How can Photonics Enable the Bandwidth Densities with Lower Energy per Bit in Emerging SIP?

Friday, June 2, 2023, 8:00 a.m. – 9:15 a.m.

Chairs: Kitty Pearsall (Boss Precision Inc., IEEE EPS President) and David McCann (Lyte)

ECTC 2023 IEEE EPS President's Panel

Chair
Kitty Pearsall
IEEE EPS President
Boss Precision Inc.

Chair
David McCann
Lyte

Panelist
Amr S. Helmy
University of Toronto

Panelist
Stefano Oggioni
AT&S

Panelist
Ritesh Jain
Lightmatter

Panelist
Ajey Jacob
University of Southern California

How can Photonics Enable the Bandwidth Densities with Lower Energy per Bit in Emerging SIP

This panel will discuss the tools, technologies and approaches that will enable the industry to enhance the bandwidth density of interconnections in SiP enabled by photonics. To be adopted, such capabilities must be provided with energy per bit which meets the roadmaps and standards targets for the interconnection protocols within the package and on chip.

IEEE ELECTRONICS PACKAGING SOCIETY

IEEE

Interconnect Solutions

- Wide range of solutions with growing energy cost as a function of communication distance

Mainly Electrical Solutions

Photonic Solutions

On-chip
- Direct gate-to-gate link
- < 0.1 pJ/bit
- ~20 Tbps/mm

In-package
- AIB, UIB, USR, HBM
- 0.1-2 pJ/bit
- ~2 Tbps/mm

On-board
- PCI-E, LVDS, DDR, SERDES
- 2-20 pJ/bit
- ~0.5 Tbps/mm

Off-board
- Directly attached copper cable or active optical cable
- >20 pJ/bit
- ~100 Gbps/mm

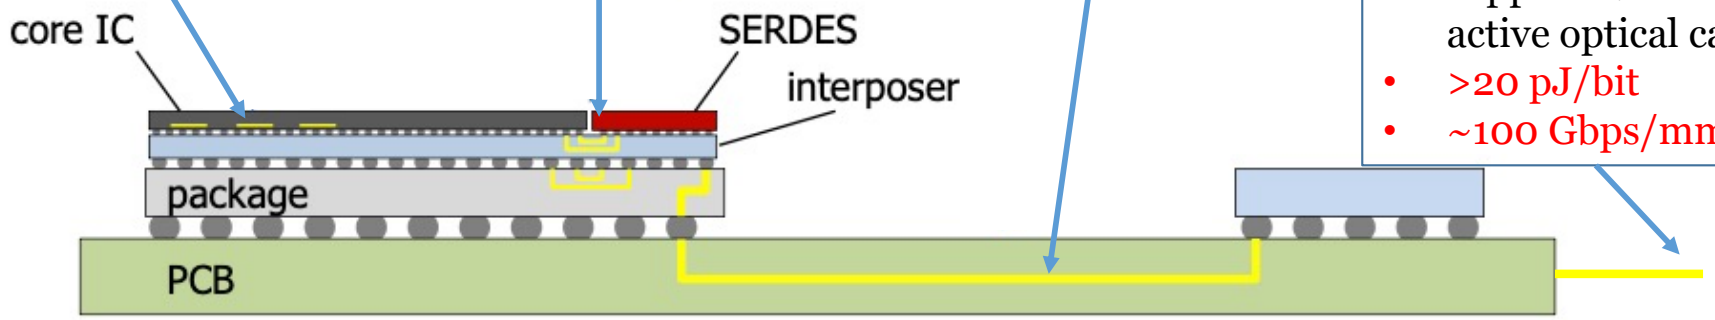

"Photonics in the Package for Extreme Scalability (PIPES)" DARPA (2018)

