Optical Interconnects for computing applications

Bert Jan Offrein

Swisslasernet Workshop, IBM October 2010

Exponential Growth in Supercomputing Power

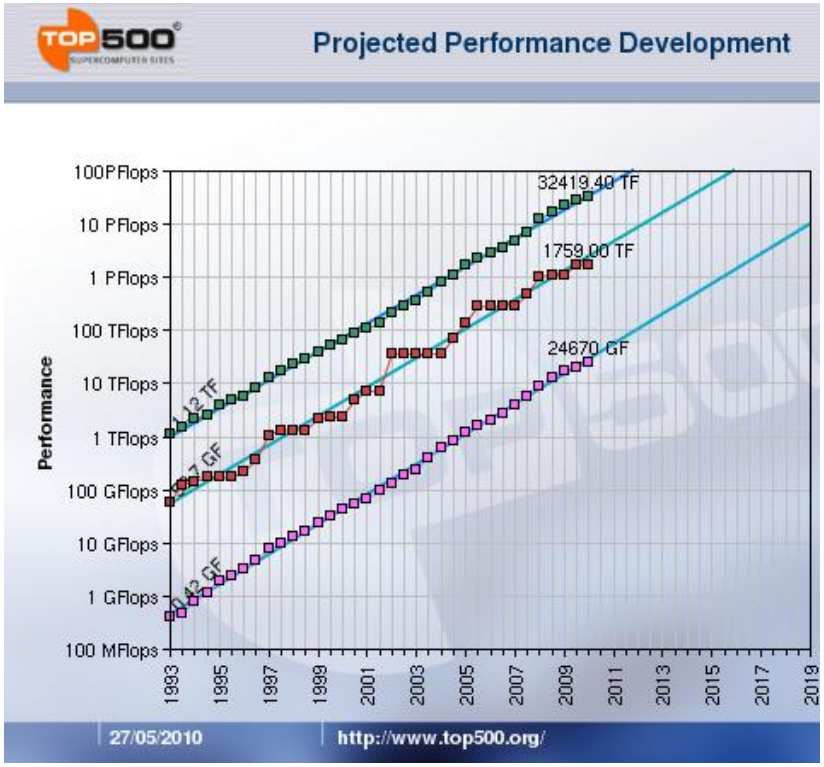

- 33% of all FLOPs are IBM
- 196 of 500 are IBM Systems
- 3 Systems in top 10 (3,5,9)

- Performance increase
  Factor 10 every 4 yrs
- Exascale Systems by 2020
  3 Orders increase
  compared to today!!!

Demands new technologies

- BW requirements must scale with System Performance, ~1B/FLOP (memory & network)
- **Requires exponential increases in communication bandwidth at all levels of the system** → Inter-rack, backplane, card, chip

Electrical BW Bottlenecks – Optics opportunities

- **Electrical Buses become increasingly difficult at high data rates (physics):**
  - **Increasing losses & cross-talk ; Frequency resonant effects**
- **Optical data transmission:**
  - **Power Efficiency , much less lossy, not plagued by resonant effects**

