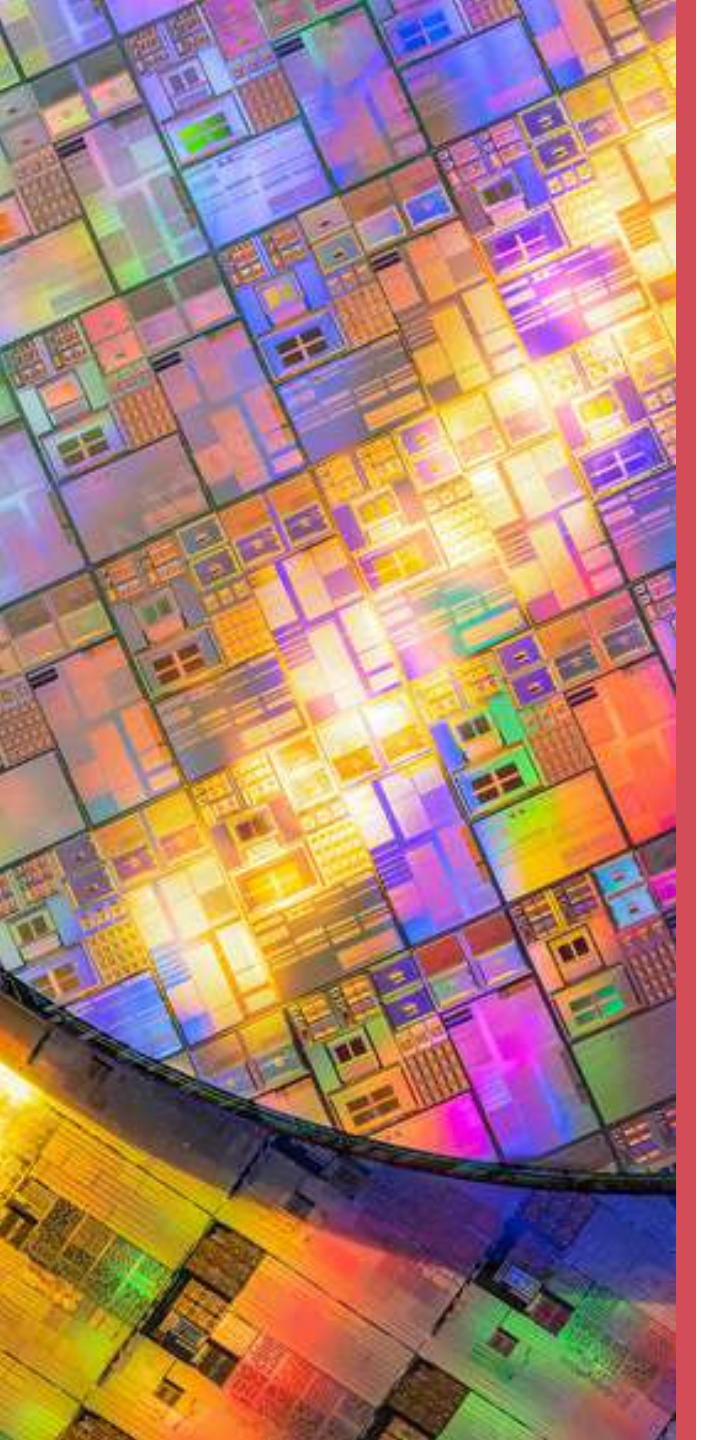

Breaking the bottleneck: the future of AI computing with optical interconnects

About me

Early days
(1992 - 2014)

Master
(2014 - 2016)

Politecnico di Torino

Grenoble INP phelma

EPFL

IBM

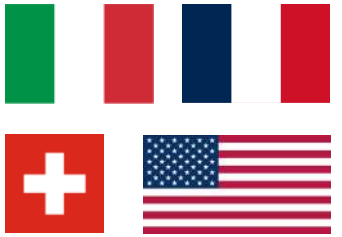

PhD
(2017 - 2021)

Université de Paris

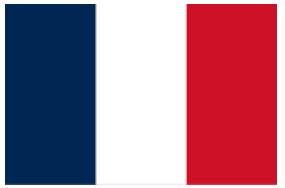

Pro life
(2021 -)

HEC Paris Challenge+

NcodiN
Illuminating the future of computing

What’s behind ChatGPT?

