RADWAN | Rate Adaptive Wide Area Networks

Rachee Singh / U. Massachusetts Amherst
Manya Ghobadi / Microsoft Research
Klaus-Tycho Foerster / University of Vienna
Mark Filer / Microsoft Research
Phillipa Gill / U. Massachusetts Amherst

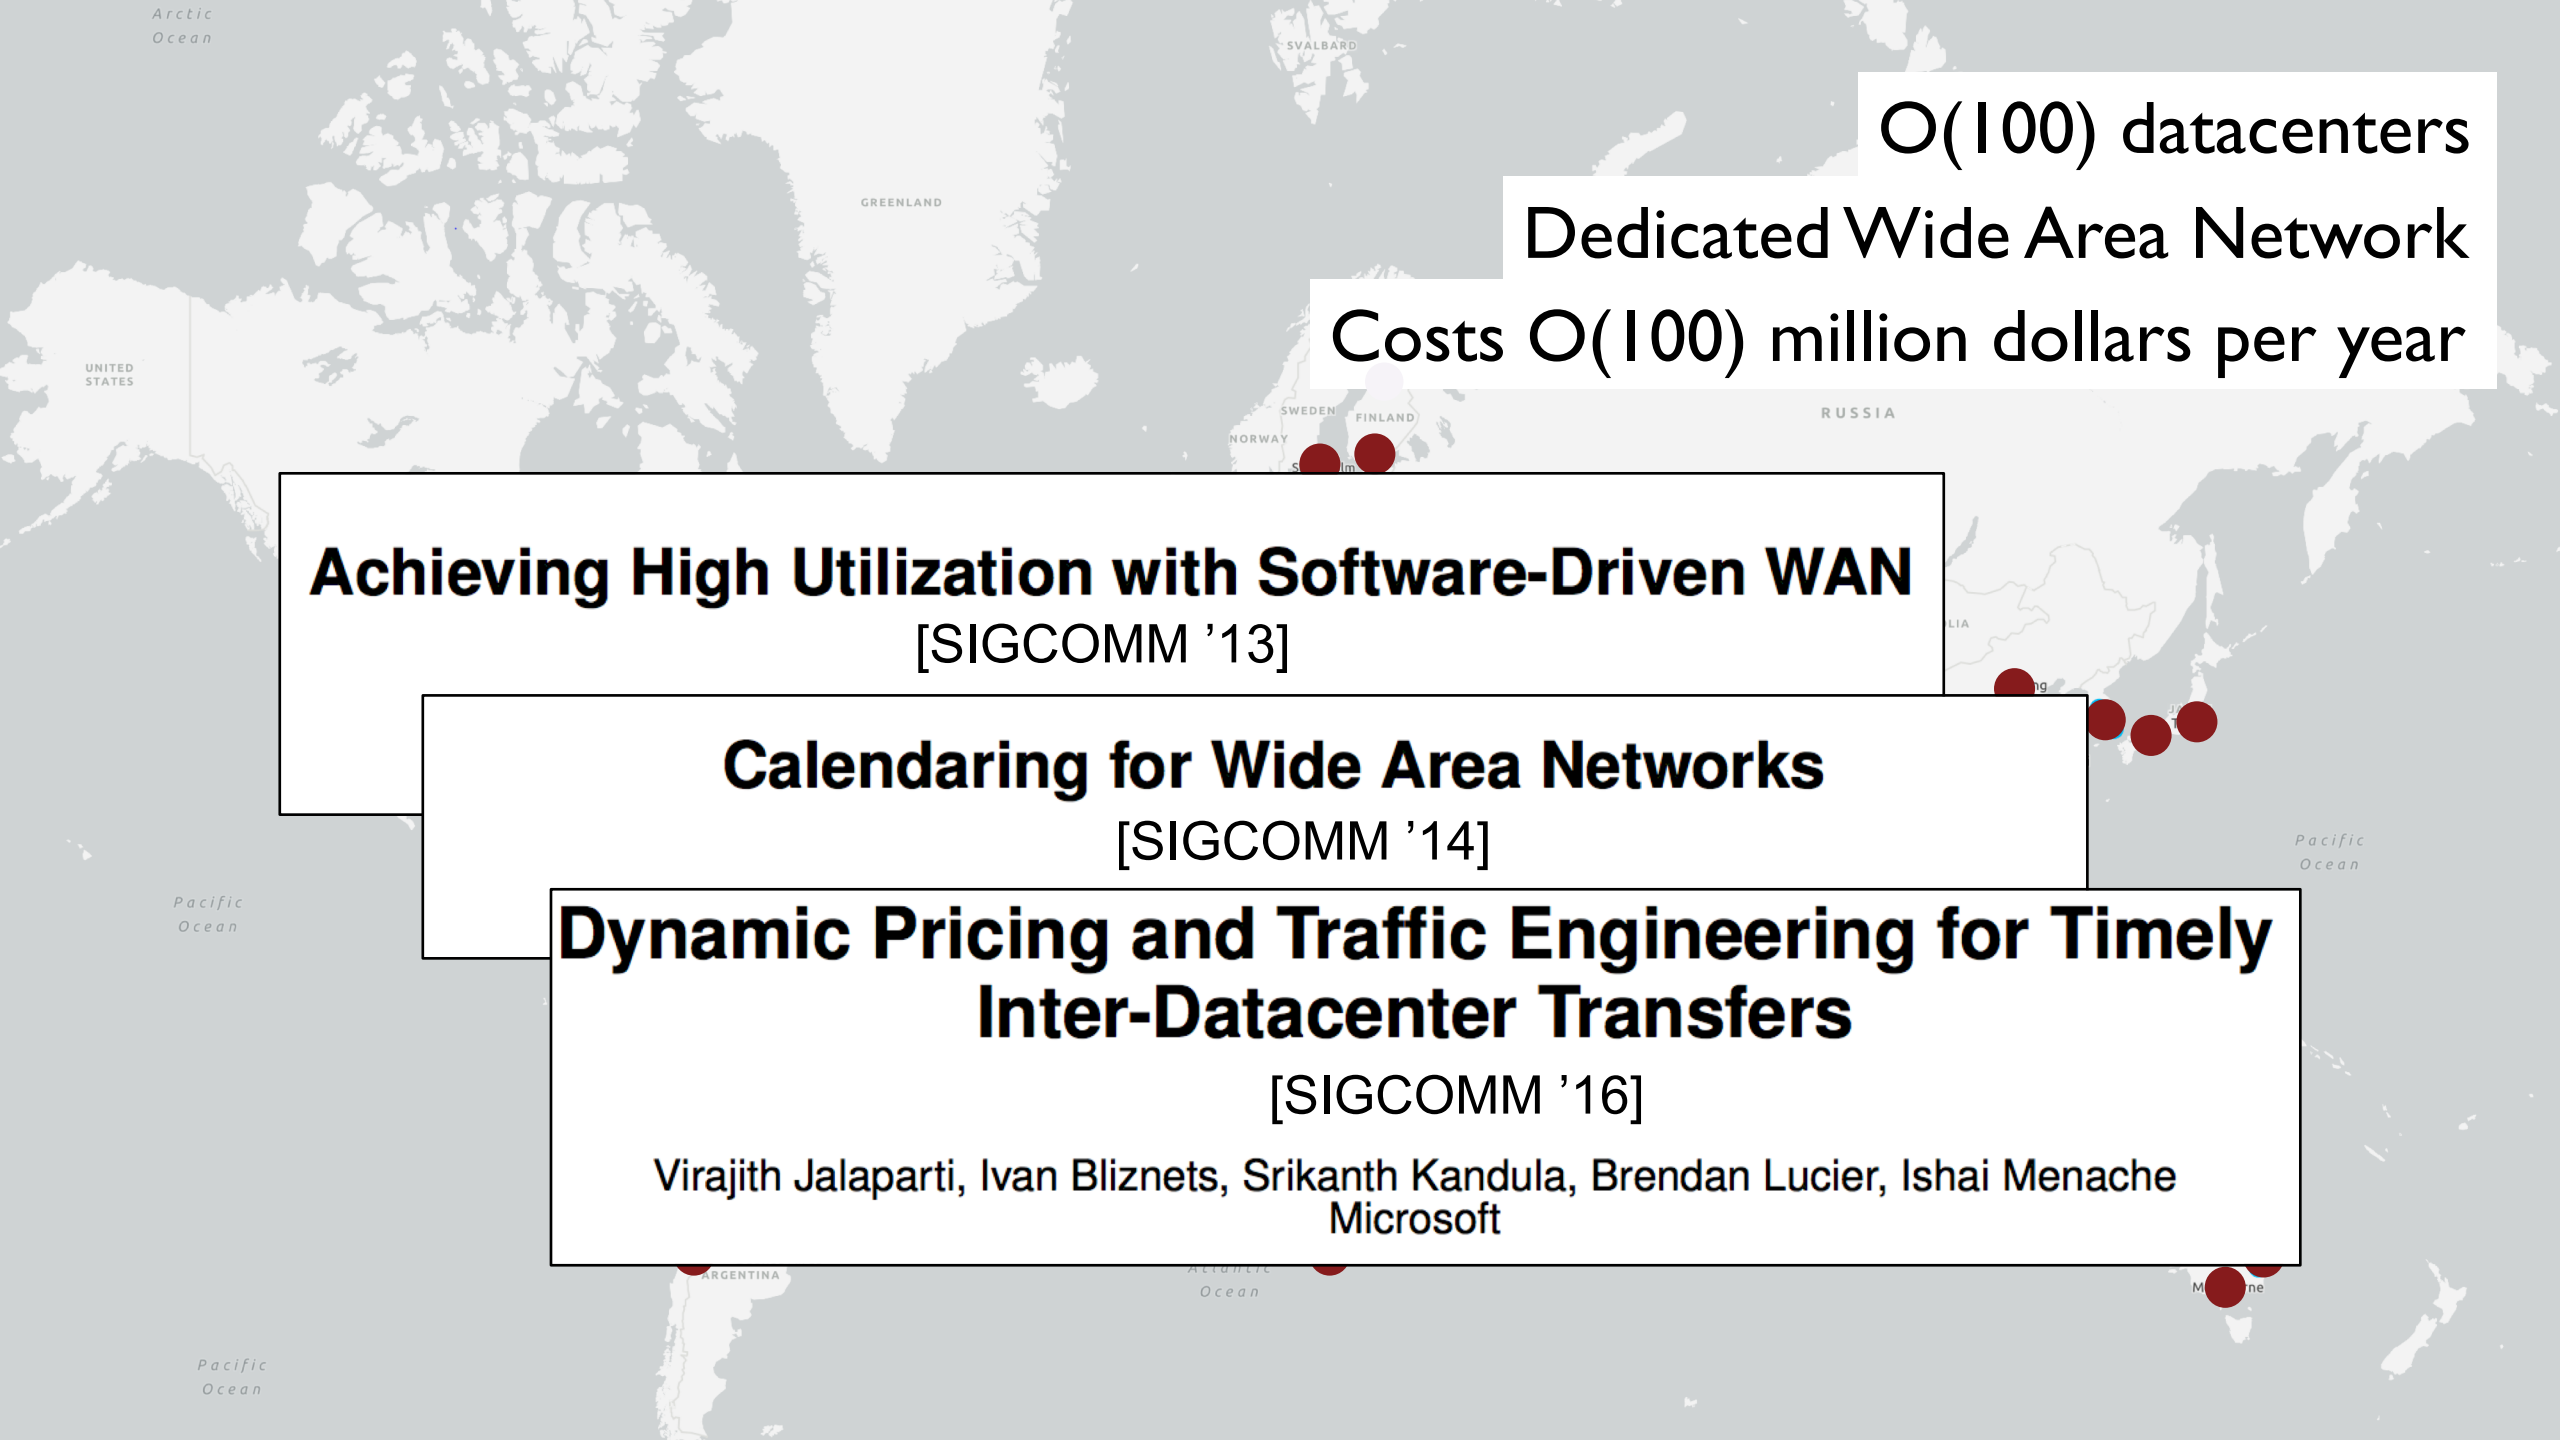
O(100) datacenters
Dedicated Wide Area Network
Costs O(100) million dollars per year
Achieving High Utilization with Software-Driven WAN
[SIGCOMM '13]
Calendaring for Wide Area Networks
[SIGCOMM '14]
Dynamic Pricing and Traffic Engineering for Timely Inter-Datacenter Transfers
[SIGCOMM '16]
Virajith Jalaparti, Ivan Bliznets, Srikanth Kandula, Brendan Lucier, Ishai Menache
Microsoft

