An Integrated Photonic Network for Multi-Processor Applications

Assaf Shacham and Keren Bergman
Columbia University

HPEC 2005
MIT Lincoln Laboratory
September 2005

Interconnection Networks for HPCS

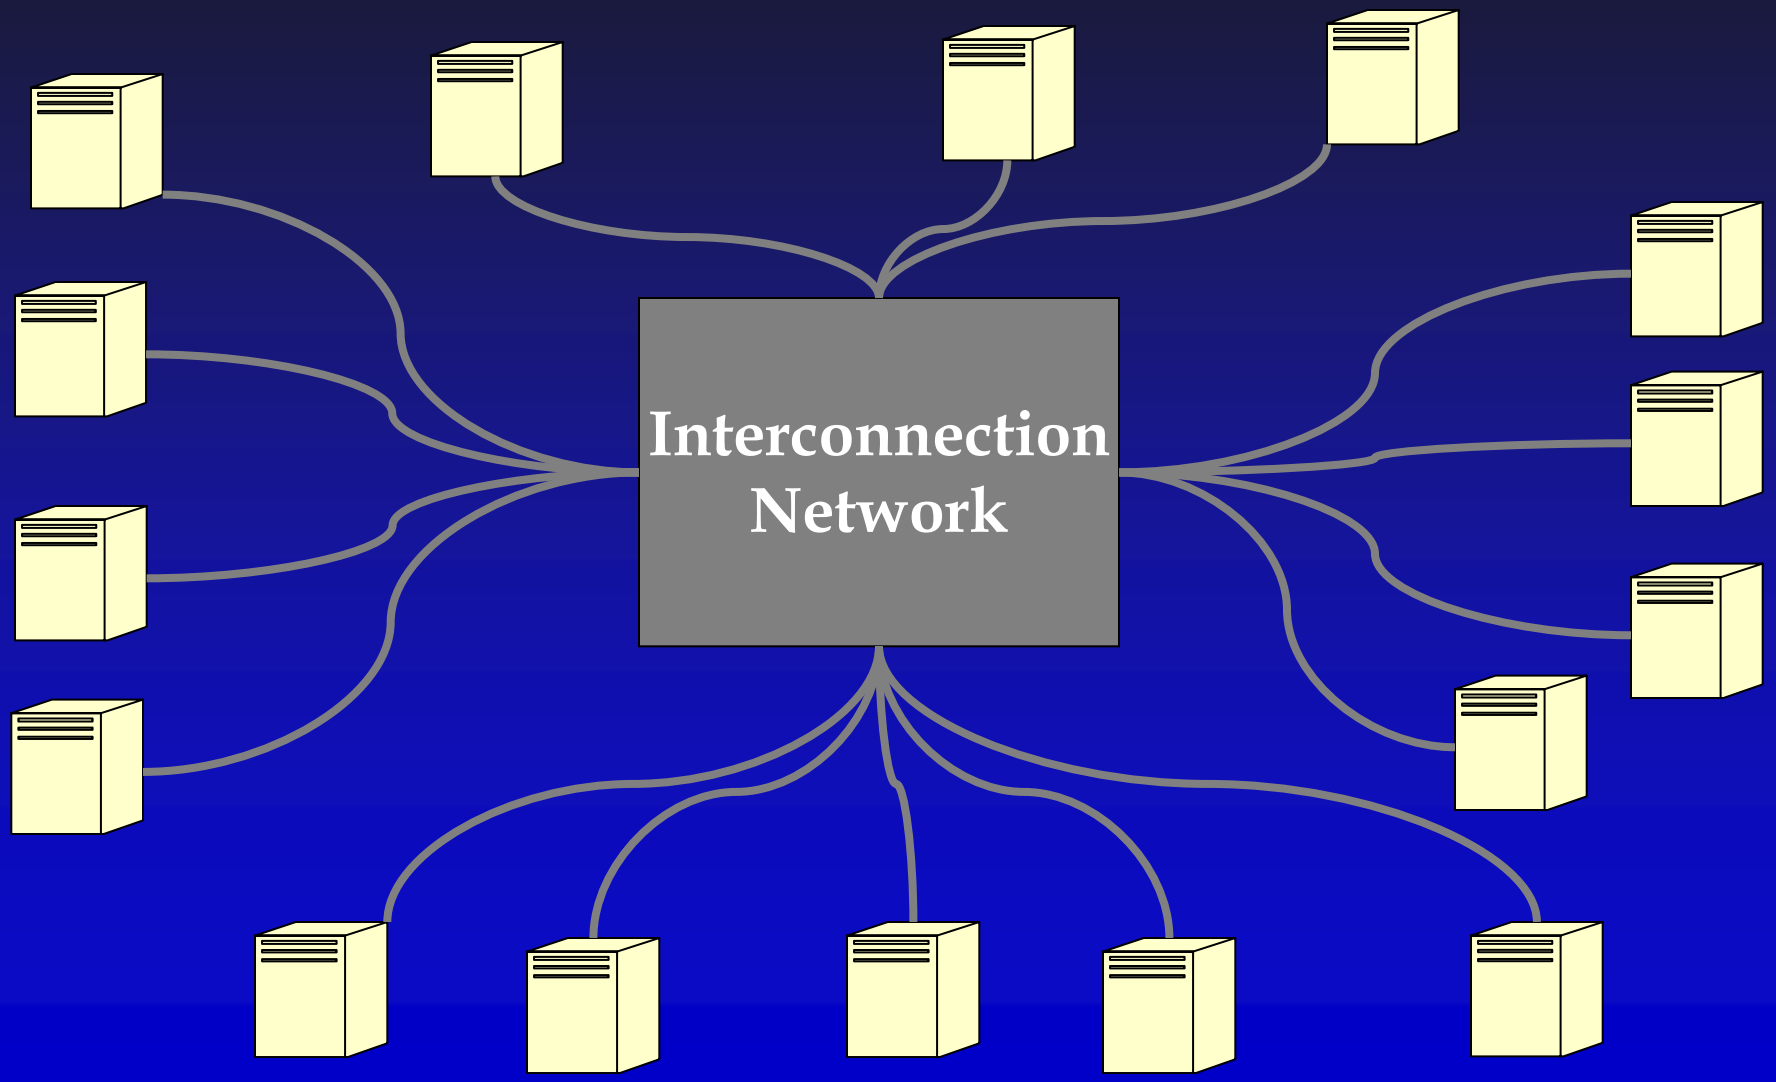

The Electronic Bottleneck

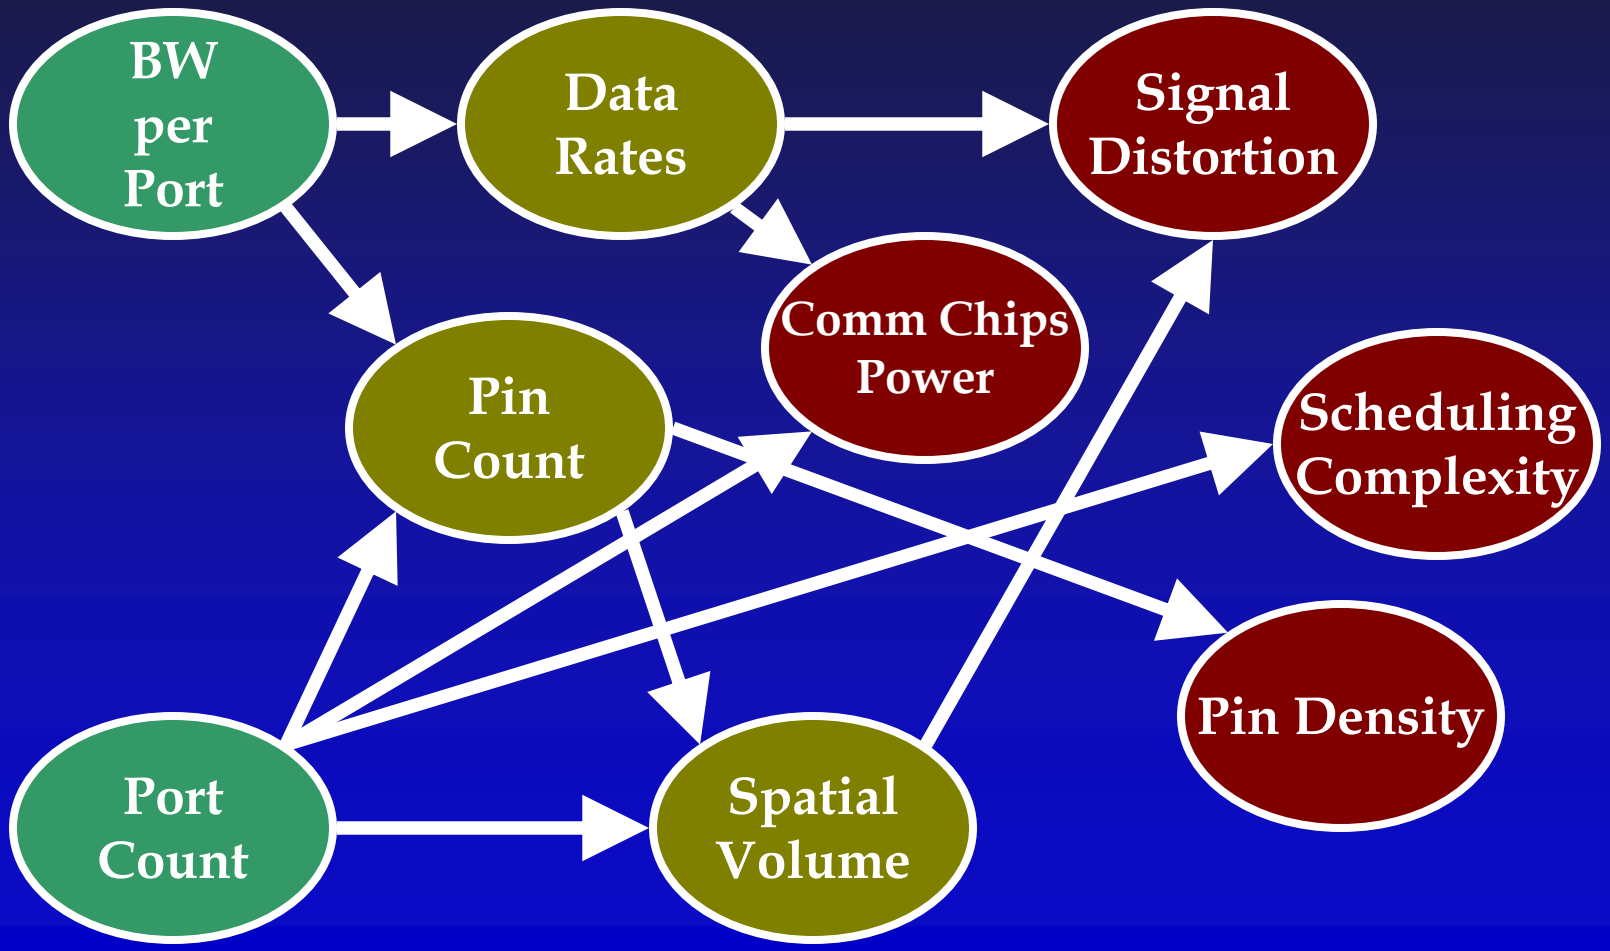

Transparent Photonic Networks

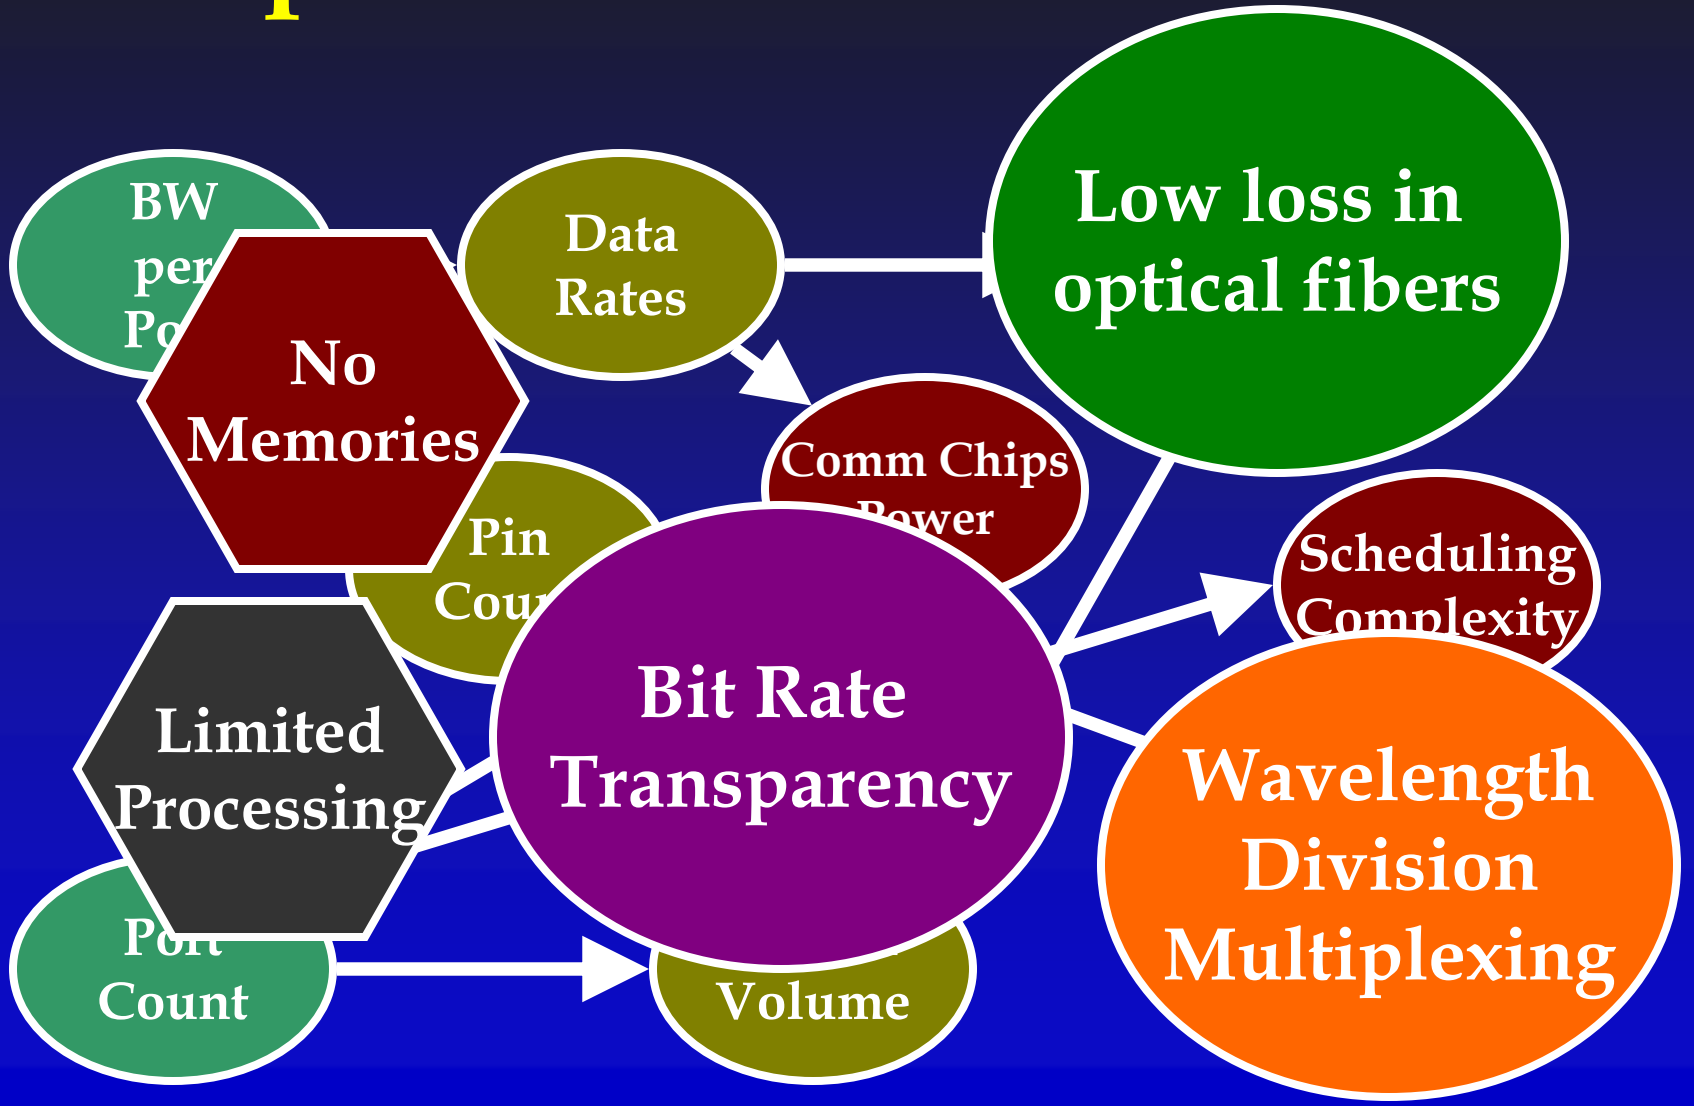

Outline

- Introduction
- Motivation: HPCS interconnects
- Photonic Integrated Networks
- SPINet – Scalable Photonic Integrated Network
  - Concepts and architecture
  - Performance simulations
  - Prototyping experiments
- Conclusions

HPCS Interconnect Requirements

- High Bandwidth
- Low Latency (< 1 μs)
- Port count scalability (> 1024)
- Flexible packet sizes
