Advanced photonic technologies for future systems:
Optical computing or optics for computing?

Bert Jan Offrein

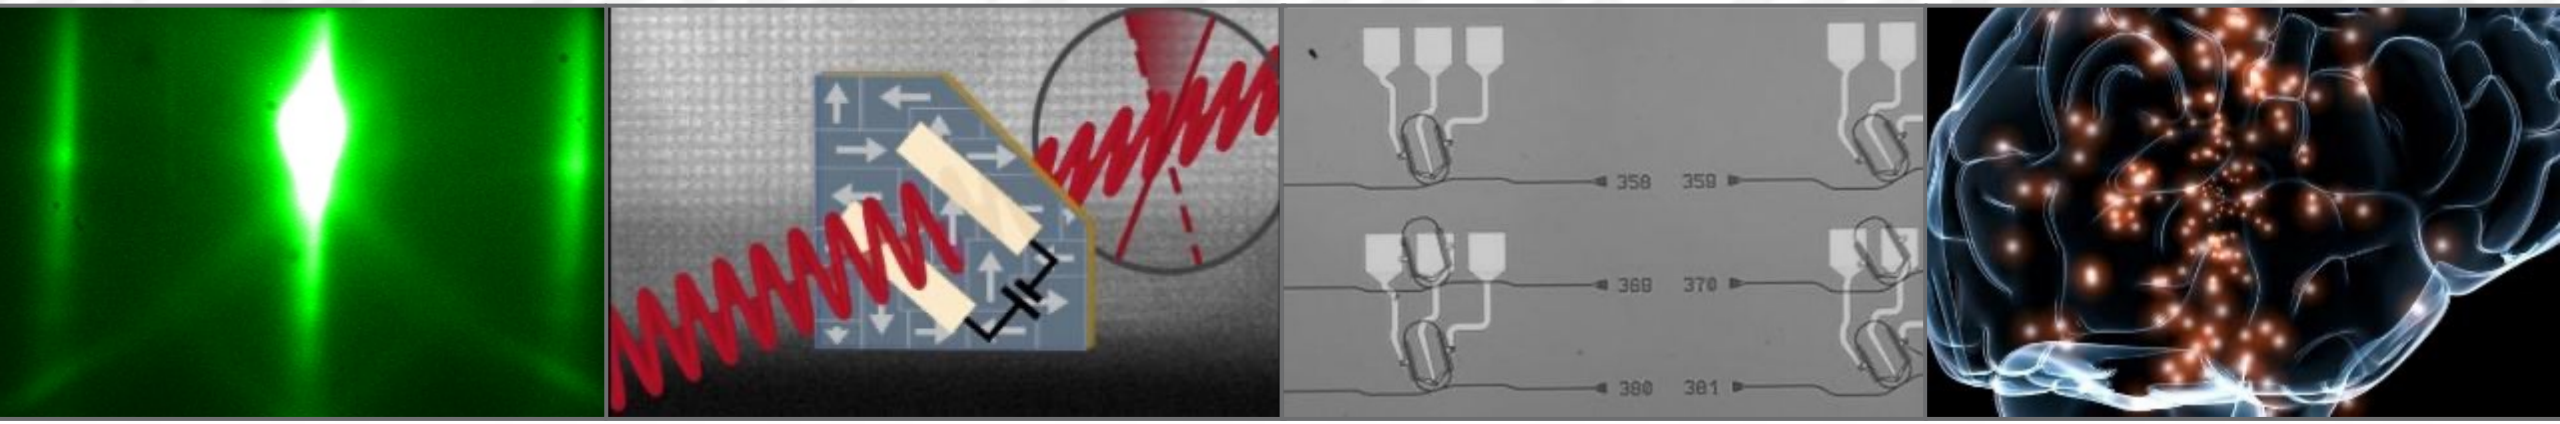

IBM Research Europe - Zurich, Rüschlikon, Switzerland

University of Twente, Enschede, The Netherlands

Co-Packaged Optics and AI Systems

The first petaflop supercomputer - Roadrunner

- Breaking the petaflop barrier
- In 2008, IBM built ‚**Roadrunner**‘, the first petaflop supercomputing system
- Why is this of interest now?

https://www.ibm.com/history/petaflop-barrier

The neural network explosion

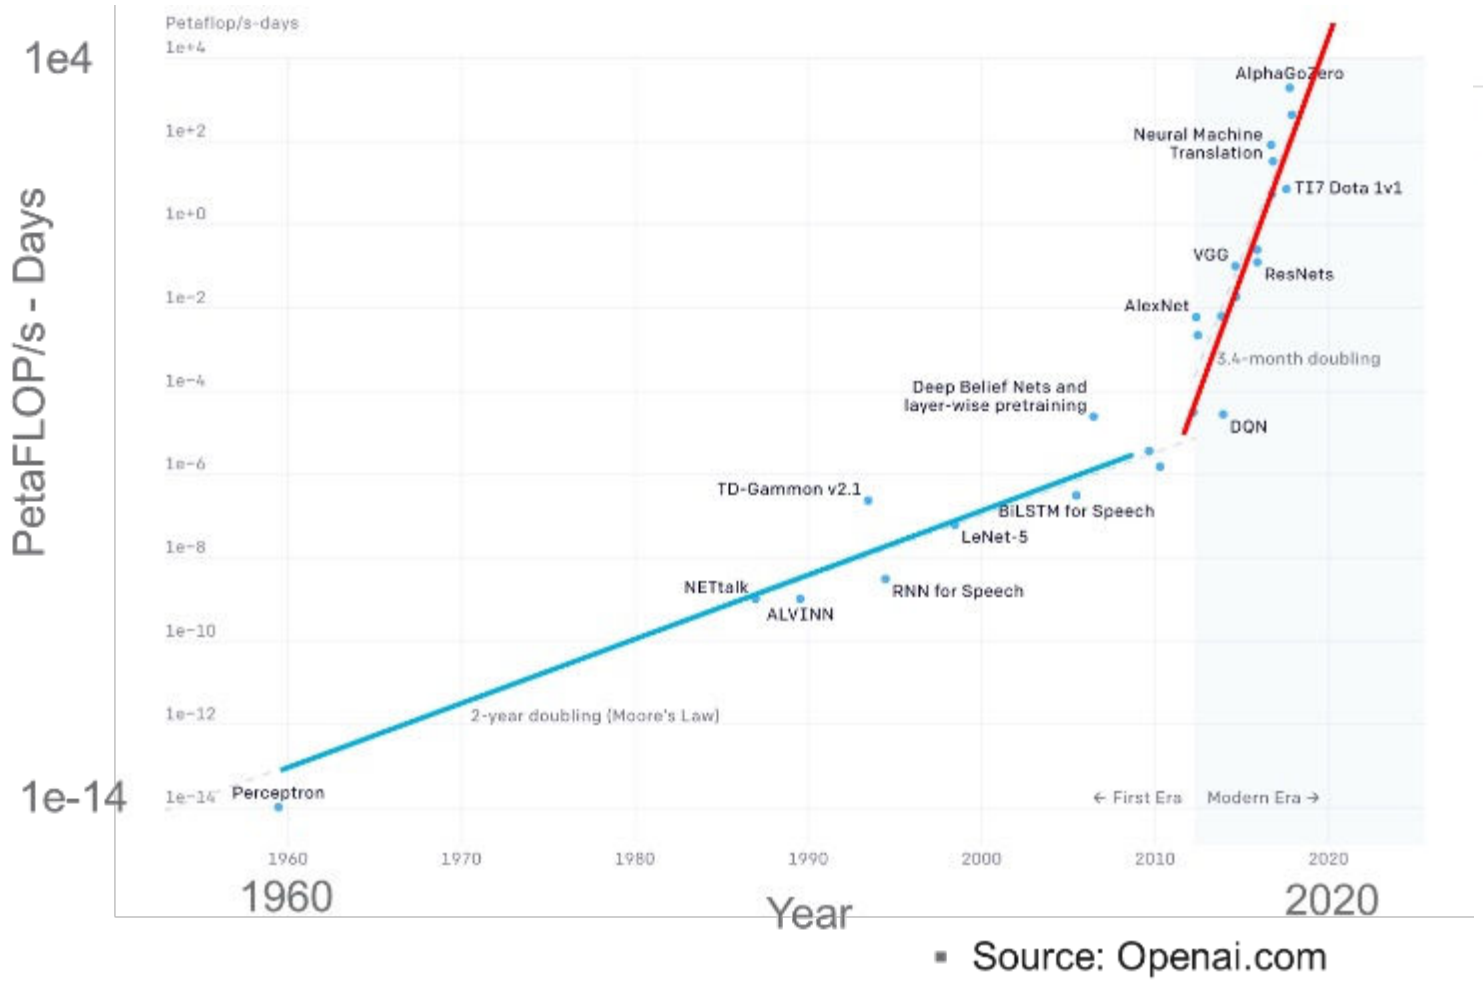

- Source: Openai.com

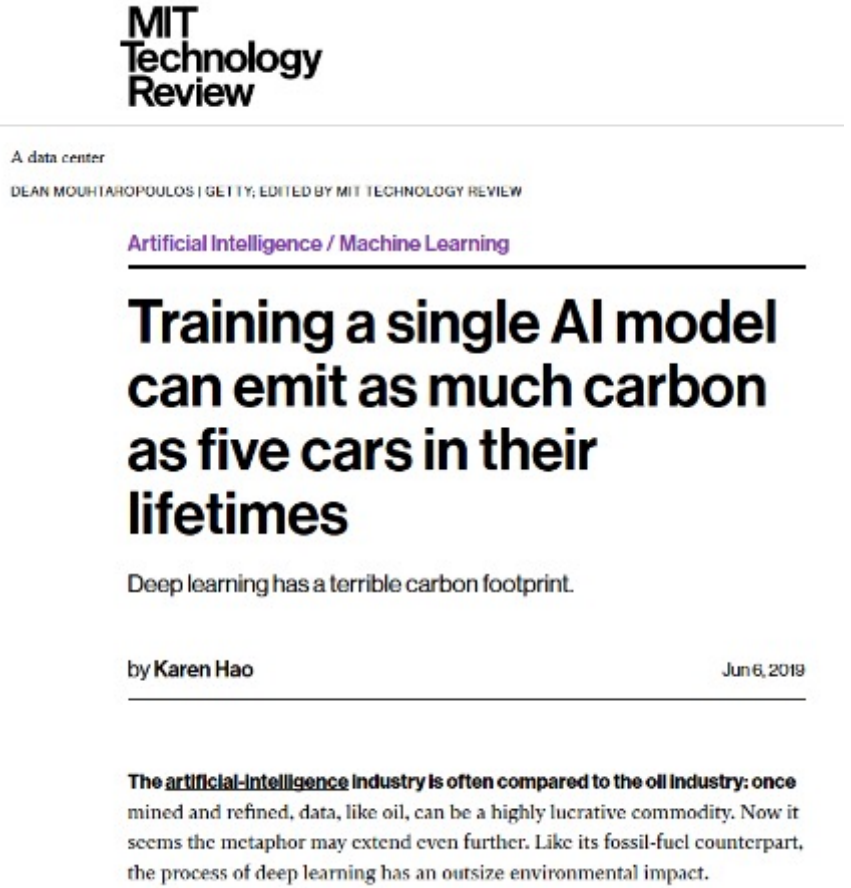

MIT Technology Review

A data center

DEAN MOUHTAROPOULOS | GETTY; EDITED BY MIT TECHNOLOGY REVIEW

Artificial Intelligence / Machine Learning

Training a single AI model can emit as much carbon as five cars in their lifetimes

Deep learning has a terrible carbon footprint.

by Karen Hao Jun 6, 2019

The artificial-intelligence industry is often compared to the oil industry: once mined and refined, data, like oil, can be a highly lucrative commodity. Now it seems the metaphor may extend even further. Like its fossil-fuel counterpart, the process of deep learning has an outsize environmental impact.

E. Strubell et al., arXiv:1906.02243

Signal processing in neuromorphic computing

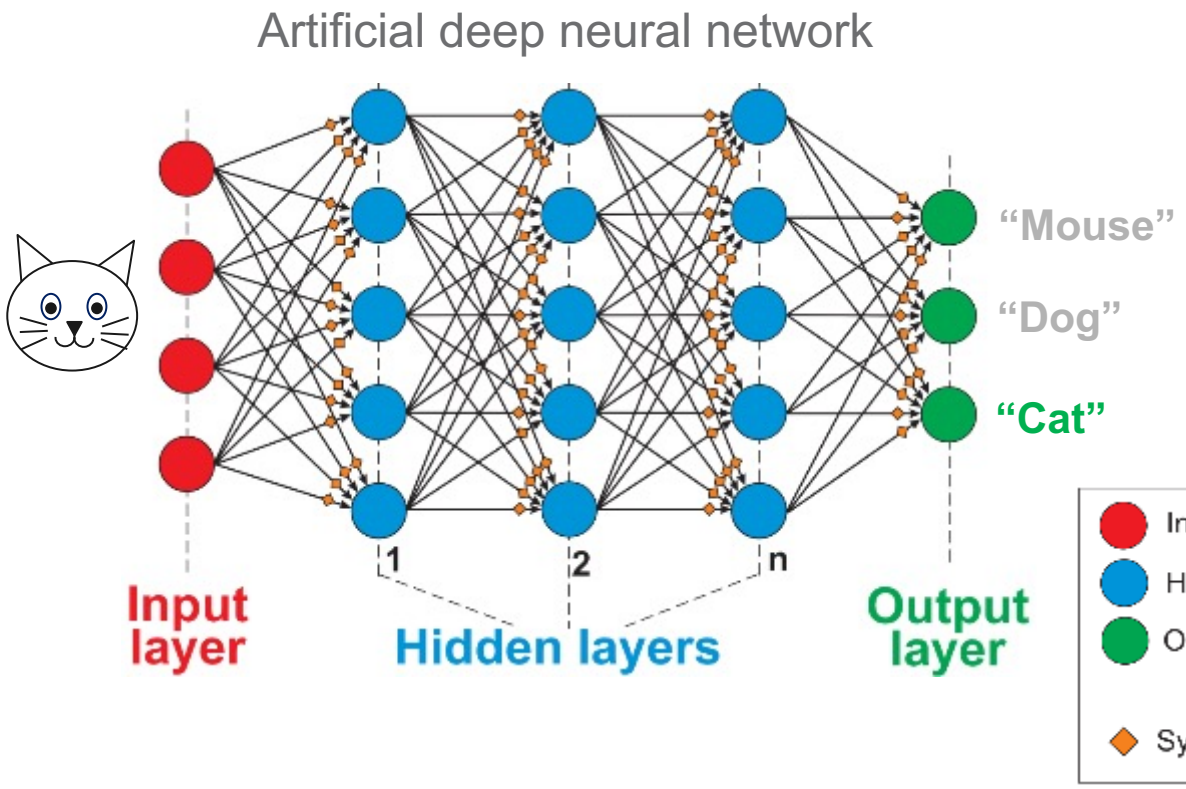

$$\begin{pmatrix} w_{11} & w_{12} & \cdots & w_{1n} \\ w_{21} & w_{22} & \cdots & w_{2n} \\ & & \cdots & \\ w_{j1} & w_{j2} & \cdots & w_{jn} \\ & & \cdots & \\ w_{mn} & w_{mn} & \cdots & w_{mn} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ \cdot \\ \cdot \\ \cdot \\ x_n \end{pmatrix}$$

