

# Nanophotonic Interconnection Networks for Energy-Performance Optimized Computing



**Keren Bergman**

Lightwave Research Laboratory

Department of Electrical Engineering  
Columbia University




# Chip Multiprocessors




Growing parallelism multi-core chip architectures straining on- and off-chip electronic interconnects; need more bandwidth with less power dissipation




Sun Niagara  
8 cores  
2005




Sony/Toshiba/IBM Cell  
9 cores  
2006




Intel Polaris  
80 cores  
2007




[Intel]



IntellaSys SEAforth 40C18  
40 cores  
2008

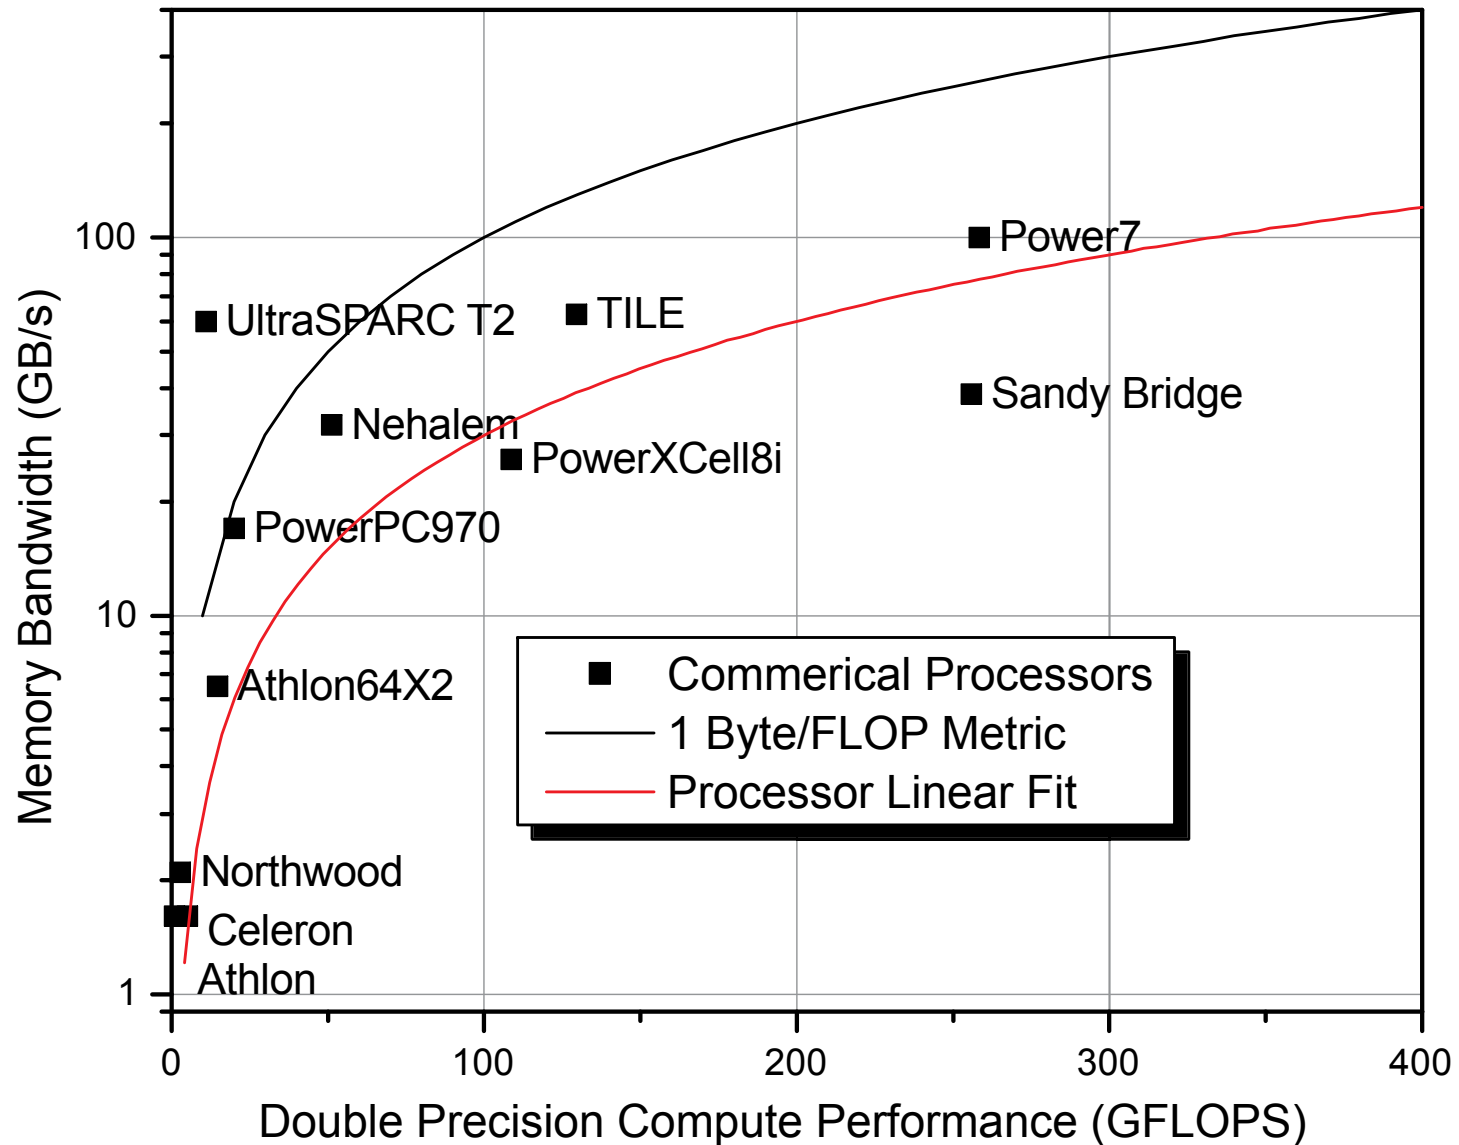


Tiler TILE-Gx100  
100 cores  
2009




NVIDIA Fermi  
512 cores  
2010

# Processor – Memory BW Performance



# Processor Pin Count Limitation




\* Target Metric of 1 B/FLOP

# Why Photonics is a good idea?

Photonics changes the rules for Bandwidth, Energy, and Distance.


