The IRIS Optical Packet Router
A DARPA / MTO Project

International Workshop
on the
Future of Optical Networking (FON)

OFC 2006

Jürgen Gripp

Bell-Labs, Crawford Hill Laboratories

jurgen@lucent.com

Acknowledgements

Lucent Bell Labs

- Jurgen Gripp, Jesse Simsarian, Pietro Bernasconi, Dimitrios Stiliadis, Liming Zhang, Larry Buhl, Jane Le Grange, Nick Sauer, Weiguo Yang, Jeff Sinsky, Martin Zirngibl

Our IRIS partners

- Lehigh University : Prof Tom Koch

Center for Optical Technologies

LEHIGH

- University of Santa California Cruz : Prof Anujan Varma

SANTA CRUZ

- Agility Communications : Dr Mike Larson

AGILITY

Overview

- Why electronic routers don't scale and how the IRIS optical packet routers fixes the problems
- IRIS architecture and technology
- Experimental Results

Overview

- **Why electronic routers don't scale and how the IRIS optical packet routers fixes the problems**
- IRIS architecture and technology
- Experimental Results

Scaling Today's Electronic Packet Routers

- Routers Growth
  - Traditionally scaled at 2× per 18 months
  - 100 Tb/s core routers in 5-6yrs!!
- Today's Electronic Routers
  - Power and heat dissipation limit density
  - Rack Capacity is Power Density Limited
    - Single rack capacity scales only at 10-20% per year in future
  - Require multiple shelves and racks of line cards
    - Nonlinear Scaling
    - Increases number of electrical to optical to electrical (E-O-E) conversions

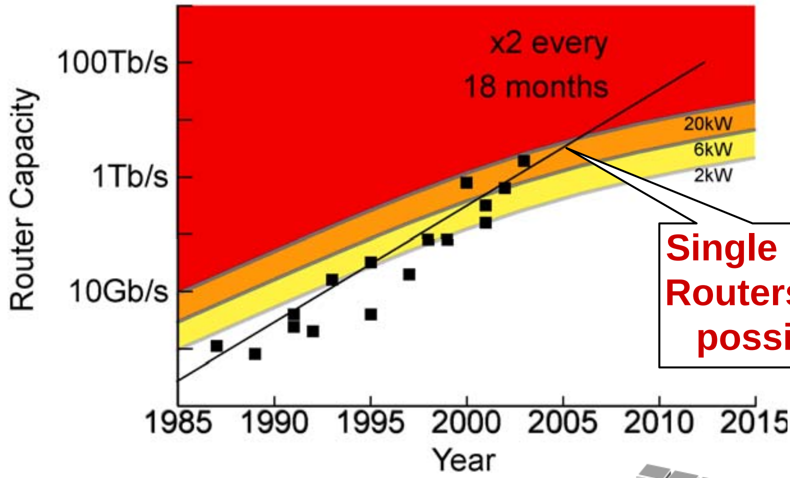

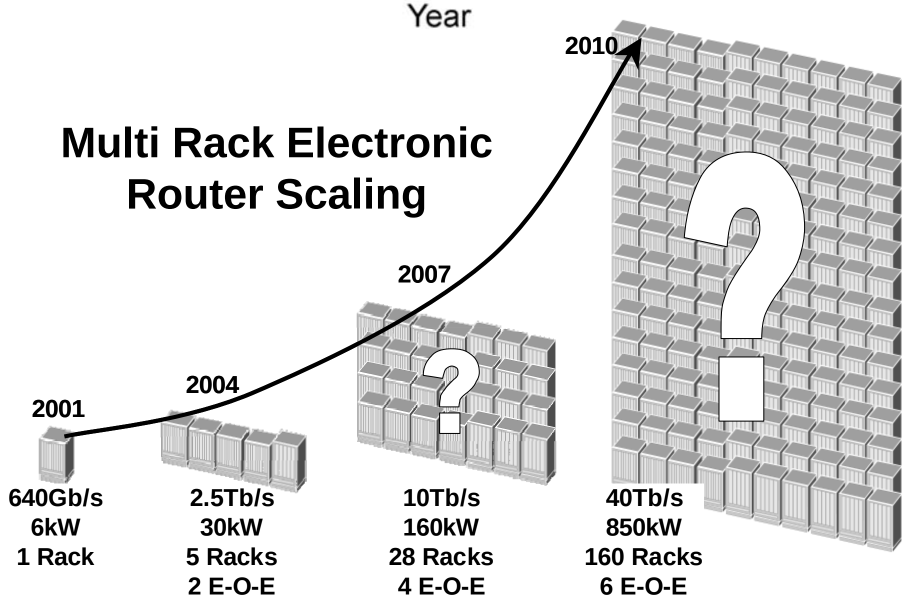

Why do current packet routers not scale?

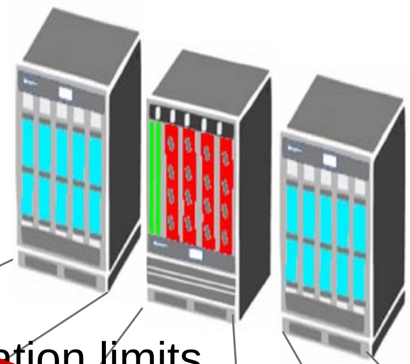

- Power and heat dissipation limit density.
  - Require multiple shelves and racks of Line Cards
  - Increases distance between Line Cards and Switch Fabric
- Centralized Switch Fabric
  - Switch Fabric alone pushes power densit… …ation limits
  - Massive wiring density between L…
- Centralized Control and Schedu… …y

Centralized functions scale poorly

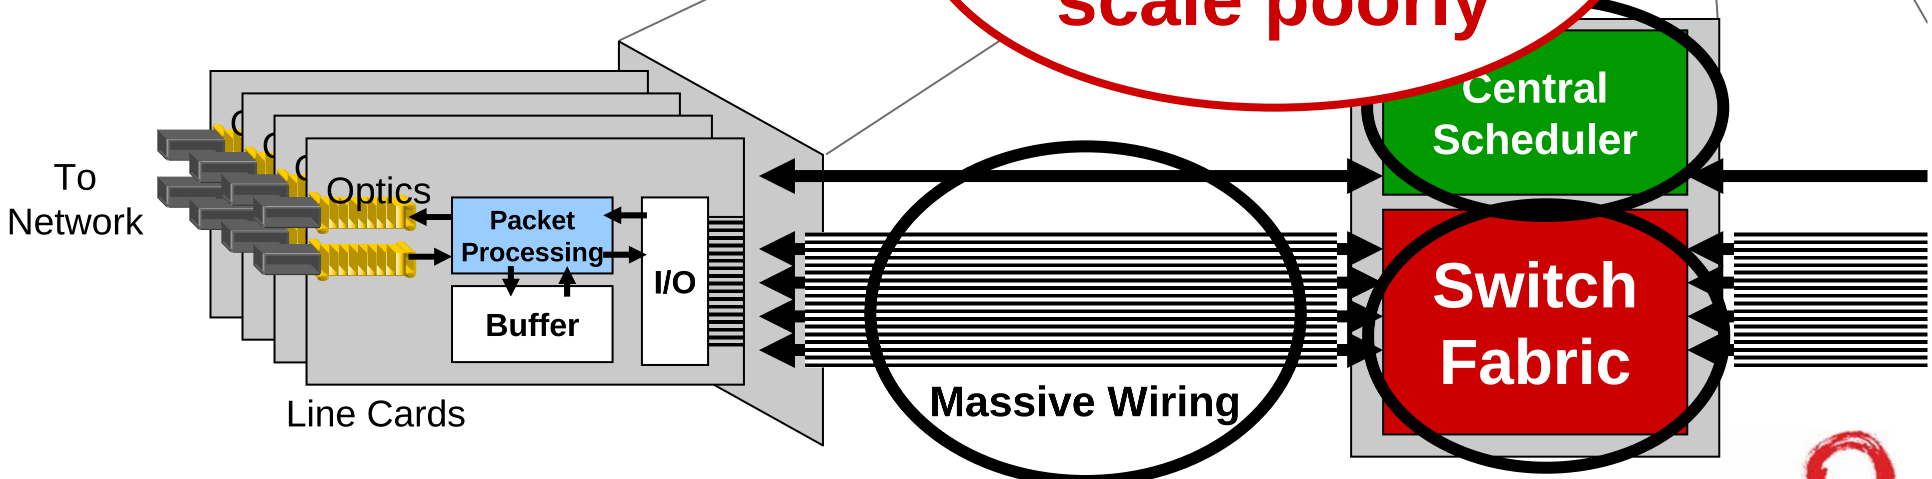

Distributed Optical Switch Fabric

- Optical fibers can eliminate massive wiring density and allow higher bandwidths
- Central Switch fabric replaced by passive optical device that directs signals according to their wavelength
  - Arrayed Waveguide Grating (AWG)
- Switching using Fast Wavelength T... -TX) distributed on each line card
  - Deployed hardware scales with
  - Passive Optical Backplane