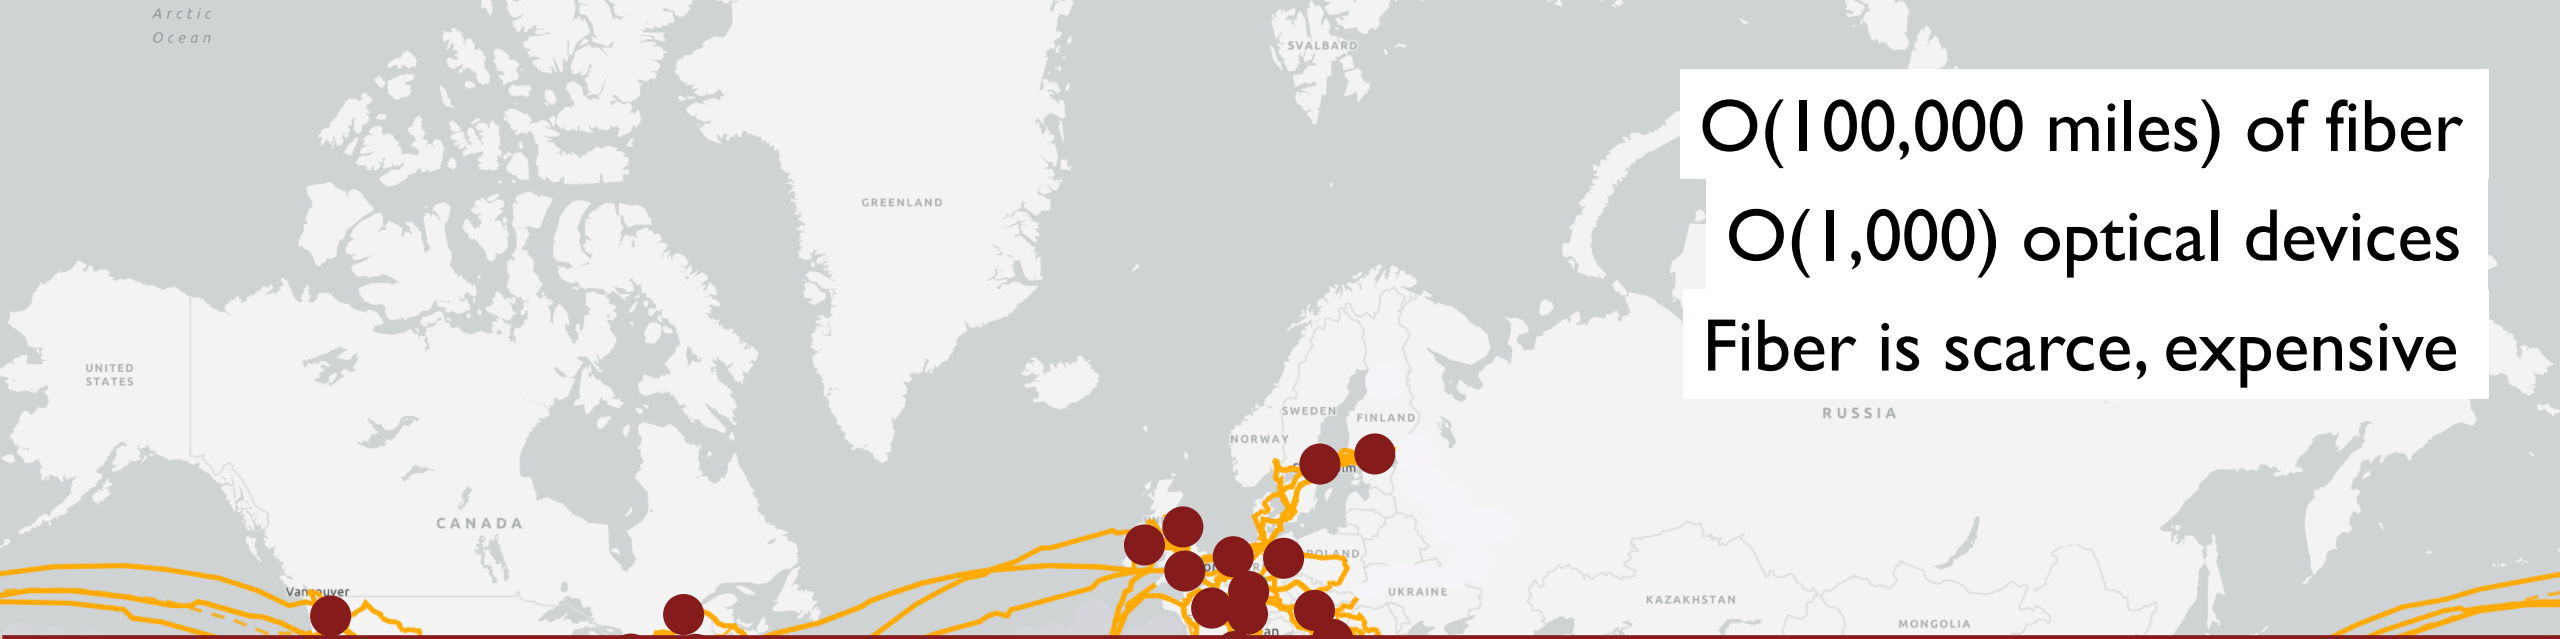

O(100,000 miles) of fiber

O(1,000) optical devices

Fiber is scarce, expensive

Identify inefficiencies in the optical backbone to gain capacity, availability at reduced cost.

This Talk

Gain 134 Tbps of capacity and prevent 25% link failures in large North American WAN.

Talk Outline

1. How inefficient are optical backbones?
2. Dynamic capacity links in WANs
3. Challenges in dynamically adapting link capacities
4. Rate Adaptive WANs

Optical Backbone Networks

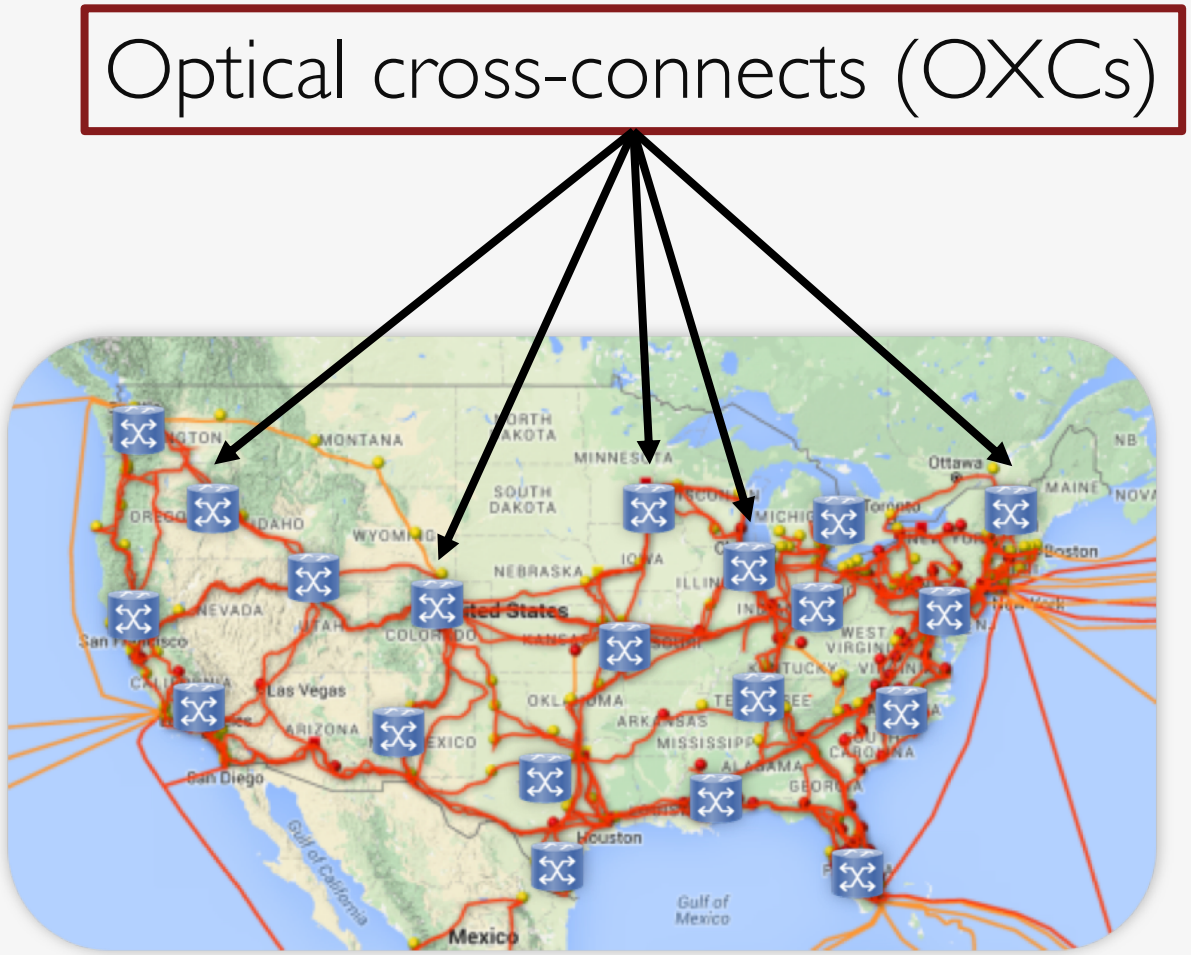

- OXC: switches optical signals

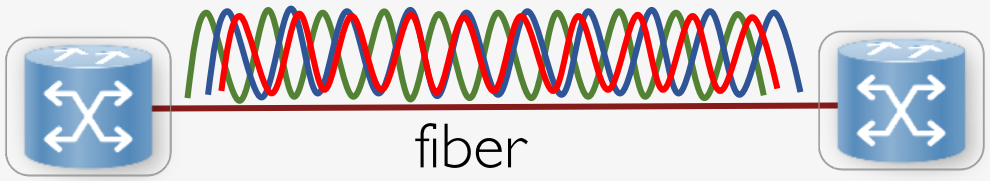

- Signal-to-noise ratio (SNR) measures signal quality
- At OXC, measure signal quality
  - 8,000 wavelengths
  - Every 15 minutes
  - February 2015 to June 2017

