Photonics Systems: Physical Layer to Scaled Architectures

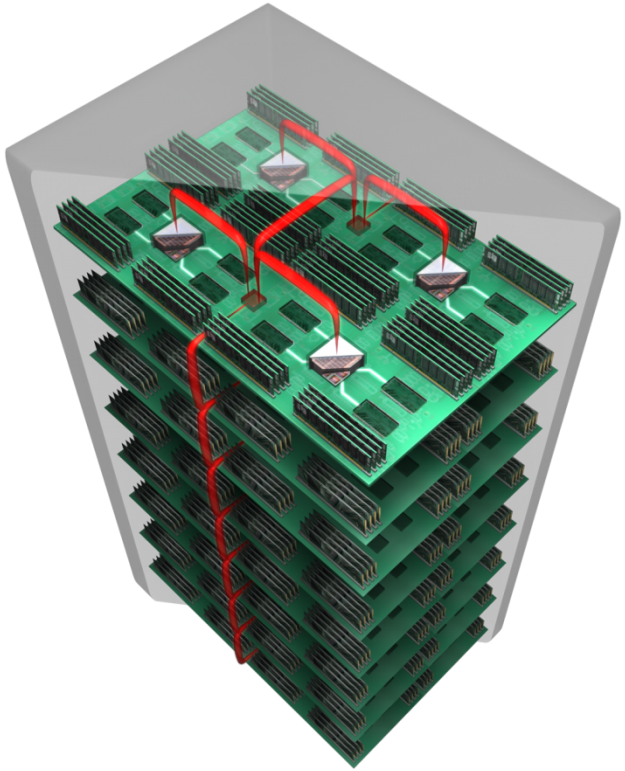

Keren Bergman
Department of Electrical Engineering
Columbia University

COLUMBIA UNIVERSITY

Computation to Communications Bound

Computing platforms with increased **parallelism** at all scales:

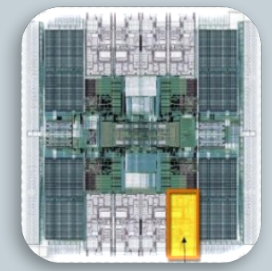

Sun Niagara
8 cores
2005

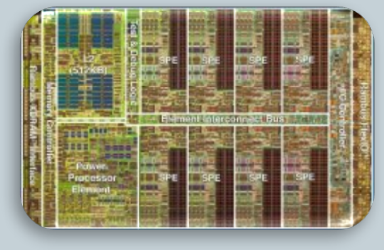

Sony/Toshiba/IBM Cell
9 cores
2006

Intel Polaris
80 cores
2007

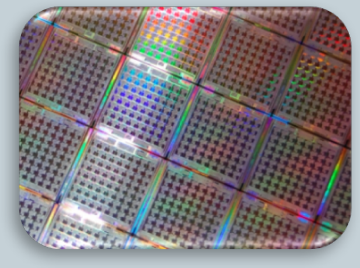

Tilera TILE-Gx100
100 cores
2009

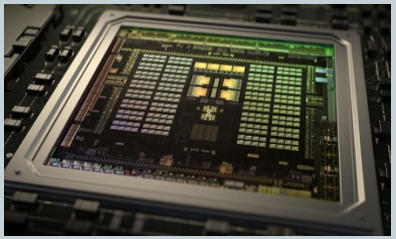

NVIDIA TEGRA X1
256 GPUs
2016

Handheld
System-on-Chip

Embedded Systems

Data Centers

Interconnect trends

- Top 10 average node level evolutions:
  - Average node compute power:
    - 31GFlops → 600GFlops
    - ~19x increase
  - Average bandwidth available per node
    - 2.7GB/s → 7.8GB/s
    - ~3.2x increase
  - Average byte-per-flop ratio
    - 0.06 B/Flop → 0.01 B/Flop
    - ~6x **decrease**
    - **Sunway TaihuLight (#1) shows 0.004 B/Flop !!**