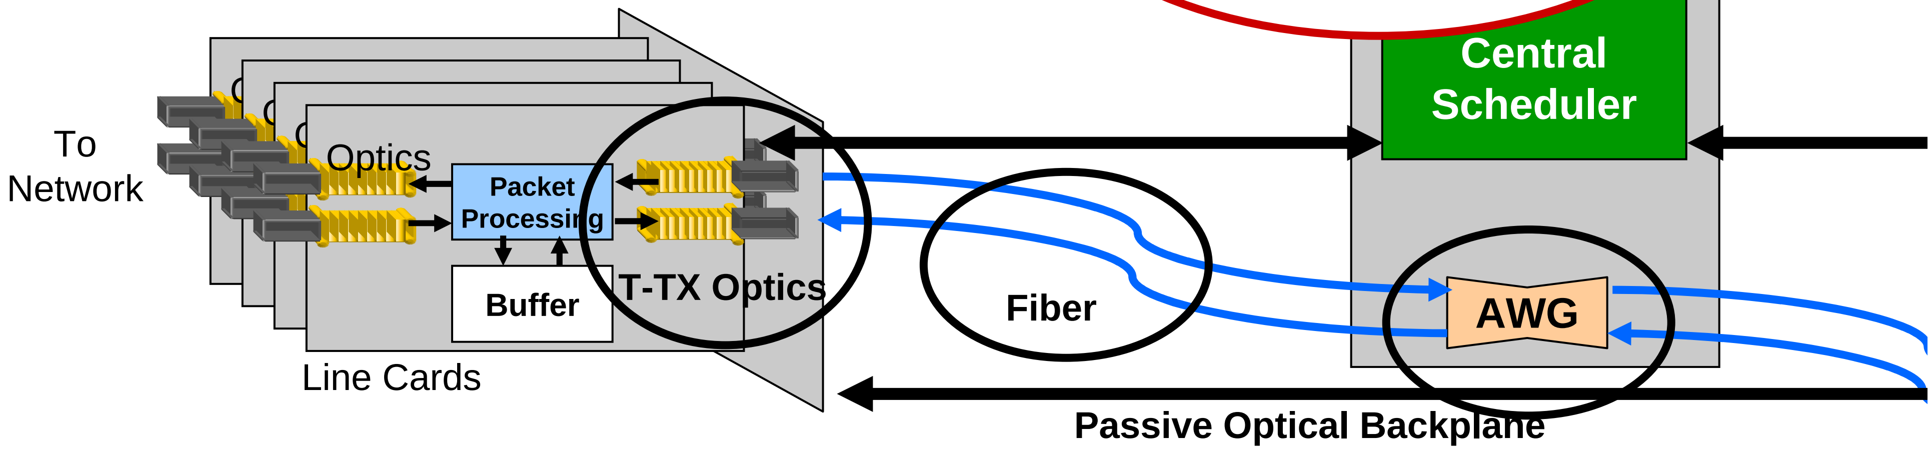

Distributed Scheduler

- Eliminate central scheduler without sacrificing throughput
  - Distribute traffic uniformly to Line Cards
  - Technique called “Load balancing”
- Each line card works like a small router with its own scheduler
  - Scales gracefully since complexity of scheduling is constant
  - Deployed hardware scales with
- Combined with Distributed O
  scaling to very large routers

Router Scales to >100Tb/s

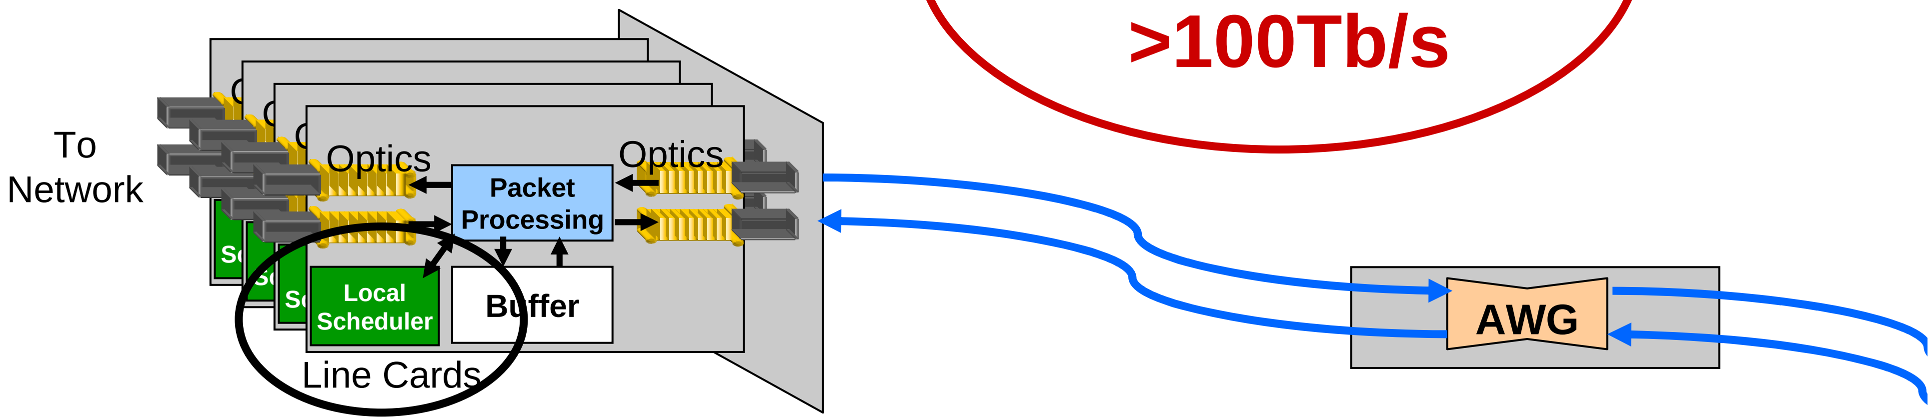

Overview

- Why electronic routers don't scale and
  how the IRIS optical packet routers fixes the problems
- **IRIS architecture and technology**
- Experimental Results

Load-Balanced Router Architecture

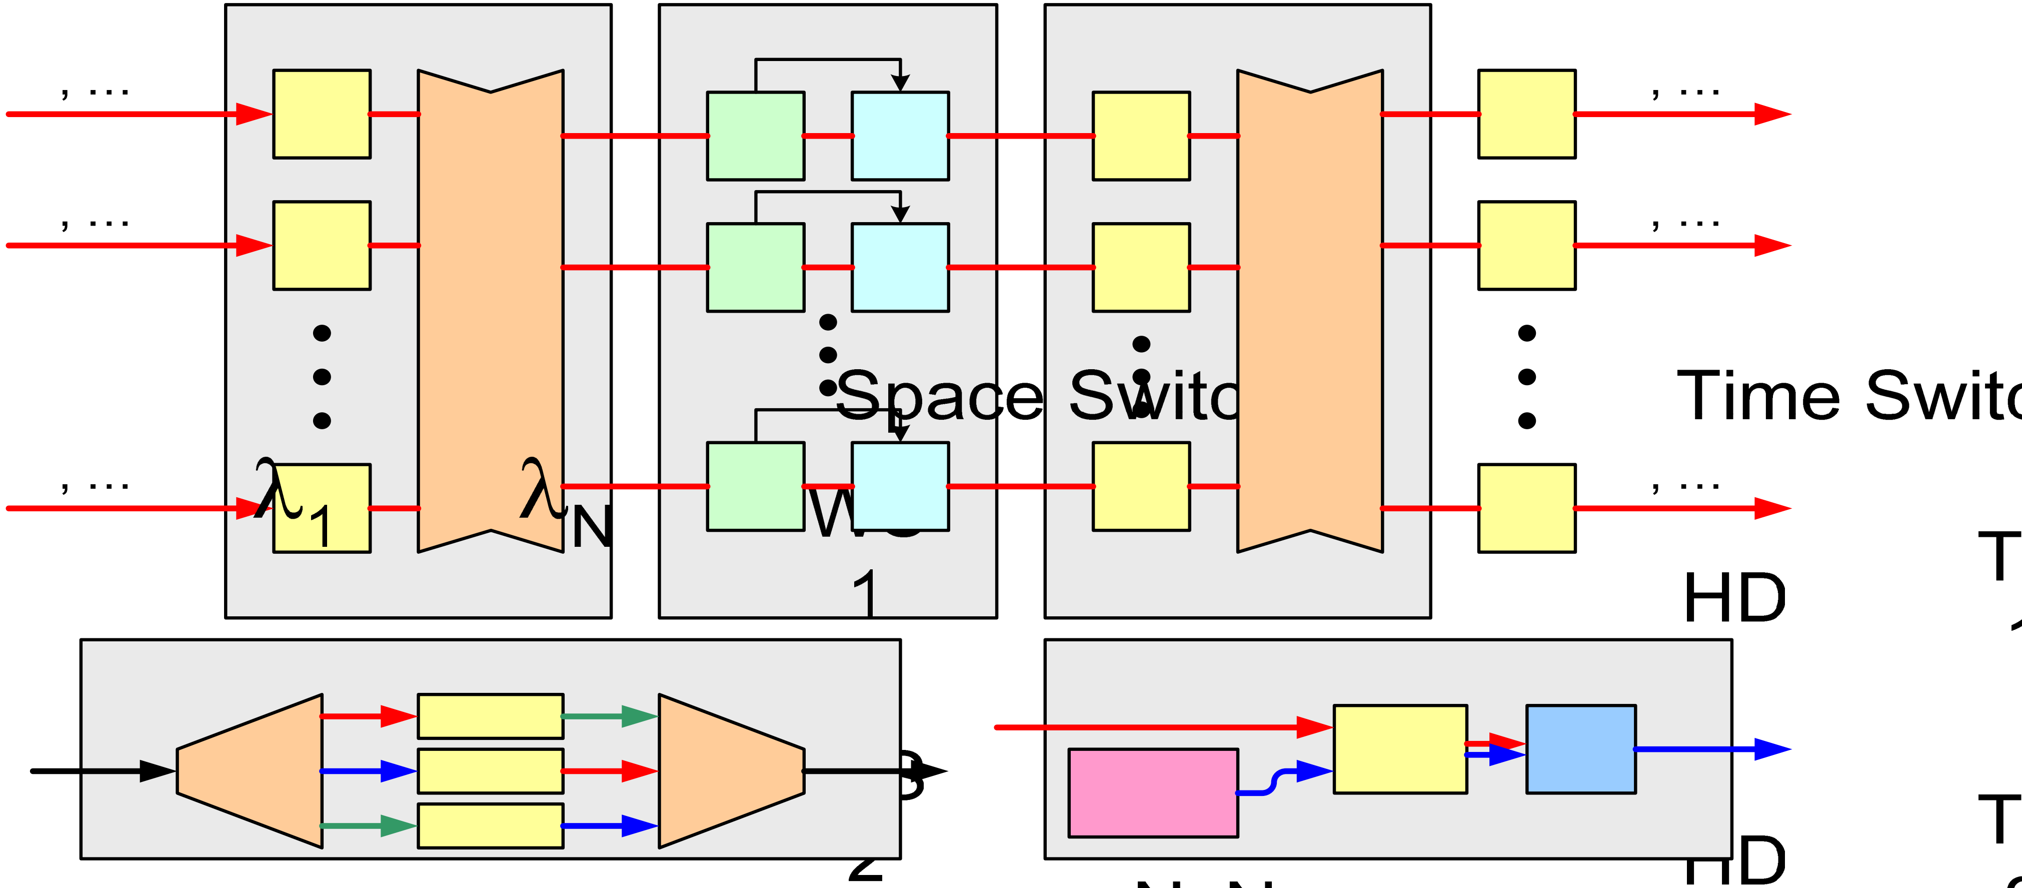

- three stage architecture (space – time – space)
- all-optical data plane, SOA-based wavelength converters
- N^2 router ports (for N AWG ports), scales to 256 Tb/s for 80x80 AWGs and 40 Gb/s data rate
- round-robin schedule for the space switches, deterministic schedule for time buffers

NxN Arrayed Waveguide Grating (AWG)

Provides strictly non-blocking cross-connectivity between N input and N output ports while using a single set of N wavelengths.