  - bulk data transfers
  - small messages (cache coherency, message passing)
- No packet reordering
- Power consumption; Spatial volume; Cost

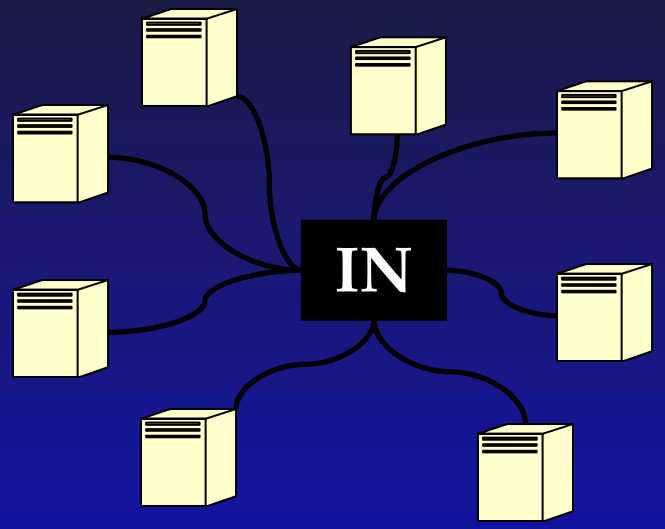

Photonic Integrated Networks

- Large Scale Photonic Integration
  - Break the optical cost barrier
  - Very low power dissipation

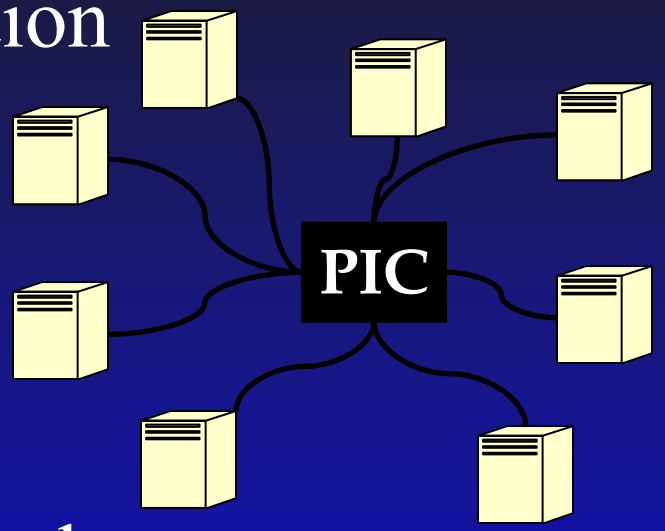

- **But:** Packet size doesn't scale down
  - Typical packet > 10 ns ≈ 2 m (silica fiber)

- **Truly Bufferless Architectures Must Be Used**

SPINet

- Scalable Photonic Integrated Network
- Multistage interconnection network (MIN).
- 2×2 switching nodes
- WDM messages
  - Bandwidth
  - Simplicity
- Instantaneous lightpaths ev
- Contentions resolved by di
- Physical layer acknowledgements