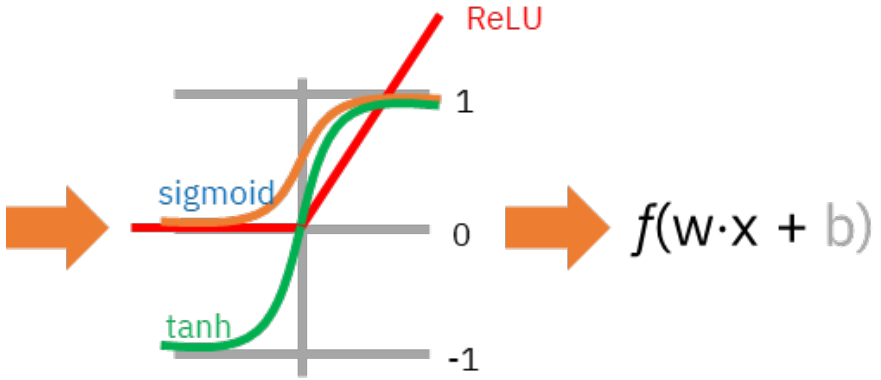

Synaptic function: Multiply accumulate → Vector matrix multiplication → $O(N^2)$
Neuron: Nonlinear activation → $O(N)$

The computing hardware

Transistors

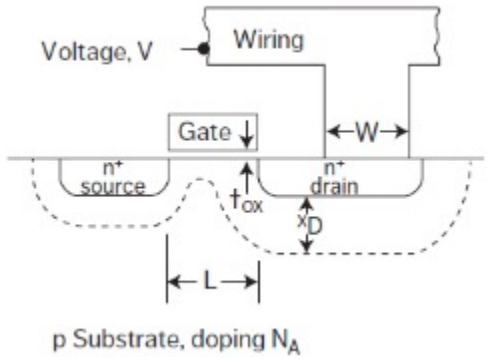

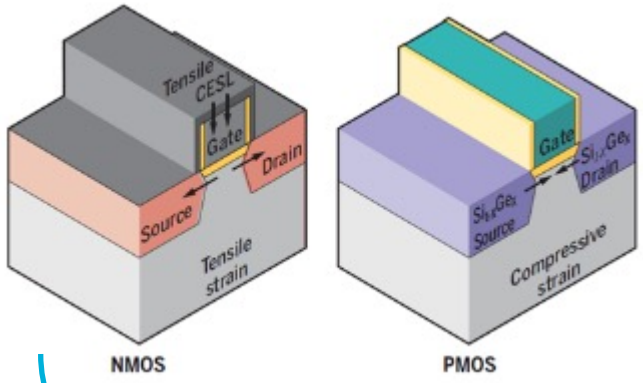

Logic & SRAM

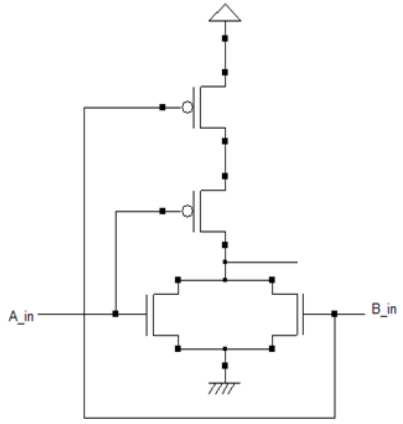

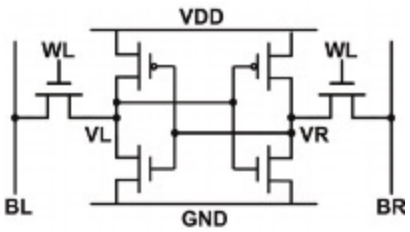

Circuits

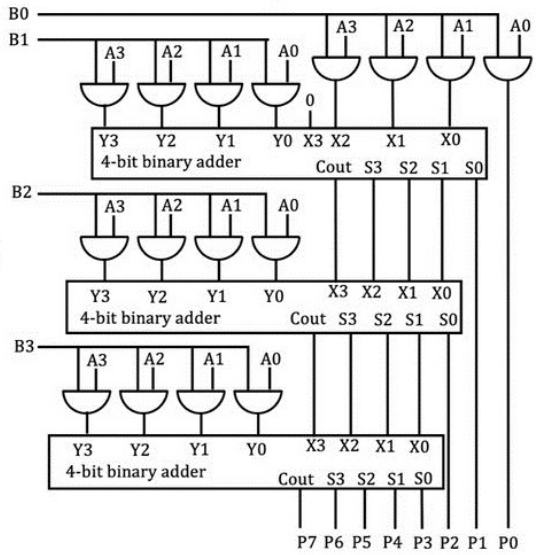

Memory & Storage

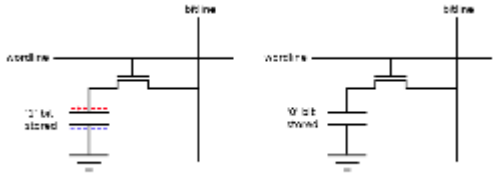

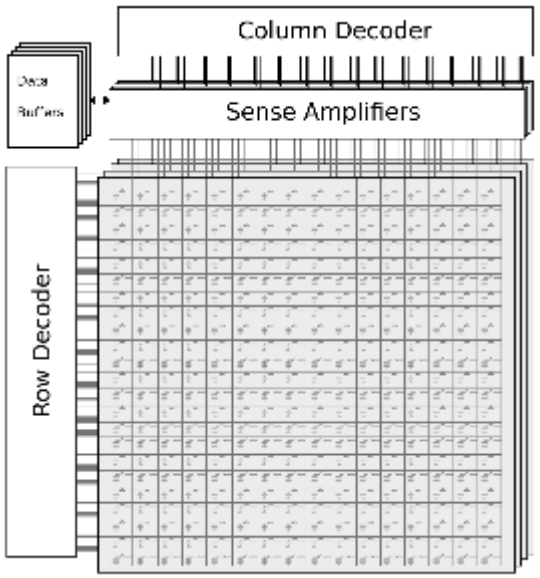

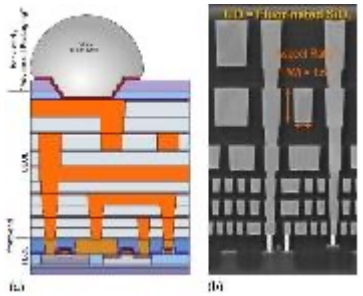

Technology scaling is continuing

Smaller & more transistors

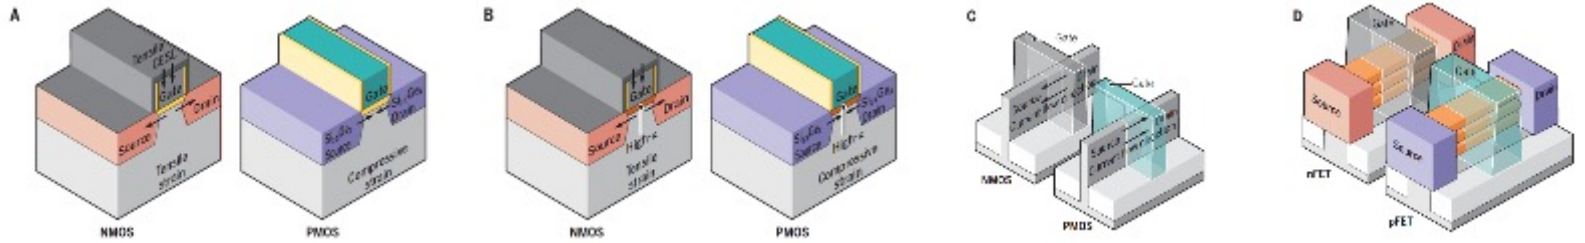

Datta et al., Science **378**, 733–740 (2022)

More & closer memory

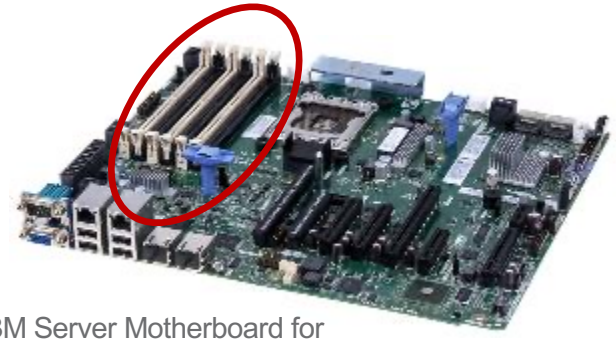

IBM Server Motherboard for X3300 M4 Laptech - The IT

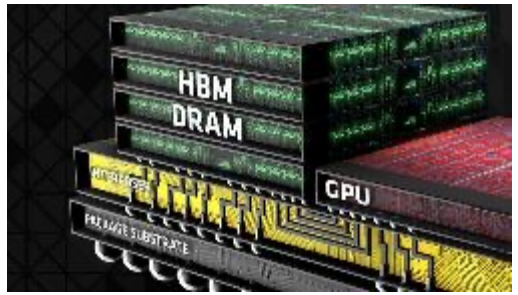

From PCPER.com

HBM memory on an AMD Radeon R9 GPU package

Higher bandwidth connectivity

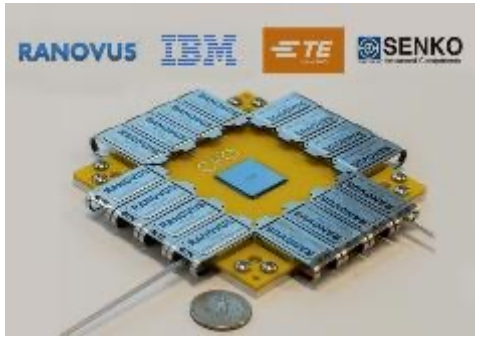

https://ranovus.com/

Technology scaling is continuing

IBM

IBM Newsroom | Press Releases | Blog | Latest News | Media Contacts | Media Center | Leadership | Global Newsrooms | About | Subscribe

IBM Unveils World's First 2 Nanometer Chip Technology, Opening a New Frontier for Semiconductors

New chip milestone to propel major leaps forward in performance and energy efficiency

May 6, 2021

IBM and Samsung Unveil Semiconductor Breakthrough That Defies Conventional Design

Dec 14, 2021

IBM and Rapidus Form Strategic Partnership to Build Advanced Semiconductor Technology and Ecosystem in Japan

Rapidus, a Newly-Formed Advanced Logic Foundry, Will Leverage IBM's Semiconductor R&D Leadership, Including 2 Nanometer Node Technology

Dec 12, 2022

Rapidus and IBM Expand Collaboration to Chiplet Packaging Technology for 2nm-Generation Semiconductors

Agreement builds on existing collaboration between the two companies for the joint development of 2nm node technology

Jun 3, 2024

IBM Brings the Speed of Light to the Generative AI Era with Optics Breakthrough

New co-packaged optics innovation could replace electrical interconnects in data centers to offer significant improvements in speed and energy efficiency for AI and other computing applications

Dec 9, 2024

Looking back – 2011!

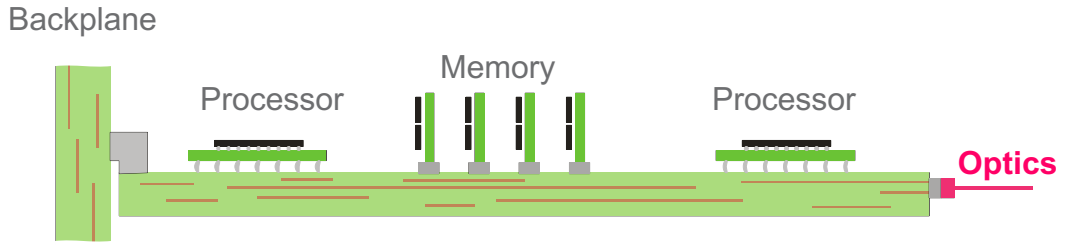

2008

Electrical system, optical fibers at card edge only

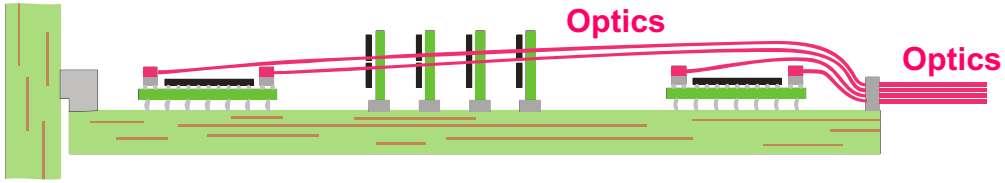

Today (2011)

Optical fibers across the boards

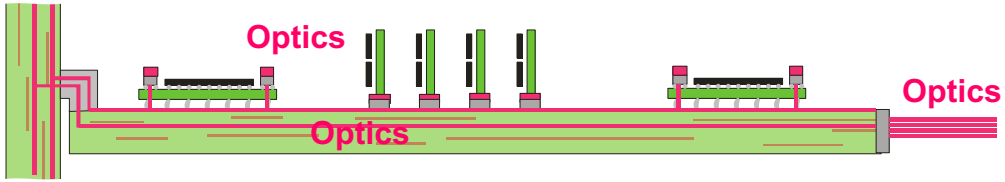

Optical waveguides in/on boards

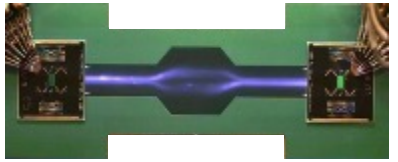

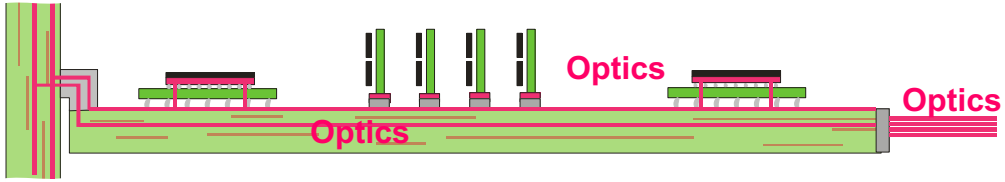

Optical interconnects integrated with the processor

Development

Research

Optical interconnects will be applied for shorter and shorter links to fulfill bandwidth and power efficiency requirements. Integration will increase bandwidth density and reduce cost.