No wavelength conflict possible.

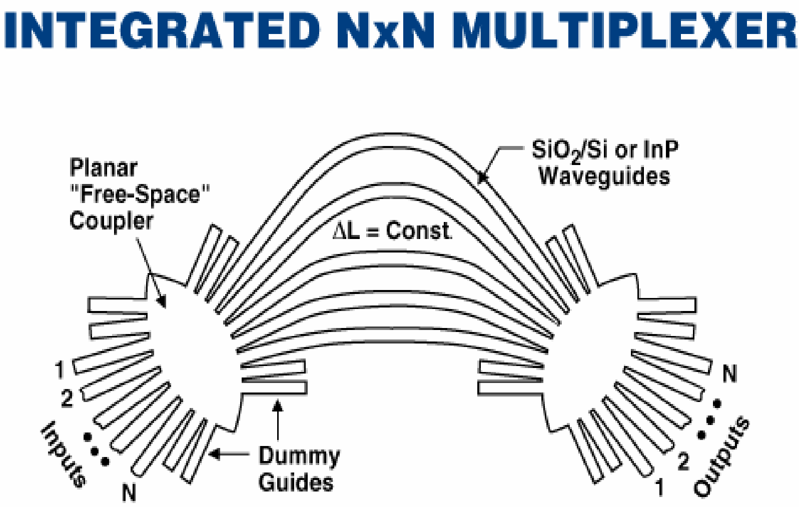

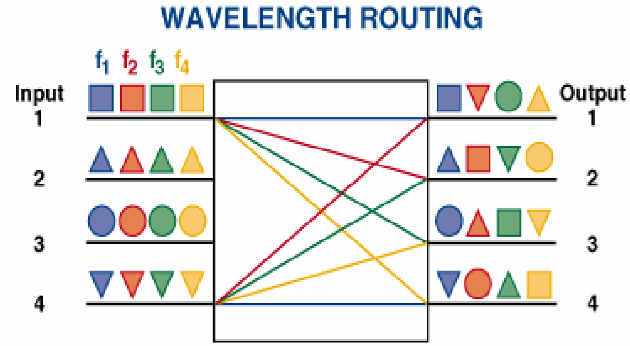

40 channels packaged device
Dimensions: 5 x 2.5 cm^2

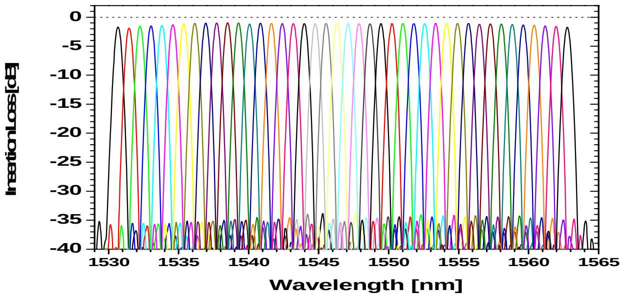

Lucent Technologies
Bell Labs Innovations

N×N AWG Switch Fabric

- Array Waveguide Grating
  - Integrated N×N multi-port grating
  - Fabricated as Planar Light Wave Circuit (PLC)
  - Converts spectral switching to spatial switching with cyclic function
- Switch Fabric
  - Use fast tunable lasers TTX on N input fibers to make N×N cross connect
  - Strictly non blocking cross connect

