Photonics TWG update

Chair Bill Bottoms MTS

Co-Chair Amr S Helmy University of Toronto

2023

Select Outcomes - Photonics TWG Chapter

- Where can photonics advance bandwidth density in interconnects
- Where can photonics advance high performance computing
- Where can photonics advance Data Centers

Co-Integration GF CLO45

- **Data rate limited by clocking circuitry. Clock distribution circuits limit the maximum reliable speed clock**

- **Using quarter-rate clocks in NRZ (most efficient architecture Today) it is possible to achieve 25Gb/s for ~6dB loss @ 12.5GHz ~2pJ/b**

  - VDD: 850mV +/- 6% 799mV
  - Temperature range: -40, 65 and 125
  - CLK Fmax: 14GHz WC SSLLLL

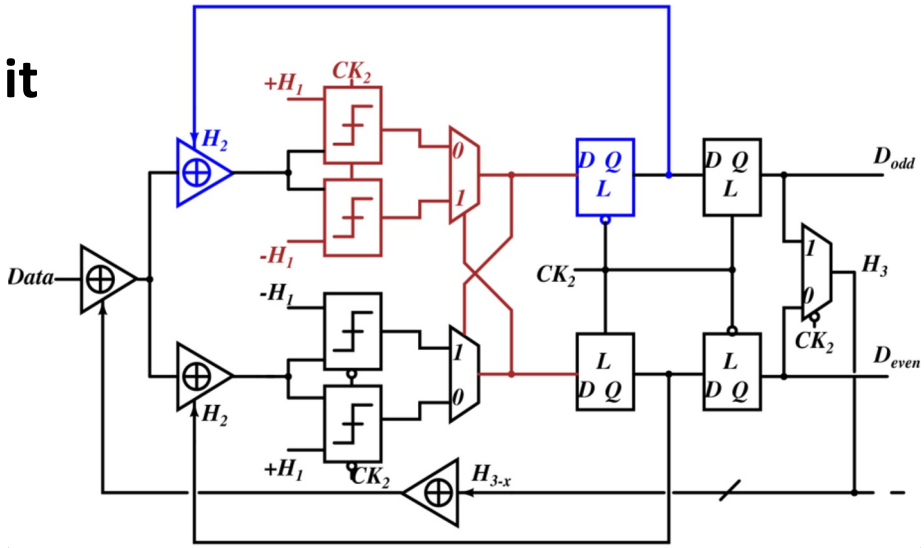

Co-Integration GF CLO45

- **Ayar Labs** : tested chiplet with 8 macros x 8λ x 25Gb/s NRZ upto AIB interface

  (~4pJ/b without laser)

- **Ranovus** : tested chiplet up to equalization. Using 22nm PAM4 interface for 100Gb/s per λ. (2.5pJ/b without laser)

- **Lightmatter**: tested chiplet with 32 macros x 32Gb/s NRZ before UCIe (~2pJ/b without laser)

- **Luminous:** (NA)