GenAI triggered extraordinary investments in supercomputing clusters

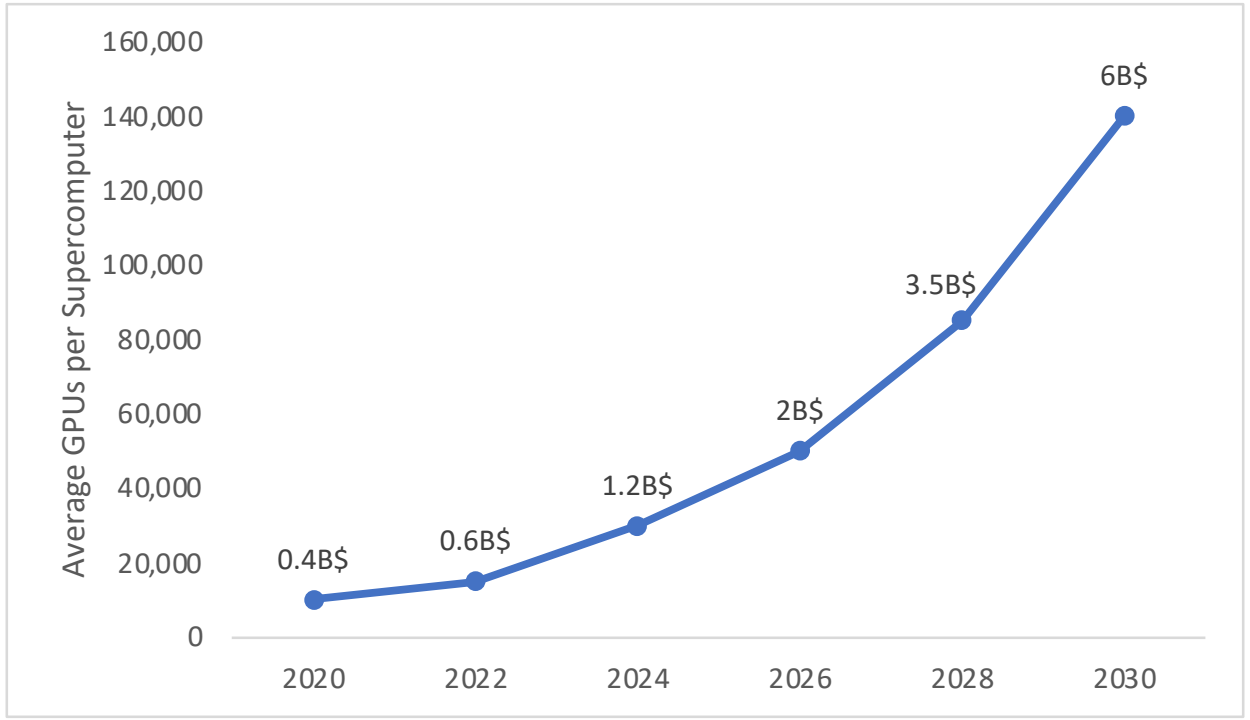

Moore's Law is over

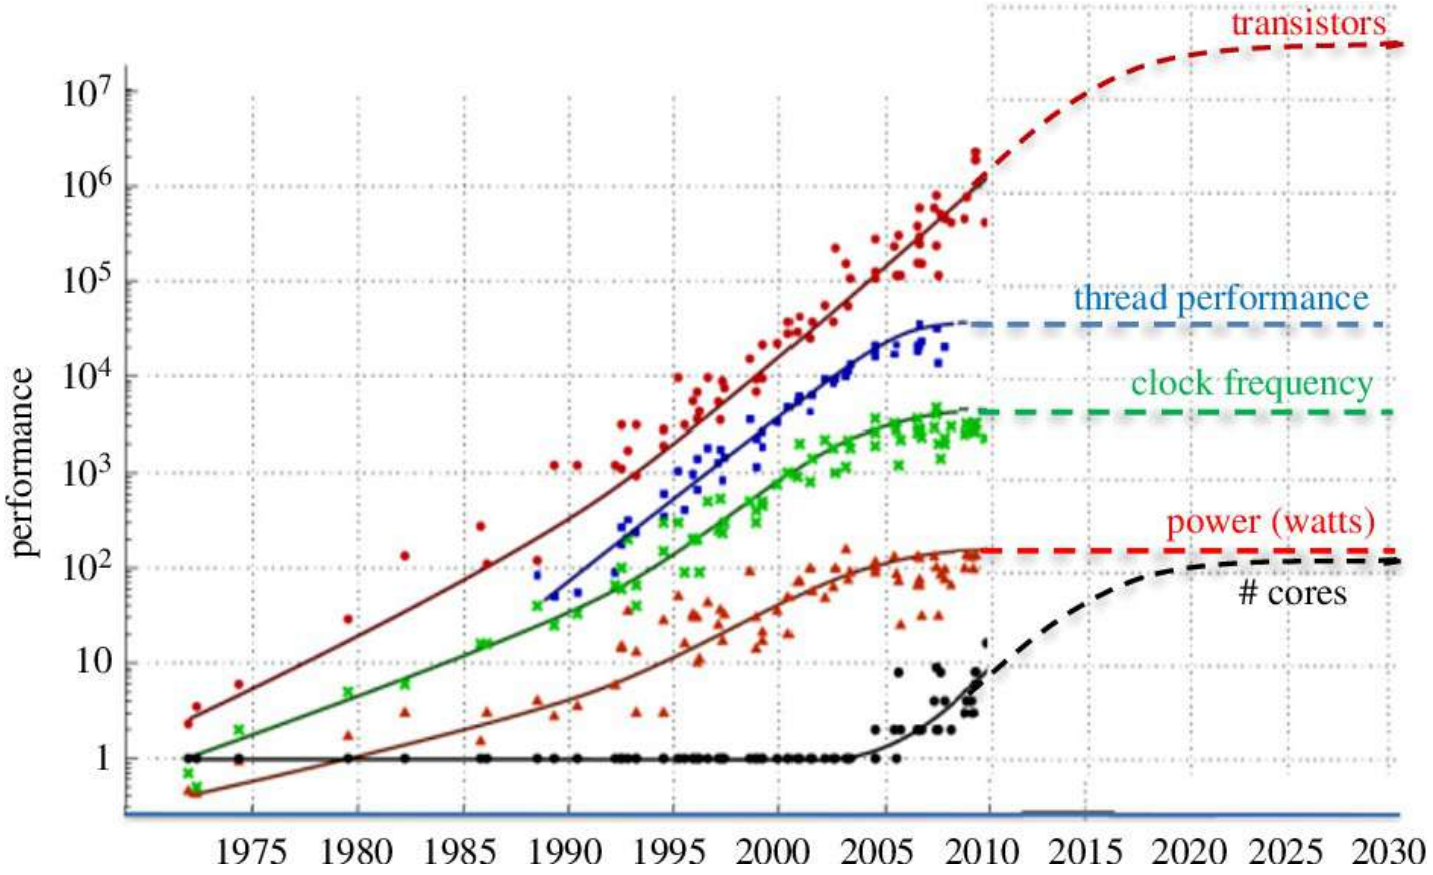

The industry is pivoting to disaggregated chiplet-based solutions

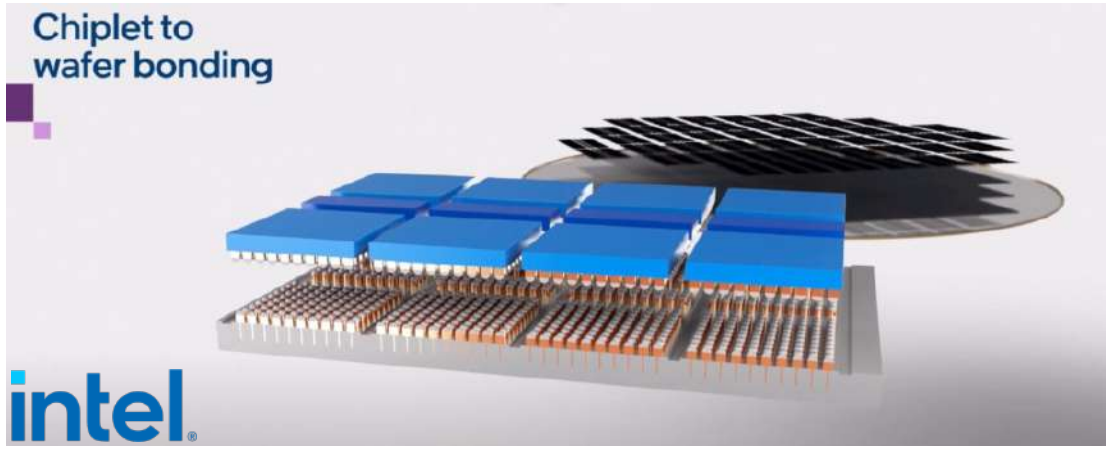

AMD MI300

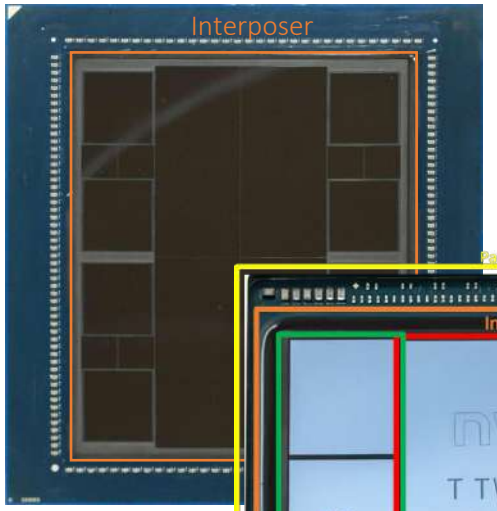

NVIDIA H100

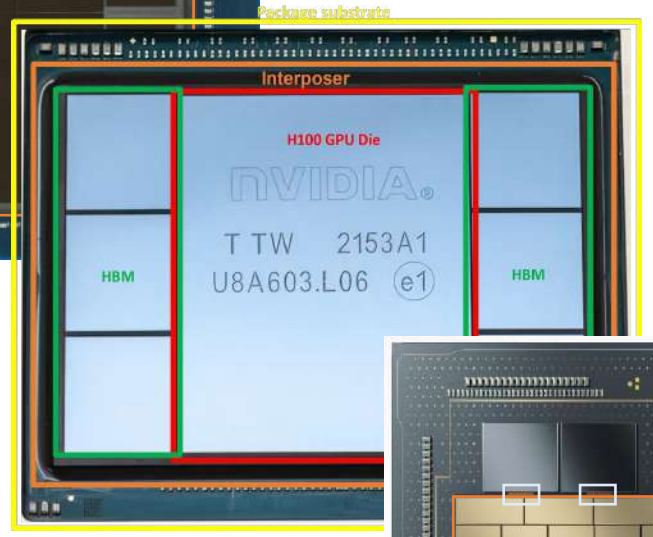

intel Ponte Vecchio

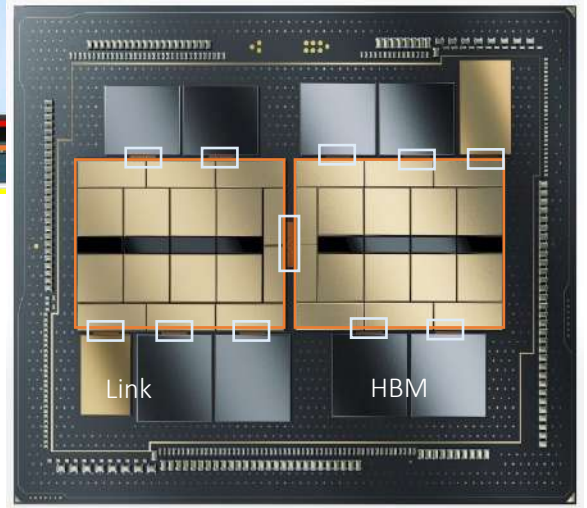

Increasing chip size comes at a cost (yield, design)

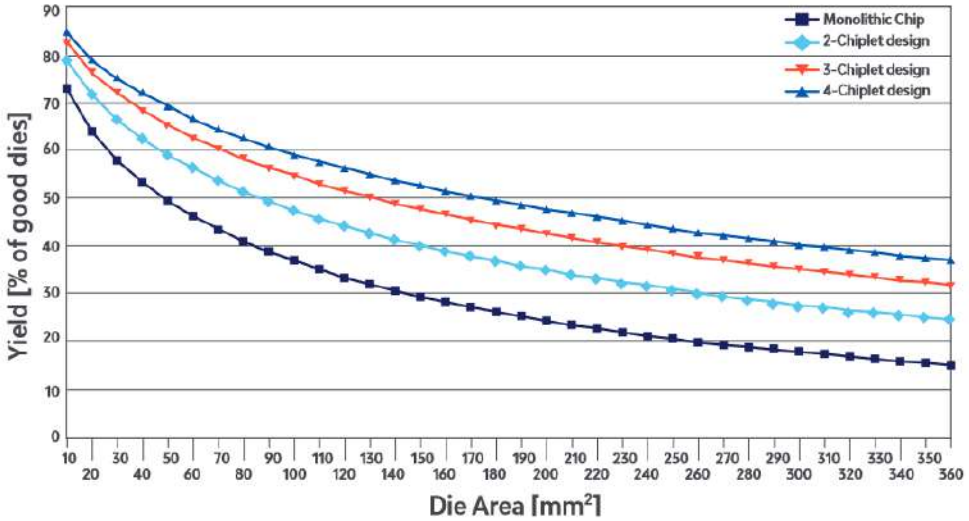

New processors need more dies (silicon) for their circuits ...

2X size!

H100

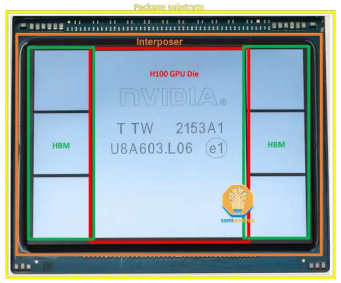

Blackwell

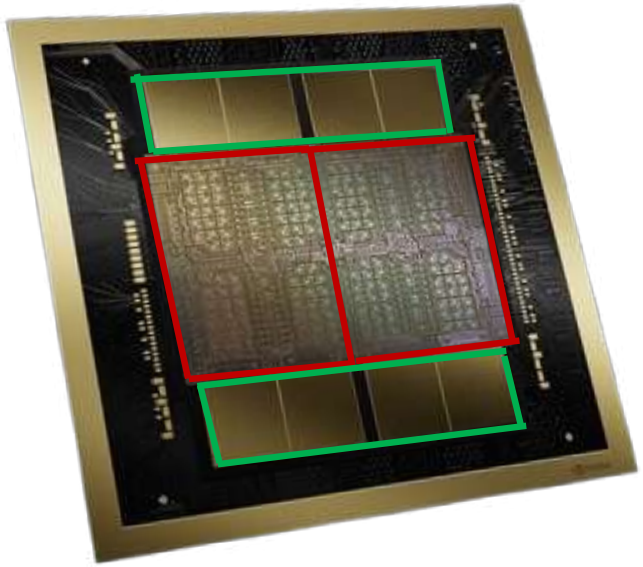

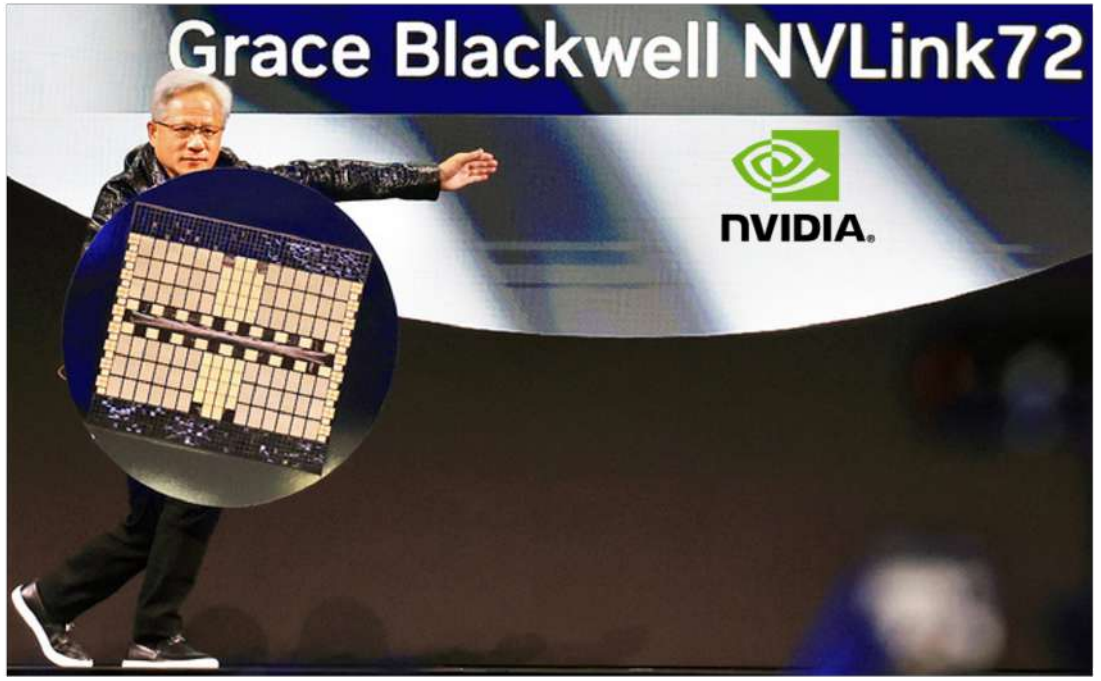

2023 2024 > 2030

... and GPUs need to be kept busy all the time!

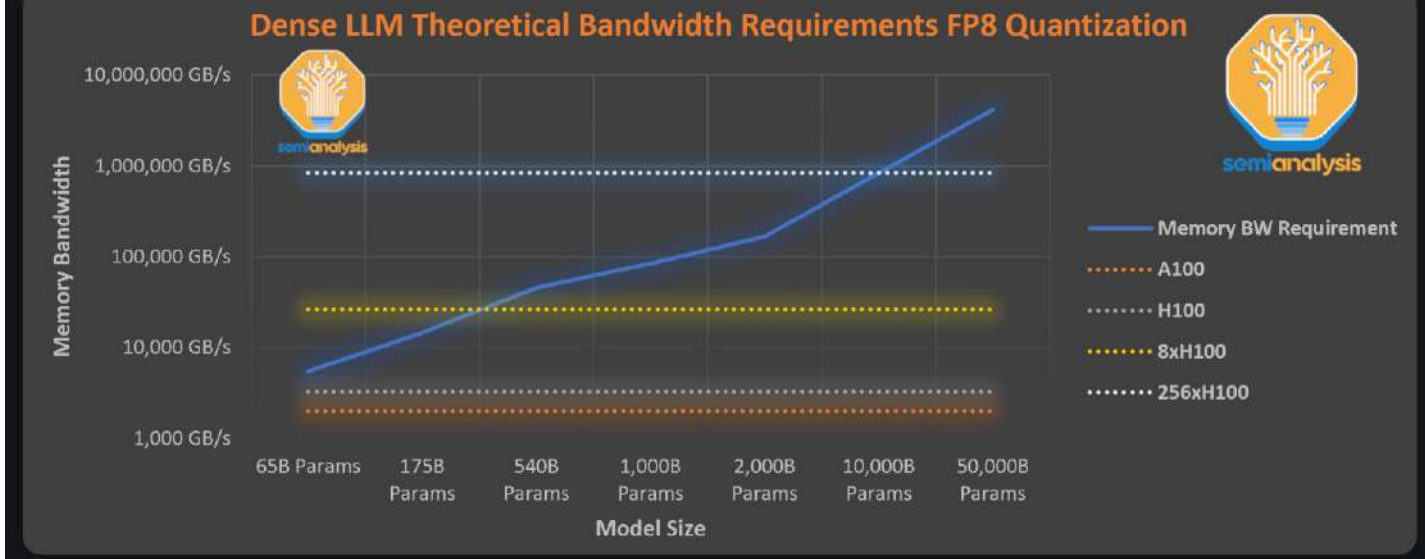

This chart assumes inefficiencies from not being able to fuse every op, memory bandwidth required for the attention mechanism, and hardware overhead are equivalent to parameter reads. In reality, even with "optimized" libraries such as Nvidia's FasterTransformer library, the total overhead is even larger.