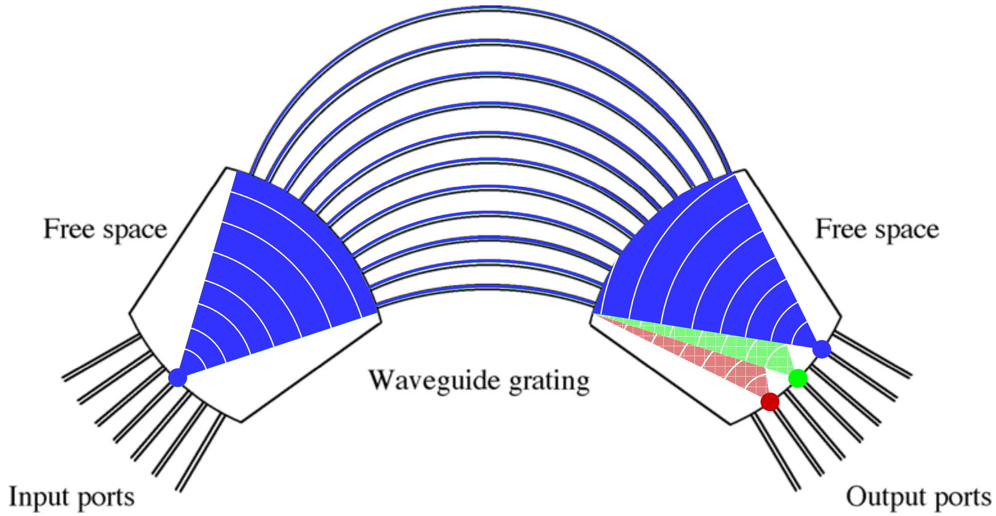

N×N periodic AWG

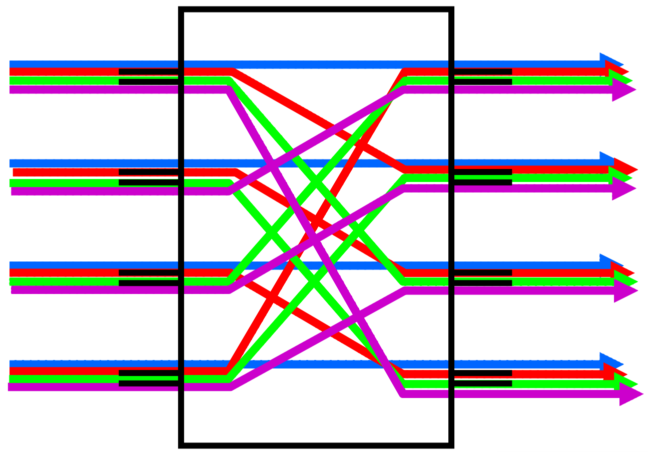

Photonic Integration in IRIS

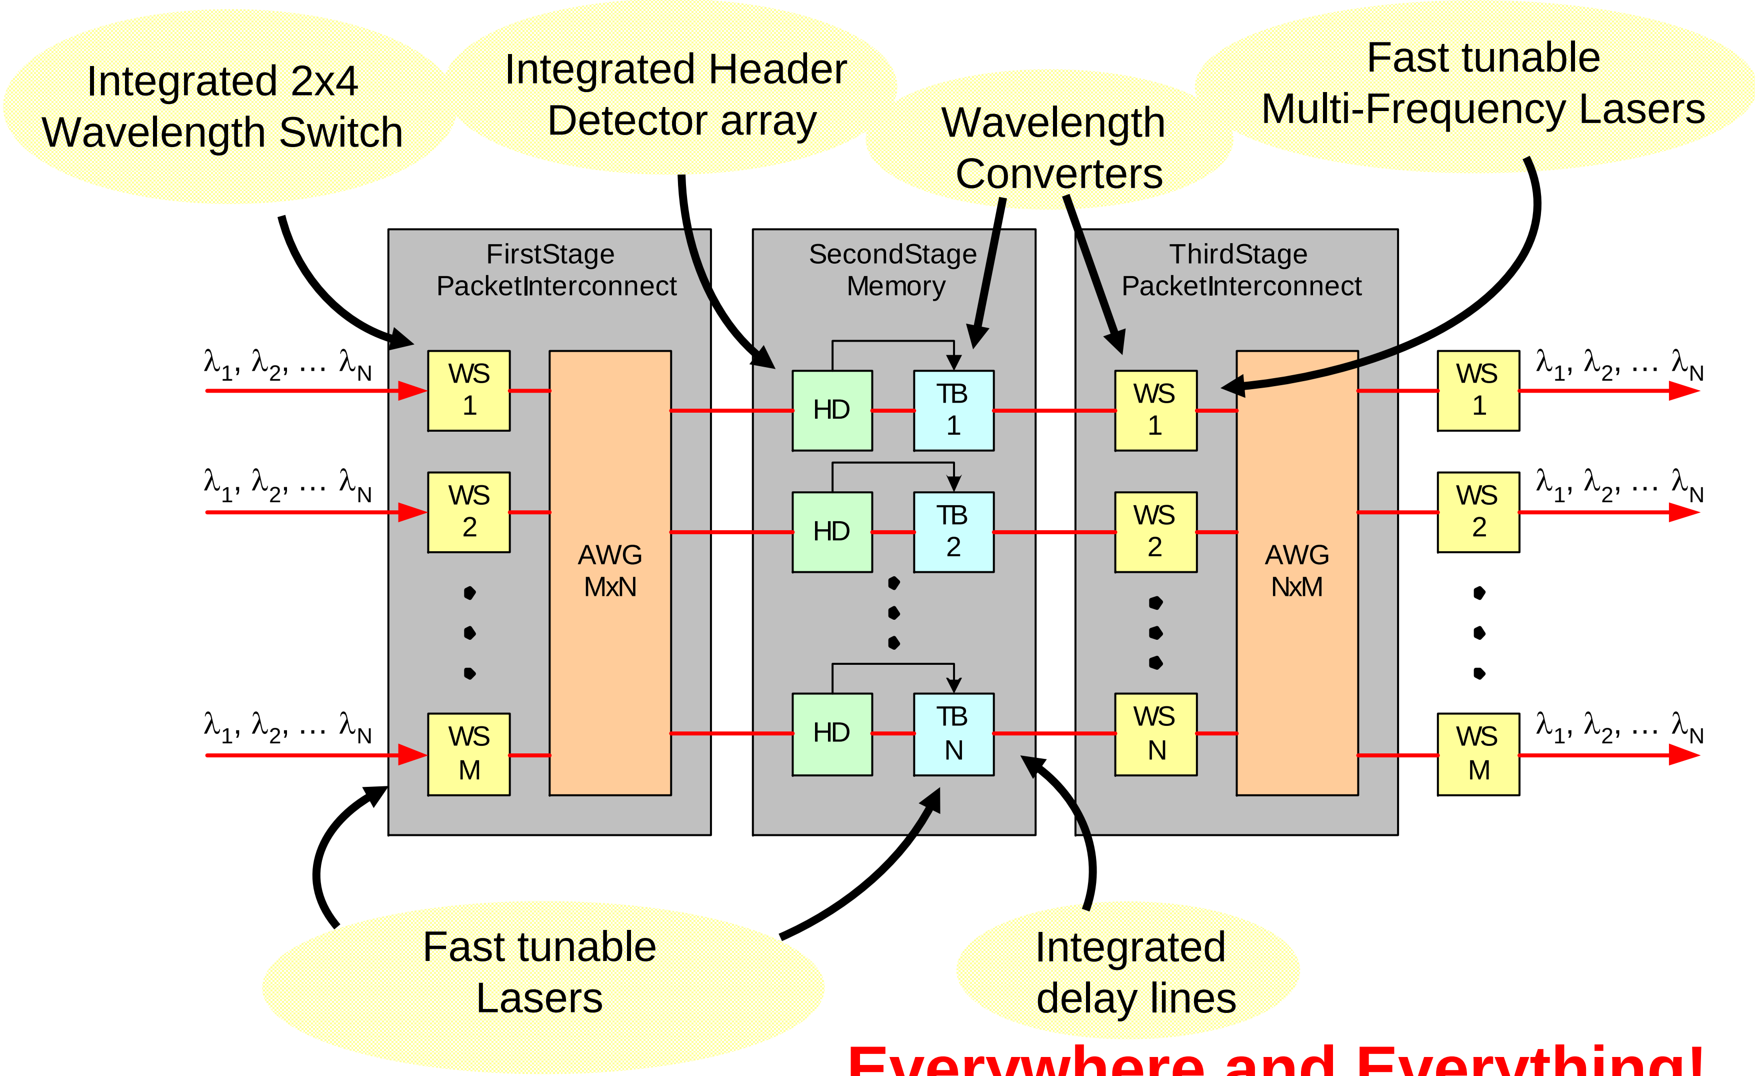

Everywhere and Everything!

Wavelength Switch Module

- The wavelength switch (WS) module consists of an array of tunable wavelength converters (TWC) each containing a tunable laser (MFL) and a wavelength converter (WC).
- The *N* incoming wavelengths are first demultiplexed and singularly converted. The converted signals are then power combined and forwarded.

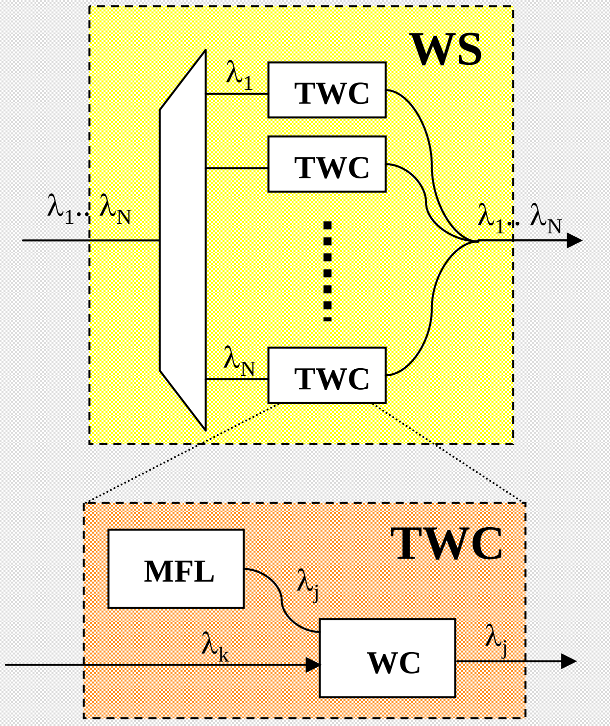

Wavelength switch module

Monolithic Integration Structure

- InP-based material structure quaternary InGaAsP
  - Low propagation losses (<0.5dB/cm) in passive ridge guides
  - 600μm bend radius at index step Δ ~ 0.8%
- Active PLCs consist mainly of SOA as amplifiers, detectors, modulators, etc.
  - MQW or bulk with active overlay or butt-joint
  - Low reflections at passive-active interface

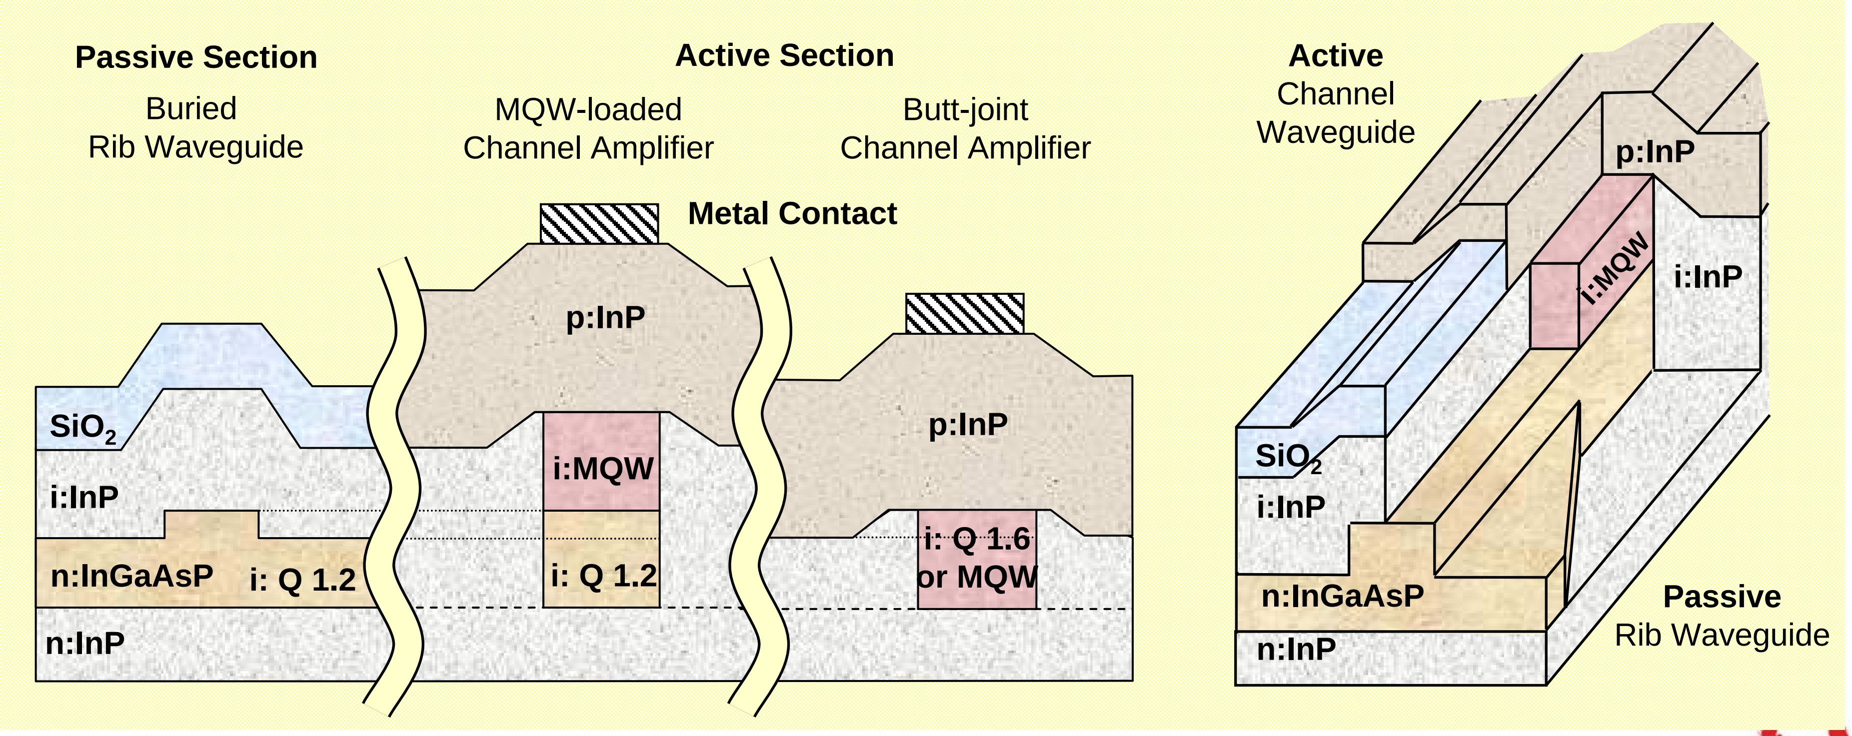

Wavelength Conversion at 40Gb/s

Two different approaches to wavelength conversion

- SOA + FIR filter
  - Single SOA but sensitive to output wavelength if operated in non-inverting mode
- MZI with differential input
  - Two SOAs but MZI works as wavelength filter for input and output not wavelength filtered