## ELECTRONICS:

- Buffer, receive and re-transmit at every router.
- Each bus lane routed independently. ( $P \propto N_{\text{LANES}}$ )
- Off-chip BW is pin-limited and power hungry.




## OPTICS:

- Modulate/receive high bandwidth data stream once per communication event.
- Broadband switch routes entire multi-wavelength stream.
- Off-chip BW = On-chip BW for nearly same power.

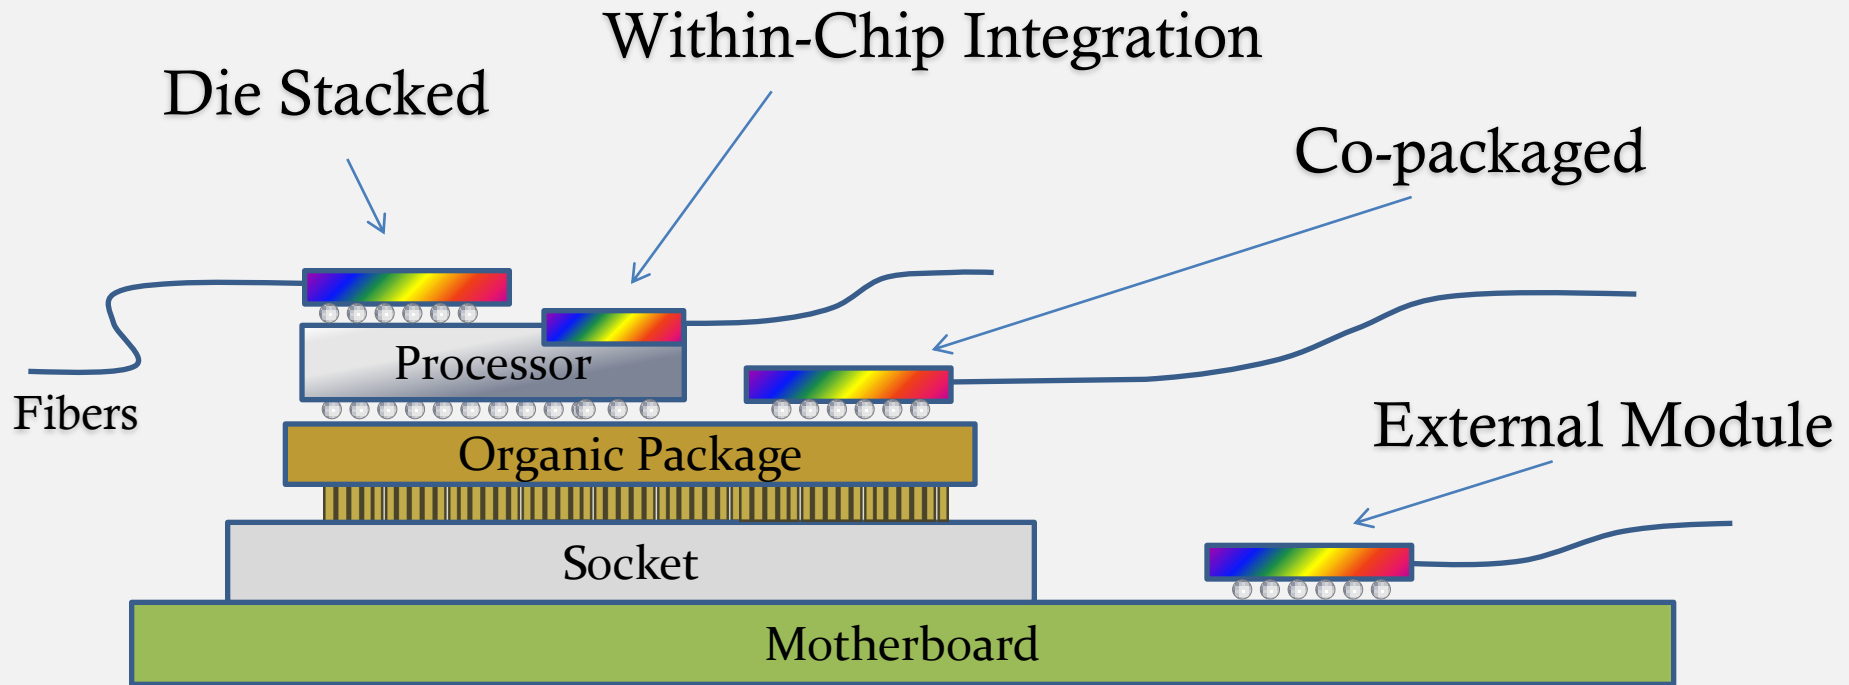


# Silicon Photonic Interconnects in Computing

- Silicon photonics: disruptive technology to address both on-chip and off-chip communication bottlenecks.
- High refraction index contrast materials allows small footprint and low power devices.
- Rules of scalability & energy efficiency differ substantially from electronic links.
- Bandwidth density:  $\sim 2 \text{ Tbps} / 20 \mu\text{m}$  pitch at chip's edge.
- Off-chip lasers – bulk of energy dissipation happens off the chip.