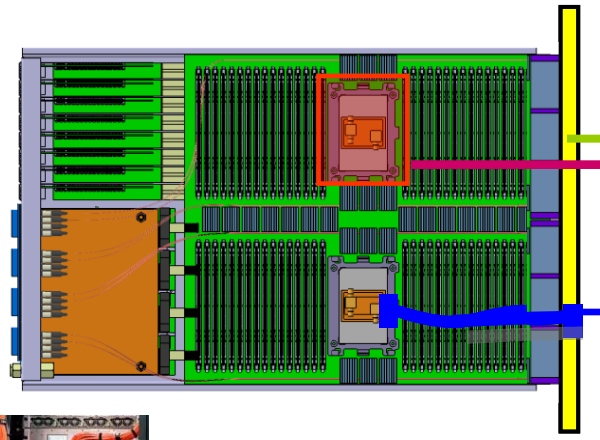

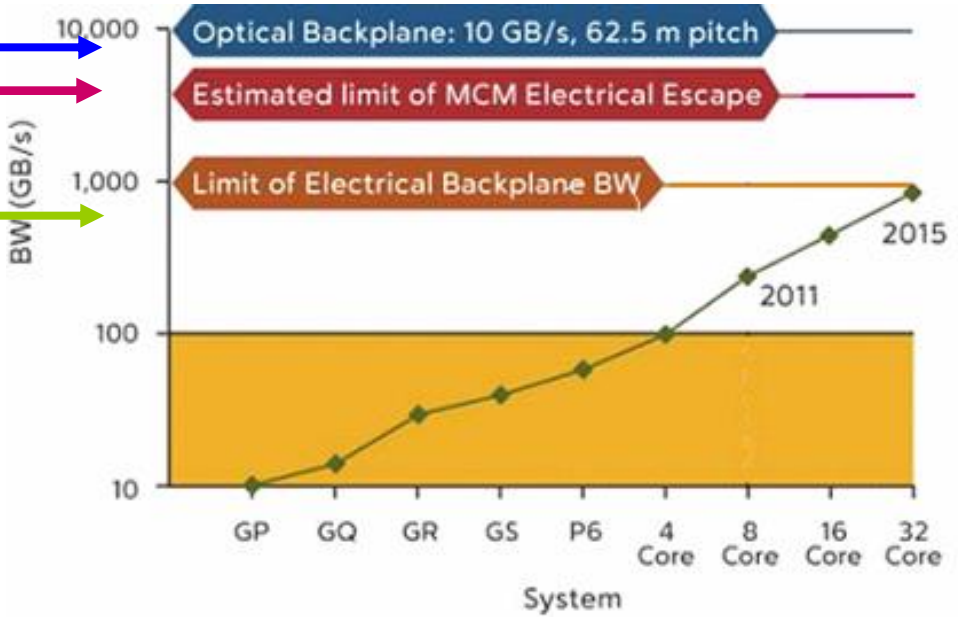

Rack

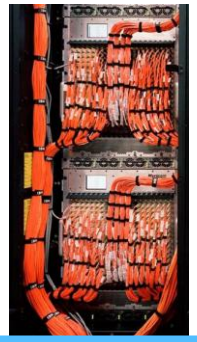

Backplane

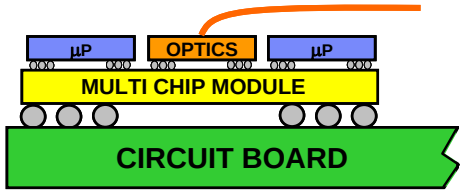

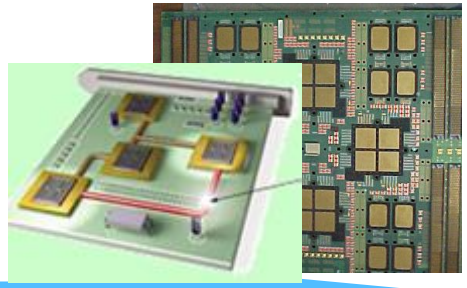

Density advantage of optics

Cables

1 m cable

IBM

Connectors

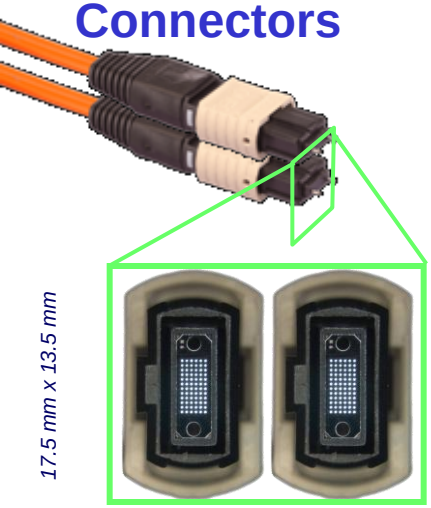

17.5 mm x 13.5 mm

144 x 10+ Gbps

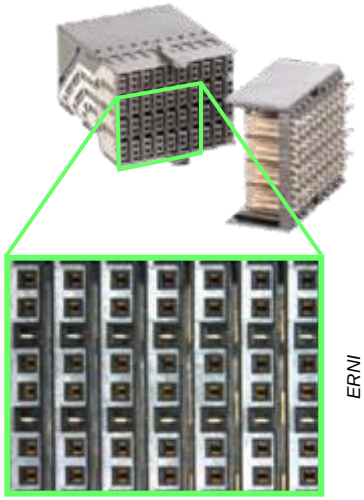

ERNI

21 x 10 Gbps

PCB-Tracks

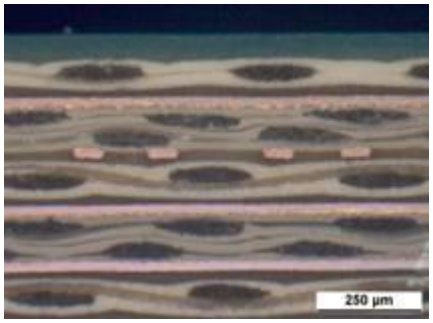

differential striplines
2 x 10 Gbps
80x17 µm tracks @ 460 µm pitch

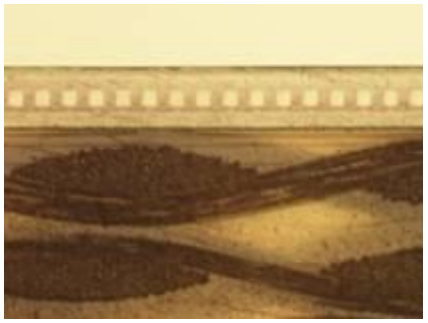

IBM

optical waveguides
16 x 10+ Gbps
35x35 µm cores @ 62.5 µm pitch

Photonics Roadmap – Optical Interconnects in Supercomputing

Backplane

Today

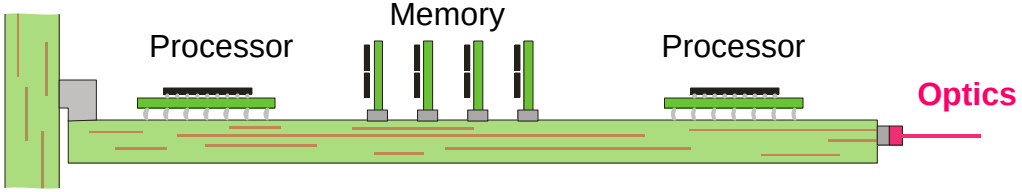

Electrical system, optical fibers at card edge only

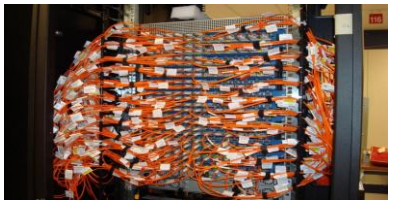

2011

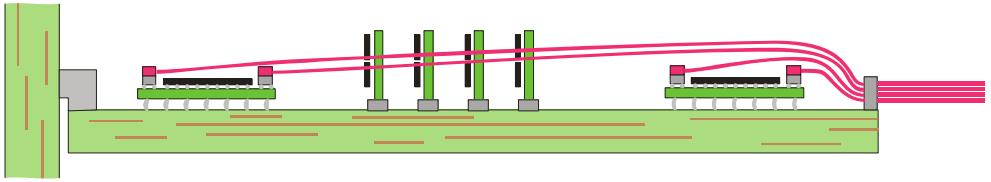

Optical fibers across the boards

~2015

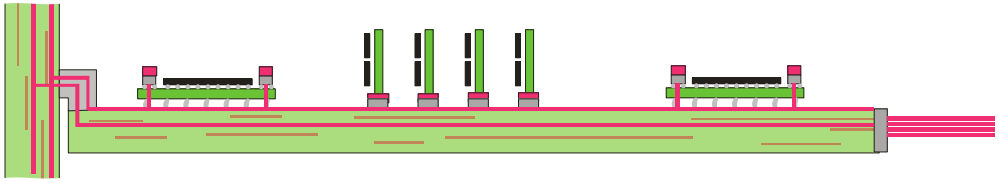

Optical waveguides in/on boards

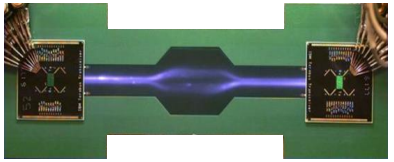

~2020

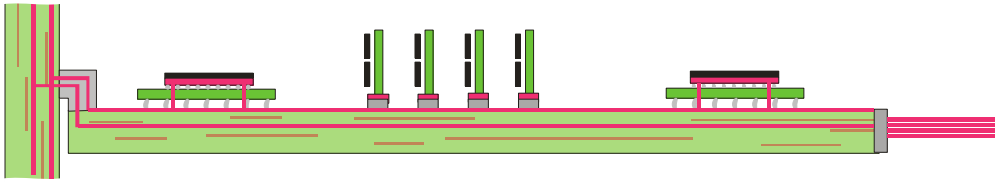

Optical interconnects integrated with the processor

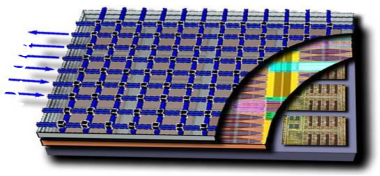

Development

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.