Longitudinal Signal Quality on Fiber

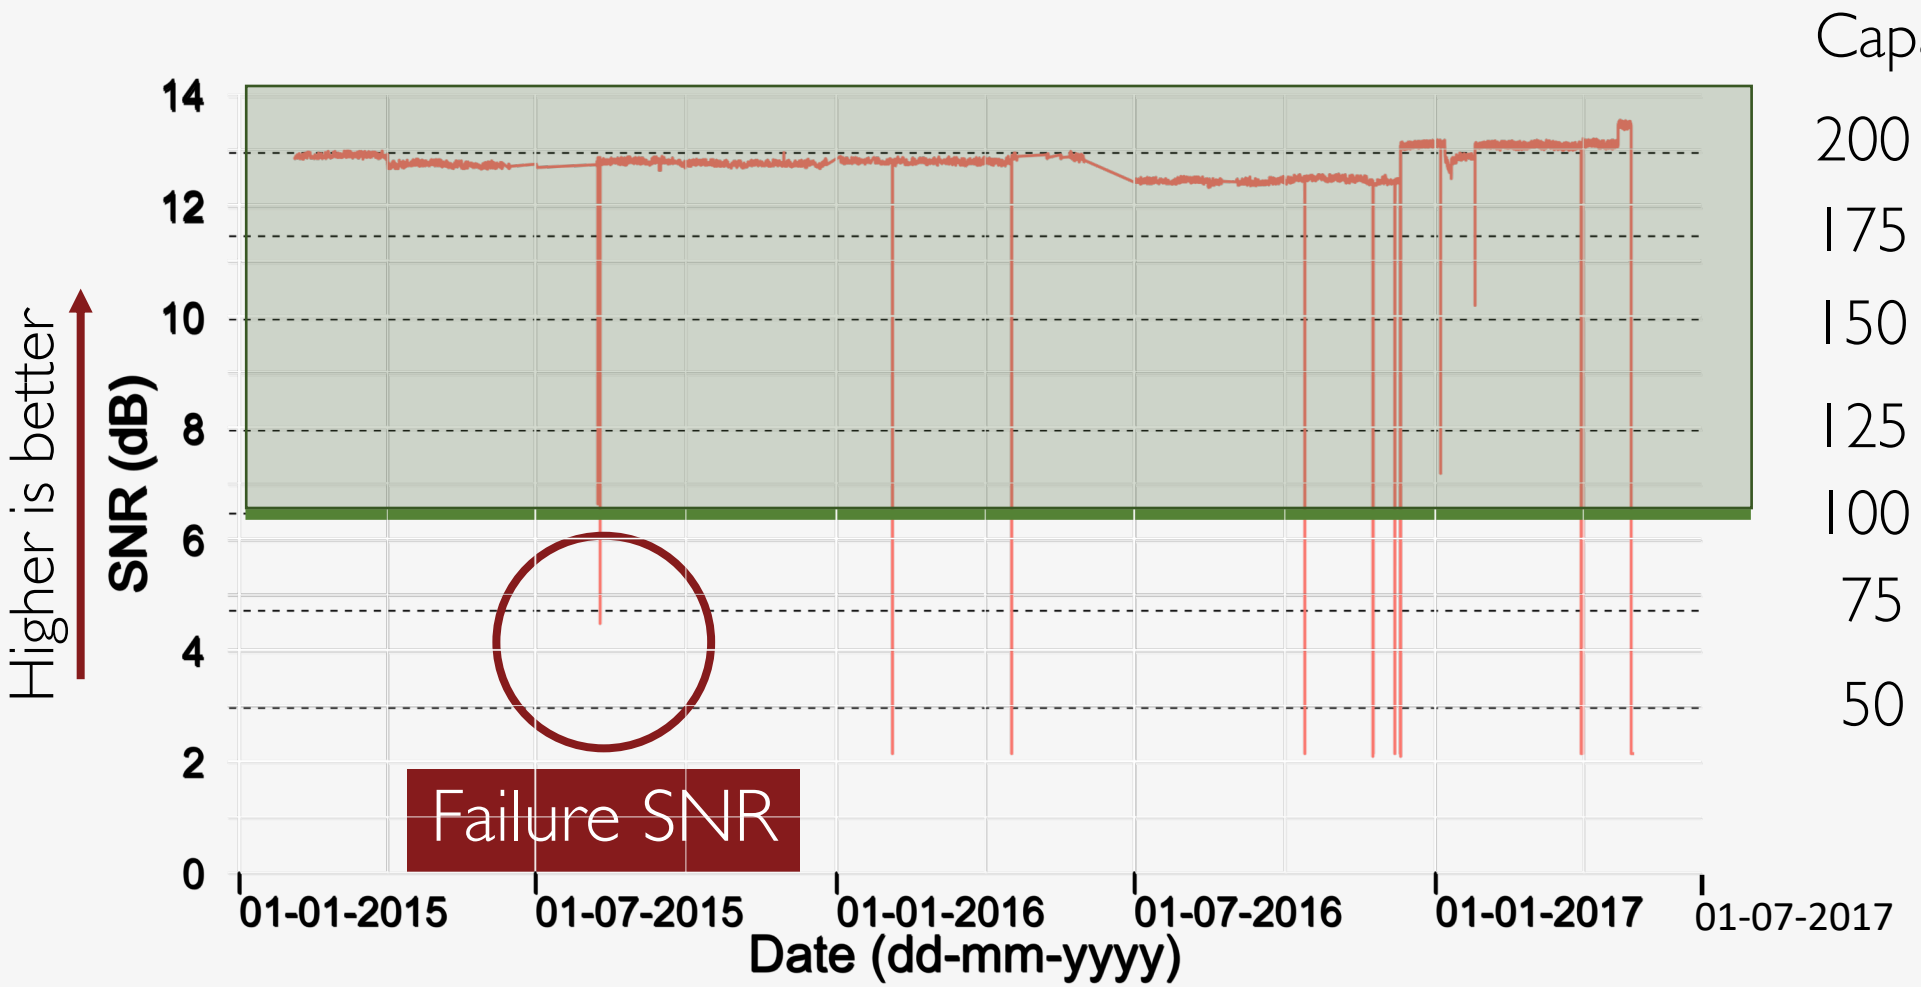

Capacity Threshold

200 Gbps

175 Gbps

150 Gbps

125 Gbps

100 Gbps

75 Gbps

50 Gbps

Opportunity for capacity gain

For 8,000 wavelengths in WAN:

- Analyze average SNR
- Compare with thresholds for link capacity

64% of optical wavelengths can operate at 175 Gbps or more.

95% of optical wavelengths can operate at higher than 100 Gbps.