Data Center Solutions

Rakesh Chopra – CISCO

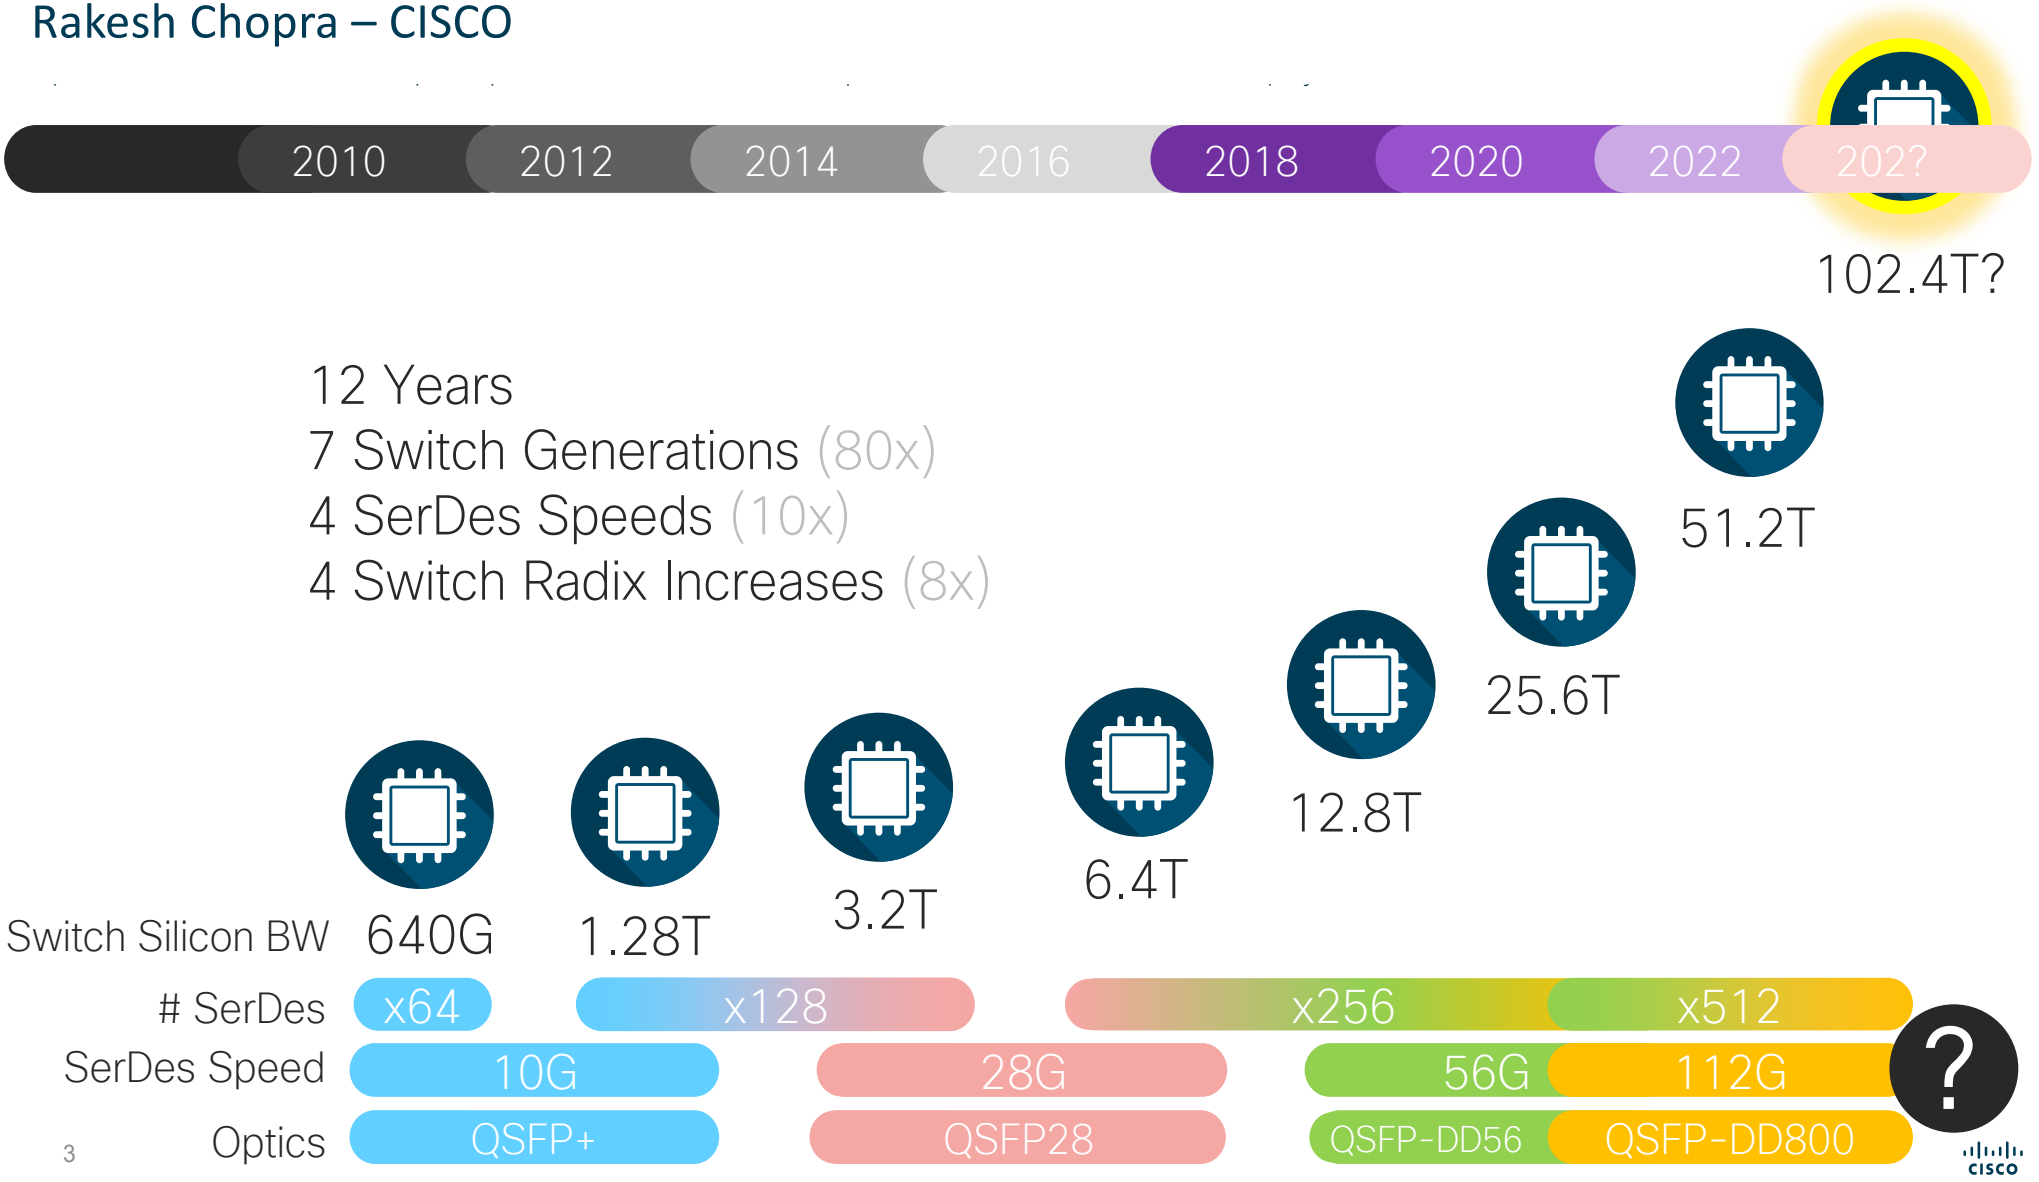

Data Center

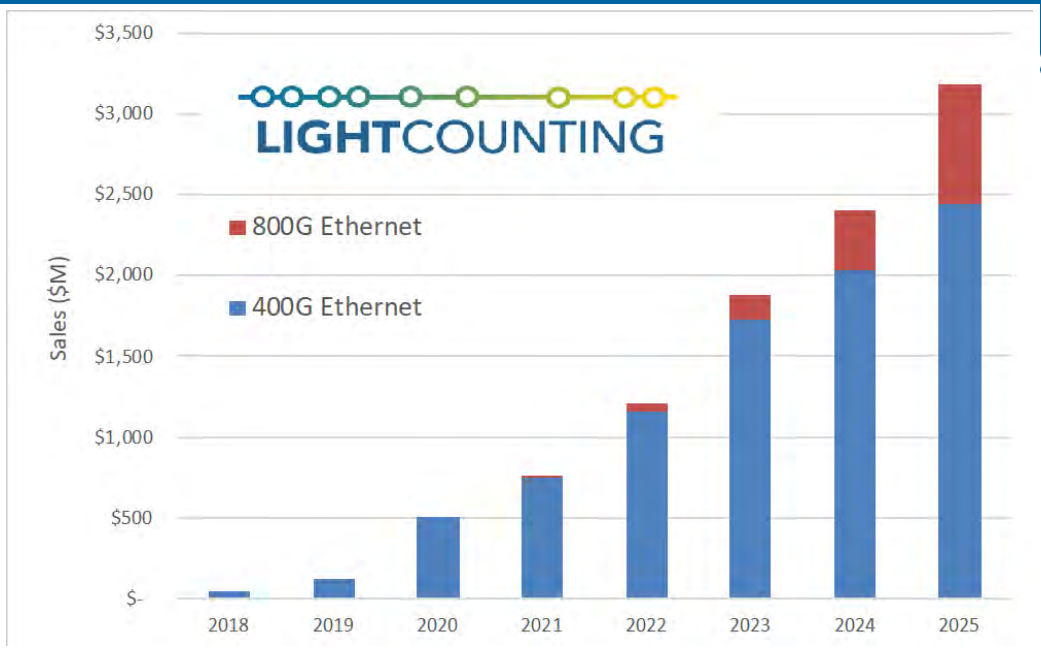

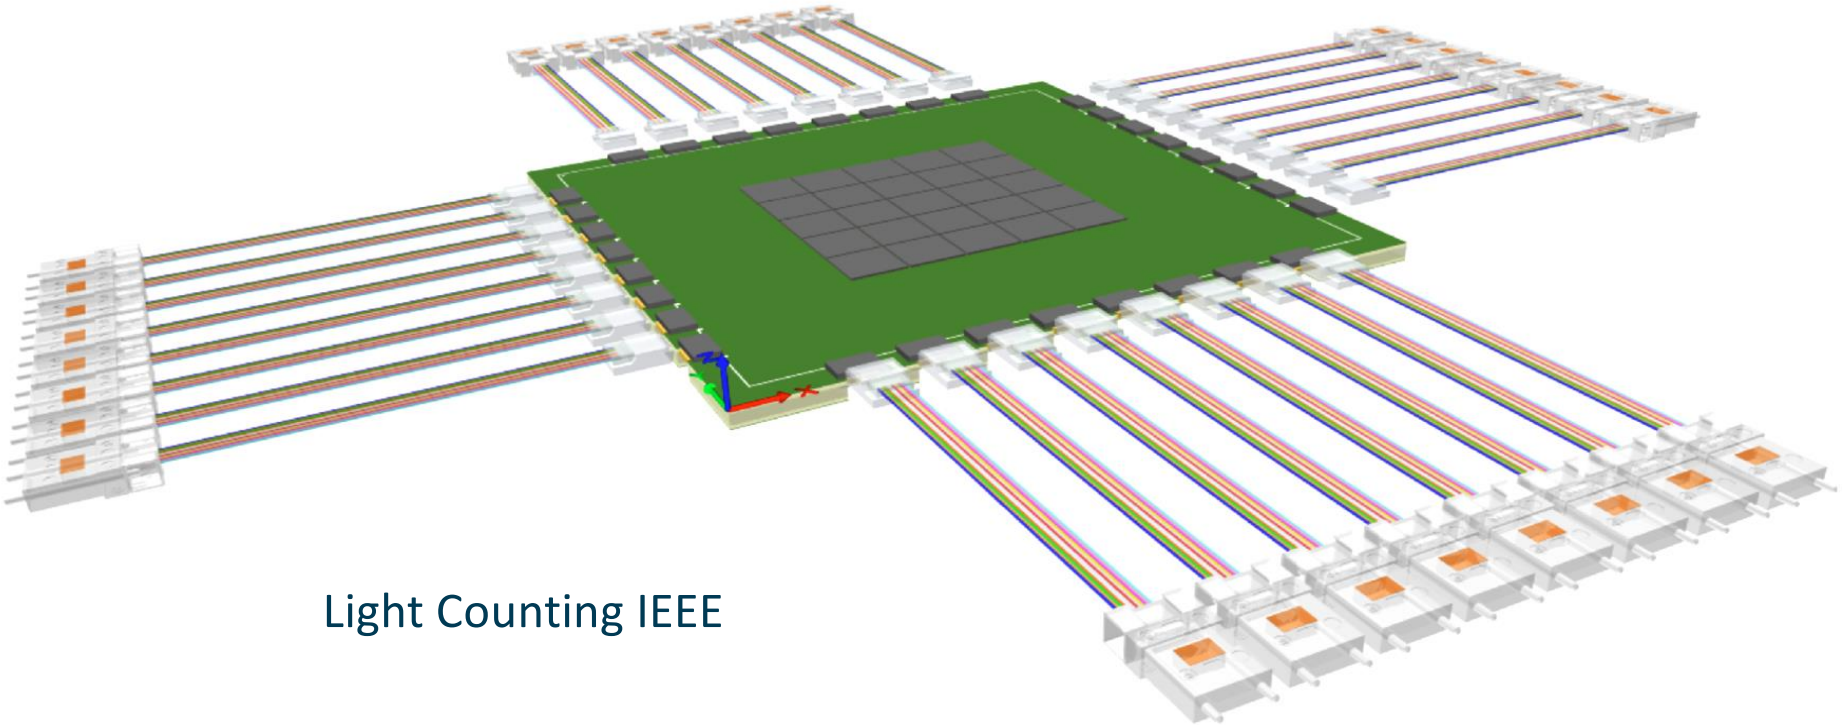

Light Counting IEEE

Data Center Solutions

Disaggregated designs will need 100x bandwidth

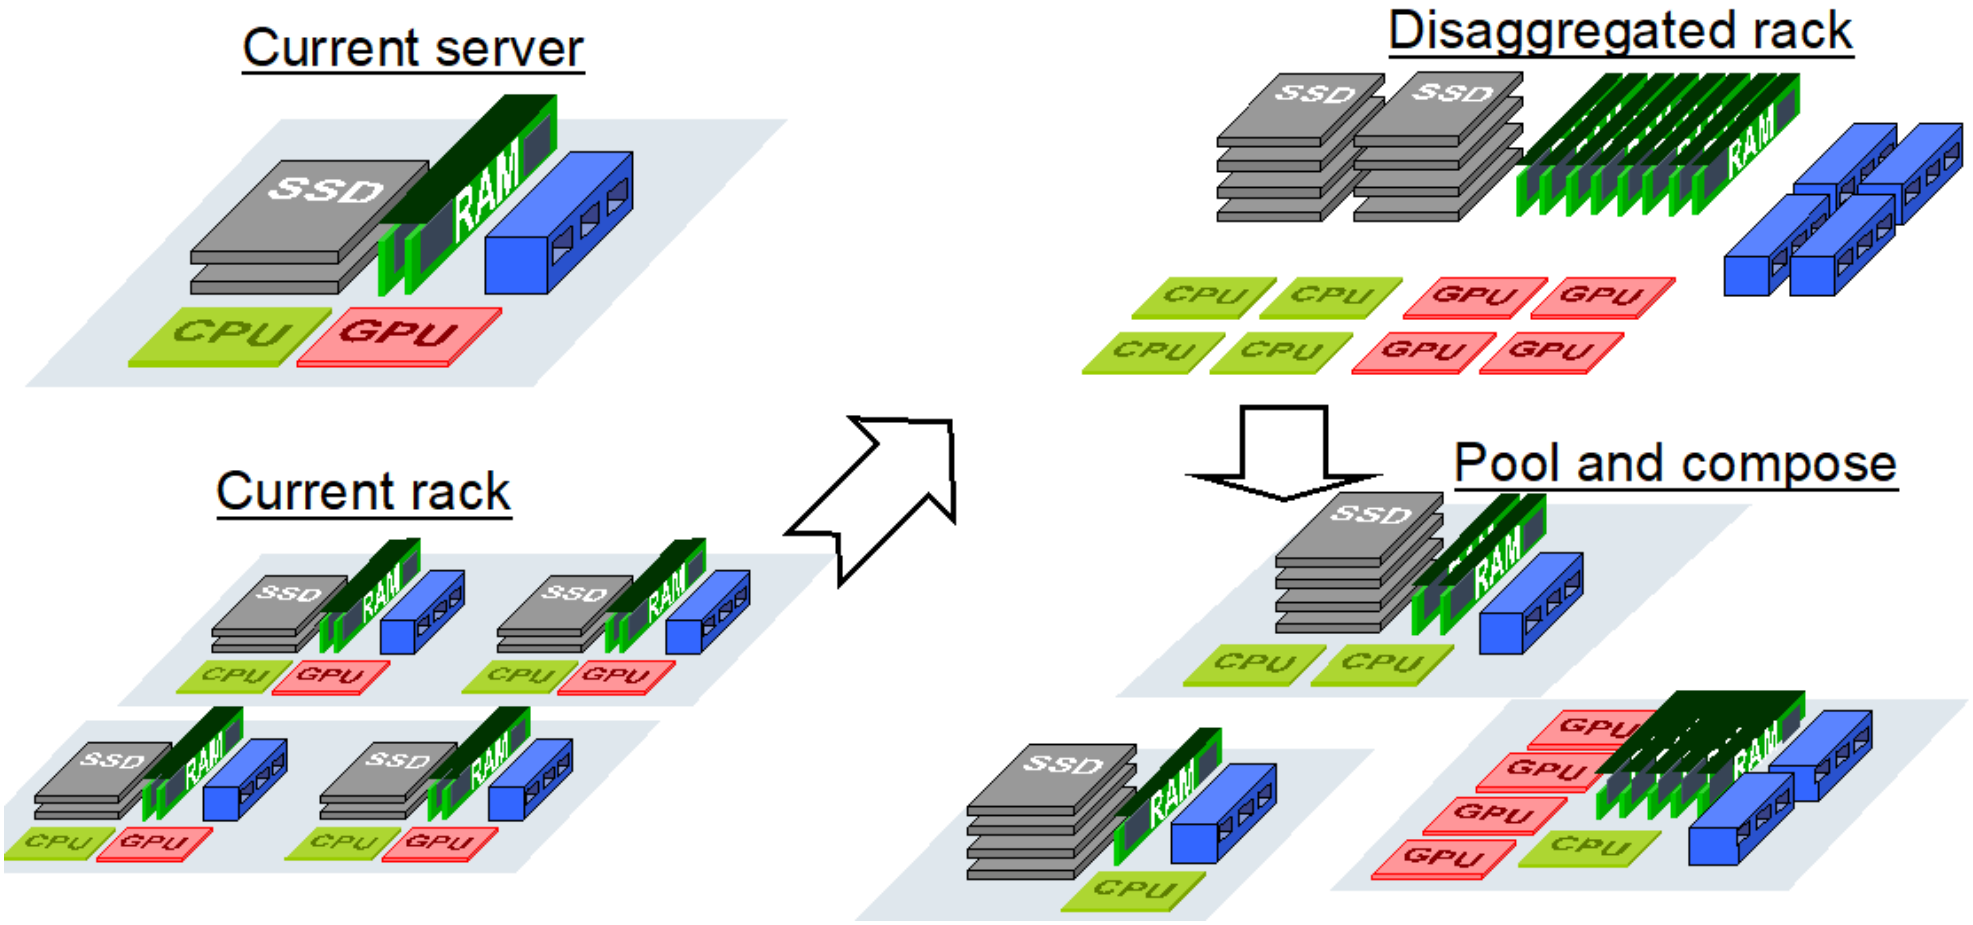

Example – NVIDIA H100

- Up to **512GB LPDDR5**
  - **6x more** than GPU HBM
- **900 GB/s** CPU-GPU BW
  - **7x PCIe Gen5** bandwidth
  - Hardware **coherent**

PCIe Gen 5 is used now; How much do Gen 6 and 7 scale?

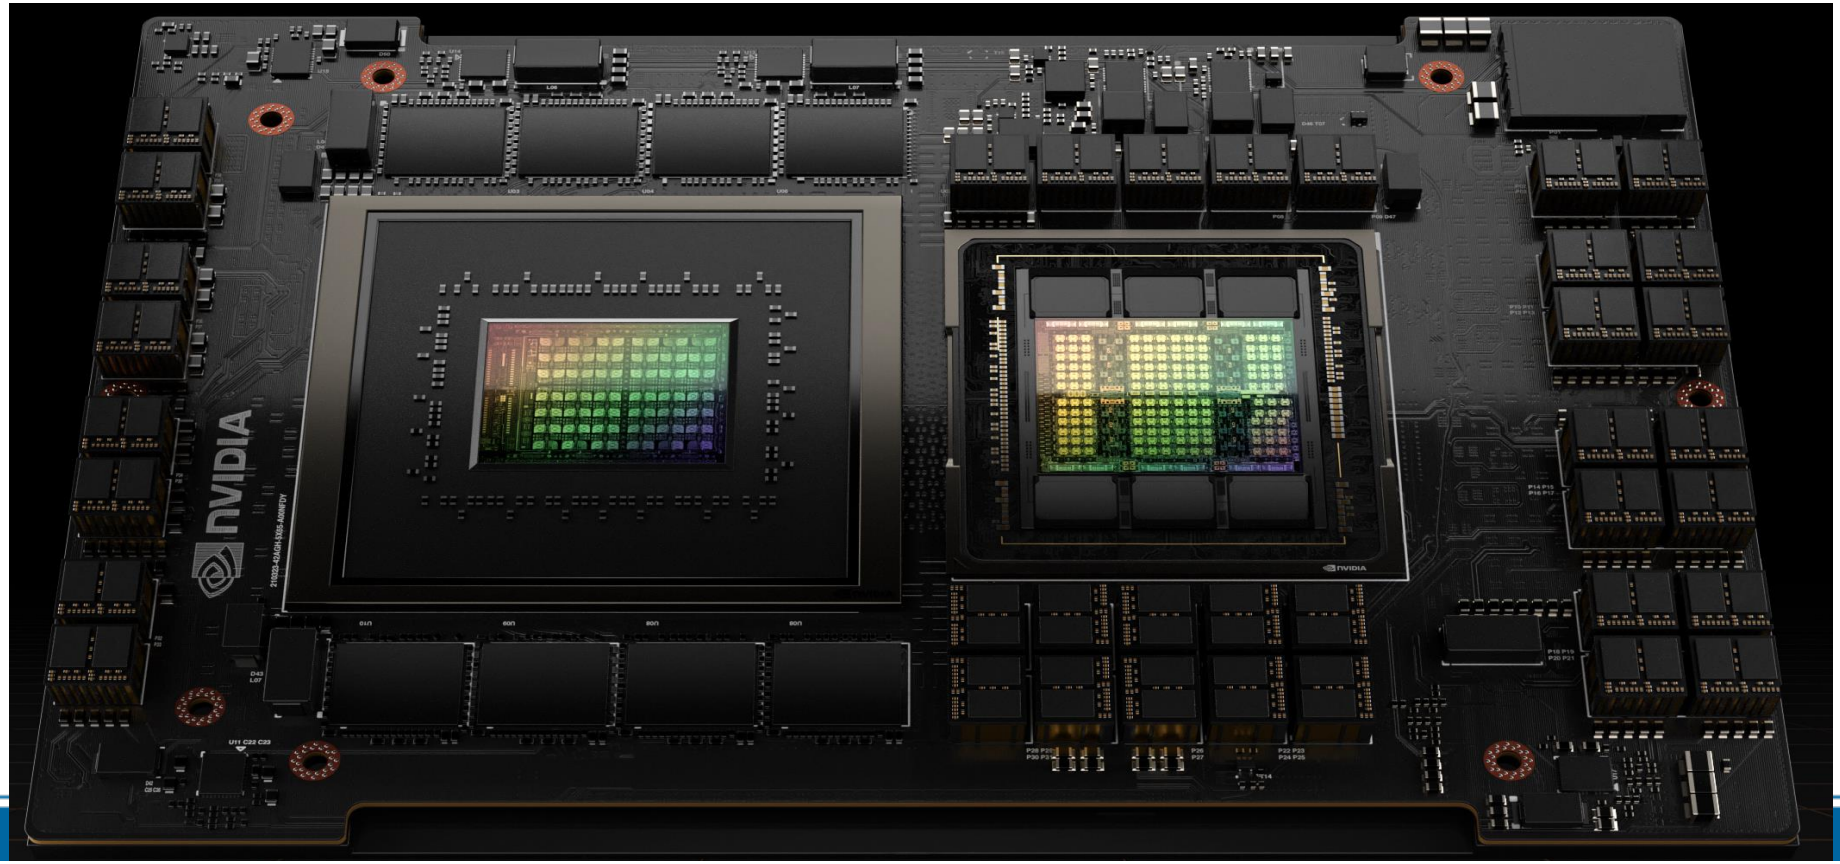

Cerebrus Design

The CS-1 wafer is a distributed-memory machine with a **2D-mesh interconnection fabric.**

The repeated element of the architecture is called a tile.

The tile contains one processor core, its memory, and the router that it connects to.