Polymer Waveguides: Processing principle

Substrate material

- FR4

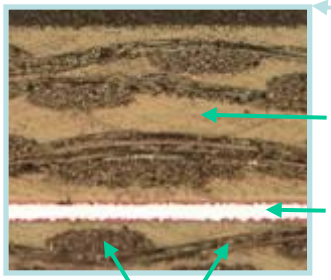

Polymer material

- Polyurethane
- Acrylate
- Polysiloxane

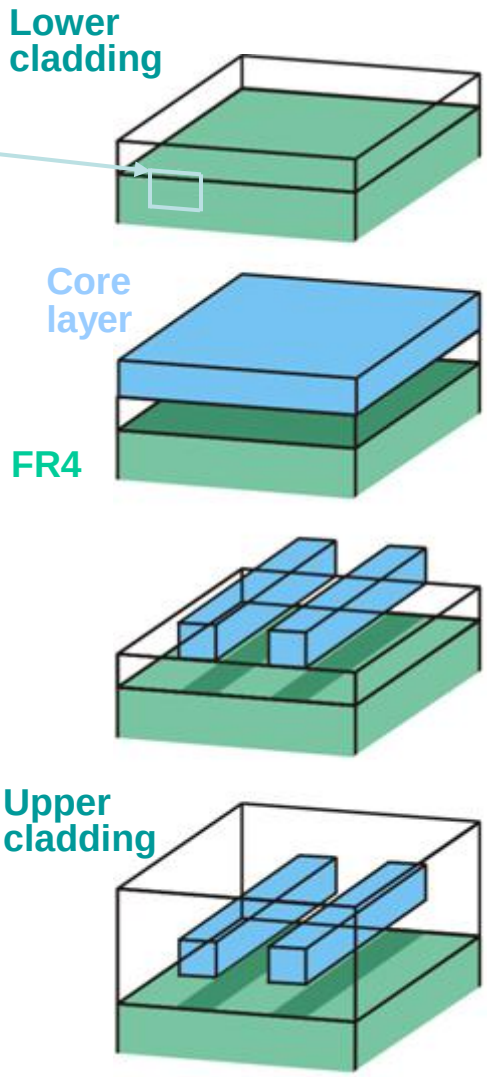

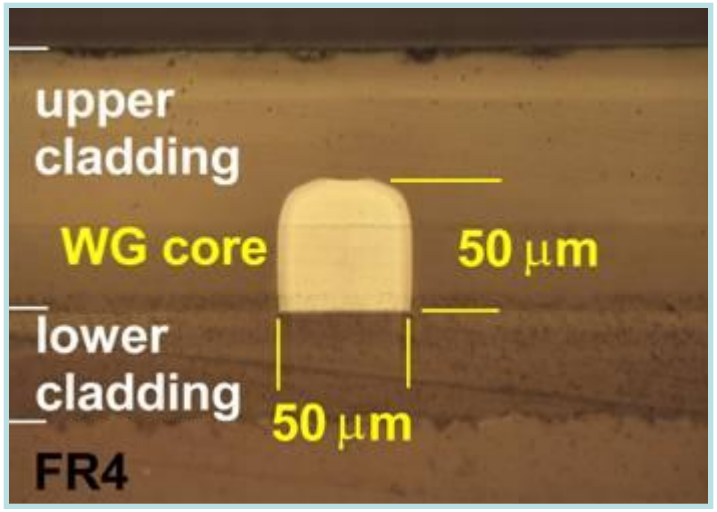

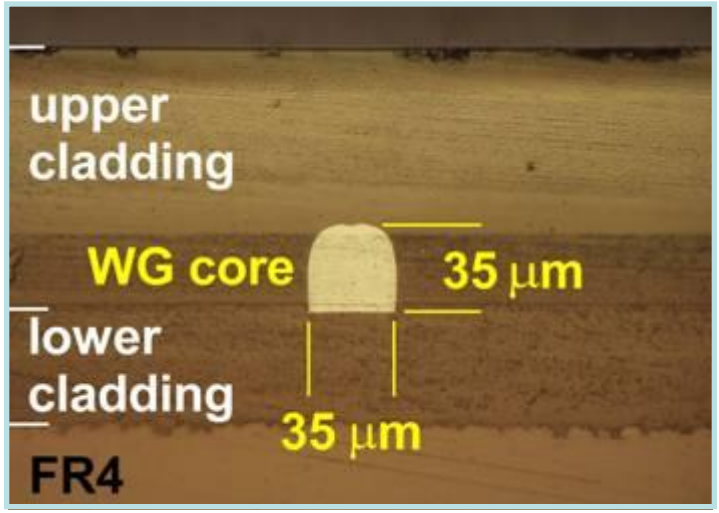

Optical printed circuit boards

In collaboration with Varioprint

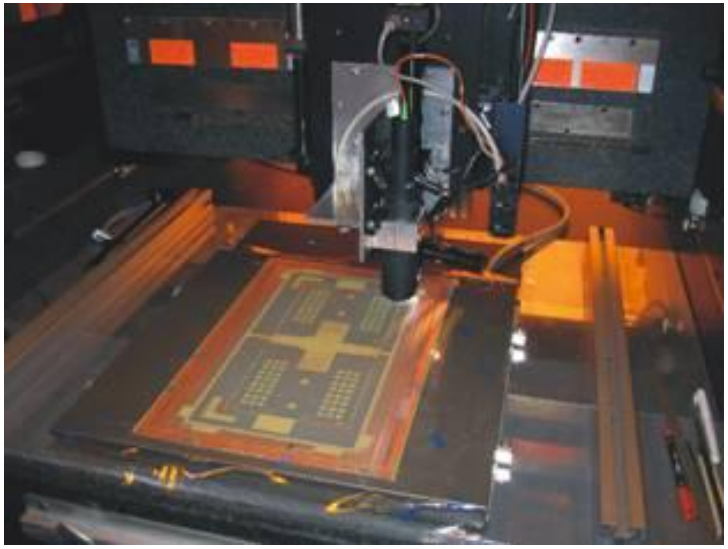

Waveguide processing on large panels,
305 mm x 460 mm

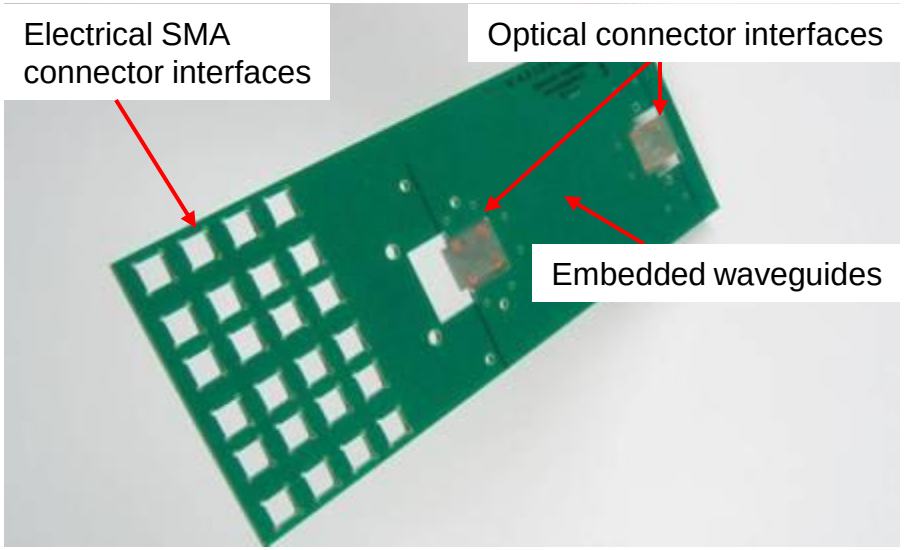