The chart above demonstrates the memory bandwidth required to inference an LLM at high enough throughput to serve an individual user. It shows that even 8x H100 cannot serve a 1 trillion parameter dense model at 33.33 tokens per second. Furthermore, the FLOPS utilization rate of the 8xH100's at 20 tokens per second would still be under 5%, resulting is horribly high inference costs. Effectively there is an inference constraint around ~300 billion feed-forward parameters for an 8-way tensor parallel H100 system today.

Copper limits reach, bandwidth and energy efficiency!

BW Density (Tbps/mm²)	2
Reach (mm)	2
Energy efficiency (pJ/bit)	0.5

BW Density (Tbps/mm²)	2
Reach (mm)	20
Energy efficiency (pJ/bit)	3.5

BW Density (Tbps/mm²)	2
Reach (mm)	200
Energy efficiency (pJ/bit)	15

2024 2028 > 2030

Copper limits reach, bandwidth and energy efficiency!

Energy efficiency (pJ/bit)	0.5
Reach (mm)	2
BW Density (Tbps/mm^2)	**2**

Energy efficiency (pJ/bit)	0.5
Reach (mm)	20
BW Density (Tbps/mm^2)	**0.3**

Energy efficiency (pJ/bit)	0.5
Reach (mm)	200
BW Density (Tbps/mm^2)	**0.07**

2024 2028 > 2030

Copper limits reach, bandwidth and energy efficiency!

