Extending Moore's Law with Integrated Photonics (and Packaging)

Dr. Nicholas Harris
CEO, Lightmatter

19TH INTERNATIONAL CONFERENCE & EXHIBITION ON DEVICE PACKAGING
FOUNTAIN HILLS, AZ · WWW.DEVICEPACKAGING.ORG · MARCH 13-16, 2023

1
TEAM

2
OFFICES

Boston

Mountain View

3
PRODUCTS

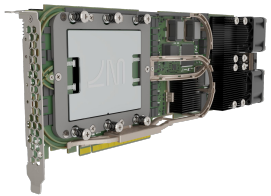

Envise

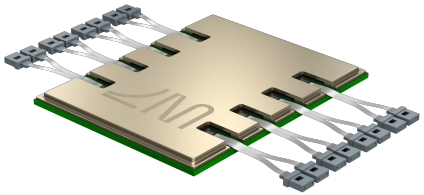

Passage

Idiom

105
EMPLOYEES

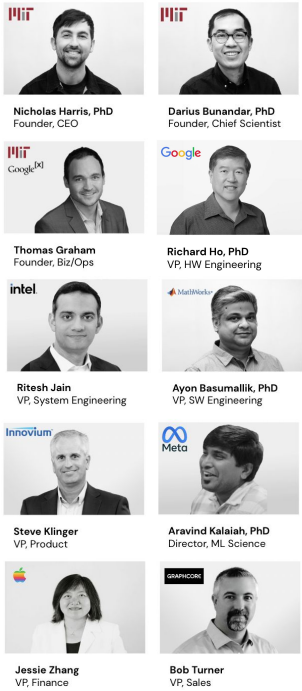

115
PATENTS

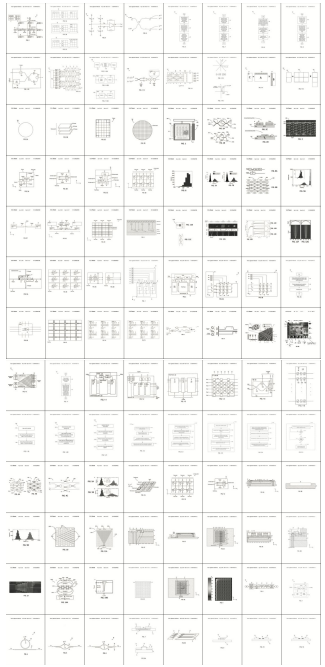

Investors

GV (Google Ventures)
Spark Capital
Matrix Partners
Viking Global
Hewlett-Packard Enterprise
Massachusetts Institute of Technology
Stanford University

Industry-Wide Challenges

Challenges

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

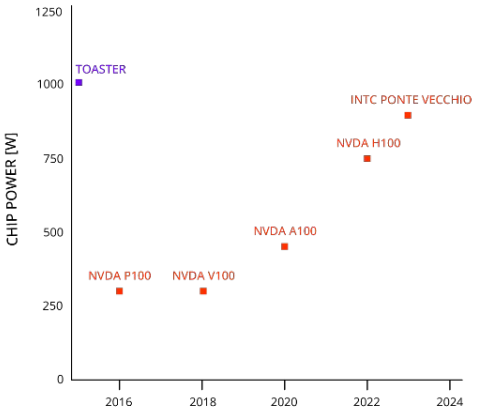

Notes
[1]Dennard Scaling article.
[2]Publicly available data.

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.

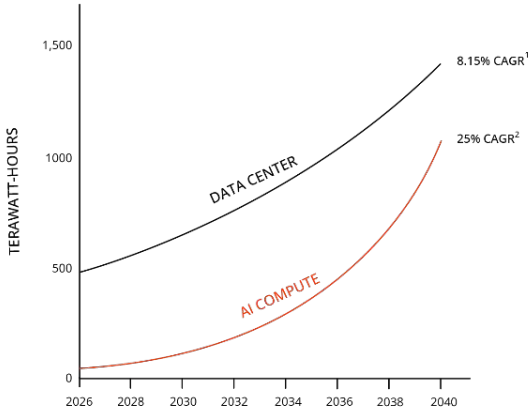

Notes
[1] Interpolation based on IEA. There are estimates from 6% to 35% annual growth rate.
[2] Averaged CAGR over 2026-2040 smoothed down from 35% growth rate from 2026 - 2030. Source: IEA
[3] Actual data center growth rate is not agreed upon in academia. Source.

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

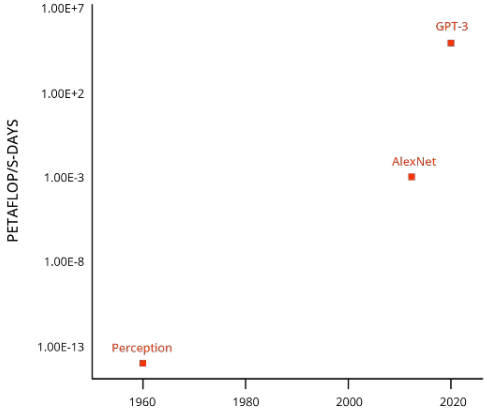

Notes
OpenAI "AI and Compute" article.

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.

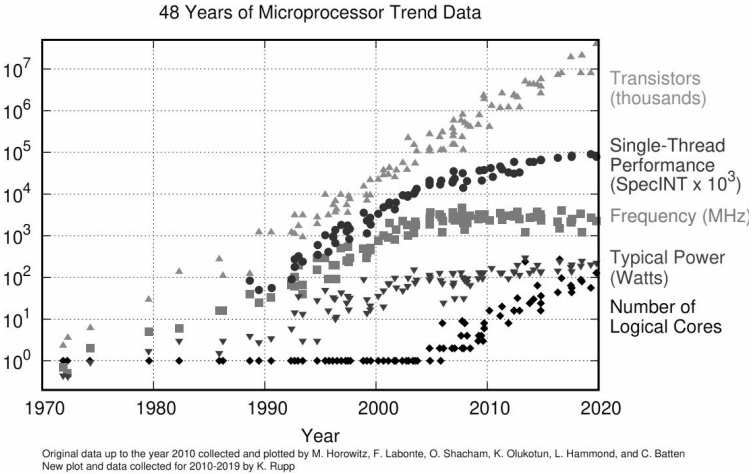

Chip design cost is exploding. Exponential growth in development costs for creating next-generation processors.