# Integration of Silicon Photonic I/O

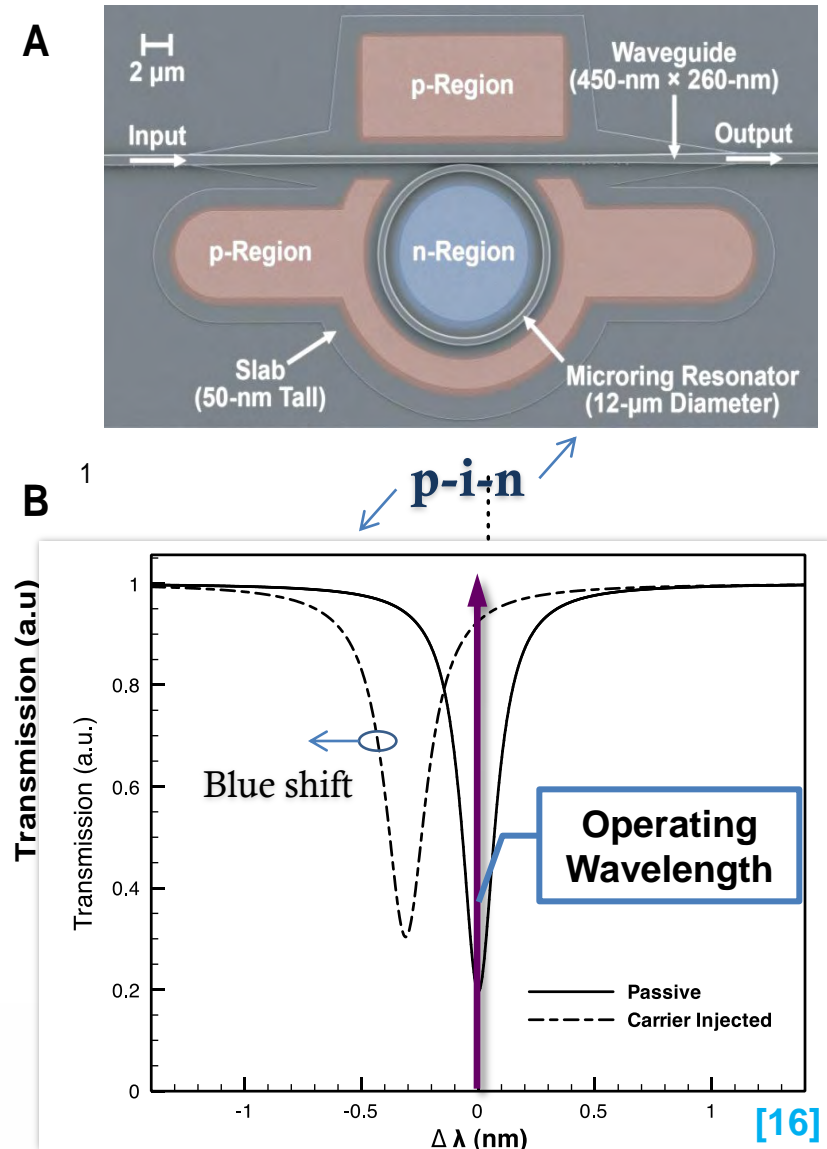


- Best performance gained by placing optic module as close as possible to electronics.



# Active / Tunable Microring Devices


- P/N-doping of silicon - diode injection (p-i-n) or depletion
- OOK modulator can be based on resonance shifts.
- Power dissipation  $\propto$  device fJ/bit.
- Integrated local heaters allow stabilization.
- Functionalities: modulators (Gbps), WDM mux / demux, filters.






# Microring-Based Comm. Links

Wavelength selectivity inherently supports WDM configuration with a single bus waveguide.




Analytical Model



# Microring WDM Link Considerations

WDM channel count limited by

- Avoid resonance overlap.  
FSR ( $\propto 1/R$ )  $\sim 50$  nm.
- Minimal channel spacing to avoid inter-modulation x-talk  
 $\sim 0.4 - 0.6$  nm




Single link bandwidth density:

- 12.5 Gbps modulation – 0.4 nm (50 GHz) channel spacing,  $Q \sim 2 \cdot 10^4$ . Allows 125 channels → **1.56 Tbps**
- 25 Gbps modulation – 0.6 nm (75 GHz),  $Q \sim 10^4$ . Allows 83 channels → **2.07 Tbps**

# Dense WDM Microring Link Design

Design driven by “best possible” single-waveguide optical link in terms of BW density and energy efficiency



## Tx Array:

- Si or SiN bus WG
- Inverse-taper edge couplers
- Depletion-mode microring modulators

## SM fiber:

- PM
- Negligible loss up to 1 km

## Rx Array:

- Thermally tuned microring filters
- Ge PD on drop ports

# Power Efficiency – 2016 projection

@ 12.5 Gbps, 125 channels (1.56 Tbps) - power dissipation:

Modulator – 0.16 pJ/bit

Demultiplexer (thermal) – 0.22 pJ/bit

Mod Driver – 0.1 pJ/bit

TIA & SA – 0.4 pJ/bit

**~0.88 pJ/bit** on optical die

Laser (~1.1 mW optical) ~ @1% efficiency: **9 pJ/bit**

@ 10% efficiency: **0.9 pJ/bit**

@ 20% efficiency: **0.45 pJ/bit**

BW density:  
**~ 1 Tbps / mm<sup>2</sup>**

**~ 1 pJ/bit** to drive signals up to and from the optical die (electronic subtask)

@ 25 Gbps, 83 channels (2.07 Tbps) - power dissipation:

Modulator – 0.09 pJ/bit

Demultiplexer (thermal) – 0.11 pJ/bit

Mod Driver – 0.3 pJ/bit

TIA & LA – 1 pJ/bit

**~1.5 pJ/bit** on optical die

Laser (~2.2 mW optical) ~ @1% efficiency: **9.6 pJ/bit**

@ 10% efficiency: **0.96 pJ/bit**


@ 20% efficiency: **0.48 pJ/bit**

BW density:  
**~ 2 Tbps / mm<sup>2</sup>**

**~ 2 pJ/bit** to drive signals up to and from the optical die (electronic subtask)


# Chip-Scale Photonic Network Architectures

Corona




[U. of W

Processor-to-DRAM  
Network



Circuit-Switching Torus  
Network




mbia]

## Architecture-focused Design

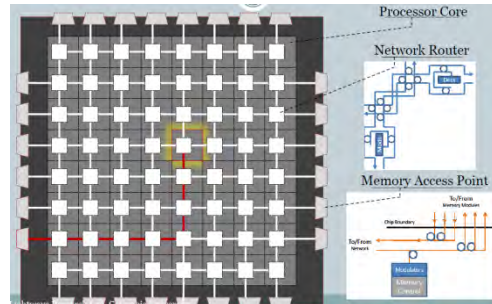
- Large number of ring resonators (over 1M)
- Missing device details for ring dimensions
- High insertion losses, due to many waveguide crossings

