Cornell University

Phastlane: A Rapid Transit Optical Routing Network

Mark Cianchetti, Joseph Kerekes, and David Albonesi
Computer Systems Laboratory
Cornell University

The Interconnect Bottleneck

- Future processors: tens to hundreds of cores
- Dire need for fast and power efficient on-chip interconnect
- Nanophotonics
  - Ultra-fast signal propagation
  - High bandwidth density
  - Low power

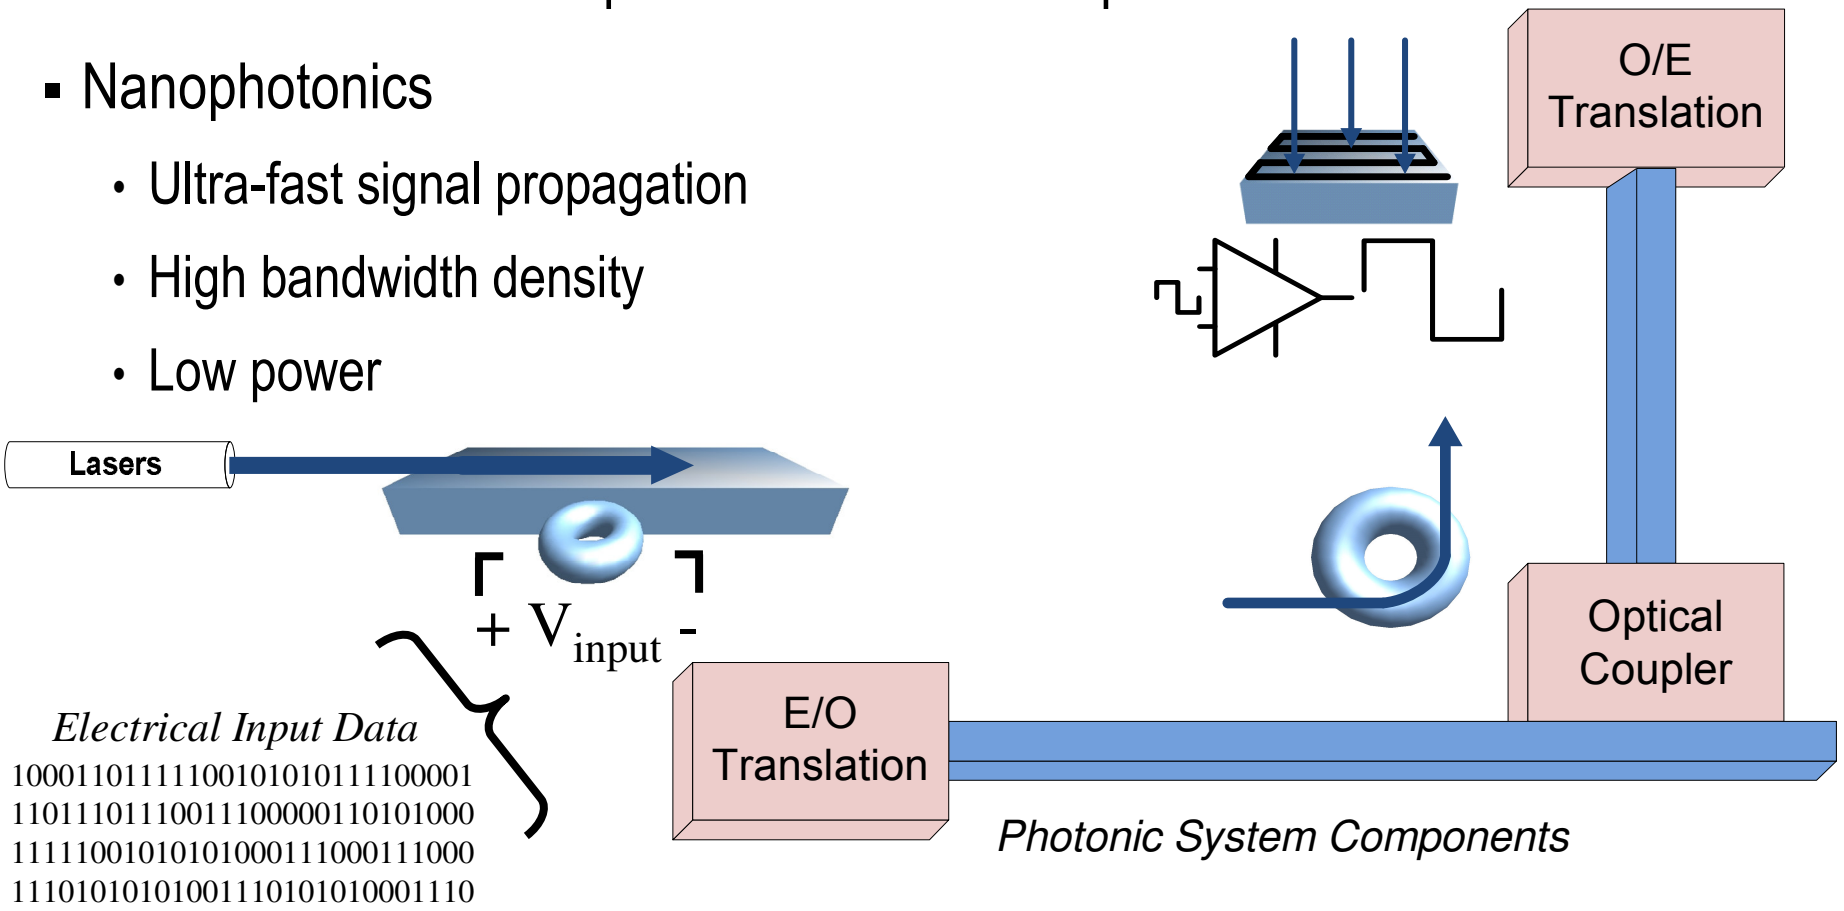

Limitations of Nanophotonics

- Lack of fundamental building blocks
  - Logic gates
  - Memory structures
- Single routing layer
  - Power loss of waveguide crossings

Traditional Approach To Exploiting Photonics

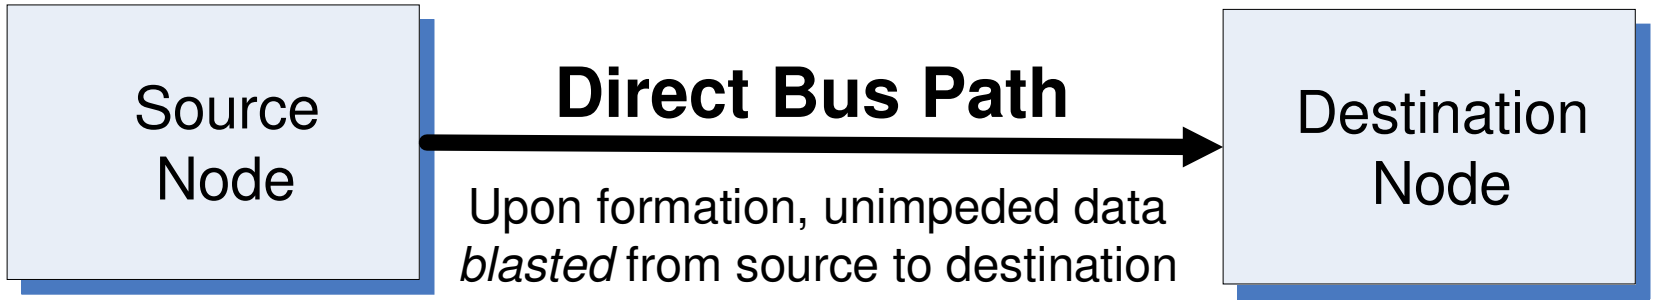

Respecting the Limitations…

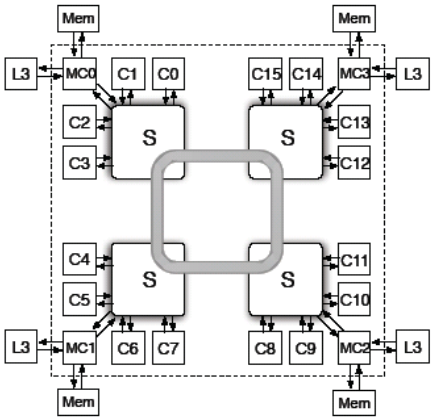

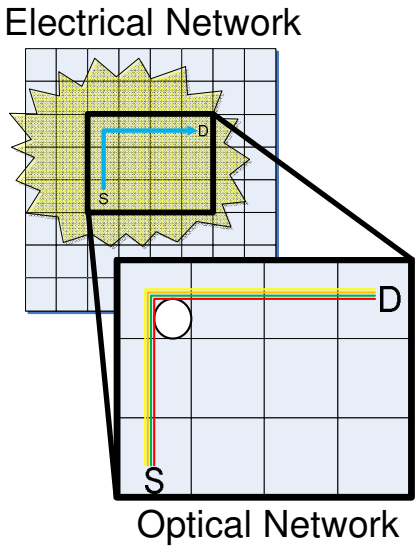

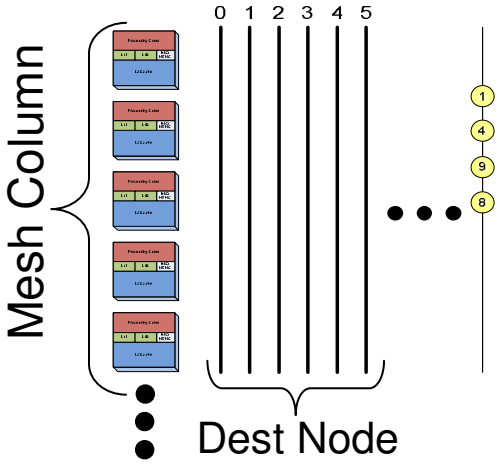

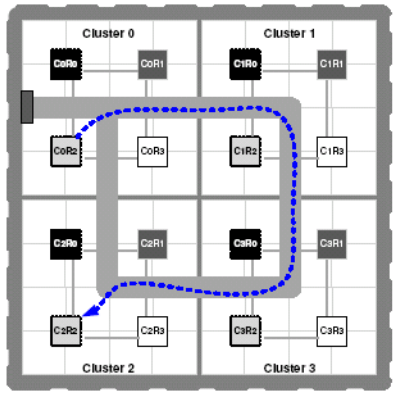

Cornell Ring Architecture *Columbia Architecture* *Corona Architecture* *FireFly Architecture*