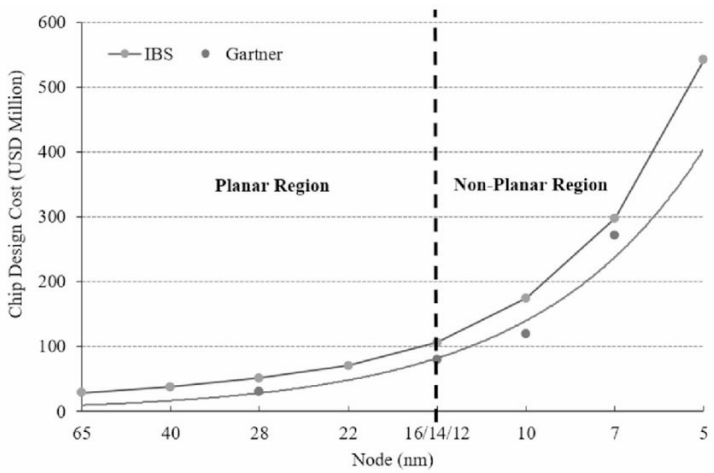

Heterogeneous Integration

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

MONOLITHIC SILICON

CHIPLETS

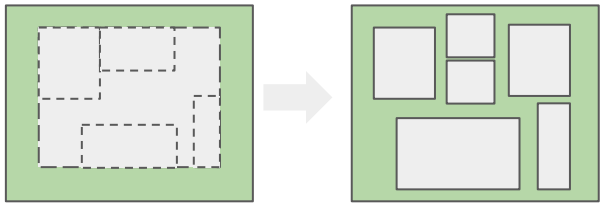

Driven by

- 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

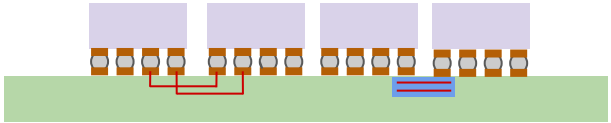

FCBGA, EMIB, RDL-BASED

3D D2D, PRIMARILY VERTICAL

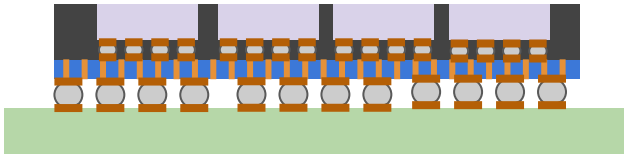

FOVEROS, HBI/V-CACHE

Chiplet ≠ Monolithic

Bandwidth Requirements

BISECTION BANDWIDTH

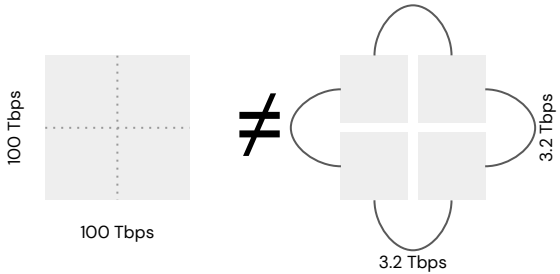

Beachfront and bandwidth are fundamentally linked in chiplet processors. Big chips are at odds with high yield.

MORE HOPS, MORE ENERGY

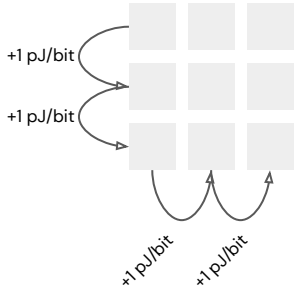

Each chiplet hop adds to communications energy consumption. Building large chiplet arrays this way will incur significant energy costs.

SIGNALING SCHEMES

DARPA The Serial vs. Parallel Question

Source: Presentation by Andreas Olofsson @ DARPA 2.5D/3D Workshop, Dec 6, 2018

	CHIPS Parallel		Serial
IP Complexity	Flip-flop + tristate		SERDES
IP Cost	Low		High (open source?)
Throughput	1Tbps/mm		1Tbps/mm?
Latency	<5ns		High
Energy Efficiency (<1000um)	0.1pJ/bit	Physics	1-10pJ/bit
Throughput per pin	2Gbps		30Gbps
Packaging Complexity	High	Economics?	Low
Packaging Cost	High		Low