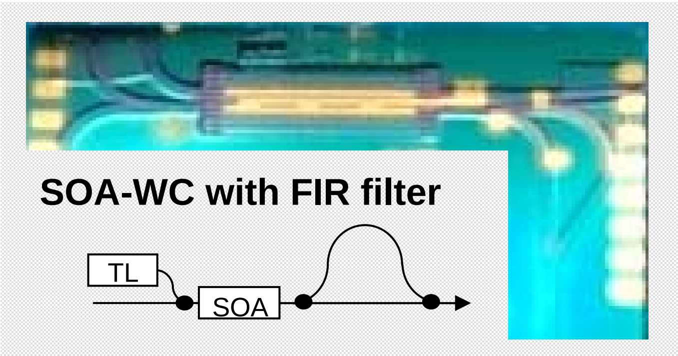

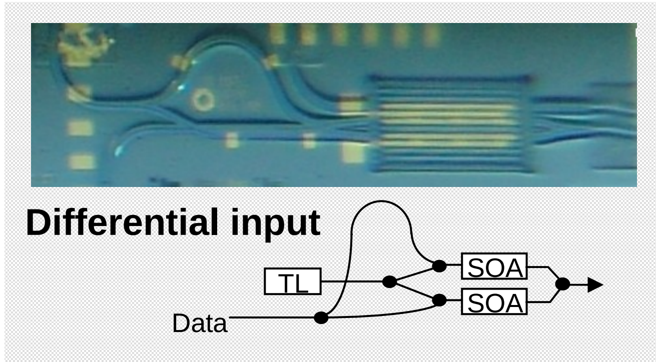

- SOA + FIR filter
  - 3 ps input pulses
  - 3 - 5.5 dB penalty at BER = 10^{-9}
- MZI with differential input
  - 8 ps input pulses
  - 6.5dB penalty at BER = 10^{-9}

Error-free wavelength conversion

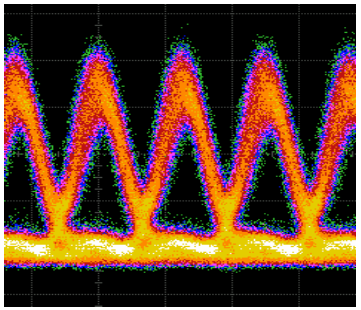

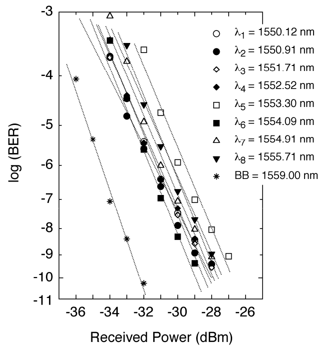

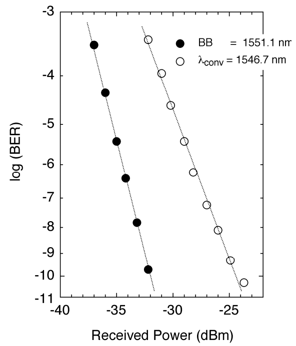

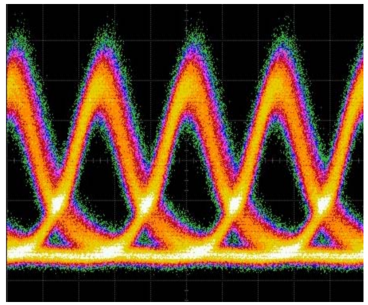

Multi-Frequency Lasers

- 8-channel MFL with double-chirped 1x8 AWG
  - Channel spacing: 100 - 200 GHz
  - Threshold currents: ~ 35-40 mA
  - Operating currents: ~ 2x 90 mA
  - On-chip power: ~ 5-7 dBm
  - SMSR: ~ 40 dB
  - Switching time: **rise ~250 ps**
    fall ~ 290-460 ps

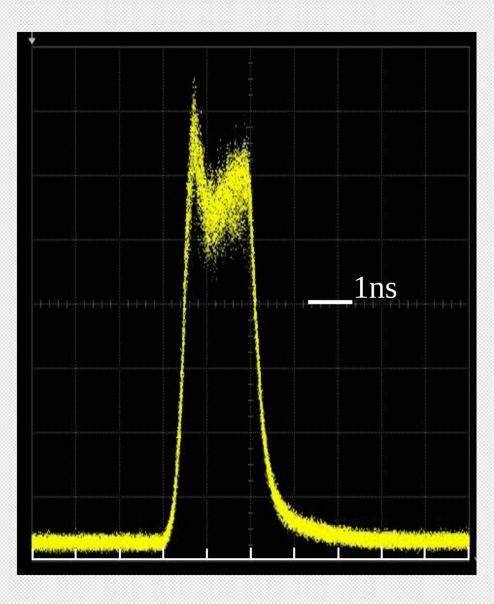

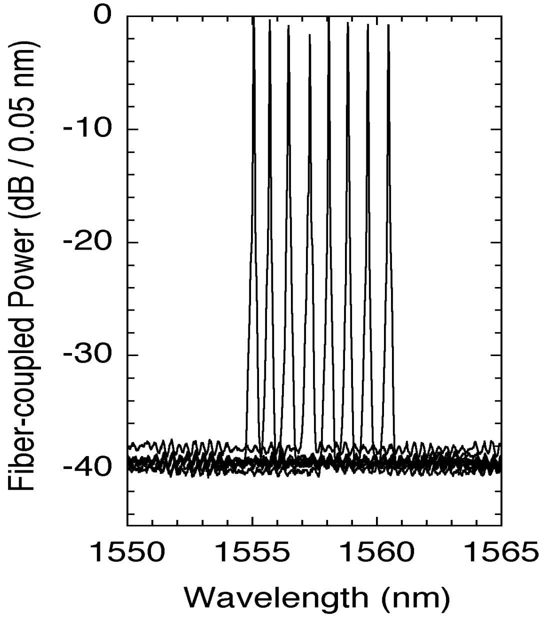

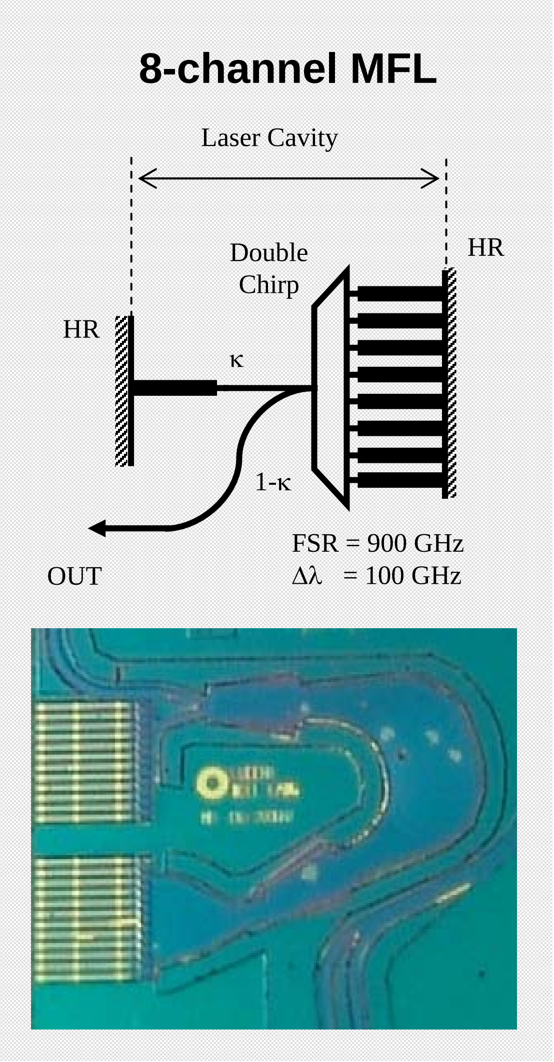

Integrated 1x8 Wavelength Switch

- Monolithically Integrated 1x8-channel wavelength switch comprising
  - 8-channel MFL with double-chirped 1x8 AWG
  - Wavelength Converter SOA + FIR filter
- Error free conversion:
  - from 1559.0 nm to 1550.1 - 1555.7 nm
- Penalty: 3 - 5.5 dB at BER = 10^{-9}

  Penalty partly due to unmatched filter before receiver and degraded OSNR due to high coupling losses
- WC-SOA: 350 mA
- MFL's SOAs: 90-150 mA

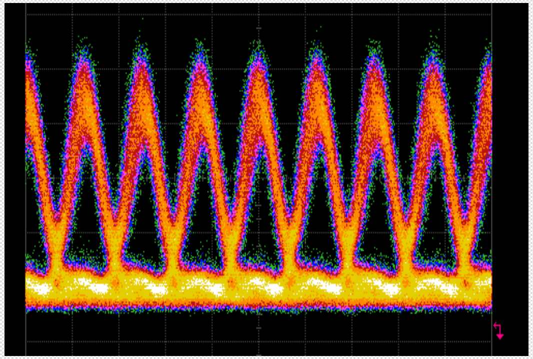

Converted optical eye diagram at 40 Gb/s (20 ps/div)