2011- IBM Power 775

- One drawer = one node = 8 “Octants”
- One switch “Hub” per octant
- Four nodes = One Super Node
- All-to-all connectivity between the 32 octants in one Super Node
- All-to-all connectivity between 512 Super Nodes → Run very large simulations
- Co-packaged optics: 56 Avago “Micro-pods” per switch Hub
- **Pluggable: Electrically and optically!**

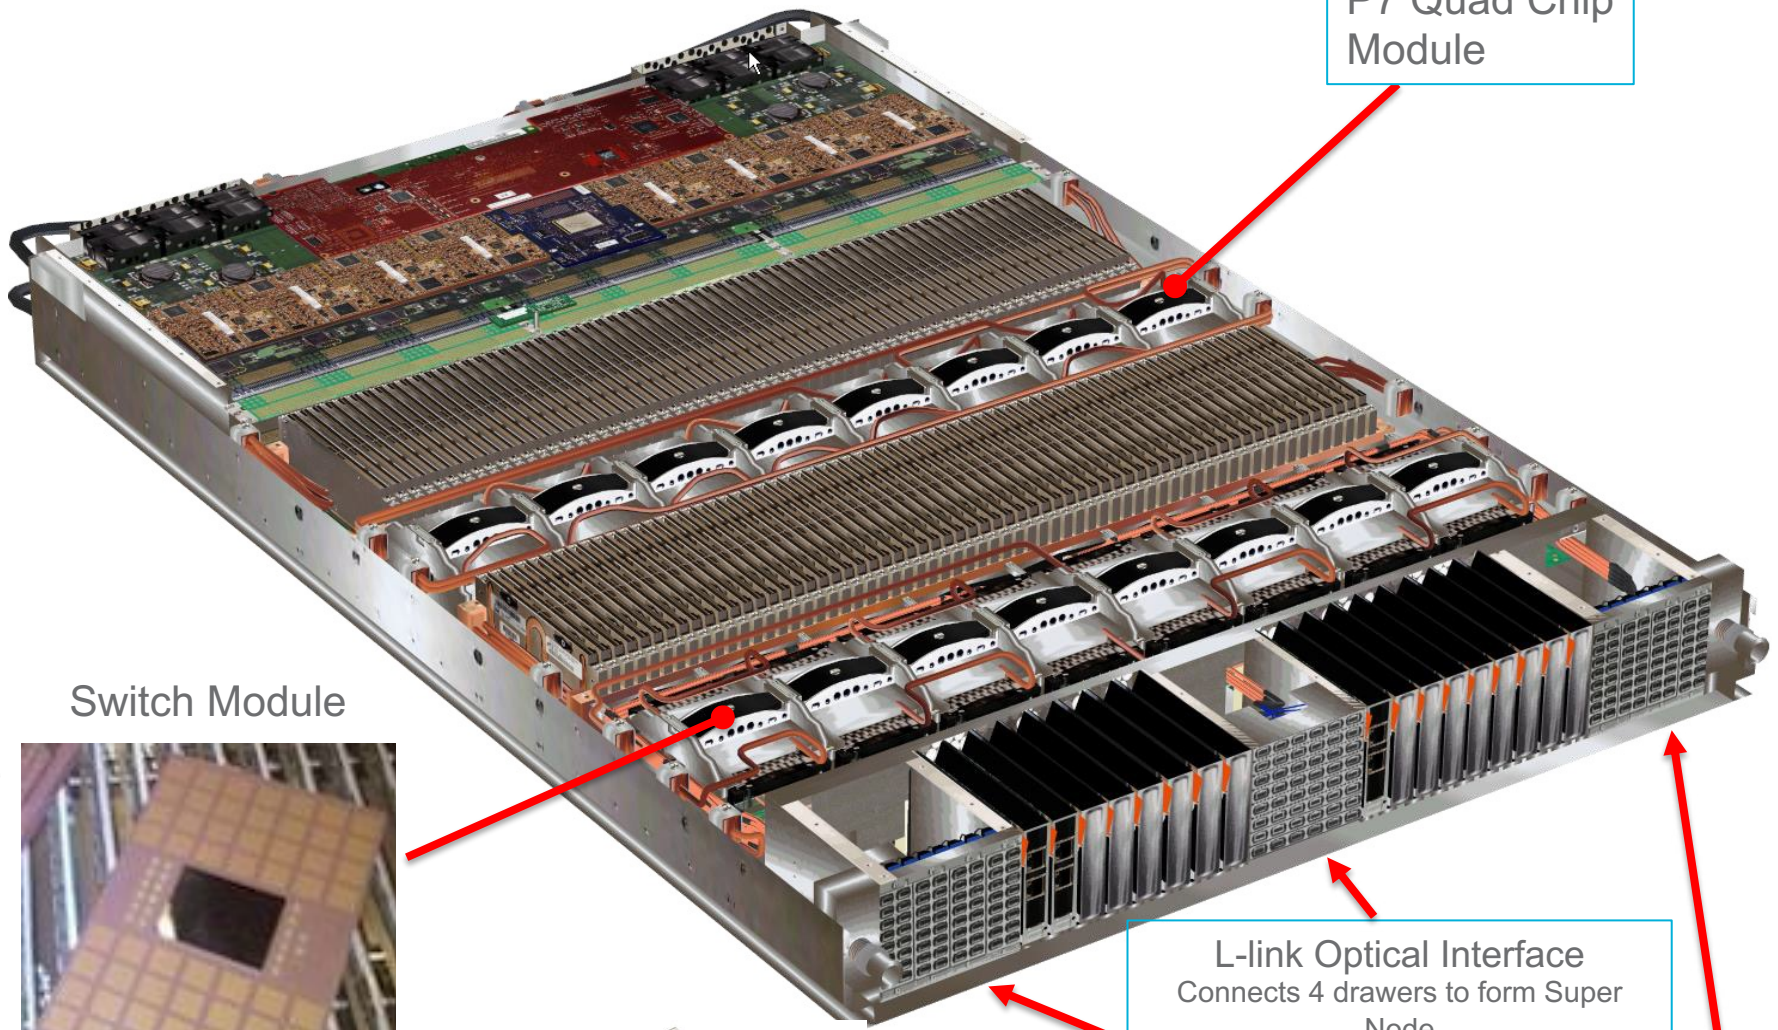

Looking back – 2011!

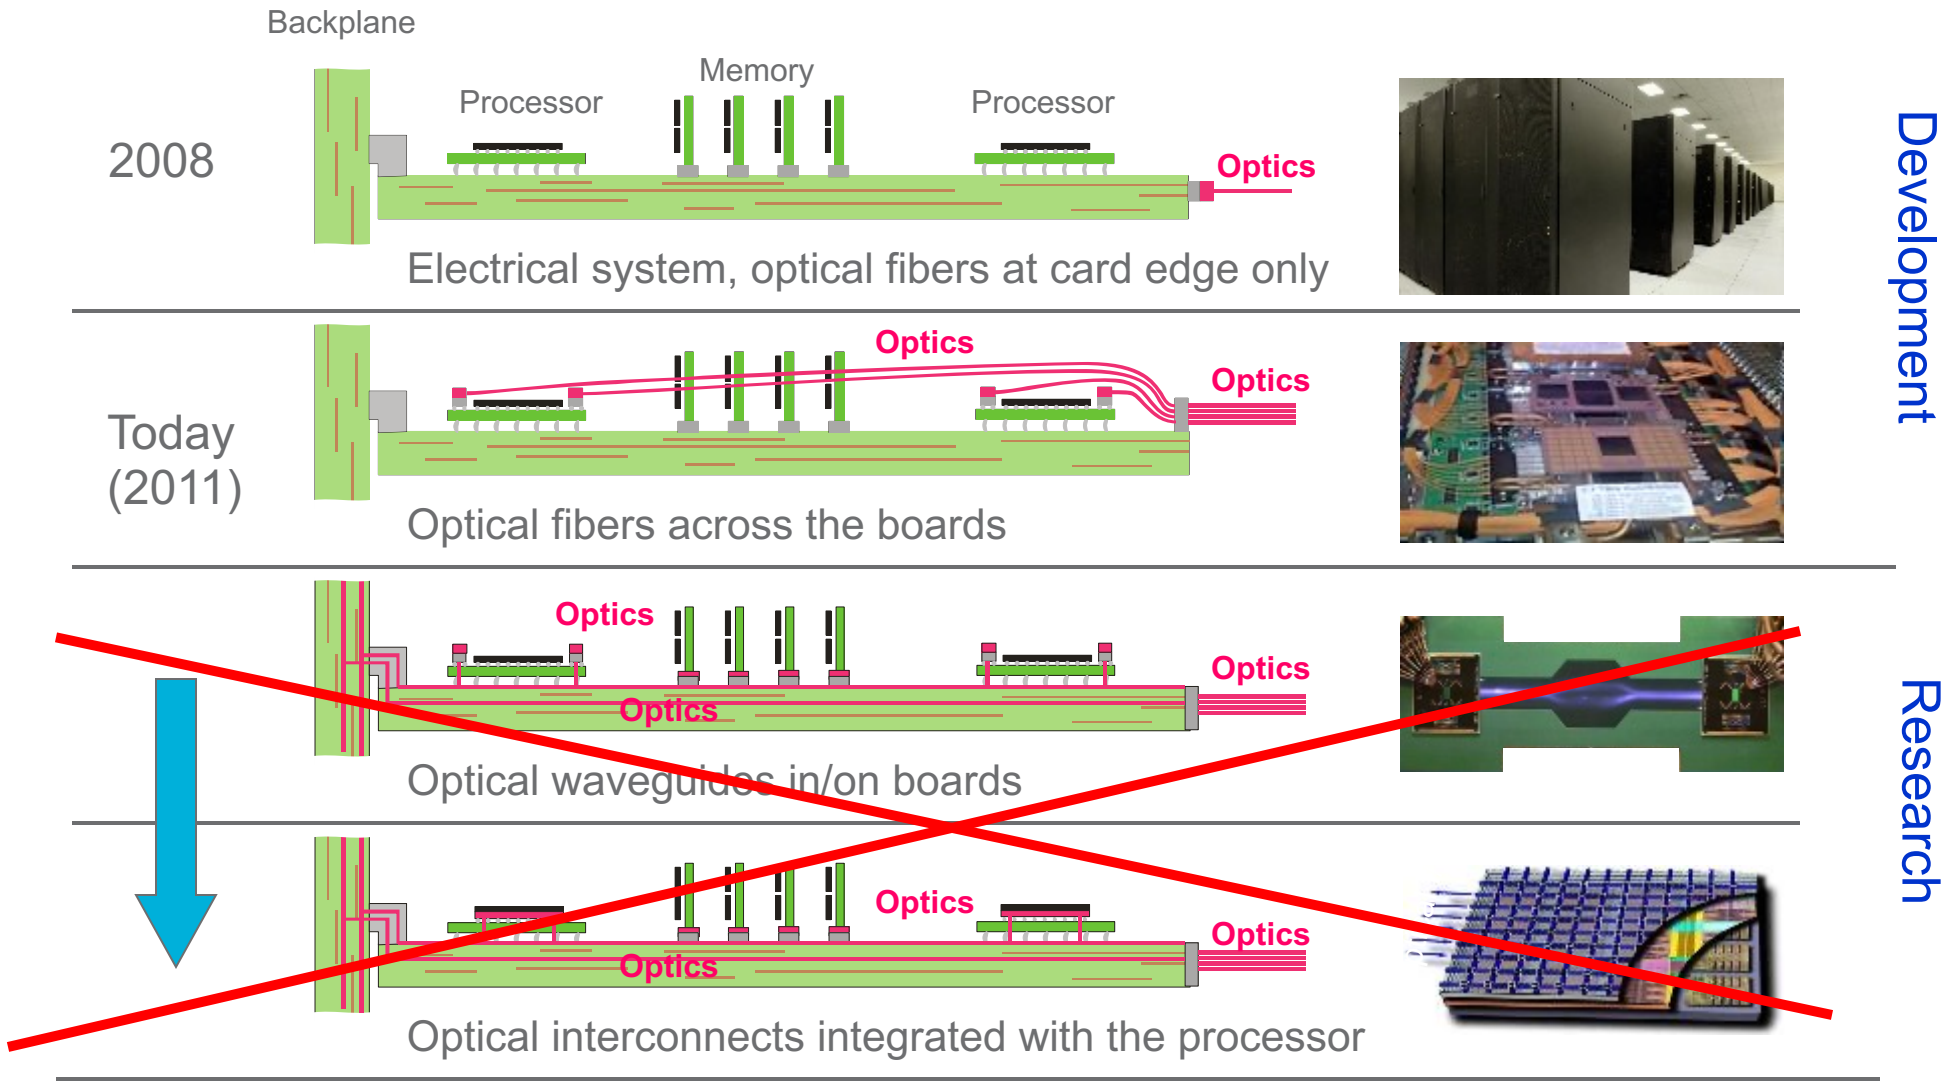

Optical interconnects will be applied for shorter and shorter links to fulfill bandwidth and power efficiency requirements. Integration will increase bandwidth density and reduce cost.