Silicon Photonics

300mm CMOS Fab

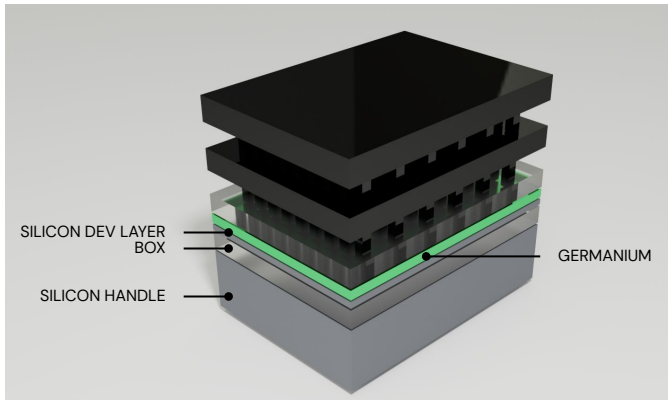

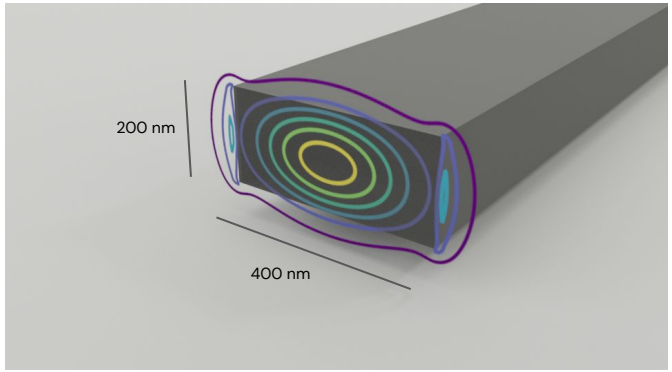

A Sense of Scale

Optical fibers versus nanophotonic waveguides

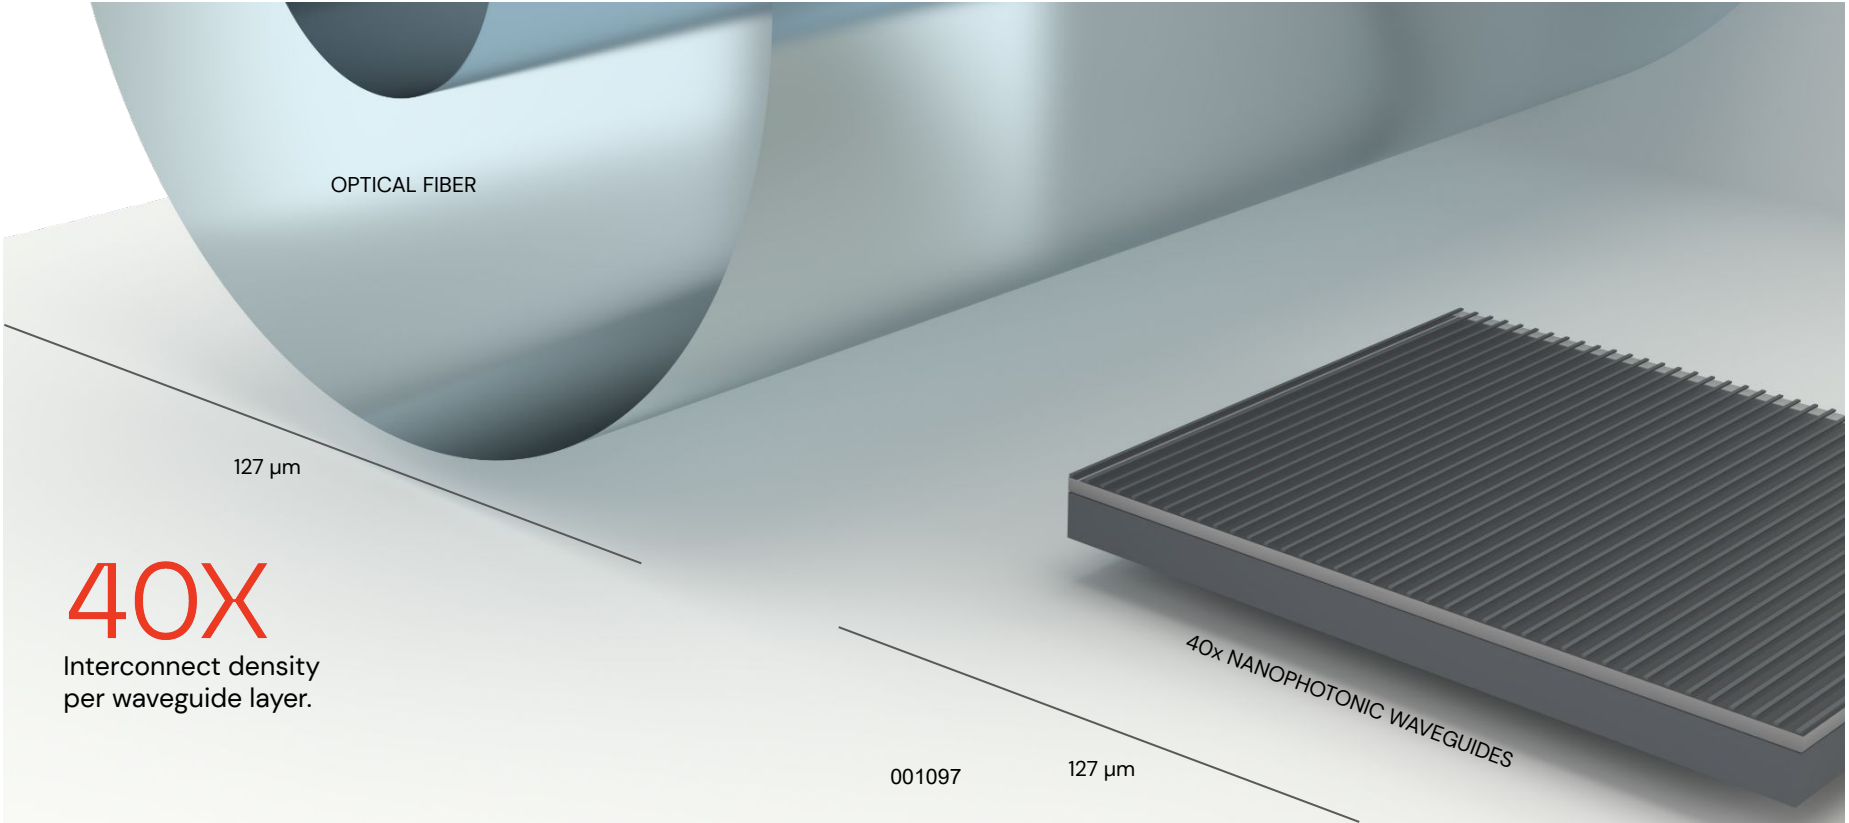

What does silicon photonics look like?

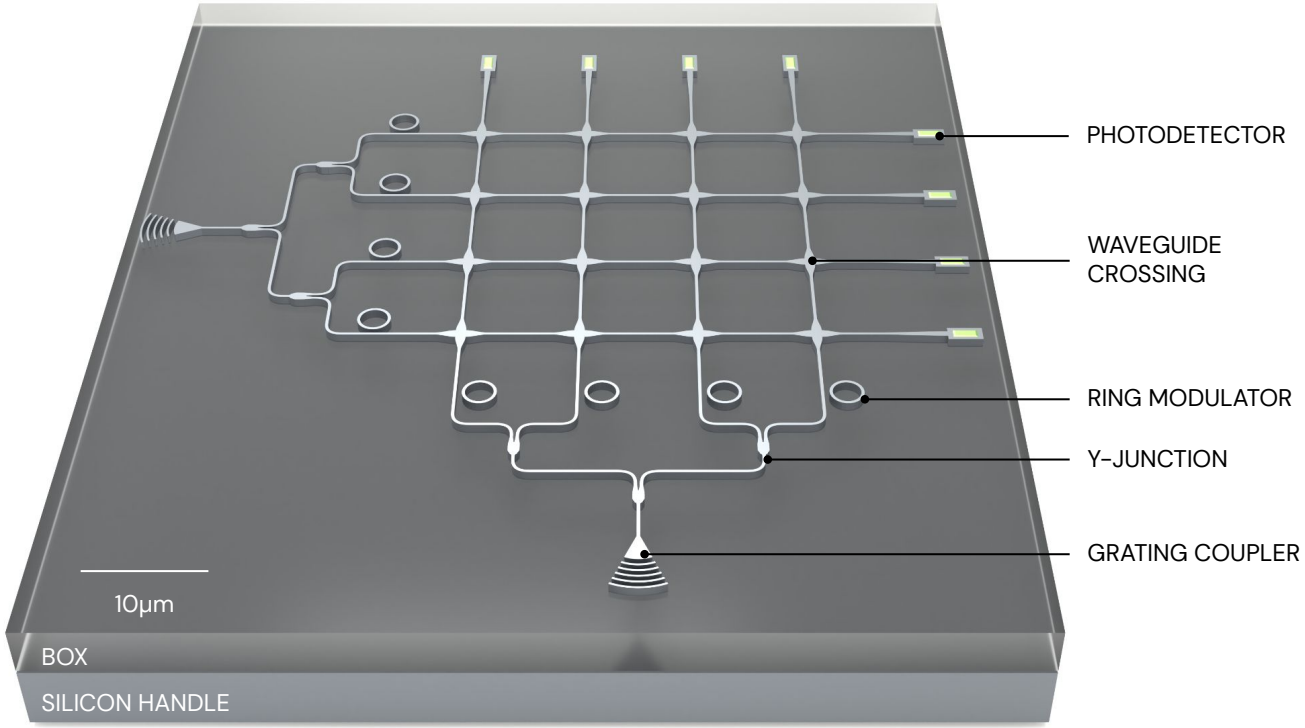

Enabling new communications technologies

Gen I
Pluggable Optics

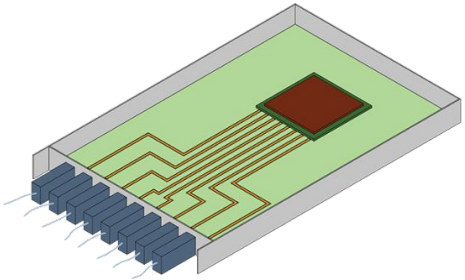

Gen II
On-board Optics

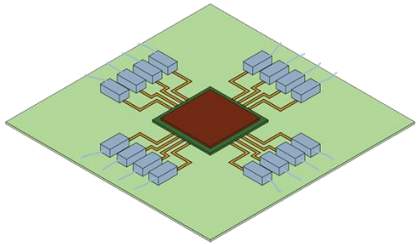

Gen III
2.5D Co-packaged Optics

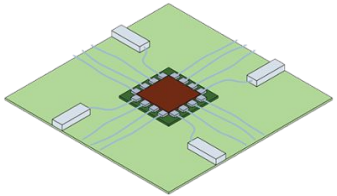

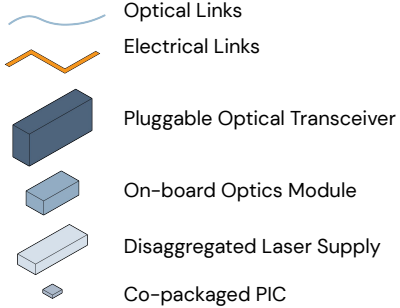

Why bring optics closer to the SoC?