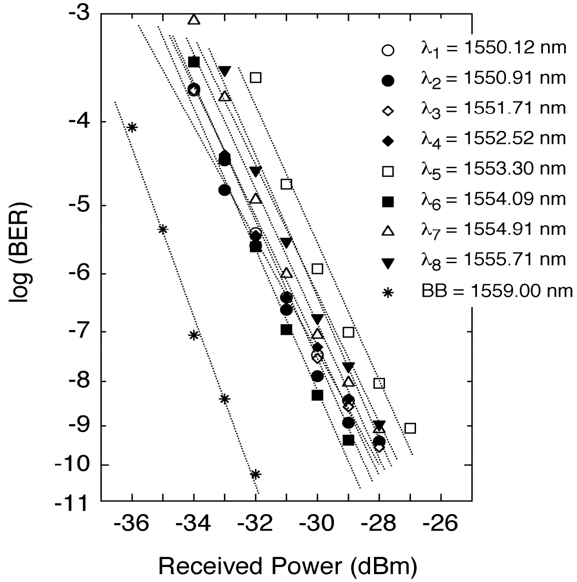

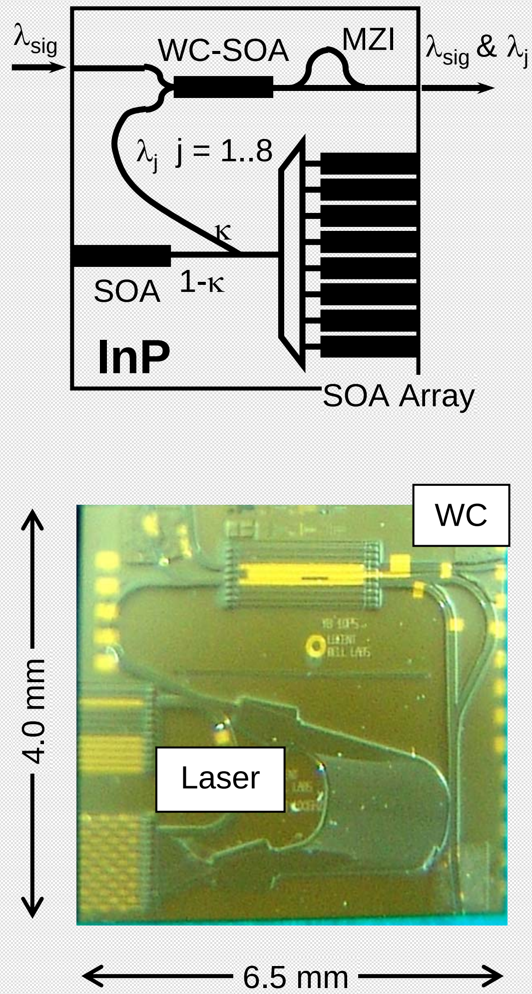

2x8 Wavelength Switch

- Monolithically Integrated 2x8-channel wavelength switch comprising
  - Two 8-channel MFLs
  - Two differential mach zender wavelength converters
  - 2-channel demux interleaver (first time ever in InP)
- MFL: channel spacing 100 GHz
  - operating currents: 80-95 mA
- Interleaver: 3-dB bandwidth ~180 GHz
  - insertion losses ~ 6-8 dB ER > 15 dB (not tuned)

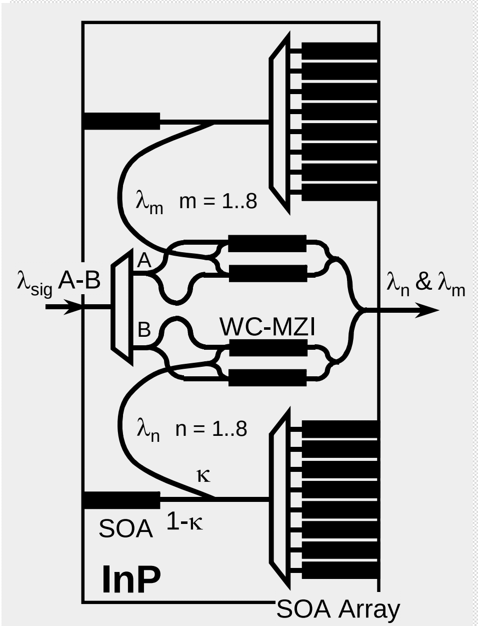

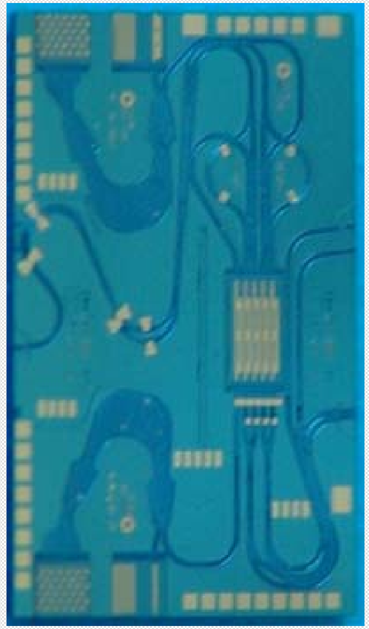

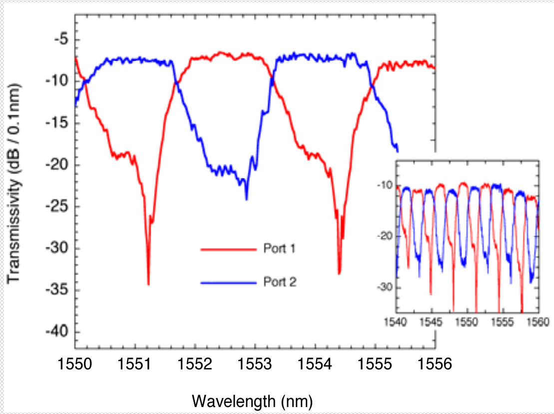

Overview

- Why electronic routers don't scale and
  how the IRIS optical packet routers fixes the problems
- IRIS architecture and technology
- **Experimental Results**