Electrical & Optical Communication

Electrical

Optical

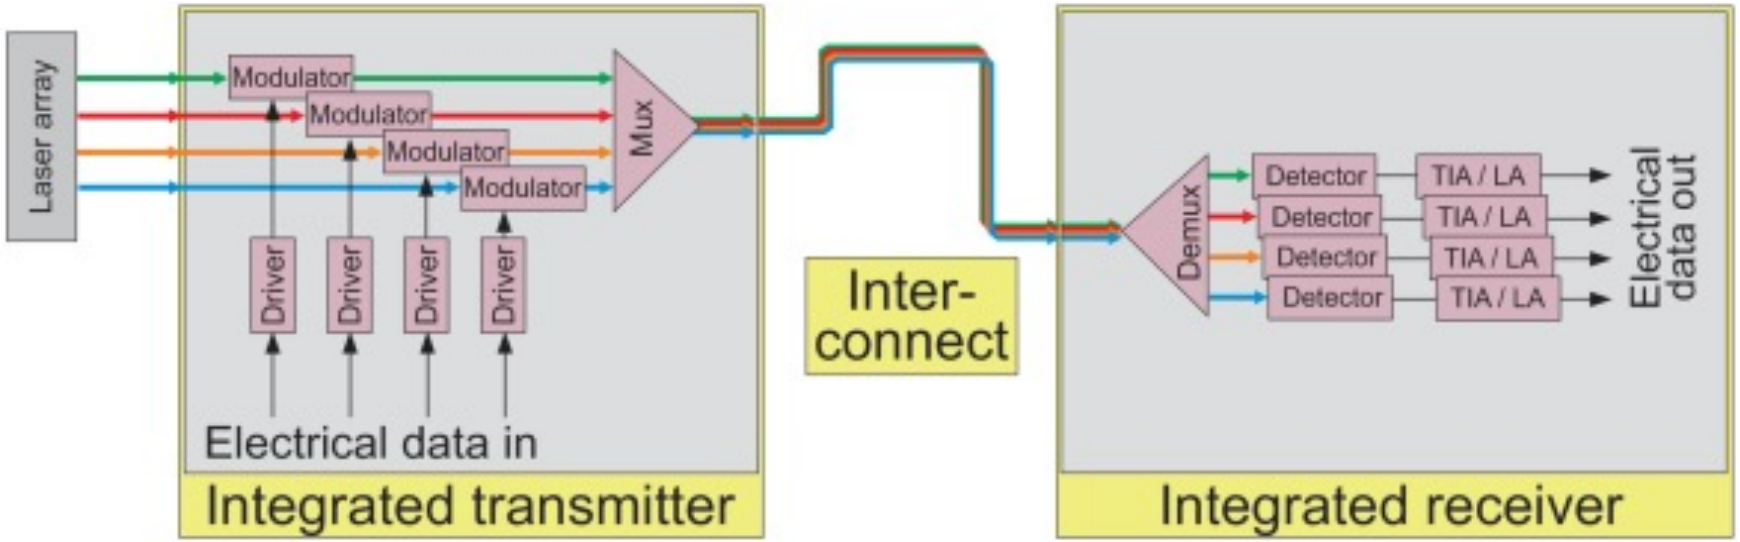

Optical interconnects enable high-bandwidth and long-distance communication
But comes with a component and assembly overhead

Scaling AI compute performance and efficiency

- **System trends**
  - More compute & memory
    - Transistor scaling
    - Dense memory
  - Dedicated optimized functions
    - Chiplets, accelerators, memory
  - Tight integration of all functions
    - Heterogeneous integration
    - Silicon interposer

- **Scaling challenges**
  - Bandwidth scaling << Compute FLOPS
    - Enhanced interest in optical interconnects
    - Driven by the data bottleneck and energy efficiency
  - Close integration of optics with ASIC for energy efficiency and bandwidth requirements

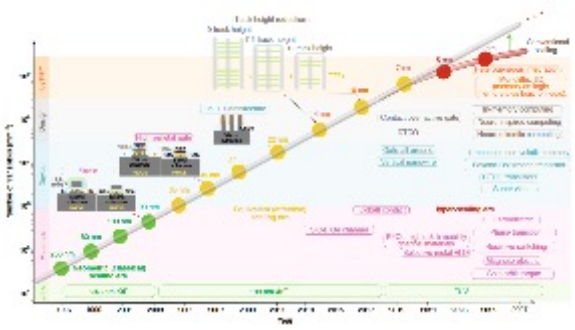

Transistor scaling
From: Datta e.a. 2018

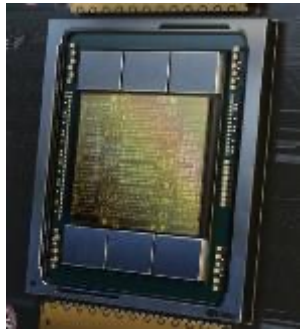

HBM on substrate
From: NVIDIA 2023

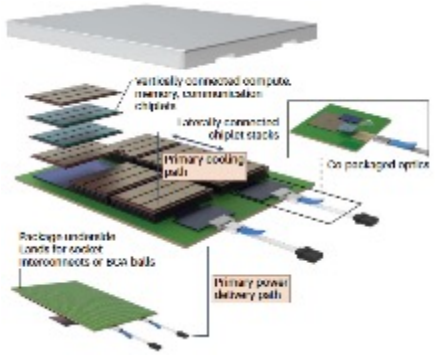

Chiplet assembly vision
From: Intel 2024

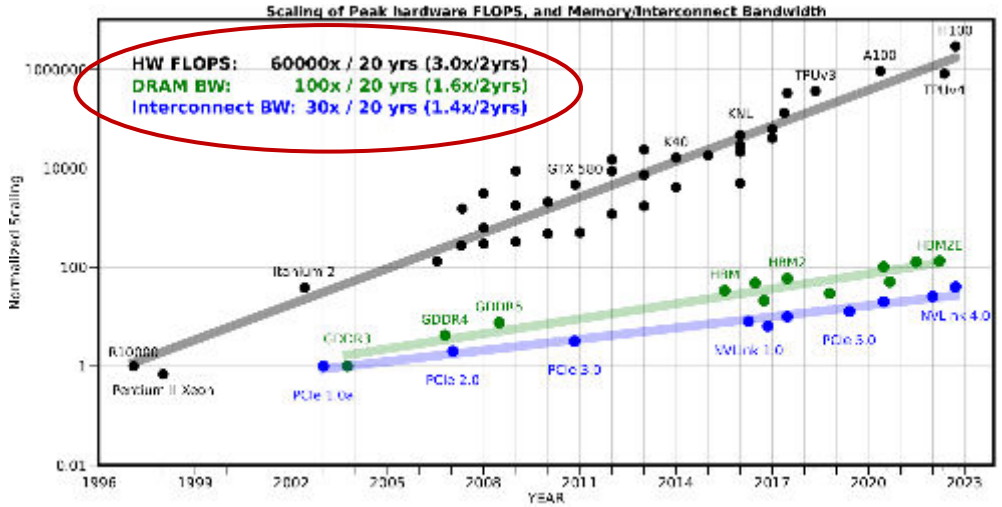

From: https://medium.com/riselab/ai-and-memory-wall-2cb4265cb0b8

Looking back – 2011!

2008

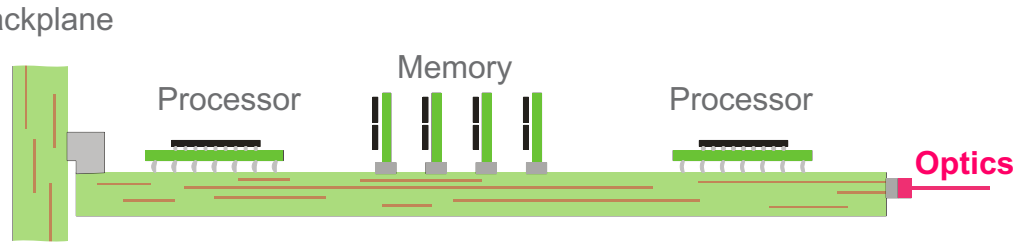

Electrical system, optical fibers at card edge only

2011

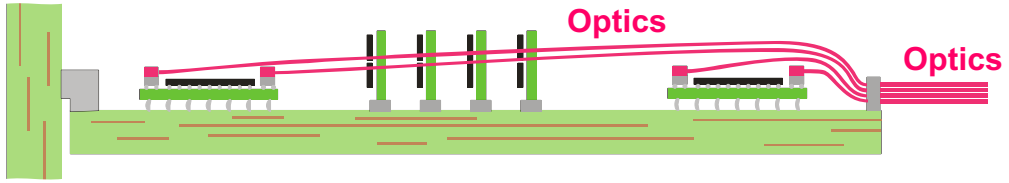

Optical fibers across the boards

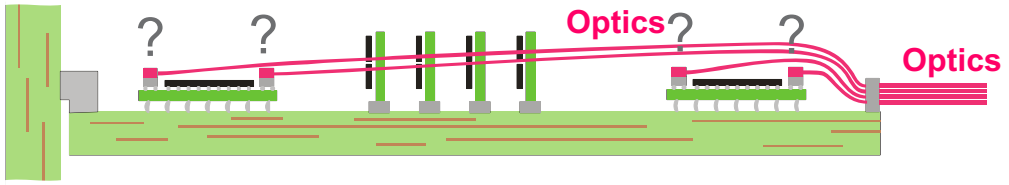

Integrated photonic chiplets

?

Overcome the overhead of optics

Development

Research

Chip-to-fiber coupling challenge

- Challenge:

on-chip SOI waveguide

0.5 x 0.2μm

spot-size conversion

SM fiber core

9μm

Tapered SOI waveguide

Polymer waveguide core

Light coupling

Taper input

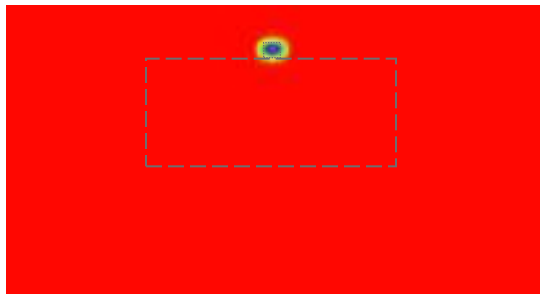

Complete light confinement in SOI waveguide

Taper center

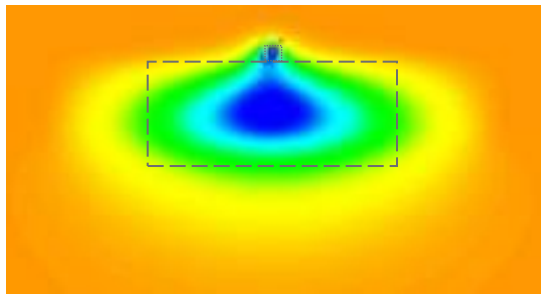

Light is confined in both waveguides

Taper output

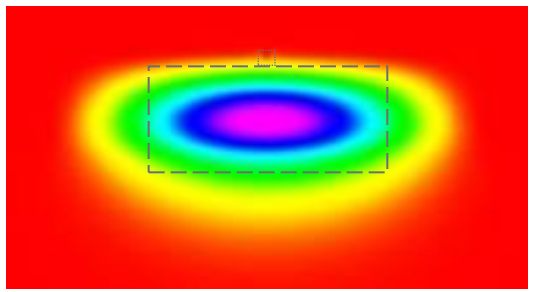

Complete light confinement in polymer waveguide

Full Module

Laminate
- 150 um pitch to PIC
- Electrical interconnection
- Micro-BGAs, 500 µm pitch

Polymer waveguide
- Pitch conversion 50 to 250 µm

Photonic circuit
- Bumping pads 150 µm pitch
- Waveguides
- 8 x10mm^2

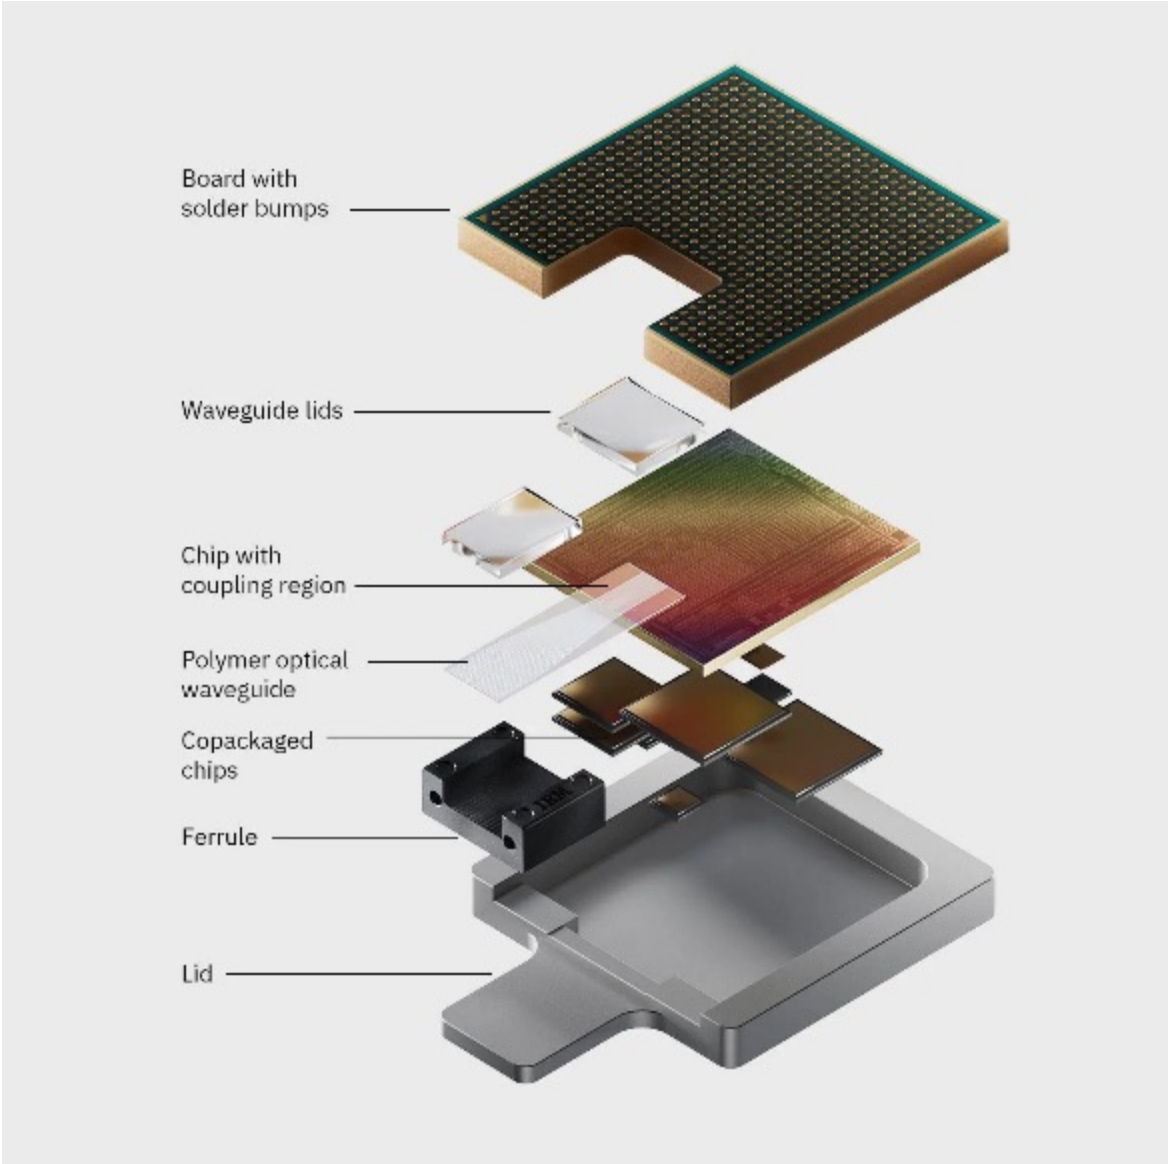

Module Assembly

PWG first

- Coupling area protected
- Optical performance monitor during build
- Square laminate