Finished optical board with optical and mechanical interfaces

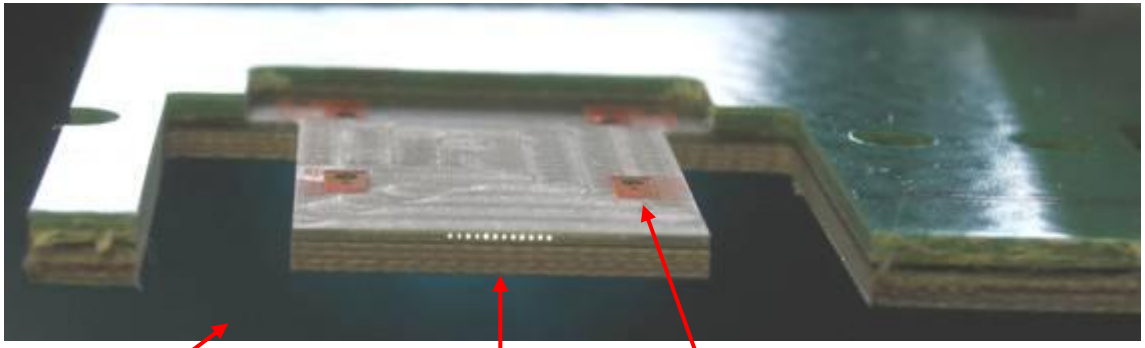

Passive alignment of optical components

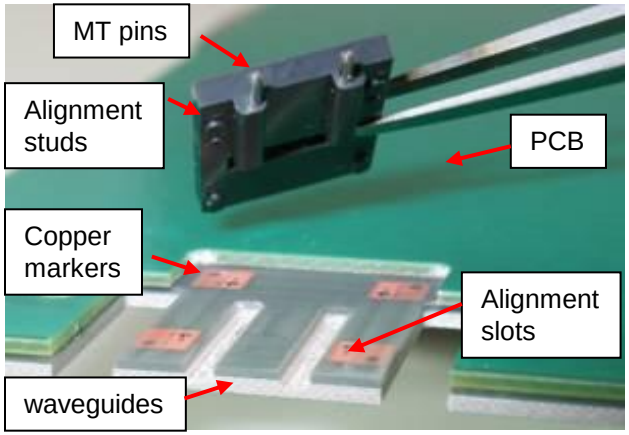

MT ferrule aligned by copper markers

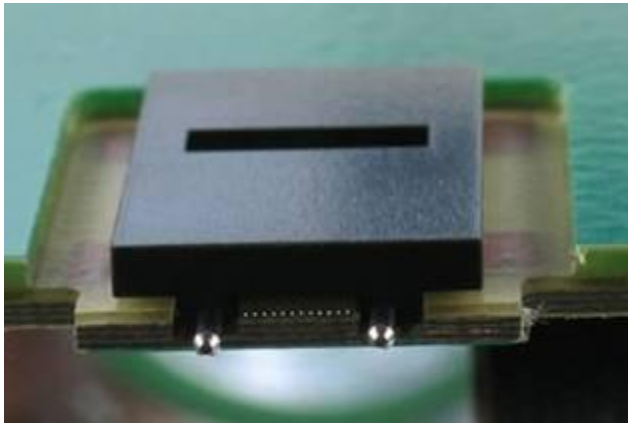

Positioned MT ferrule to polymer waveguides

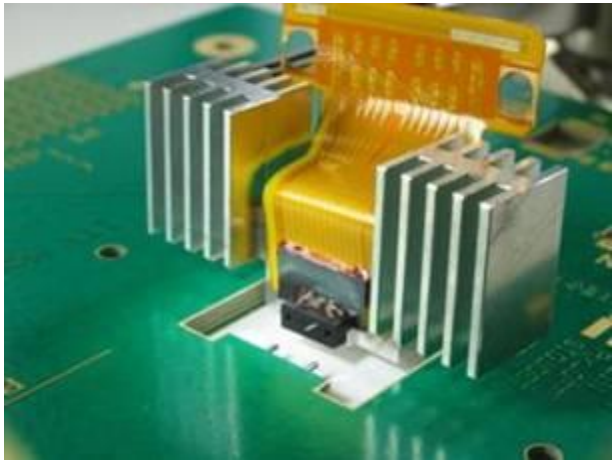

Connection of 120 Gb/s transceiver module to the MT ferrule

Assembled transceiver module
Transceivers realized in collaboration with Intexys Photonics