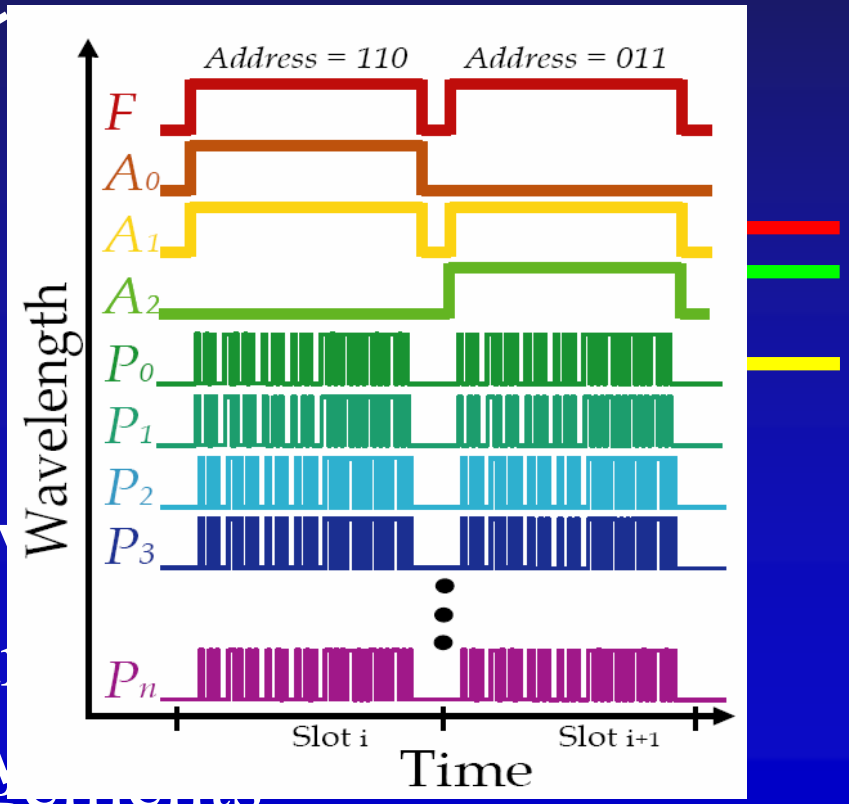

2×2 Switching Node

- Routing according to wavelength encoded address
- No buffering
- Six switching states

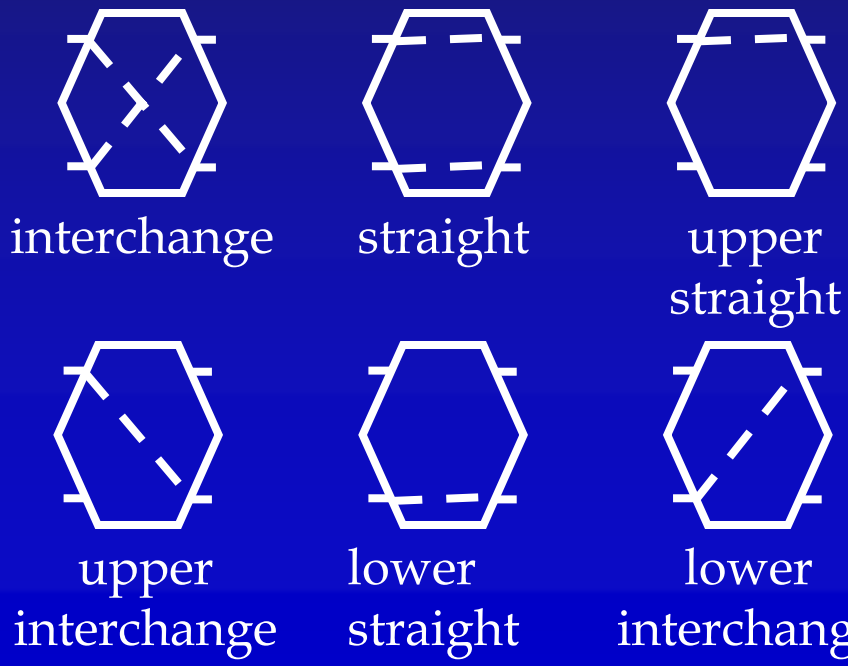

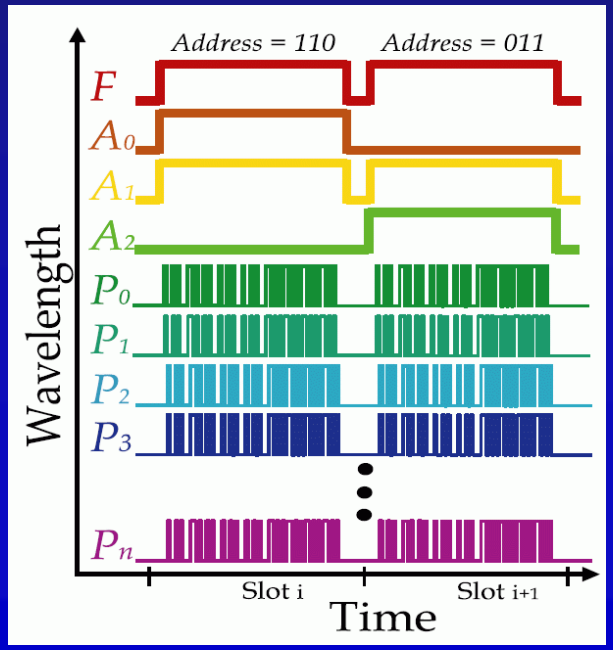

Demonstration

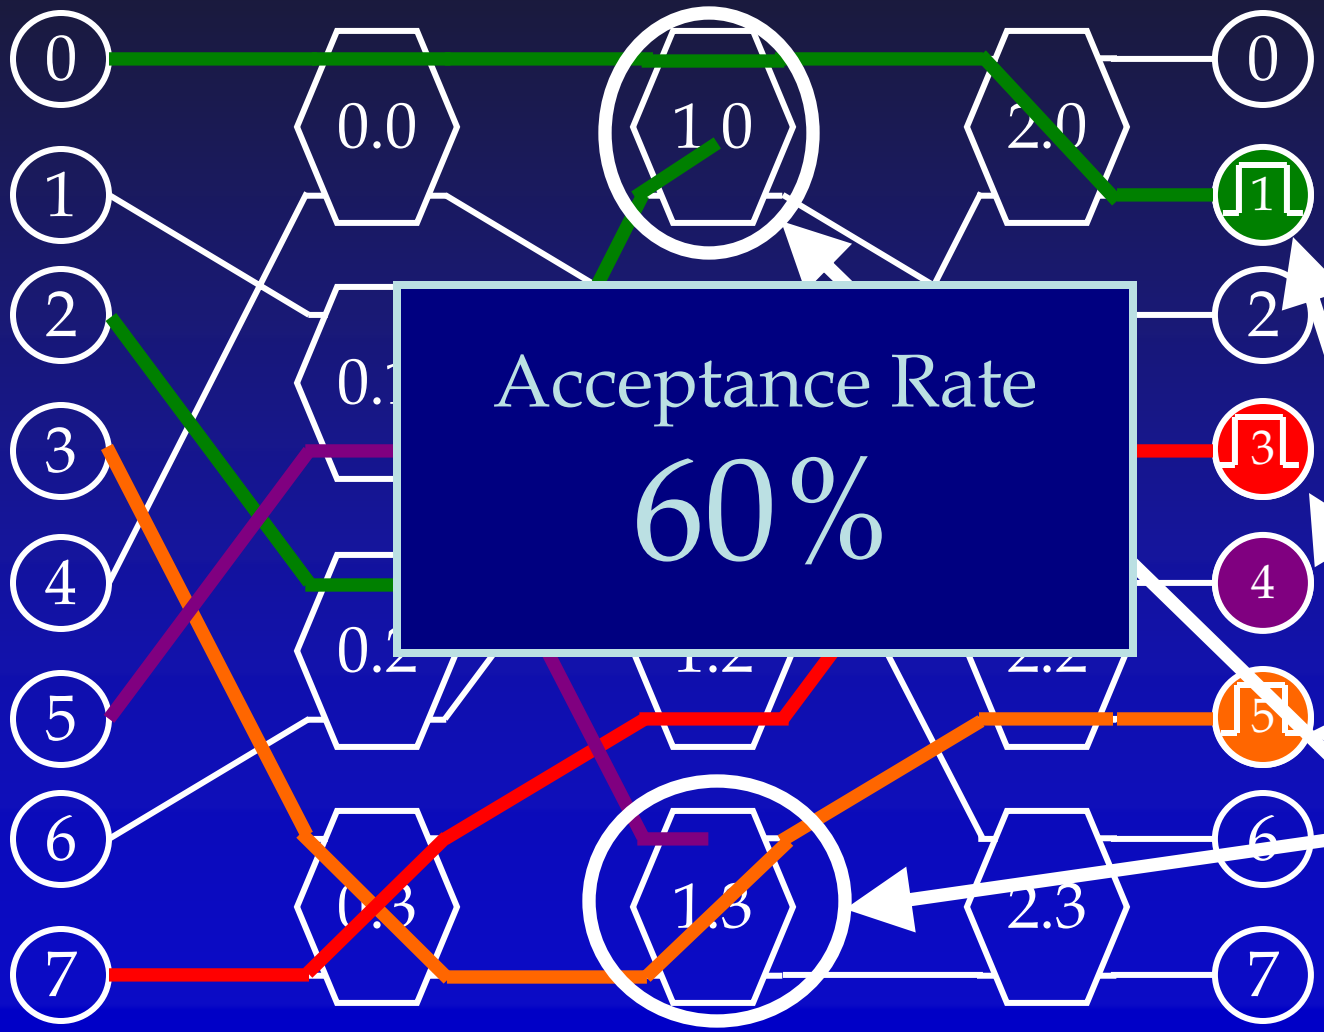

Messages:

Slot is over

0 → 1

2 → 1

Ready for the next slot:

3 → 5

5 → 4

7 → 3

Acks

Contention

Paths torn down

resolved

Acceptance Rate

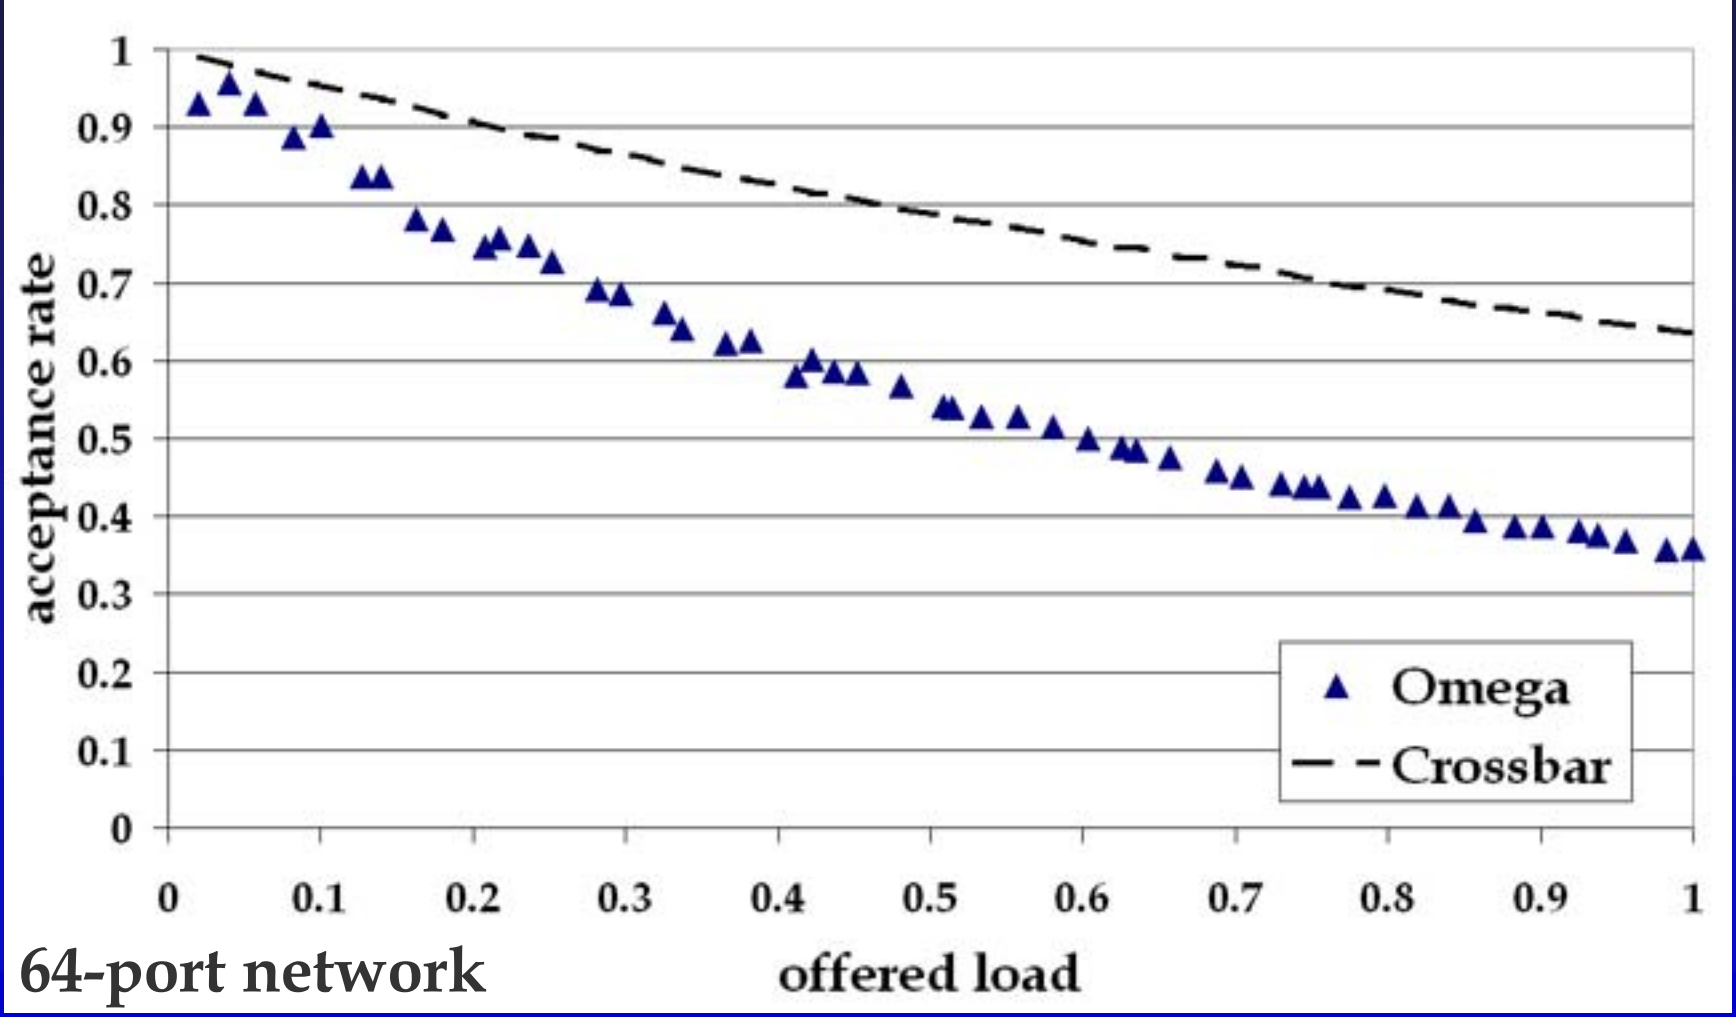

Enhanced Omega Network

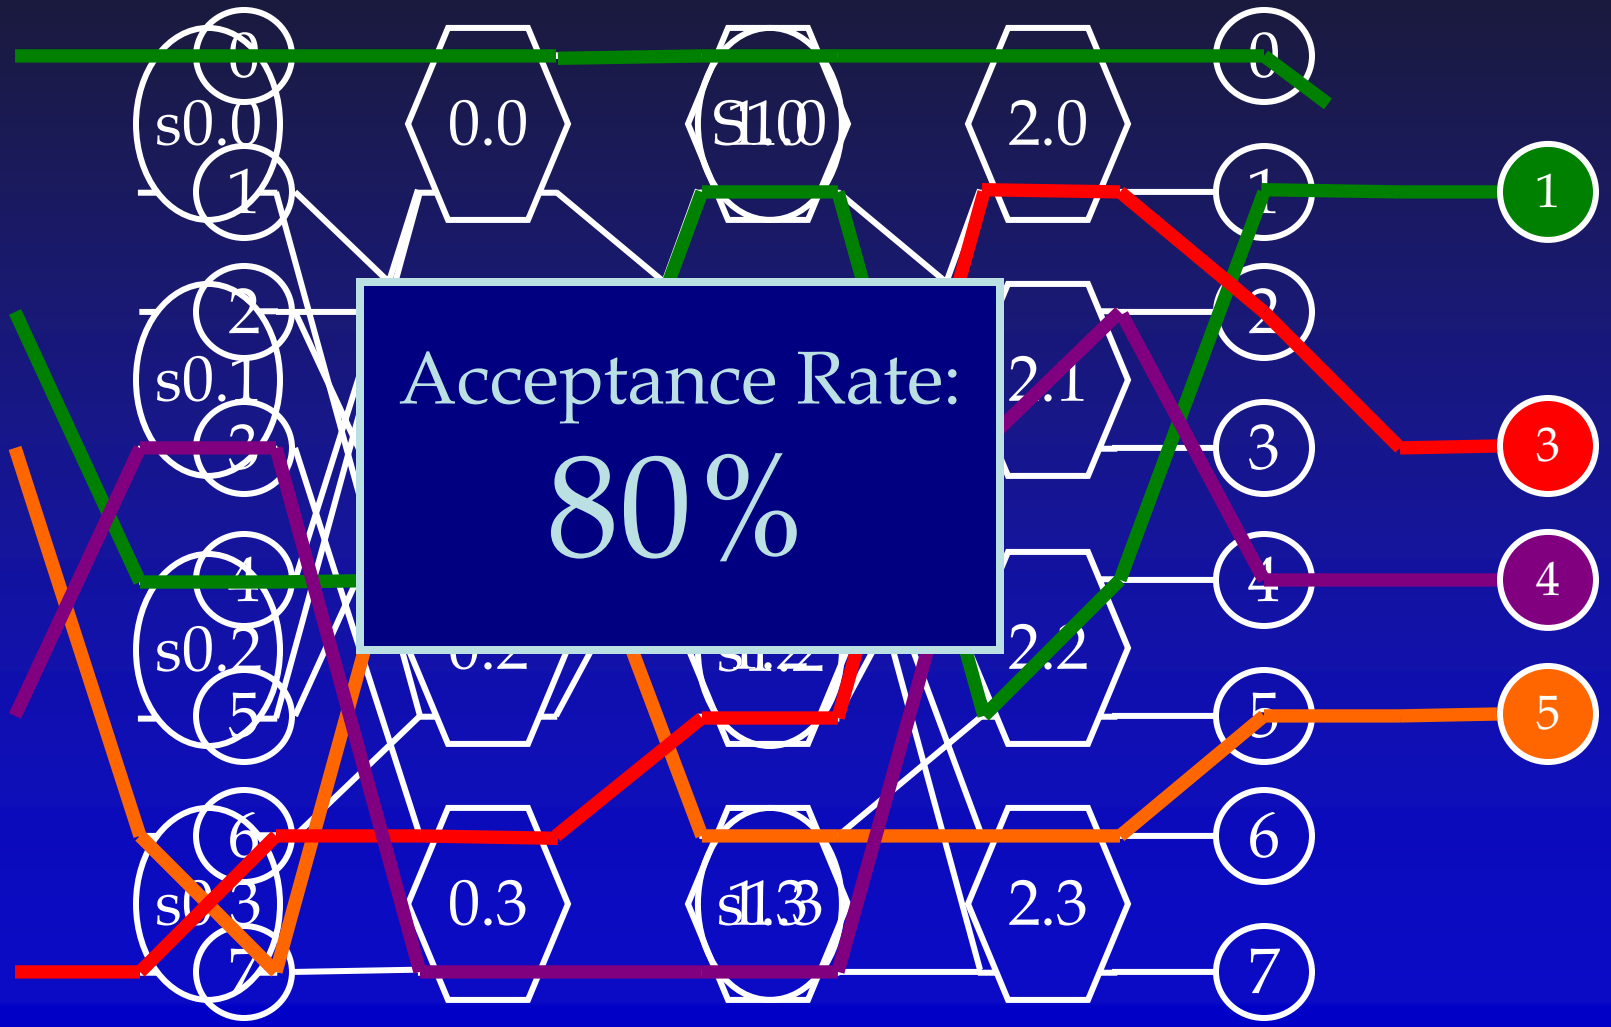

Acceptance Rate Revisited

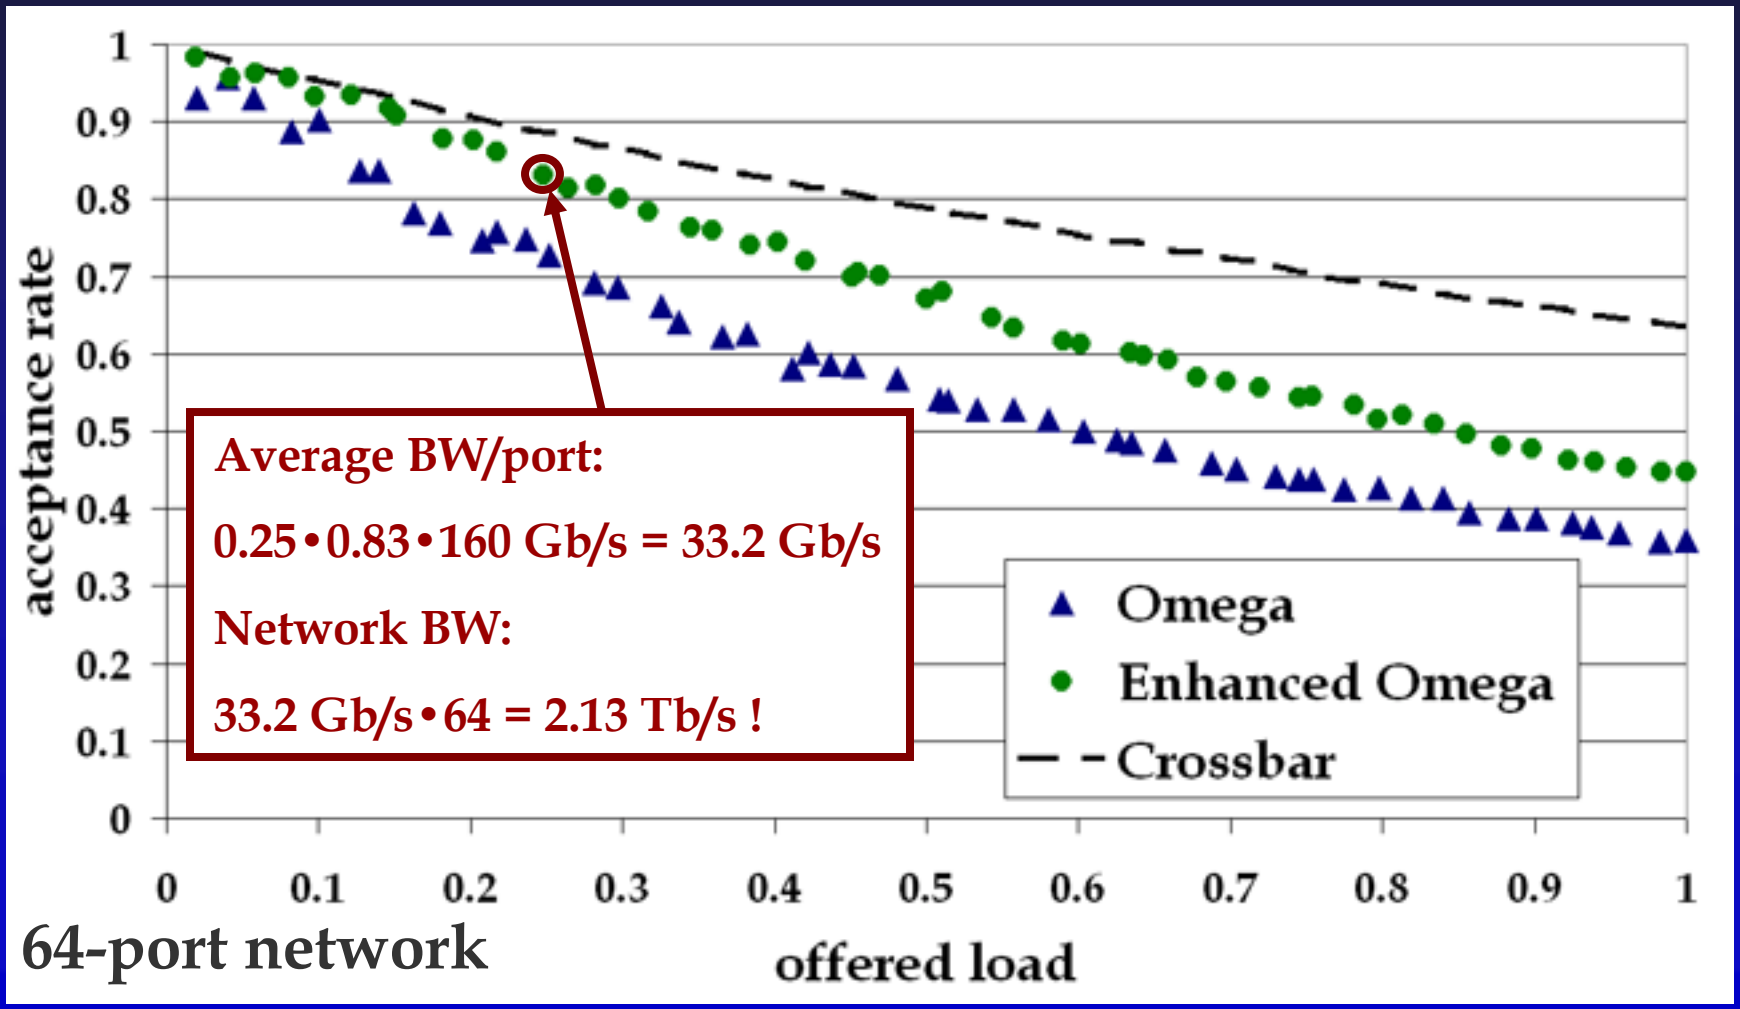