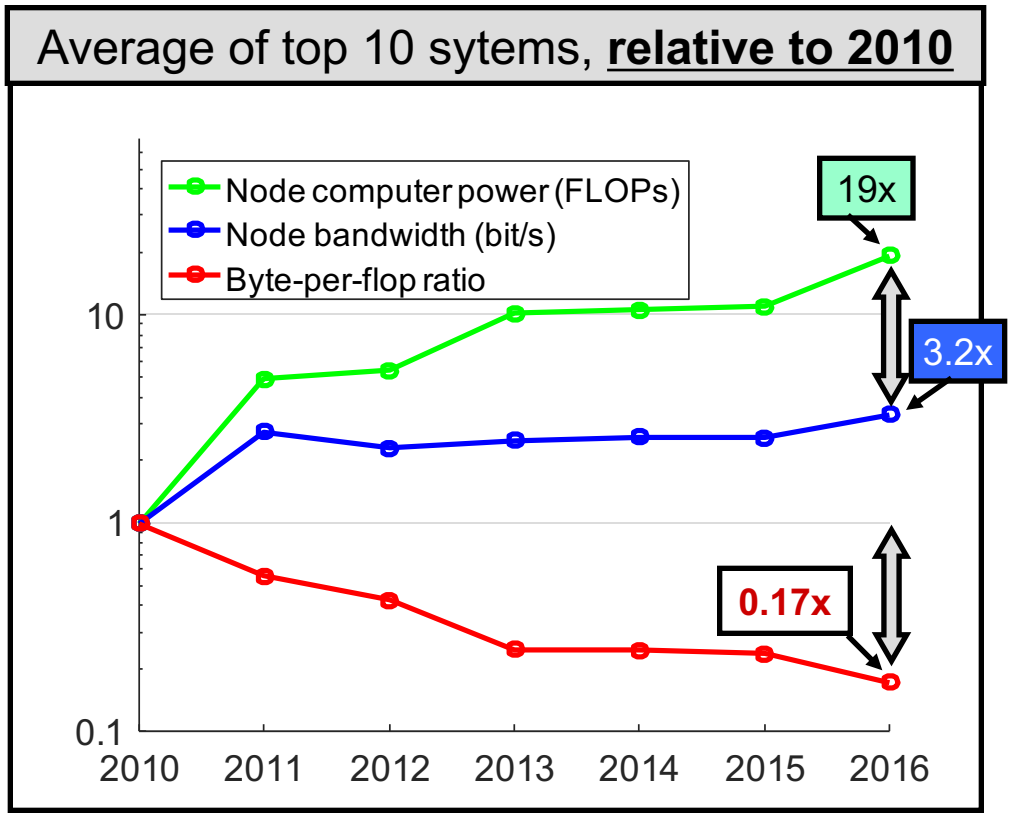

[top500.org, S. Rumley, et al. Optical Interconnects for Extreme Scale Computing Systems, Elsevier PARCO 64, 2017]

→ **Growing gap in interconnect bandwidth**

Mid 2000s: Silicon Photonics

Silicon-on-insulator (SOI) platform photonic building blocks:

High index contrast enables high confinement, low-loss propagation, virtually lossless bending

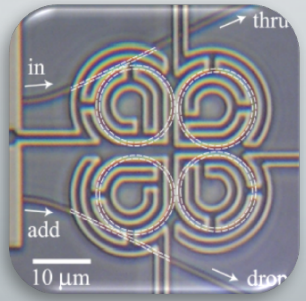

MIT

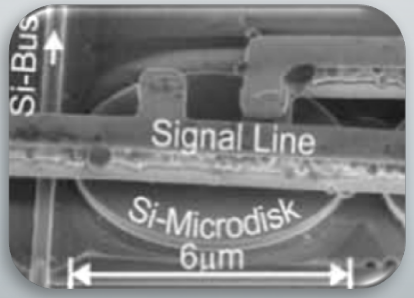

Sandia/MIT

Ghent

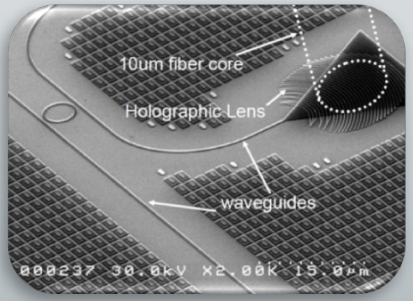

Luxtera

IBM

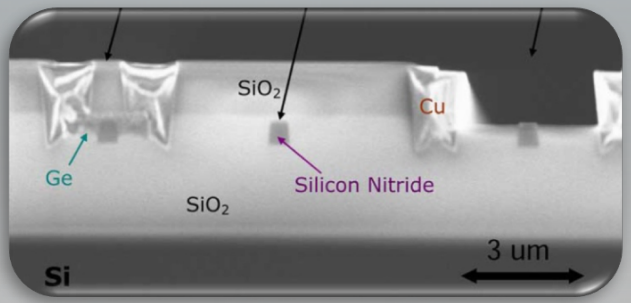

Intel

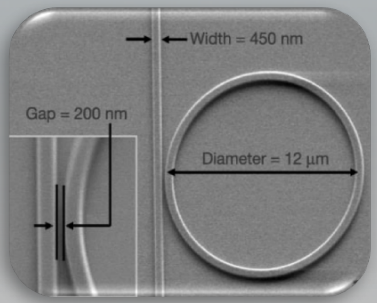

Cornell

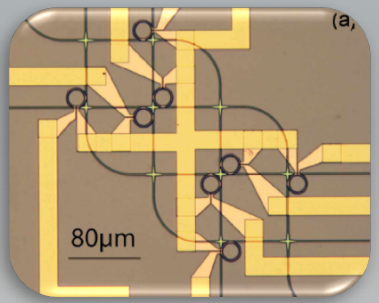

Cornell/Columbia

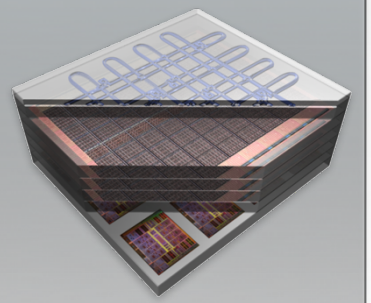

Columbia

Silicon Photonics: all the parts

- Silicon as core material
  - High refractive index; high contrast; sub micron cross-section, small bend radius.
- Small footprint devices
  - 10 µm – 1 mm scale compared to cm-level scale for telecom
- Low power consumption
  - Can reach <1 pJ/bit per link
- Aggressive WDM platform
  - Bandwidth densities 1-2Tb/s pin IO