System Setup

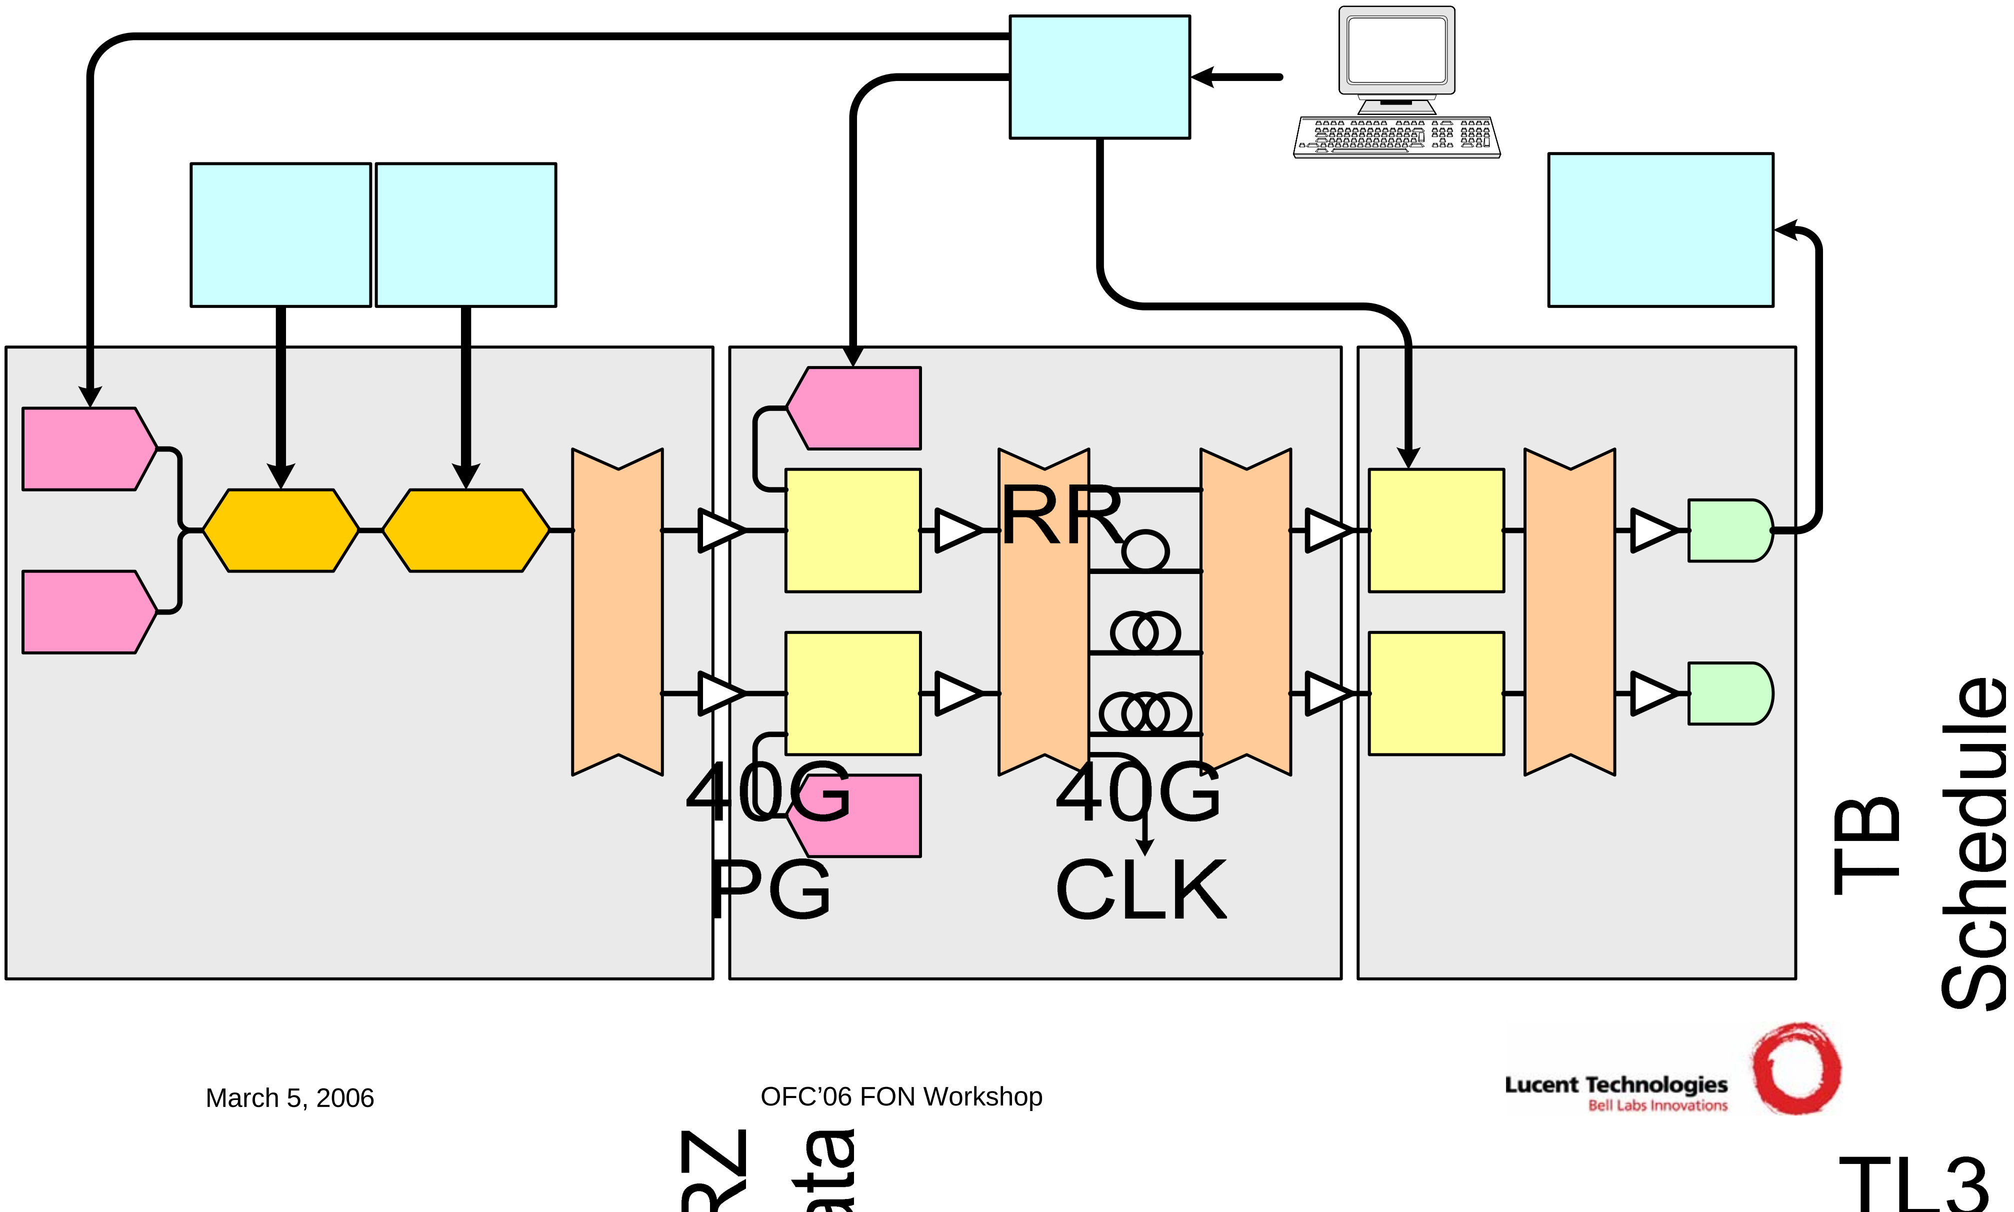

Operation of the Router (Animation)

From Tx1	From Tx2	Delay Tx1	Delay Tx2
→Rx1	→Rx1	0	1
→Rx1	→Rx1	2	1
→Rx2	→Rx1	1	3
→Rx1	→Rx1	(4) drop	1

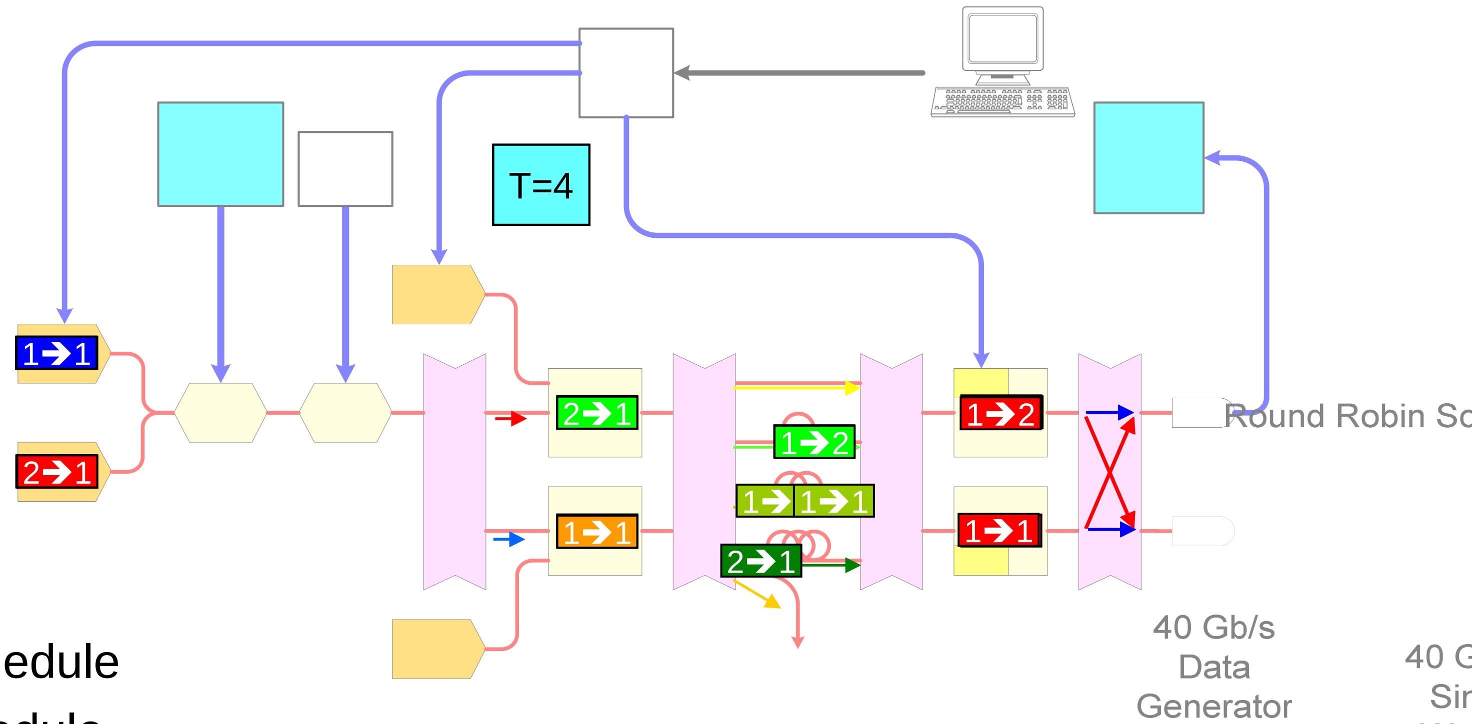

- 1. stage: round robin schedule
- 2. stage: time buffer schedule
- 3. stage: round robin schedule

- all stages switch based on a schedule that simulates random traffic with adjustable traffic load and imbalance

BER performance of the Wavelength Converters

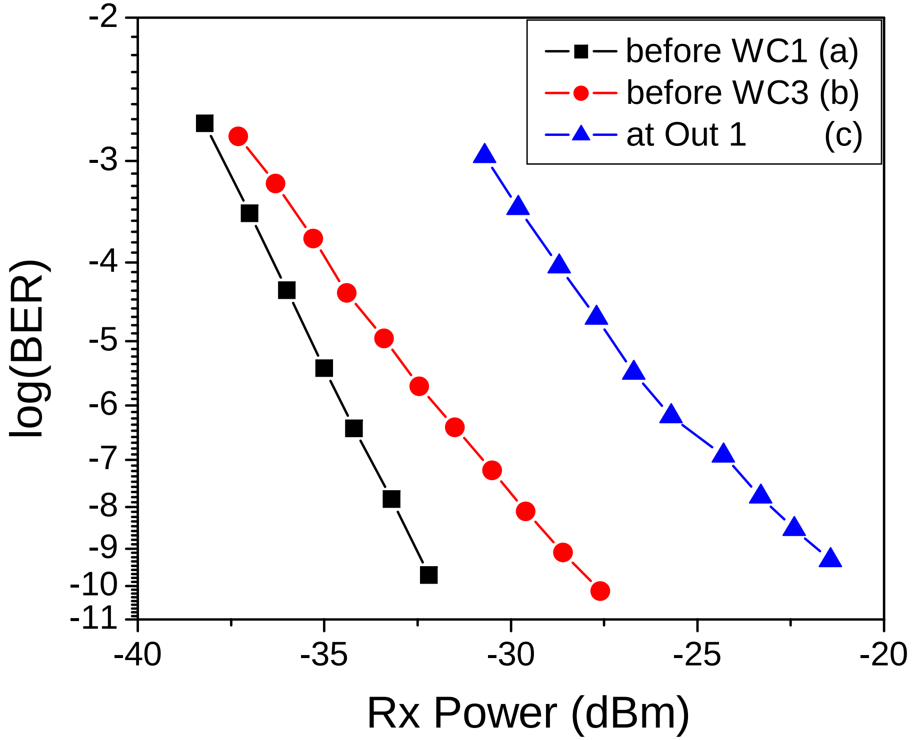

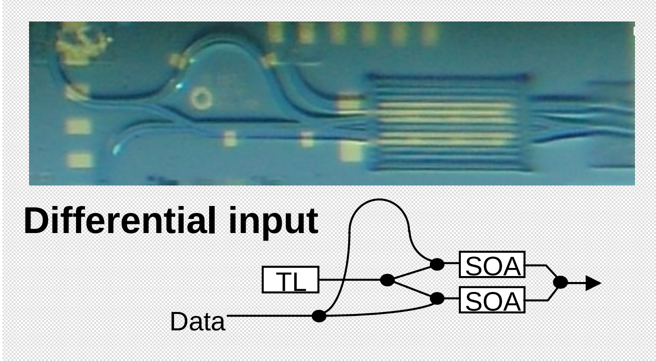

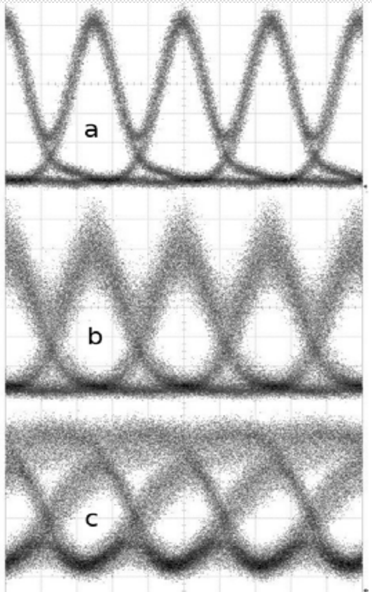

- 4 dB penalty for the first conversion and AWG filtering in the time buffers
- 6.5 dB penalty for second conversion and final AWG pass
- second converter aligned for inverted pulses
- reduced coupling losses and better filter matching should decrease penalties

