

# RADWAN

Rate Adaptive Wide Area Networks

Rachee Singh / U. Massachusetts Amherst (now at MSR)

Manya Ghobadi / Microsoft Research (now at MIT)

**Klaus-T. 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]

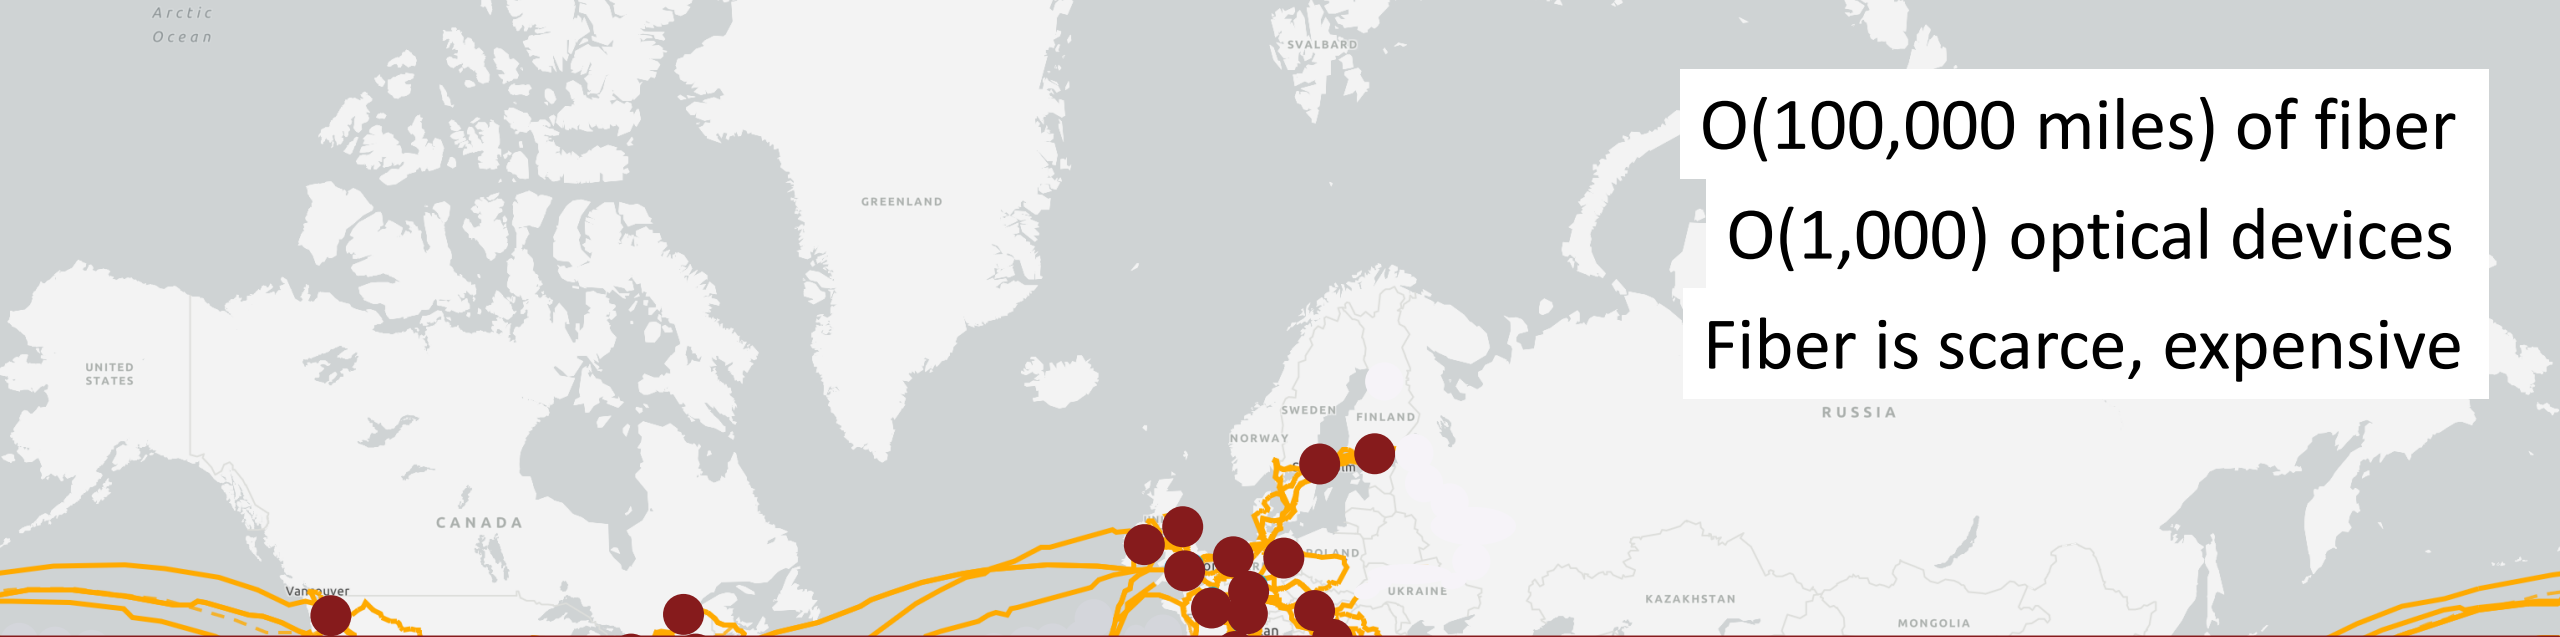
## **Calendaring for Wide Area Networks**