- PCB trace losses too high for long, high-speed links
- Data closer to compute
- Disaggregation
- SoC die sizes max'd out, can't accomodate large #'s of LR SerDes
- Reduction in interconnect power
- Increased bandwidth over pluggables
- Higher bandwidth density at die edge
- Bypass intermediate stages (PCIe)

*Used with permission from Dr. John Bowers and adapted from an article in Applied Physics Letters (2021)

Challenges with Silicon Photonics

Optical Fiber Attach

Expensive, low throughput

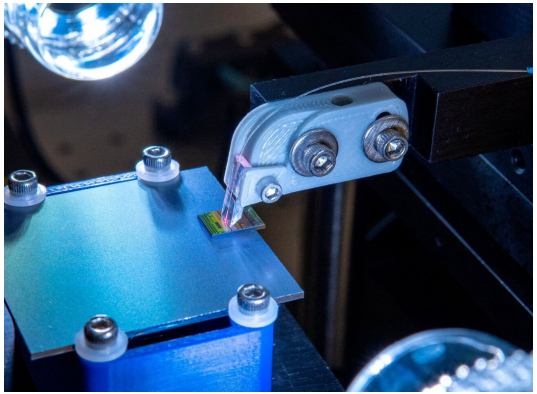

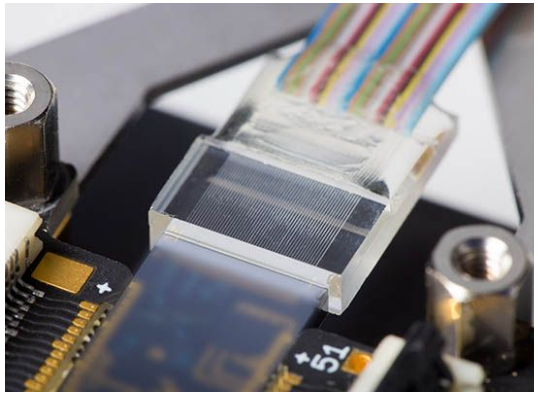

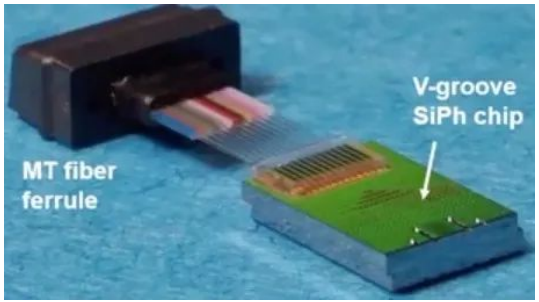

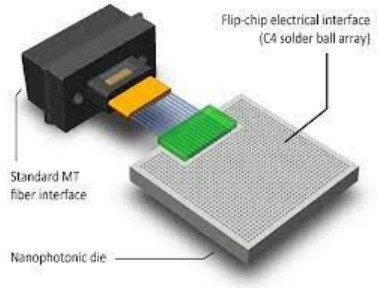

Optical Fiber Beachfront Density

Fibers are massive

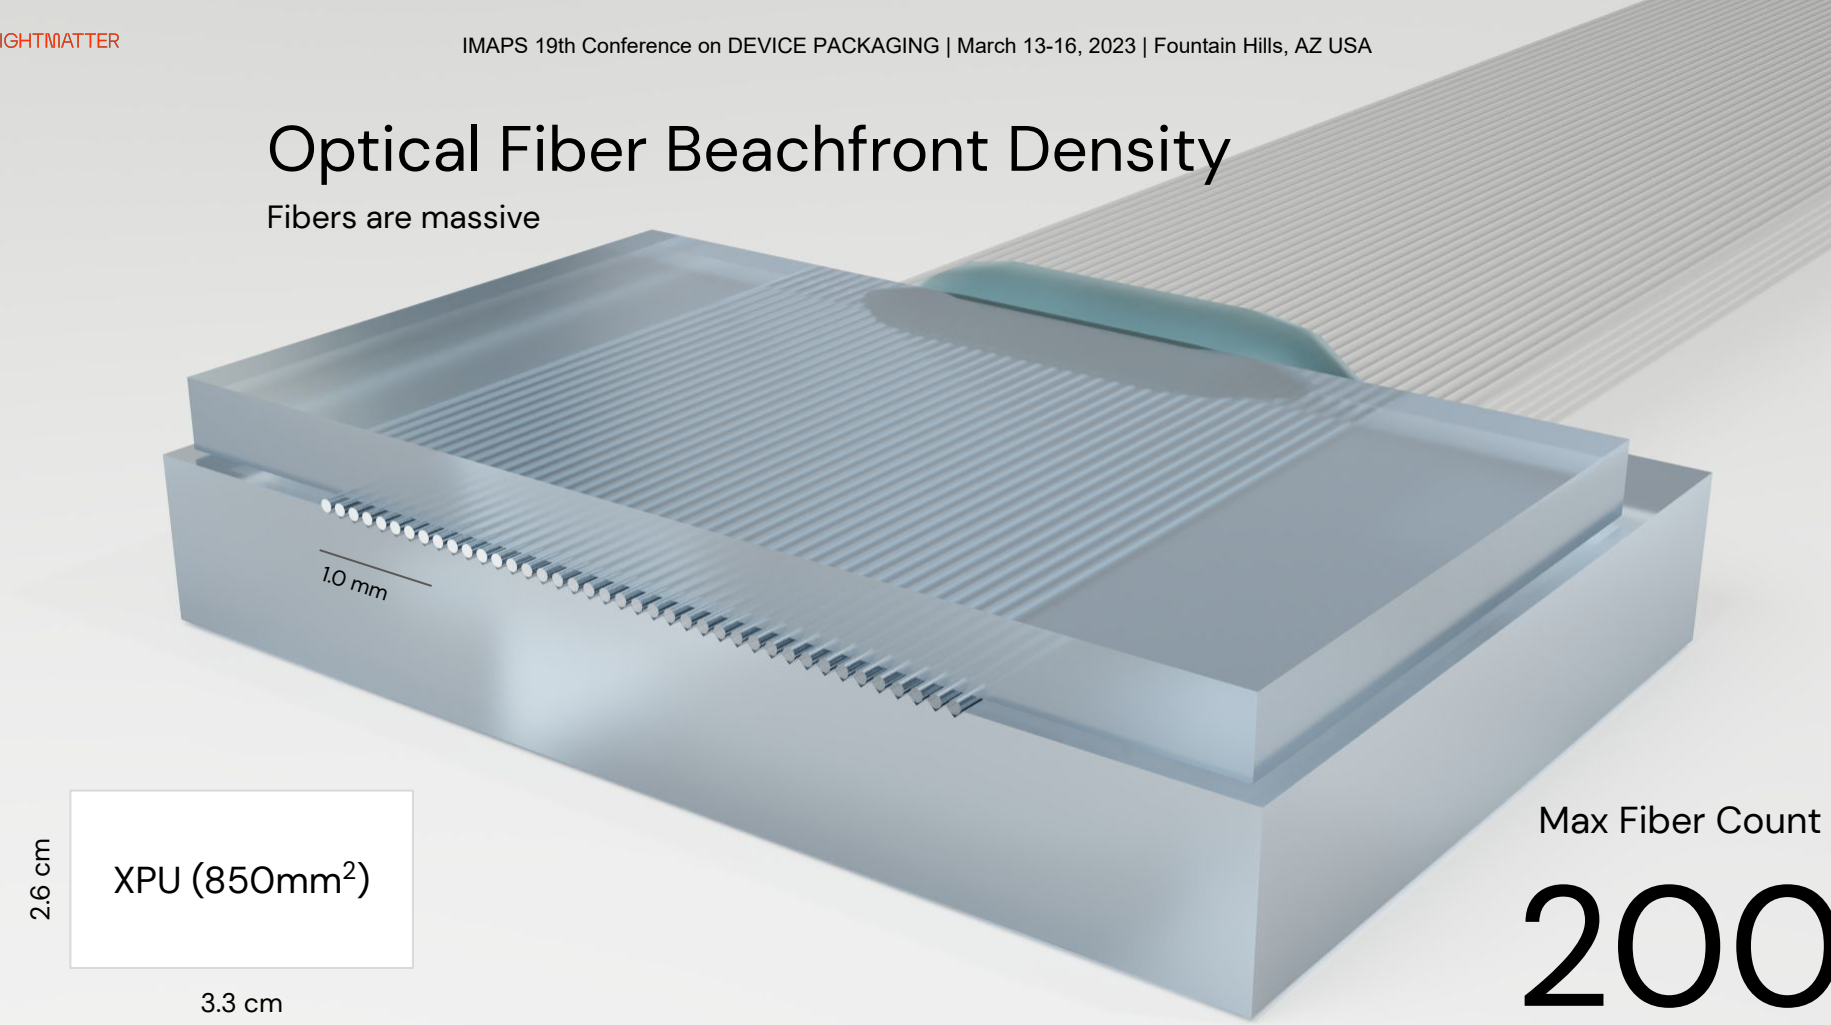

2.6 cm

XPU (850mm^2)

3.3 cm

Max Fiber Count

200

Chiplets and Co-packaged Optics

Scaling challenges

CHIPLET BISECTION BANDWIDTH

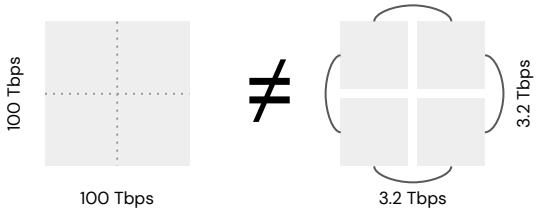

Beachfront and bandwidth are fundamentally linked in chiplet processors. Big chips are at odds with high yield.