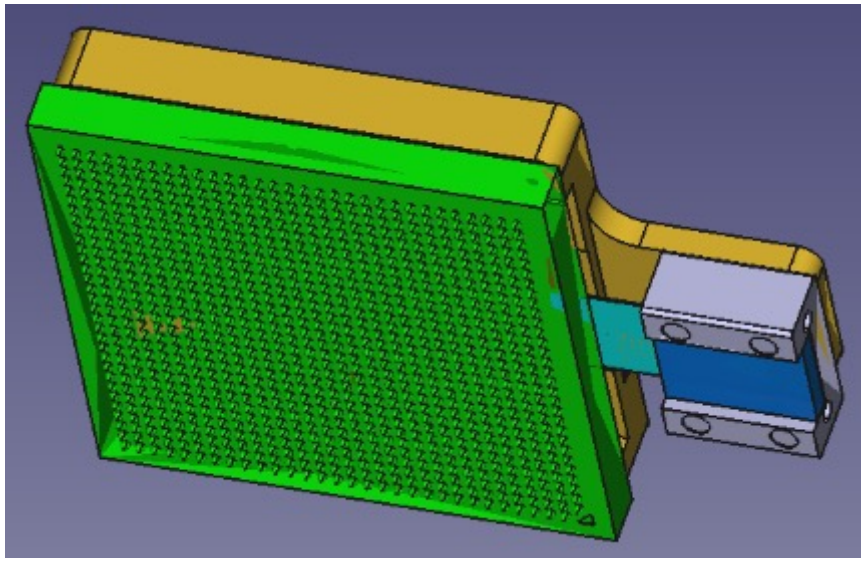

PWG last

- Easy handling during build
- Avoid chip to laminate join temperature excursion

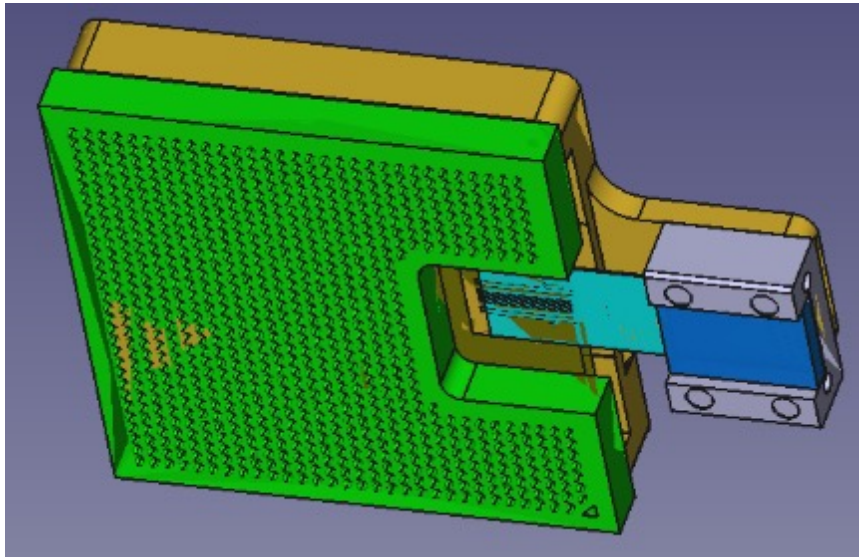

- 17x17mm^2 square shape
- 7 mm protrude length

- Reflow-compatible
- Adiabatic coupling PIC to PWG
- O-band

Typical insertion loss spectra

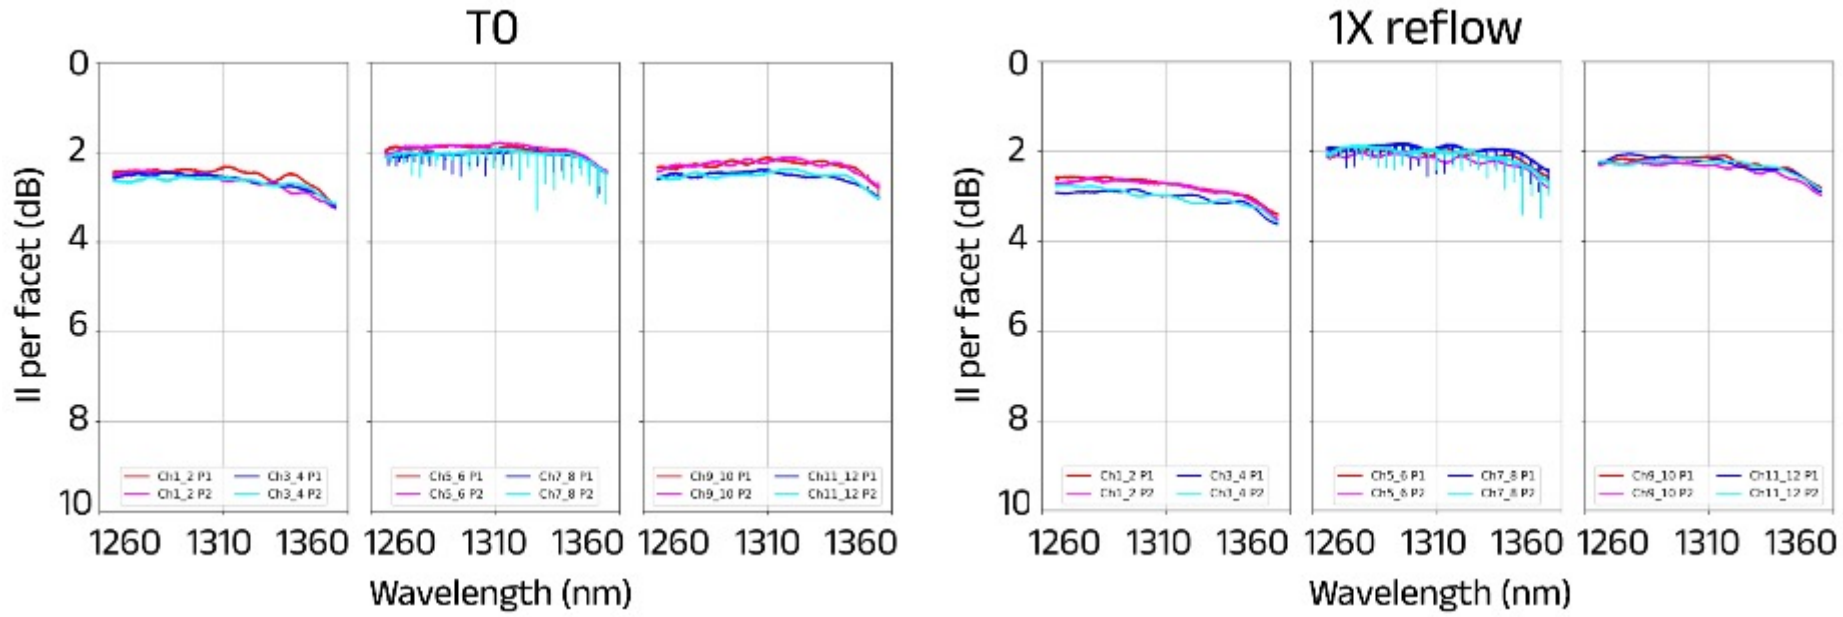

O-band insertion loss results:

- 6 loopbacks in SiN waveguides
- Resonant ring as polarization reference
- Loss per facet, includes fiber-PWG as well as PWG-PIC interfaces
- Loss as low as 1.2 dB at 1310 nm also observed

Digital signal processing

- The Von Neumann architecture
  - Memory for programs and data, a bus for memory access, an arithmetic unit & a program control unit

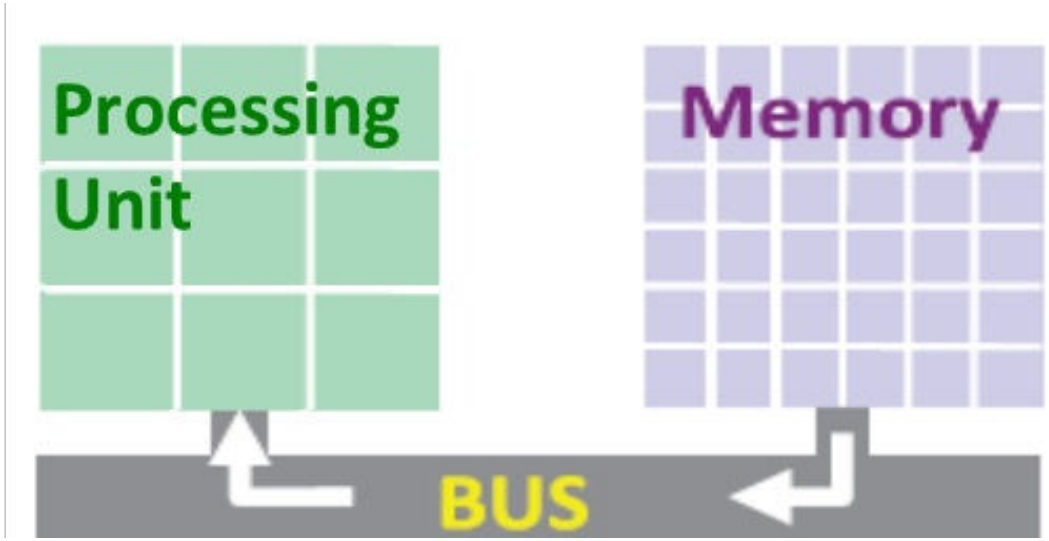

Analog signal processing for scalability

- **Limiting factors**
  - Memory access
  - Sequential operations
  - Digital signal processing

- **Overcome by**
  - In-memory computing
  - Parallel operations
  - Analog signal processing

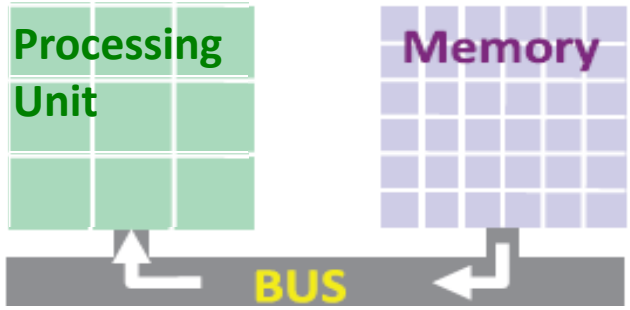

Compute effort ~O(#Neurons^2)

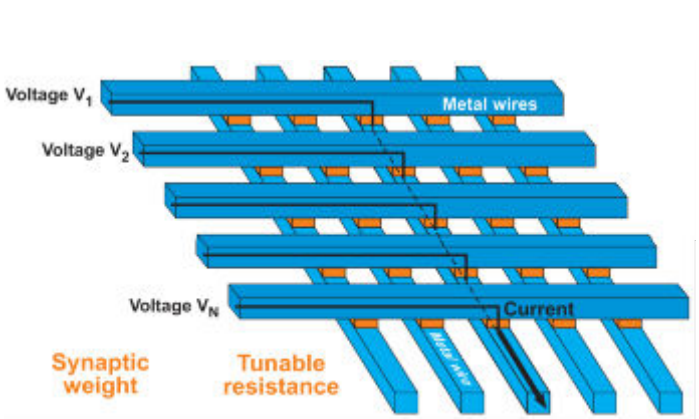

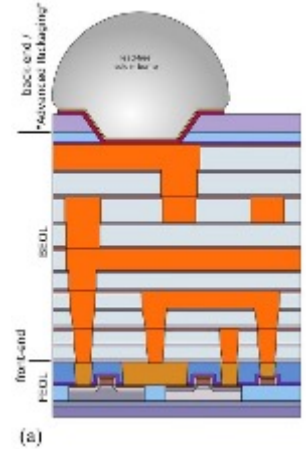

Compute effort ~O(1)

Electrical and optical solutions are viable candidates

Analog signal processing systems

- Electrical

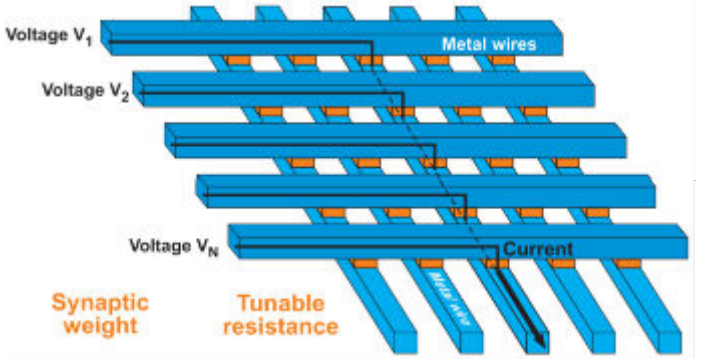

Ohm's and Kirchhoff's law

- Integrated photonics

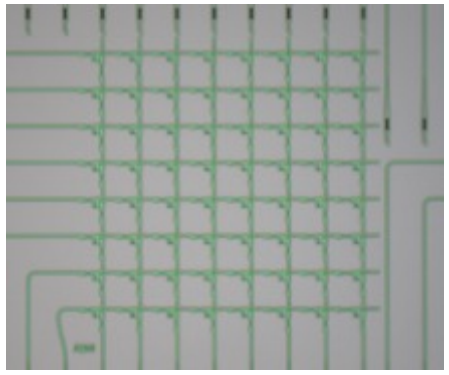

From: EU PHOENICS (U Oxford).

