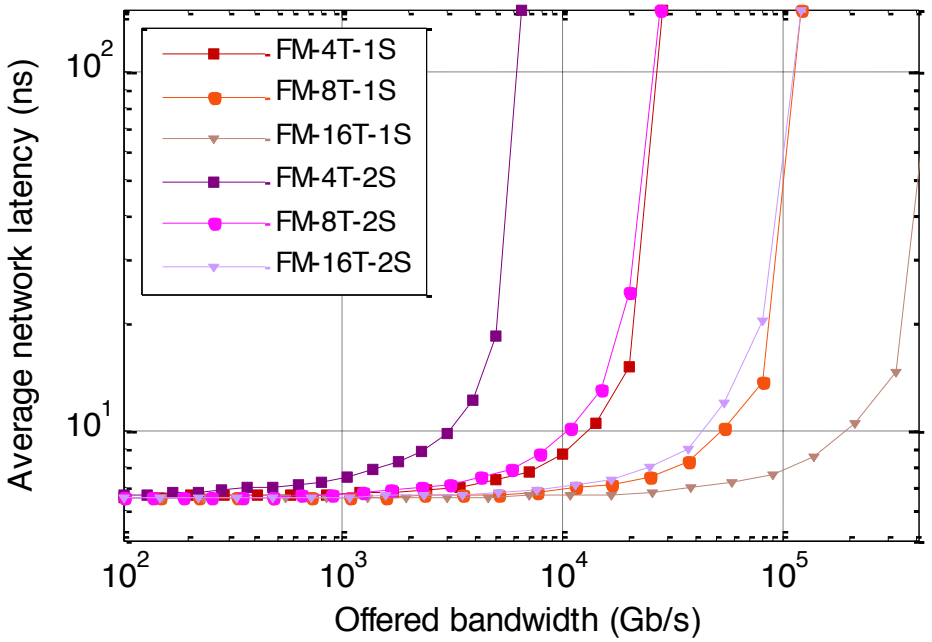

Impact of layout on energy per bit

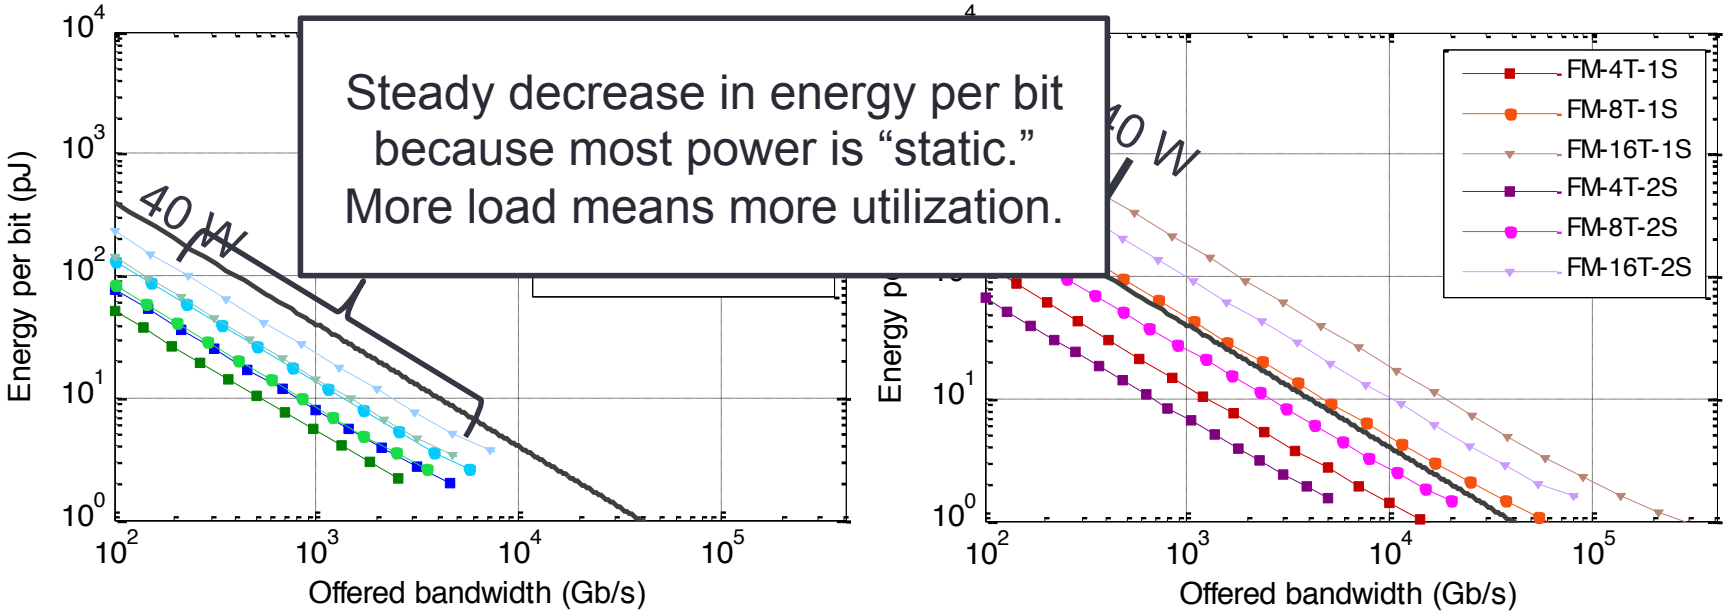

- The best configuration in terms of energy per bit depends on offered load
- However, these figures hide latency

