Optical On-Chip Networks for High-Performance, Energy-Efficient Multi-Core Architectures

Cornell University

Columbia University

David H. Albonesi, José F. Martínez, Keren Bergman, Luca Carloni, Matthew A. Watkins, Meyrem Kırman, Nevin Kırman, Assaf Shacham

Cornell University Computer Systems Lab
http://csl.cornell.edu

Columbia University
http://lightwave.ee.columbia.edu

Objective:

Co-architect high bandwidth optical switch architectures and fine grain Chip Multi-Processors to speed up high bandwidth-demand applications.

Primary Research Focus:

- Design a high bandwidth grid network implemented on a dedicated optical layer composed of very low power optical components
- Co-architect a fine-grain CMP system of low-power cores that enables learning algorithms to fully exploit high bandwidth, low power communication

Building Blocks

Photonic Switching Element (PSE)

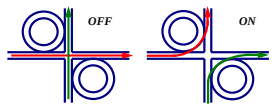

- Broadband ring-resonator switch
- ON state:
- carrier injection
  - coupling into ring
  - signal switched
- OFF state
  - passive waveguide crossover
  - negligible power.

Broadband ring-resonator switch

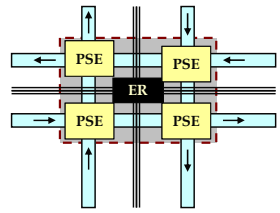

- 4 PSEs grouped with electronic control
- 4 waveguide pairs I/O links
- Electronic router
  - High speed, simple logic for fast arbitration time
  - Electric links optimized for low latency
- Small footprint
  - 70 μm x 70 μm (including spacing)
  - Area < 0.005 mm^2
- Nearly no power consumption in OFF state

Network Architecture

- Optical data network grid
  - 2x spatial overprovisioning (i.e. two 4x4 switches per gateway per dimension)
  - Message generation occurs once and then messages are routed for free.
  - Larger networks consume the same amount of power

- Electronic control network
  - Manages photonic network
  - Path-setup packets sent in advance of data to reserve path for photonic messages
  - Distributed adaptive routing algorithms to minimize latency and power consumption

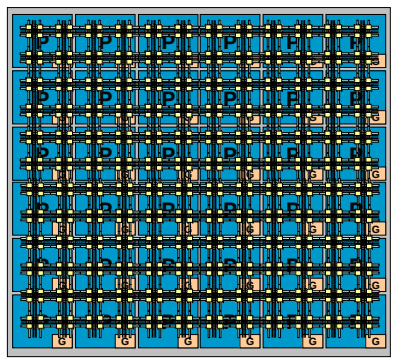

- Traditional electronic on-chip networks:
  - Huge amounts of data are buffered and regenerated at every hop.
  - The larger the network – more hops, more regenerations, and more power!

Architecture

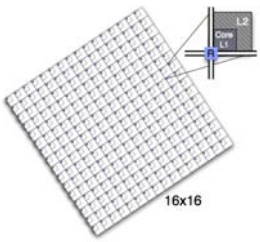

- 22 nm process technology
- 256 PowerPC-464-like 2-issue superscalar cores on a die
- 1MB L2 cache per core for a total of 256 MB of on-chip cache
- 4GHz core frequency

Interconnect Characteristics

4GHz, α = 1 @ 22nm	Proposed Optical Network		Electrical baseline (32 bits)
	Setup (16 bits)	Optical part	
Area	5.38mm^2	23.3mm^2 *Active layer 25.5mm^2 *Optical layer	11.67mm^2
Power	28.8W	37W	57.6W
Bisection BW	256GB/s	10.24TB/s	512GB/s
Latency *4KB data	60 cycles	103 cycles	1084 cycles

• **Routers:** 5 input, 5 output ports (E,W,S,N,Local), 4 virtual channels, 4-entry input buffers

For equal power → 20x bandwidth improvement
→ 10x latency improvement

4GHz, α = 1 @ 22nm	Proposed Optical Network		Electrical baseline (640 bits)
	Setup (16 bits)	Optical part	
Area	5.38mm^2	23.3mm^2 *Active layer 25.5mm^2 *Optical layer	933.5mm^2
Power	28.8W	37W	1152W
Bisection BW	256GB/s	10.24TB/s	10.24TB/s
Latency *4KB data	60 cycles	103 cycles	112 cycles

• **Routers:** 5 input, 5 output ports (E,W,S,N,Local), 4 virtual channels, 4-entry input buffers

For equal bandwidth → 17.5x less power
→ 173.5x less area

Applications

Data Mining

- Mine medical, financial, communication, travel, sensor, scientific, spatial, temporal data
- Find structure, rules, associations, clusters
- Algorithms:
  - –Clustering : K-means, k-medoids, agglomerative, divisive
  - –Frequent item set mining
  - –Graph mining
  - –Association rules mining

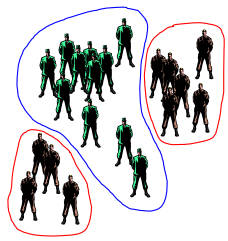

Strategic Data Analysis

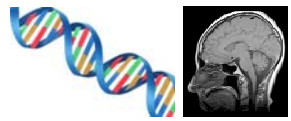

Medical & Scientific Profiling

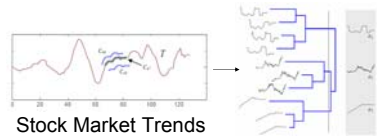

Stock Market Trends

Satellite Image Recognition

Speech Recognition

Moving Object Recognition

Supervised Learning

- Learn to predict target from inputs
- Algorithms:
  - –Artificial Neural Networks
  - –Decision trees
  - –K-Nearest neighbor
  - –Support vector machines

Autonomous Vehicle Control