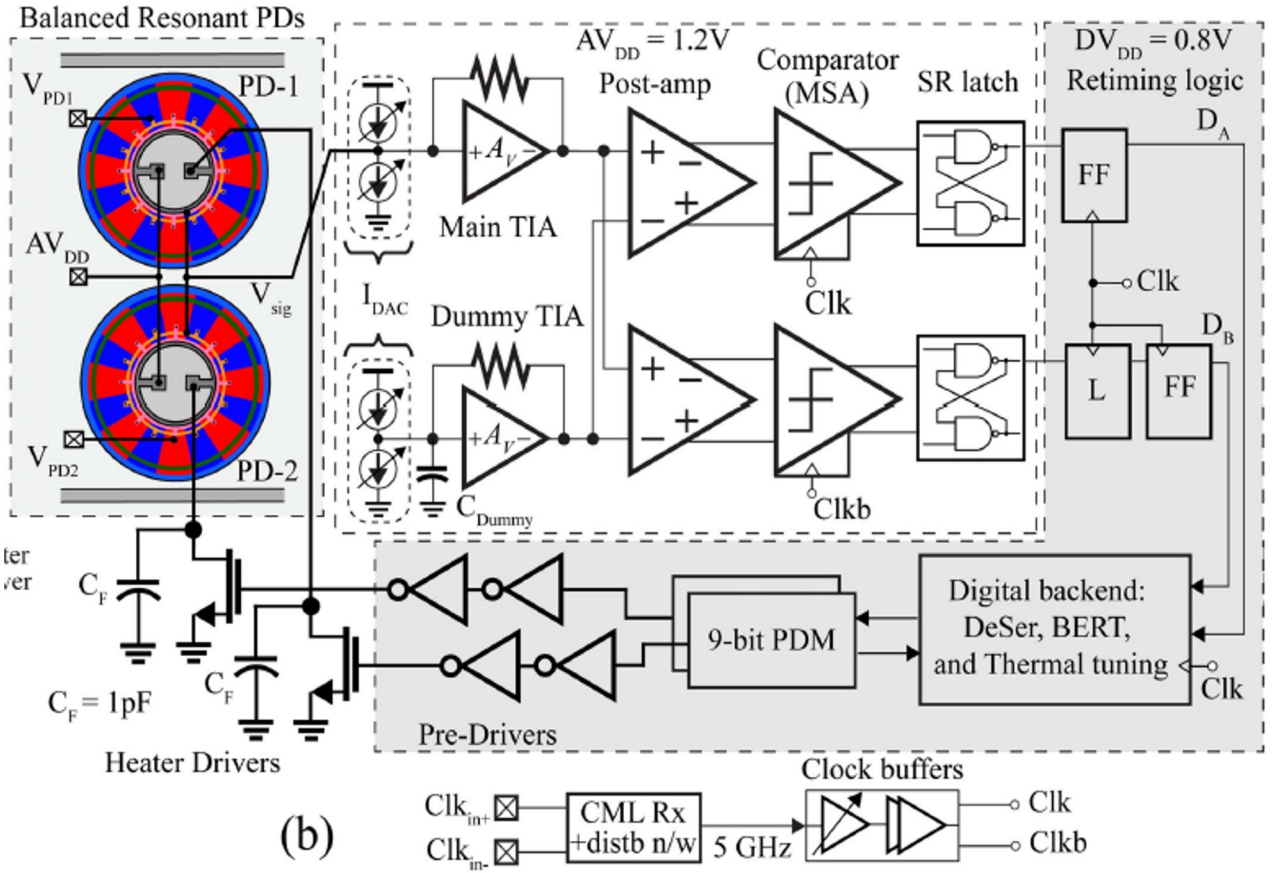

Optically Capable Interposer

HETEROGENEOUS INTEGRATION ROADMAP

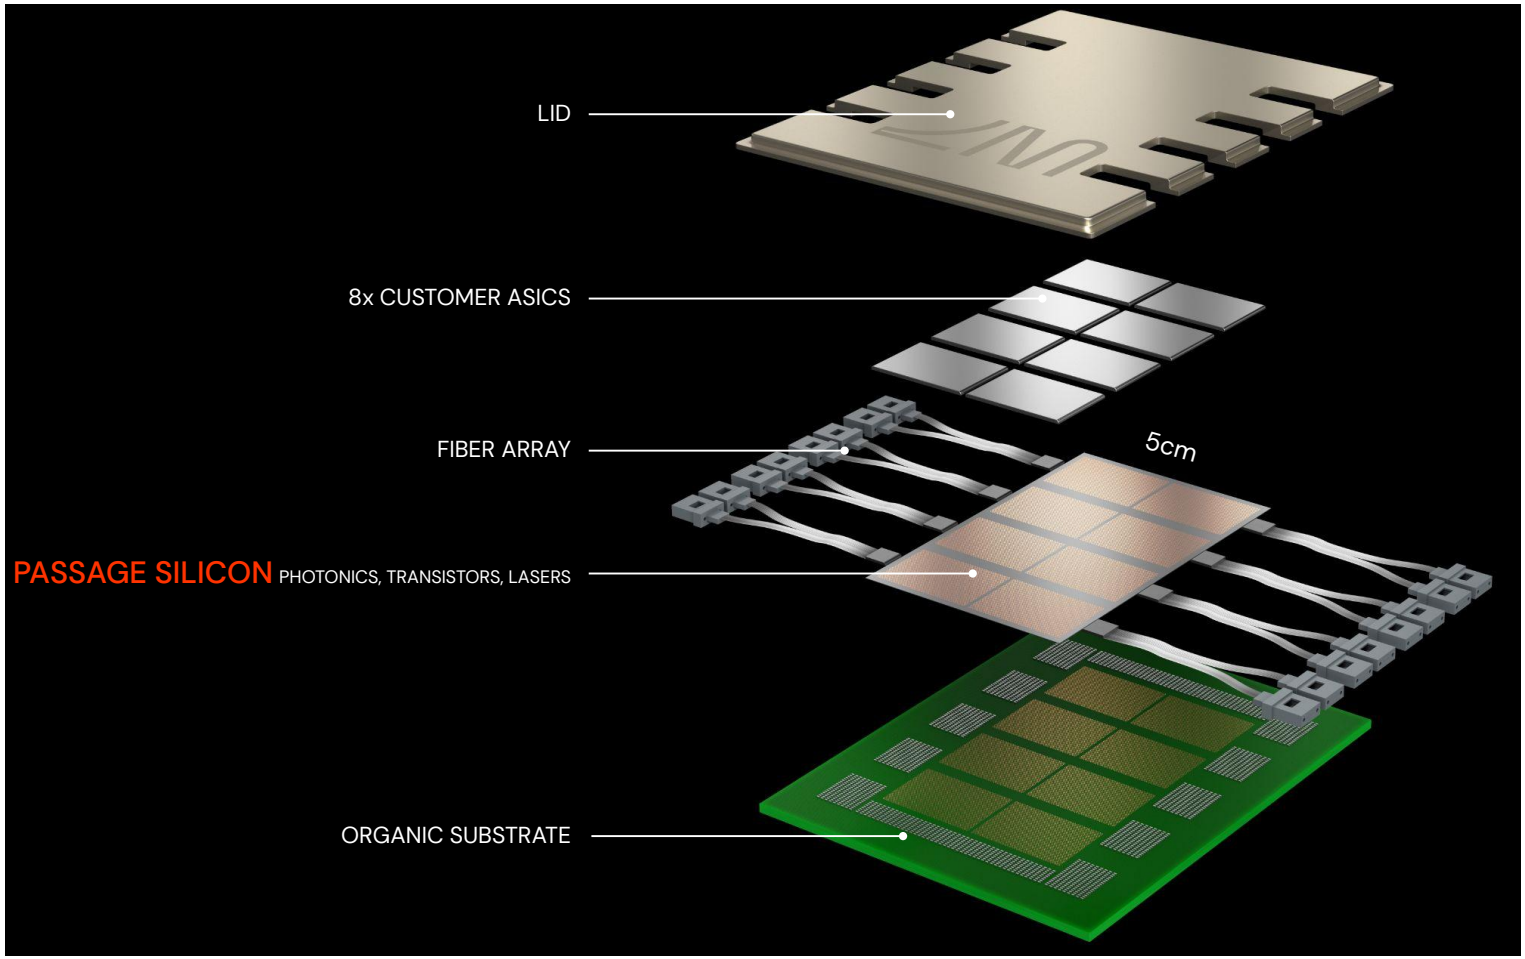

IEEE

IEEE ELECTRON DEVICES SOCIETY

Optical Capable Interposers

HETEROGENEOUS INTEGRATION ROADMAP

Solutions LM is Driving

A variety of applications

DISAGGREGATION

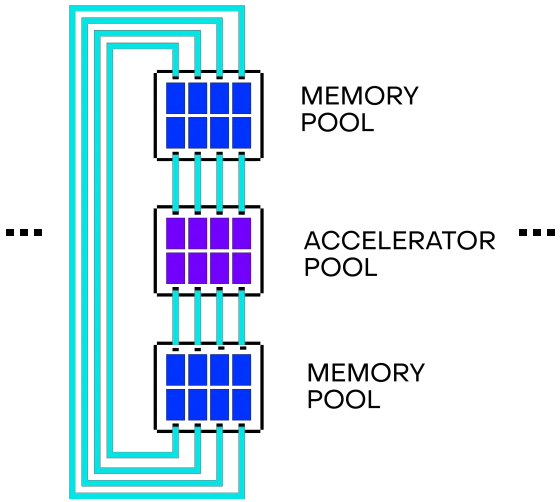

DYNAMIC COMPUTE ALLOCATION & AIR GAP ISOLATION

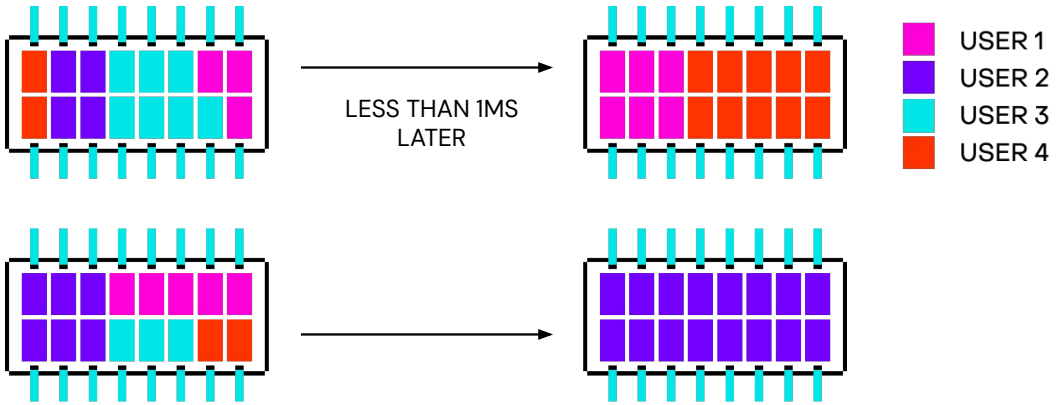