Impact of layout on energy per bit

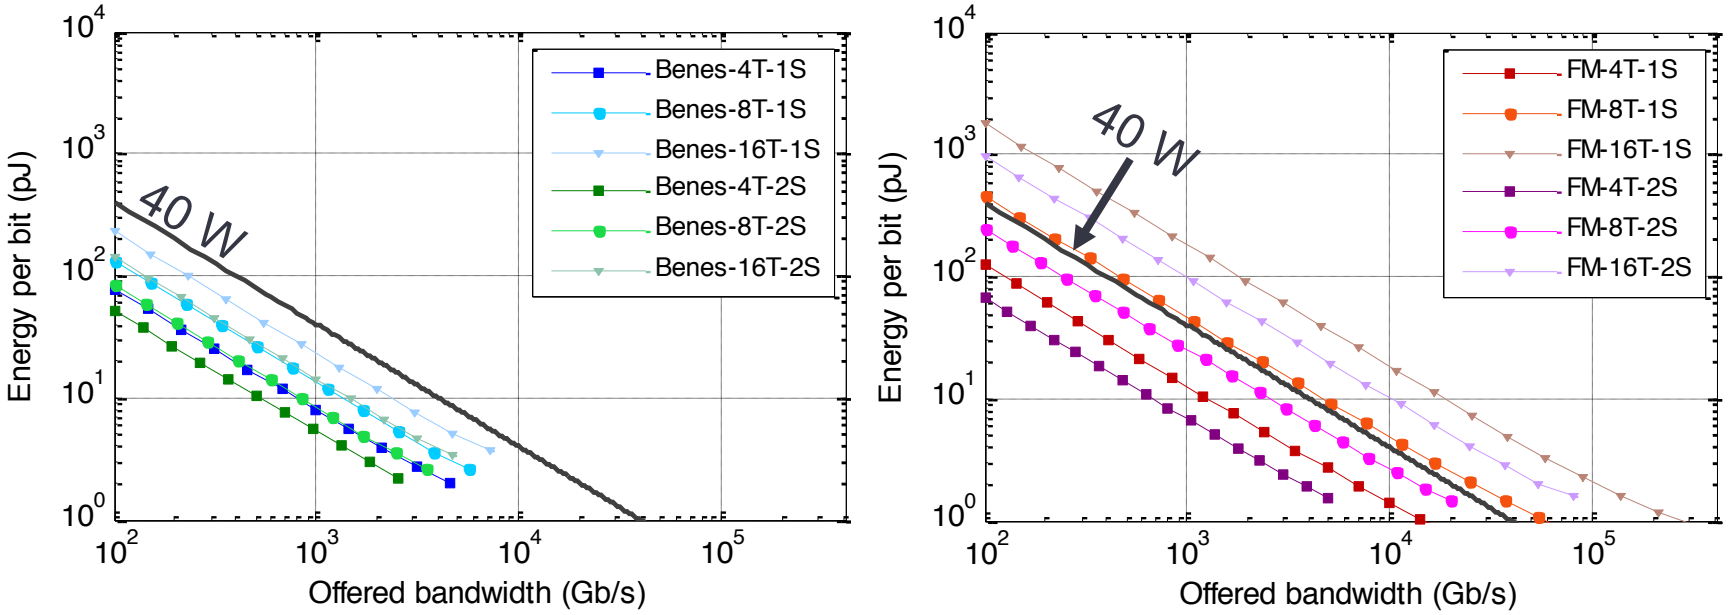

- The best configuration in terms of energy per bit depends on offered load
- However, these figures hide latency