The routers link to the routers of the four neighboring tiles.

The wafer contains a 7×12 array of 84 identical “die.” A die holds thousands of tiles.

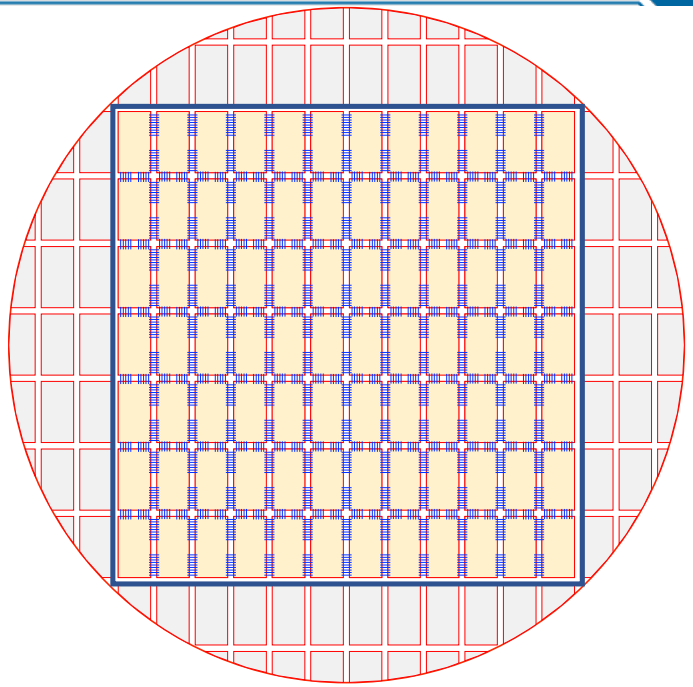

Existing Cerebrus system offer 1.2 Terabits/sec using electrical System IO over 12x standard 100 GbE

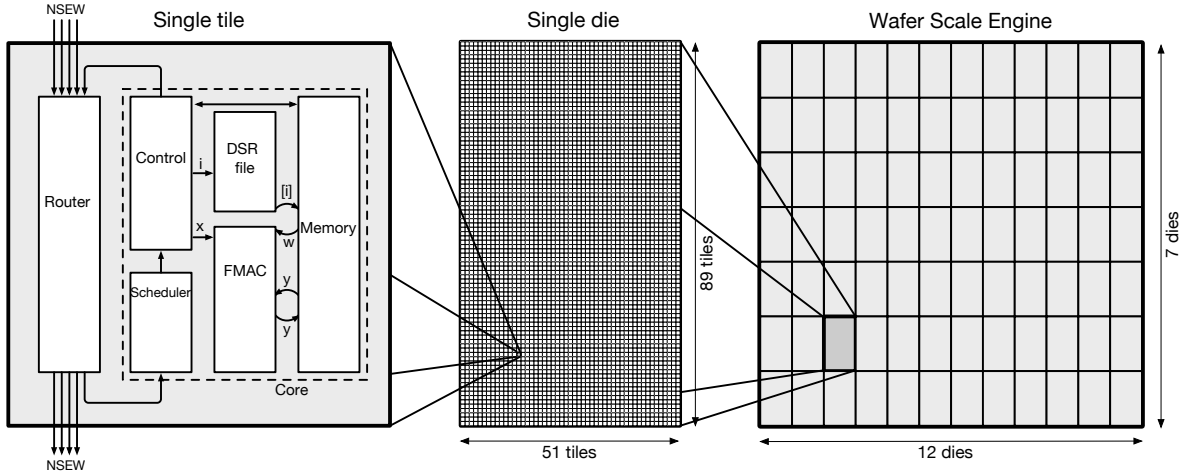

Fig. 2. CS-1 Wafer Scale Engine (WSE). A single wafer (rightmost) contains one CS-1 processor. Each processor is a collection of dies arranged in a 2D fashion (middle). Dies are then further subdivided into a grid of tiles. One die hosts thousands of computational cores, memory and routers (leftmost). There is no logical discontinuity between adjacent dies and there is no additional bandwidth penalty for crossing the die-die barrier. In total, there are 1.2 trillion transistors in an area of 462.25 cm^2.