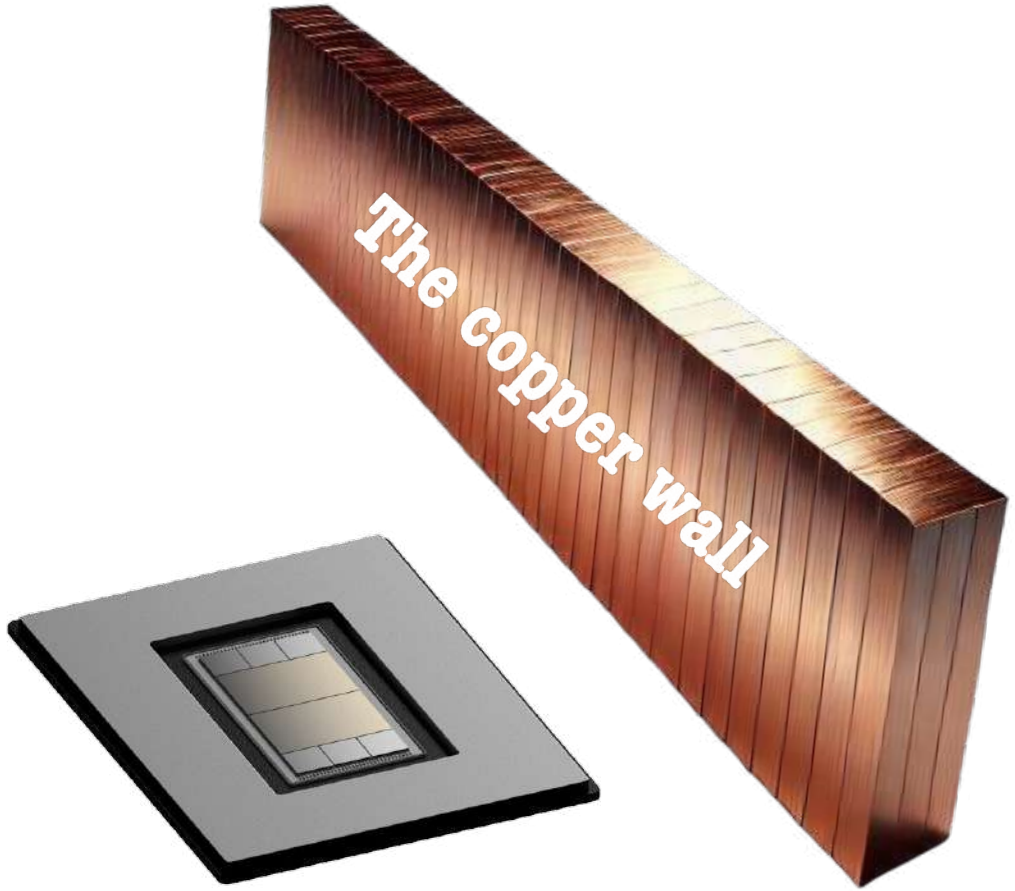

2024 2028 > 2030

Photonics is THE way

Transoceanic communication (80s)
> 10^3 km range

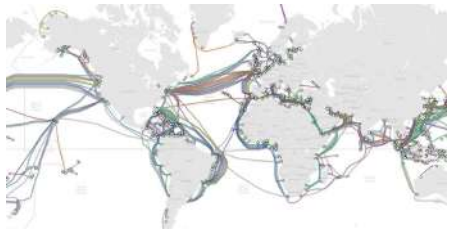

Local Area Network (early 2000s)
~km range

Intra-datacenter pluggable transceivers (2010s)
~m range

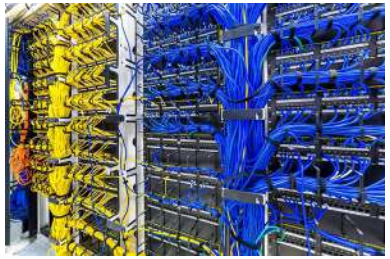

Co-packaged optics (now)
~m range

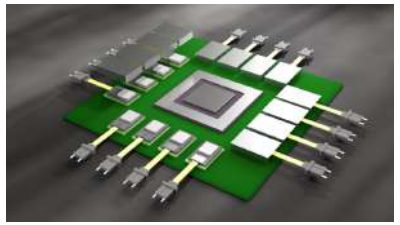

SiP die-to-die connectivity (~2030)
≤1 cm range

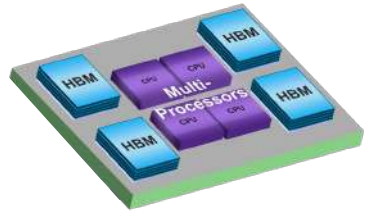