Pareto optimality of topologies

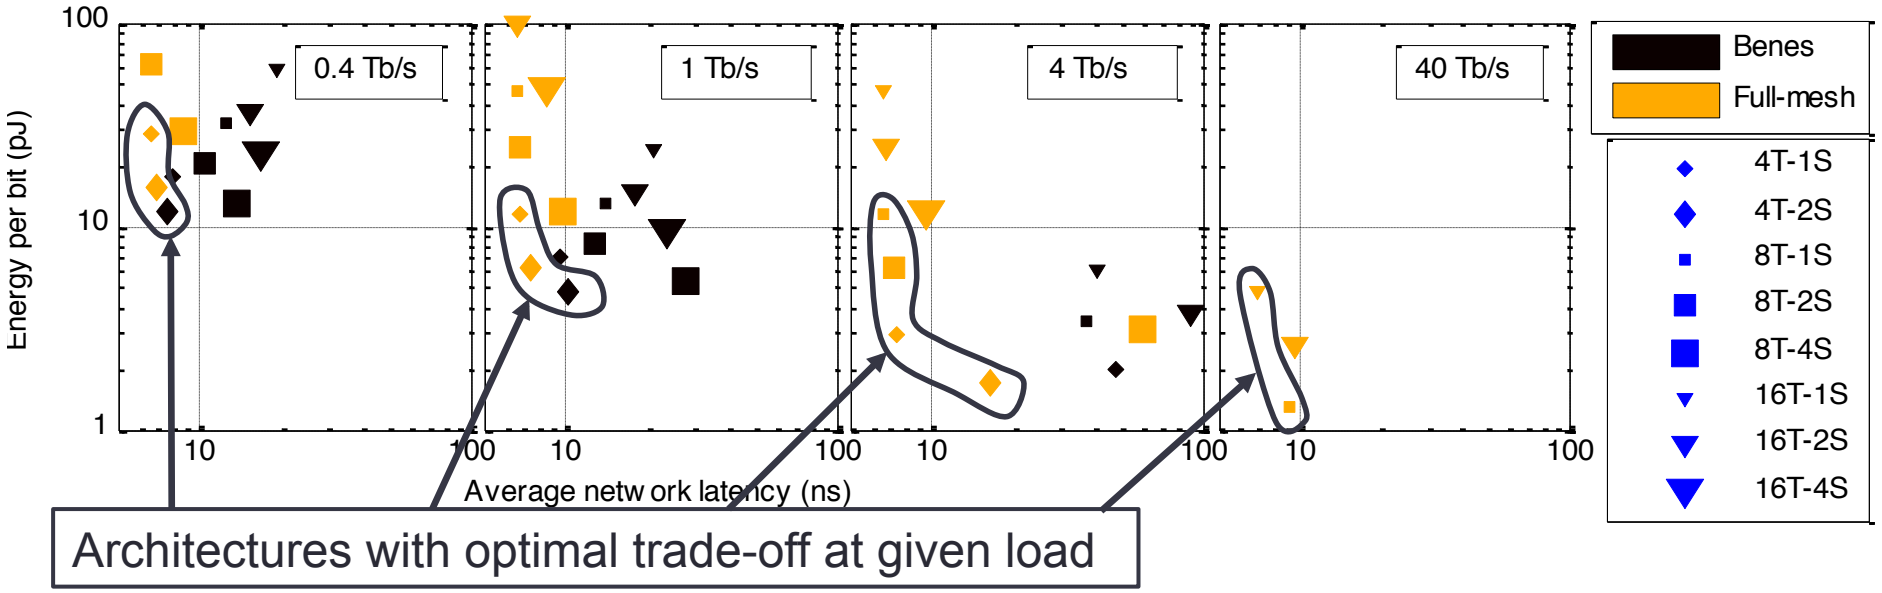

- Can move to a more power-consuming topology to improve latency, or vice versa

Pareto optimality of topologies

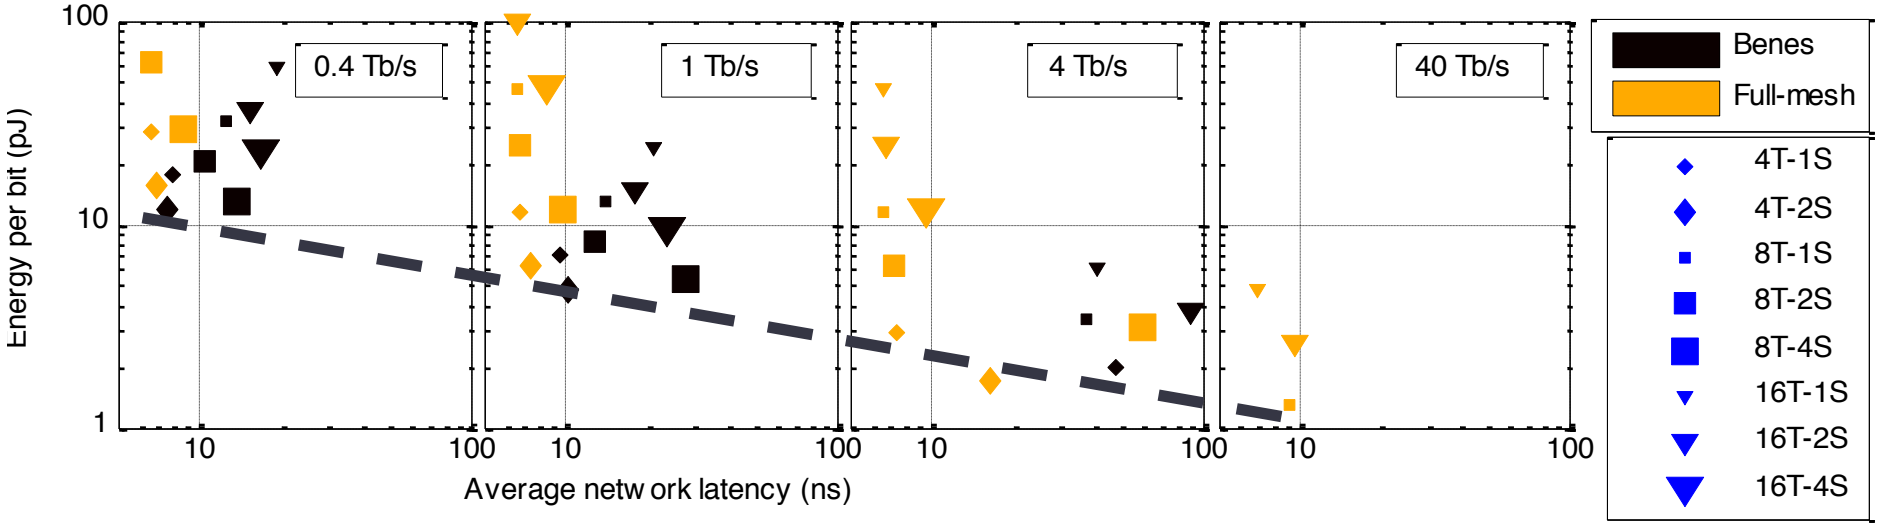

- Low loaded networks inevitably suffer from higher energy per bit

Conclusions

- Developed methodology for navigating design space
- Using cross-layer analysis, we characterized an upper bound on the energy efficiency of silicon photonic networks at the chip-to-chip scale
- Trend: *For (relatively small scale) silicon photonic networks,* the mechanisms that accommodate for low loads (i.e. resource sharing) degrade energy efficiency