120 Gb/s Board-to-Board Optical Link Demonstrator

Optical TX/RX board as building block

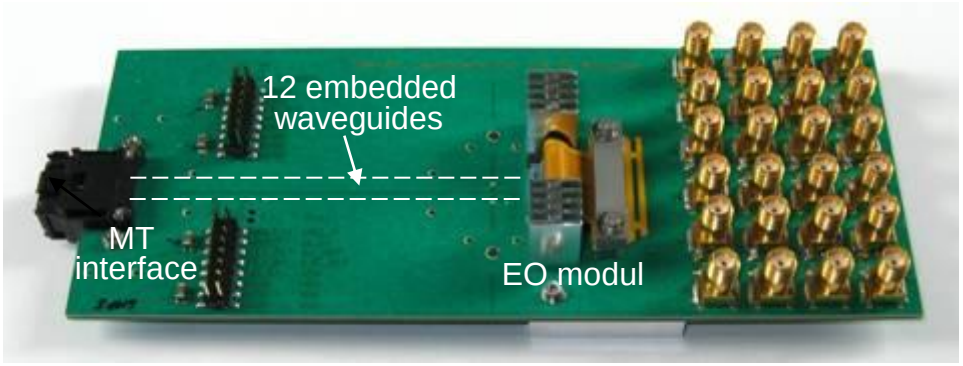

- Embedded polymer waveguides (12 channels)
- Passive alignment of MT standard based connectors
- MT interface as standard interface for WG, fiber bundles/optical flexes and transceivers
- Pluggable TX/RX module (butt coupling)

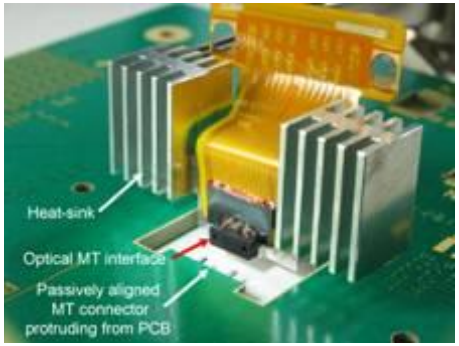

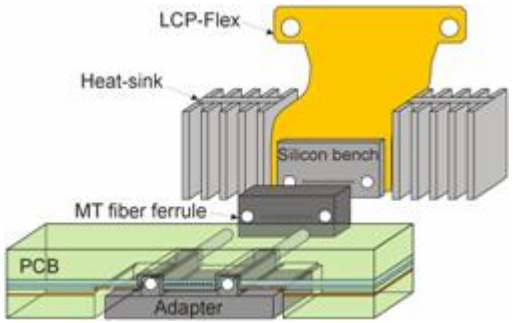

Complete 12x10 Gb/s link demonstrator

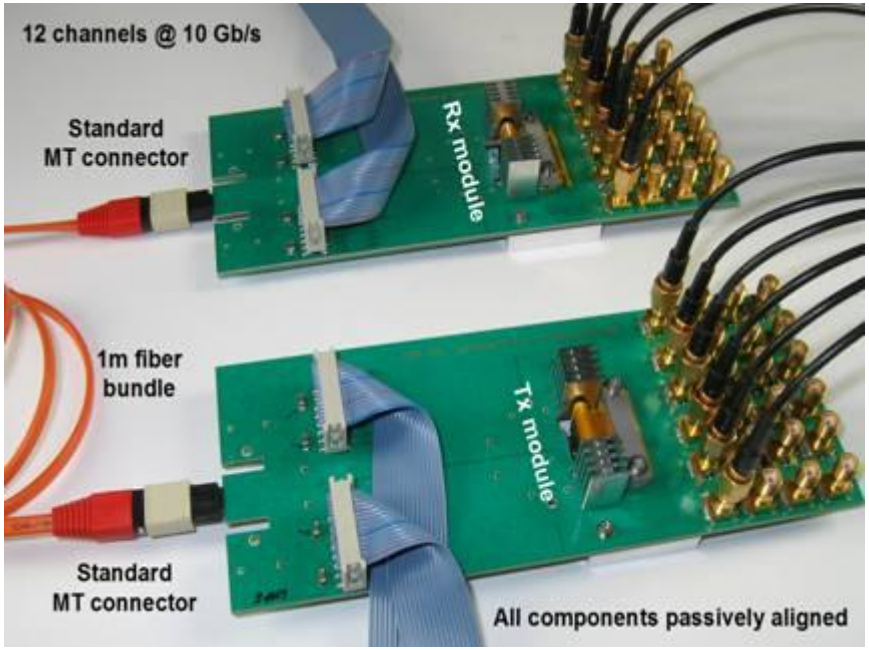

Eye diagrams for 2 channels at 10 Gb/s

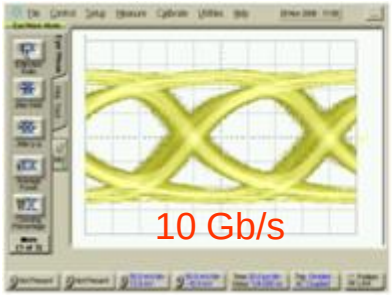

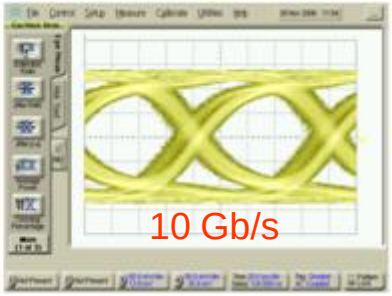

TERABUS Link Demonstrator

Optocard

- Waveguides on top of PCB
- Ultra-high density (62.5 μm channel pitch)
- Coupling with 45° mirrors