Bandwidth demand

Introducing NConnect: the optical network powered by the world's smallest laser

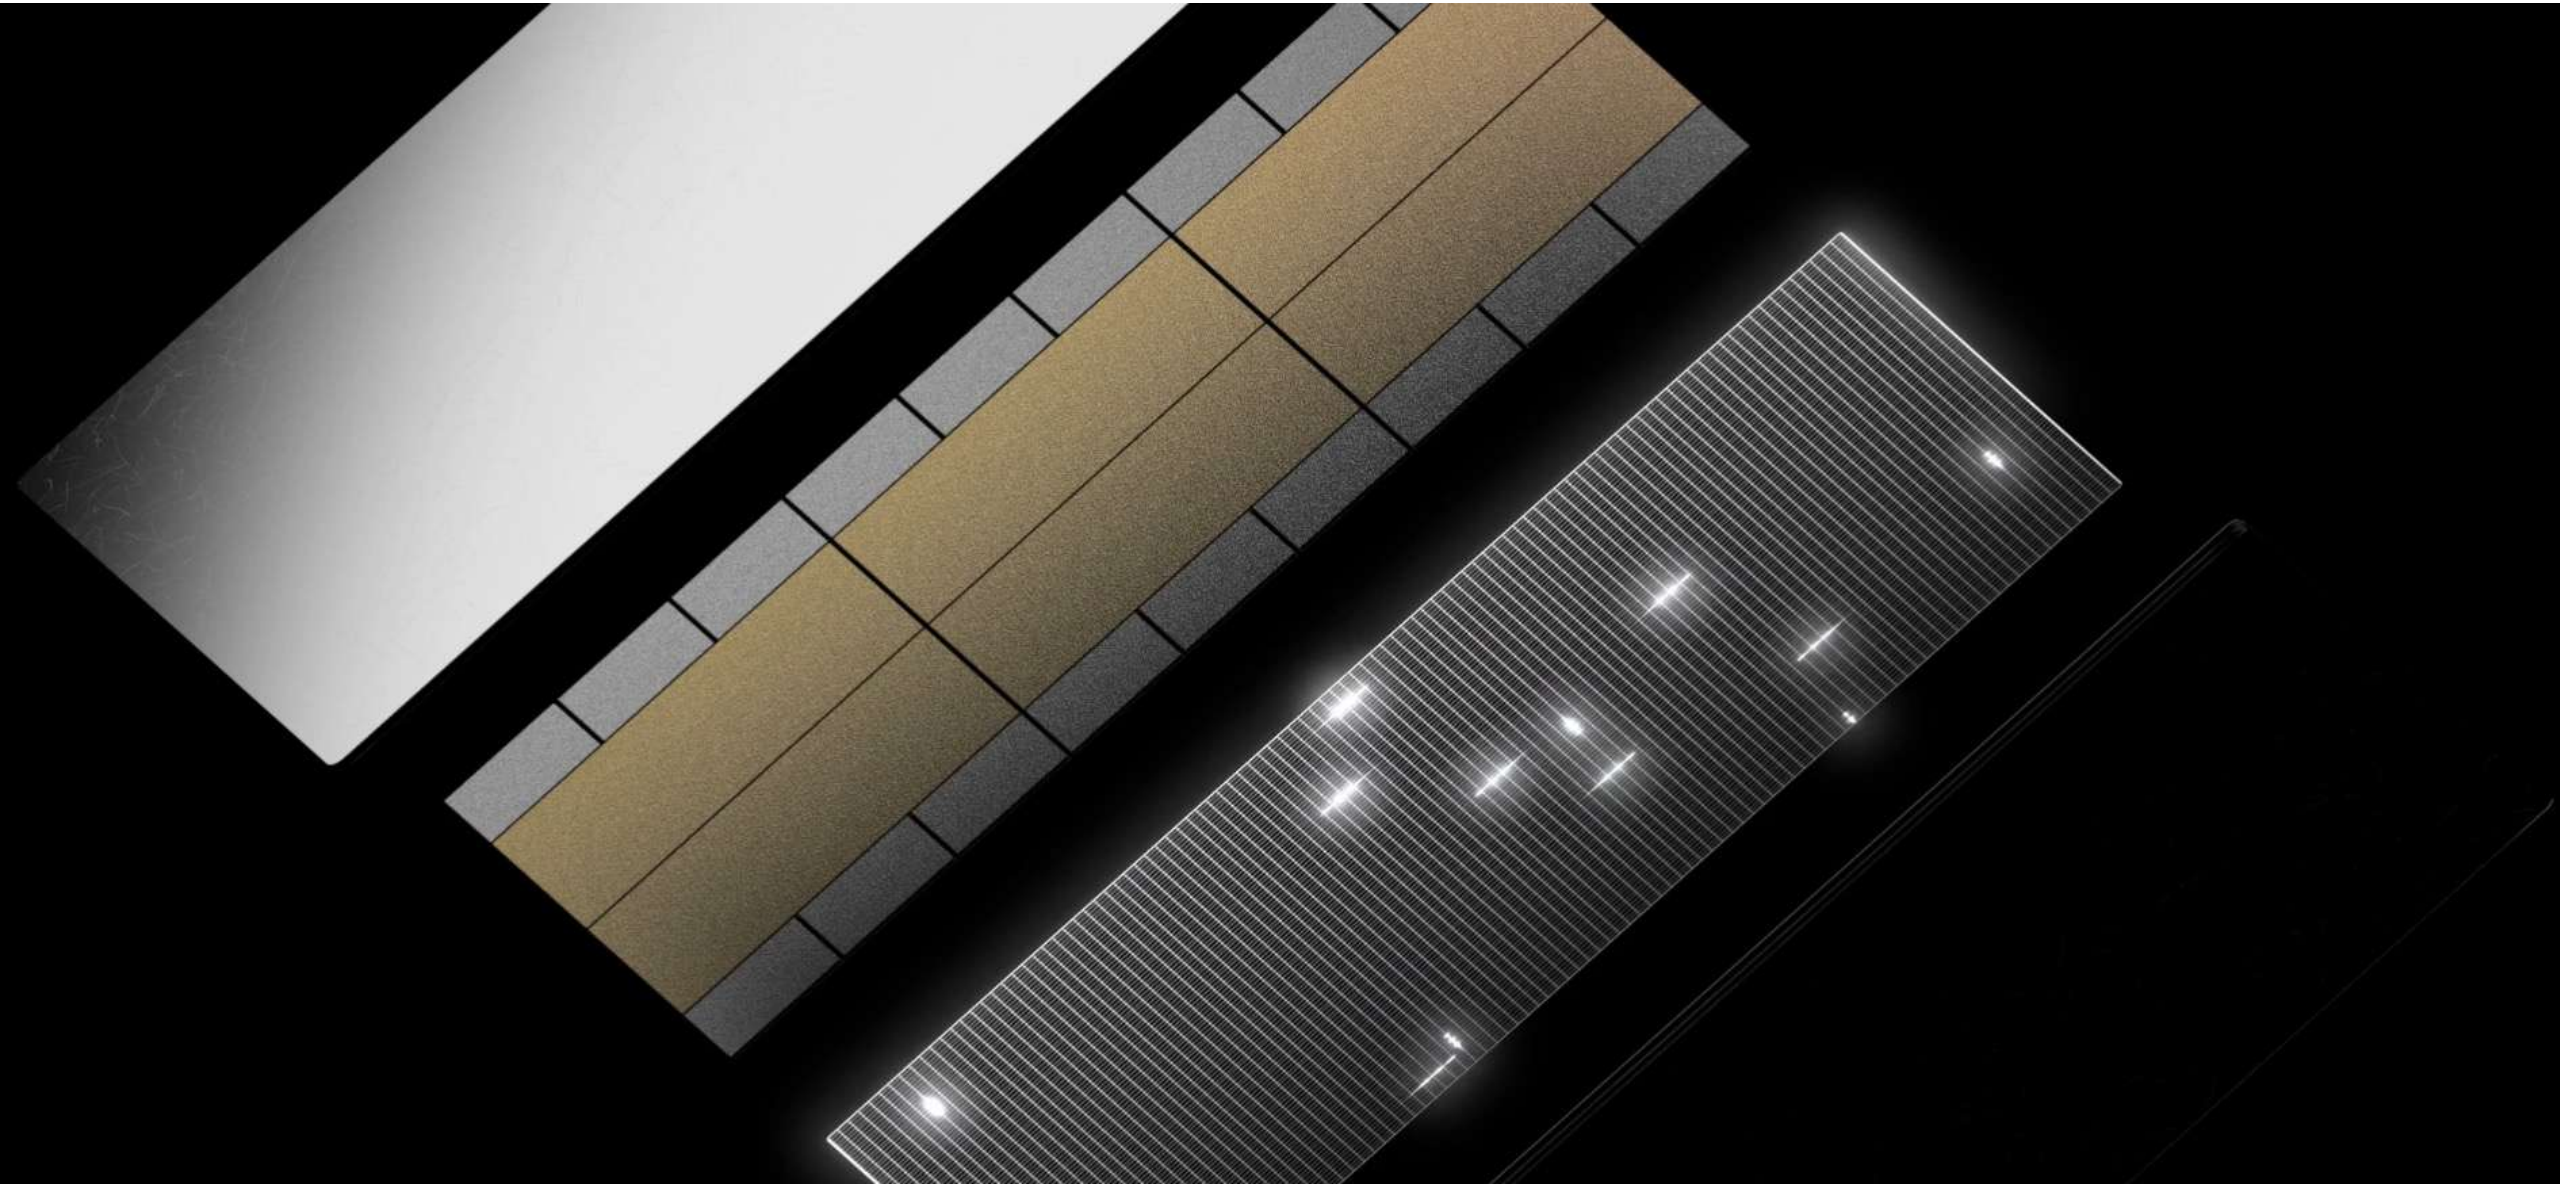

Our unique technology results from >15 years of research @ C2N-CNRS

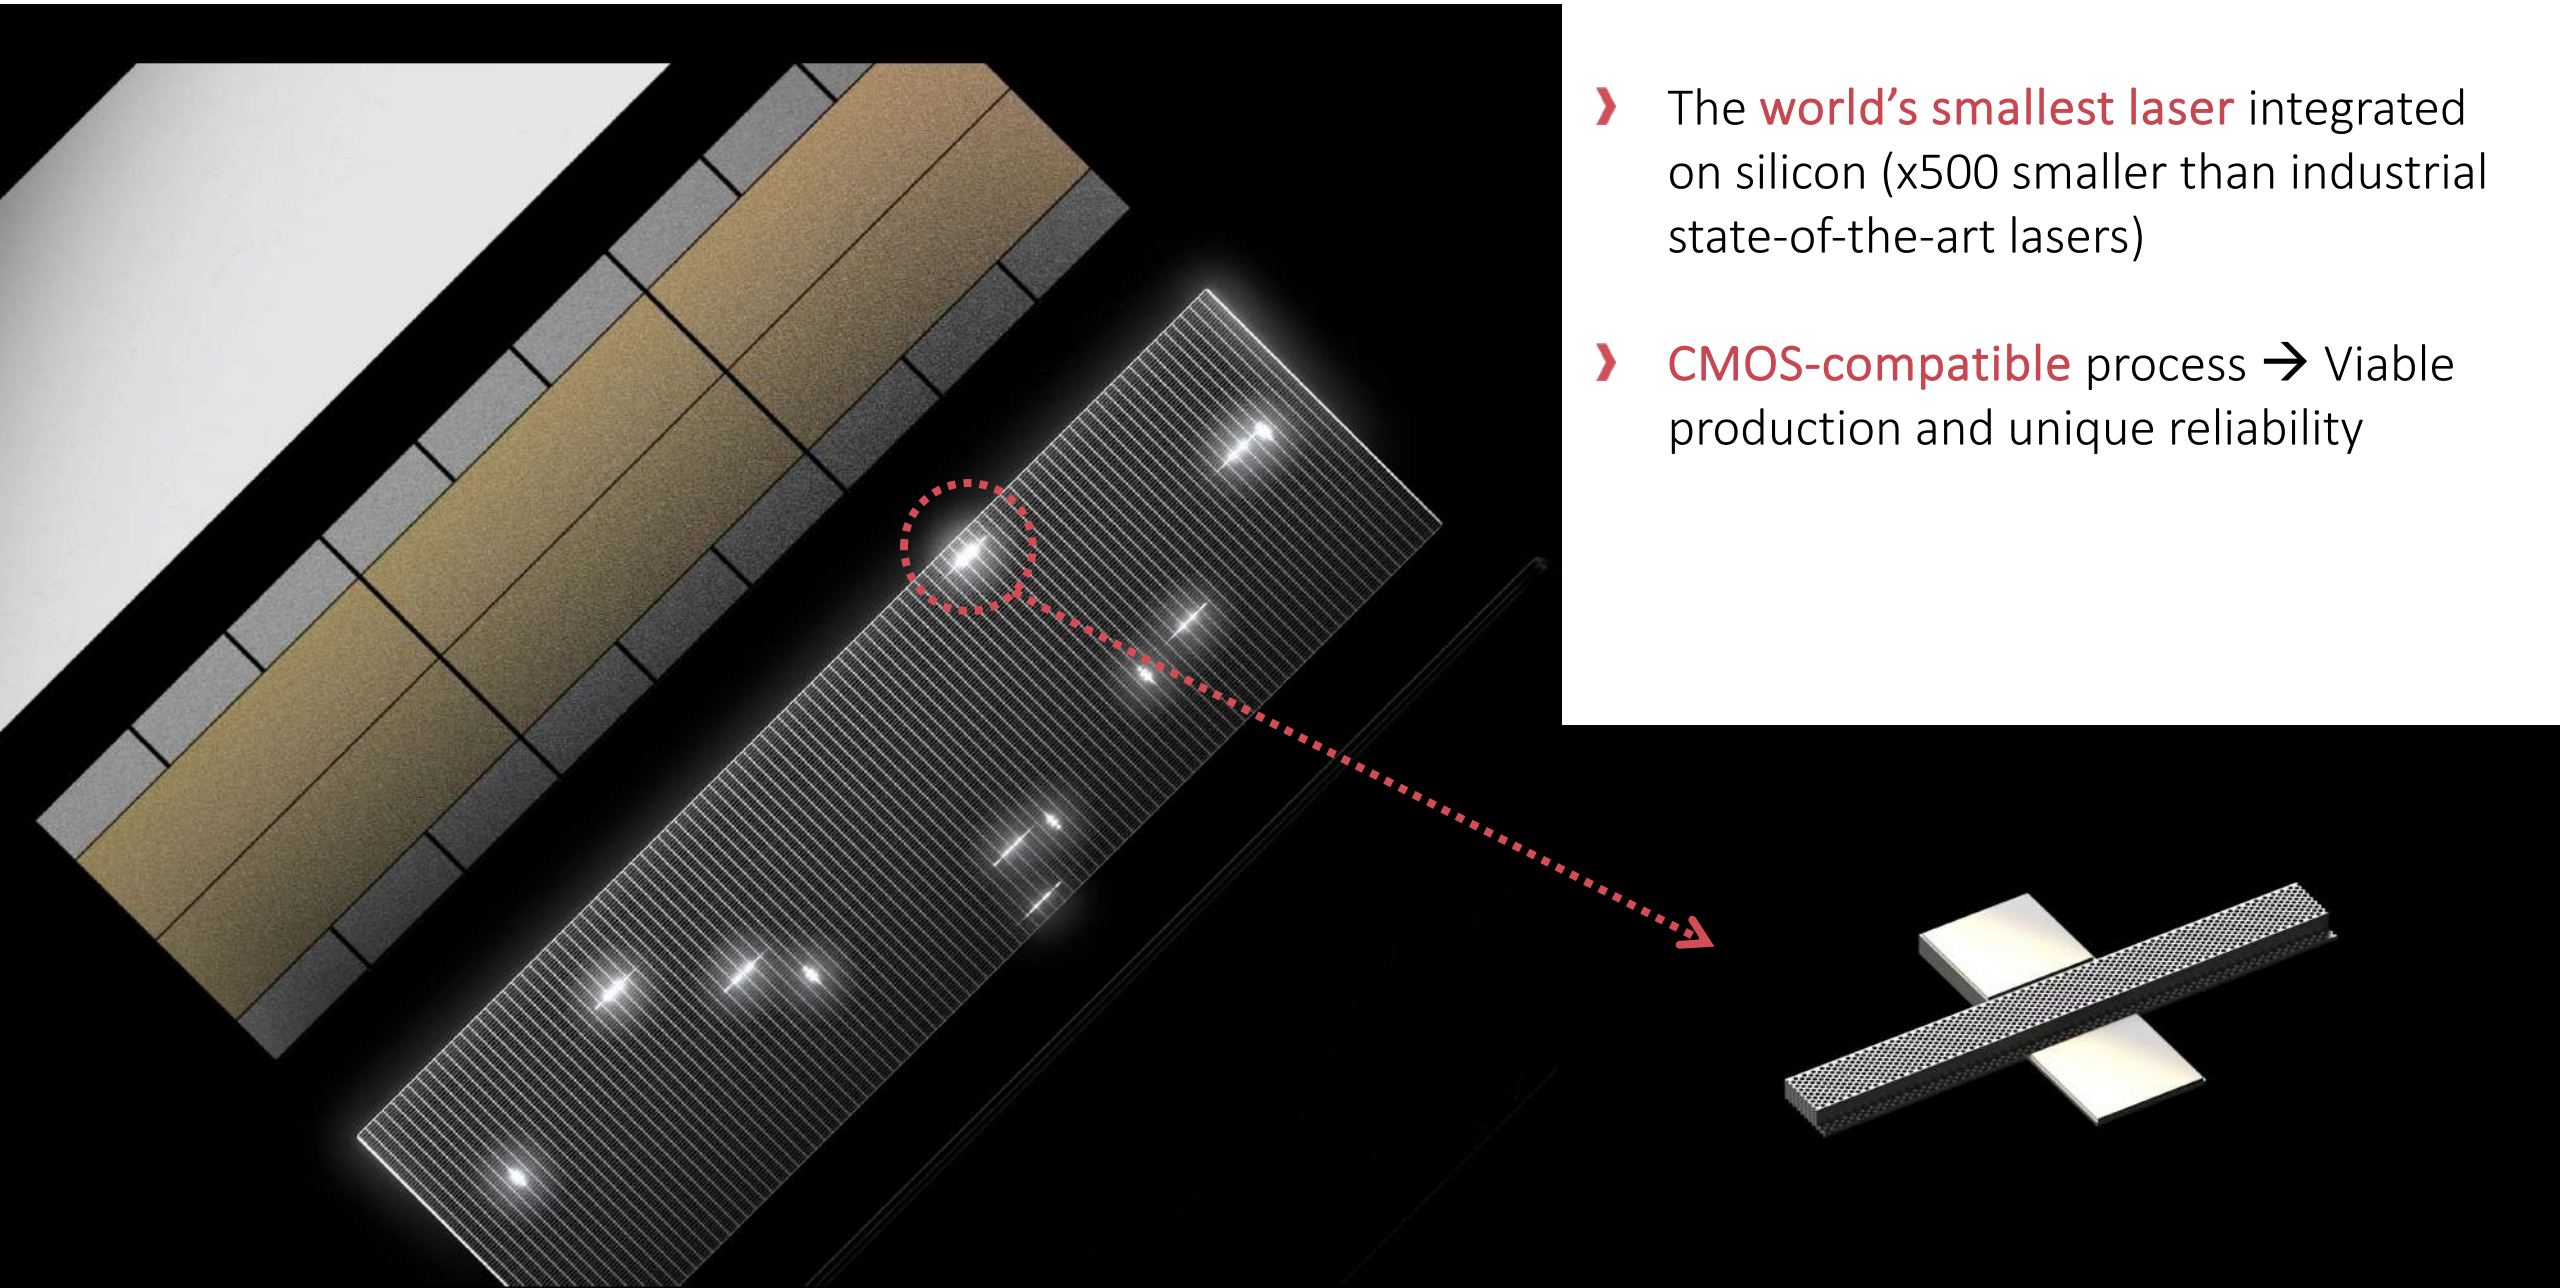

- The world's smallest laser integrated on silicon (x500 smaller than industrial state-of-the-art lasers)
- CMOS-compatible process → Viable production and unique reliability

How does it work?