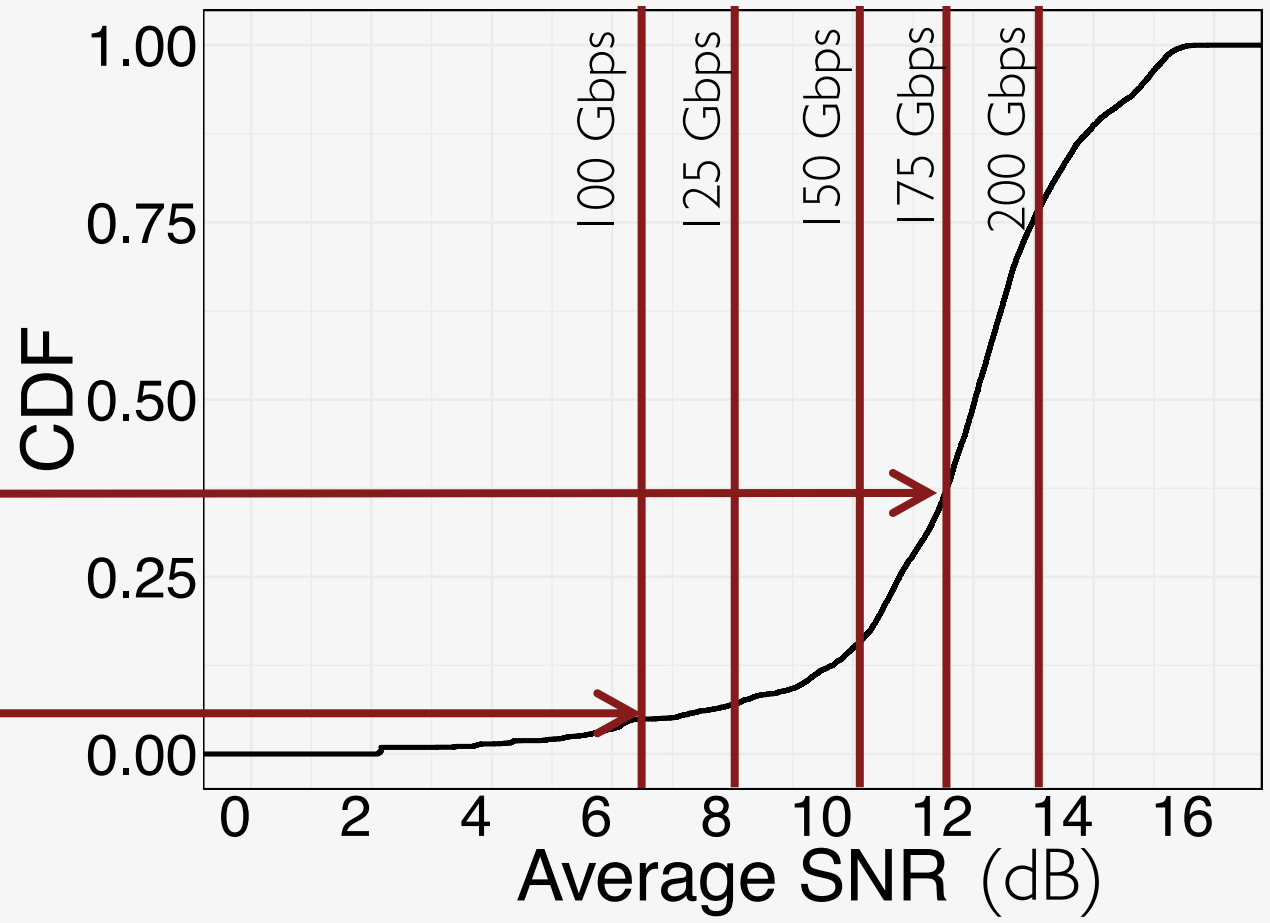

Opportunity for availability gain

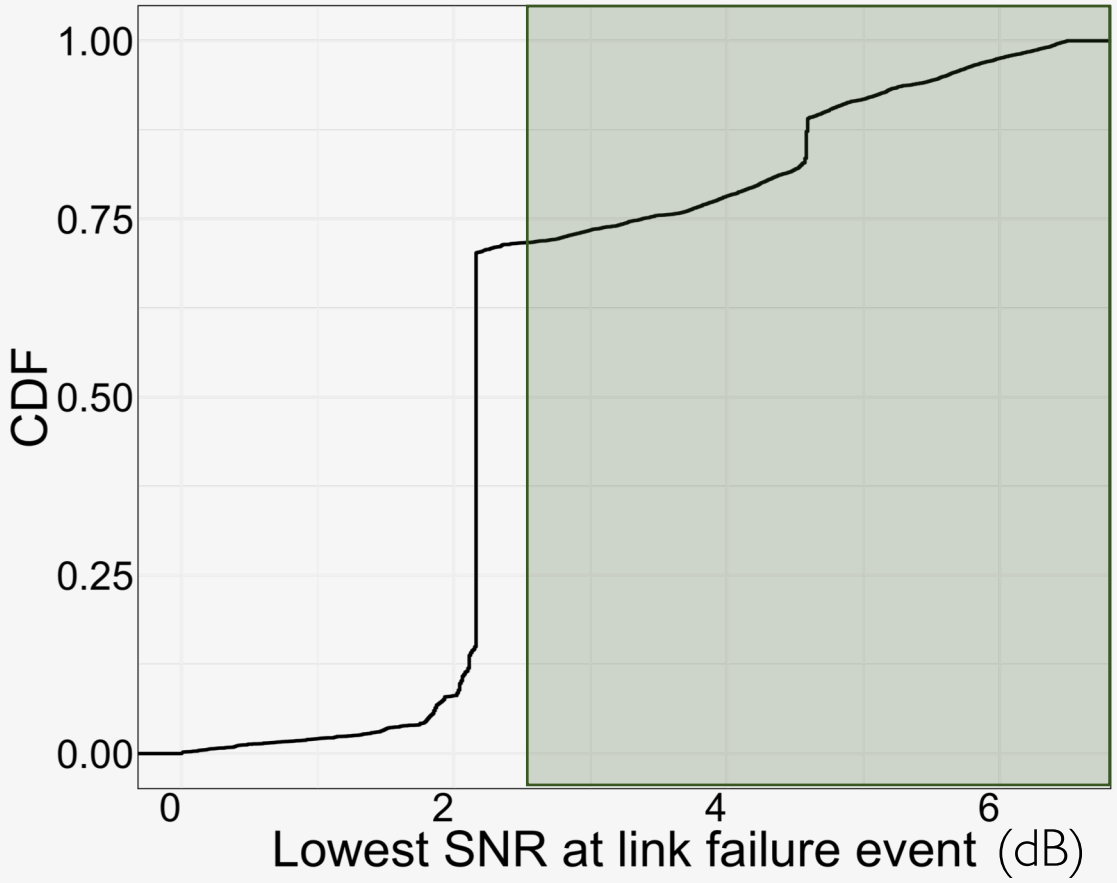

- Distribution of link failure SNR
  - Across WAN links
  - For 2.5 years

25% of failures have SNR > 2.5dB

These failures can be prevented by reducing link capacity to 50 Gbps

Our Proposal

Dynamically adapt link capacities in response to changes in SNR.

Gain 134 Tbps capacity	Prevent 25% link failures
By increasing link capacity when high SNR	By reducing link capacity when low SNR

Talk Outline

1. How inefficient are optical backbones?
2. Dynamic capacity links in WANs
3. Challenges in dynamically adapting link capacities
4. Rate Adaptive WANs

Challenges in dynamically adapting link capacities

1. Requires hardware support for capacity reconfiguration