# Spectrum of Feasibility






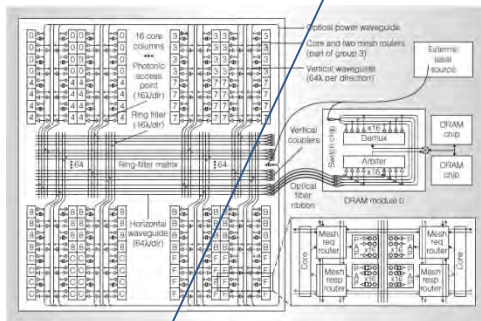
Corona [HP]



PhotonicMesh [Columbia]

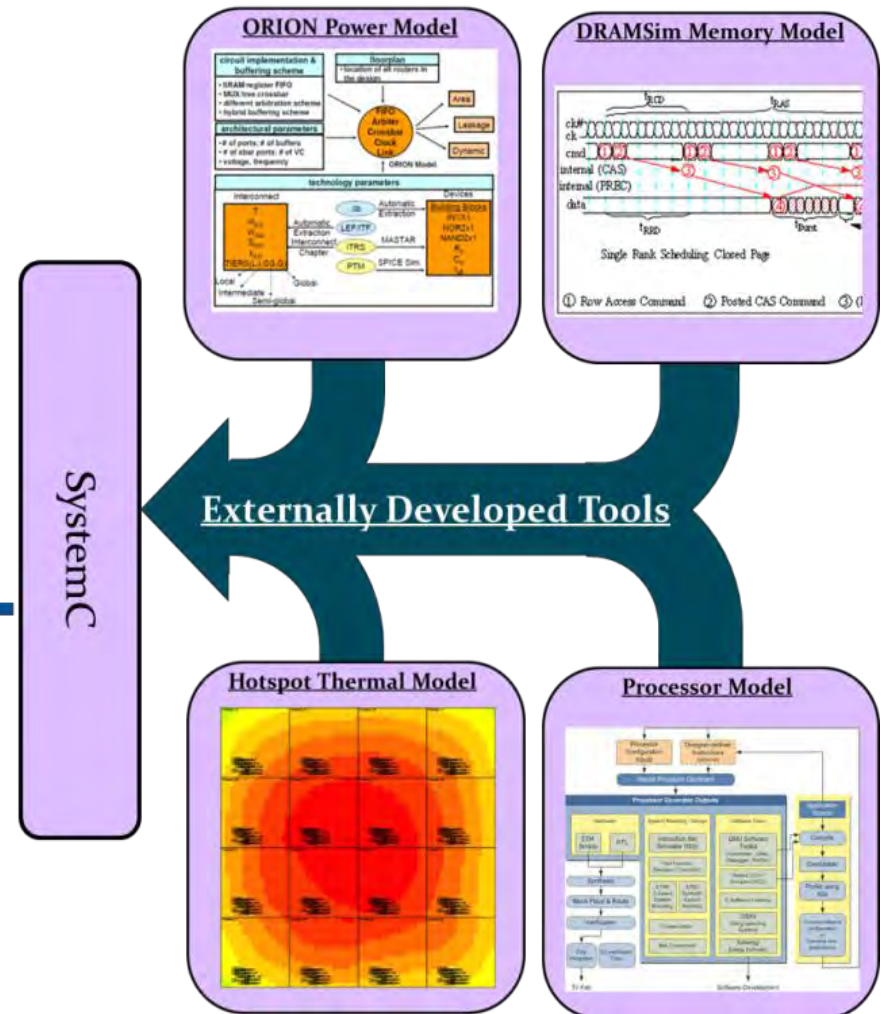
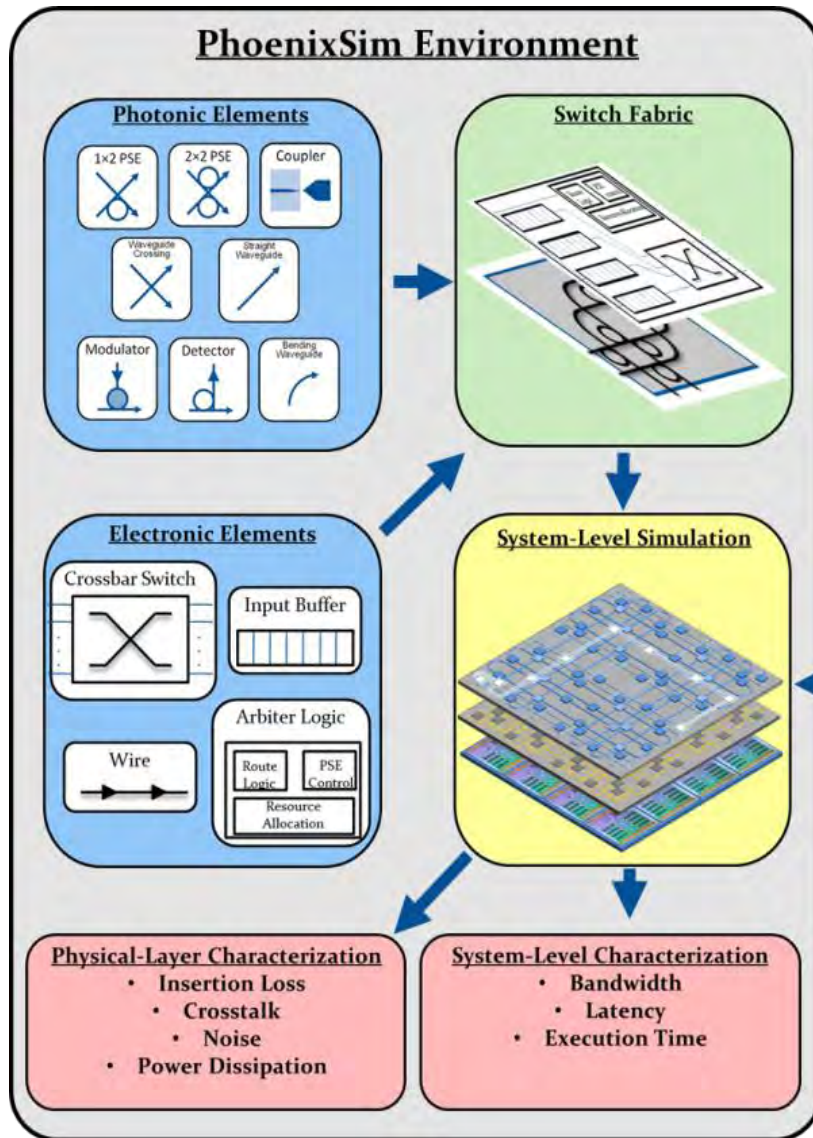