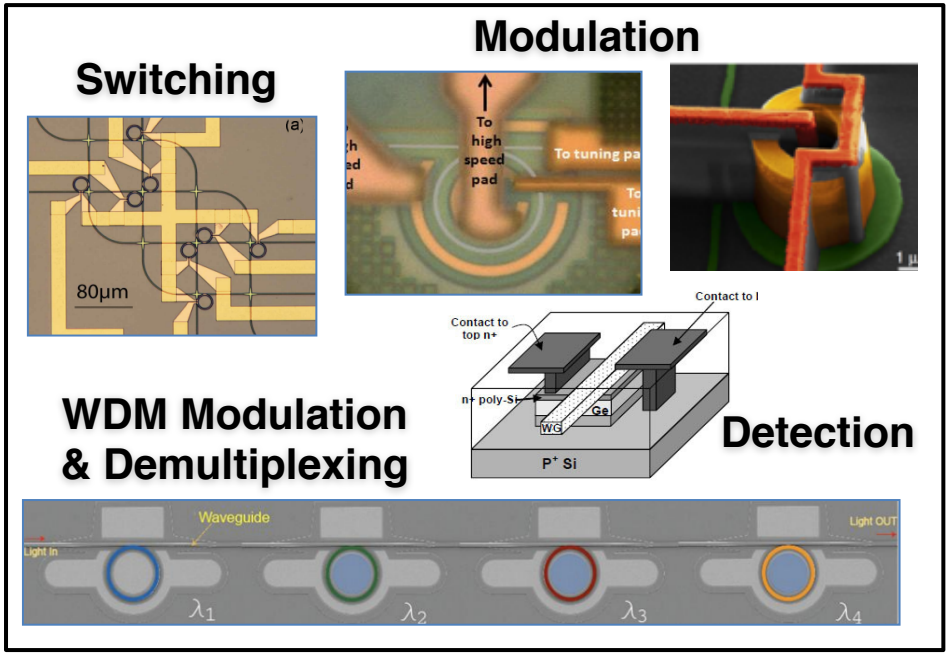

- Silicon wafer-scale CMOS
  - Integration, density scaling
  - CMOS fabrication tools
  - 2.5D and 3D platforms