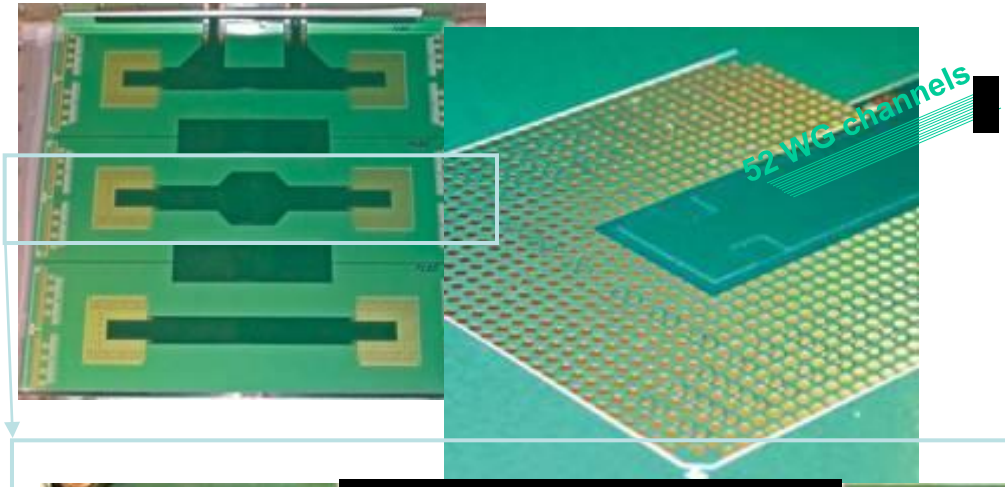

Optomodule with Optochip

- 16 TX + 16 RX
- Minimized OEIC footprint (17 mm²)
- High-speed (up to 15 Gb/s/ch)
- Ultra low power ICs
  (5 mW/Gb/s per unidirectional link)

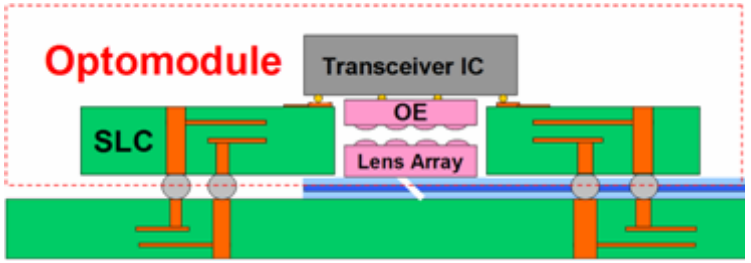

16 channels TRX1 ⟶ TRX2 at 10Gb/s

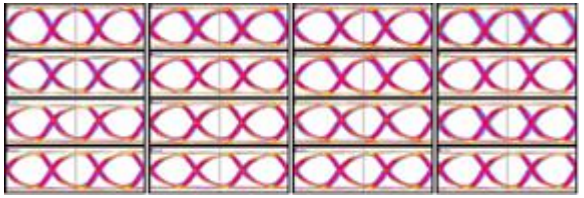

16 channels TRX1 ⟵ TRX2 at 10Gb/s

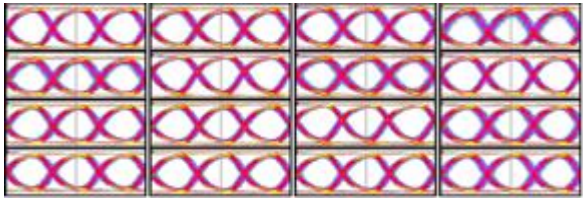

Multi-layer flexible waveguide connector

Established technology:

Single layer on **rigid** boards

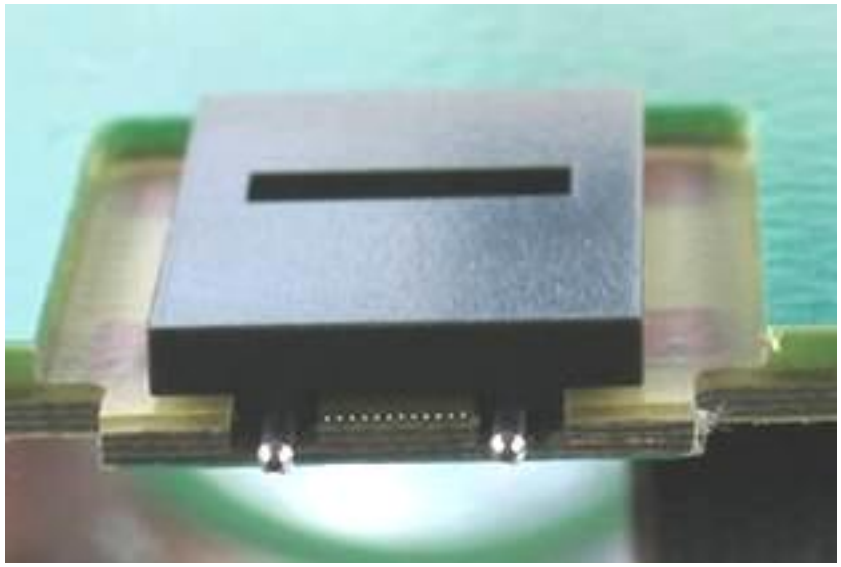

- Rigid board allows:
  - a good horizontal alignment of the WGs
  - robust alignment features

New situation:

Multiple flexible layers

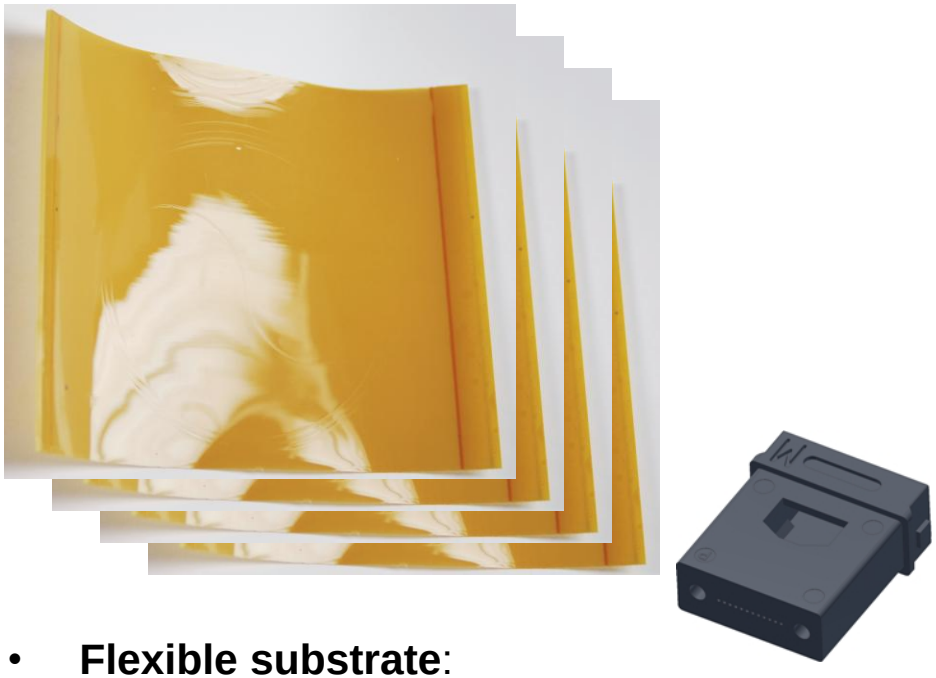

- **Flexible substrate**:
  - WG alignment critical (curling, etc)
  - Mechanically weak
- **Multi-layer** alignment of polymer waveguides