2. Requires re-thinking IP layer traffic engineering

Key question

Can we use commodity hardware for changing link capacities?

Bandwidth Variable Transceiver

- Supports higher order modulations (QPSK, 8-QAM, 16-QAM)
- Link capacity of 100G, 150G, 200G

Arista 7504 linecards

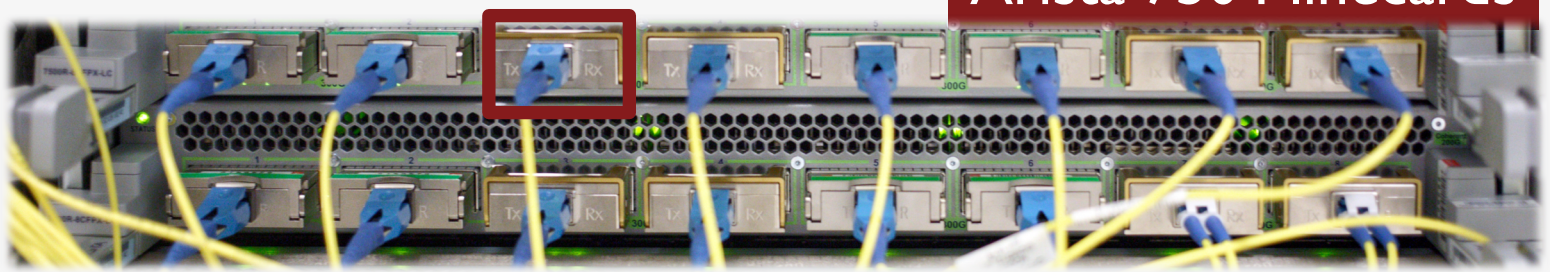

Challenge 1: Adapting capacity on commodity h/w

Arista 7504 Chassis

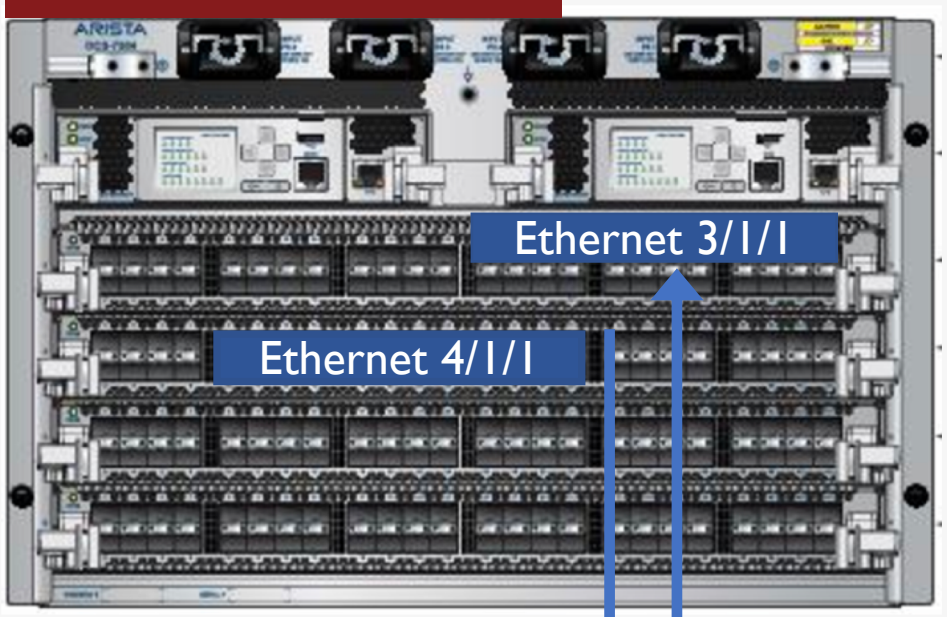

Variable Optical Attenuator

Increasing noise from attenuator

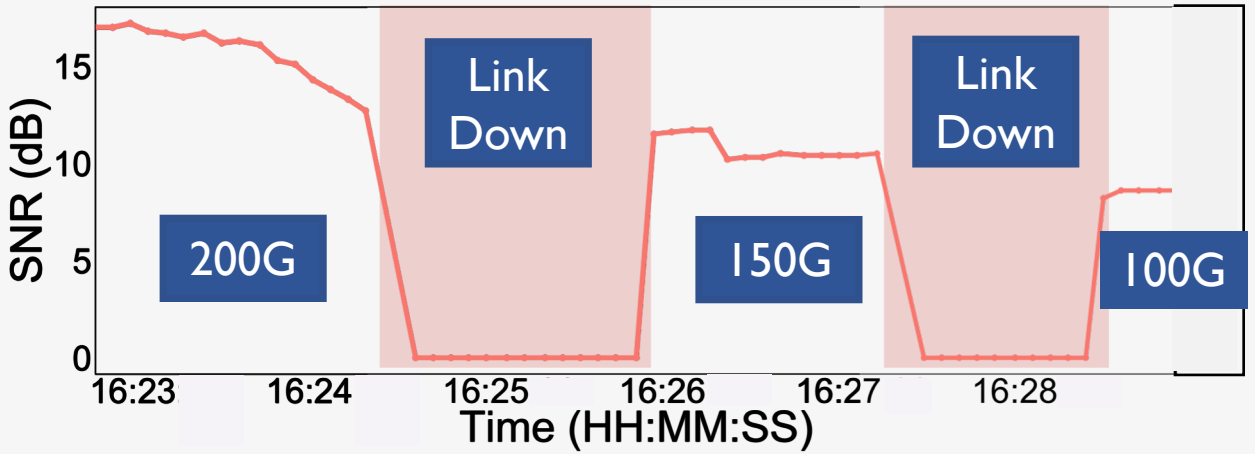

Capacity Downgrade to 150G

Capacity Downgrade to 100G

Takes over 1 minute to change capacity → link downtime

Problem

Commodity hardware is not optimized for dynamically adapting link capacity.