S. Rumley et al. "Silicon Photonics for Exascale Systems", IEEE JLT 33 (4), 2015.

Silicon Photonics for Computing

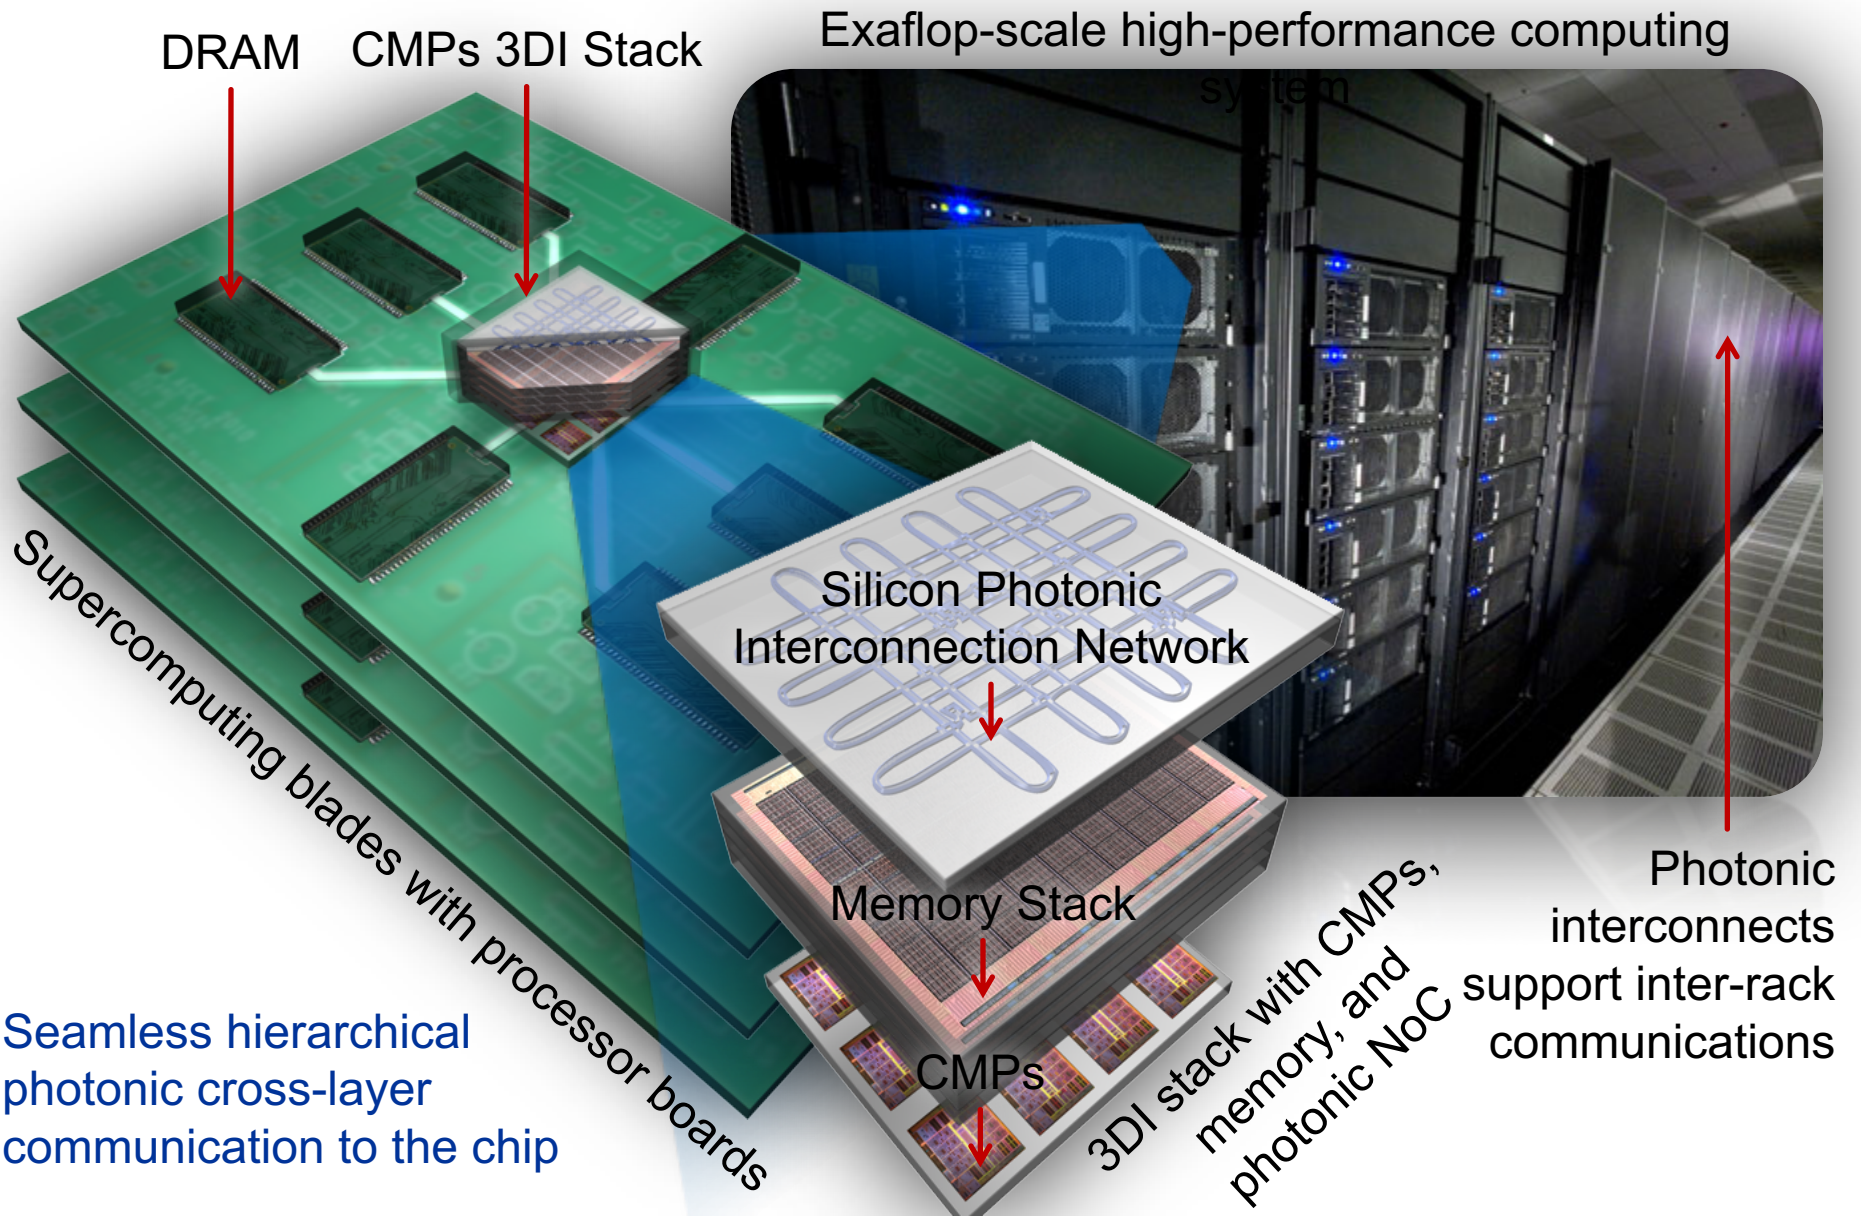

Seamless hierarchical photonic cross-layer communication to the chip

The Integrated Photonics Manufacturing Institute's Core Hubs - Albany

- 1.3M ft² facility
- cutting edge 300/450mm toolset
- 135k ft² of class 1 capable cleanroom
- processing capability span 65nm - 7nm