[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.



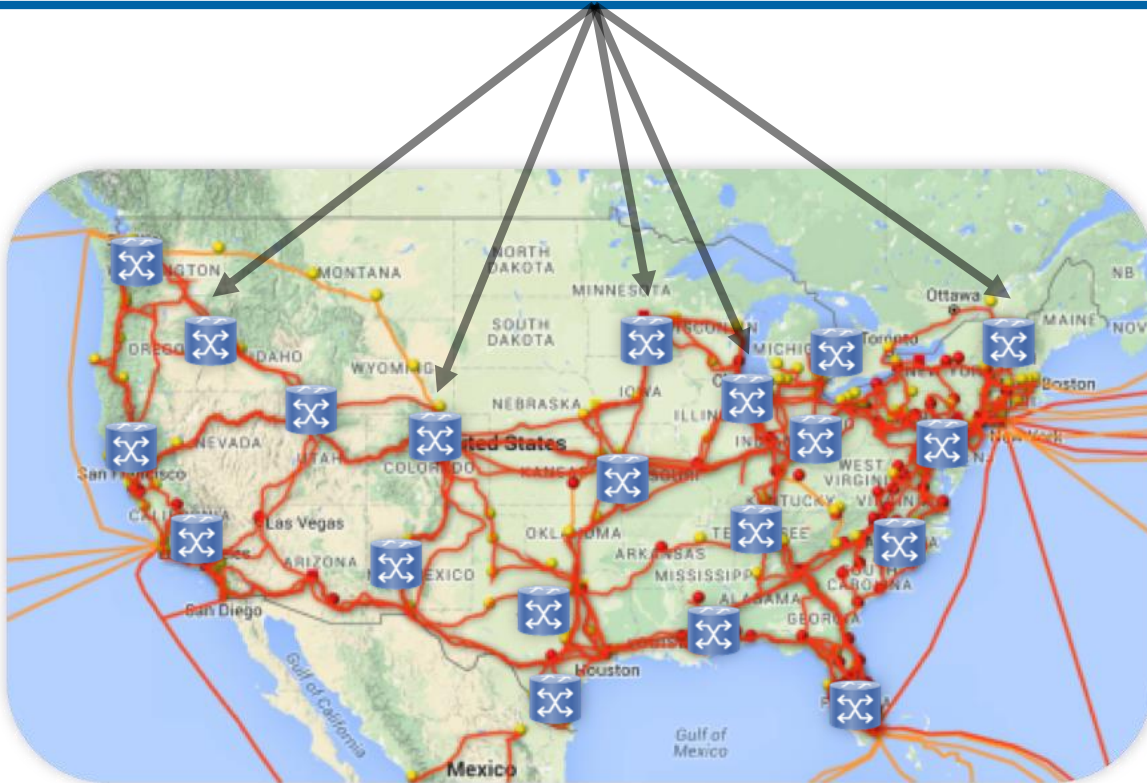
**Gain 134 Tbps of capacity and prevent 25%  
link failures in large North American WAN.**