Packet Switching with Time Buffers & Drop

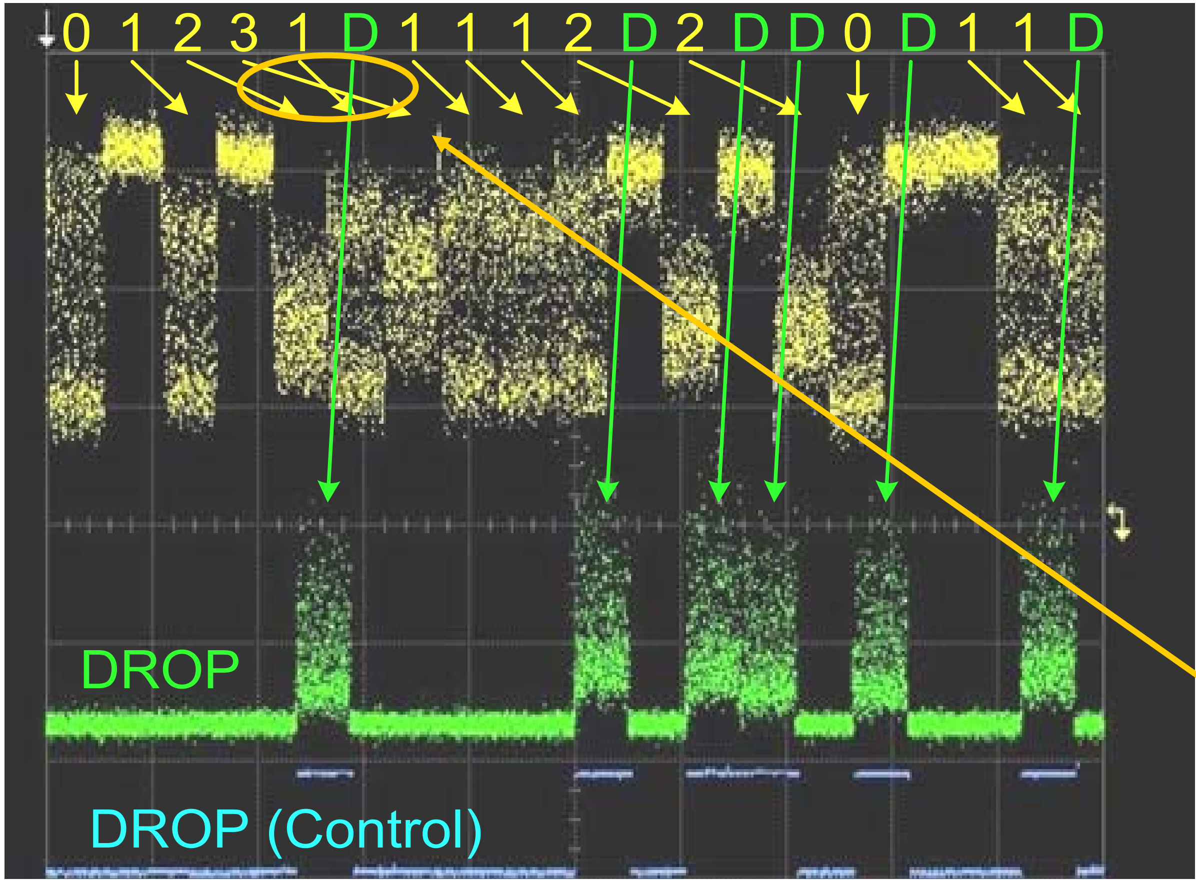

- 1 μs packets, 50 ns guard band, 25 ns haders
- random schedule
  - 20 packets long
  - 80% traffic load
  - 6 packets are dropped
- time buffers shift packets into later time slots to avoid contention
- BERT gates on packets going along one path
- packet re-ordering observed

Lucent Technologies Bell Labs Innovations

Packet Loss Rate and Latency

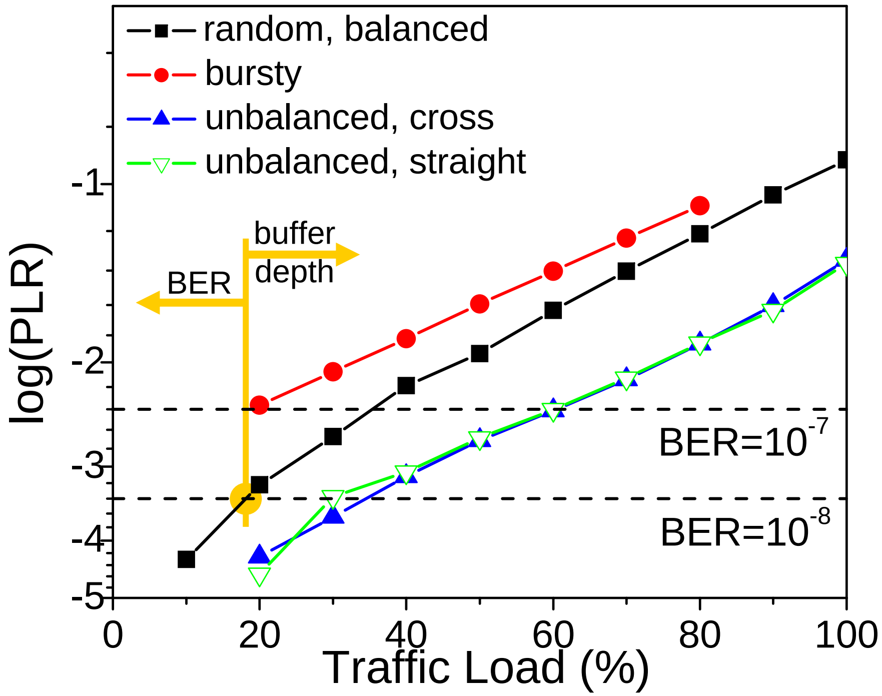

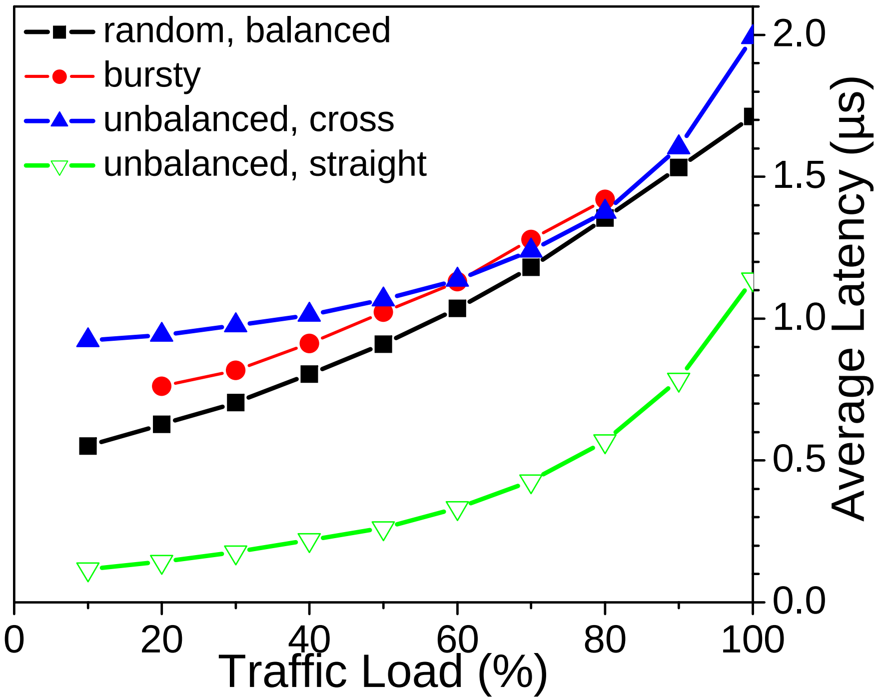

- For high traffic load, the PLR is dominated by the buffer overflow
- For low traffic load, the PLR is dominated by the BER Graph shown for 1 µs packets at 40 Gb/s (40 000 bits)
- Schedule ~ 4000 packets long

$$\text{PLR} < \text{BER} * \frac{\#\text{ of bits}}{\text{packet}}$$

- bursty traffic uses bimodal traffic load distribution (50% at TL – 18%, 50% at TL + 18%)
- unbalanced straight (cross) traffic sends 90% of TL from In1 to Out1 (Out2)

Summary

- Today's electronic router do not scale, Optical Packet Routers can help
  - Power dissipation, fabric and scheduler scalability problem
  - IRIS provides an optical solution to these scalability problems
  - Less O-E-O conversions, passive optical backplane, load-balancing

- Demonstrated a load-balanced all-optical packet router
  - 2x2 ports at 40 Gb/s ➔ 160 Gb/s router
  - complete data plane, including two switching stages and buffering
  - time-buffers handle random traffic, control based on simulation
  - Switching with 50 ns dead-time, 25 ns training pattern and 1 μs payload
  - Data can pass through two wavelength converters without errors
  - tested the router with arbitrary traffic distributions and determined packet loss rates and average latencies