- Previous proposals largely bus-based
- Direct photonic links between source and destination
- Data blasted in a single optical transmission

Overview of Previous Work

Feature	Cornell Ring	Corona	Columbia	Phastlane
Network Topology	Ring-bus	Snake XBAR	Torus	Mesh
Network Operation	Snoopy Bus	Fully Connected XBAR	Electrical Setup	Packet Switched
Shared Resources	None	Destination Bus	Router Channels	Router Channels
Shared Resource Arb.	WDM	Token	Per Hop	Per Hop
Unit of Data Transfer	Cache Line	Cache Line	>> Cache Line	Cache Line

Phastlane Contributions

- Novel nanophotonic router architecture
- Hybrid optical/electrical packet-switched, mesh network
  - Nanophotonics: High-speed, multi-hop packet transfers
  - Electrical: Logic and buffering to handle packet conflicts

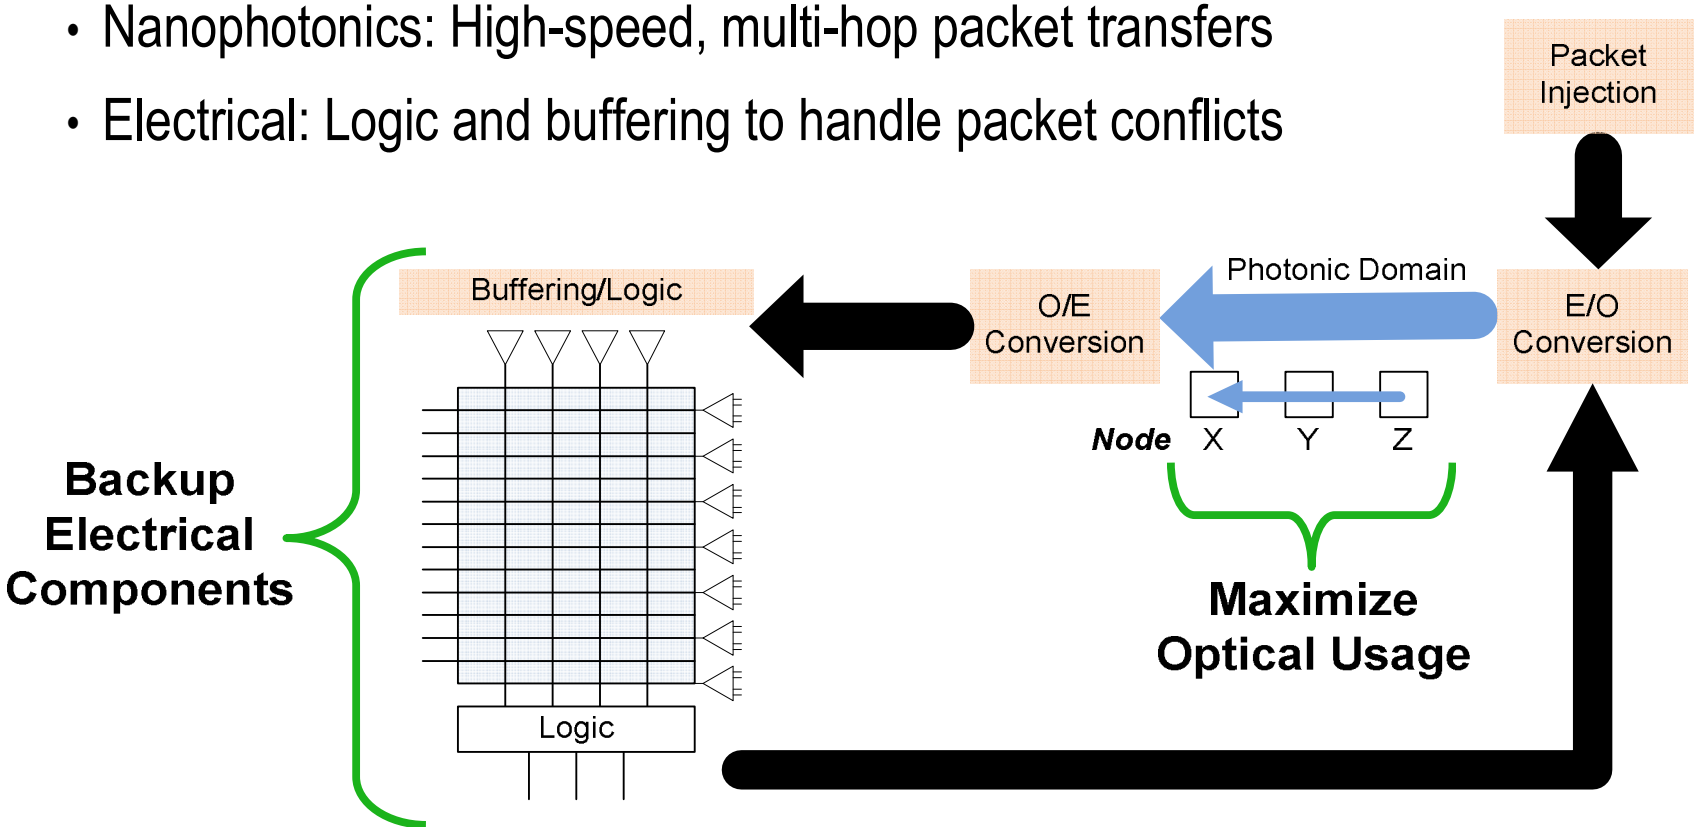

Phastlane Architecture

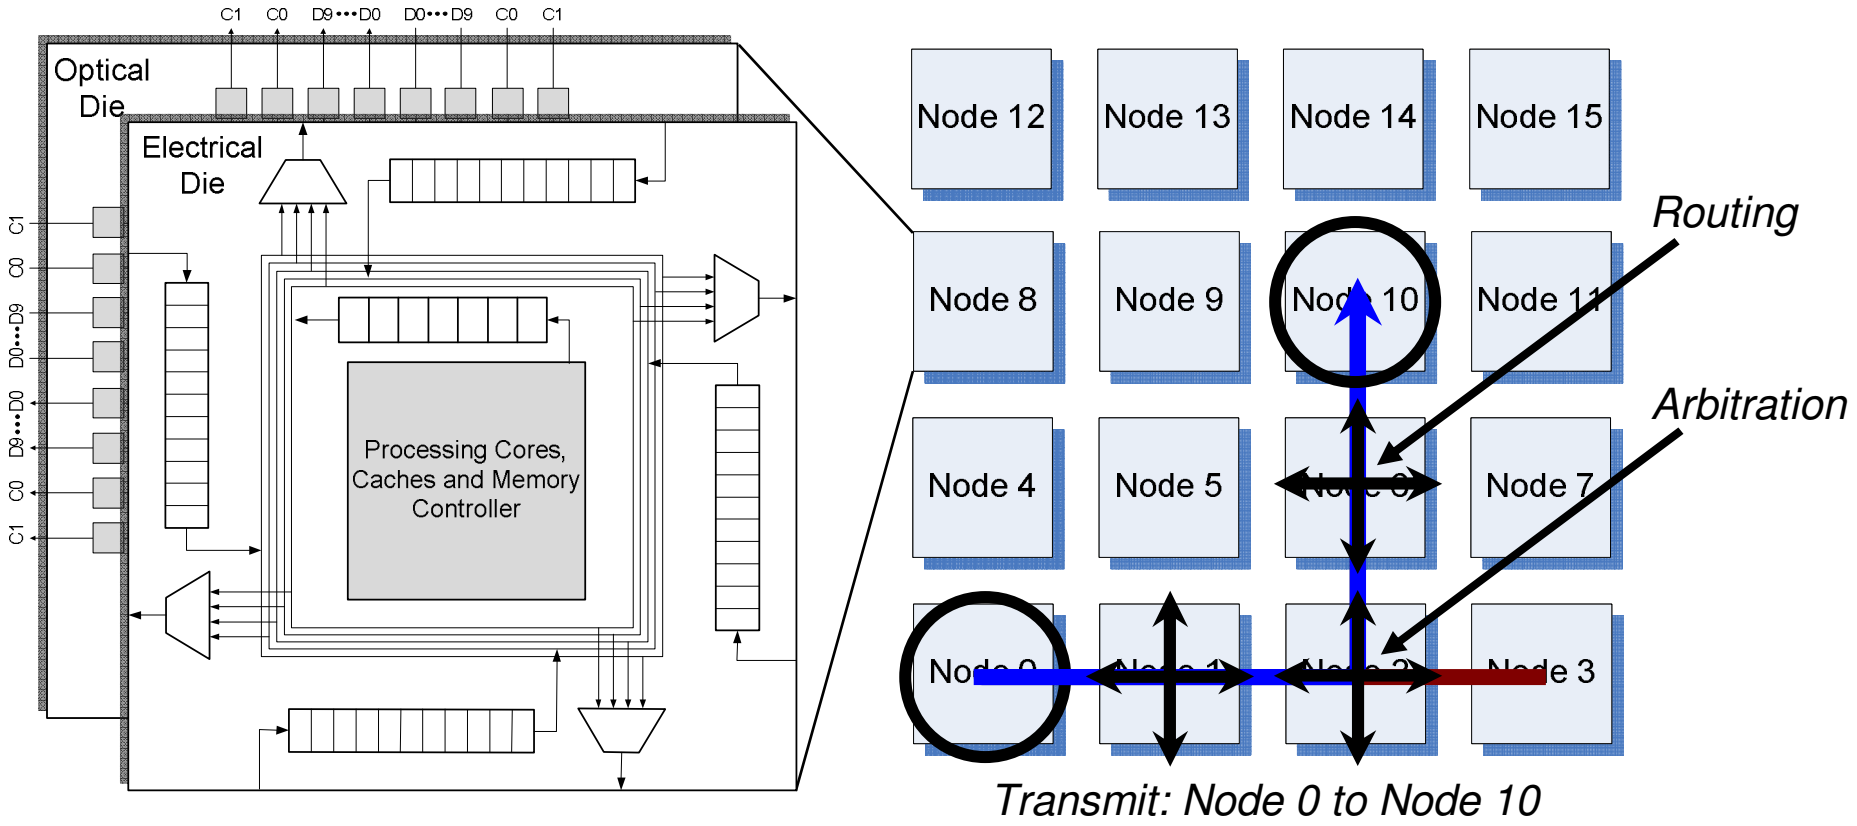

- Minimally impede flow of light from source to destination
  - Simplified routing, arbitration
- Dual die configuration