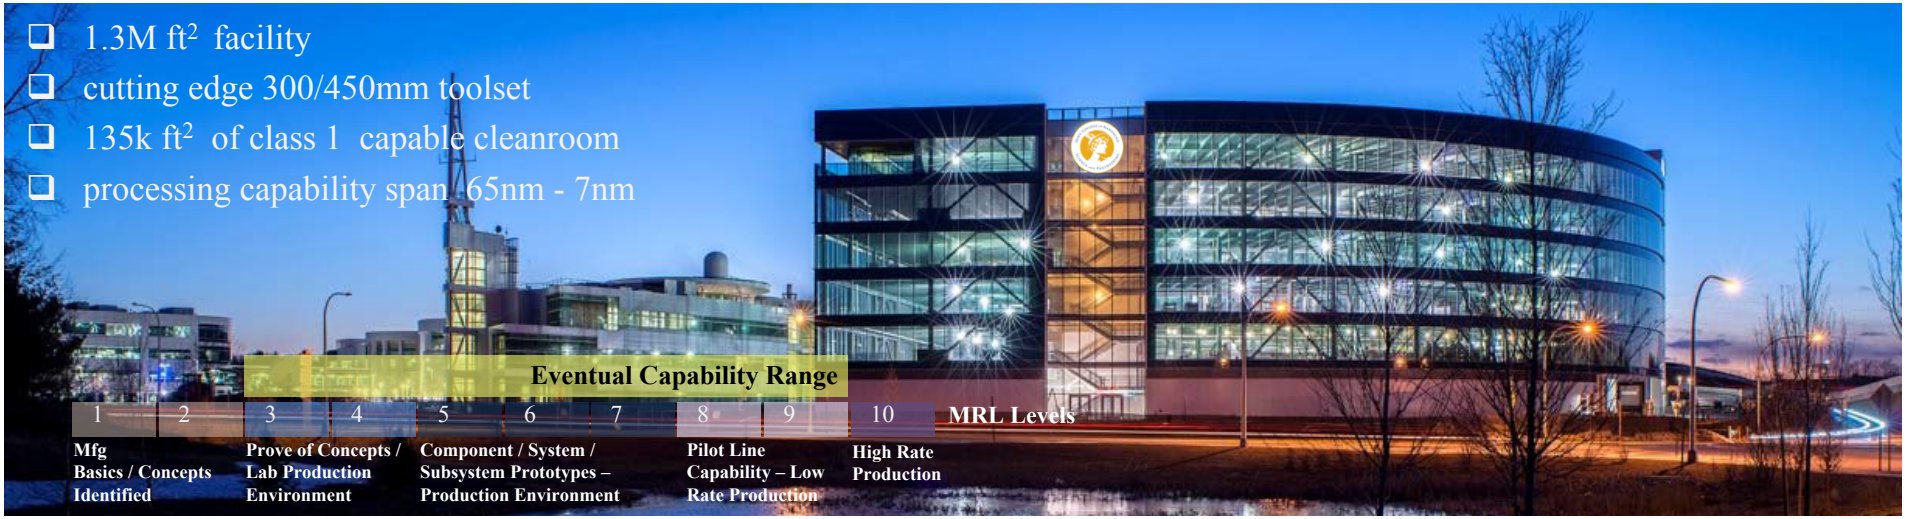

- years of proven silicon photonics results – multiple government & industry projects
- 300mm tools provide unprecedented quality photonics
- unmatched 3D stacking w/CMOS
- partnerships drive continuous revitalization investments

SUNY POLYTECHNIC INSTITUTE

300mm Si Photonics Wafer

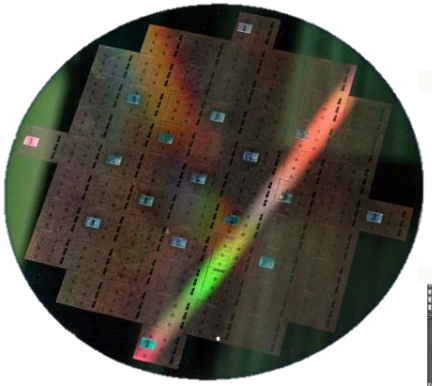

Continuously Tunable
Optical Orbital Angular Momentum Generator

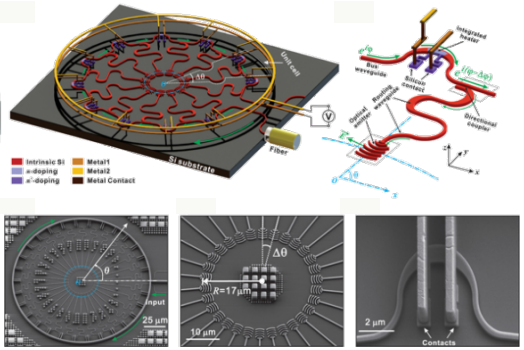

Erbium Laser

Post S-MLD

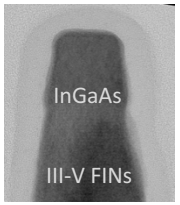

IV-IV 65nm Si CMOS

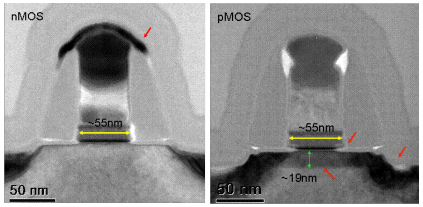

Undamaged
III-V FIN

Prior | Post CMP

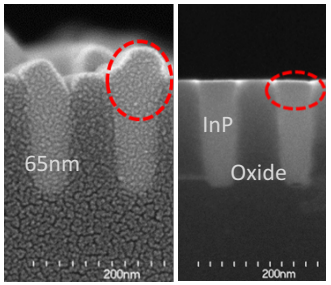

95nm gap in Si

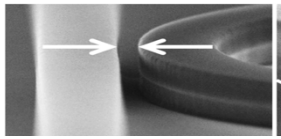

95nm Si_3N_4 Taper on
Si Waveguide

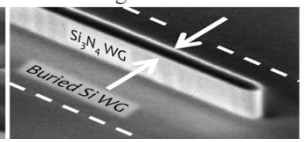

Market needs Common Manufacturing Technology Platforms to unfold its Potentials / Opportunities …

market vehicles are the means to drive the manufacturing technology platform

- industry sets the pace:
  exploring technology potentials,
  revolutionizing the data and sensing
  market
  - ➢ market is materializing
  - ➢ stunning growth rates projected
  - ➢ it's the technology of the future
- scientific/defense leverages
  industry pace:
  building on solid ground, adding
  uniquely required functionalities