Kamil Rocki, Dirk Van Essendelft, Ilya Sharapov, Robert Schreiber, Michael Morrison, Vladimir Kibardin, Andrey Portnoy, Jean Francois Dietiker, Madhava Syamlal and Michael James, Fast Stencil-Code Computation on a Wafer-Scale Processor, SC20, November 9-19, 2020.

opto.utoronto.ca

a.helmy@utoronto.ca

ECTC EPS Presidents Panel

Ritesh Jain
VP, HW Engineering (Systems & Packaging)
Lightmatter

LIGHTMATTER

IEEE ELECTRONICS PACKAGING SOCIETY

Challenges

Chip power is exploding. The death of Dennard scaling[1] is evident. Water cooling is widely deployed. Immersion is up next.

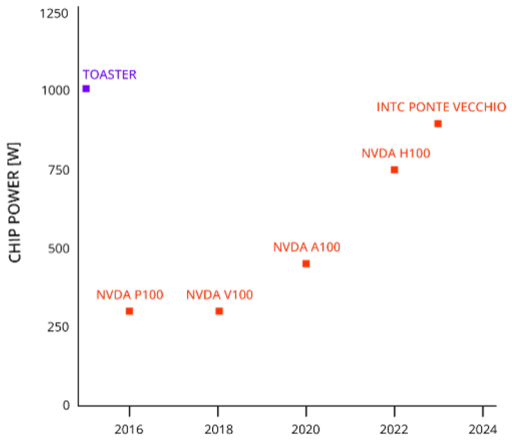

Notes
[1]Dennard Scaling article.
[2]Publicly available data.

AI will consume 80% of data center power by 2040. Overall data center power usage is growing much slower than AI hardware deployments. Both growth rates are massive.

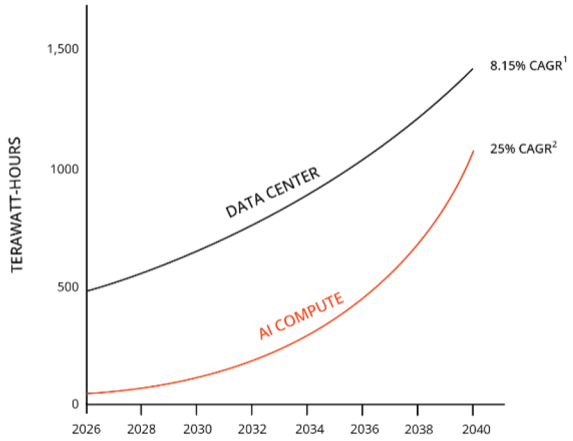

Notes
[1]Interpolation based on IEA. There are estimates from 6% to 35% annual growth rate.
[2]Averaged CAGR over 2026-2040 smoothed down from 35% growth rate from 2026 - 2030. Source: IEA
[3] Actual data center growth rate is not agreed upon in academia. Source.

AI models will be 1000x larger in 2030. Deployed AI models double in size every year. Research models double 6x faster.

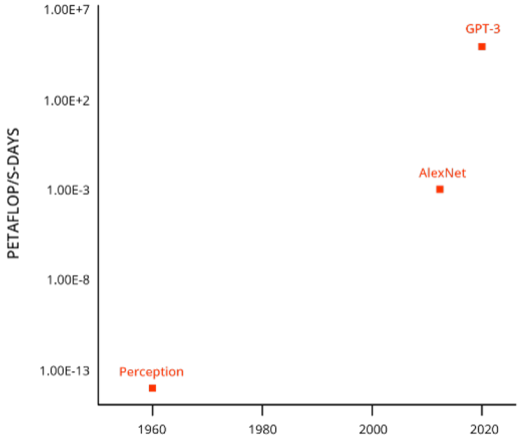

Notes
OpenAI "AI and Compute" article.

LIGHTMATTER

IEEE ELECTRONICS PACKAGING SOCIETY

Compute requirements for AI will increase 1000x over the next decade, conservatively.

Performance and Cost Challenges

Transistor performance improvements are slowing. Compute performance is bound by thermal limitations at the package level.

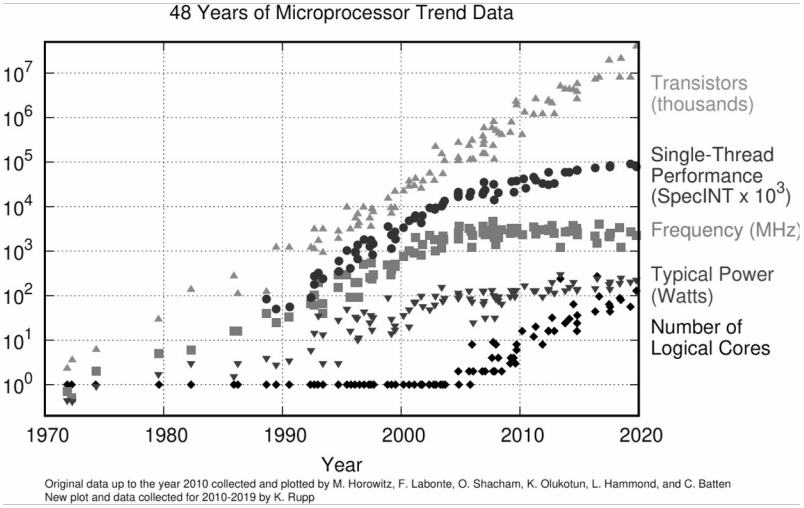

Chip design cost is exploding. Exponential growth in development costs for creating next-generation processors.

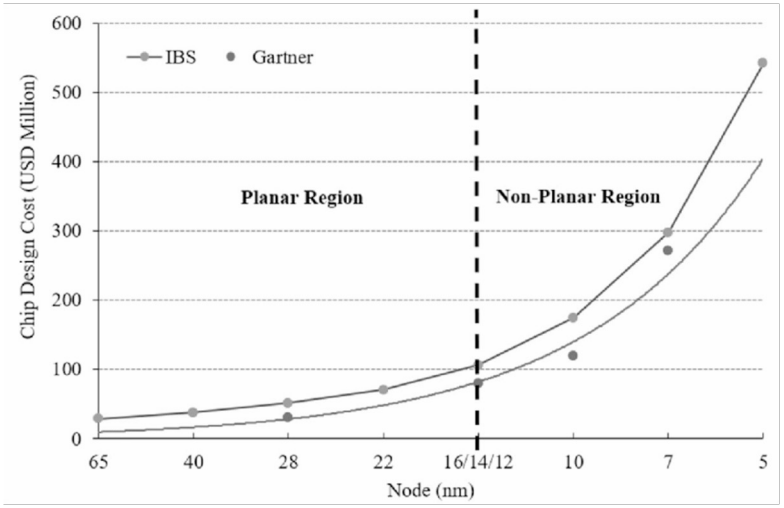

LIGHTMATTER

IEEE ELECTRONICS PACKAGING SOCIETY

Heterogeneous Integration

Okay, transistors have issues. Let's use packaging to help.

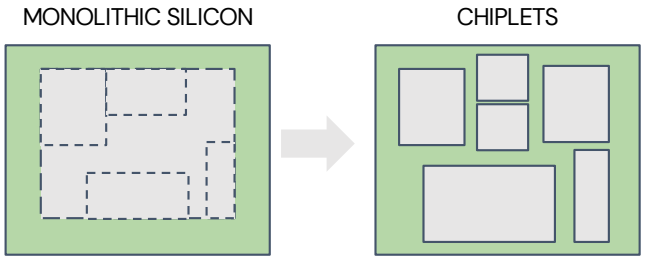

Driven by

- Dev. and wafer cost
- Integration of multiple nodes, IP
- Si-node yield resiliency
- TTM
- Higher memory density (HBM)
- Energy efficiency
- Compute scaling

2D D2D, PRIMARILY HORIZONTAL

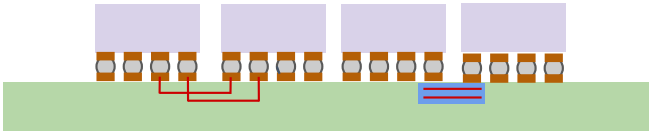

FCBGA, EMIB, RDL-BASED

3D D2D, PRIMARILY VERTICAL

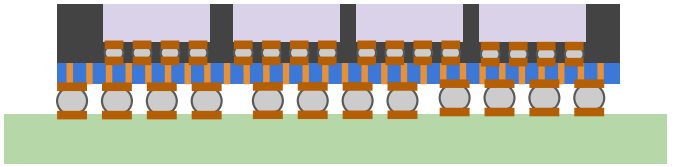

FOVEROS, HBI/V-CACHE

Chiplet ≠ Monolithic

Bandwidth Requirements

BISECTION BANDWIDTH

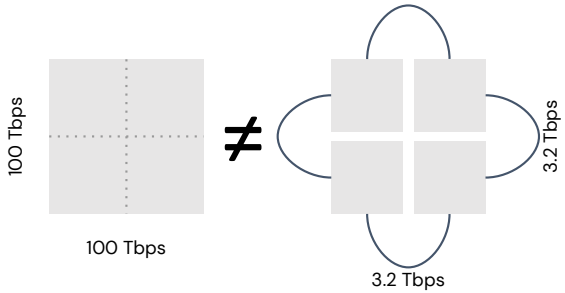

Beachfront and bandwidth are fundamentally linked in chiplet processors. Big chips are at odds with high yield.