Optical Node Operation

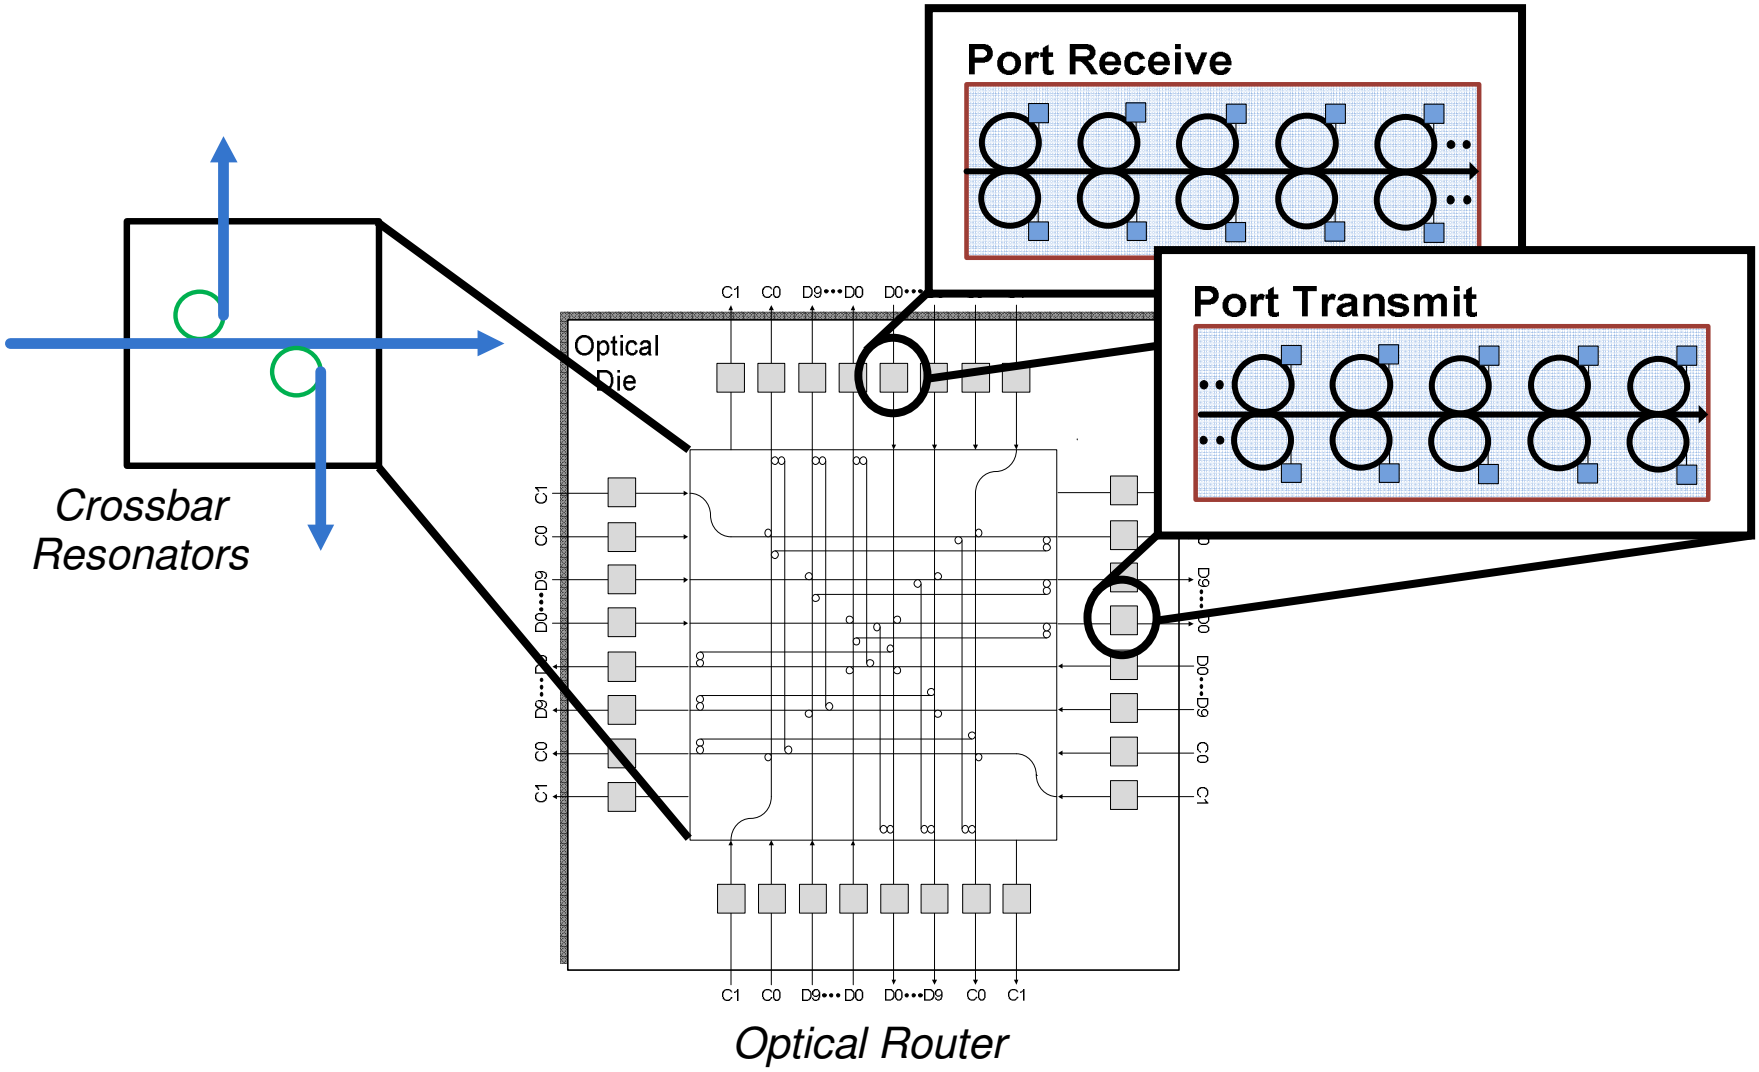

Electrical Node Operation

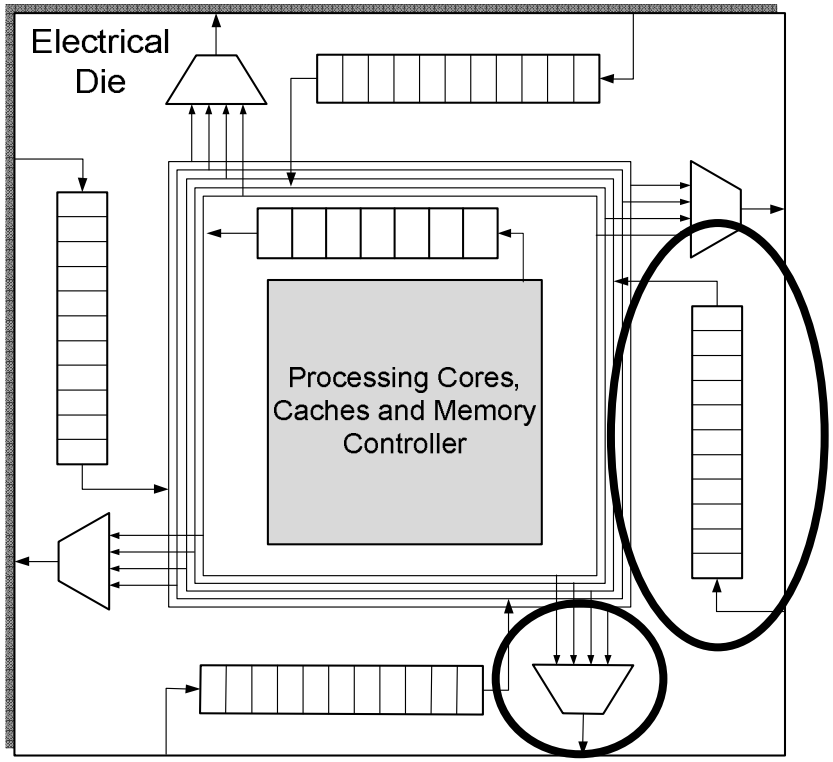

Electrical Node

- Blocked packets are buffered locally
  - Per port buffers, single processor buffer
- Output multiplexers connect to output port optical transmitters