Latency

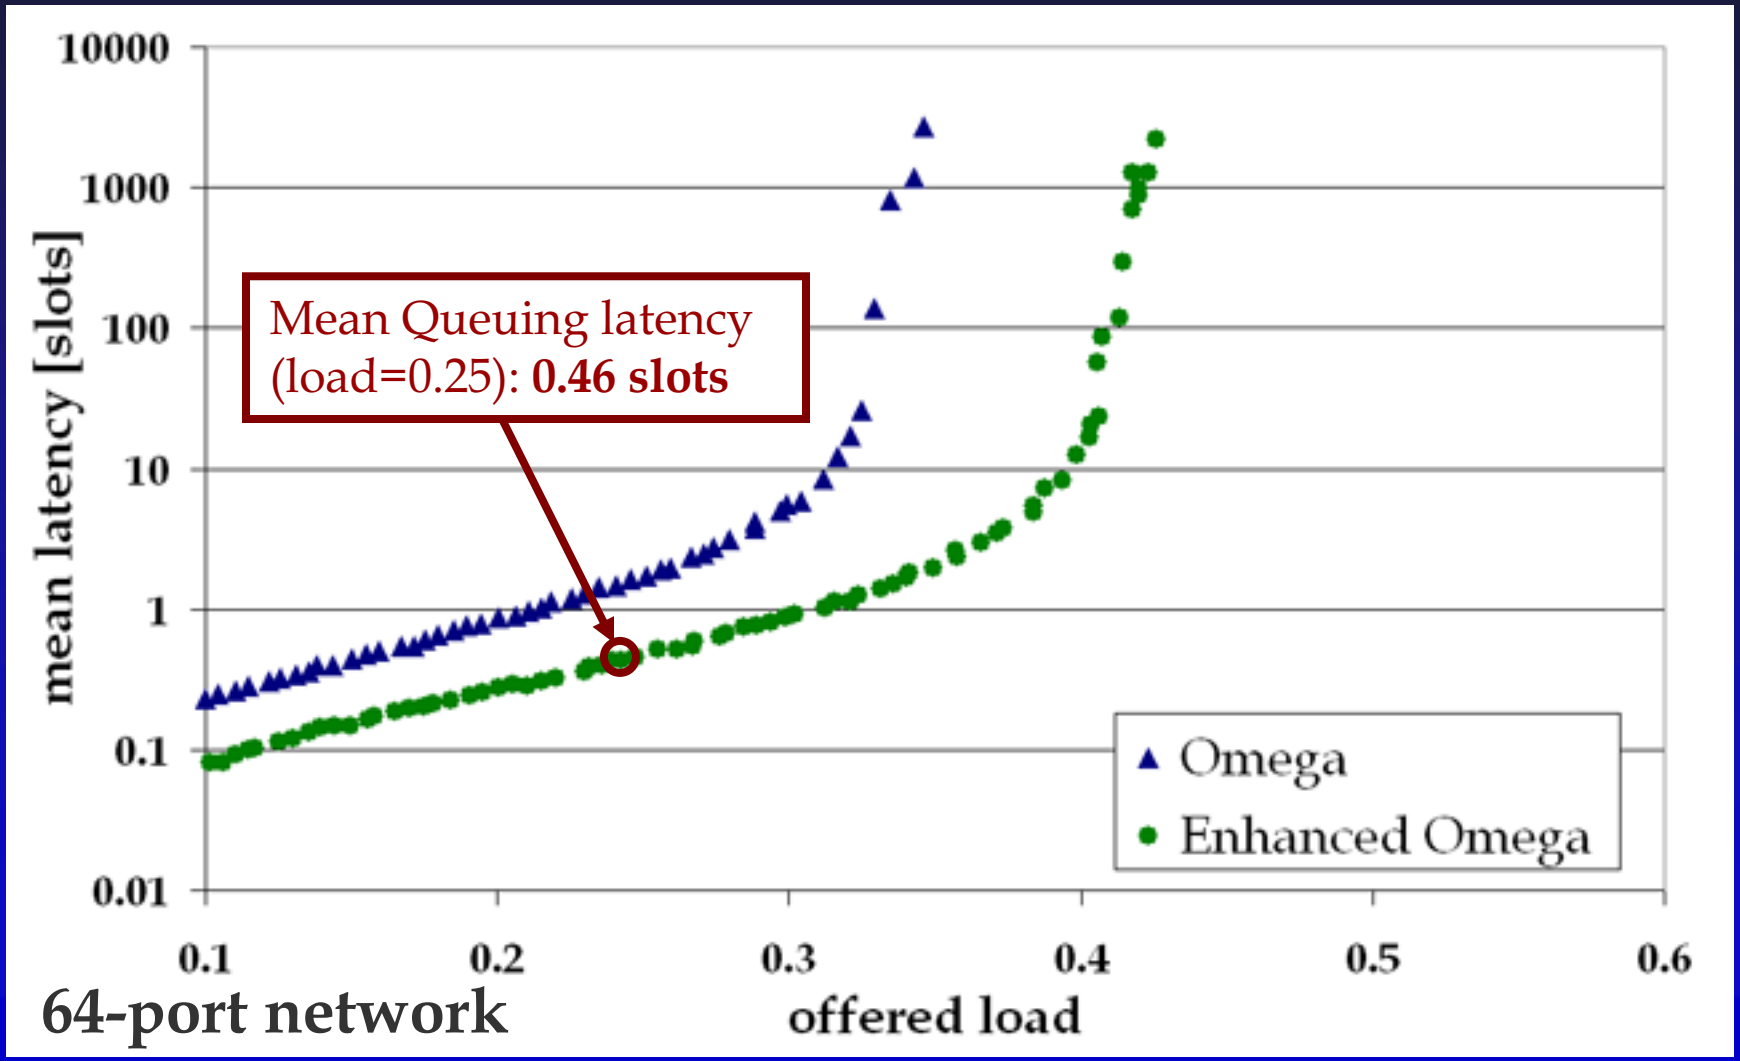
10000
1000
100
10
1
0.1
0.01
mean latency [slots]
Mean Queuing latency (load=0.25): **0.46 slots**
Omega
Enhanced Omega
0.1
0.2
0.3
0.4
0.5
0.6
offered load
64-port network