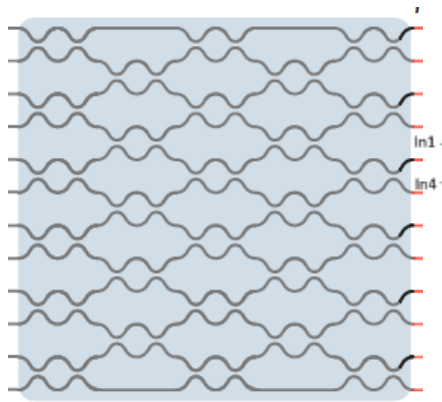

From: Y. Shen et al., doi: 10.1038/nphoton.2017.93.

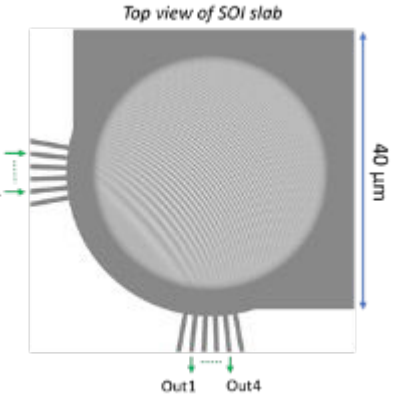

From: F. Horst, IBM.

Attenuation, interference, diffraction

- Memristive devices in a crossbar
  - PCM
  - OxRAM
  - FERAM
- Various device concepts and materials
  - Crossbar
  - Mach-Zehnder interferometer
  - Diffractive

For inference & training

Pipelining data

- **Neural network architecture = hardware architecture**

- **Avoid memory access**
  - ~100x better inference power-efficiency
  - Non-volatile analog weights

  - For training even more enhancement
  - Controlled weight update required

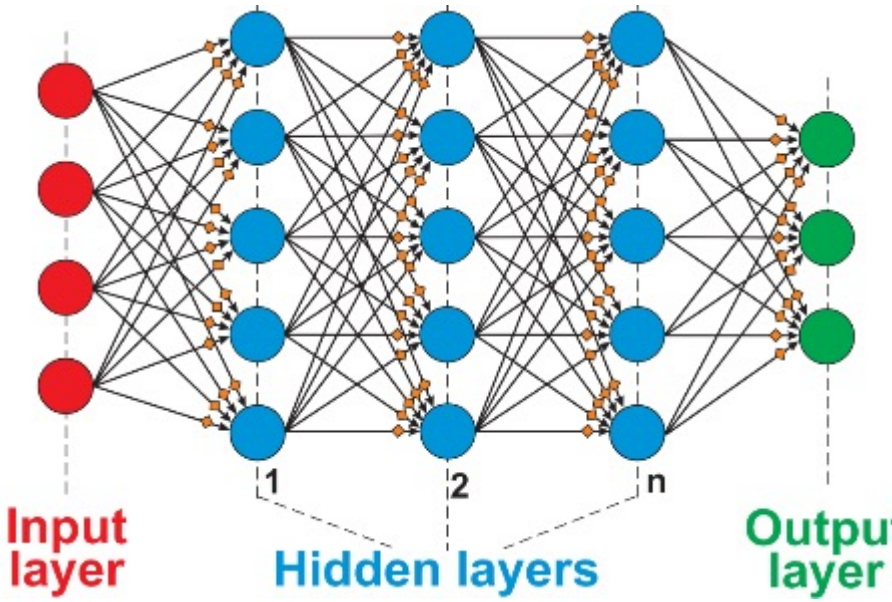

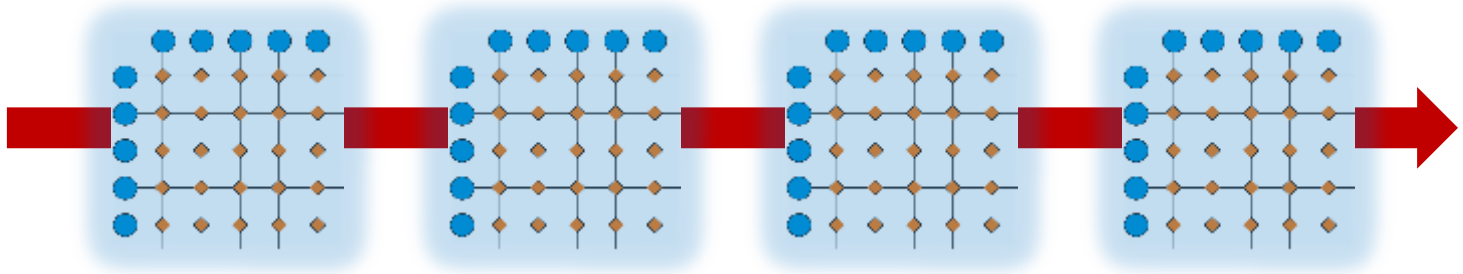

Analog signal processing systems

- Electrical

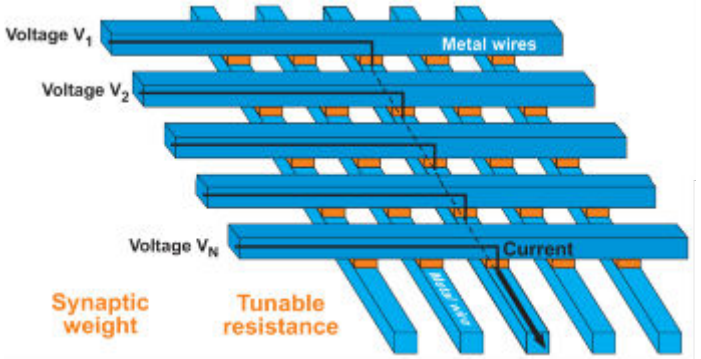

Ohm's and Kirchhoff's law

- Integrated photonics

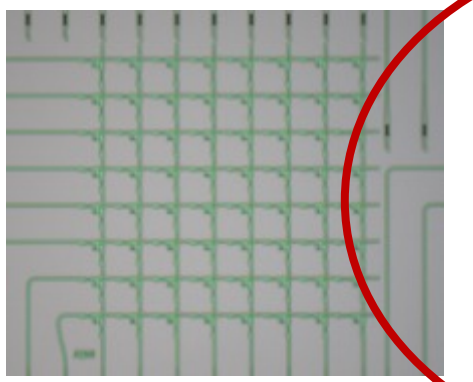

From: EU PHOENICS (U Oxford).

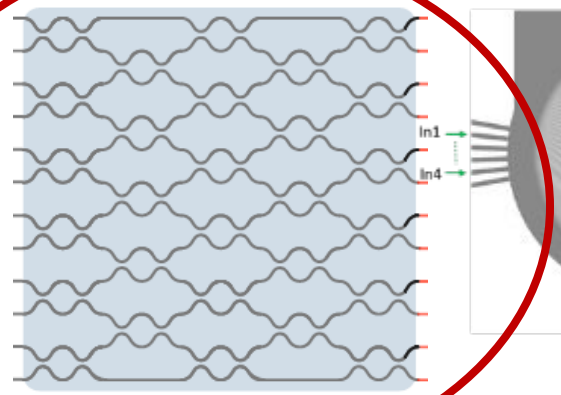

From: Y. Shen et al., doi: 10.1038/nphoton.2017.93.

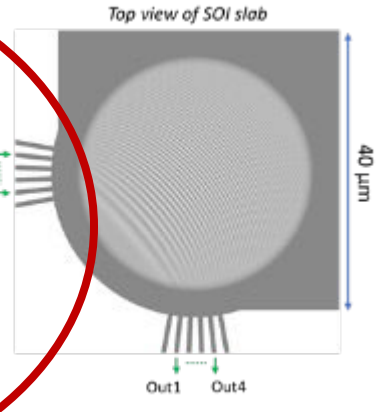

From: F. Horst, IBM.

Attenuation, interference, diffraction

- Memristive devices in a crossbar
  - PCM
  - OxRAM
  - FERAM
- Various device concepts and materials
  - Crossbar
  - Mach-Zehnder interferometer
  - Diffractive

For inference & training

Analog signal processing systems

- **Optical convolution filter** with fully **programmable** complex **kernel** elements w_x
- **Foundry** technology (SiN waveguides)
- **High-bandwidth electro-optic modulators**
- **Low-loss circuit** by daisy-chained 2-element convolution stages
- w determined by phase shifters $\{\theta_{DL}, \varphi_{MZI}\}$ (scattering matrices)