Simplified Packet Routing

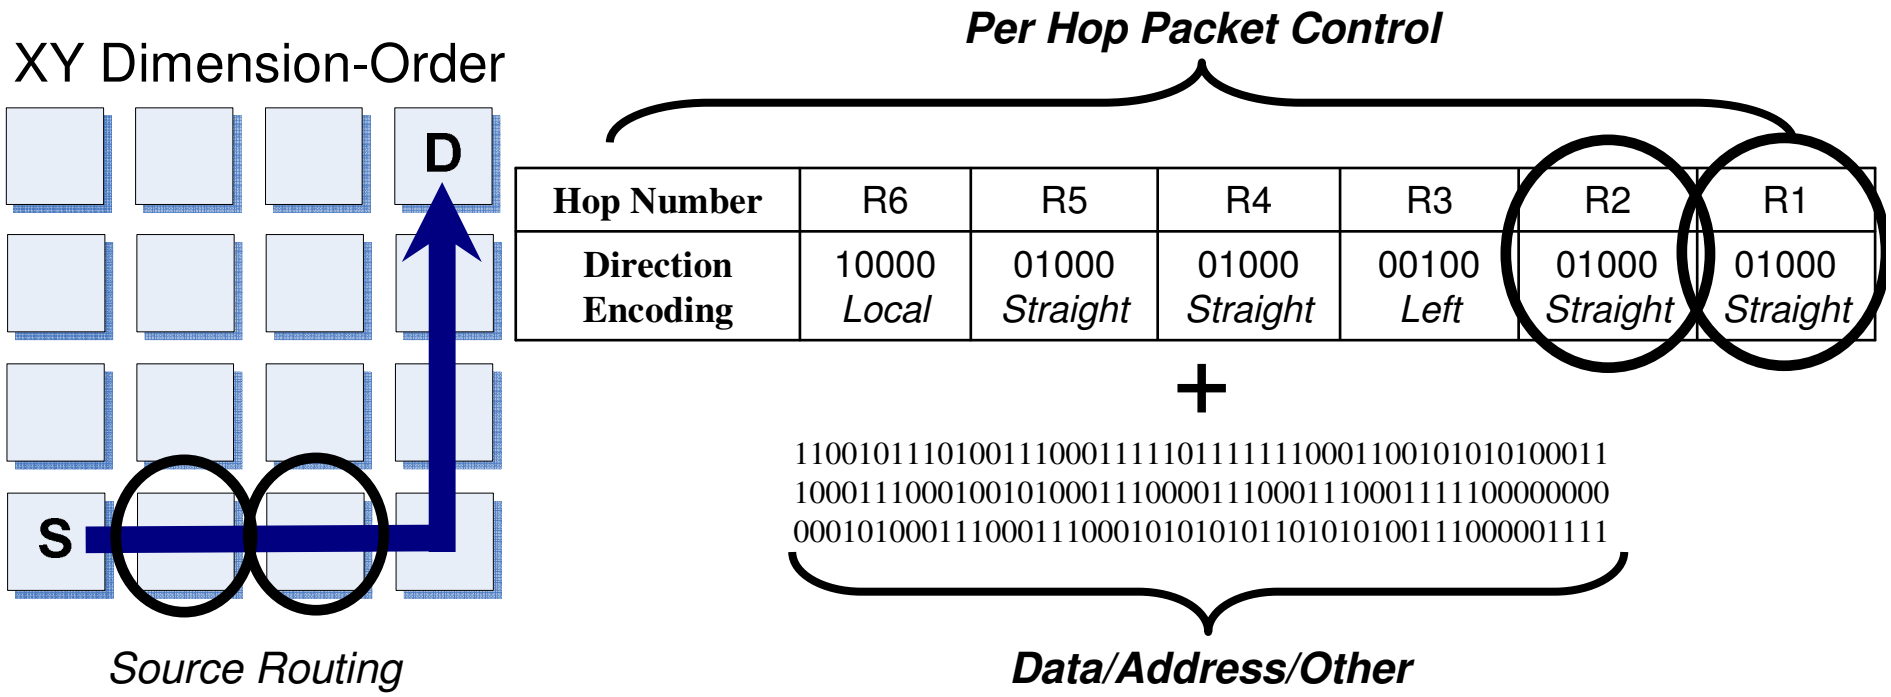

- Packets are dimension-order, source routed
  - Portion of packet holds pre-computed routing bits
- Per hop control bits enable near-instant switch traversal

Per Hop Control Bits

Hop Number	R6	R5	R4	R3	R2	R1
Direction Encoding	10000 *Local*	01000 *Straight*	01000 *Straight*	01000 *Straight*	00100 *Left*	01000 *Straight*

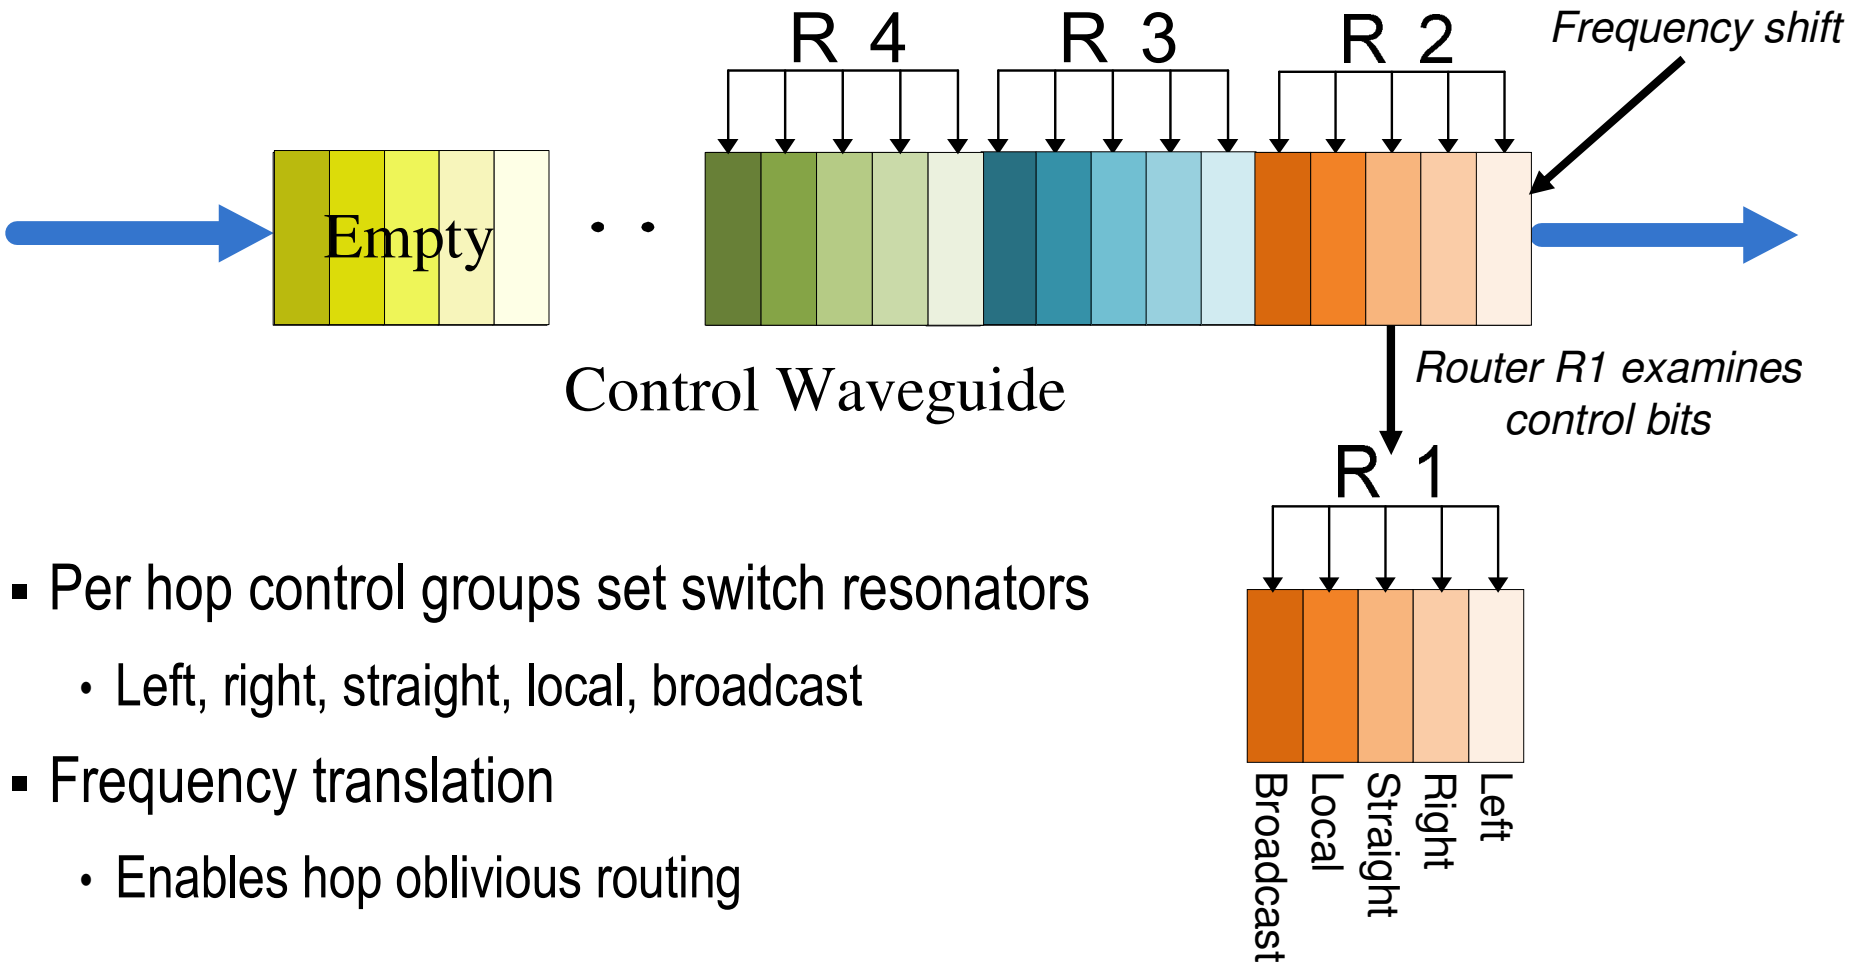

- Per hop control groups set switch resonators
  - Left, right, straight, local, broadcast
- Frequency translation
  - Enables hop oblivious routing