MORE HOPS, MORE ENERGY

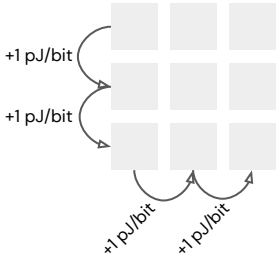

Each chiplet hop adds to communications energy consumption. Building large chiplet arrays this way will incur significant energy costs.

CHIPLET XPU & CPO

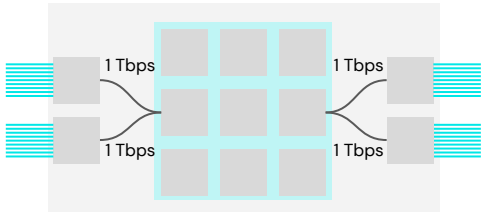

- Fibers have low beachfront density
- More wavelengths, more BW
- Static interconnect

System Level

Serviceability, manufacturability, yield

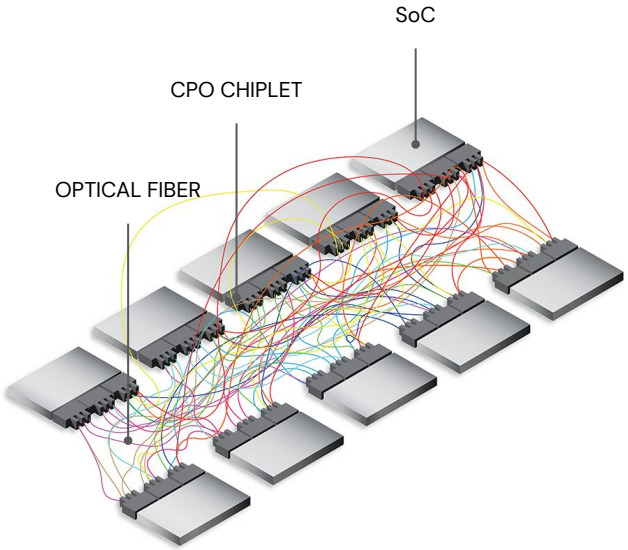

CO-PACKAGED OPTICS ALL-ALL

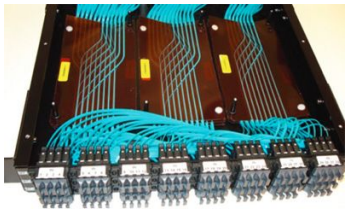

How can we solve these challenges leveraging state-of-the art packaging technologies?