Optical Capable Interposers

Passage™ Alpha Link Performance

Model, loss, and data rates

32 links per tile, 32 Gbps per link (NRZ), 48 Tiles -> 48 Tbps

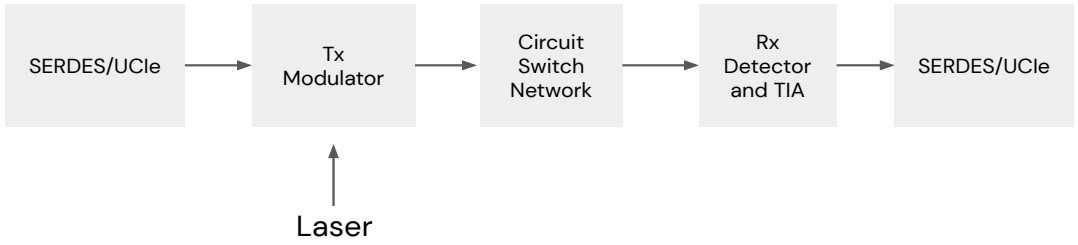

Link Characteristics

- 32 Gbps per channel NRZ
- 5 mW laser power at each Tx
- 0.5 dB/cm waveguide loss
- 0.08 dB per MZI
- 0.028 dB per crossing
- 0.004 dB per reticle crossing
- 1.1 A/W detectors

