LEVERAGING INTEGRATED SILICON PHOTONICS FOR A STREAMLINED GPU ARCHITECTURE

Luca Ramini[*], Jinsung Youn[*], Marco Fiorentino[*], Steven Dean[**], and Raymond G. Beausoleil[*]

[*]Hewlett Packard Labs, Milpitas, CA, USA, [**]Hewlett Packard Enterprise, Chippewa Falls, WI, USA

In Graphic Processing Unit (GPU)-Accelerated System there are two key communication links:
1) **GPU-GPU link** and 2) **GPU–high-bandwidth memory (HBM) link**.

We propose a streamlined GPU architecture using an integrated Si Photonics via an optical network-on-chip (ONoC) to:
- reduce the interconnect bottlenecks in GPU-to-GPU and GPU-to-HBMs due to copper-based electrical lines in Si interposer
- reduce high design complexity and degraded signal integrity due to parallel lines (> 1,024) per a GPU-HBM link

Our streamlined GPU Architecture
- Coarse wavelength division multiplexing **(CWDM)** with **wavelength-routed optical ring network** for GPU-GPU link
- Dense wavelength division multiplexing **(DWDM)** with serializer/de-serializer **(SerDes)** for GPU-HBM link
- **No modification** in GPU and HBM PHY by integrating SerDes and photonic devices inside SiPh interposer

