Designing Multi-Socket Systems Using Silicon Photonics

Scott Beamer, Chris Batten, Ajay Joshi (MIT), Krste Asanović, Vladimir Stojanović (MIT)

Scott Beamer

Abstract

A manycore system will require tremendous bandwidth to memory, and current and projected electrical interconnects for off-chip communication may be unable to meet this challenge. In this work, we use monolithically integrated silicon photonics to construct a scalable interconnect between many cores and memory within a multi-socket system. Our interconnect also makes it possible to consider using smaller dies for cost reasons. Our design can be extended to a general template which can be configured to a particular die size and number of sockets.

Silicon Photonics Introduction

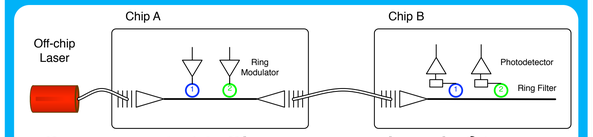

- Can transmit 10 Gbps per wavelength (λ)
- Can fit 64 λ per waveguide per direction

Silicon-Photonic Components

Waveguide

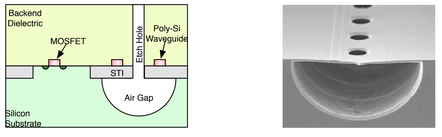

- Poly-silicon on top of shallow trench isolation

Resonant Ring

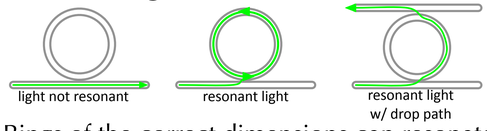

- Rings of the correct dimensions can resonate
- Can use heaters to tune rings (battle variation)

Ring Modulator

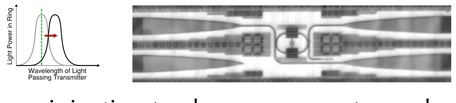

- Charge injection to change resonant wavelength

Ring Filter

- Filter used for wave-division multiplexing
- Use two cascaded rings to get additional frequency selectivity

Technology Comparison

		Electric	Photonic
Energy	On-chip	50fJ/mm	150fJ*
	Off-chip	5000fJ	150fJ*
Delay	On-chip	100ps/mm	200ps + 10ps/mm
	Off-chip	50ps + 5ps/mm	200ps + 5ps/mm
BW Density	On-chip	5 Gb/s/um	160 (320) Gb/s/um
	Off-chip	0.2 Gb/s/um	13 (26) Gb/s/um

- Electrical based on 22nm projection

Silicon Photonics Summary

- Photonic links excel over a distance since conversion cost is constant
- Photonic links are energy efficient off-chip
- Once conversion cost is paid, link can go seamlessly on and off chip
- Path layout on-chip can greatly impact power

Design Goals

- Connect cores to memory controllers
- High and uniform bandwidth
  - Use photonic advantage to overbuild
- Single photonic hop - best use of technology
- Co-design on and off chip networks to best use seamless links

Topology Insights

Concentation

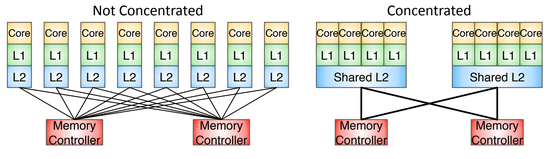

- Concentration improves link utilization and reduces serialization latency

Connections Off-Chip

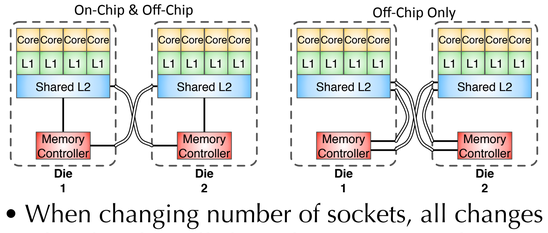

- When changing number of sockets, all changes localized to off-chip where easier to change
- Enables scalable system to have uniform bandwidth
- Group fibers by source (memory controller or cores)