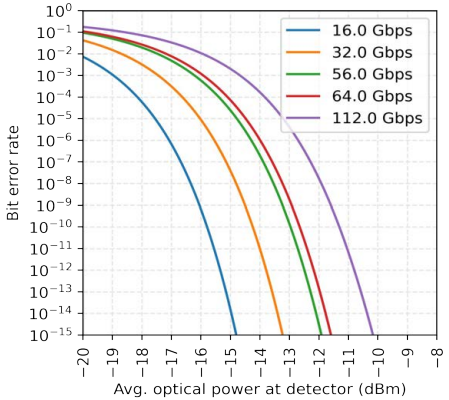

UCIe

Supports high-rate UCIe including 16, 32 Gbps

AMD Instinct MI300 (AI-HPC)

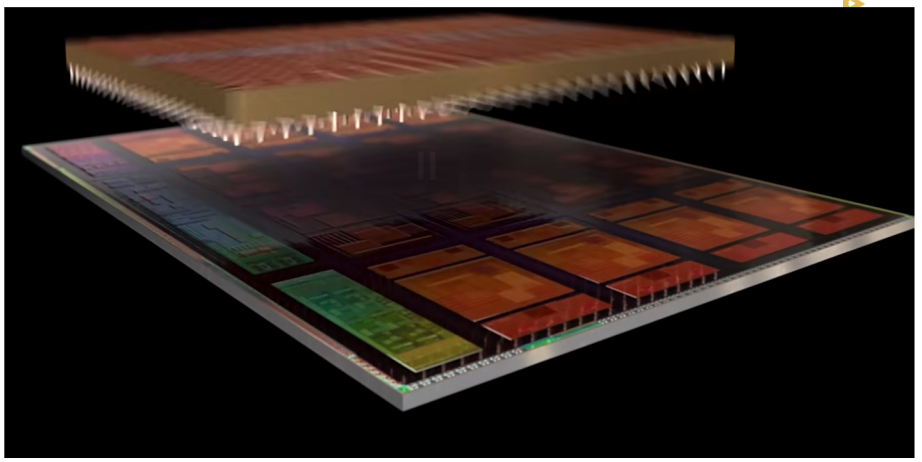

Application Driver

- ChatGPT model requires 175B parameters ~ at 32b = 700GB or 5.6Tb of storage
- ChatGPT will require 20 dies of 32GB HBM

128GB HBM3 / 8 dies 16GB per die @6.4Gbs - 3pJ/b?)
16 x 64-bit channels → 819.2GB/s → 6.5Tb/s per die!
(1~2ns latency) → vertical cache!