Question

What causes latency of capacity reconfiguration?

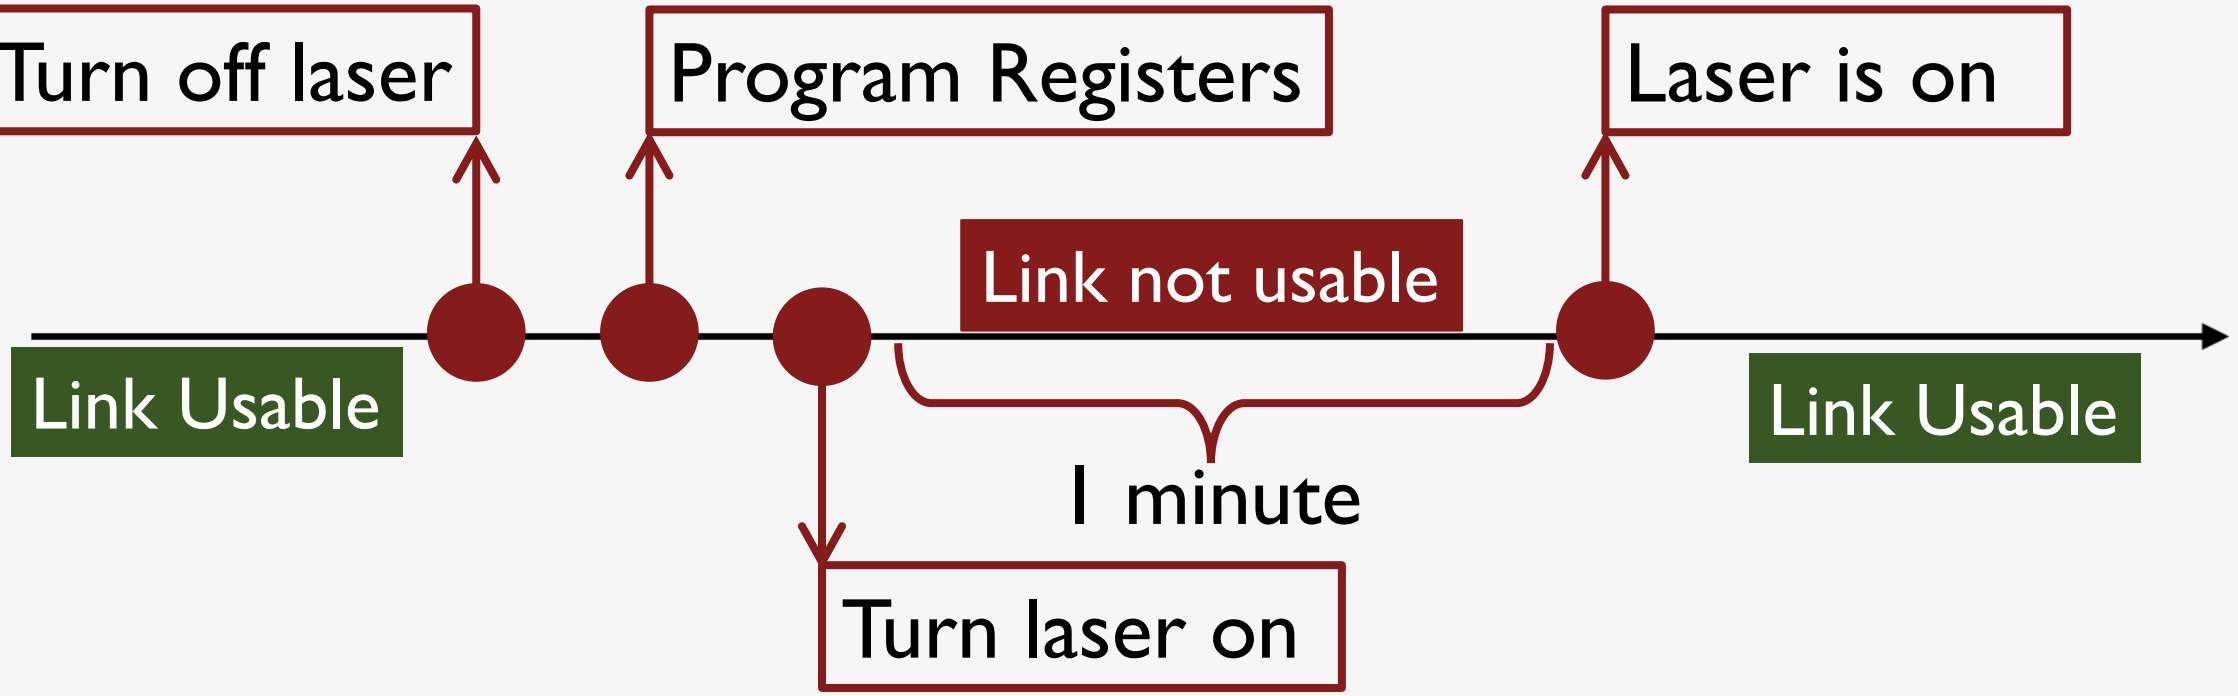

Majority of time spent in turning laser on.

Question

Can we reduce the latency of capacity reconfiguration by not turning off the laser?

Question

Can we reduce the latency of capacity reconfiguration by not turning off the laser?

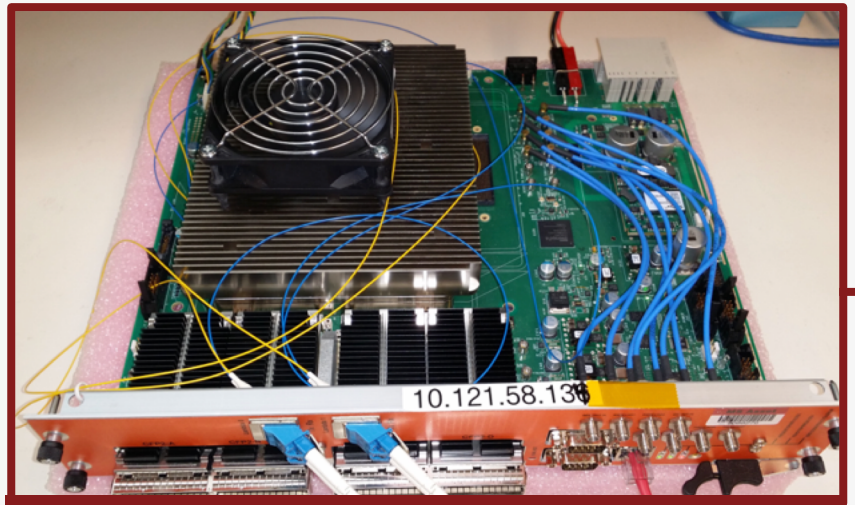

Acacia BVT Evaluation Board

Program registers for modulation change

Do not turn off laser in the evaluation board

Repeat experiment 200X

If the laser is left on, the outage is only 35ms to change capacity

Key question

How should traffic engineering incorporate dynamic capacity links?

Question

How should traffic engineering incorporate dynamic capacity links?

Capacity changes cause links to be unavailable for carrying traffic.

Capacity changes lead to network churn and can be disruptive.

Talk Outline

1. How inefficient are optical backbones?
2. Dynamic capacity links in WANs
3. Challenges in dynamically adapting link capacities
4. **Rate Adaptive WANs**

Solution

We design the Rate Adaptive Wide Area Network (RADWAN) traffic engineering controller.