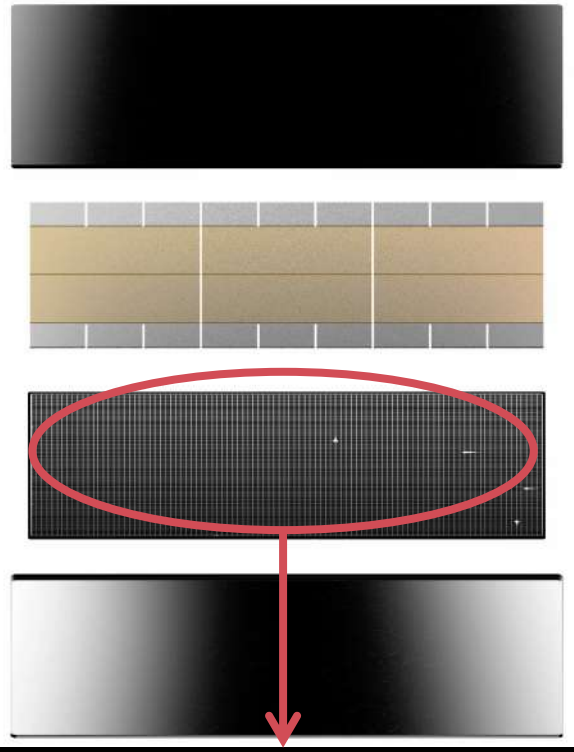

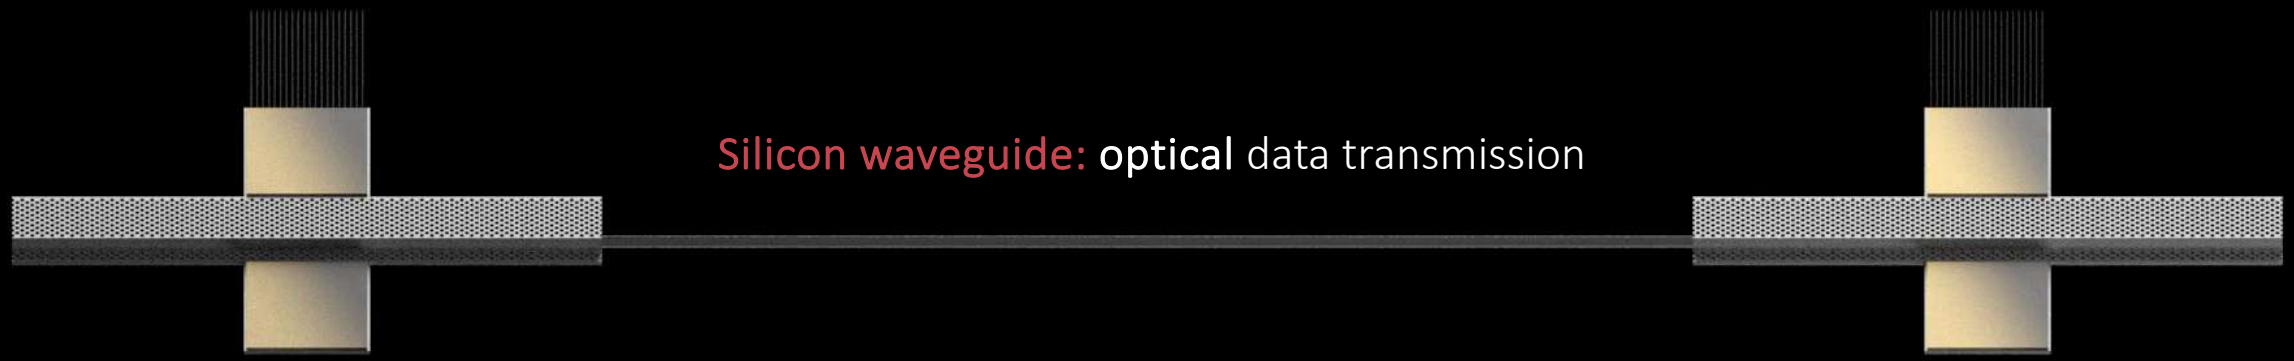

Silicon waveguide: **optical** data transmission

Nanolaser: electro-optical signal conversion

Nanodetector: opto-electronic signal conversion

Our technology works with existing chiplet architectures

OPEN Compute Project

Bunch of Wires (BoW) PHY Specification

The Open Domain-Specific Architecture BoW Workstream

CHIPS ALLIANCE

Advanced Interface Bus (AIB) Specification

UCIe™ Universal Chiplet Interconnect Express

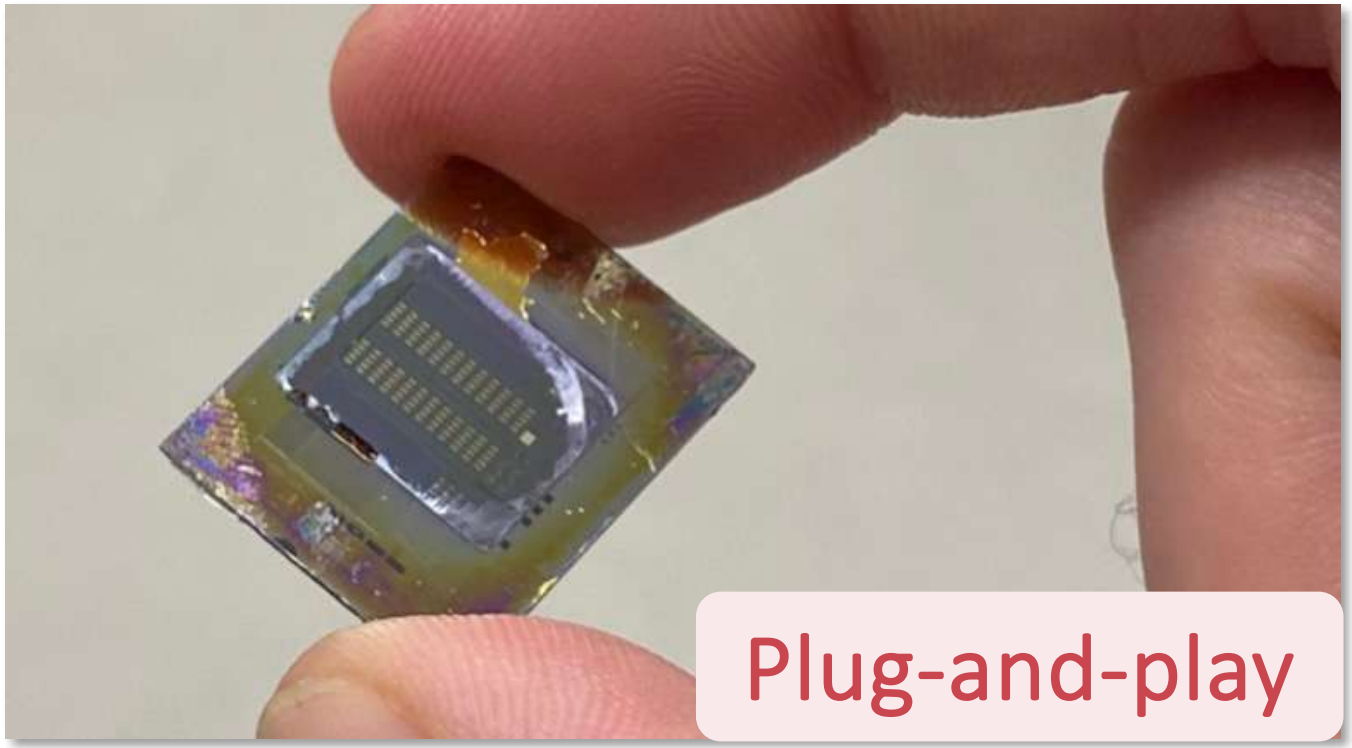

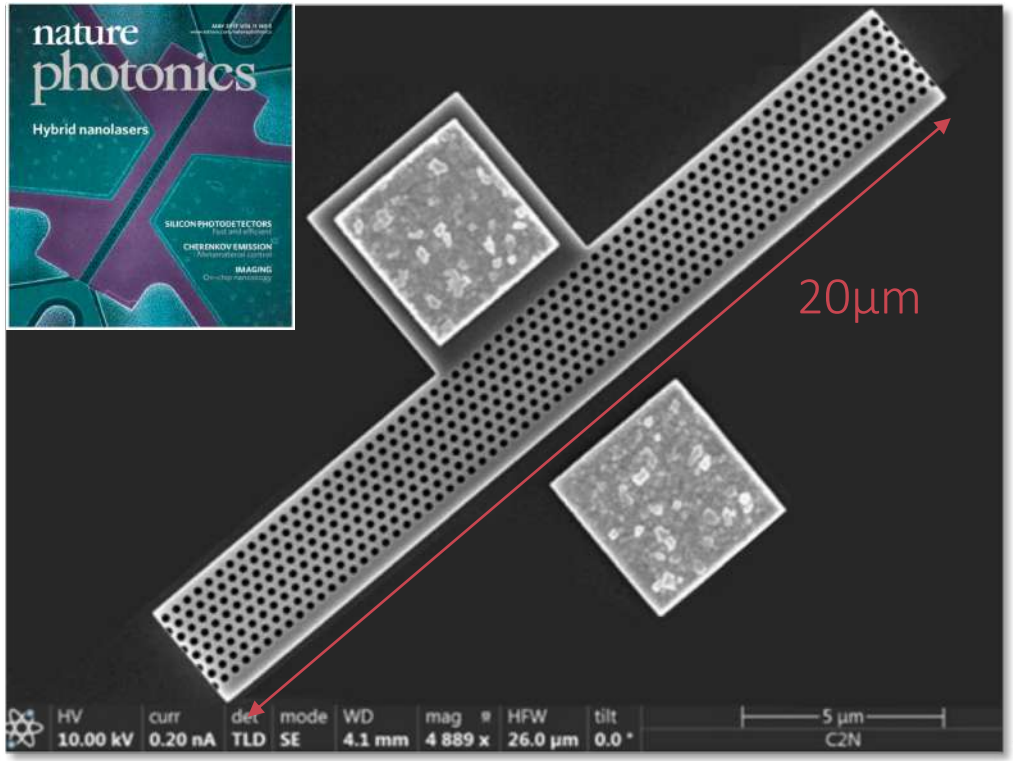

NcodiN solves the interconnect bottleneck limitations all at once …

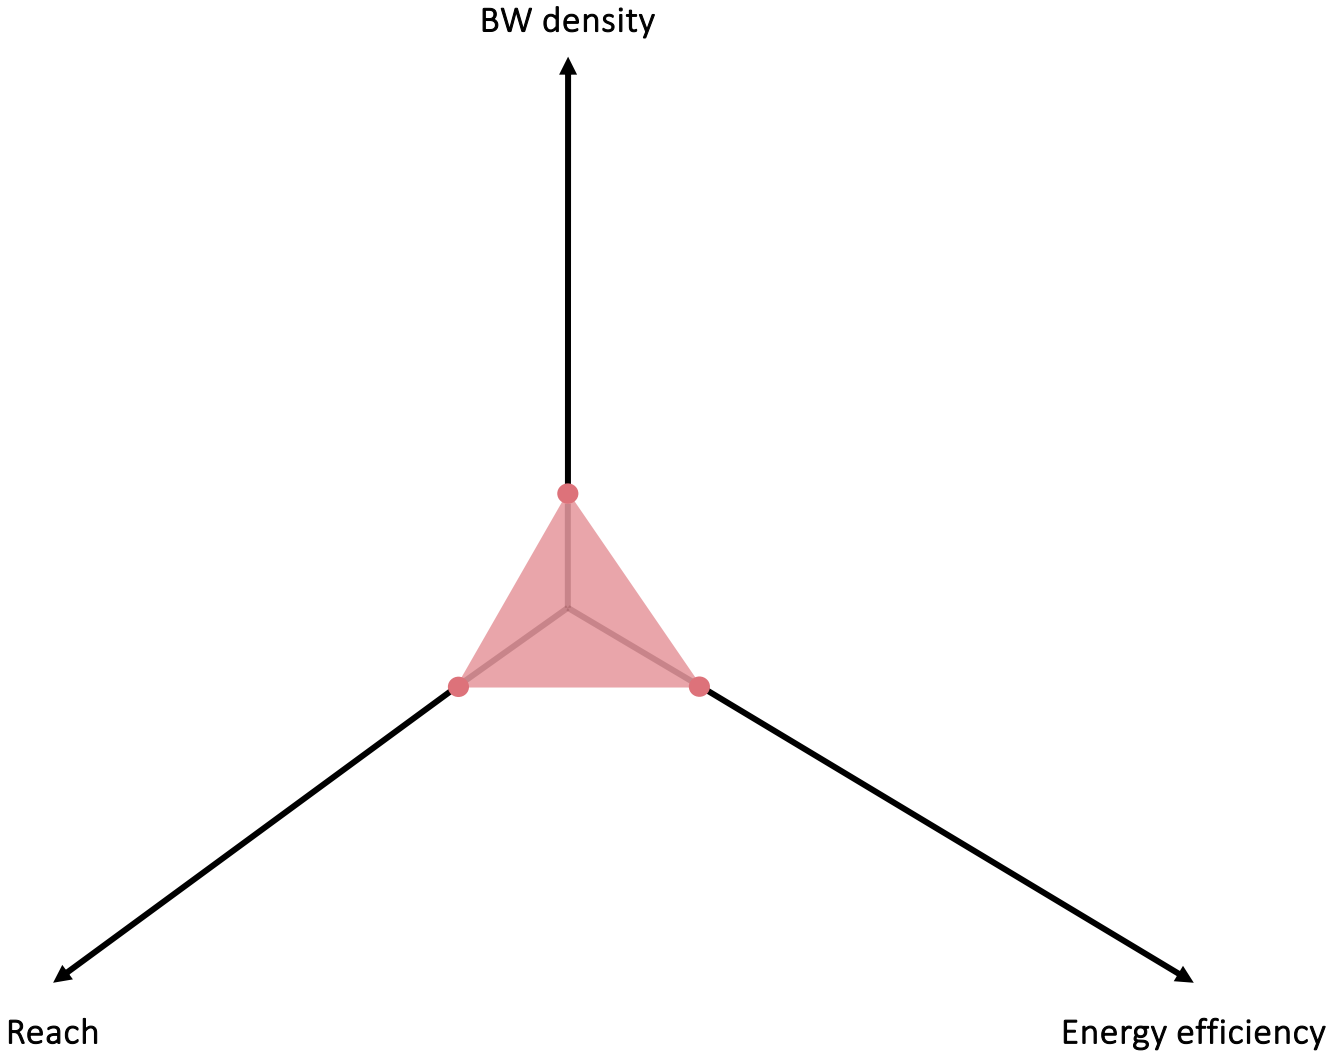

Figure of Merit	Copper
BW Density (Tbps/mm^2)	2
Reach (mm)	2
Energy efficiency (pJ/bit)	0.5
Communication system	Point-to-Point (bidirectional)

… through an Optical Network-On-Chip with superior performance

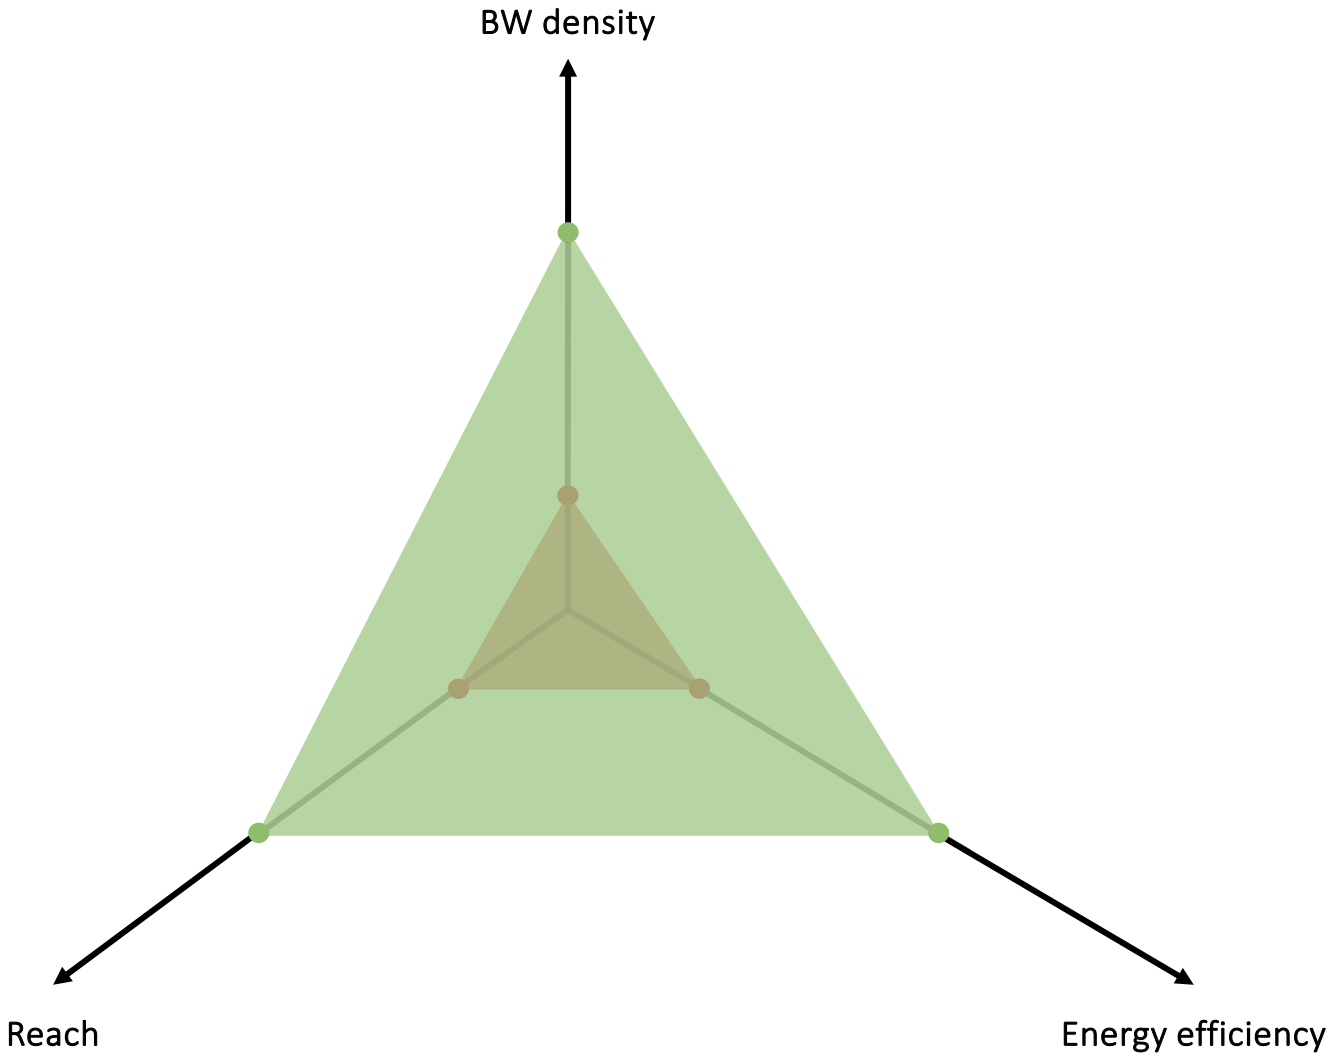

Figure of Merit	Copper	NcodIN
BW Density (Tbps/mm^2)	2	≥60
Reach (mm)	2	≥50
Energy efficiency (pJ/bit)	0.5	<0.1
Communication system	Point-to-Point (bidirectional)	Point-to-All (broadcasting)

We enable new horizons in the design space

Elec I/O-based standard package

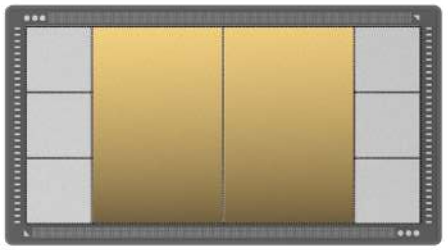

+

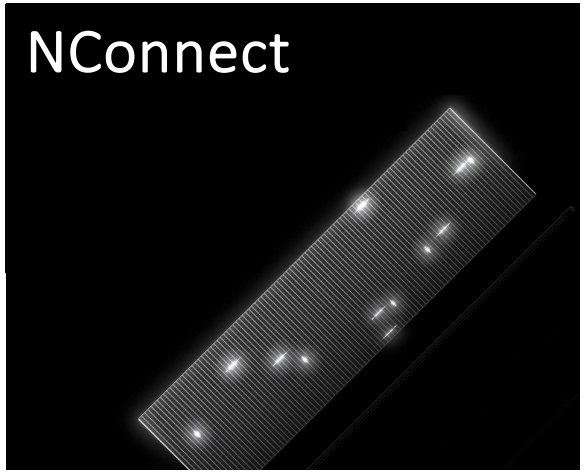

=

1) “Super-package”