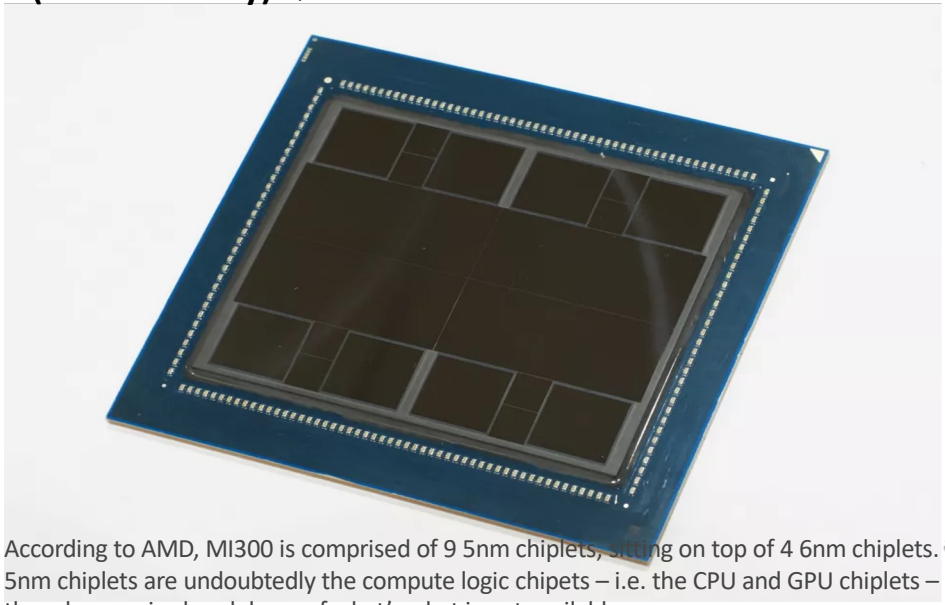

According to AMD, MI300 is comprised of 9 5nm chiplets, sitting on top of 4 6nm chiplets. The 5nm chiplets are undoubtedly the compute logic chipets – i.e. the CPU and GPU chiplets – though a precise breakdown of what's what is not available.

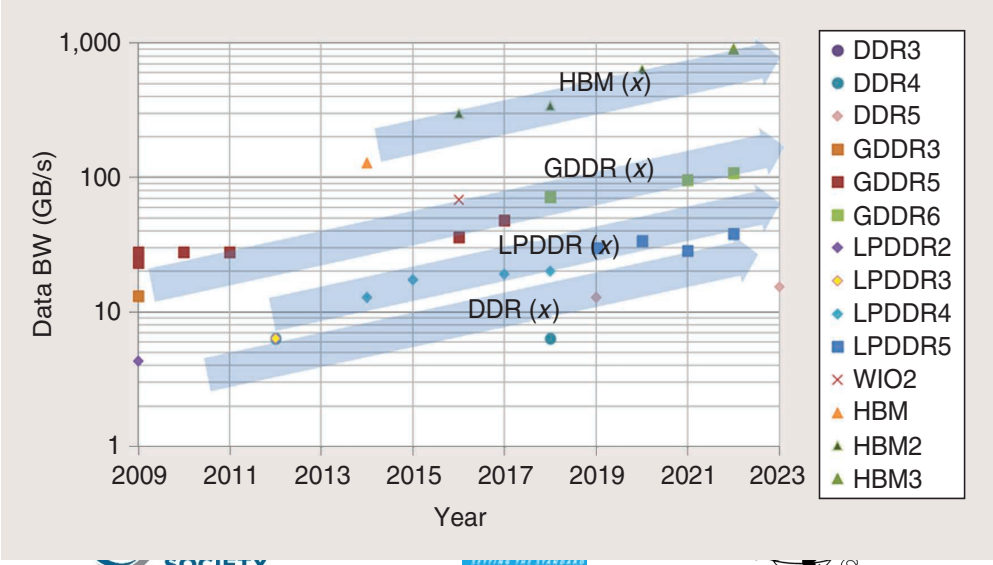

FIGURE 26: The DRAM data BW growth. LPGDDR: low-power GDDR; WIO: wide I/O.