## Outline


- How inefficient are optical backbones?
- Dynamic capacity links in WANs
- Challenges in dynamically adapting link capacities
- Rate Adaptive WANs

## Optical Backbone Networks

### Optical cross-connects (OXCs)

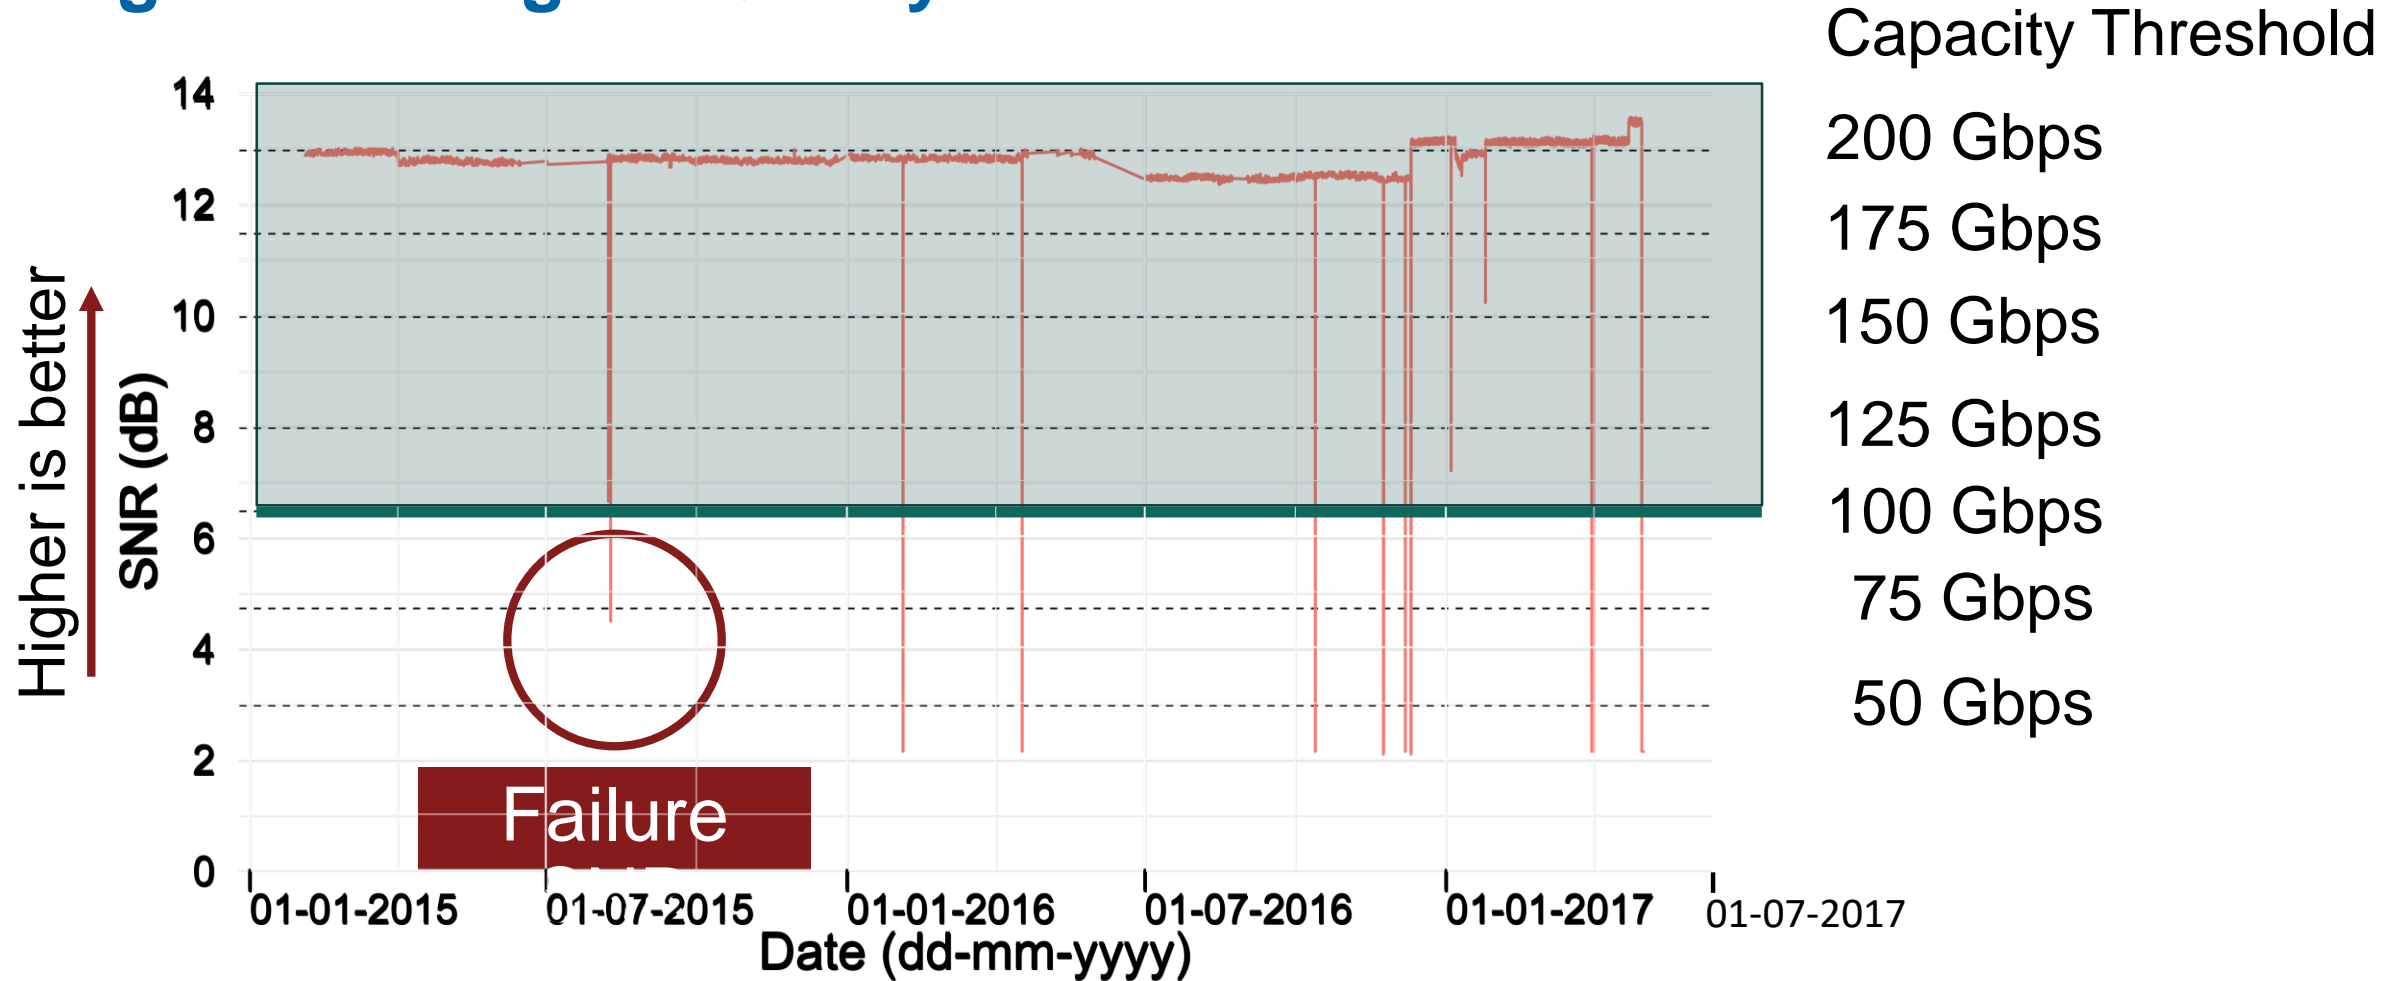


- OXC: switches optical signals