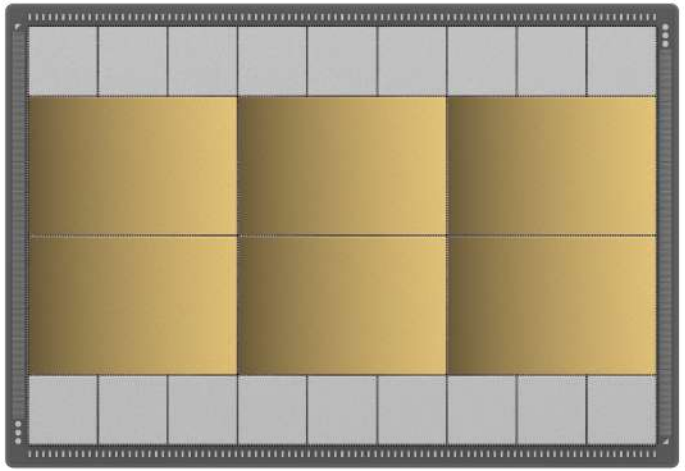

$/performance ↑

2) Larger distance

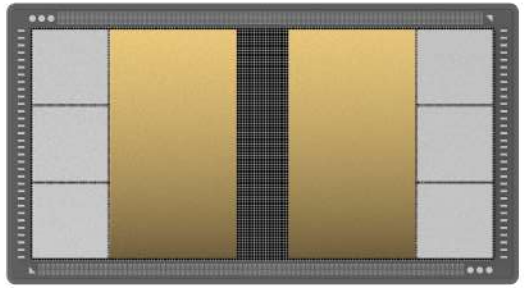

Warpage ↓

3) More memory

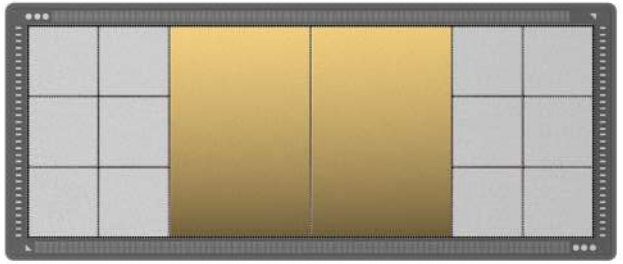

Memory capacity ↑

HBM warpage ↓

4) Enhanced disaggregation

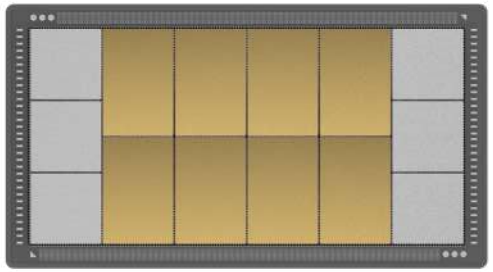

Yield ↑

Design flexibility ↑

The only photonic solution allowing to scale on-interposer copper

In-package interconnects (2.5D die-to-die)	NcodiN + Copper		LIGHTMATTER	celestial AI	AVICENA
Manufacturing viability	✅	✅	❌	❌	❌
Plug-and-play capability	✅	✅	❌	❌	❌
Scalability	✅	❌	❌	❌	❌

> *"You have the most suited solution to break the power/reach tradeoff for on-interposer copper links.*
>
> *Other photonics players cannot be competitive on your use-case: if a solution helps increase the bandwidth inefficiently in terms of cost and power, it will never be adopted."*

Director Photonics Products NVIDIA

The only photonic solution allowing to scale on-interposer copper

From today's packaging …

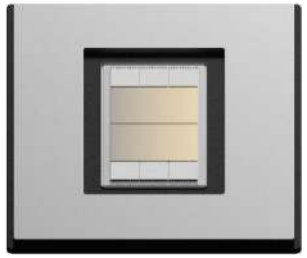

	Copper
Bandwidth demand (Tbps)	80
Power (W)	25
Cost ($)	500

In-package interconnects (2.5D die-to-die)	Ncodin + Copper		LIGHTMATTER	celestial AI	AVICENA
Manufacturing viability	✅	✅	❌	❌	❌
Plug-and-play capability	✅	✅	❌	❌	❌
Scalability	✅	❌	❌	❌	❌

> *"You have the most suited solution to break the power/reach tradeoff for on-interposer copper links.*
>
> *Other photonics players cannot be competitive on your use-case: if a solution helps increase the bandwidth inefficiently in terms of cost and power, it will never be adopted."*

Director Photonics Products — NVIDIA

The only photonic solution allowing to scale on-interposer copper

From today's packaging …

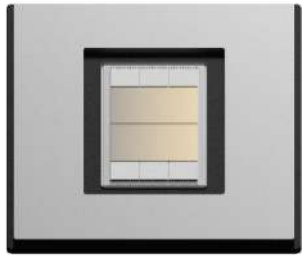

	Copper
Bandwidth demand (Tbps)	80
Power (W)	25
Cost ($)	500

… to scaled architecture

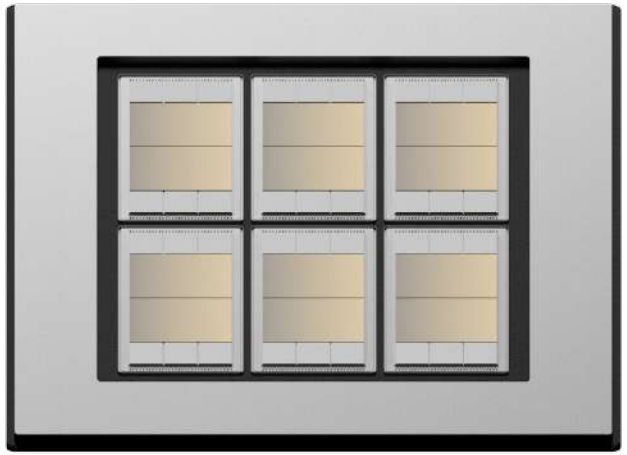

	Ncodin + Copper	LIGHTMATTER	celestial AI	AVICENA
Bandwidth demand (Tbps)	300			
Power (W)	15 + 50	1500	900	300
Cost ($)	1000 + 1000	150 000	90 000	30 000

In-package interconnects (2.5D die-to-die)	Ncodin + Copper		LIGHTMATTER	celestial AI	AVICENA
Manufacturing viability	✅	✅	❌	❌	❌
Plug-and-play capability	✅	✅	❌	❌	❌
Scalability	✅	❌	❌	❌	❌

> *"You have the most suited solution to break the power/reach tradeoff for on-interposer copper links.*
>
> *Other photonics players cannot be competitive on your use-case: if a solution helps increase the bandwidth inefficiently in terms of cost and power, it will never be adopted."*

Director Photonics Products — NVIDIA

Ncodin — Illuminating the future of computing

- ✓ X100 Energy efficient
- ✓ X100 Cost efficient

A complementary team of nanophotonics pioneers …

Francesco Manegatti

CEO and co-founder

PhD in physics, optoelectronic nanodevices
5+ years experience in nanophotonics

Bruno Garbin

CTO and co-founder

PhD in physics, Neuromorphic Optical Systems (NOS)
10+ years on experimental optics and NOS

Fabrice Raineri

CSO and co-founder

Full Professor, 20+ years in nanophotonics
Pioneering research on optoelectronic nanodevices

Carlo Guareschi

Head of BizDev

30+ years experience in semiconductor industry

Yacine Halioua

Head of ProdDev

10+ years experience in semiconductor industry

Team highlights

Team size

18

R&D, BizDev, Prod. Dev., HR

Upcoming SOON

2

R&D, Prod. Dev.

… working with expert advisors and partners

Léo Apotheker
International executive
previous CEO at
HEWLETT PACKARD SAP
& board member at
Ncodin

Eric Meurice
International executive
previous CEO at
ASML
& board member at
intel

Jean-Pascal Tricoire
International executive
previous CEO at
Schneider Electric
& board member at
Qualcomm

Ashkan Seyedi
Director of Silicon Photonics
product
NVIDIA

Supported by

elaia OVNI EARLYBIRD VERVE VENTURES
intel Ignite bpifrance cnrs C2N Centre de Nanosciences et de Nanotechnologies
WILCO 21ST BY CENTRALESUPÉLEC HEC PARIS Agoranov

An ambitious and exciting roadmap

	2024	2025	2026
R&D	Optical link POC	Optical interposer advanced demos (multi-link, high-T, high-f)	
	Tech. pre-industrialization in CMOS pilot line (optical interposer industrial pilots on 300mm wafers)		
	→ 18	→ 37	→ 54
	€4.5m Pre-Seed (dilutive + non-dilutive) ✅	€15m Seed	

elaia EARLYBIRD OVNI VERVE VENTURES

- **Advanced demo** under development
- **Pre-industrialization** ongoing
- **Collaboration** with big scaler under definition
- **Seed round** under preparation