Data Centers

Rakesh Chopra, Cisco Systems

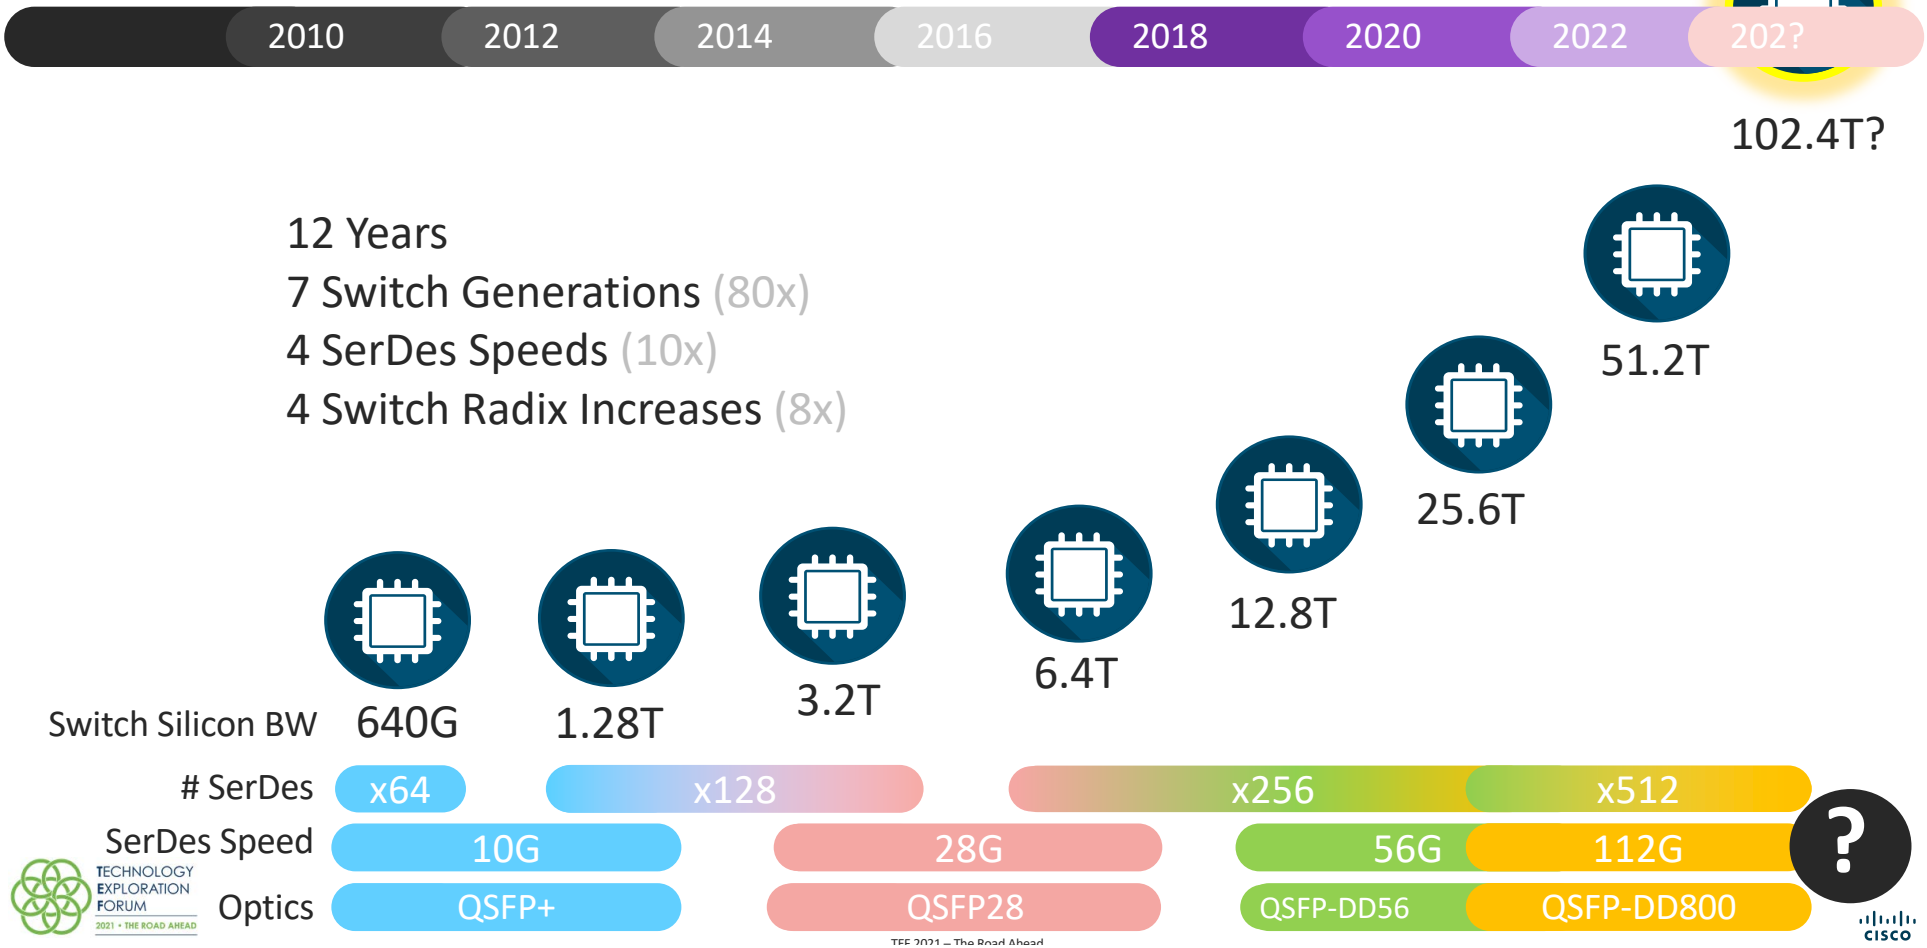

Data Centers

Rakesh Chopra, Cisco Systems

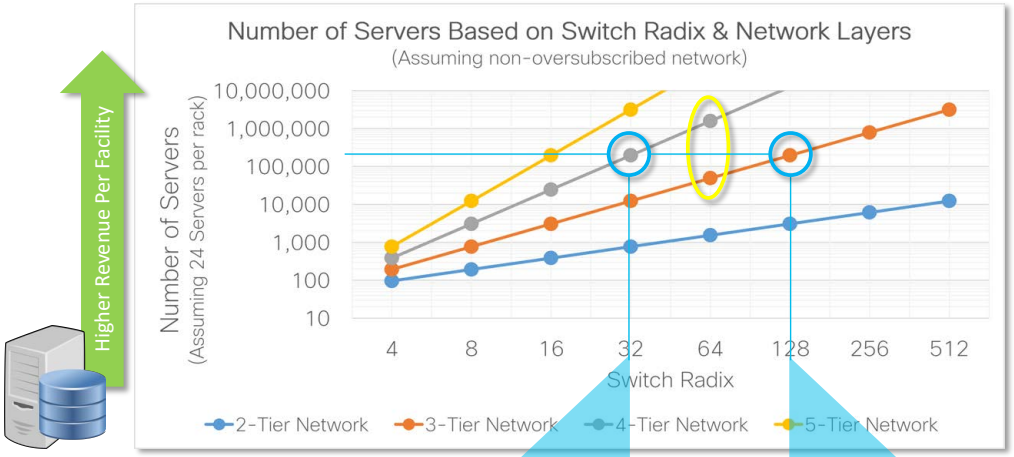

Graph concept leveraged from R. Nagarajan, Ilya Lyubomirsky, "Next-Gen Data Center Interconnects: The Race to 800G"
Adjusted to hold servers per rack constant

Switch BW	SerDes	Radix x32	Radix x64	Radix x128
12.8T	56G	400GE x8	200GE x4	100GE x2
25.6T	112G	800GE x8	400GE x4	200GE x2
51.2T	112G	1.6TE x16	800GE x8	400GE x4
102.4T?	224G	3.2TE x16	1.6TE x8	800GE x4

Wider Radix - Scale Out
More Layers – Scale Up

Improved Power Efficiency
Improved Link Utilization

- x32 and x128 radix are prominent today
  - Ethernet rates are lagging for x32 radix
  - Will x32 networks migrate to x64?

Radix 64

- **Potential** need for 800GE with 8x112G Lanes
  - 51.2T
  - 64 x QSFP-DD800 (carrying 1x800GE) – 2RU
- **Potential** need for 1.6TE with 8x224G Lanes
  - 102.4T
  - 64 x QSFP-DD1600 (Carrying 1x1.6TE) – 2RU

Radix 128

- **Clear** need for 800GE with 4x224G Lanes
  - 102.4T with 128-Radix
  - 128 x QSFP-800 (carrying 1x800GE) – 4RU

    or
  - 64 x QSFP-DD1600 (carrying 2x800GE)-2RU

Data Centers.