Multi-layer waveguide optical backplane

In collaboration with Varioprint

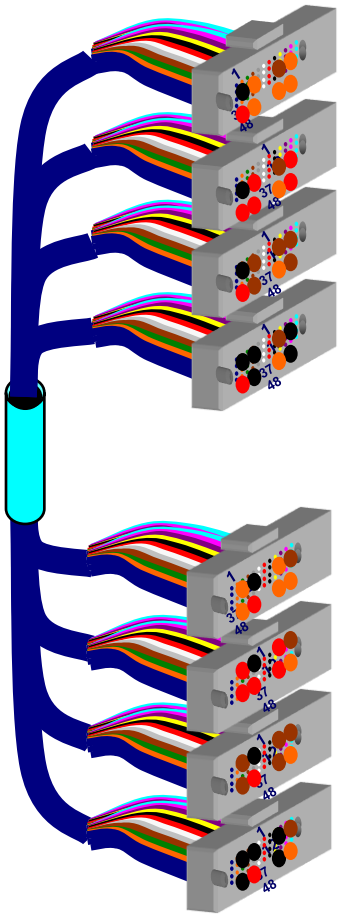

One ‘quad L-link’, stack of 8 waveguide flex sheets, 192 waveguides, 8 connectors

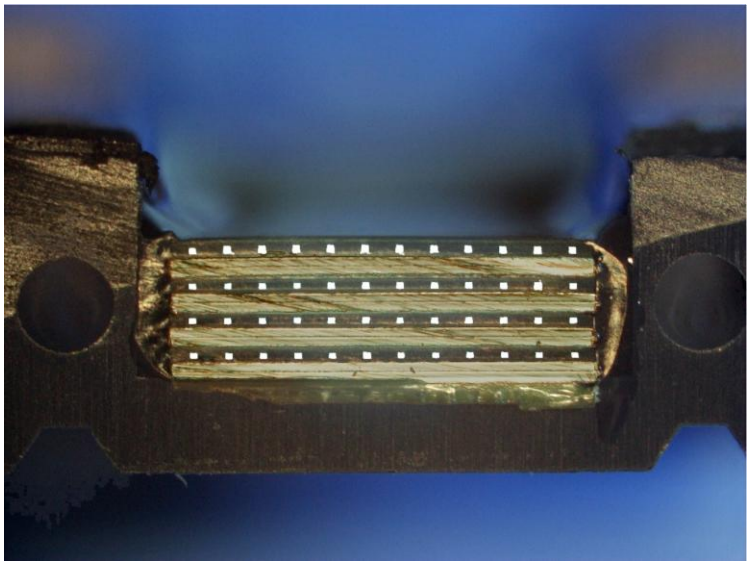

4x12 = 48 waveguide connector

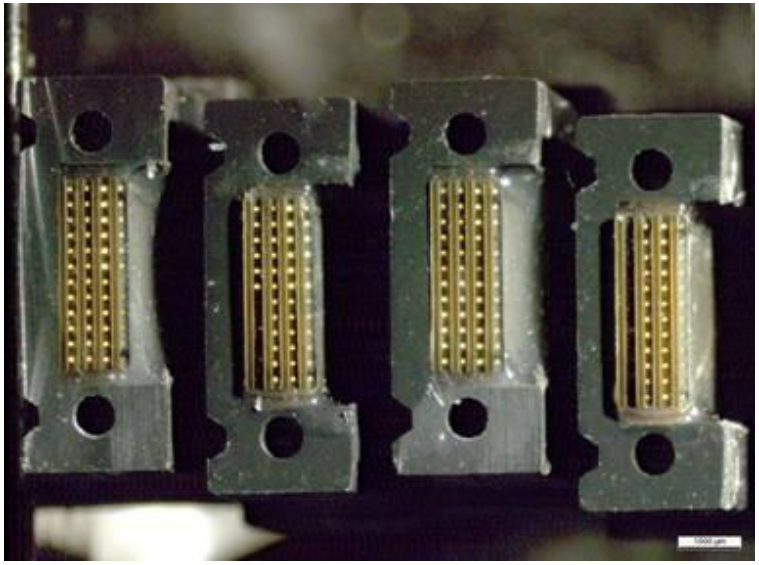

4 connectors with 48 waveguides each

Demonstrated feasibility of optical waveguide technology
for complex and high channel count interconnect solutions in future computing systems

Hybrid Optical Interconnect Technology

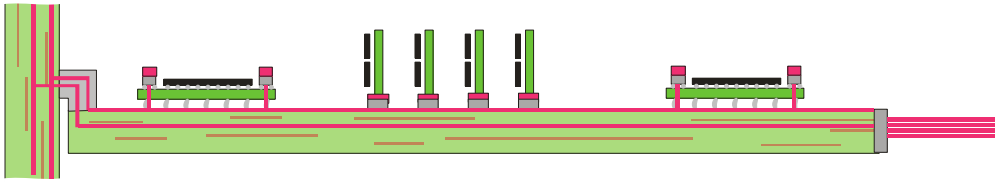

Optical waveguides in/on boards

Optical printed circuit board technology allows
- low cost processing and assembly
- higher density optical links