Star Fiber Coupler

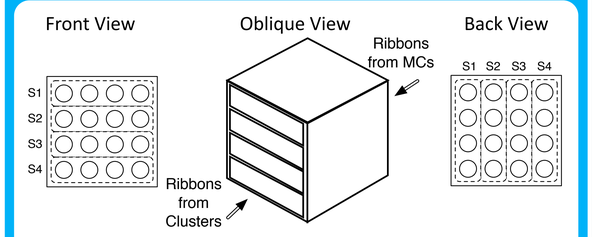

- Simplifies off-chip fiber wiring
- Same number of fibers, but now less ribbons

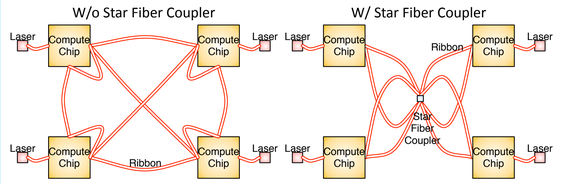

Physical Packaging

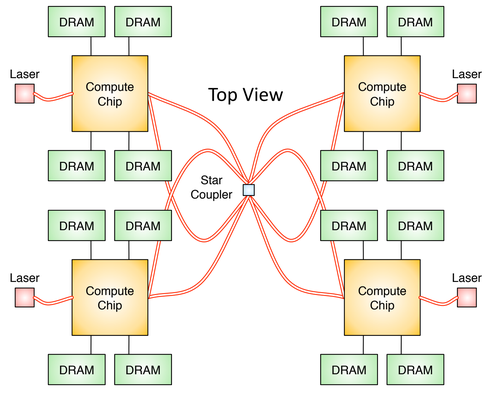

On-Chip Network Layout

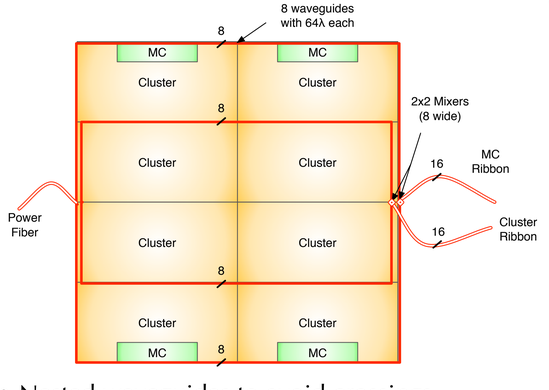

- Nested waveguides to avoid crossings
- Each cluster evenly loads nearby waveguides

Disintegration

Since all connectivity is done off-chip, why not build smaller dies?

- Smaller dies will increase yield which decrease total cost for the same amount of silicon
- A single reusable design will increase volume
- Smaller dies allow for binning on finer granularity to combat process variation
- A more spread out system will lower the power density, saving on cooling

Template Summary

- Not all scaling components presented here
- Template can handle:
  - Cores/die: 16, 32, 64, 128, 256
  - Cores/system: 64, 128, 256, 512, 1024

Network Power

Optical Power

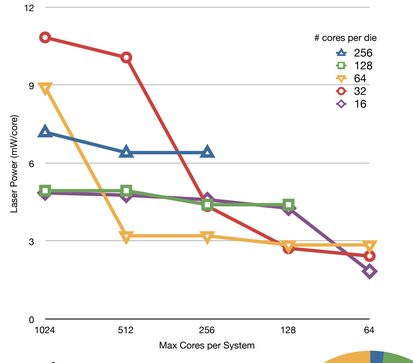

- Optical power 25-30% of electrical power needed to make it

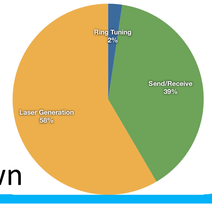

Power Breakdown

Conclusions

- Multi-socket systems are a place where silicon photonics could offer a large performance improvement
- Seamless photonic links encourage co-design of on-chip and off-chip networks
- A photonic interconnect could even allow for smaller dies (cost incentives)

Future Work

- Multi-hop networks to enable larger designs
- Coherence protocols (needed for real systems)
- Going closer to DRAM with photonics