MORE HOPS, MORE ENERGY

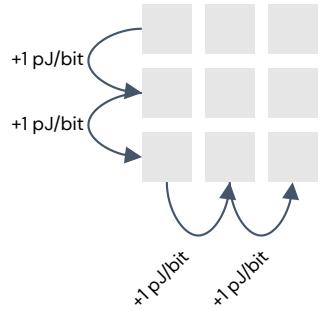

Each chiplet hop adds to communications energy consumption. Building large chiplet arrays this way will incur significant energy costs.

How can we solve these challenges leveraging state-of-the art packaging technologies?

PASSAGE

Wafer-Scale Programmable Optical Interconnect

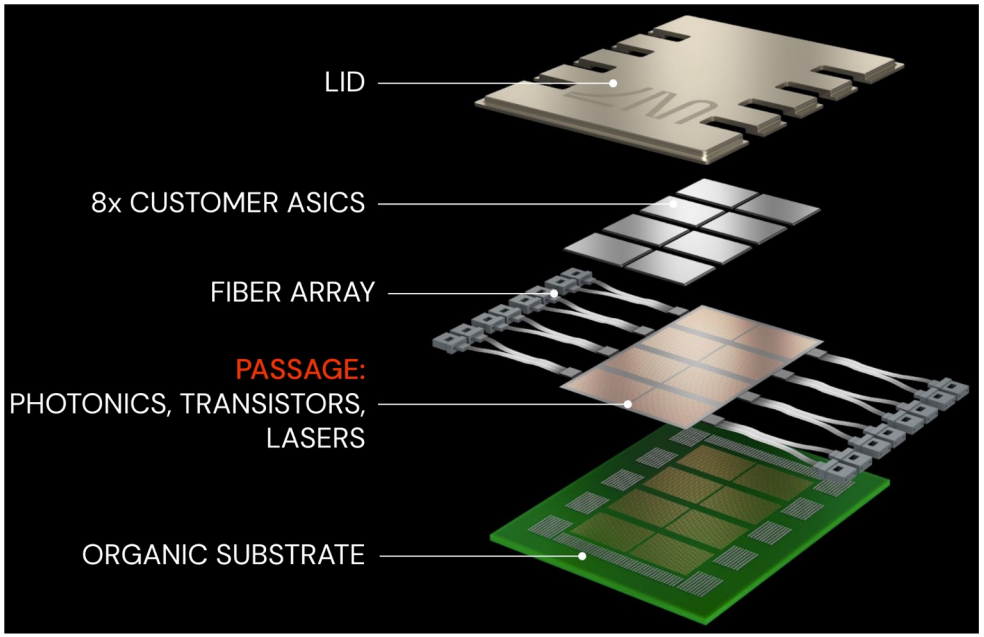

Technology Highlights:

- Enable arrays of heterogeneous chips to communicate optically
- Device Agnostic (Use any market GPU's, CPU's, FPGA's, Neuromorphic, HBM, etc)
- Customize # of nodes
- Protocol agnostic (UCIe, AIB, SerDes, etc)
- Dynamic interconnect topology via 1ms programming (All-to-All, 1D Ring, 2D Hypertoroid, etc)
- Single hop anywhere in 2ns max latency

Cross Section

Chip-on-wafer Packaging

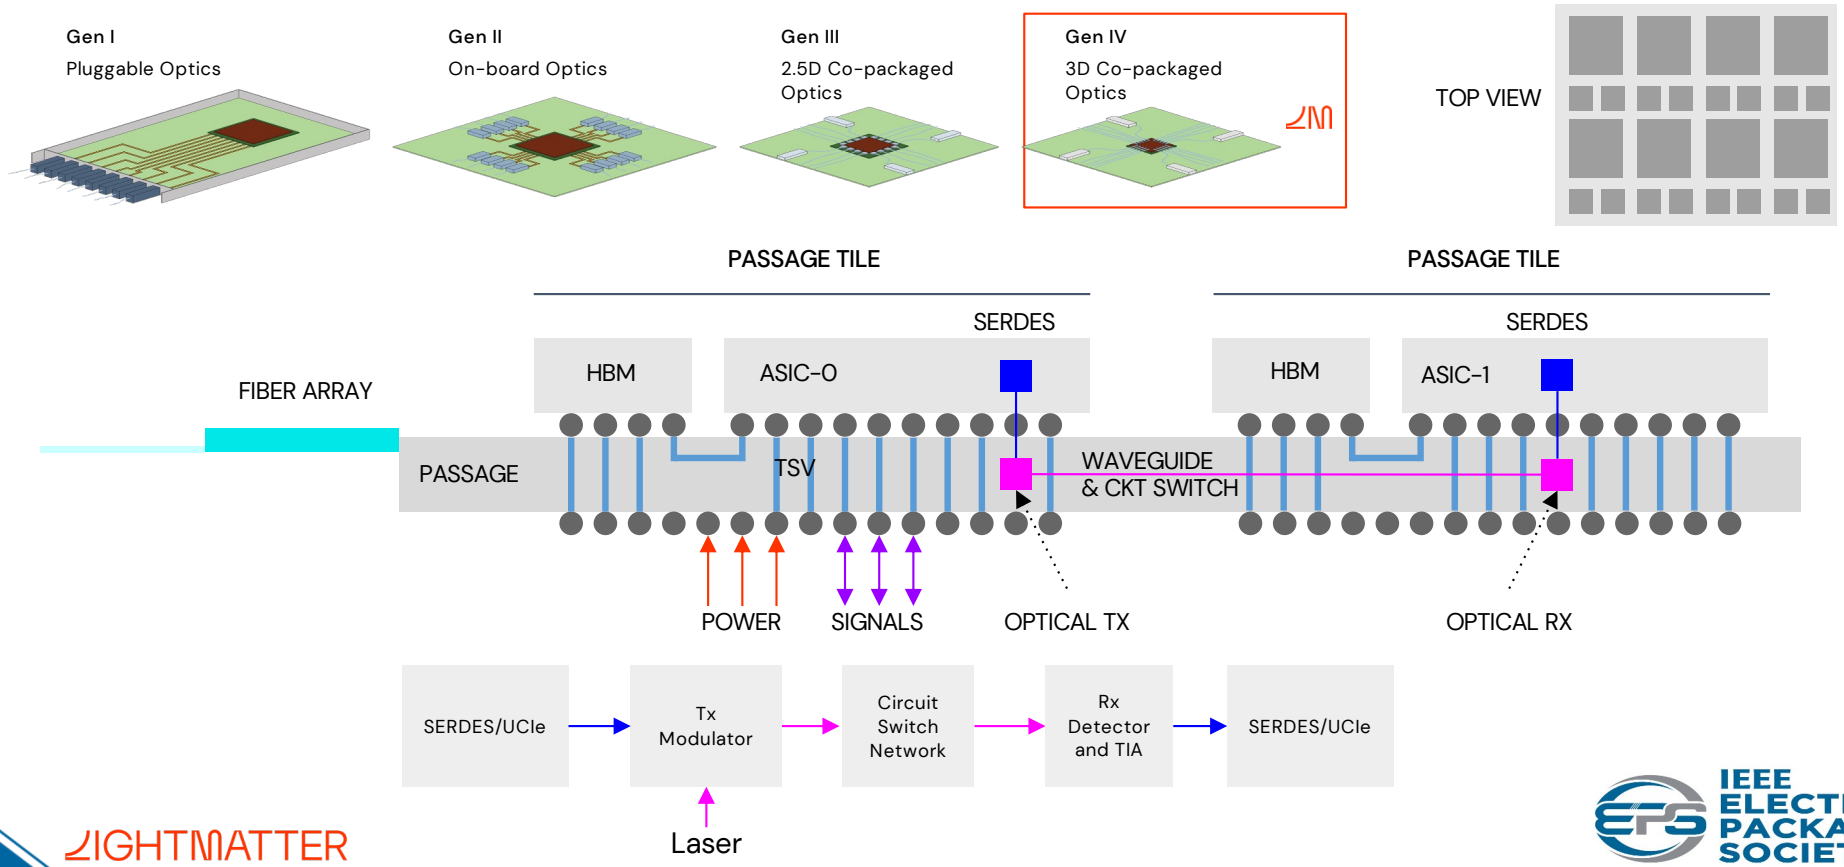

LIGHTMATTER

IEEE ELECTRONICS PACKAGING SOCIETY

Call to Action

Systems and Packaging

- Heterogeneous packaging technologies provide a platform for photonic integration and open opportunities to innovate in:
  - Power delivery
  - Thermal management at package/system level
  - Fiber attach process development to improve throughputs
  - Scaling physical dimensions
- Drive standardization of ecosystem and consolidate supply chain
- Pluggable connectors for fiber attach

Energy efficient, high data rate, low latency, and Scalable SIP enabled by photonics

Ajey Jacob
Director | Application Specific Intelligent Computing Lab (ASIC Lab)
Information Sciences Institute (www.isi.edu)
University of Southern California (www.usc.edu)