Per Hop Control Bits Continued

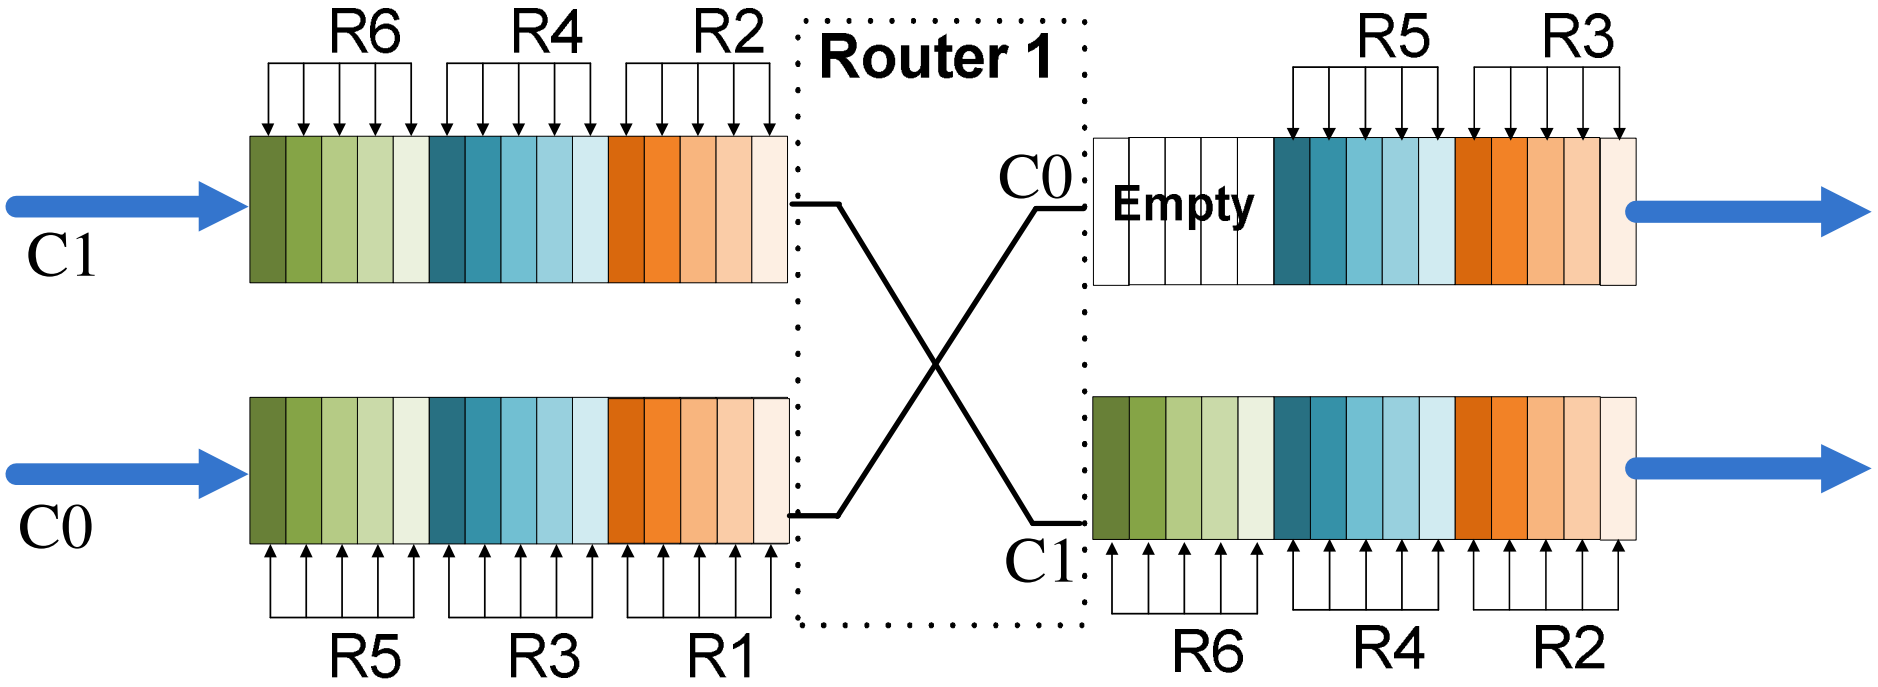

- Control waveguide split into C0 and C1
  - WDM limitations
- Frequency and physical translation
  - Enables hop oblivious routing

Competing Packets: Resolution

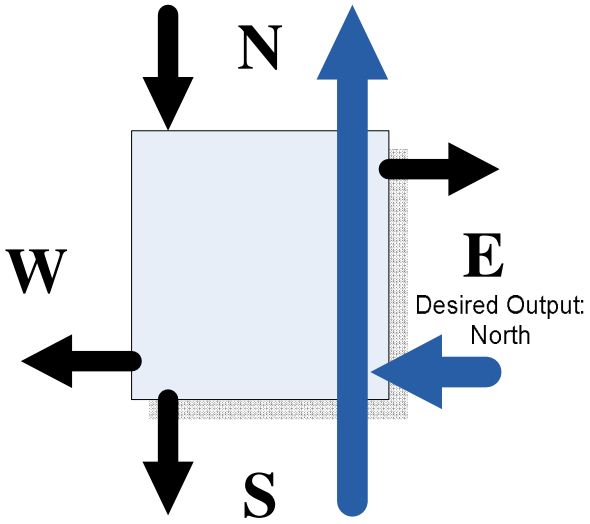

Two packets collide

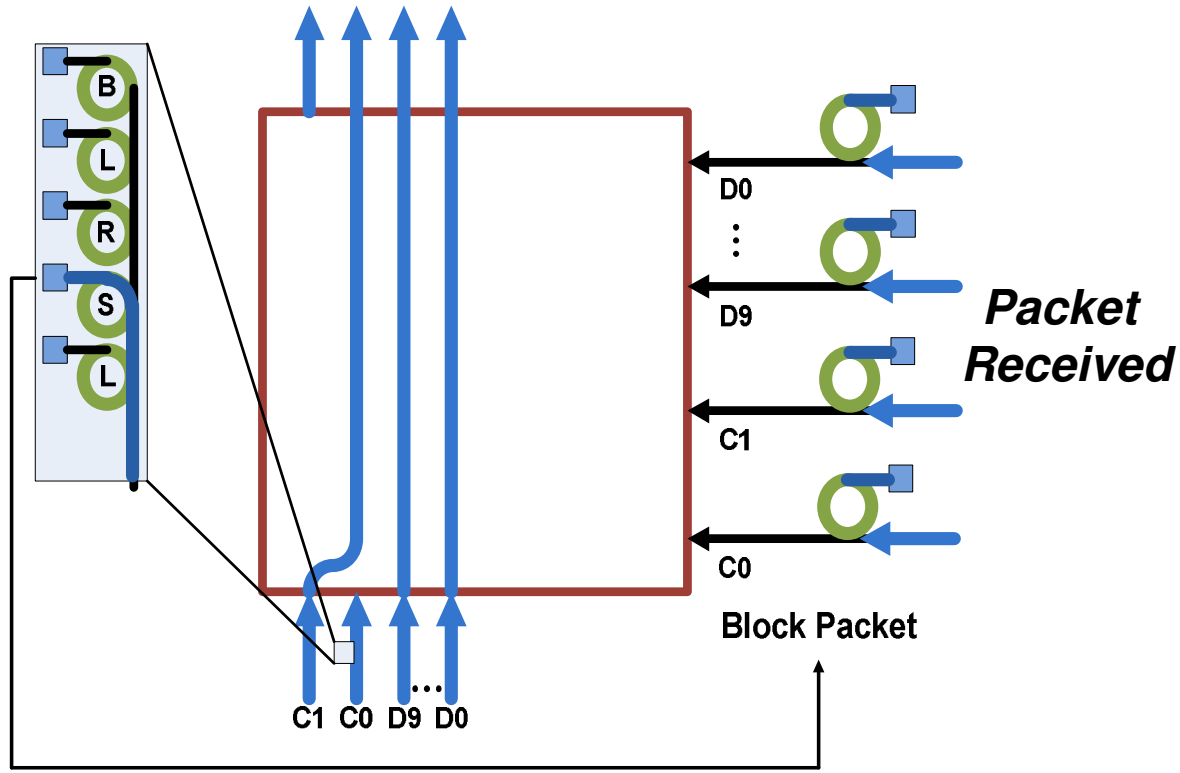

Collision resolution

- Packet collisions avoided by packet blocking
- Collision resolution performed “on the fly”