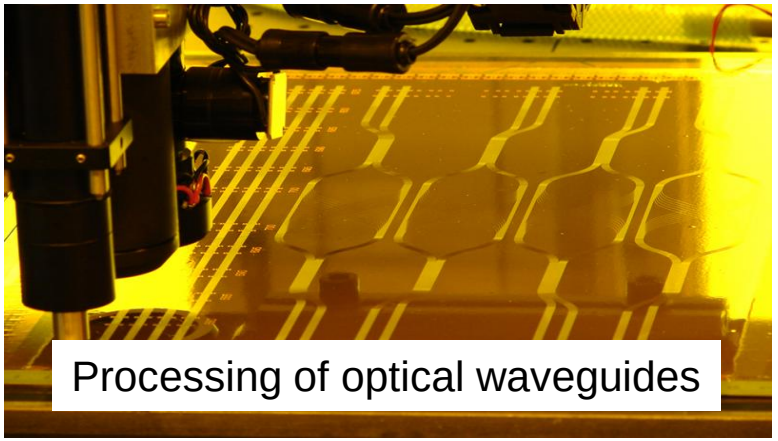

Processing of optical waveguides

Waveguide based multi-layer optical backplane

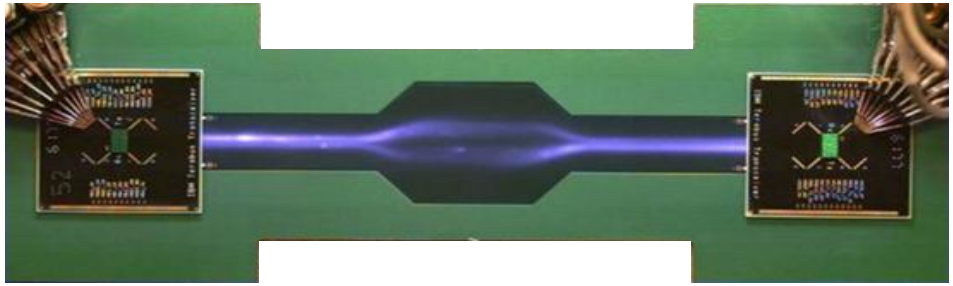

Optical link with waveguides between processors

Integrated Optical Functions at Processor-Level

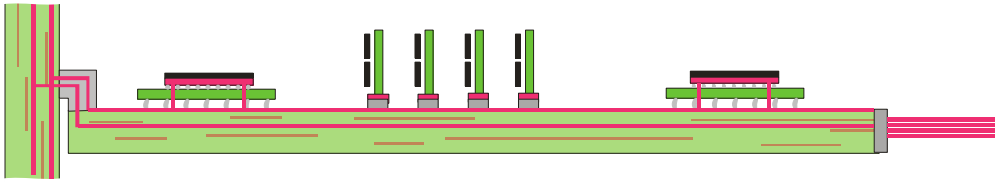

Optical interconnects integrated with the processor

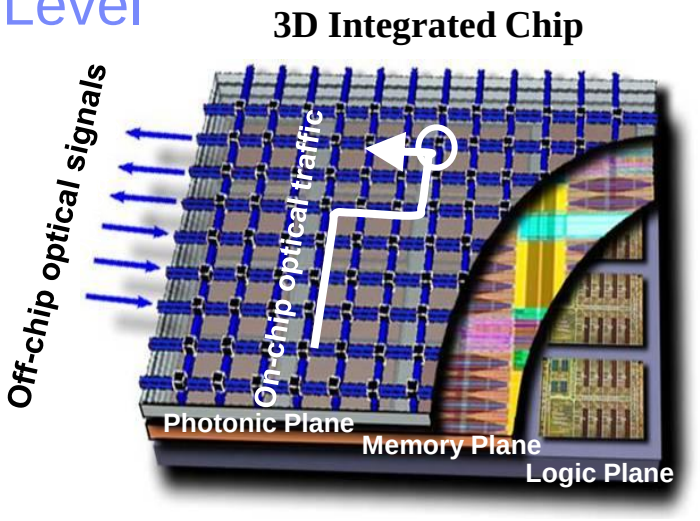

3D Stack, combining logic, memory and optical functions

Silicon photonics
- enables the integration of optical functions in silicon
- allows a tight integration between logic and optics
- brings massive improvements in bandwidth density, power efficiency and cost

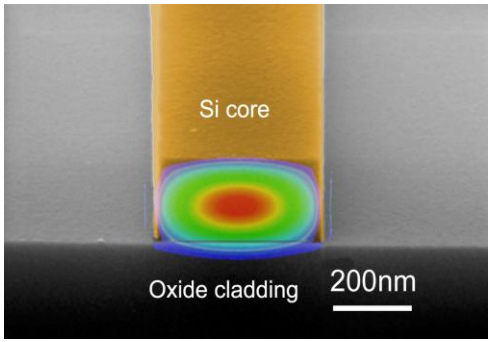

Silicon waveguide

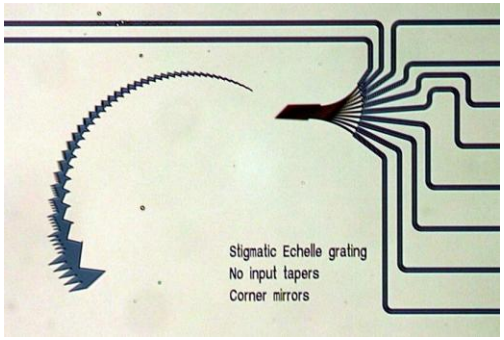

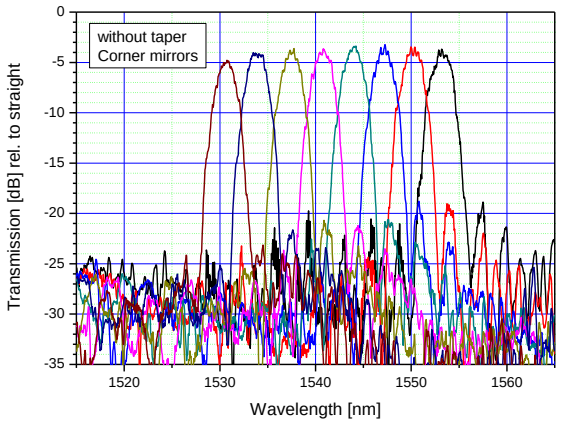

Optical wavelength multiplexer for ultra-high bandwidth data transfer

Conclusions

- Optical interconnects will play an important role in future computing systems
- Multimode optical printed circuit board technology
  - Avoids tedious fiber handling
  - Simplifies electro-optical assembly
  - Increase interconnect density
  - Reduces cost
- Single mode silicon photonics
  - Is the next scaling step
  - Increases the bandwidth per channel by using WDM
  - Enables a direct integration with the CMOS logic
- Exciting future ahead !

Thank you for your attention

- Bert Jan Offrein
- ofb@zurich.ibm.com