The Case for Photonics enabled SIP Platform

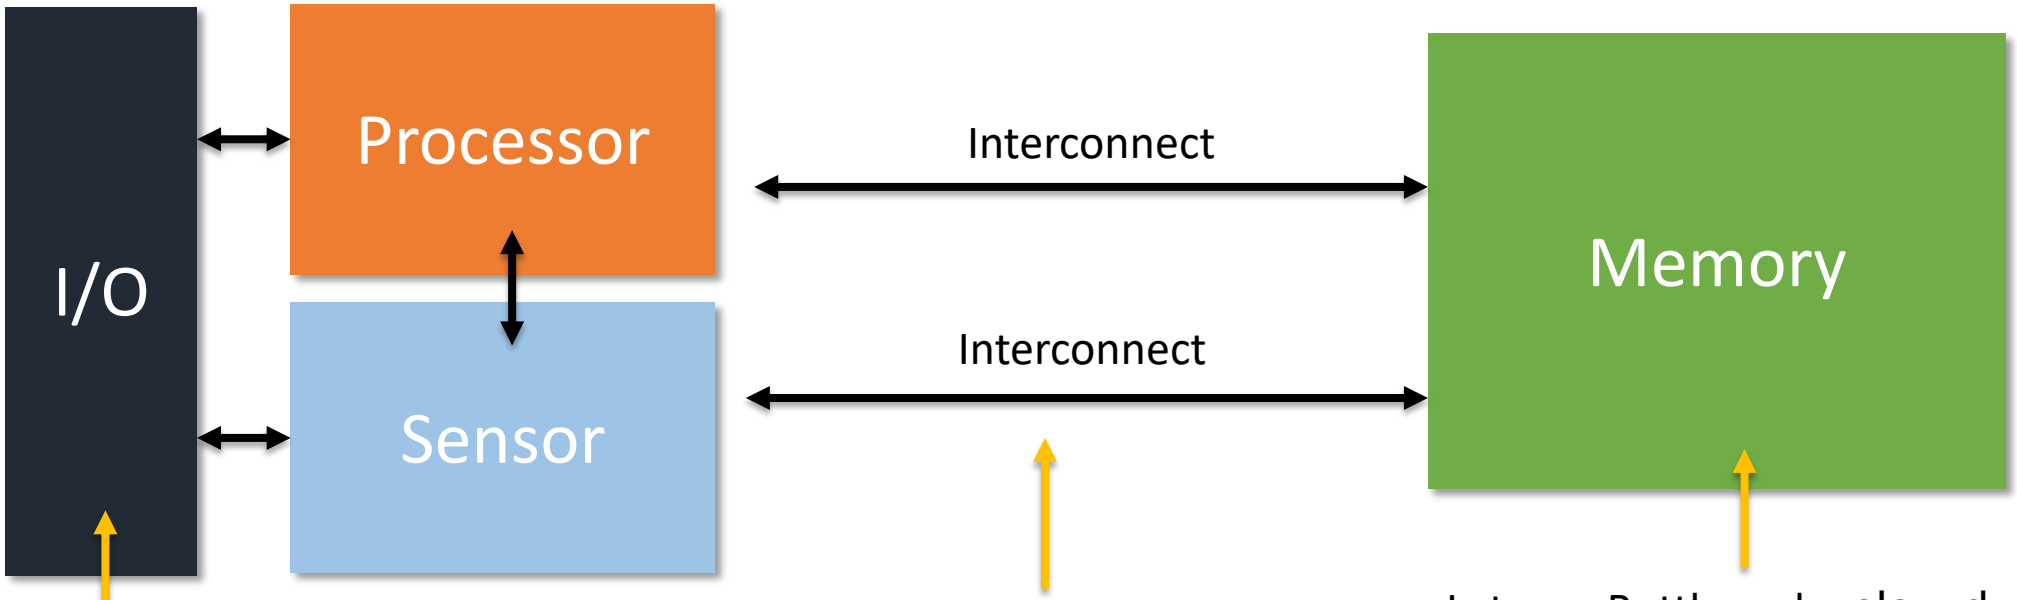

- Limited bandwidth
- High energy consumption
- Area constraints

Solution: Novel optical I/O

- Limited Bandwidth
- High Energy Consumption
- Parasitics

Solution: Optical Interconnects

- Latency Bottleneck - slow due to parasitics (RC delays)
- limited wafer scale integration
- Memory bandwidth/throughput Bottleneck

Solution: Optical-SRAM

Optical Memory Solution

Wafer-scale chip with optical memory and optical interconnects

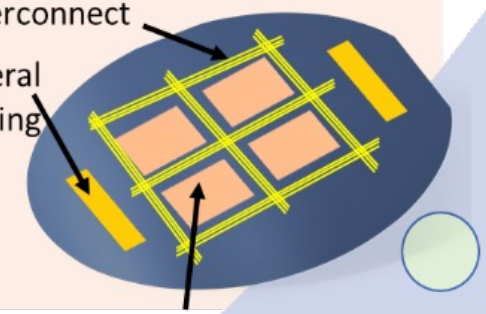

3D stacked wafer chips

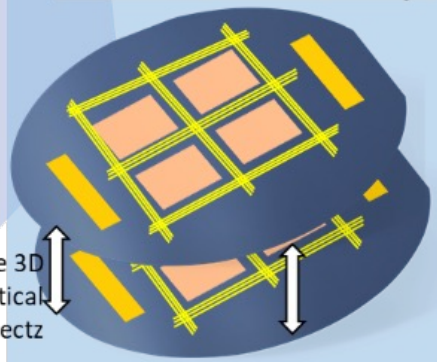

Highlight

Optical free-space vertical interconnects

End Goal

- **Fast, minimal to no memory-I/O bottleneck**
- **Suitable for data intensive applications**

High density magnetic non-volatile optical RAM

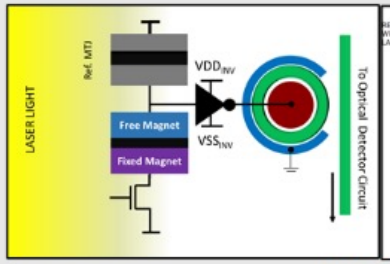

High speed SRAM like optical RAM

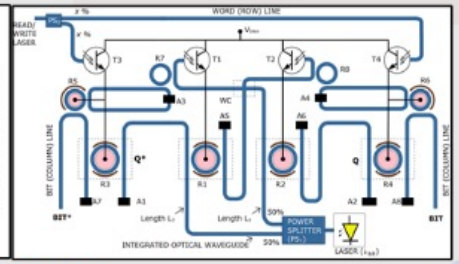

Highlight

Creating novel optical memory hierarchy

Very large-scale optical memory array

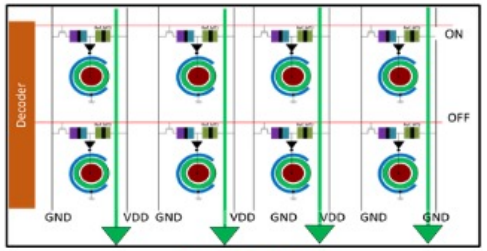

Highlight

Largest and fastest semiconductor memory array

Problem Statement

- **Slow, memory and I/O bottleneck**
- **Unsuitable for data intensive applcaitions like drug discovery and deep learning, etc**

Interconnect Solution

Problem Statement: The integration of chiplets, particularly when they are located at significant distances from each other, poses a challenge in achieving reliable and high-speed communication. An effective solution to address this challenge can be found through the implementation of analog optical interconnects.

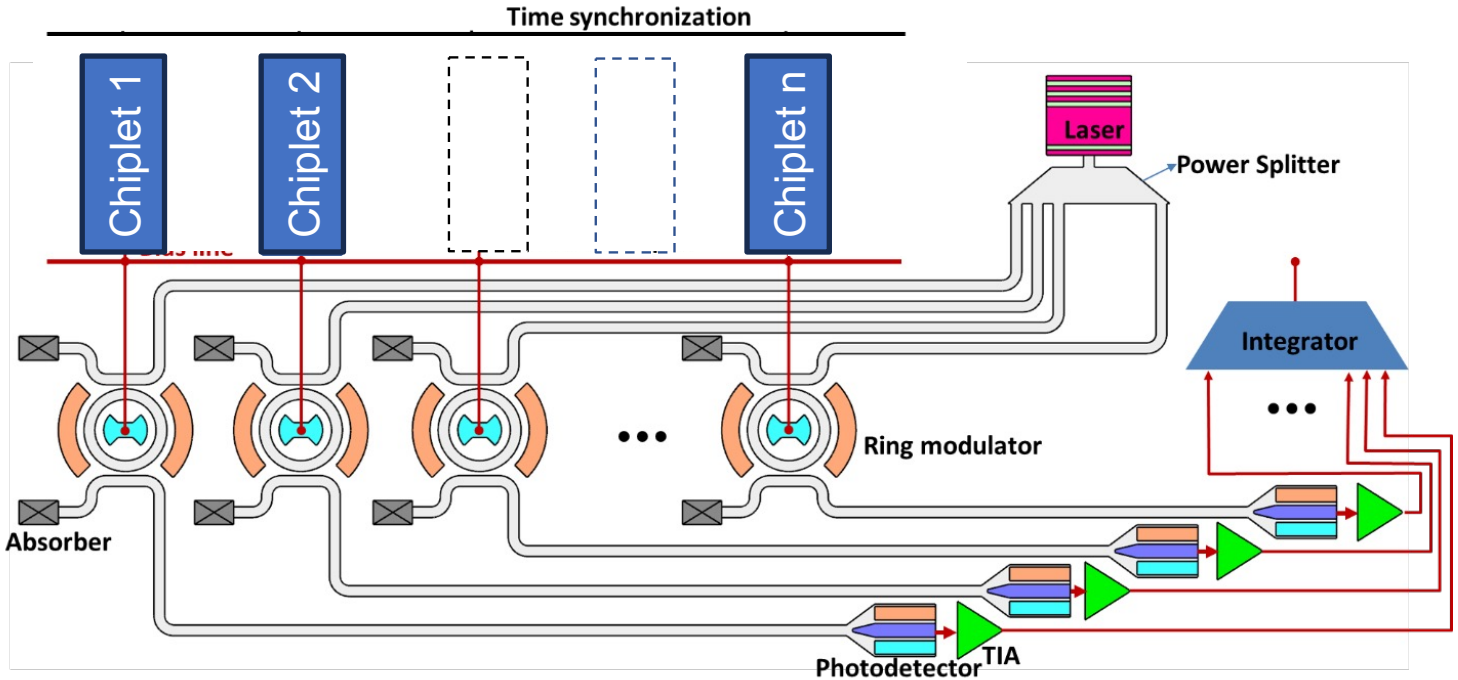

Schematic of an analog optical on-chip interconnect for
(1) Wafer Scale High Performance Computing (HPC)
(2) extreme edge SIP applications

Analog Optical Interconnects provides

- Longer transmission distances
- High Bandwidth
- Low Latency
- Power Efficiency
- Noise immunity
- Scalability and Integration
- Reliability and Robustness
- Hybrid integration with electronics
- Manufacturability and cost