PASSAGE

PASSAGE™

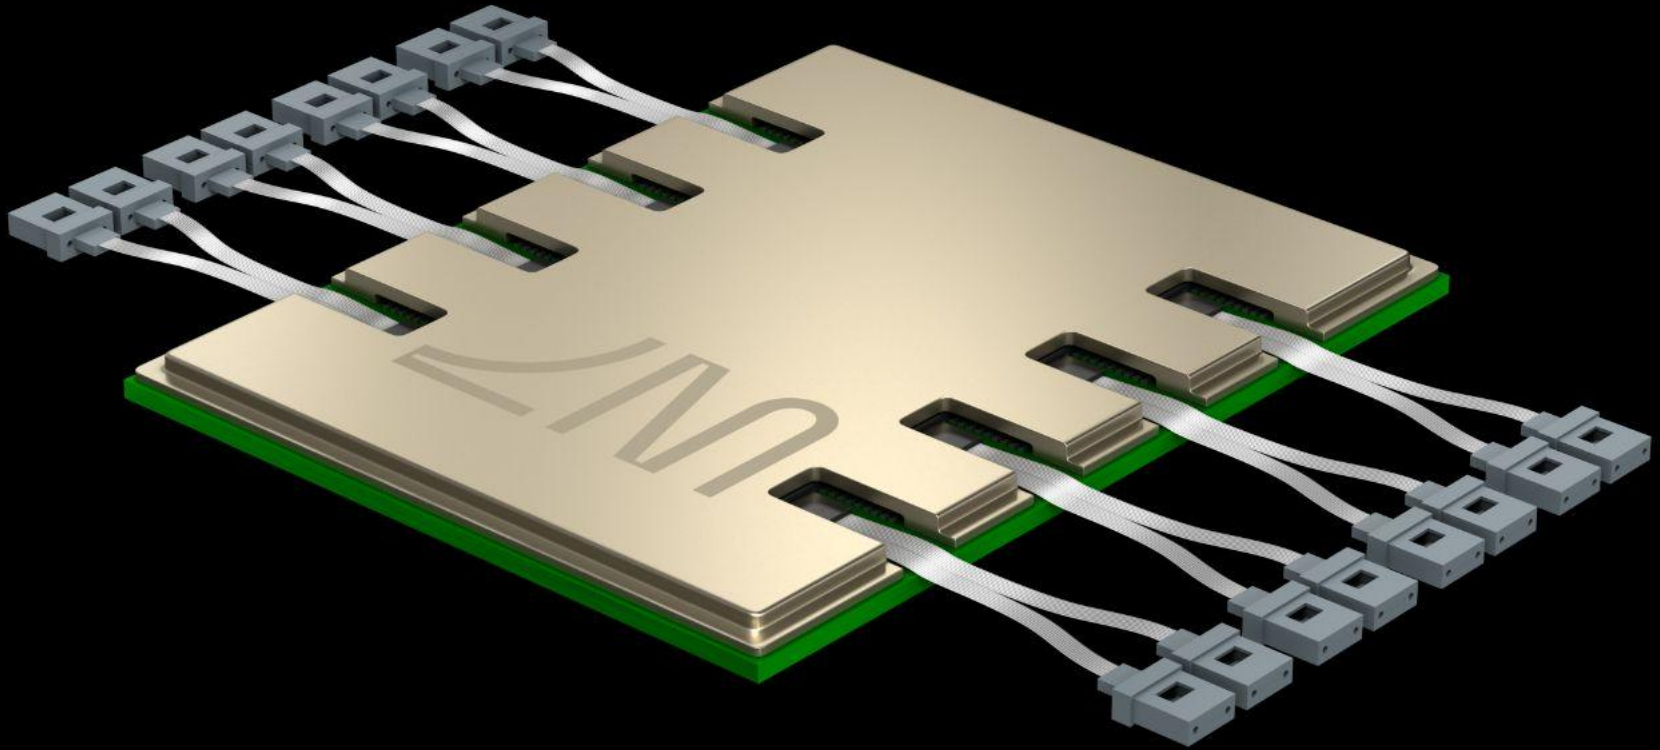

LIGHTMATTER

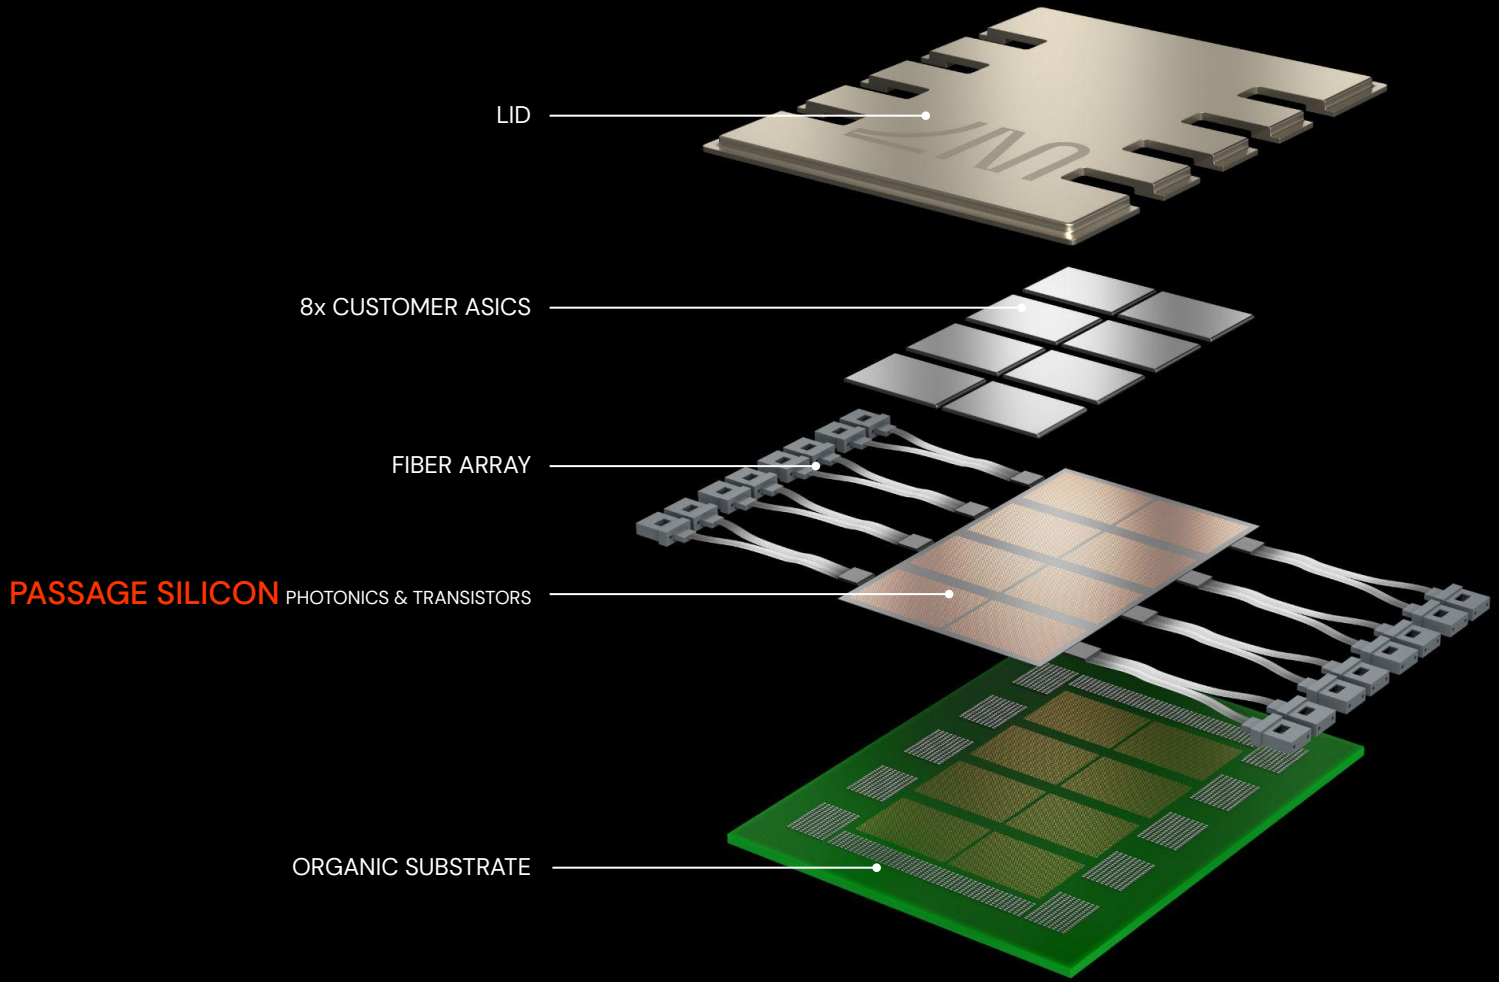

Let's put it into context

Lightmatter, the next generation

Gen I
Pluggable Optics

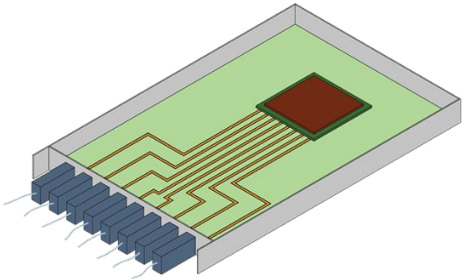

Gen II
On-board Optics

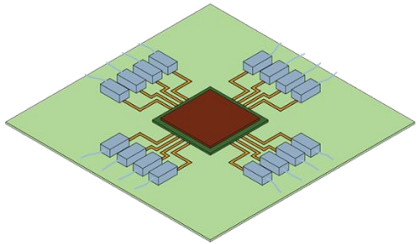

Gen III
2.5D Co-packaged Optics

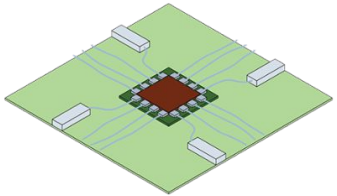

Gen IV
3D Co-packaged Optics

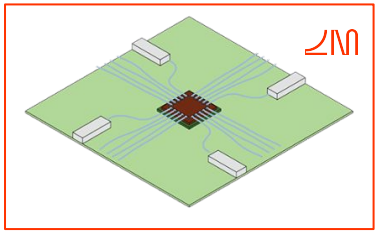

*Used with permission from Dr. John Bowers and adapted from an article in Applied Physics Letters (2021)