[MIT]






# Simulation Methodology and PhoenixSim Environment



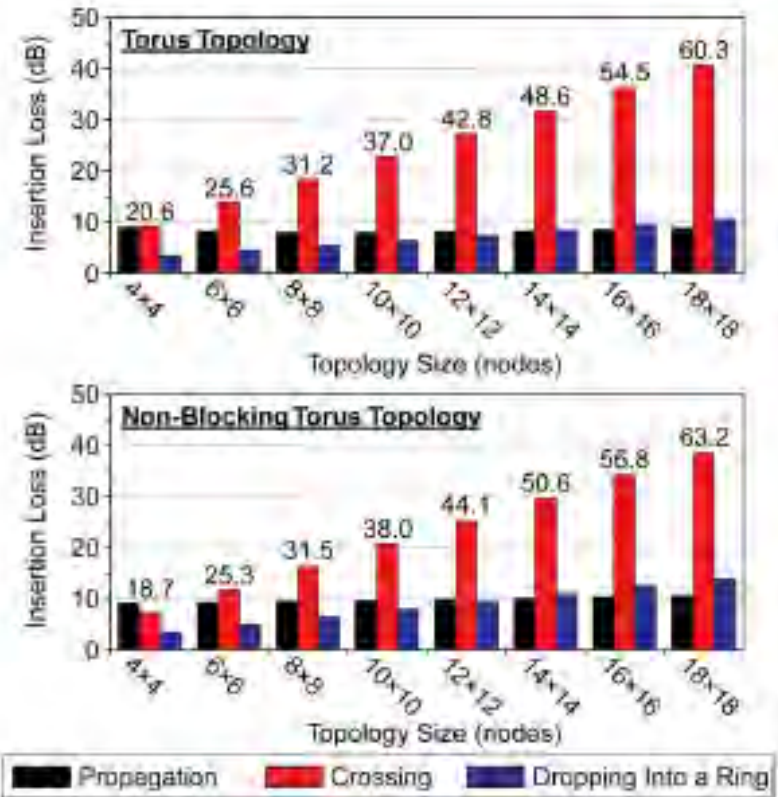


# Multilayer Deposited Architectures

- Mesh single-layer NoC designs' scalability limited by **crossing losses**



[J. Chan *et al.*, *JLT*  
28 (9) (2010)]



- Deposited photonics allows creation of multilayer designs
  - SiN for low loss bus waveguides
  - Poly-Si for E-O active components (modulators, switches)


# Summary Table of Silicon Photonic Materials

Material	Propagation Loss	Electrically-Active	Deposit Capability	Stable	CMOS Compatibility
c-Si	1.7 dB/cm	Yes	No	Yes	Yes
Poly-Si	6.45 dB/cm	Yes	Yes	Yes	Yes
Si <sub>3</sub> N <sub>4</sub>	0.1 dB/cm	No	Yes	Yes	Yes
a-Si:H	2 dB/cm	No	Yes	No	Yes


- Crystalline silicon (c-Si)
- Polycrystalline silicon (poly-Si)
  - Silicon nitride (Si<sub>3</sub>N<sub>4</sub>)
- (Hydrogenated) Amorphous Silicon (a-Si:H)

# Crystalline SOI Silicon Photonics


- Bulk of silicon photonics research based on etching SOI to form optical structures.



Cornell/Columbia, 2010




Sandia, 2008



IBM, 2010




- High optical-quality crystalline-Si. Thick BOX SOI provides optical isolation from substrate. Single-plane architectures.
- In-package integration
  - Die stacking / bonding – fine pitch electrical connectors allow high bandwidth but fabrication complexity & cost significant.
  - Monolithic - transistor real-estate, incompatibly-thick BOX (or thin BOX with undercut), significant manufacturing process modification.

## Cross-Sectional View of 3D Silicon Photonic Stack







# Multilayer Integration for Photonic NoC

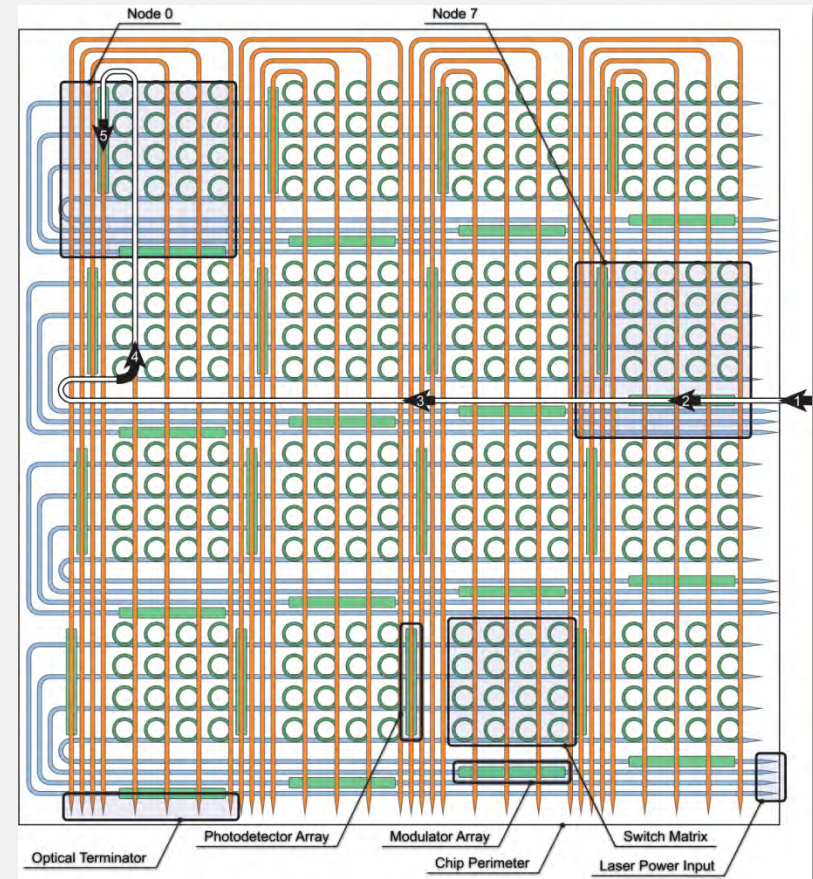


# Scalability of Multilayer Design



[A. Biberman *et al.*, *OFC 2011 OMM2*]