	DataCom / Telecom	Analog RF Applications	PIC Sensors	PIC Array Technologies
Industry	 Internet / Datacenters	 CATV / Cellular	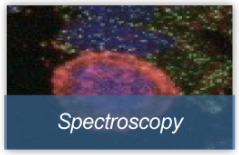 Spectroscopy	 LIDAR / 3D Displays
Increase	Capacity / Flexibility	Quality / Reliability	Sensitivity	Safety
Enable	Internet Scaling	Subcarrier Applications	Environmental / Proactive Healthcare	3D Augmented Reality
Decrease	Power / Cost	Cost	Healthcare Cost	Cost
Market	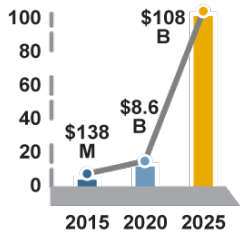	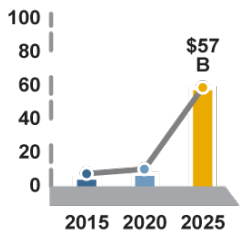	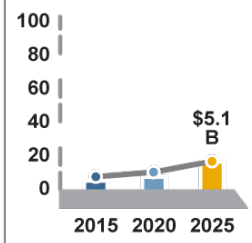	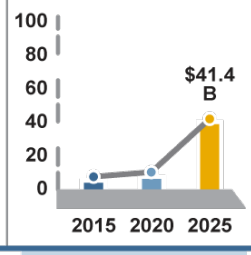	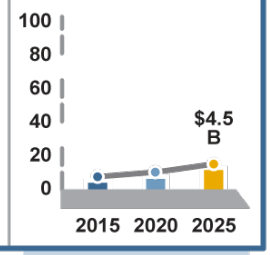
Scientific / Aero / Defense	 Enterprise Networking / Supercomputing	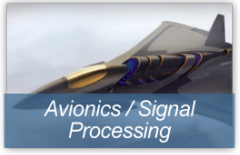 Avionics / Signal Processing	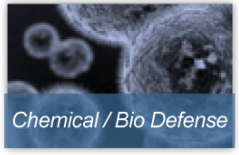 Chemical / Bio Defense	 Safety / Performance
Increase	Capacity / FLOPs	Bandwidth & Reliability	Safety & Health	Detection Accuracy
Enable	Reallocatable Bandwidth in Theater / Multiscale Modeling	Enhanced Capability	Warfighter	Augmented Reality
Decrease	Power / Cost	Weight / Cost / Power	Cost	Cost

… Holistic Integrated Photonic Manufacturing Technology Platforms

- modular value chain – design, wafer fabrication, assembly, to packaging and test