Cross Section

Chip-on-wafer Packaging

TOP VIEW

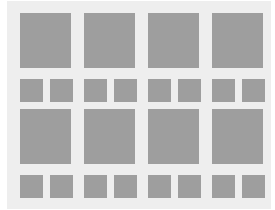

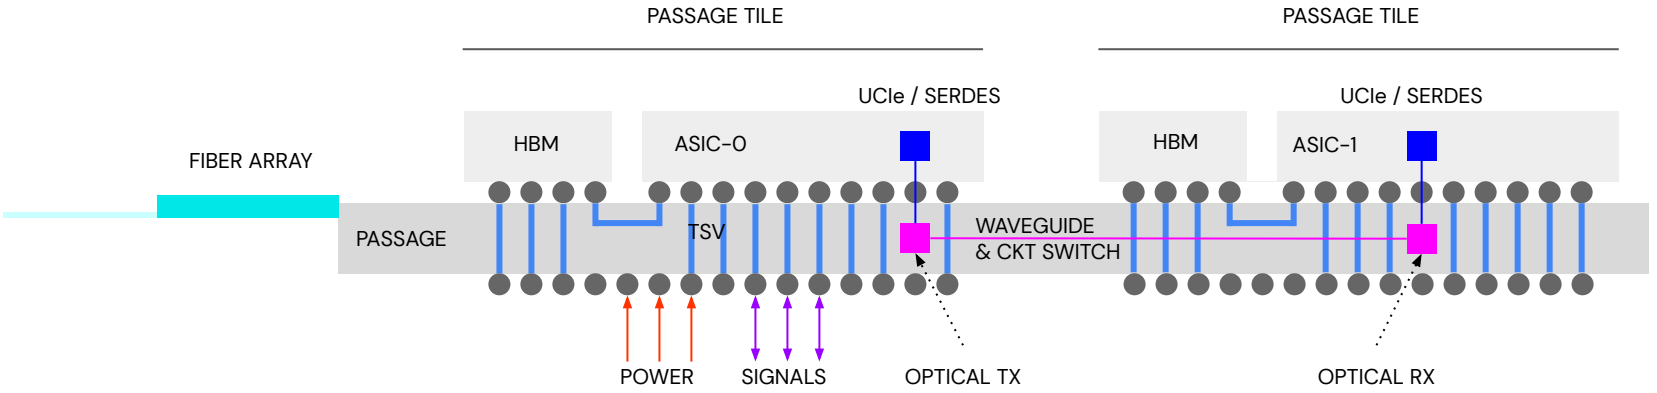

- Cooling solution depends on tile ASIC TDP.
- Up to 0.9W/mm^2 via TSVs

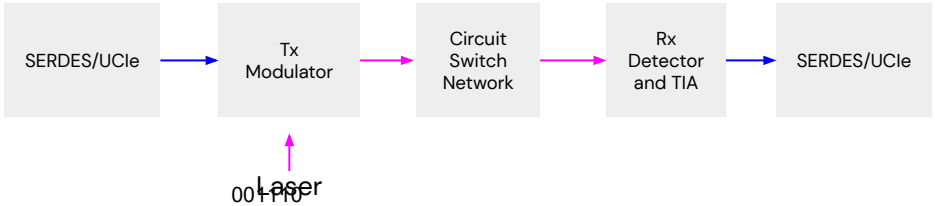

Intra (inside)-System Comparison

Passage

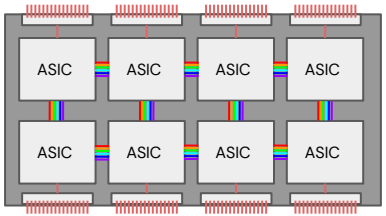

Top View

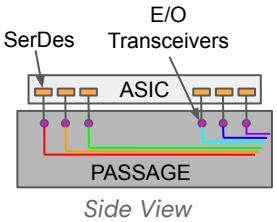

Side View

CPO

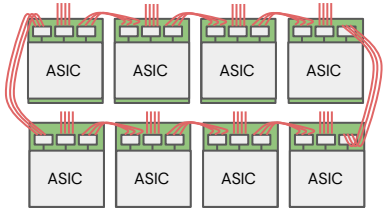

Top View

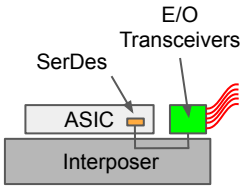

Side View

Electrical / Pluggable Optics

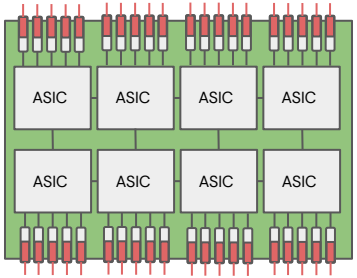

Communication	Photonic Waveguides	Optical Fibers	Electrical Traces + DAC/AOC
Energy efficiency	< 2 pJ/bit	< 6 pJ/bit (ASIC interface + Optical SerDes)	MR SerDes < 30 pJ/bit
Latency	SerDes latency (<5ns) + time-of-flight + (possibly FEC)	SerDes latency (<5ns) + time-of-flight + (possibly FEC)	SerDes latency (<5ns) + time-of-flight + (possibly FEC)
Bandwidth density	200 Tbps/mm	0.2 Tbps/mm	0.02 Tbps/mm
Max Chip I/O	Area-limited Assuming 1/3 area of 400mm^2 dedicated to UCIe 16GT I/O: ~670 Tbps	Shoreline-limited Assuming entire shoreline of 20x4 mm: 80 mm x 0.2 Tbps/mm = 16 Tbps	Shoreline-limited Assuming entire shoreline of 20x4 mm: 80 mm x 0.02 Tbps/mm = 1.6 Tbps <assuming 400Gbps>
SerDes compatibility	USR/XSR SerDes	USR/XSR SerDes	VSR or MR SerDes

Inter-System Single-Tile Fiber Export Comparison

= UCIe
= MR SerDes

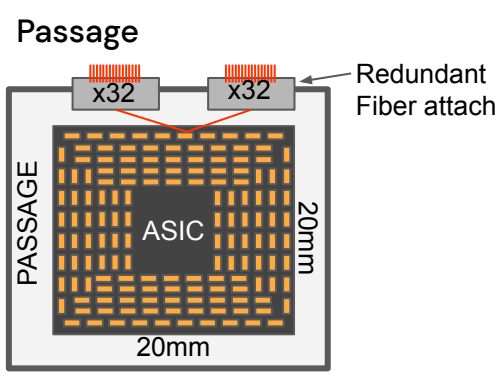

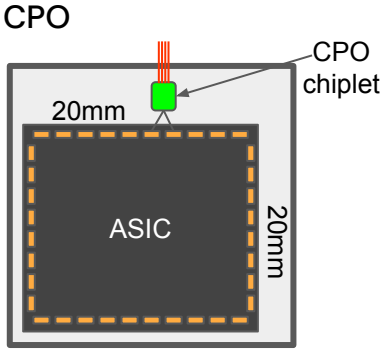

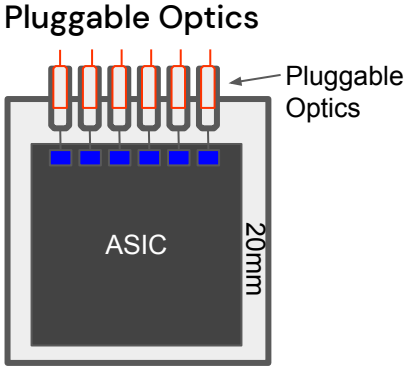

Communication	Optical fibers via Edge Attach	Optical Fibers	Plugable Active Optical cables
Energy	< 2 pJ/bit (ASIC interface + Optical SerDes)	< 6 pJ/bit (ASIC interface + Optical SerDes)	< 30 pJ/bit (ASIC interface + Optical SerDes)
I/O Density	Max Tile export, assuming 32 fibers per 20mm edge = 0.72 Tbps/mm 2x4 Passage I/O Fiber Escape for ⅓ area = 115.2 Tbps * a 1x1 passage tile w/ x32 fibers on 4 edges yields 57.6 Tbps @ 56Gbps modulation and 8 lambda ** 96% fiber yield, via attach redundancy, results in high I/O density (i.e. more fibers v. CPO)	Max Chip I/O for 20mmx20mm (4 edges) = 0.2 Tbps/mm	Max Chip I/O w/ 400G DAC = 20Gbps/mm
Reach / Ease of scaling out	<1 km via fibers and optically circuit switched. Packet switching capable (more lasers required)	<1 km via fibers	20 cm for VSR., 50 cm for MR.; Standard packet switching technology available.
Fiber Losses	0.2 dB/km	001112 0.2 dB/km	loss: 20 dB @ 28 GHz (VSR), 40 dB @ 28 GHz (MR), Depends on modulation frequency.