Optical I/O Solution

Problem Statement

– Coupling light from a fiber to a waveguide

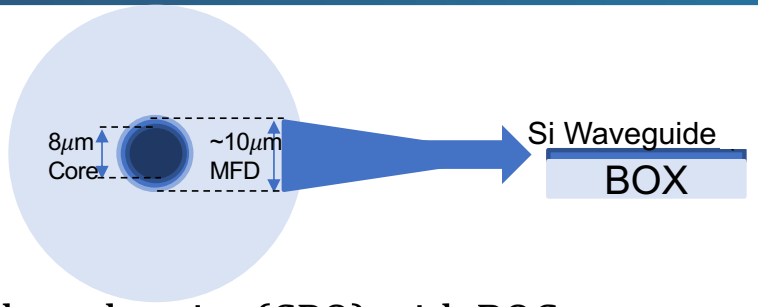

1. Backside Optical Coupler (BOC)

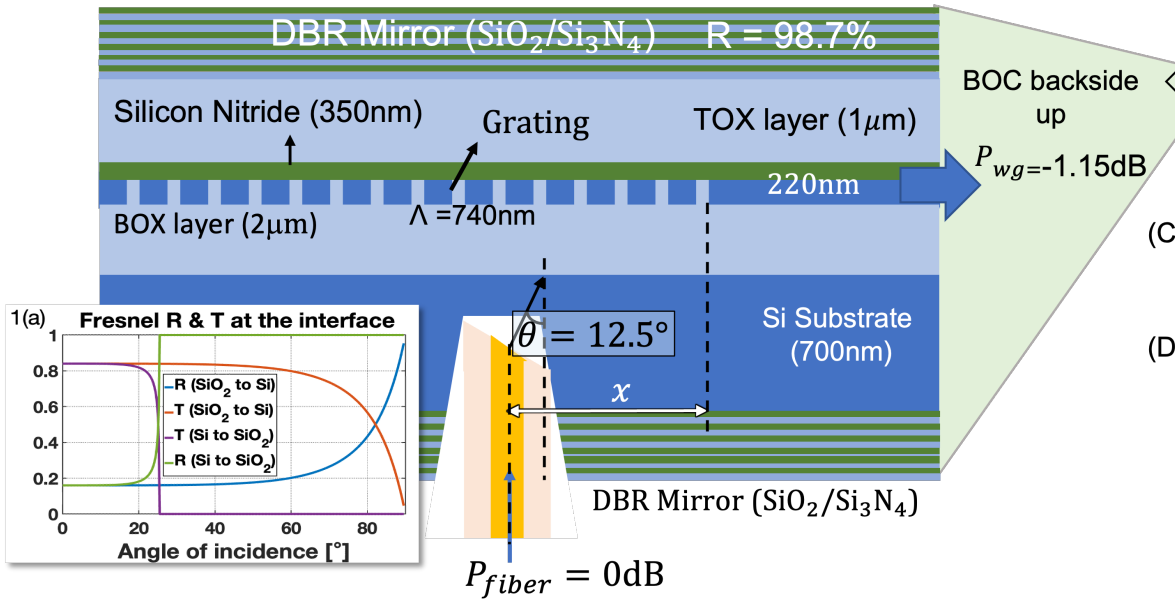

2. Co-Packaged optics (CPO) with BOC

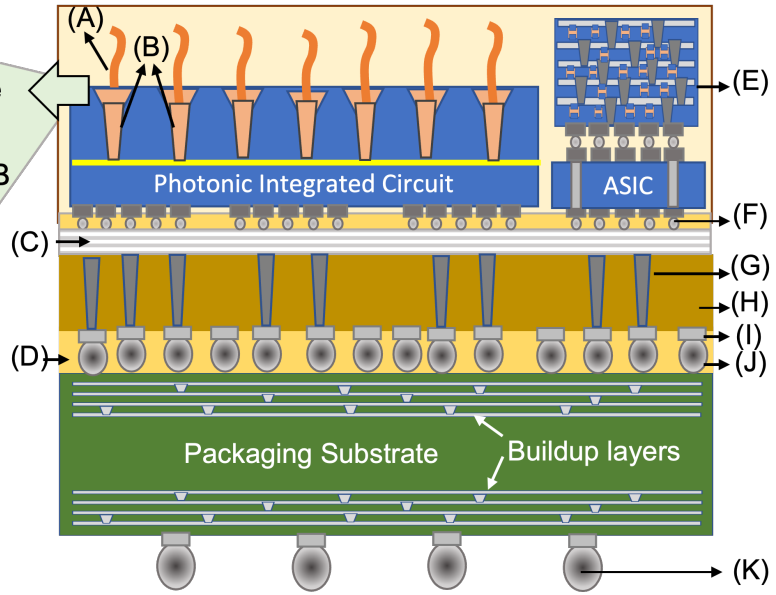

1. Cross-sectional image of a Backside Optical Coupler (BOC) with fiber attached to the V-groove formed on the backside of the wafer with DBR mirrors on the top and bottom surfaces, 1(a) Fresnel reflectance and transmittance values for Si/$Si0_2$ interface as a function of incidence angle showing the need to avoid light incidence on the high contrast interface.

2. Cross-sectional image of a CPO with BOC, (A) Optical I/O (B) V-Groove (C) RDL redistribution and communication layers, (D) underfill, (E) Memory, (F)Copper bumps, (G)Through Silicon Via (TSV), (H) Silicon Photonics Interposer, (I) Ultra Bump Metallization, (J) Solder Bumps, (K) Solder balls

Transparency

A Point of View from the Substrates Industry about the readiness level for the photonics packaging integration

IEEE ELECTRONICS PACKAGING SOCIETY

By: Stefano Oggioni, Technology Manager
AT&S, AUSTRIA

IEEE

Substrates Research Focuses

Packaging allows integration of functions for System performance

Miniaturization
Sustain the increasing demand for computational power, fast data processing, wiring density, Chip to Chip communication bandwidth,

Modularization
More functionality at same or reduced space

Increased speed / Low latency
Communication of high data volumes (5G, Autonomous Driving)

Increased power / Power efficiency
Reducing impacts of the package/substrates to the electrical performance

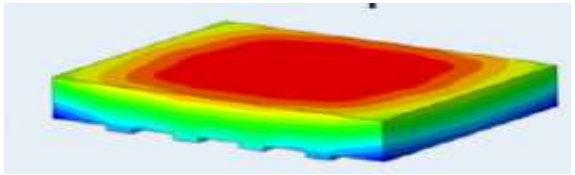

Thermal management/ Cooling efficiency
Adapting different or innovative and more efficient Cooling Solutions

IEEE ELECTRONICS PACKAGING SOCIETY

IEEE

Future Trend for Advanced Computing

Substrate (organic/inorganic) may enable innovative solutions

- Node scaling alone does not provide desired performance

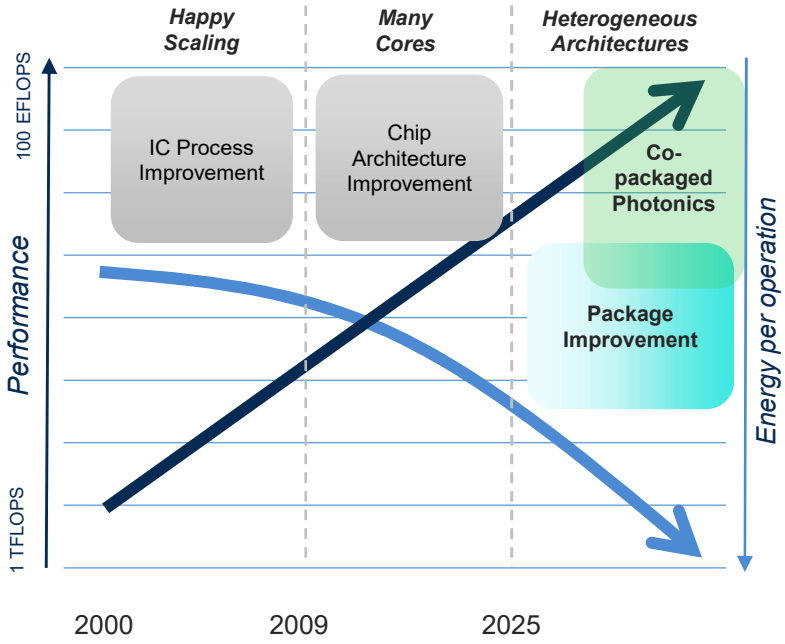

Advanced Substrates & Packaging

New Challenging Functions in the Substrate

- Co-packaged optics
- Larger Body Sizes (Chiplets, HIR)
- CTE Mismatch (Chiplets, Interposer, Photonics)
- D2D Connection (BW, low latency)
- Power Delivery (PDN/Voltage conversion)
- Thermal Dissipation (Efficient Cooling)

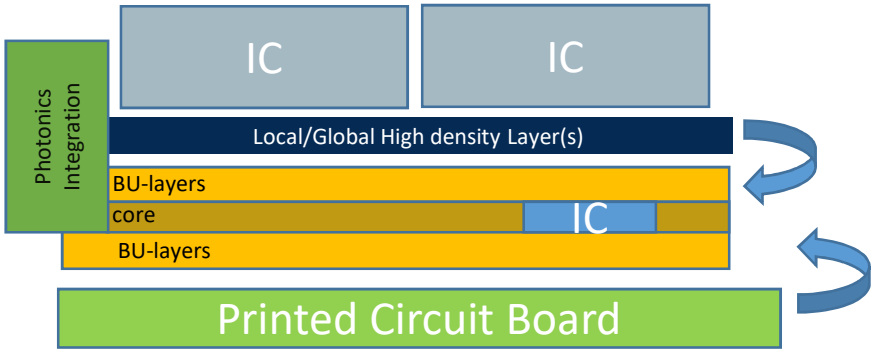

Which functions can/should be optimized & where in the Package?