- Core Manufacturing Locations

AnalogPhotonics — Massachusetts Institute of Technology — Massachusetts

SUNY Polytechnic Institute — UCSB — infinera — New York — California

University of Rochester — New York — samtec — Colorado

SUNY Polytechnic Institute — Columbia University — New York

- Workforce Coordination

Massachusetts Institute of Technology — Massachusetts

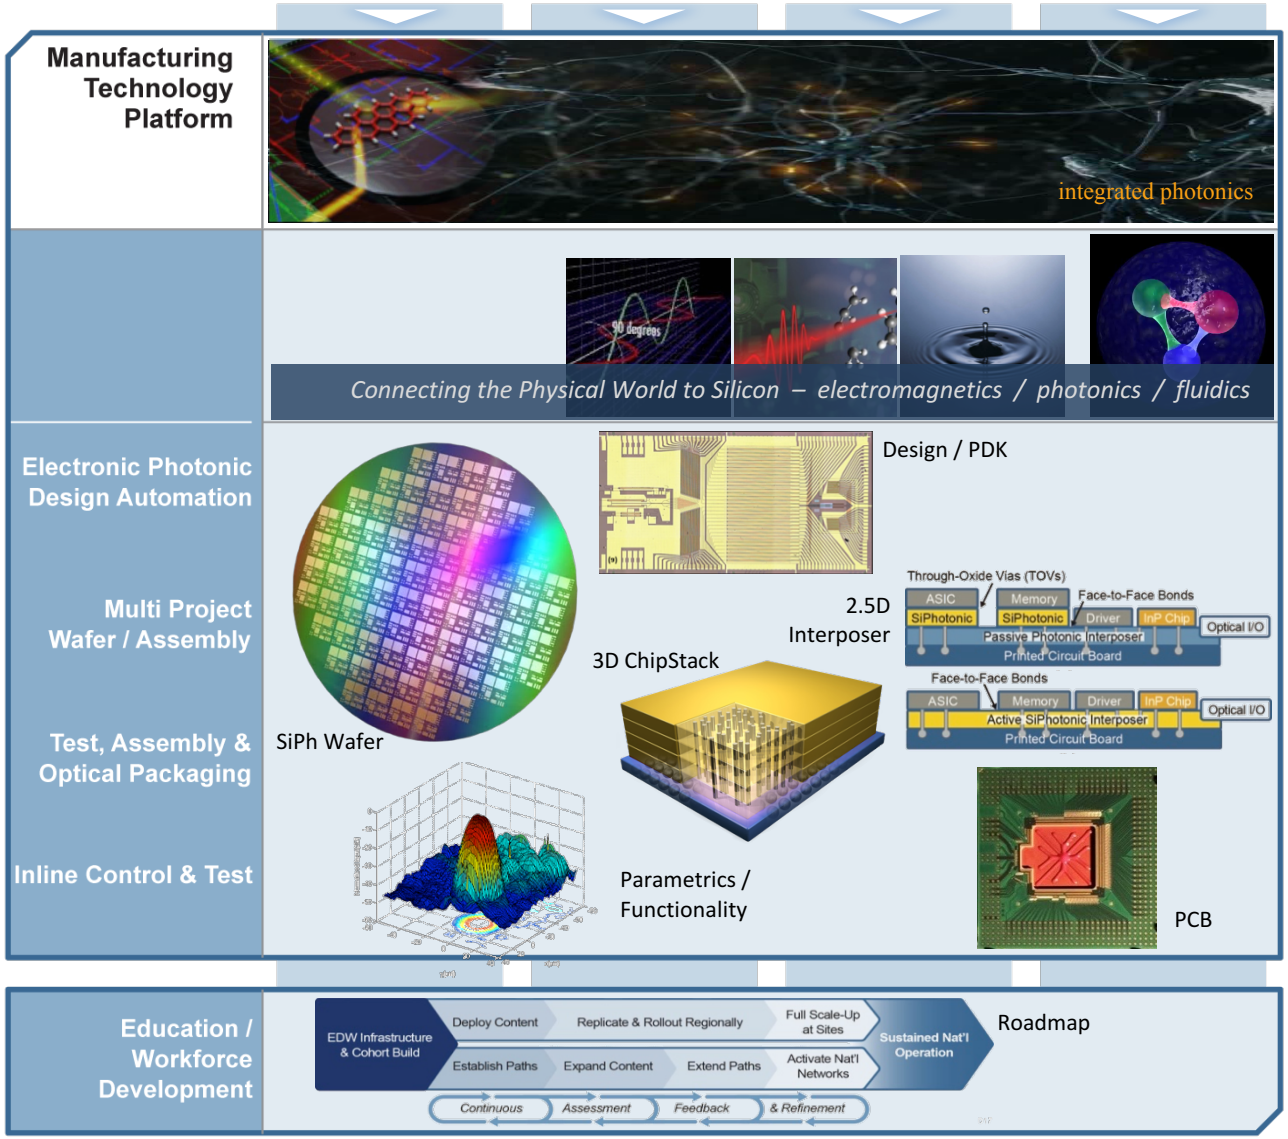

Supercomputing Performance

- Current World Top Supercomputers are Petascale:

  #1) TaihuLight (China) Peak: **125 PetaFLOPs (PF**)
  #2) Tianhe-2 (China) 55 PetaFLOPs (PF)
  #3) Titan (US) 27 PetaFLOPs (PF)

- Worldwide drive to reach Exascale in next few years
- Need a **10x** improvement factor to Exascale

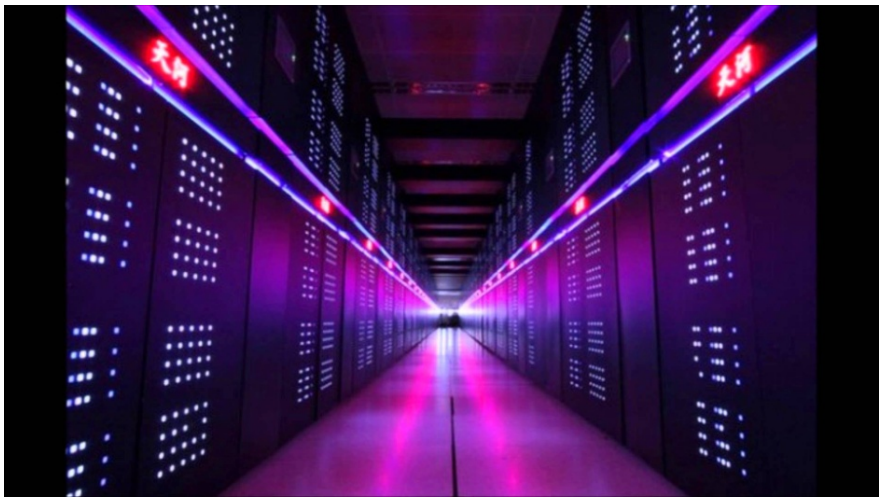

The Major Lag in Data Communications…

Top 10 Supercomputers computation capabilities over past 6 years:

- Vast increase in parallelism…but bandwidth is stagnated

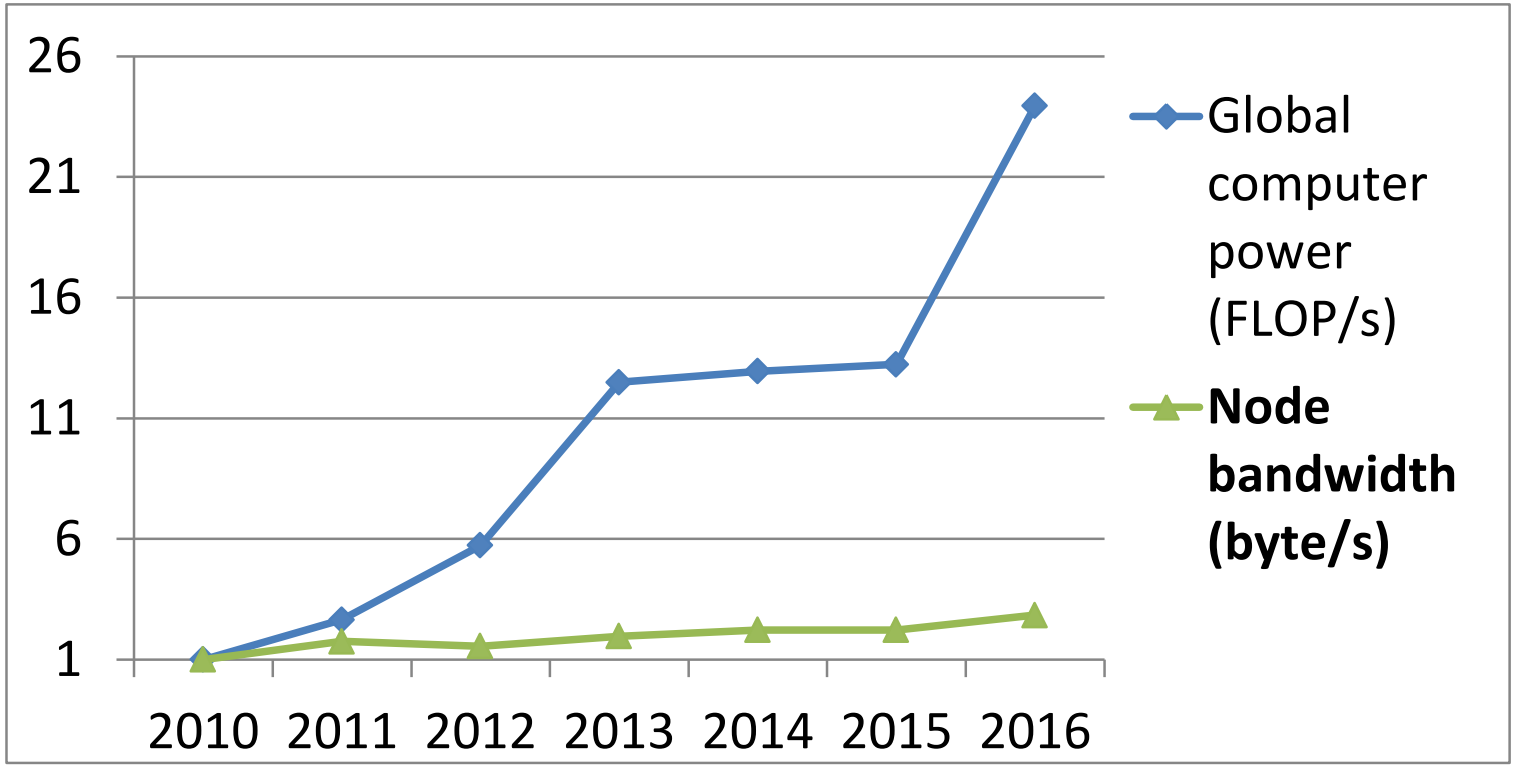

- While system compute power grows by **23X**
- Node I/O bandwidth increases by only **< 3X**

→ **Data-movement is too expensive! ($ and Energy)**

The Photonic Opportunity for Data Movement

- Energy efficient, low-latency, high-bandwidth *data interconnectivity* is the core challenge to continued scalability across computing platforms
- Energy consumption completely dominated by costs of data movement
- Bandwidth taper from chip to system forces extreme locality

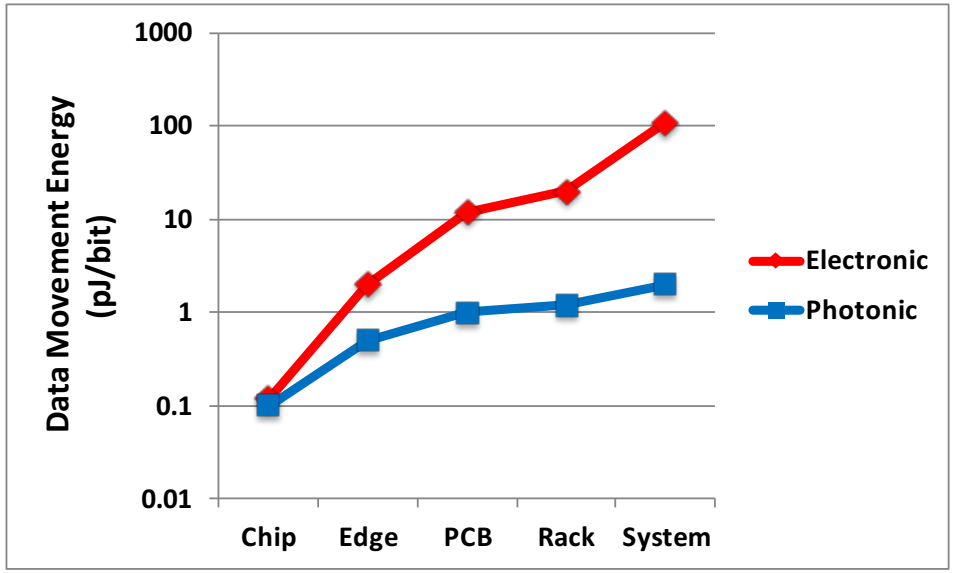

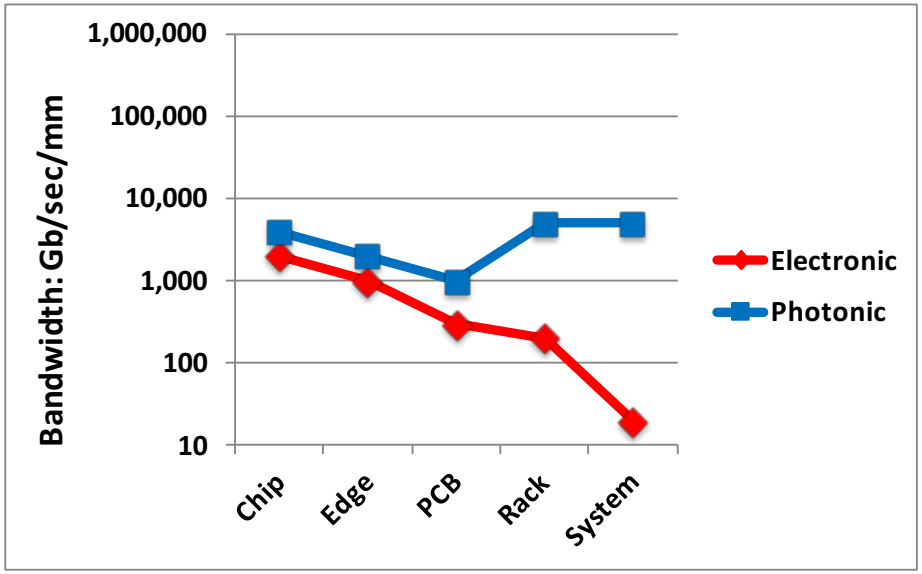

Reduce Energy Consumption

Eliminate Bandwidth Taper