Electrical Approach

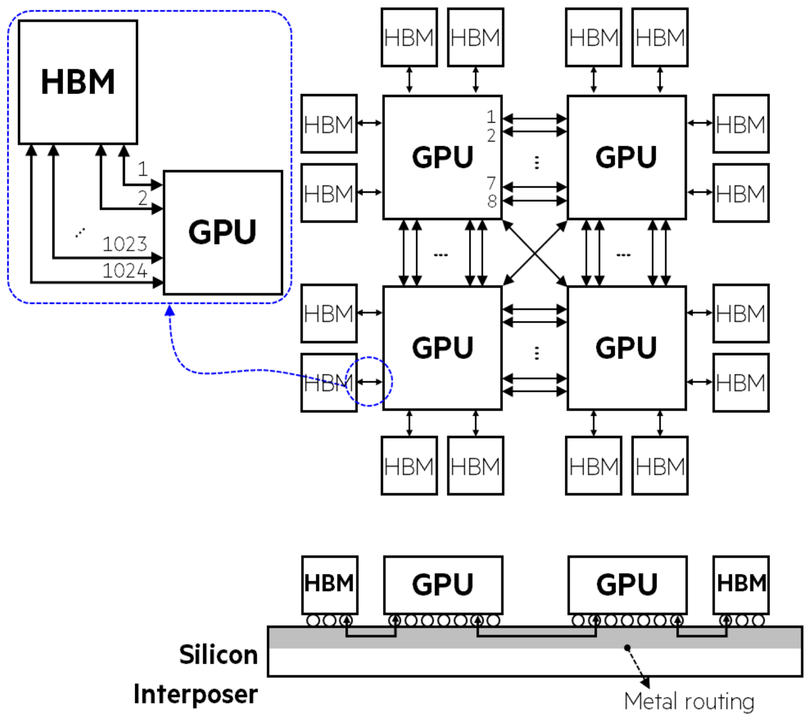

Our idea: Silicon Photonics Approach

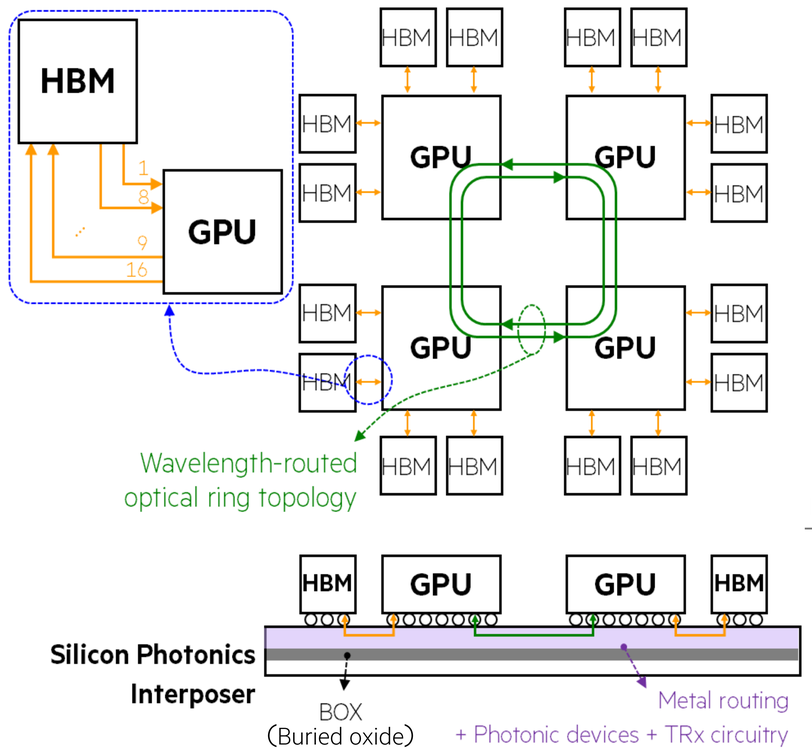

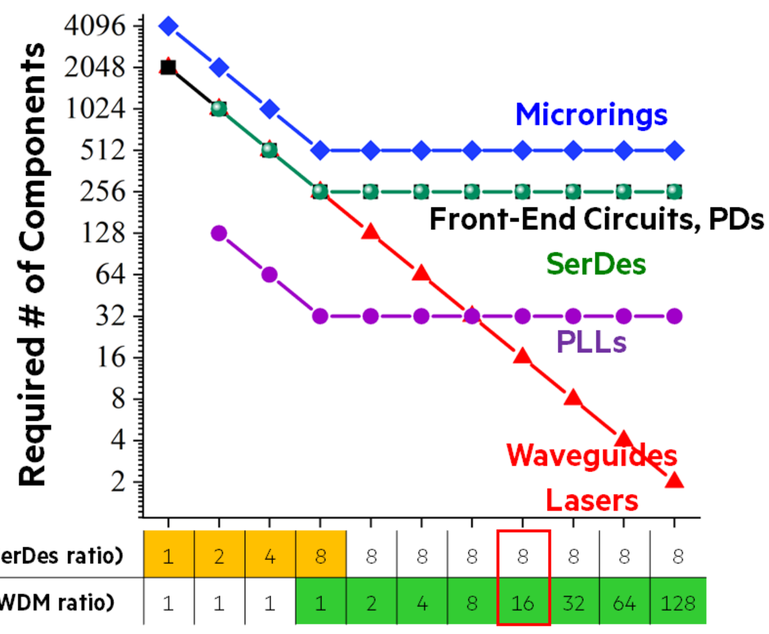

	Electrical Approach	Proposed Silicon Photonics Approach
Interposer	Silicon Interposer	**Silicon Photonics Interposer**
GPU-GPU Link	> 48 electrical routings	**2 waveguides** with CWDM
	20 Gbps per lane	**160 Gbps** per waveguide
GPU-HBM Link	16,384 electrical routings	**256 waveguides** with DWDM **98 % routing complexity reduction**
	2 Gbps per lane	**16 Gbps** per lane

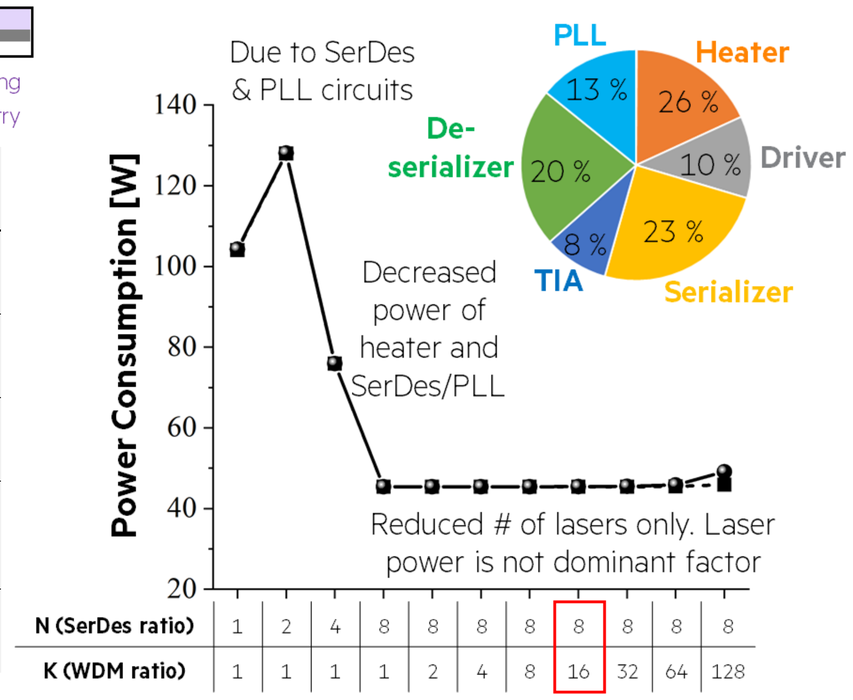

*Assumption of 5-mm GPU-HBM distance

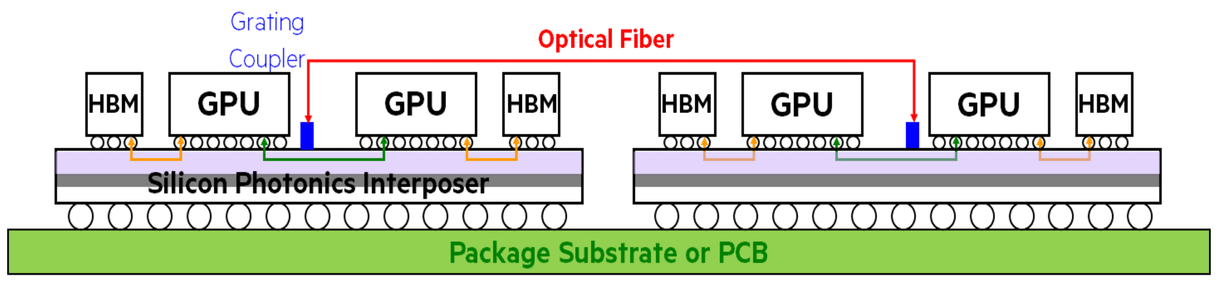

Conclusion
- Expanding our idea for **all-to-all (A2A) connectivity** in scalable GPU system
- Planning to further investigate system requirements (size, cost, power, etc.) and energy-efficient GPU architecture