Rakesh Chopra, Cisco Systems

Higher data rates and distance drive the move from copper to optics

Future innovations will only be possible with **silicon** and **optical** integration

Si Silicon + **Op** Optics

51.2T

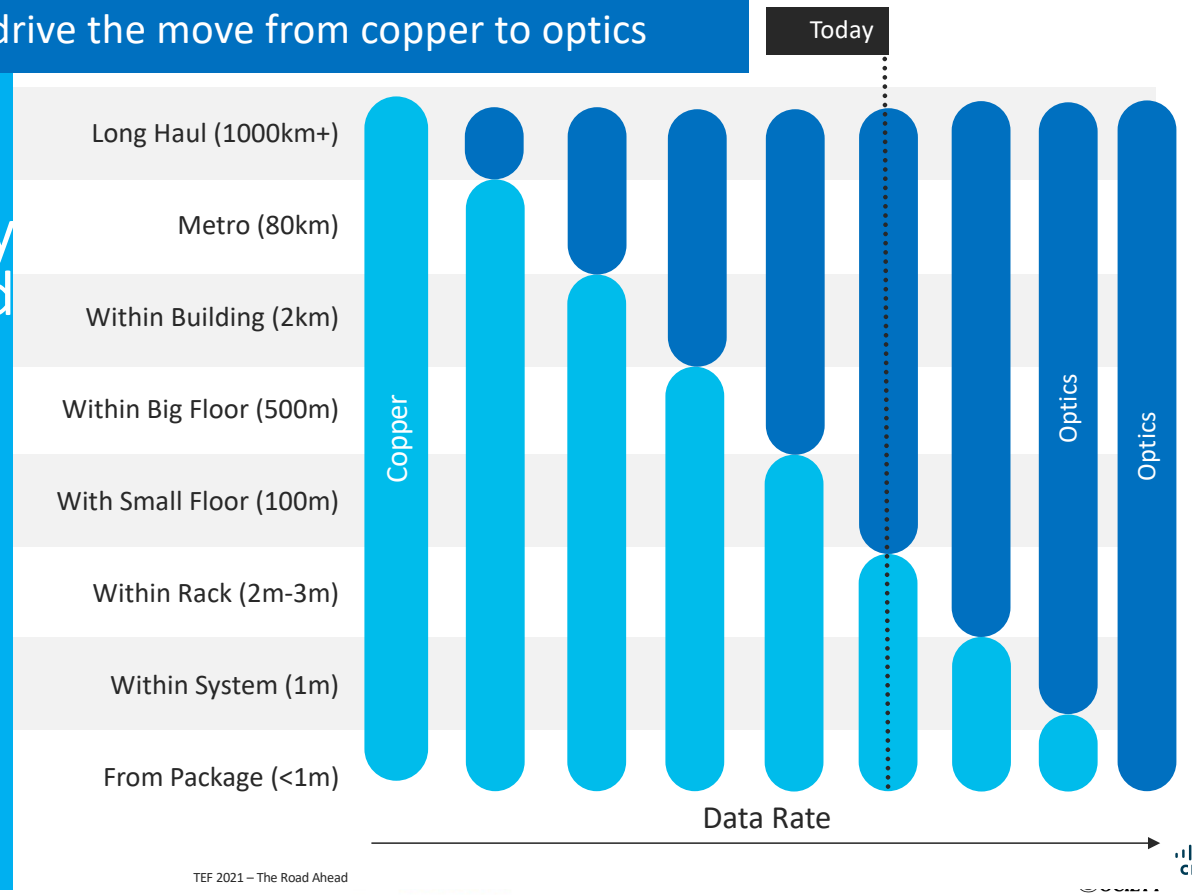

Conclusion

Integrated photonics is NOW REPLACING CMOS

- delivering the low power, low cost, low latency and increased bandwidth density and performance required to meet the demands of the data and data traffic explosion that will continue rapid growth.

The physical bandwidth density needed

- to support big data, data analytics, artificial intelligence and neuromorphic & quantum computing is increasing exponentially.

This will only be possible by getting optics / logic in the same package

Thank You

- Please get involved
- Emails us or leave comments
- a.helmy@utoronto.ca

Integrated Photonics TWG Membership

- Bill Bottoms Chair
- Amr Helmy Co-Chair
- Guillermo Carpintero-del-Barrio
- Thomas Brown
- Mayank Bulsara
- David Butler
- Benjamin Fasano
- Matthew Grandbois
- Thomas Marrapode
- Lei Shan
- Mingxui Sun
- Ravi Tummidi
- Ehrenfried Zschech
- Lei Wang