## Photonic Matrix Crossbar




[A. Biberman *et al.*, *JETC 7 (2) (2011)*]


- Multilayer design eliminates waveguide crossings to allow improved scalability



# Optically connected memory with deposited photonics



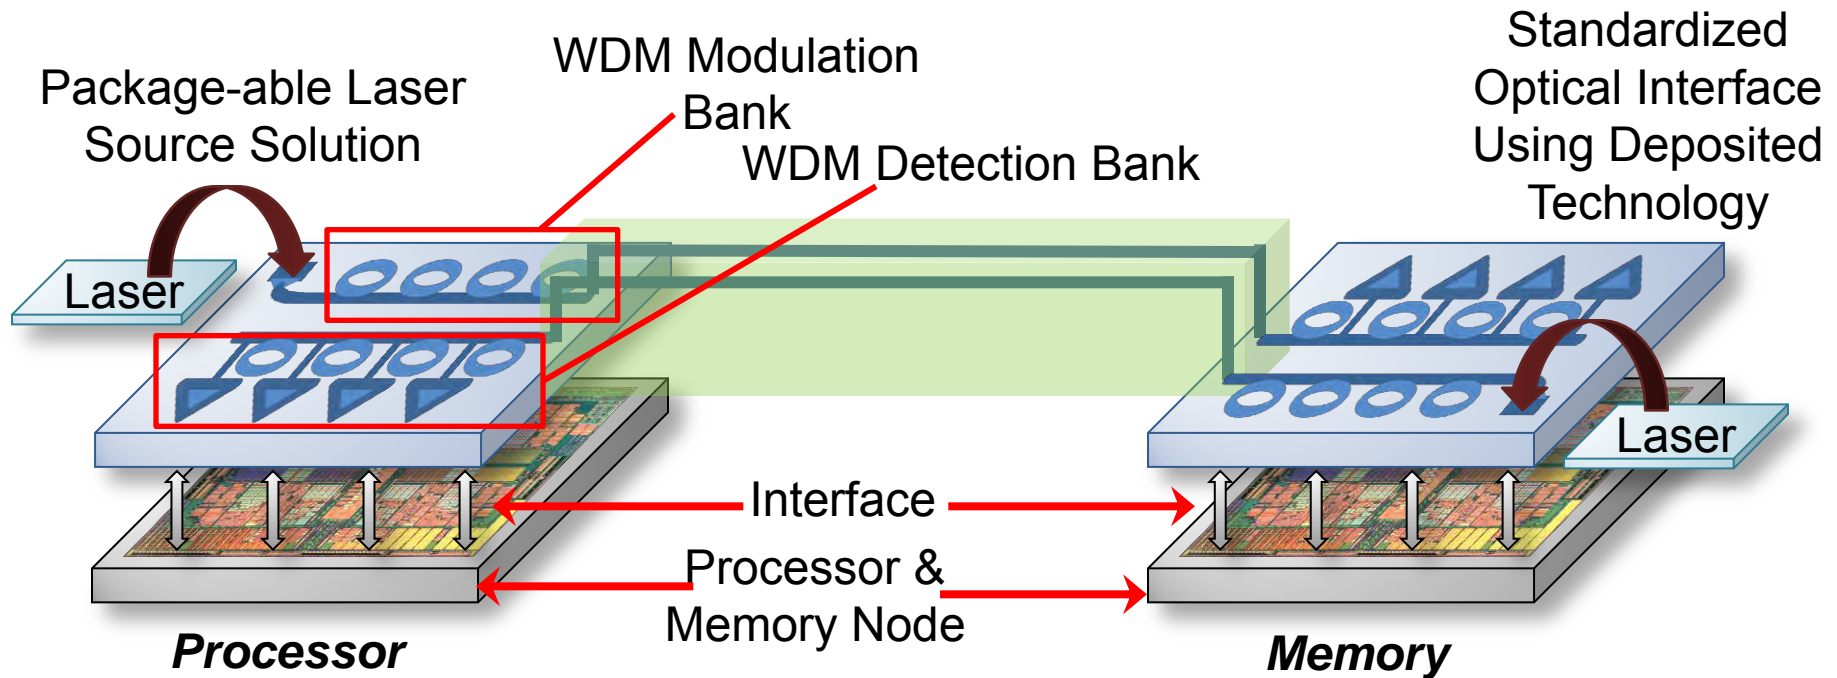
CMOS Back-End Photonic Deposition



Lipson, Cornell