High-Speed Drop Network

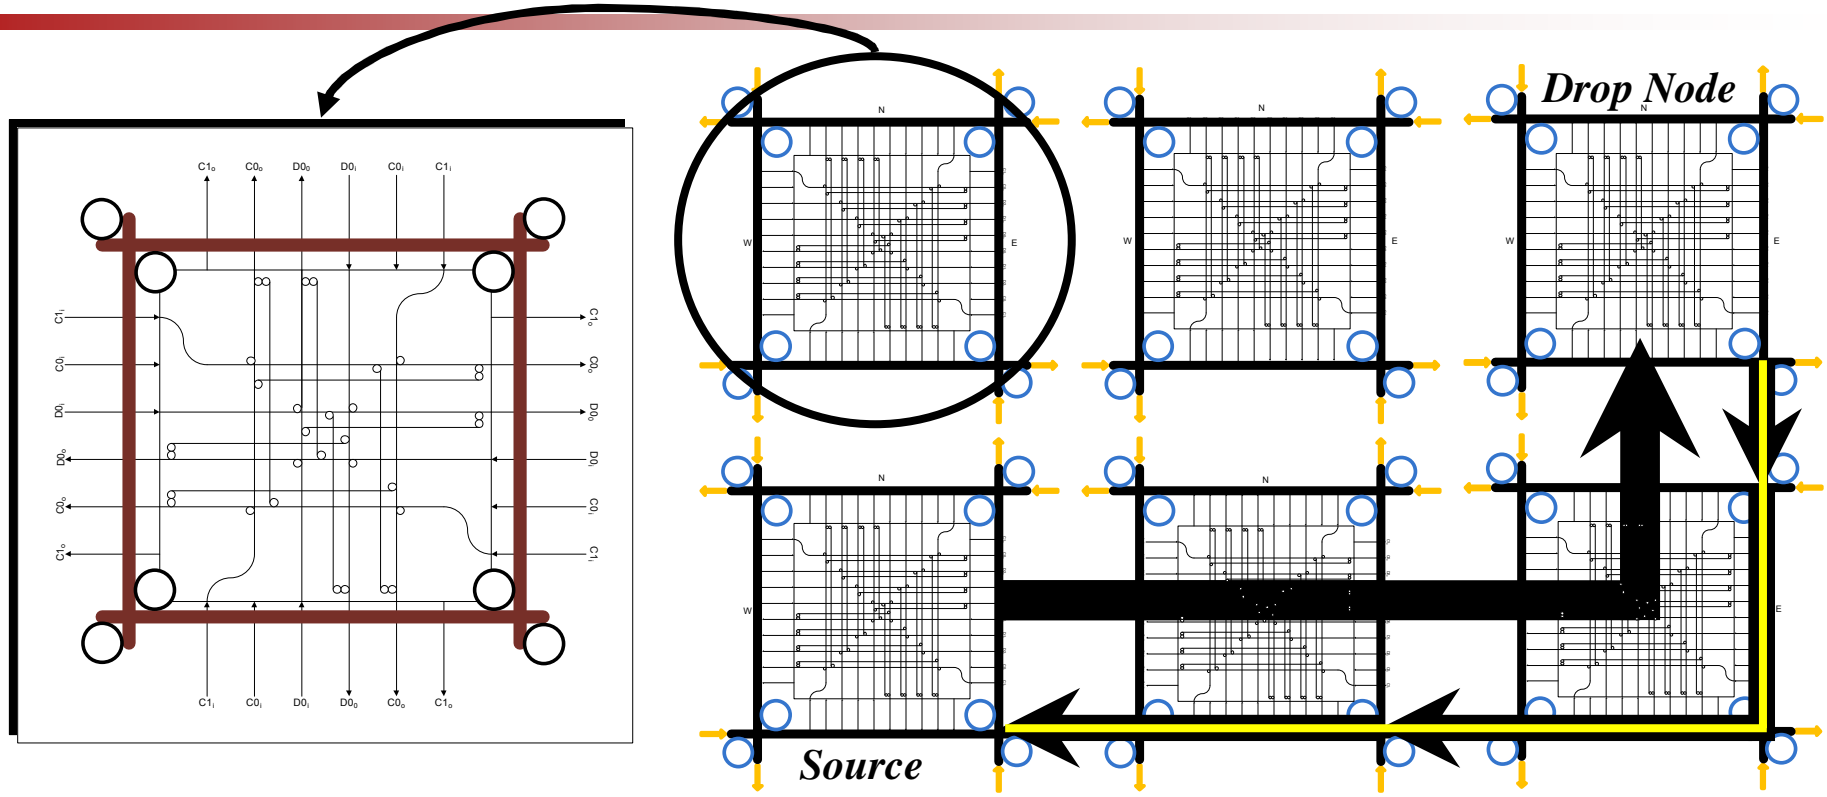

- Drop path gradually formed as packet passes through network
- Packet dropped if downstream buffer full
- Drop signal travels opposite data packet

Router Design Space Exploration

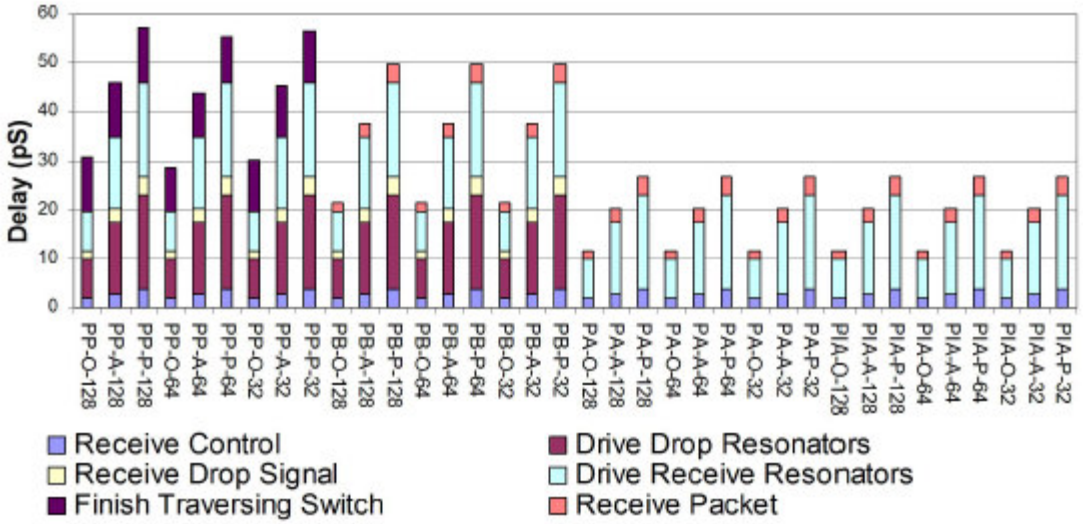

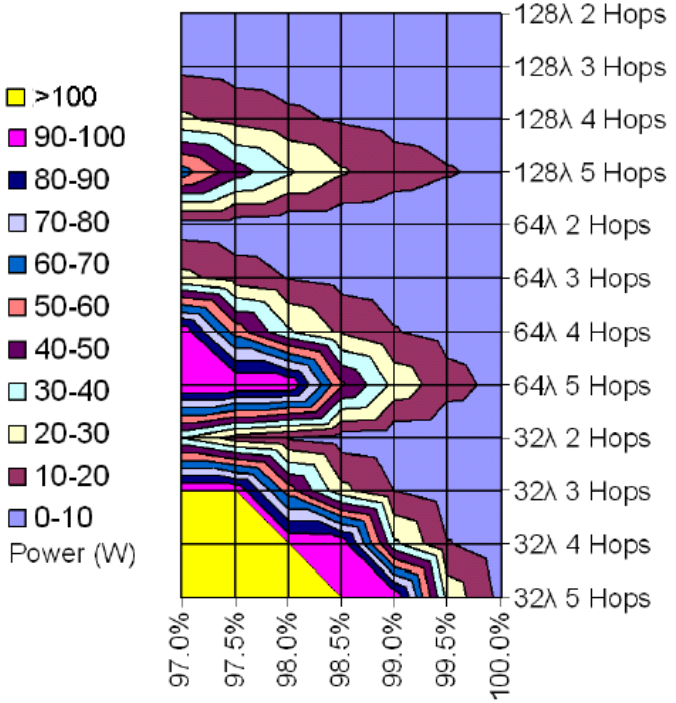

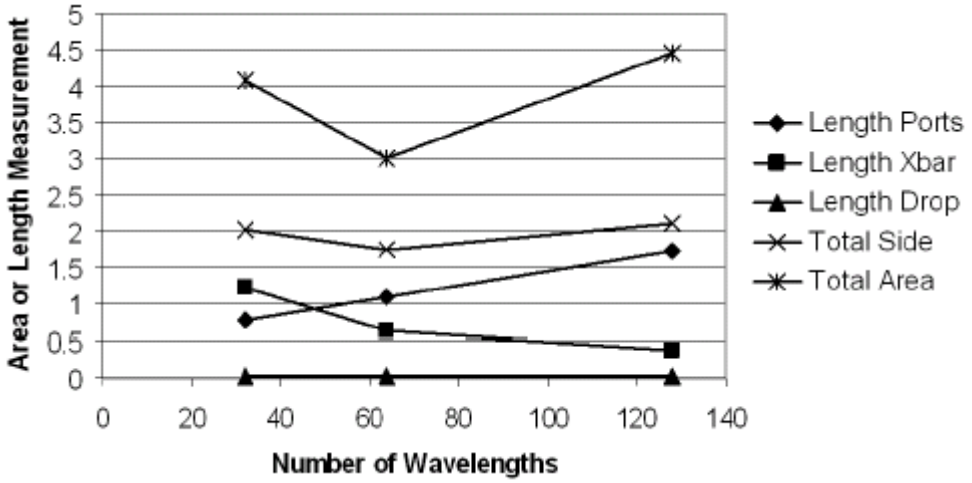

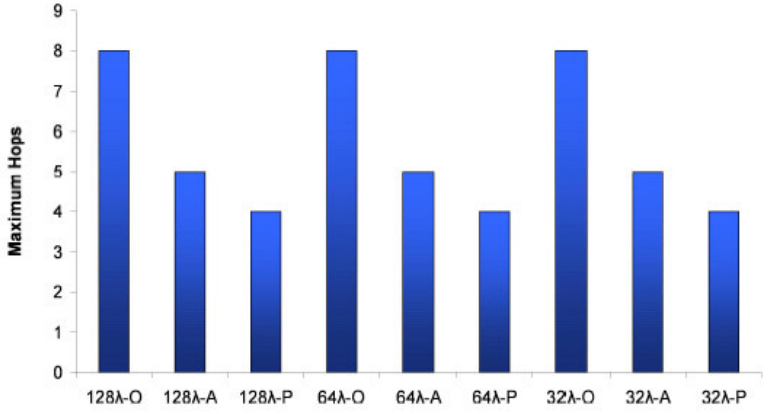

Evaluation Methodology

Routing Function	Dimension-Order
Number of VCs per Port	10
Number of Entries per VC	1
Wait for Tail Credit	YES
VC_Allocator	ISLIP
SW_Allocator	ISLIP
Total Router Delay	2 or 3 cycles