Scaling the Network

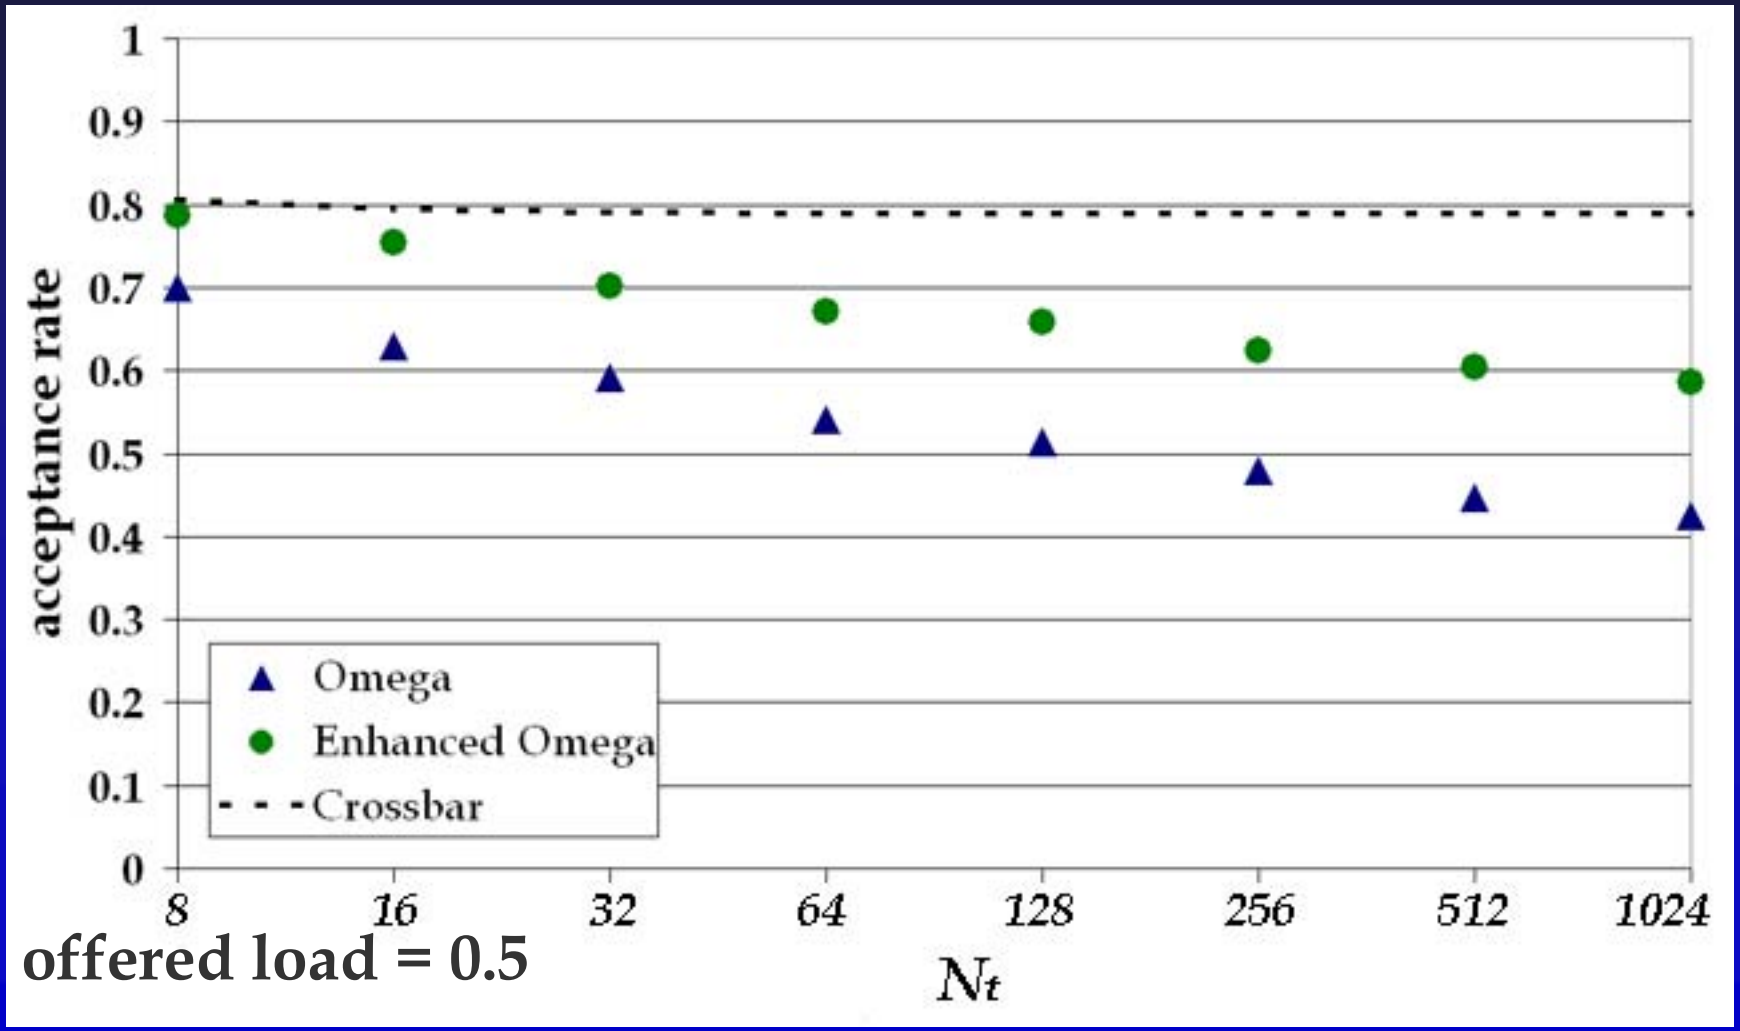

Part II: Implementation and Experimental Results

Implementation: 2x2 Switching Node

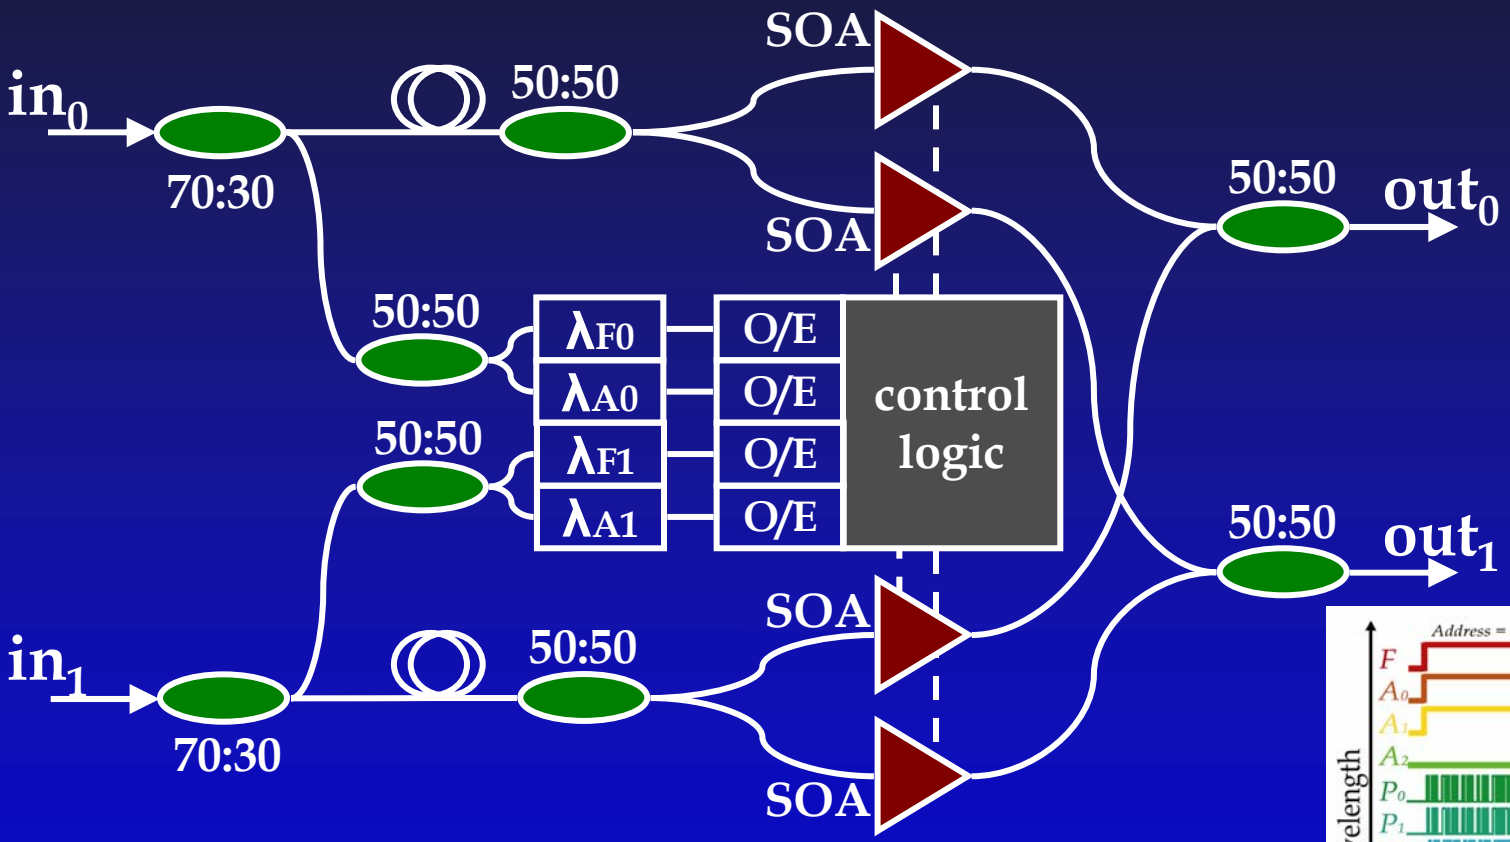

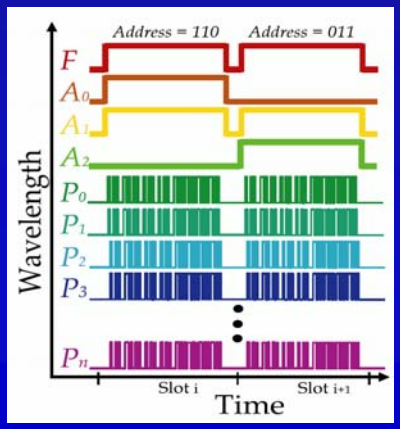

Prototype Switching Node

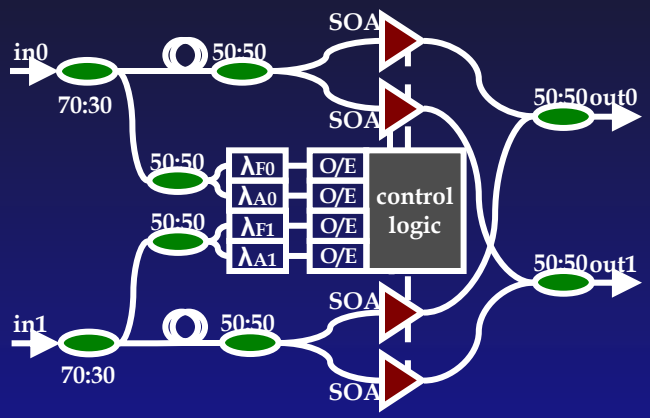

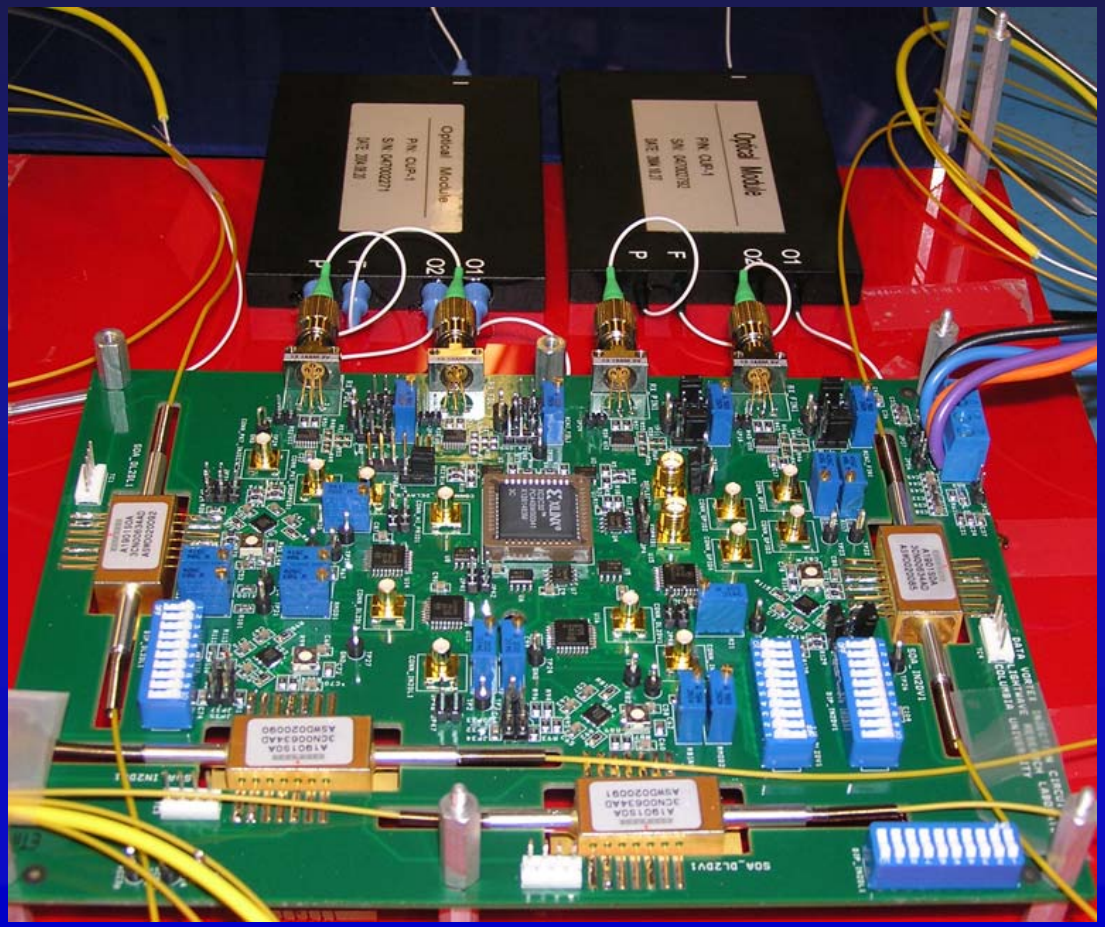