Many Options for a Single Strategy

Not a Winner yet, where to invest to be the right partner

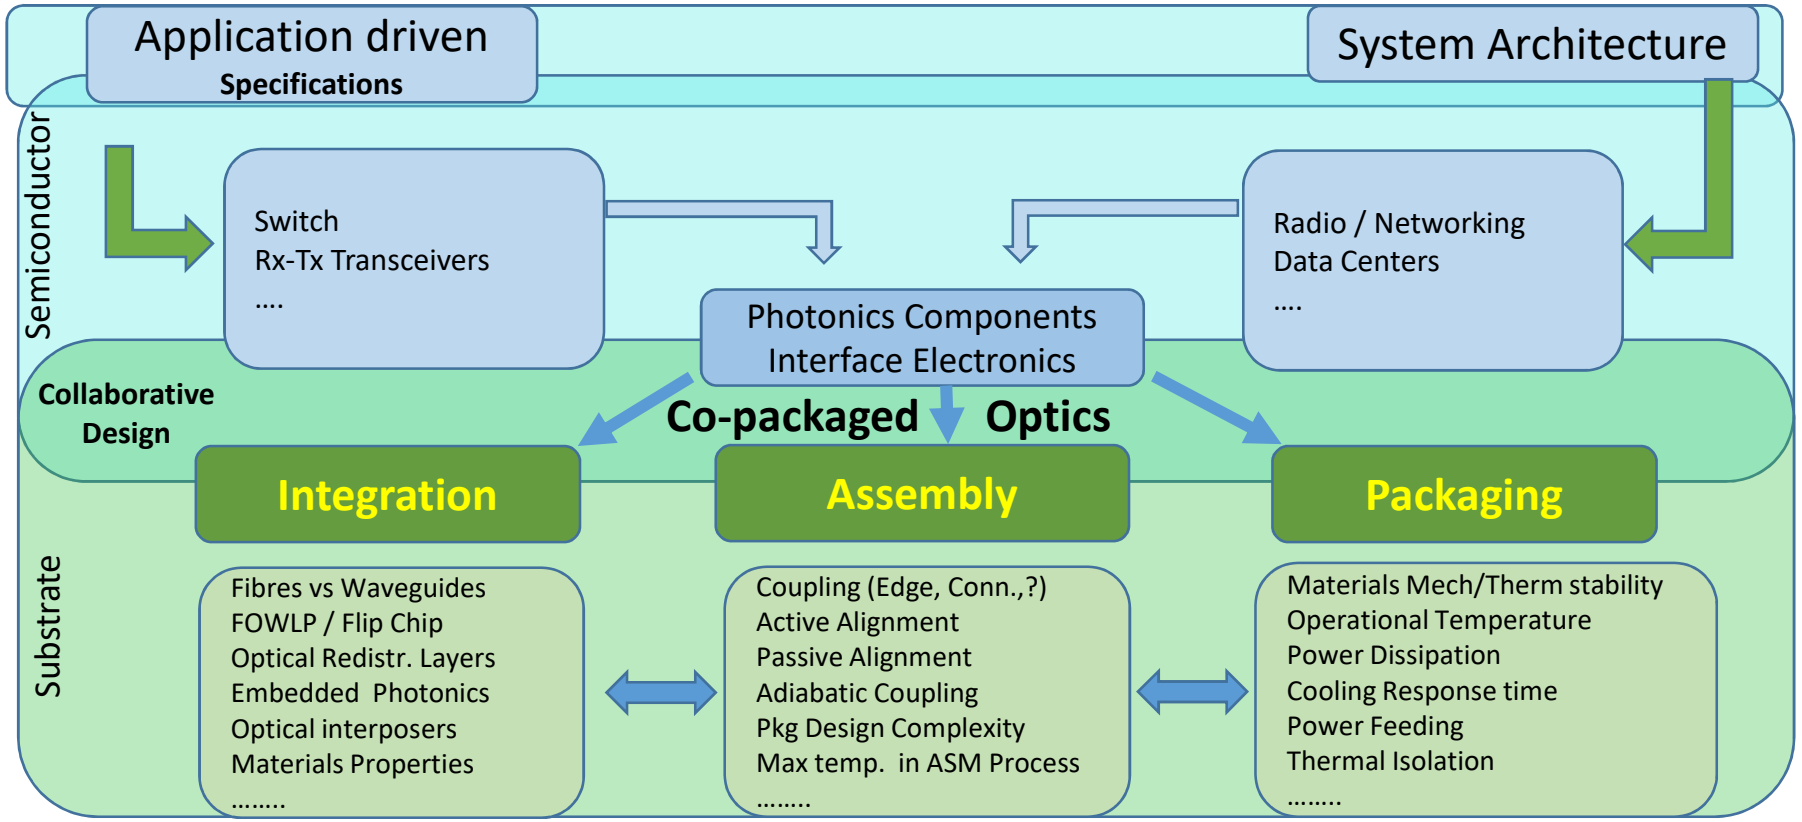

IEEE ELECTRONICS PACKAGING SOCIETY

IEEE

WiP Optical Interconnect Solutions

Optical Paths Integration in Substrates/Packages

Optical Advantages:

- Low-loss transmission for high speed signals
- Capable to provide large bandwidths (THz range)
- Almost no crosstalk between channels

Optical RDL

High Performance Computer, Server, Client, ASIC

Integrated optical waveguides

Glass waveguides

Co-Packaging Optics using glass waveguide interface

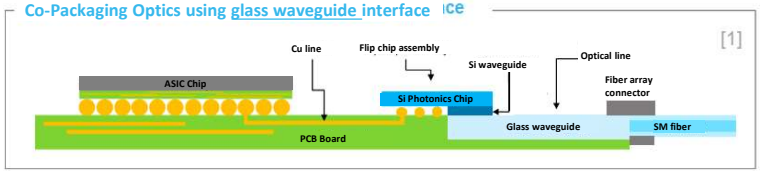

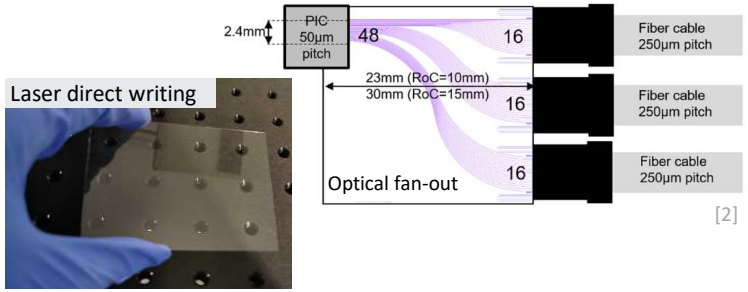

Optical RDL: Increase optical IO density from 250 µm pitch (fiber array) to 50 µm pitch (PIC)

Propagation loss: 0.08 dB/cm @1310 nm
0.14 dB/cm @1550 nm

Polymer waveguides

Co-Packaging Optics using polymer waveguide interface

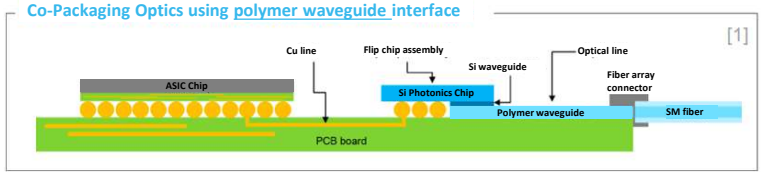

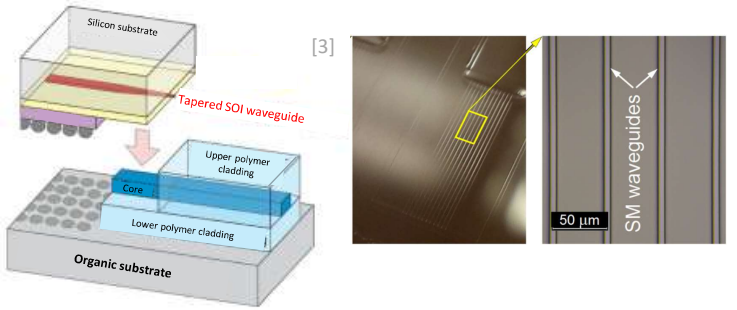

Evanescent coupling: SOI waveguide direct contact with polymer waveguide

Propagation loss: 0.2 dB/cm @1310 nm
0.6 dB/cm @1550 nm [4]

[1] adapted from Yole, Co-packaged optics, Yole Development, 2022.
[2] L. Brusberg et al., "Glass Substrate With Integrated Waveguides for Surface Mount Photonic Packaging," J. Lightwave Technol. 39, 912-919 (2021).
[3] R. Dangel et al., "Polymer waveguides for electro-optical integration in data centers and high-performance computers," Opt. Express 23, 4736-4750 (2015).
[4] https://www.microresist.de/wp-content/uploads/2020/02/PI_waveguides_materials_2015-1.pdf

Readiness for Co-packaged Optics

CPO may come in steps but with collateral requirements in packaging

- Building an R&D Network for substrate/packaging/module solutions
- Photonics Packaging will materialize soon, but it will requires a holistic approach to the Packaging

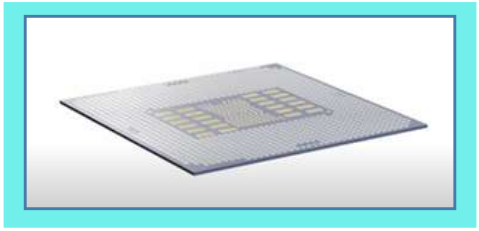

Advanced IC Substrates

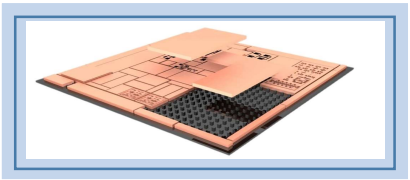

System Understanding

- Understand requirements of future Heterogeneous Integration (e.g. SIP, Chiplets)
- Focus on organic interconnect solutions

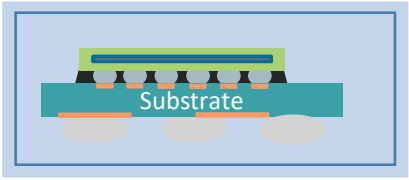

Die mounting (chip last)

- Hybrid bonding on IC Substrate
- Developing capabilities

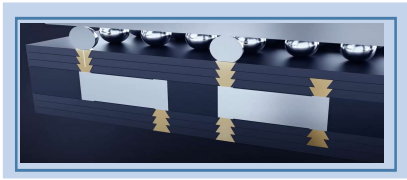

Multilayer / functional cores

- Embedding components
- System integration
- Test and verification

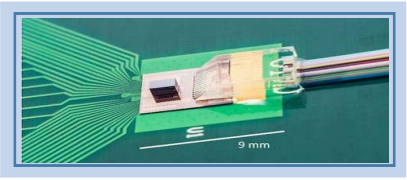

Silicon photonics

- Bridging optics to substrates
- Functional integration
- Test and verification

Dedicated Research Center for Advanced Substrates and Electronic Packaging

Summary

AT&S

Transparency

- Co-packaged Optics implementation require an extended cooperation across System Architecture & Packaging Design
- Miniaturization, Modularization, Hybrid Comps Integration, Power Delivery & Thermal Management remain a valid contributors from the supporting Industry Partners
- To achieve the earliest implementation it requires the contribution from the all players in the Supply Chain
  - From Optical Devices to Materials, Design methodologies, Simulation SW, Substrates construction, ASM Processes, Test methods, …
- The Supply Chain Readiness can be construed only with collaborations in the early phases of projects
- The Result should be “Transparent” to Systems performances

Disclaimer

This presentation is provided by AT & S Austria Technologie & Systemtechnik Aktiengesellschaft, having its headquarter at Fabriksgasse 13, 8700 Leoben, Austria, or one of its affiliated companies (“AT&S”), and the contents are proprietary to AT&S and for information only.
AT&S does not provide any representations or warranties with regard to this presentation or for the correctness and completeness of the statements contained therein, and no reliance may be placed for any purpose whatsoever on the information contained in this presentation, which has not been independently verified. You are expressly cautioned not to place undue reliance on this information.

This presentation may contain forward-looking statements which were made on the basis of the information available at the time of preparation and on management‘s expectations and assumptions. However, such statements are by their very nature subject to known and unknown risks and uncertainties. As a result, actual developments, results, performance or events may vary significantly from the statements contained explicitly or implicitly herein.
Neither AT&S, nor any affiliated company, or any of their directors, officers, employees, advisors or agents accept any responsibility or liability (for negligence or otherwise) for any loss whatsoever out of the use of or otherwise in connection with this presentation. AT&S undertakes no obligation to update or revise any forward-looking statements, whether as a result of changed assumptions or expectations, new information or future events.
This presentation does not constitute a recommendation, an offer or invitation, or solicitation of an offer, to subscribe for or purchase any securities, and neither this presentation nor anything contained herein shall form the basis of any contract or commitment whatsoever. This presentation does not constitute any financial analysis or financial research and may not be construed to be or form part of a prospectus. This presentation is not directed at, or intended for distribution to or use by, any person or entity that is a citizen or resident or located in any locality, state, country or other jurisdiction where such distribution, publication, availability or use would be contrary to law or regulation or which would require any registration or licensing within such jurisdiction.

IEEE