Baseline Configuration

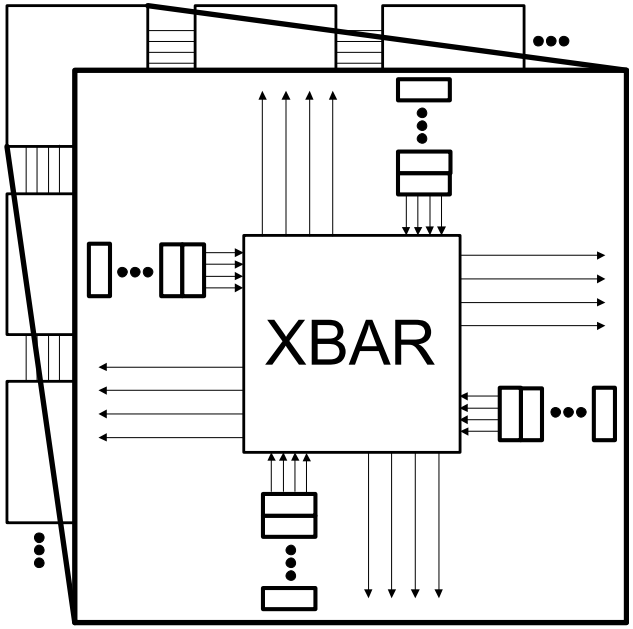

Baseline Node

- 64 node, 8x8 mesh network topology
- Low latency, high bandwidth electrical network baseline
  - Multi-port receive for destination packets
- Packet size: 80 bytes, single flit
- Synthetic and Splash benchmarks