- 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



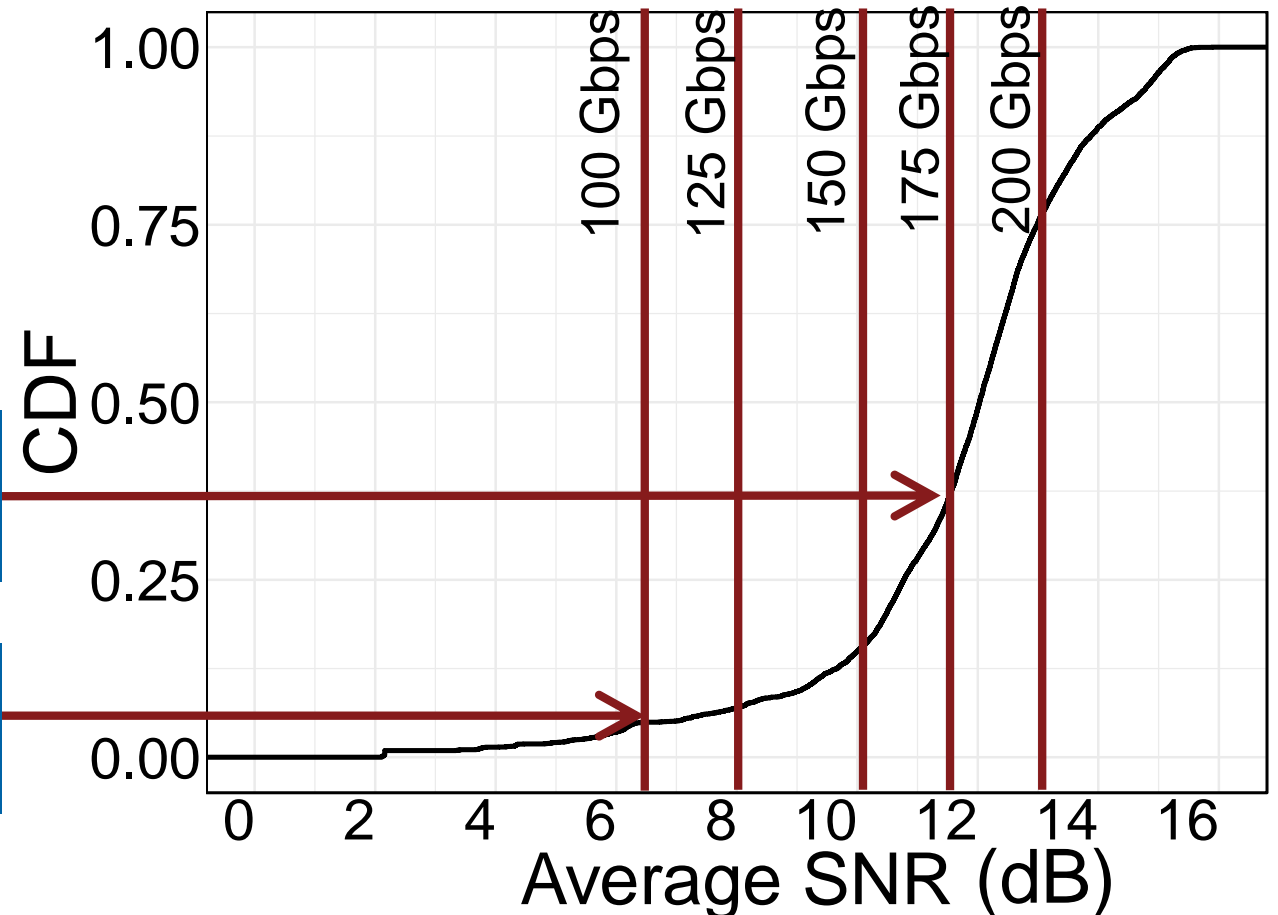
## Opportunity for capacity gain

For 8,000 wavelengths in WAN:

- Analyze average SNR
- Compare with thresholds for link capacity

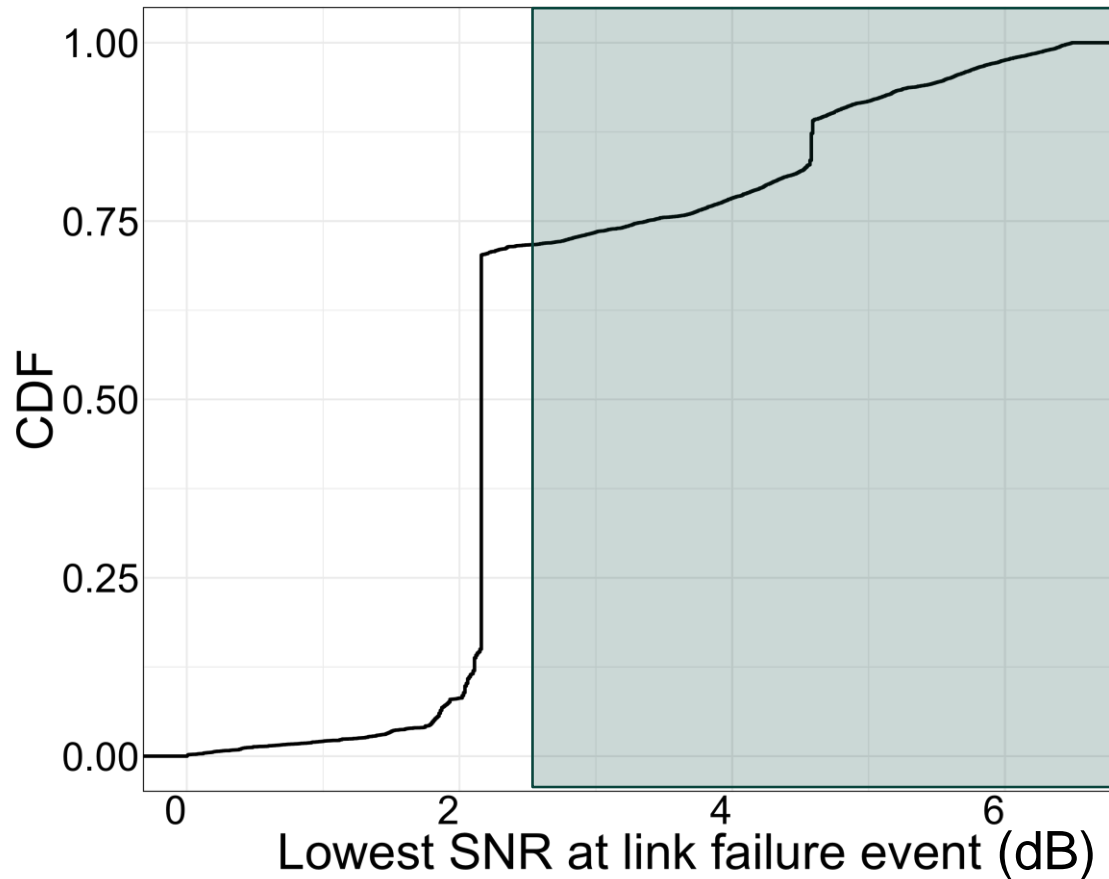
64% of optical wavelengths  
can operate at 175 Gbps

95% of optical wavelengths  
can operate at higher than 100





## 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**

By increasing  
link capacity  
when high SNR

**Prevent 25%  
link failures**

By reducing link  
capacity when  
low SNR

## Outline

- How inefficient are optical backbones?
- Dynamic capacity links in WANs
- Challenges in dynamically adapting link capacities
- Rate Adaptive WANs

## Challenges in dynamically adapting link capacities

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

## Hardware support for capacity reconfiguration

Small scale lab experiments show:

- Commodity hardware takes over **1 minute of link downtime** to change capacity
- Able to reduce to 35ms with evaluation board

## 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.**

## Outline

- How inefficient are optical backbones?
- Dynamic capacity links in WANs
- Challenges in dynamically adapting link capacities
- Rate Adaptive WANs