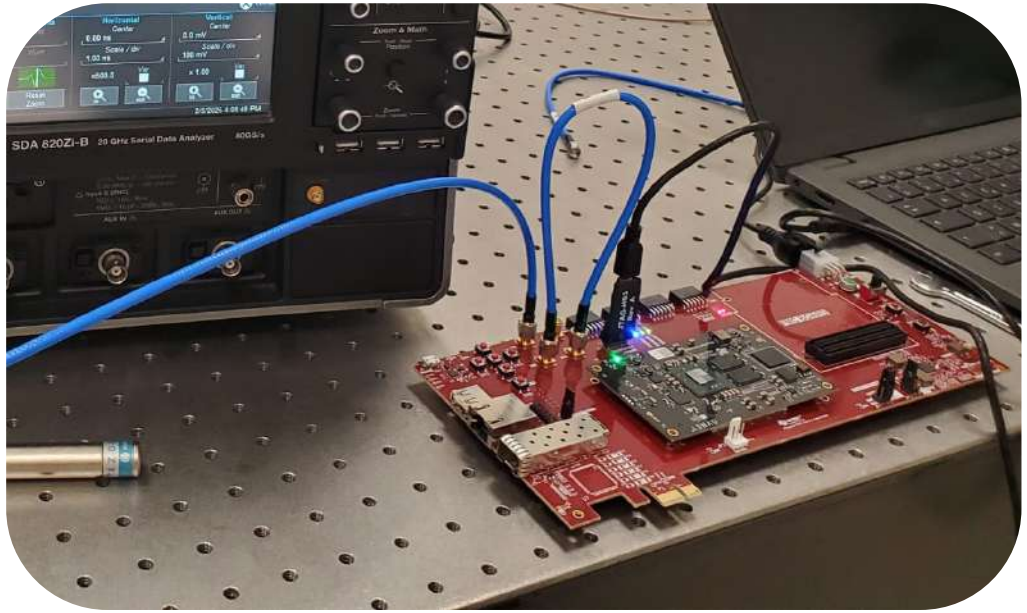

The revolution has started: be part of it

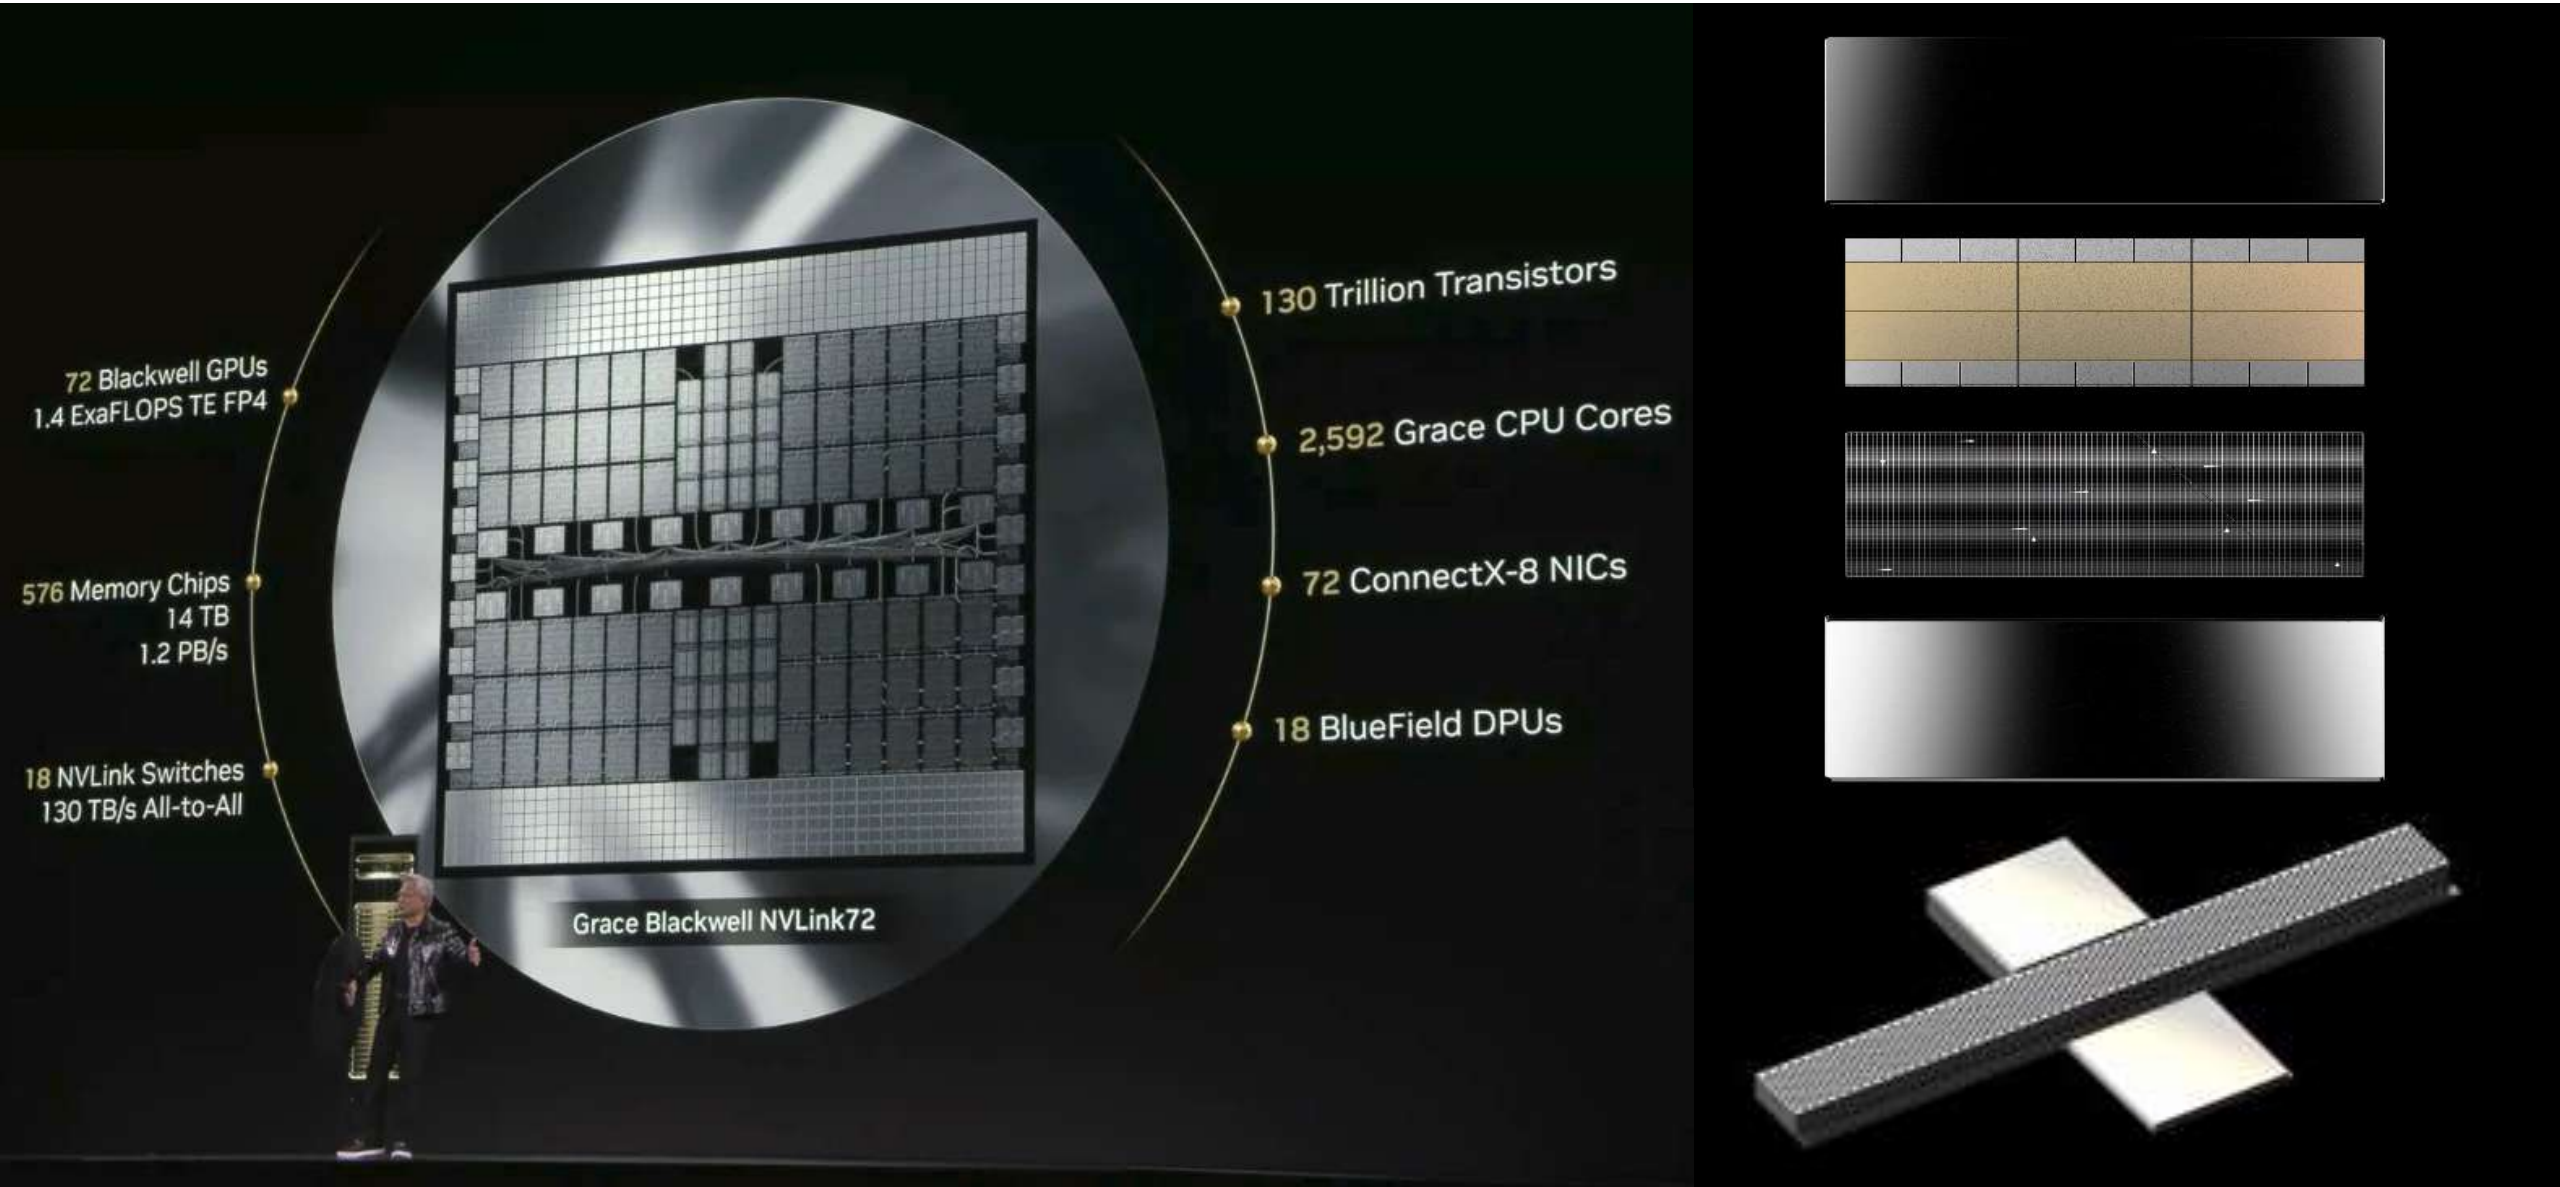

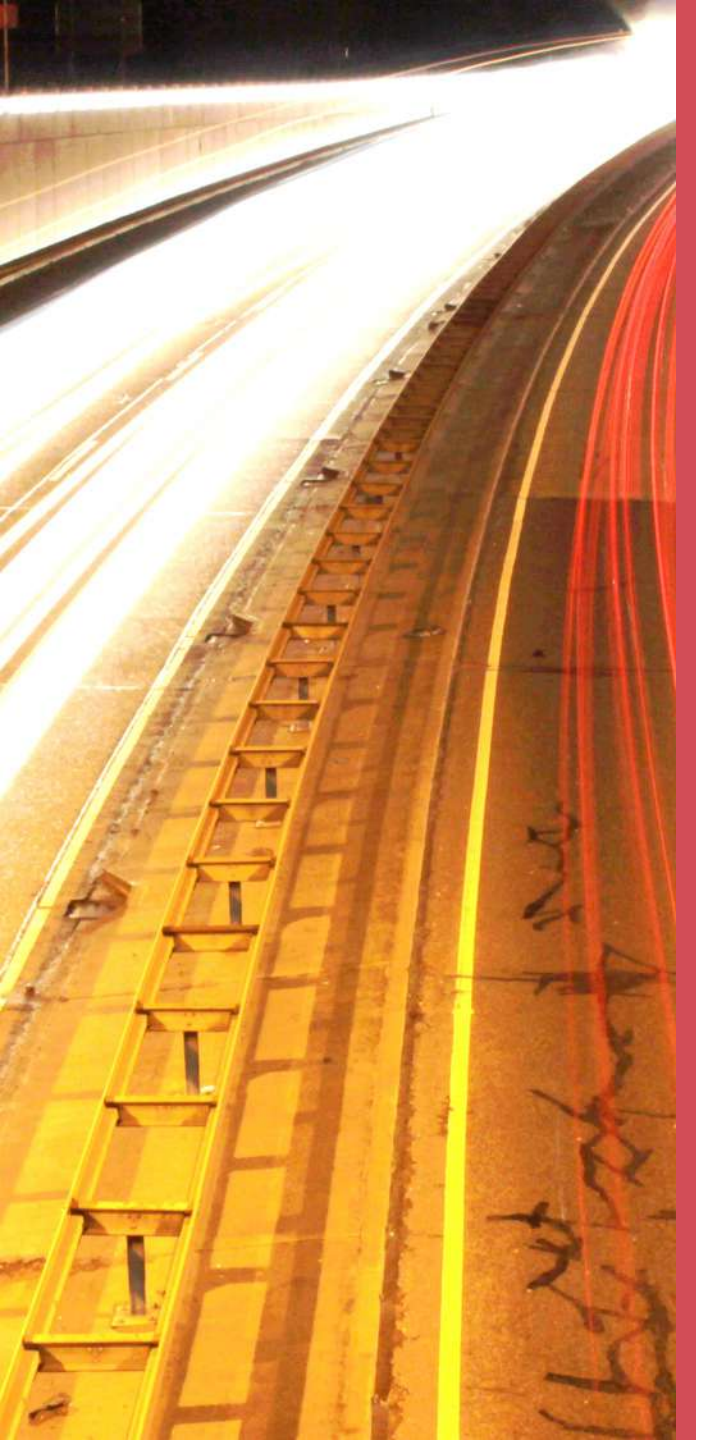

NcodiN

Illuminating the future of computing

Francesco Manegatti

francesco.manegatti@ncodin.com | +33 7 69 08 61 02

www.ncodin.com