First Silicon Success

The world's first photonic wafer-scale interconnect

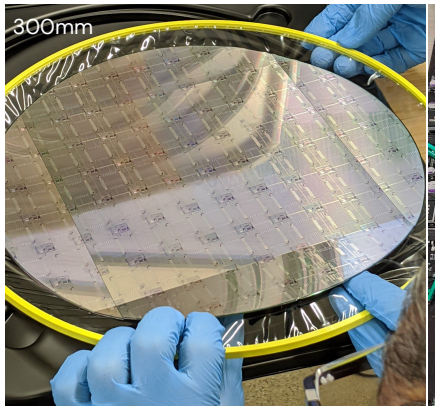

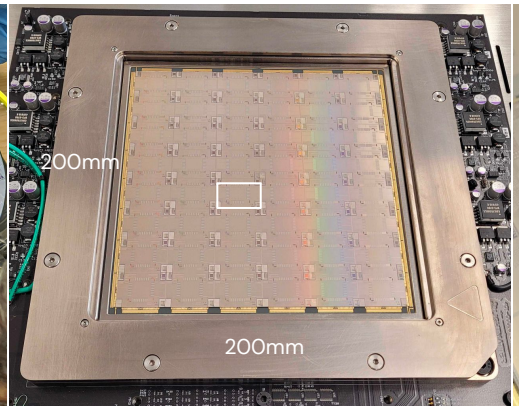

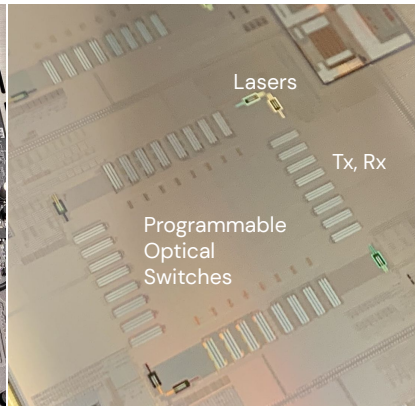

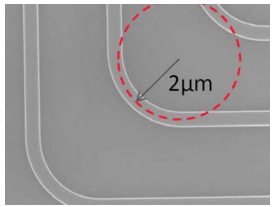

Photonic waveguides with ~4 µm pitch.

Passage™ Alpha Silicon

- <50 Watts
- 32 channels per site, 1.024 Tbps
- 32 Gbps per channel NRZ
- 48 x 800mm^2 tiles
- 288x 50 mW Lasers
- 6,144 DACs
- 6,144 MZIs
- 150,000 photonic components
- JTAG interface
- Integrated Lasers, transistors, photonics
- Programmable interconnect topologies

Solutions LM is Driving

A variety of applications

DISAGGREGATION

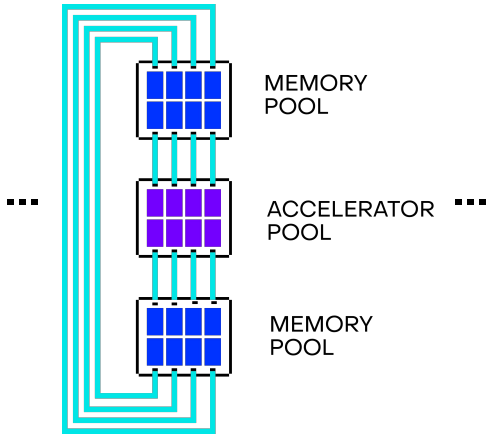

DYNAMIC COMPUTE ALLOCATION & AIR GAP ISOLATION

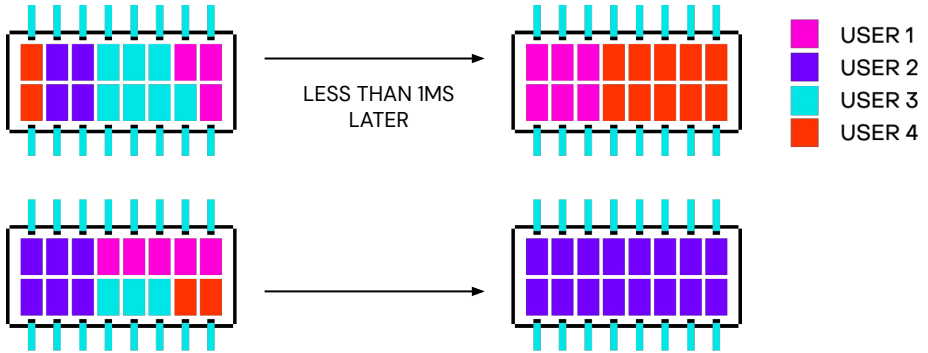

Call to Action

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

Conclusion

Rapid advances in AI and HPC enabled by photonics.

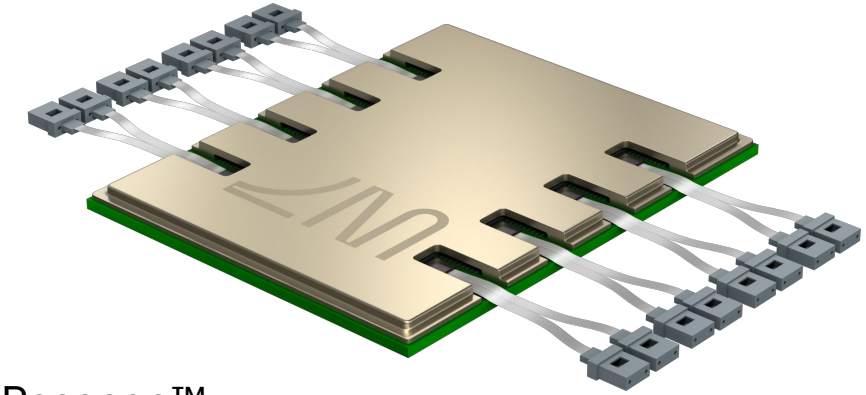

Passage™

3D programmable photonic interconnect that pushes the limits of high bandwidth and low latency… and it makes your deployment simpler.

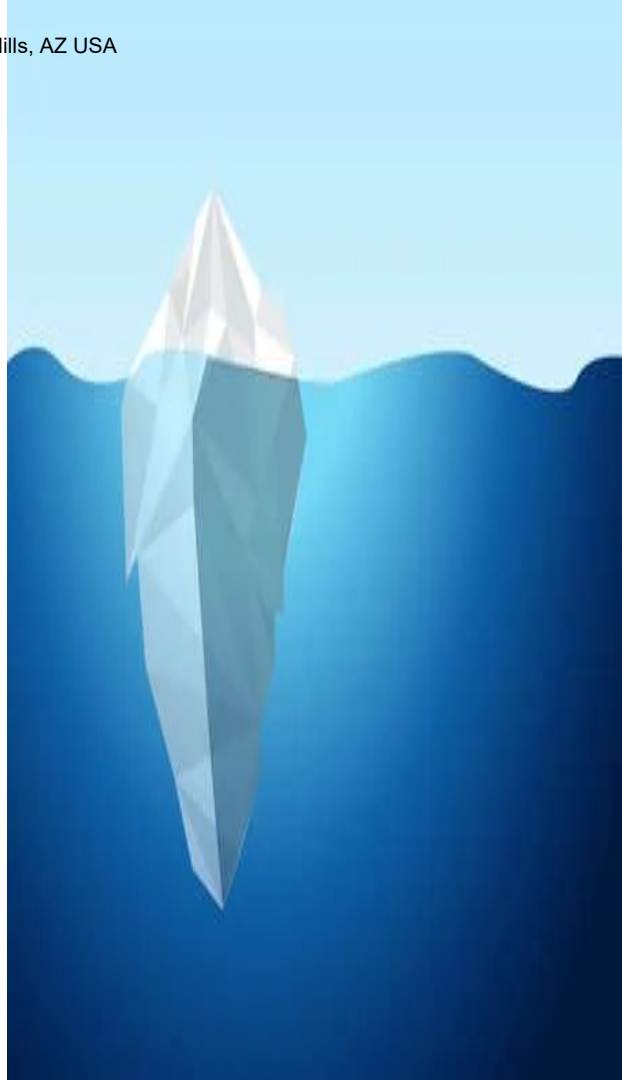

THANK YOU

Q & A