→Tune phases to perform convolution with kernel w

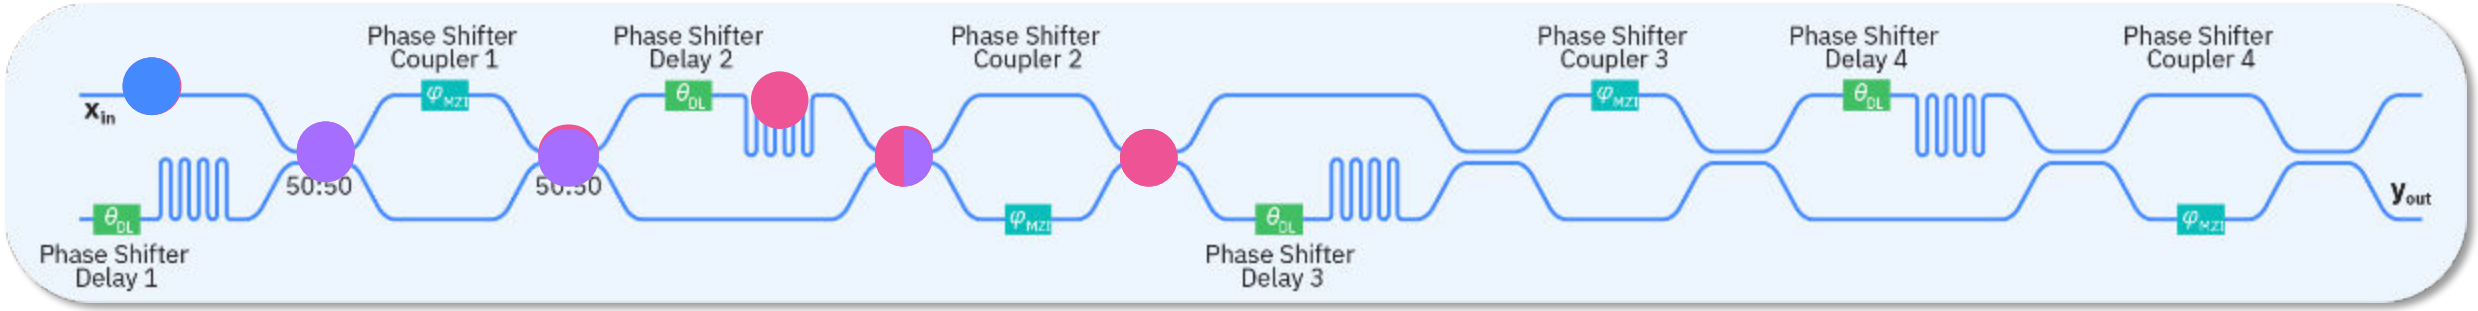

Layout and setup

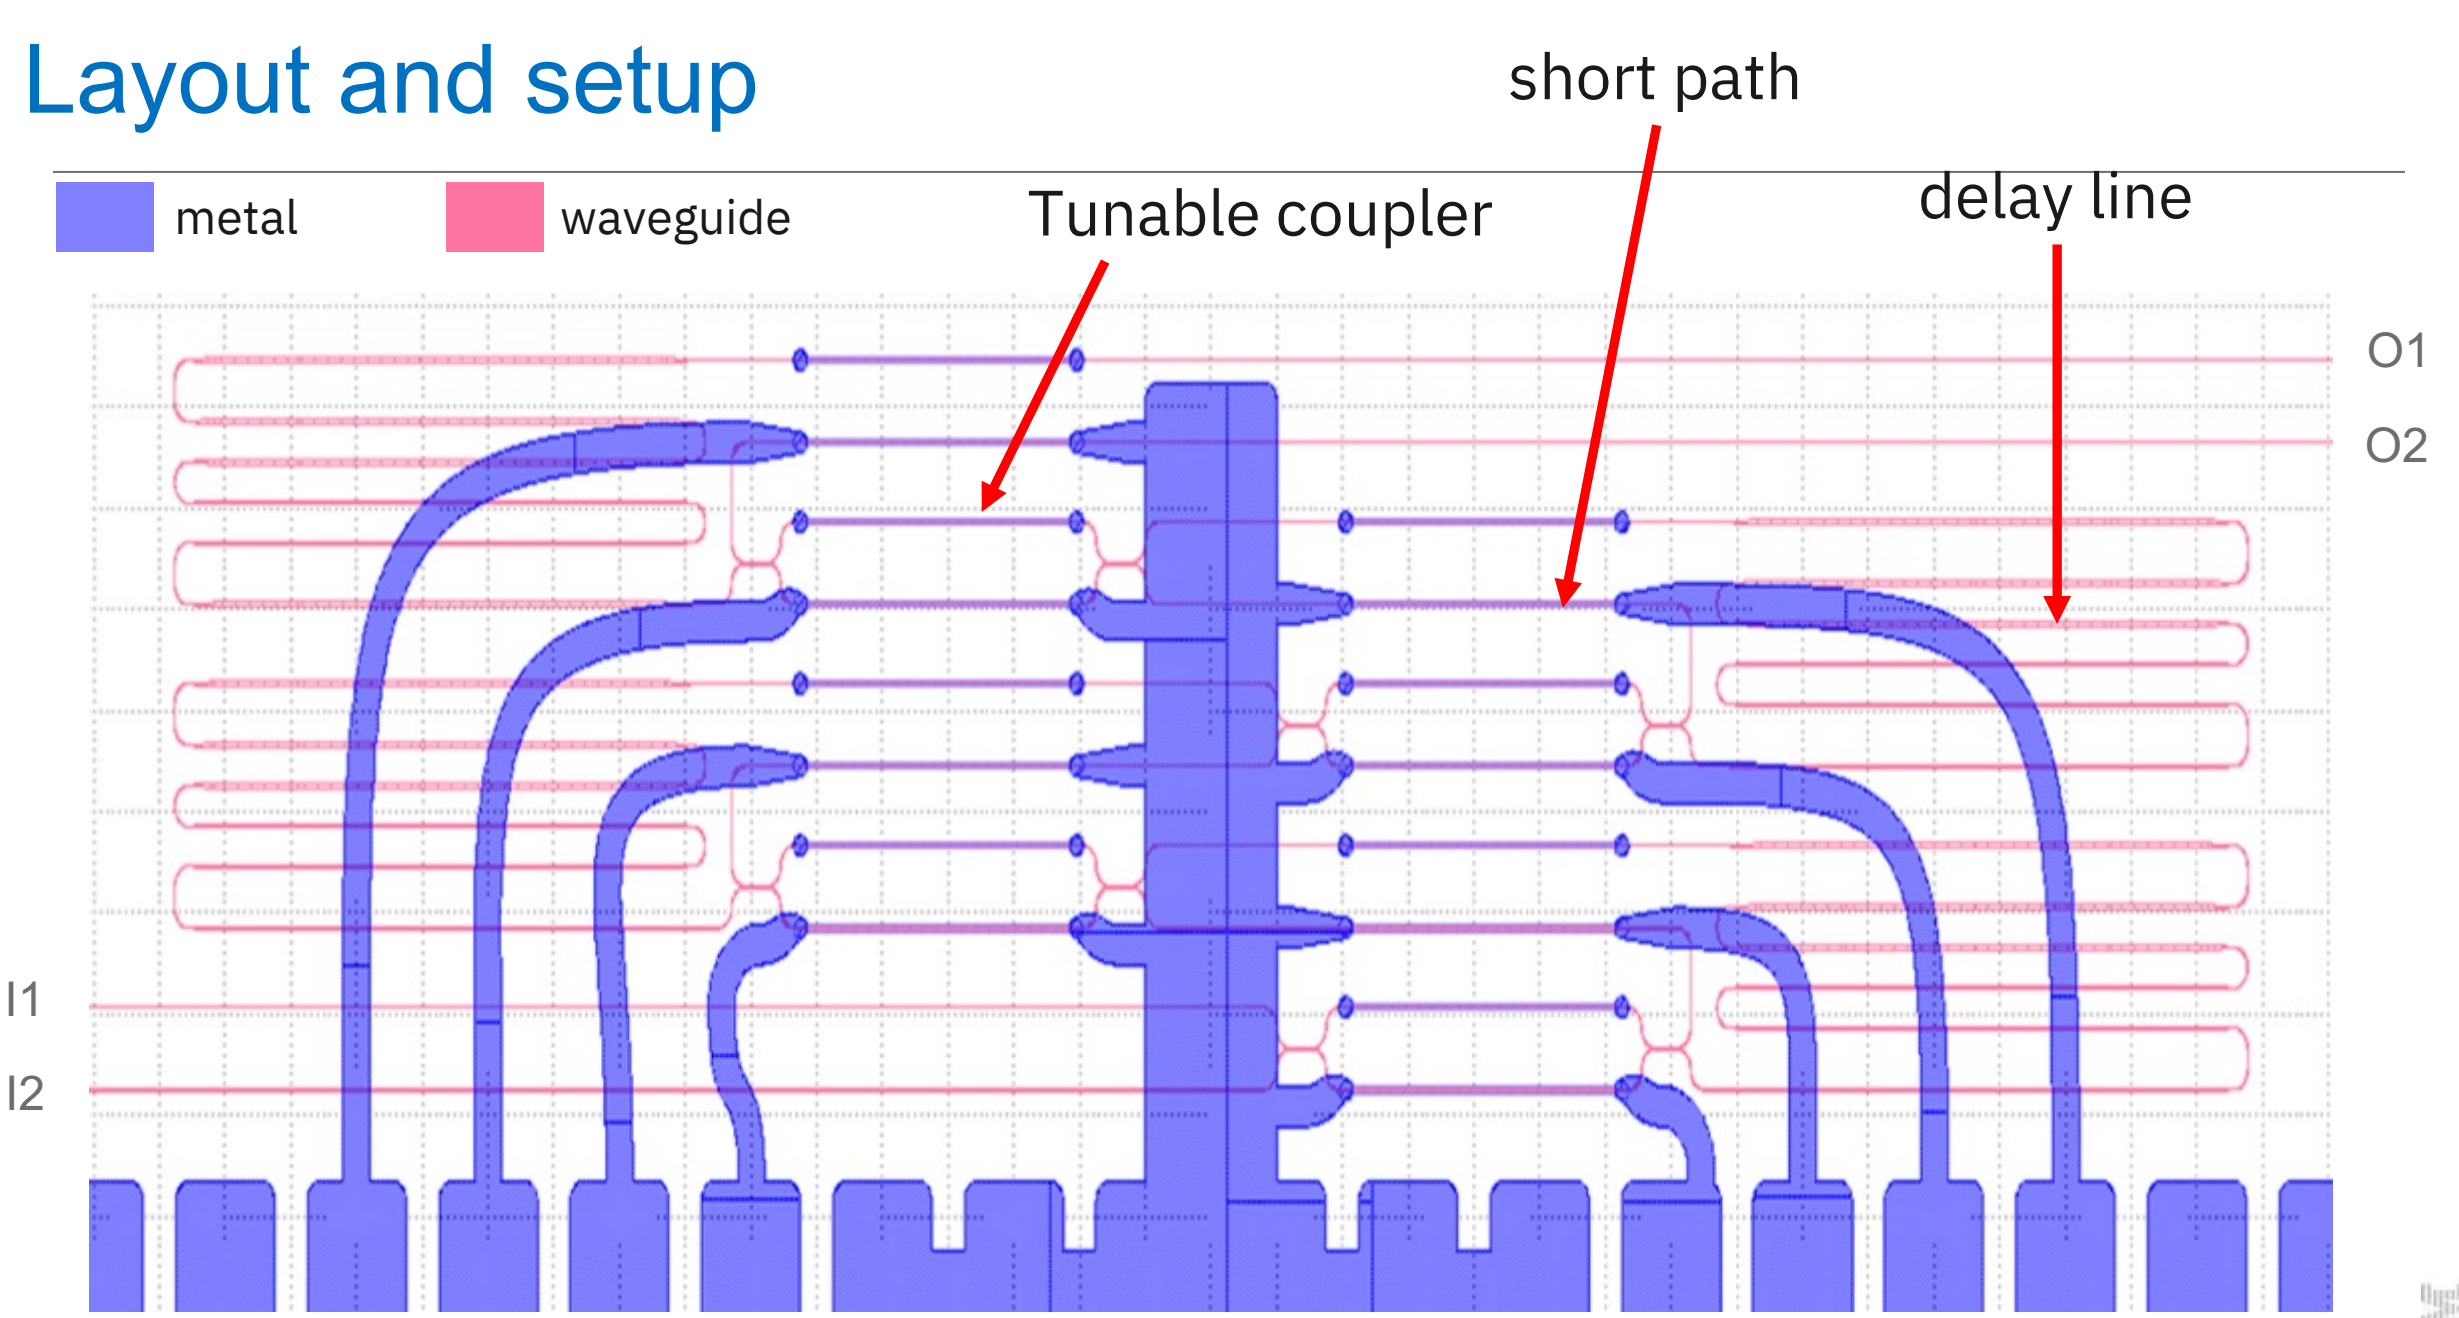

Optical convolutional signal processor

- Photonic implementations, volatile weights but well controlled and fast set

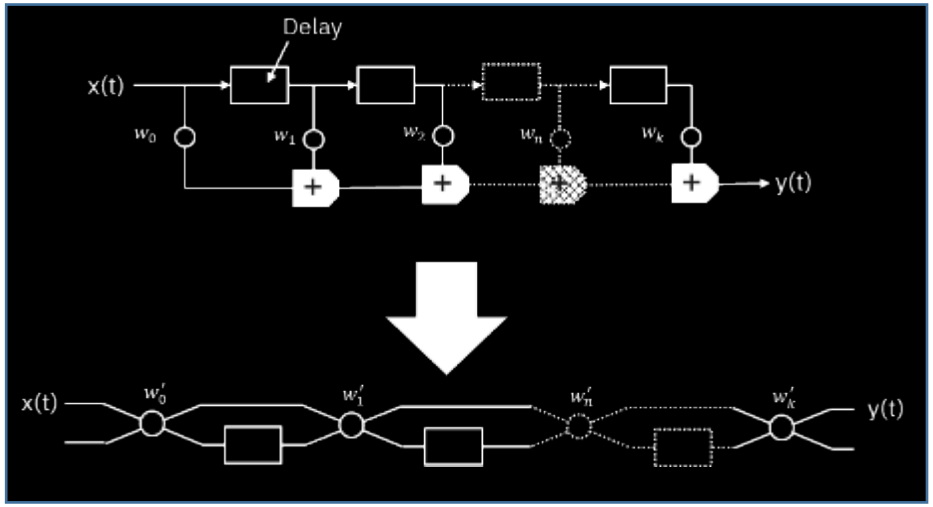

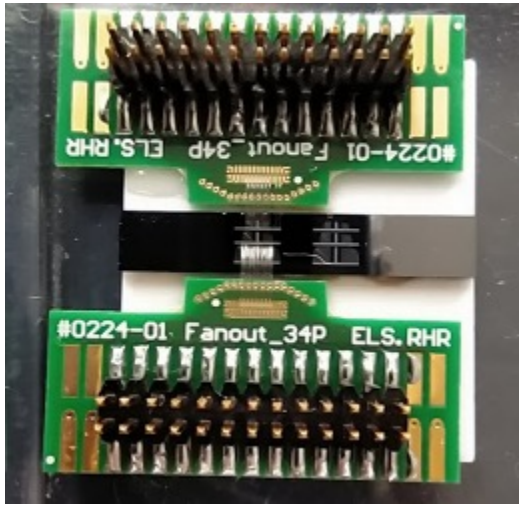

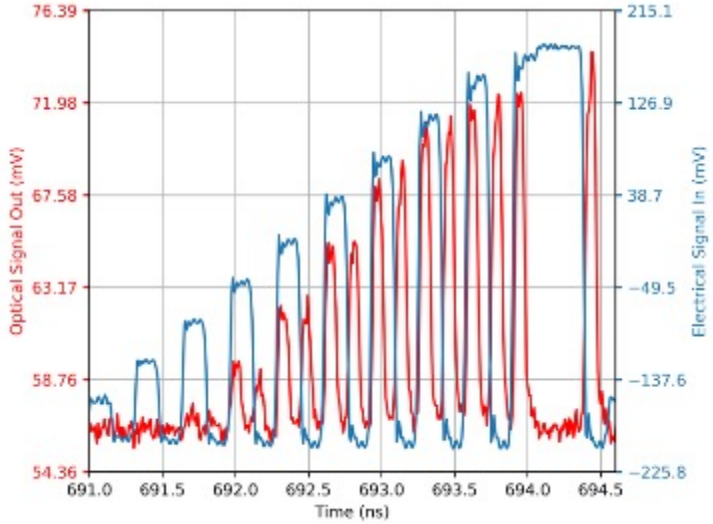

Measurements by Pascal Stark

- Time domain operation
- High-speed signal processing (12.5 GSample/s)
- Fast and efficient reconfiguration (electro-optic modulators)

Lattice filter: Link and Power budget

Link and Power budget calculations:

- **Scaling limits for the lattice filter:**
  - **Stage loss**
  - **Control complexity**

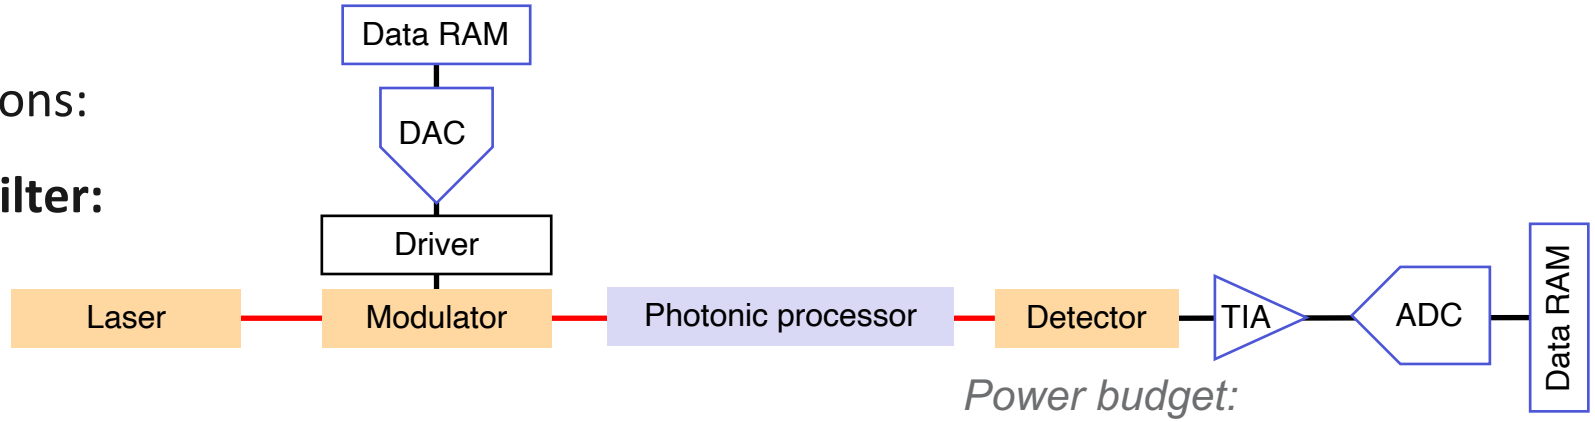

Link budget:

Parameter	Value	Unit
CW laser launch power	**13.0**	**dBm**
Laser to chip coupling	0.2	dB
Modulator insertion loss	3	dB
Lattice filter loss: **9 stages** @ 0.58 dB/stage	5.2	dB
Kernel normalization loss	2	dB
Detector coupling loss	0.2	dB
Optical power at photodetector	**2.6**	**dBm**
Power penalties (jitter, crosstalk, ISI etc.)	1.7	dB
Effective optical power at photodetector	**0.9**	**dBm**
Optical Sensitivity for a resolution of 4 bits, at 32 GSps	**-2.4**	dBm
Available link margin	**3.3**	**dB**

Power budget:

Building Block	Power in mW at 32 GSps:
Data RAM (read)	11
High Speed DAC	67
Driver and Modulator	70
Detector and TIA	6
Output ADC	115
Results RAM (write)	10
CW Laser	200
Sum of Power	476
Efficiency TOps/Watt	**2.49**

- Comparable to existing digital hardware, but
  - High-speed, low latency / Real time
  - Can do complex data and kernels
  - Room for further improvements

Analog signal processing systems

- Electrical

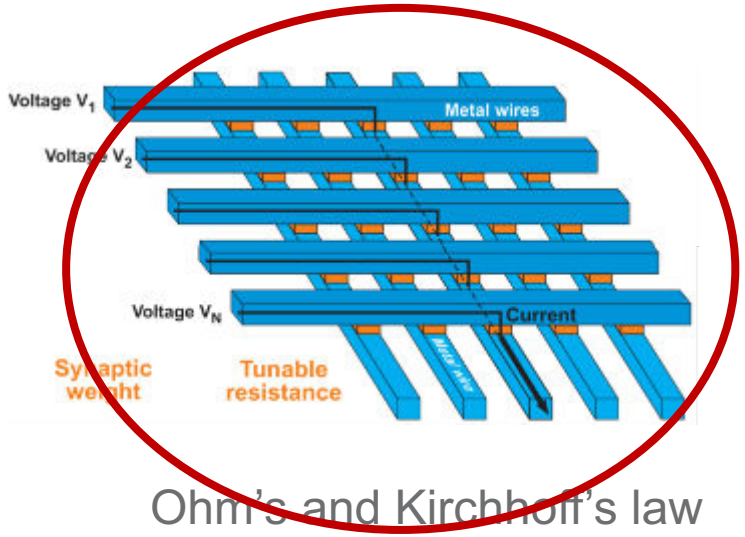

Ohm's and Kirchhoff's law

- Integrated photonics

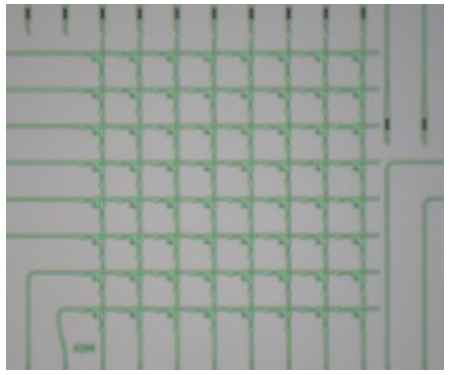

From: EU PHOENICS (U Oxford).