Synthetic Benchmark Results

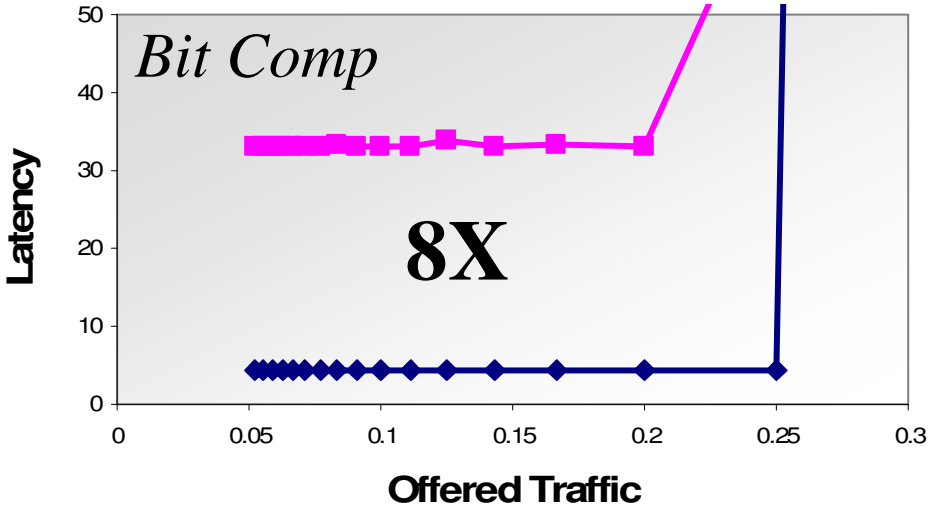

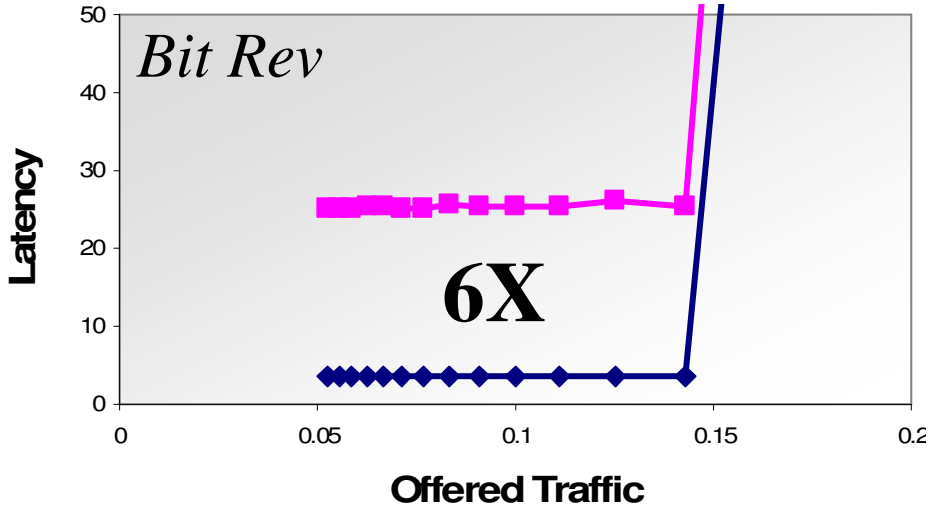

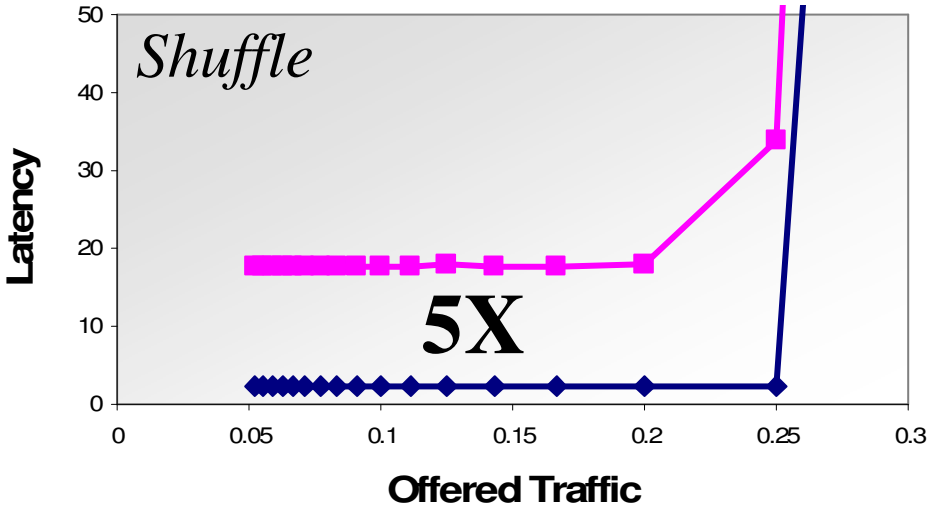

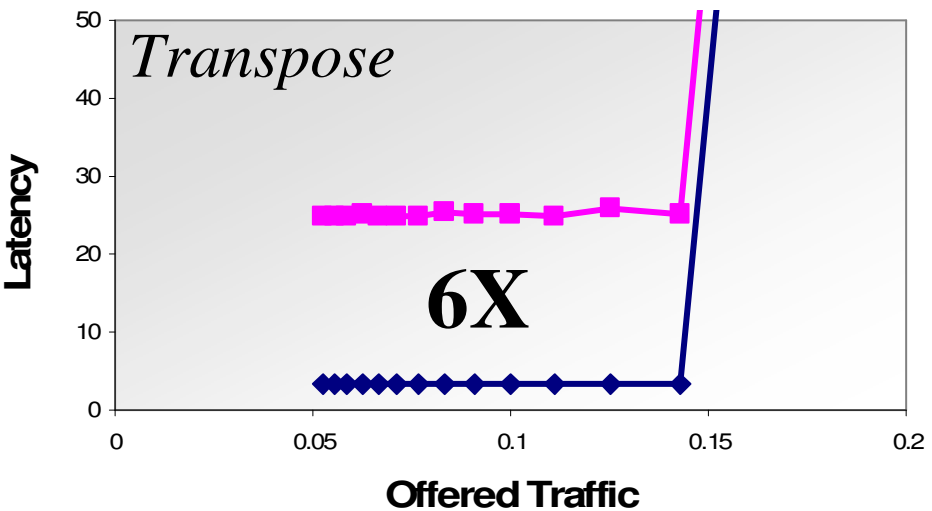

Splash Performance Analysis

- 2X speedup across all benchmarks

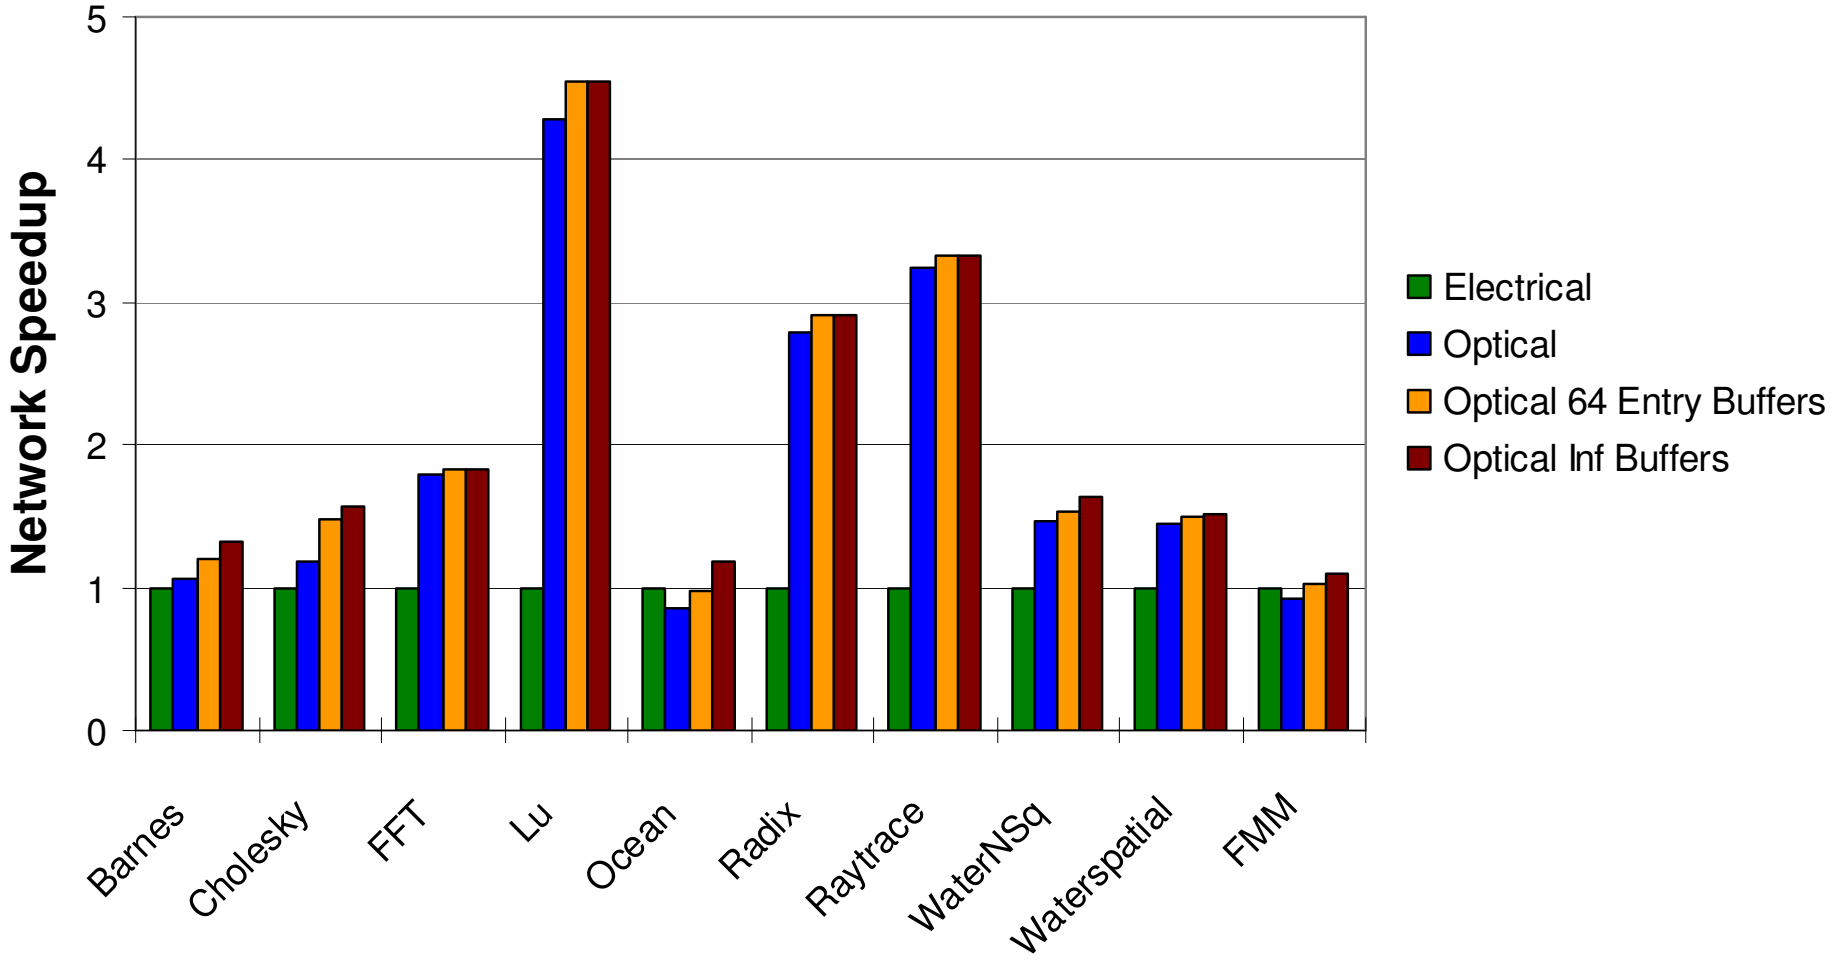

Splash Power Analysis

- 80% reduction in power across all benchmarks

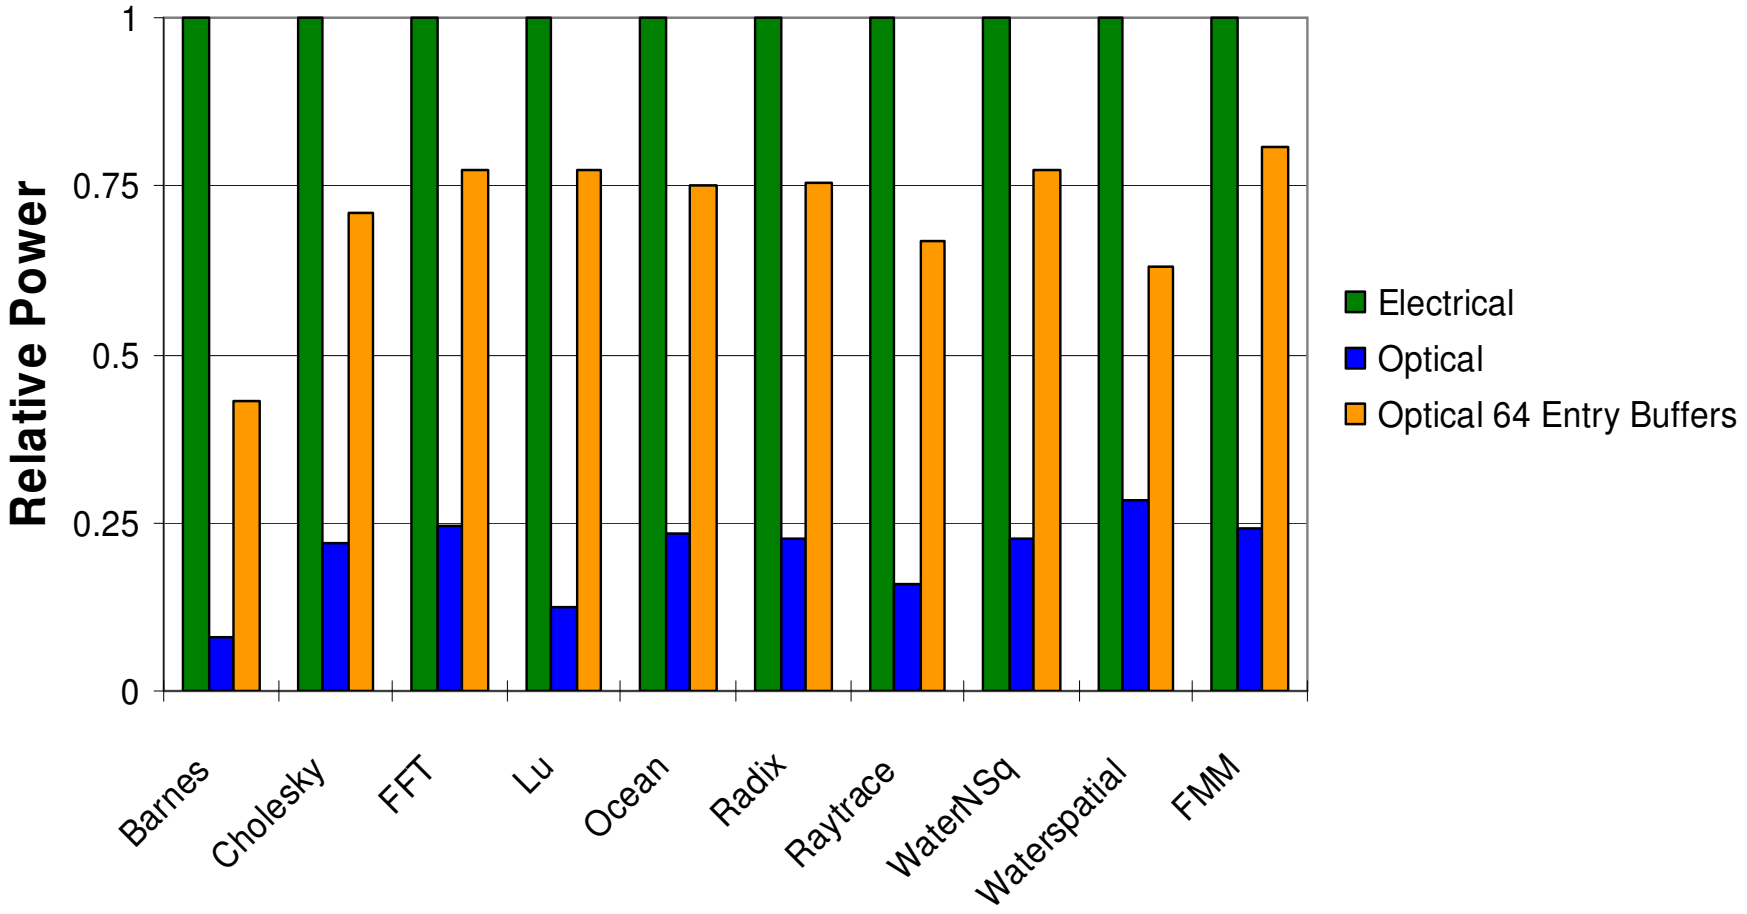

Conclusions

- Novel nanophotonic router architecture
- Packet-switched, hybrid optical/electrical mesh network
- Up to 8X performance improvement for synthetic workloads
- Up to 4X performance improvement, 80% power reduction, for Splash
- Future work
  - Lower power broadcast scheme
  - Improved allocation and flow control