SNR-aware	Rate Adaptive	Minimally disruptive
Knows possible capacity gain of each link	Adapts link rates to meet demands and improve availability	Reconfigure capacity while minimizing network churn

RADWAN Traffic Engineering Formulation

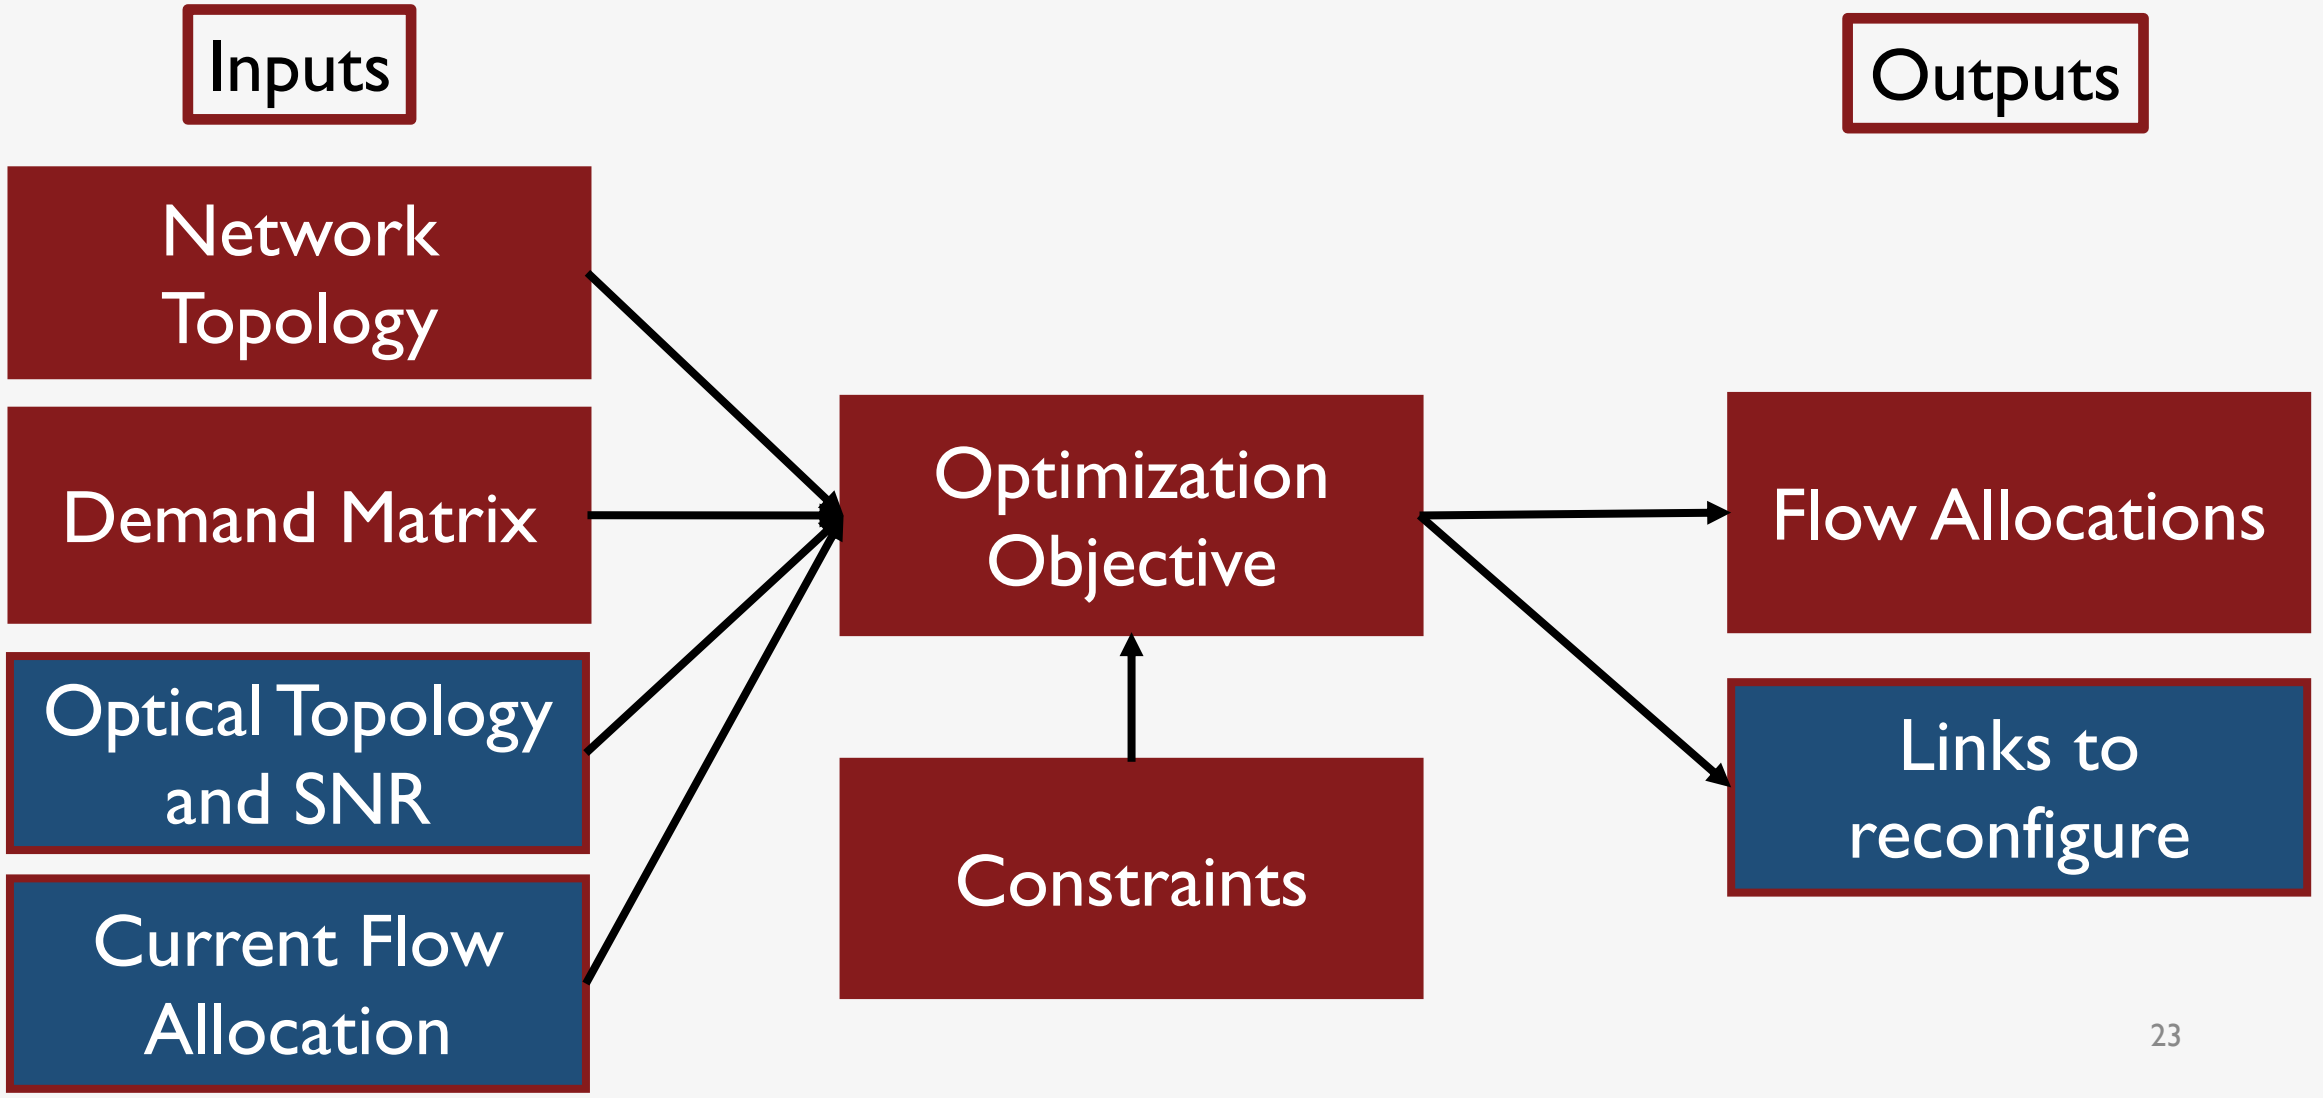

Proof of concept: RADWAN

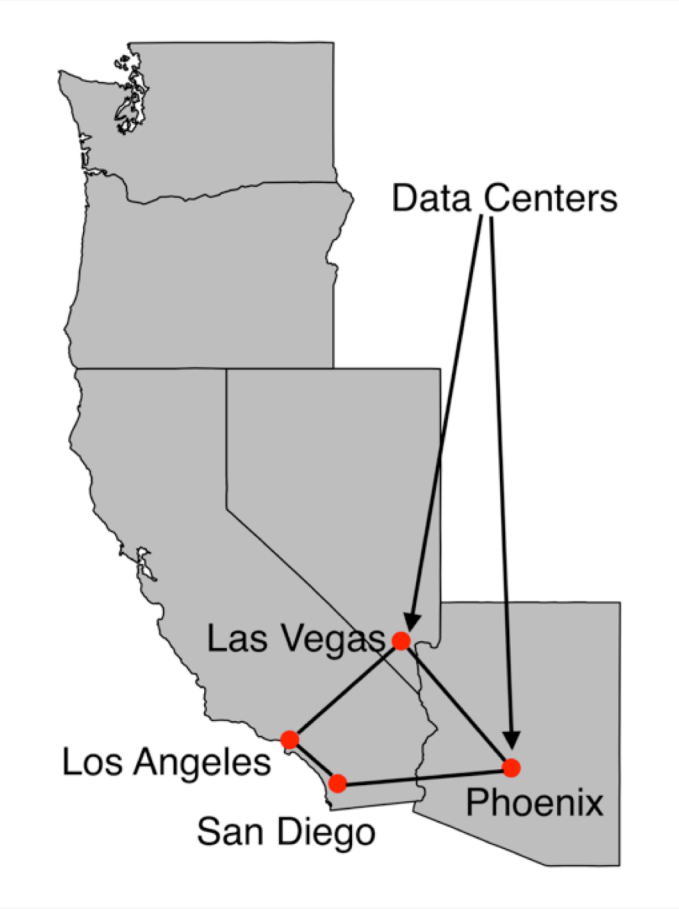

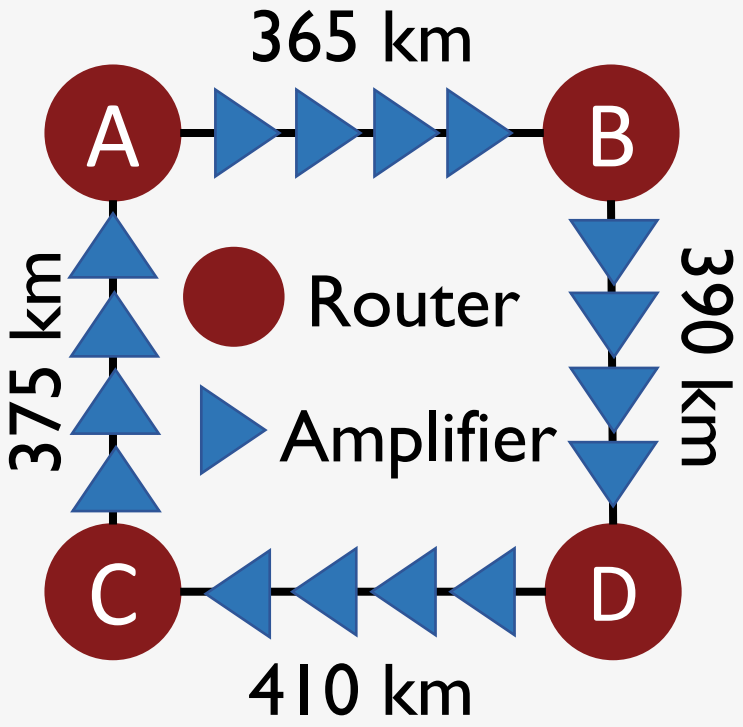

Throughput Gains with RADWAN

RADWAN-hitless

RADWAN

SWAN-150

SWAN [SIGCOMM '13]