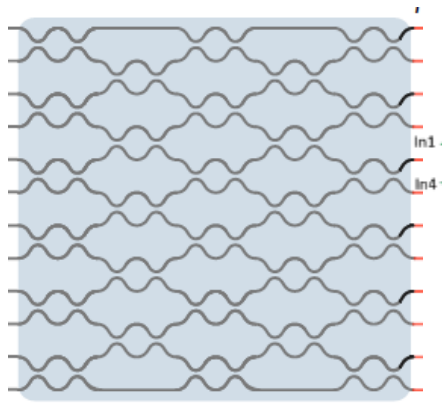

From: Y. Shen et al., doi: 10.1038/nphoton.2017.93.

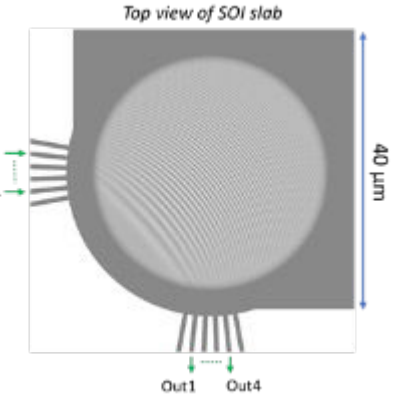

From: F. Horst, IBM.

Attenuation, interference, diffraction

- Memristive devices in a crossbar
  - PCM
  - OxRAM
  - FERAM
- Various device concepts and materials
  - Crossbar
  - Mach-Zehnder interferometer
  - Diffractive

For inference & training

MO-Resistive RAM technology

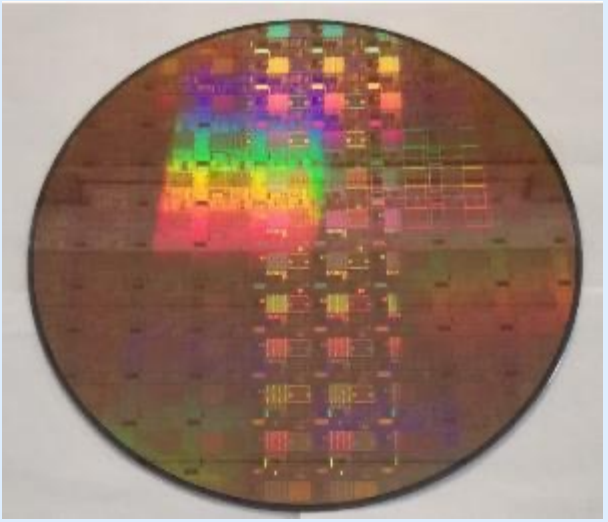

Crossbars on 300mm 14nm CMOS wafer

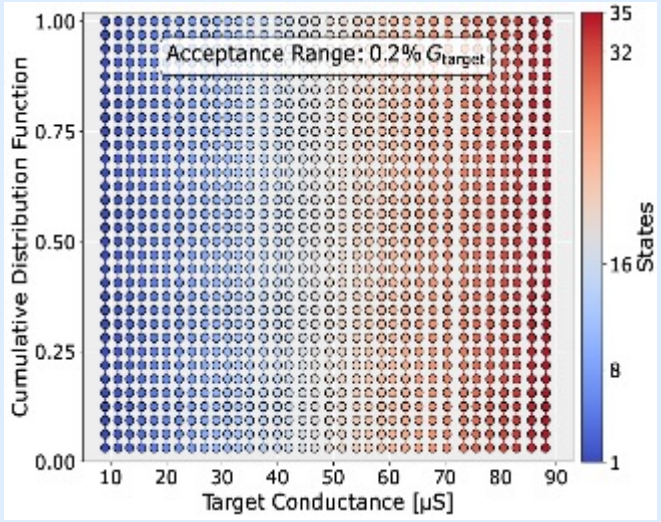

Multibit capability

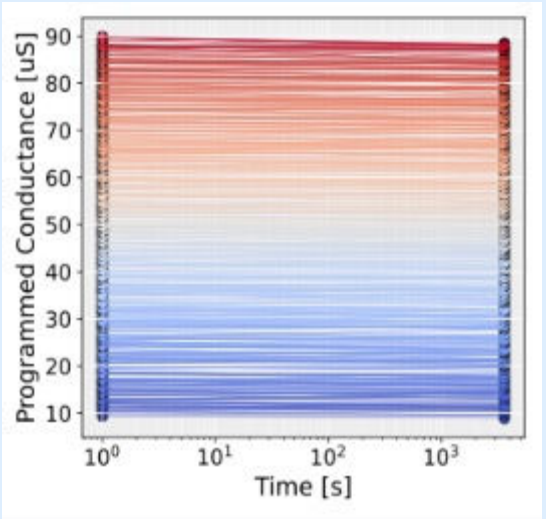

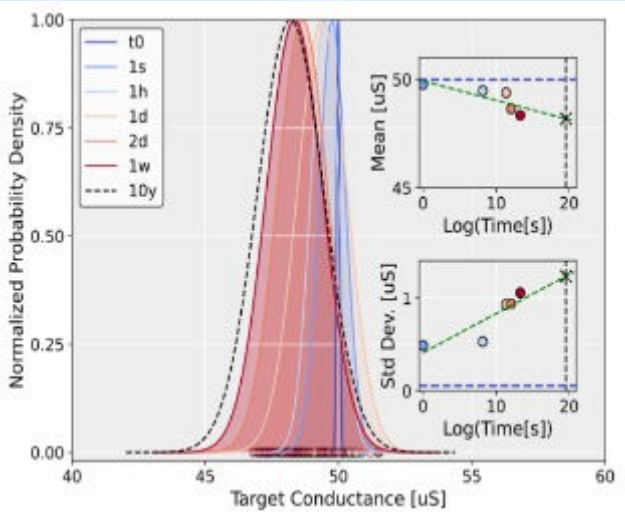

Short-, and long-term retention

BEOL MO-RRAM devices on CMOS:

- Integration on 14nm CMOS demonstrated
- Scaling: <70x70nm^2
- Multi-level capability: >35 states
- Long term retention: 10 states (@10 yrs)
- Good endurance: >10^8

So, what is my answer …

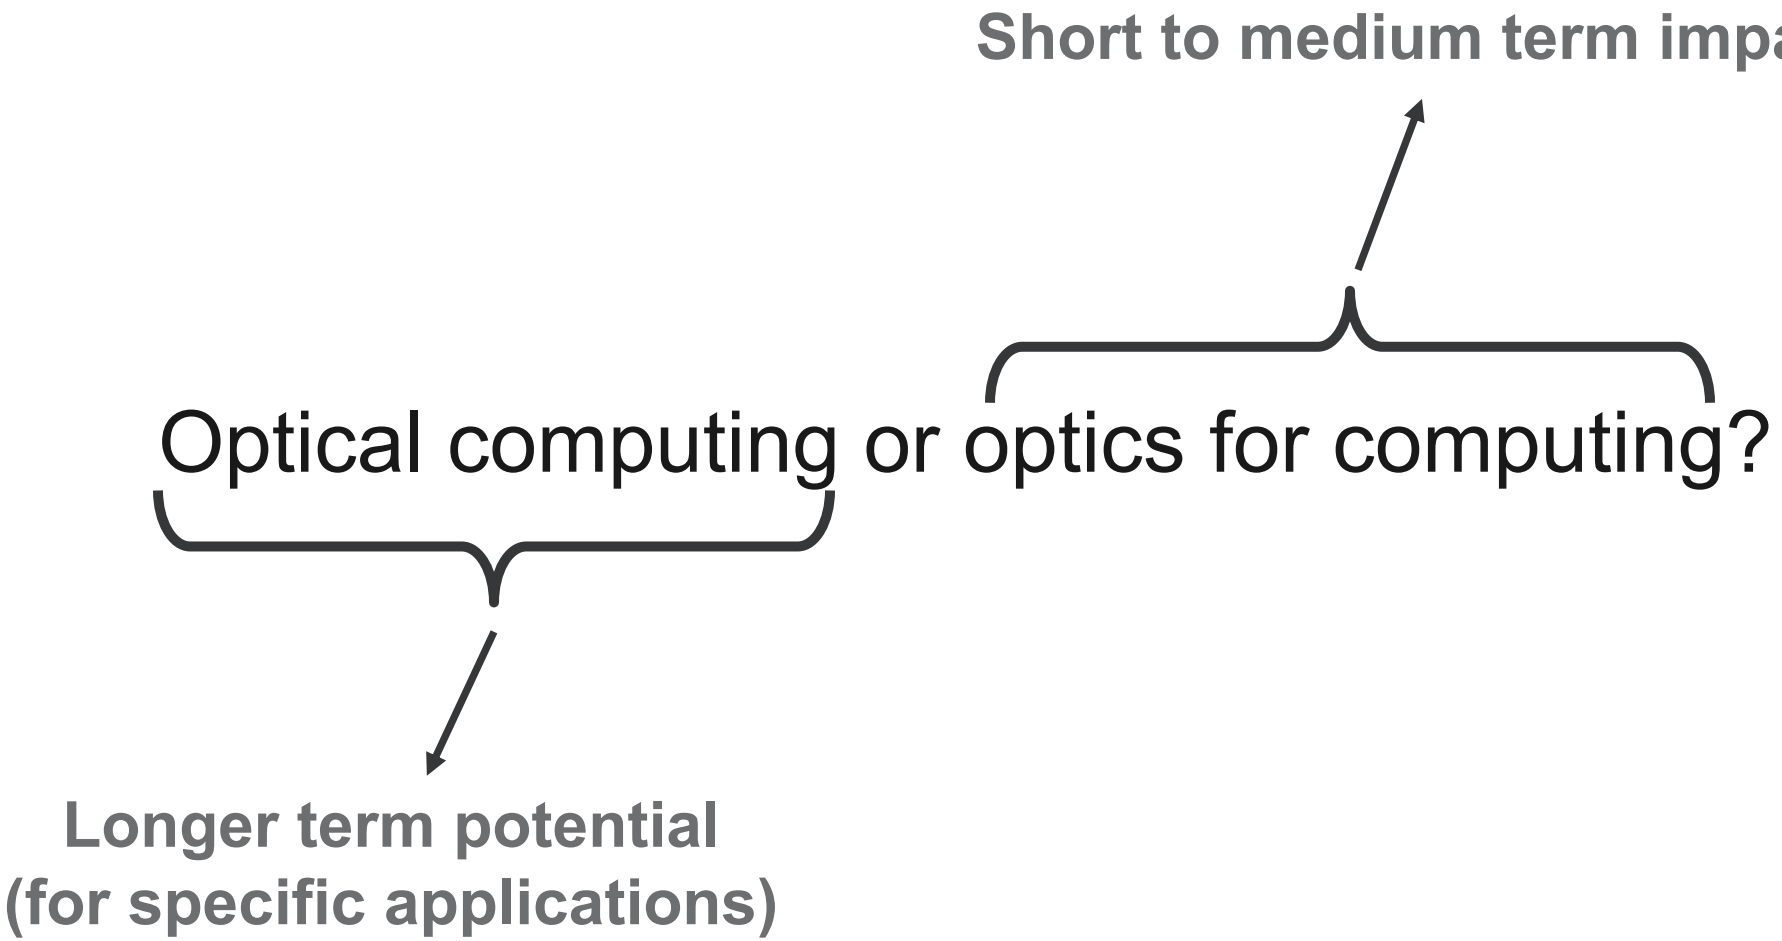

Discussion

- Energy consumption of AI
  - Minimize memory access
- Don’t forget about traditional technology scaling!
  - Is continuing at all levels
- Analog signal processing
  - Map the technology to the Neural Network architecture
  - Electrical for density and scalability
  - Optical for throughput and reconfigurability
- Know the application and the technology prospects

Co-Packaged Optics and AI Systems - The team

Acknowledgments

IBM Research Europe – Zurich, Switzerland
Neuromorphic Devices and Systems group

IBM Research, worldwide research teams

The IBM BRNC cleanroom opteam

Co-funded by the European Union Horizon 2020 Programme and the Swiss National Secretariat for Education, Research and Innovation (SERI)

EU & CH-SERI
PHOENICS, PHOENIX, PROMETHEUS

PHASTRAC, FIXIT, CONCEPT, TOPOCOM

Thank you for your attention!

OFB@zurich.ibm.com