Leveraging Optical Technology in Future Bus-based Chip Multiprocessors

Cornell University

Nevin Kırman **Meyrem Kırman** **Rajeev K. Dokania** **José F. Martínez**

Alyssa B. Apsel **Matthew A. Watkins** **David H. Albonesi**

Cornell University Computer Systems Lab

www.csl.cornell.edu

Published in The 39th Annual International Symposium on Microarchitecture (MICRO-39) – December 2006

Objective:

To assess the potential of on-chip optical interconnects to enhance the performance of future large scale chip multiprocessors.

Motivation:

- Tens or hundreds of cores on-chip in future technologies
- Need to provide sufficient data to each core
- Problems with metal global interconnects
  - Delay does not scale with technology
  - Metal wire sizing and more layers in metal hierarchy
  - Power and area consumed by repeaters
  - Increased latency from latch insertion
  - Crosstalk
- Benefits of optical interconnects
  - Speed-of-light propagation
  - Large bandwidth
  - Very low power
  - Good electrical isolation
  - Immune to crosstalk
  - Ease of design

Optical Interconnects

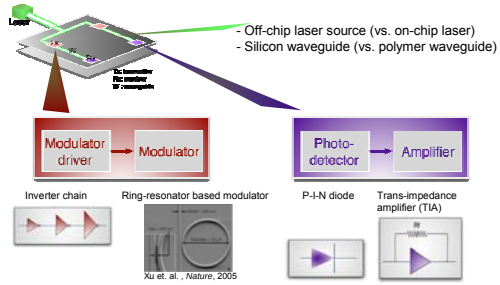

Considerations in applying optics

- Wave Division Multiplexing
  - Parallel data transmission by sending numerous wavelengths on a single waveguide
  - Cost-effective bandwidth boost
  - Increases with technology scaling
- Light power loses
  - Losses due to on-chip coupling, bending & crossing losses, modulator insertion losses, & receiver trapping losses influence necessary optical power

Interconnect Architecture

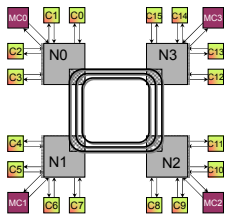

- Loop-like structure
  - Simple
  - Minimum waveguide crossings and bendings
- Nodes
  - Each serves one or multiple processors

Multiplex-by-node

- Each node transmits request using 1 wavelength & receives messages on all other wavelengths

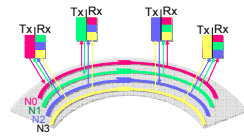

+ **No** arbitration required
- **#** of nodes limited to ***w*** at best
+ Relative position of transmitting node **known** at design time
+ In unidirectional transmission "sink"ing **easily achieved**
+ Num of transmitters O(w)
+ possible to optimize light power at design time through individual coupling-ratio tuning at each detector

System Parameters:

- 32 nm technology node
- 400 mm² die area
- 64 4-way cores running at 4 GHz
- 16 2MB L2 caches, each shared by 4 cores
- Optical off-chip communication providing 3 Tbits/s (384 GB/s) off-chip bandwidth
- 4 on-chip L3/Memory controllers

Area Estimation (mm²)

λs/node	1	2	3	2	3
Snoops/bus clk	4	8	12	4	4
Switches	1.71	2.72	4.00	1.93	2.13
Tx/Rx	0.39	0.78	1.17	0.56	0.72
Metal wiring	15.21	24.42	33.64	15.21	15.21
Optical layer	33.68	34.10	34.51	33.86	34.04

Power Estimation (W)

λs/node	1	2	3	2	3
Snoops/bus clk	4	8	12	4	4
Electrical comp.	14.57	23.62	33.00	14.88	15.19
Optical comp.	1.39	2.77	4.16	1.98	2.58
Total (α=1)	**15.56**	**25.60**	**35.98**	**16.30**	**17.03**

Architectural Evaluation

Electrical Baseline

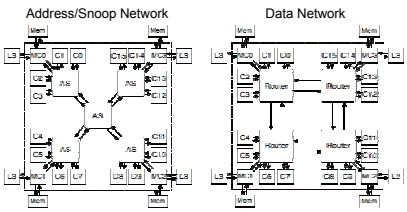

- Address/Snoop Network – Hierarchical tree
- Data Network – Bidirectional ring
- Choose 2 snoops/bus cycle for comparison as it most closely matches power consumption & active area of 12λ optical configuration

Snoops/ bus clk	Area (mm²)		Power (W)		
	Switches/Routers	Wiring	Switches/Routers	Wiring	Total (α=1)
2	1.47	15.9	1.42	13.40	14.82
4	1.66	22.81	1.68	19.23	20.91

Results

- 11 SPLASH2 Benchmarks
- Optical interconnect provide significant reduction in data transfer time with 2 or 3 wavelengths per node
- Speedups up to 1.71 for high bandwidth workloads like radix and ocean
- Average speedup of 1.14 over all applications

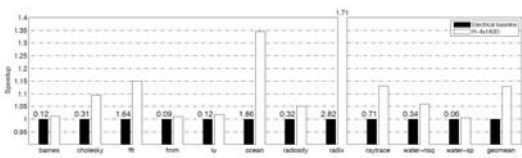

Areas for Future Research

- Best use of optical waveguides and wavelengths
  - Many wavelengths, crucial for performance
- Larger scale CMPs
  - May require different organizations
- Greater use of optics in the bus protocol
  - Fully optical instead of hierarchical
- Polymer as a waveguide material
- Temperature management of the optical components
- Other interesting topologies
  - E.g. optical switched networks

Challenges

- CMOS-compatible on-chip integration
- Low manufacturing cost
- High yields

Conclusions

- The proposed interconnect architecture provides a 14% speedup on average over all benchmarks and a maximum speedup of 71% compared to an aggressive electrical counterpart.
- Optical Interconnects have the potential to provide low latency, high bandwidth communication in future large scale CMPs