RADWAN has 40% Higher network throughput compared to SWAN

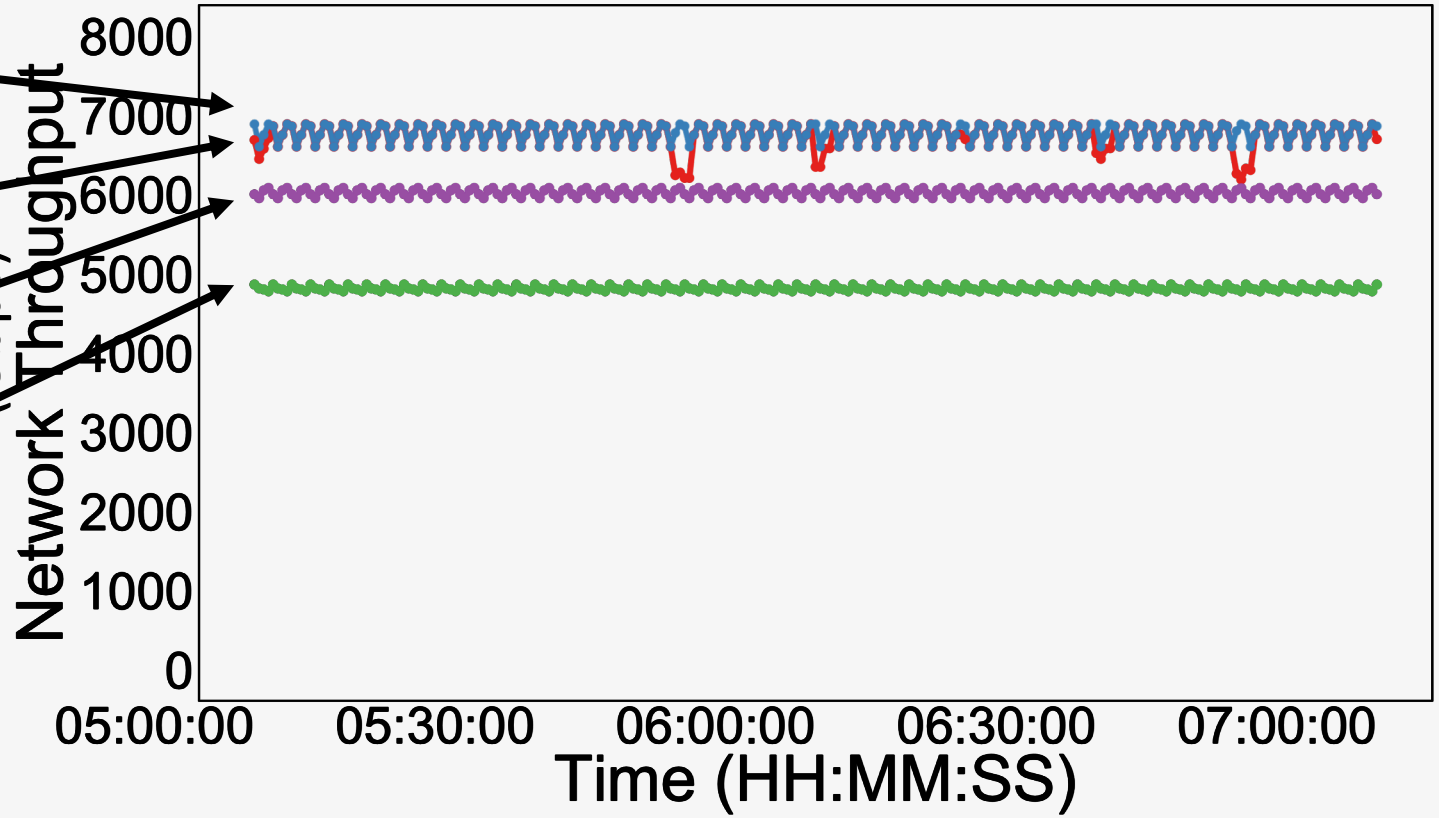

Conclusion

- Physical layer today is configured statically
- We show that this leaves money on the table, in terms of
  - Network performance capacity
  - Link availability
  - Equipment cost ($/Gbps)

- RADWAN introduces programmability in Layer 1
  - Improves network throughput by 40%
  - Reduces link downtime by a factor of 18
  - Reduces equipment cost ($/Gbps) by 32%