# Optically interfacing processor with memory

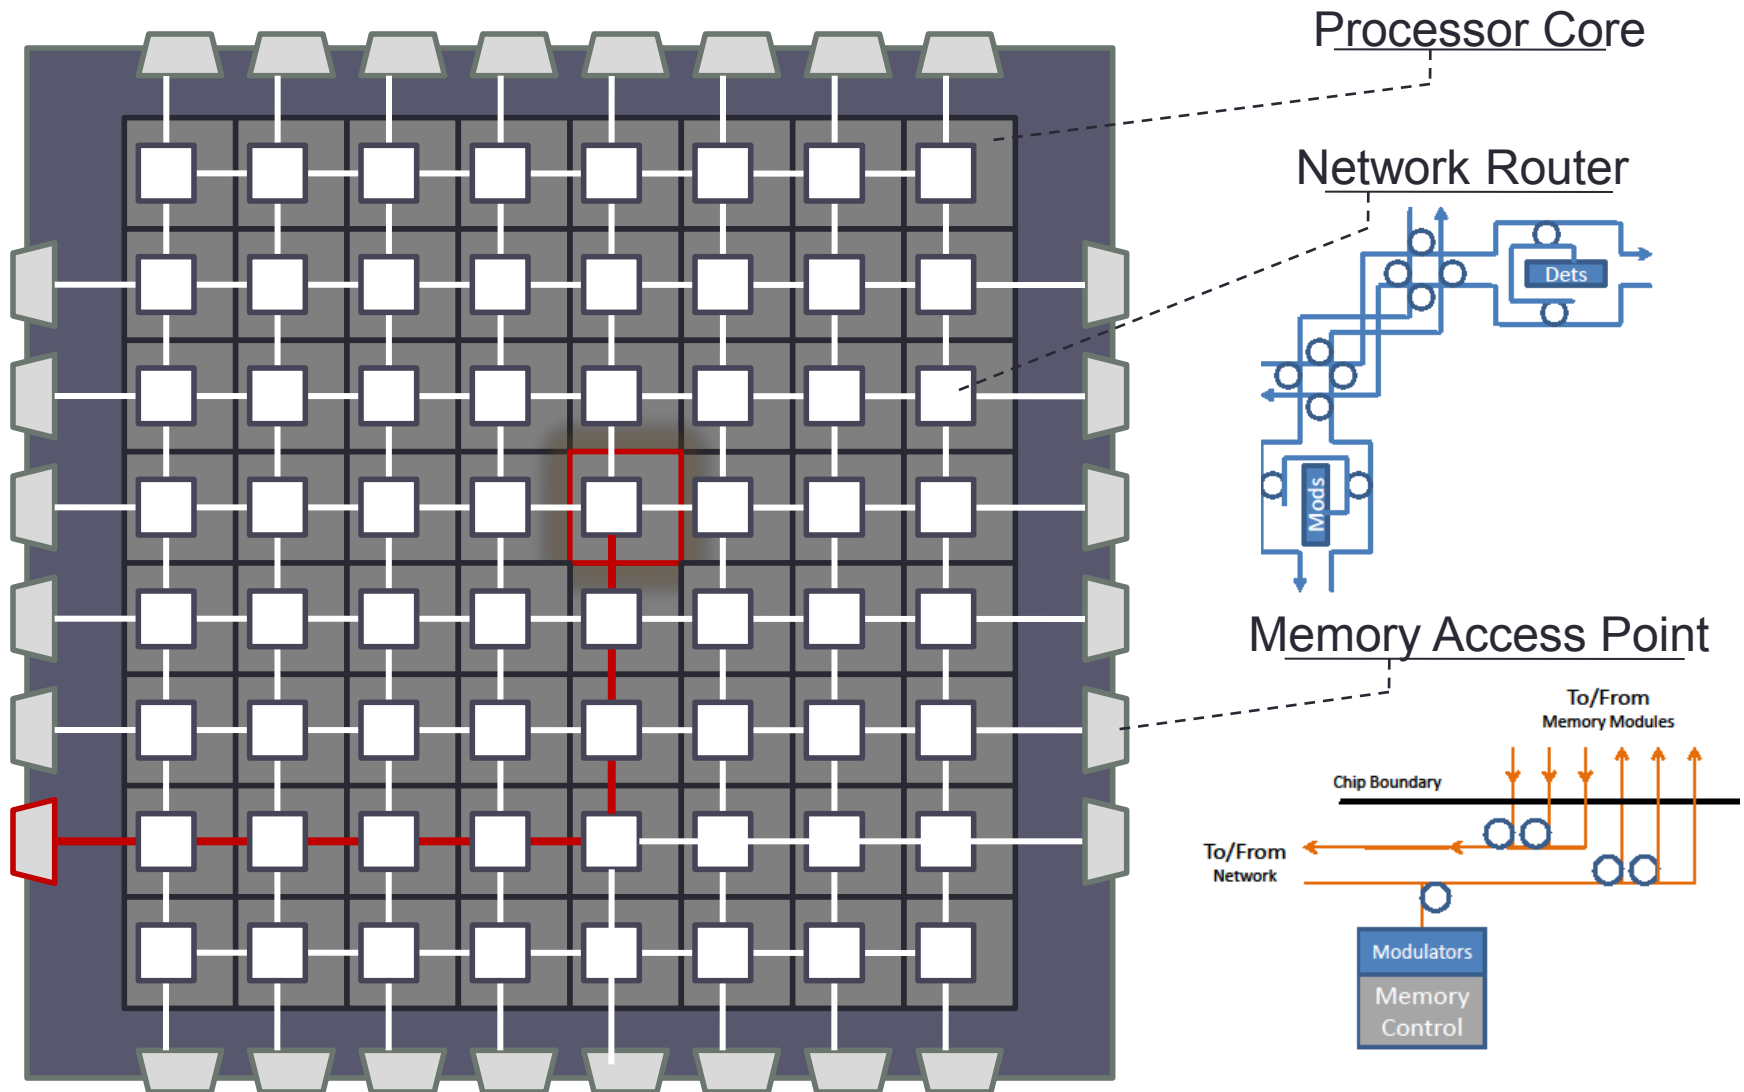


- Processor-to-memory communication using optical I/O to go off-chip
  - Photonic integration on CMOS
- WDM modulators and detectors to exploit wavelength parallelism
  - Enhanced performance boost (Gb/s, pJ/bit) with SerDes integrated into HMC