Experiment I: Address Resolution

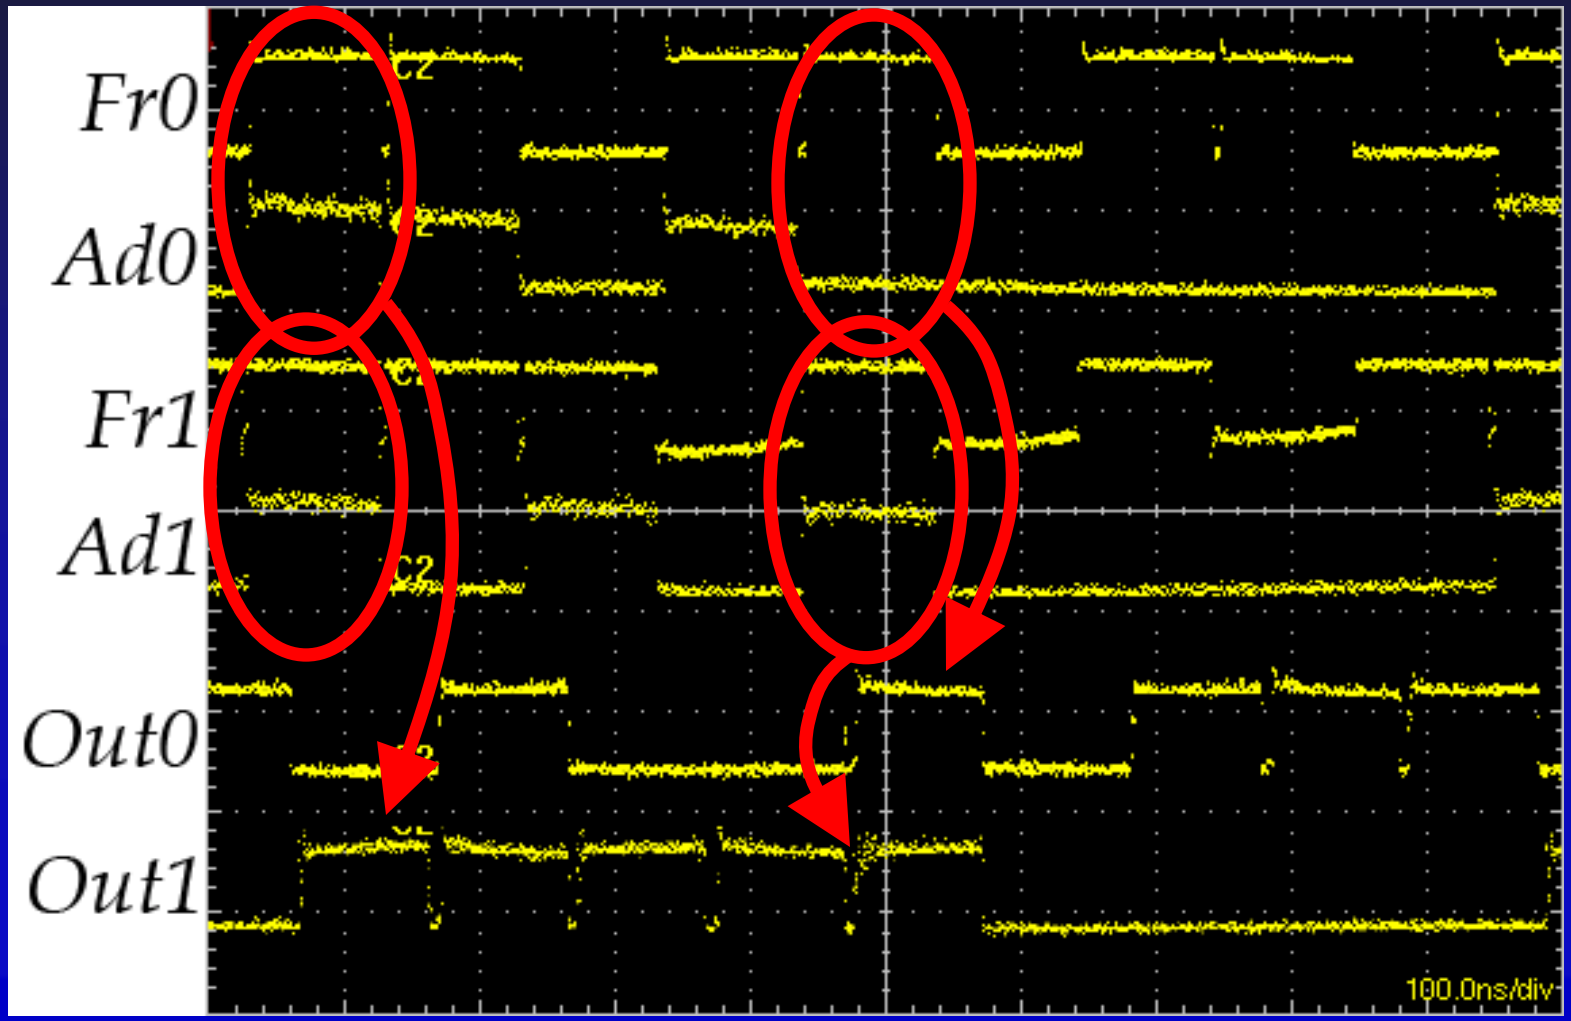

Experiment II: Ack Transmission

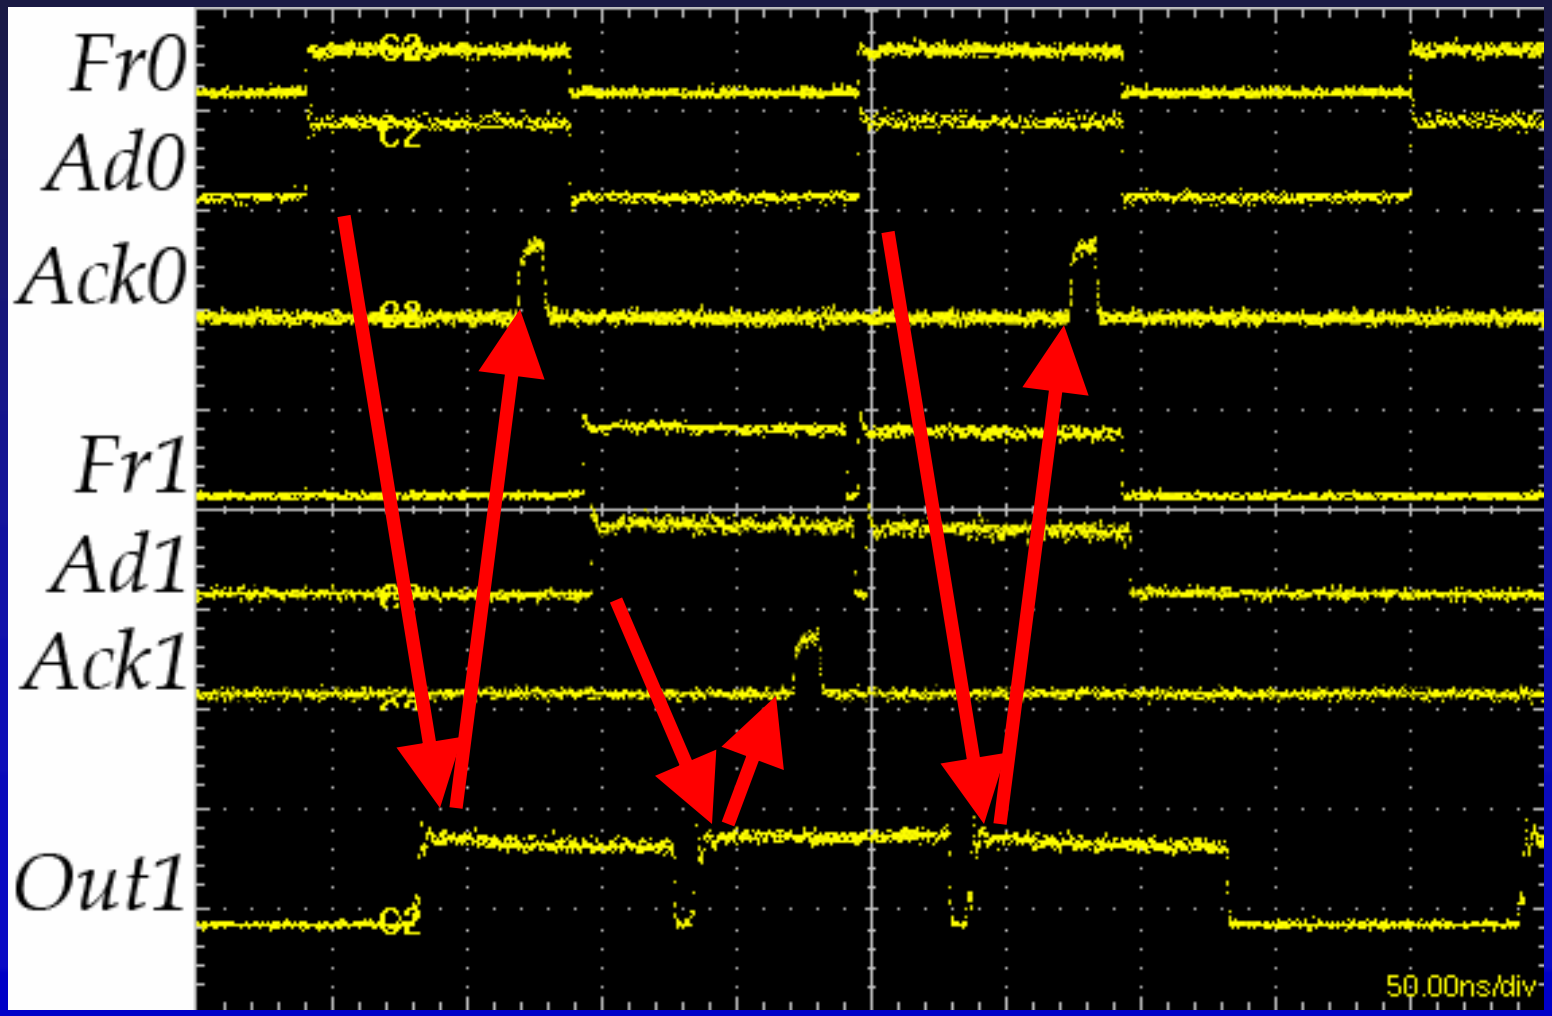

Prototype Experiments (cont.)

- 160 Gb/s peak BW (16 wavelength x 10 Gb/s)
- Bit Error Rate < 10^{-12}

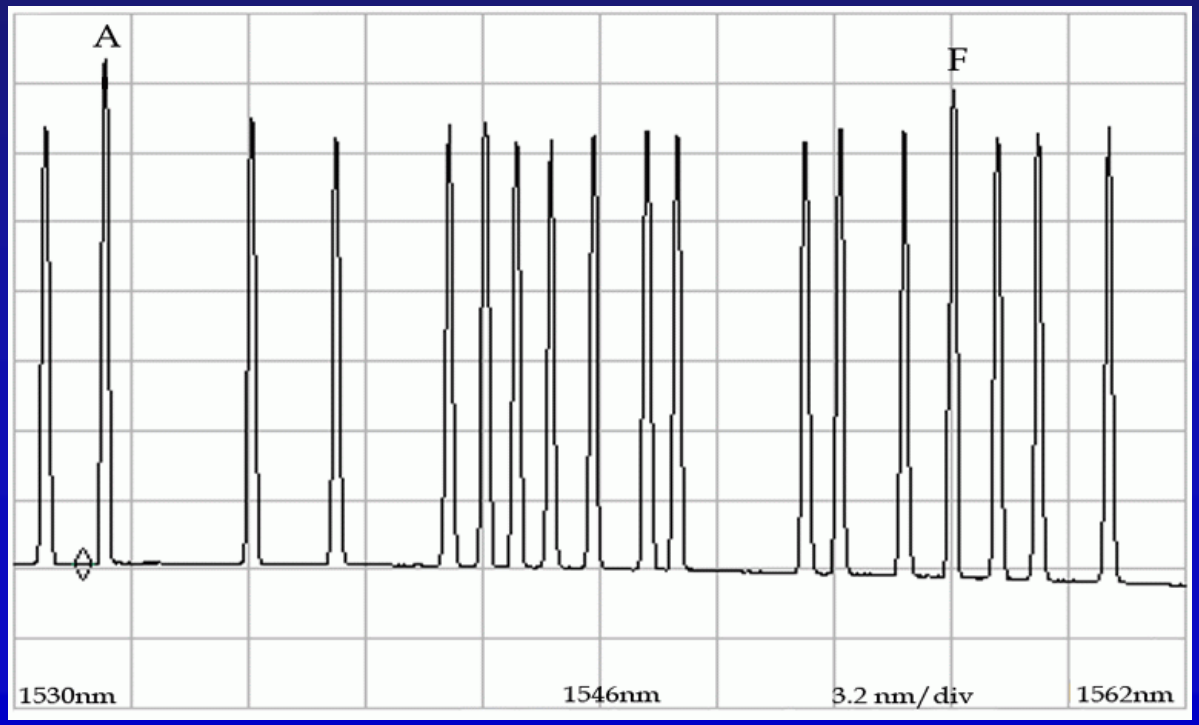

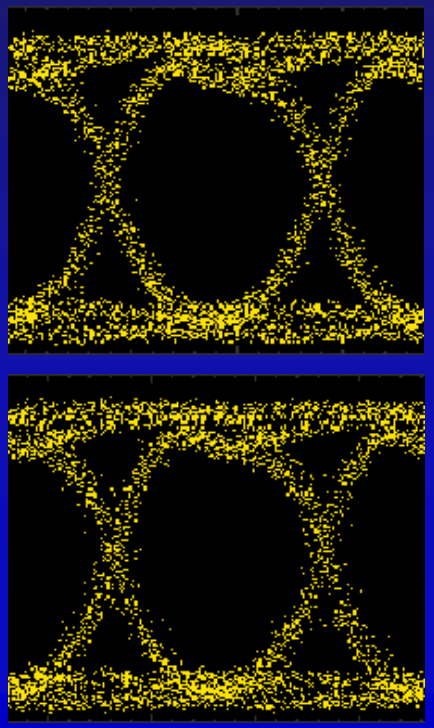

Conclusions

- SPINet: A new Photonic Interconnect Architecture
  - Ultra low latency and High bandwidth
  - Scalable and Integratable
  - Address decoding and routing
  - Contention resolution
- Future research
  - Increase utilization and decrease $P_{[blocking]}$
  - Implementation technologies

“Will Software Save Moore's Law?”