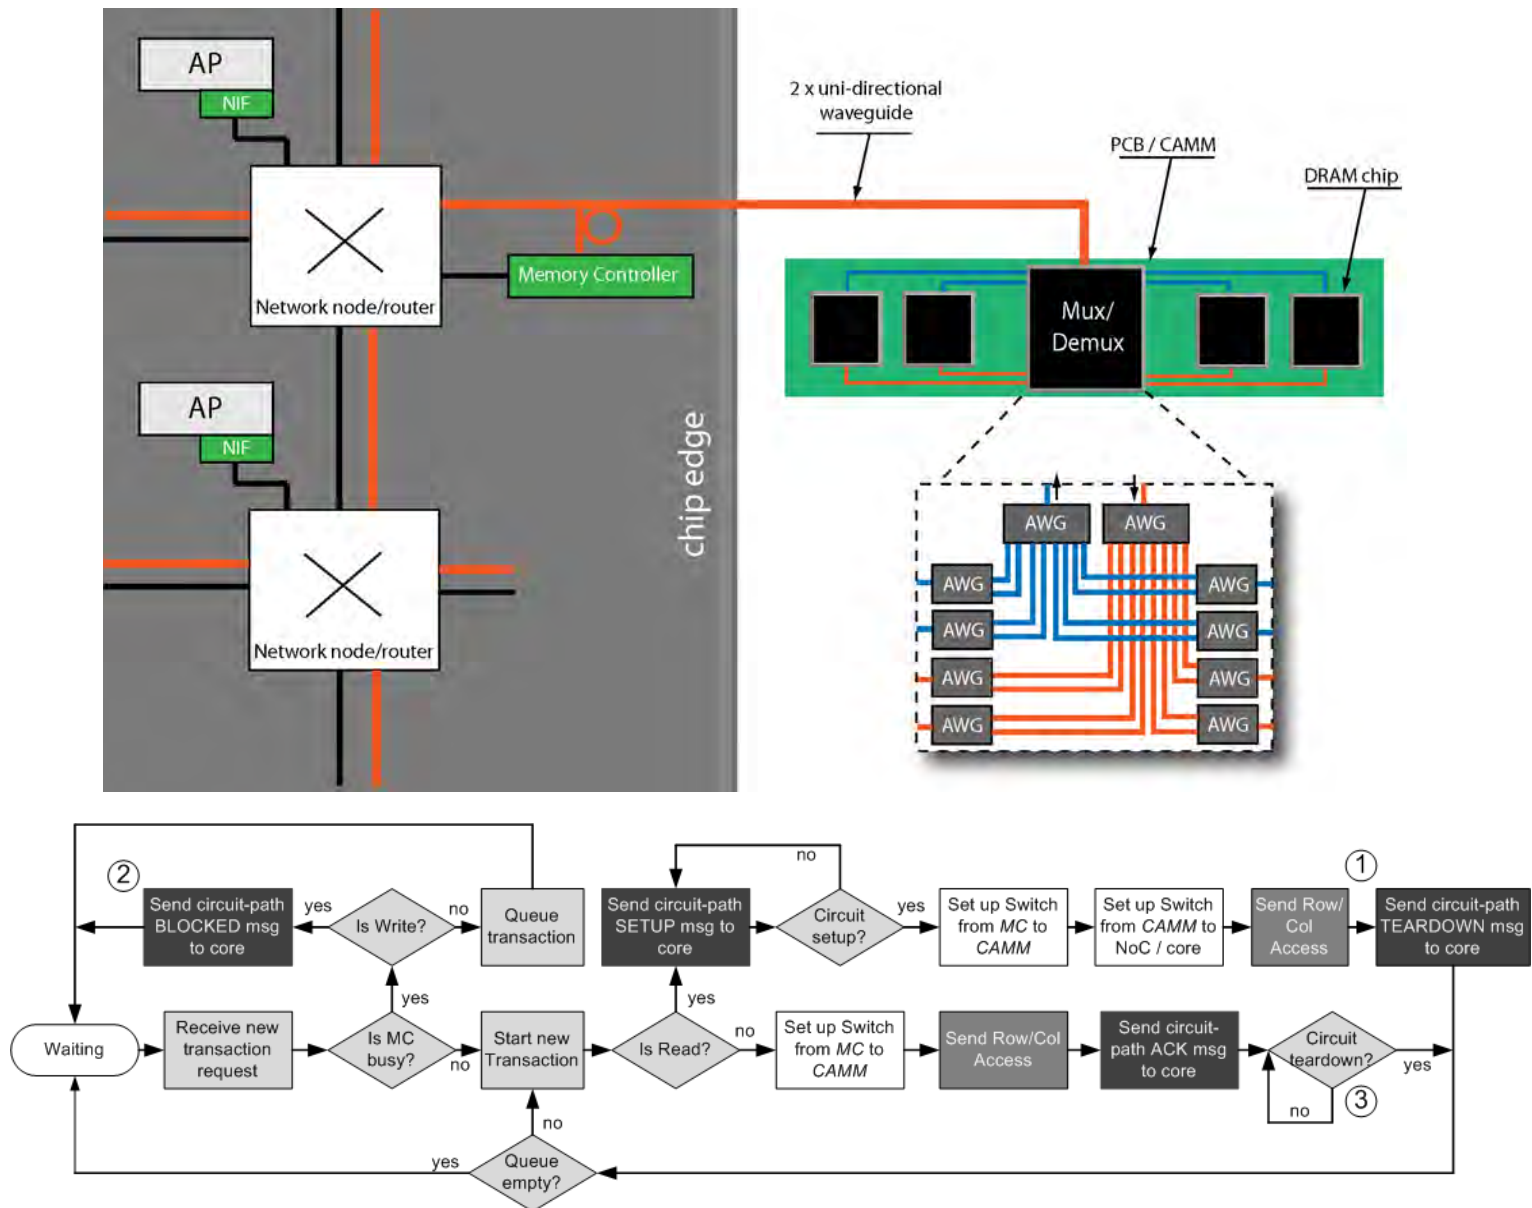
# Photonic Network Design Exploration



# Photonic Circuit-Switched DRAM Access




# Circuit-Switched Memory Access



# Circuit-Switched Memory Access

- Image processing application kernels



Network	Projective Transform			Matrix Multiply			FFT		
	Power (Watts)	Perf. (GOPS)	Impr. (GOPS/W)	Net. Pow. (Watts)	Perf. (GOPS)	Impr. (GOPS/W)	Power (Watts)	Perf. (GOPS)	Impr. (GOPS/W)
Emesh	11.2	1.04	1x	11.1	0.78	1x	11.4	1.75	1x
EmeshCS	19.0	47.3	26.9x	15.8	31.82	29.01x	11.2	4.74	2.82x
PS-1	4.37	27.80	68.6x	4.35	26.51	87.64x	4.28	4.32	6.72x
PS-2	2.21	17.76	86.7x	2.17	13.48	89.33x	2.15	3.12	9.67x

# Summary



- Silicon photonics shaping to be the prime candidate to address chip I/O limitations: system wide high-bandwidth density with extreme energy efficiency.
- Optically connected memory - very compelling application in near future
- Deposition-based silicon-photonics fabrication: “leap frog” solution for close integration with electronics.
- Not just a “wire” replacement...potential for revolutionizing across compute and memory architectures.