## Solution

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

### SNR-aware

Knows possible capacity gain of each link

### Rate Adaptive

Adapts link rates to meet demands and improve availability

### Minimally disruptive

Reconfigure capacity while minimizing network churn

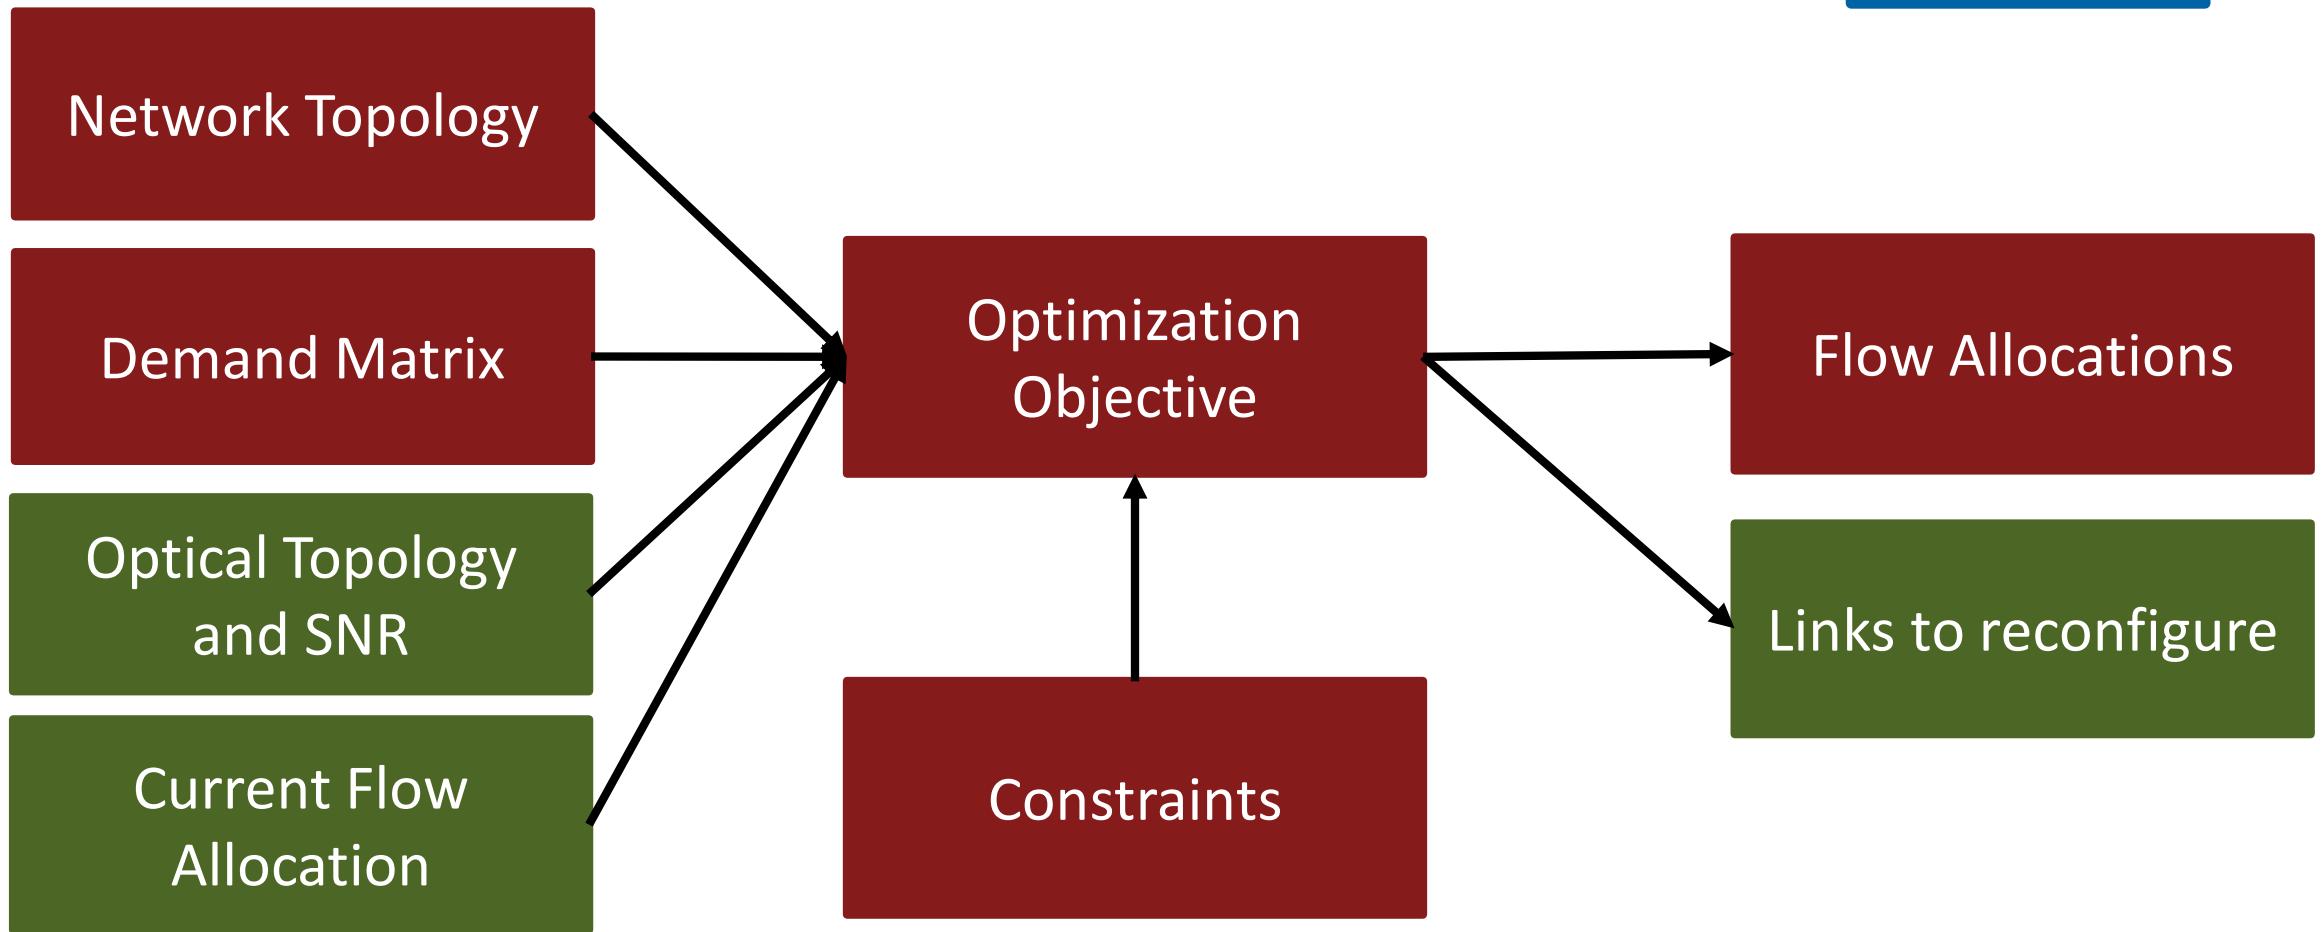





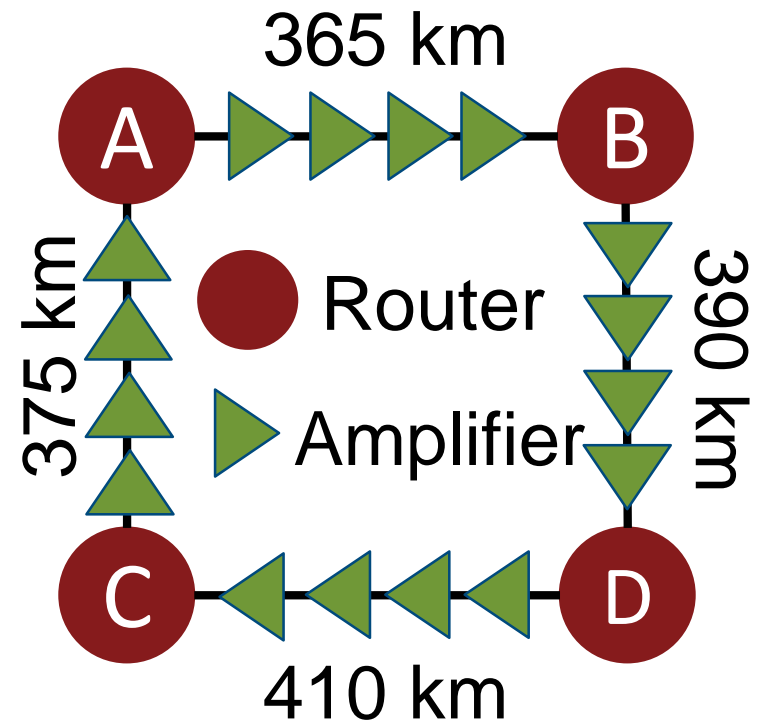

# RADWAN Traffic Engineering Formulation

Inputs

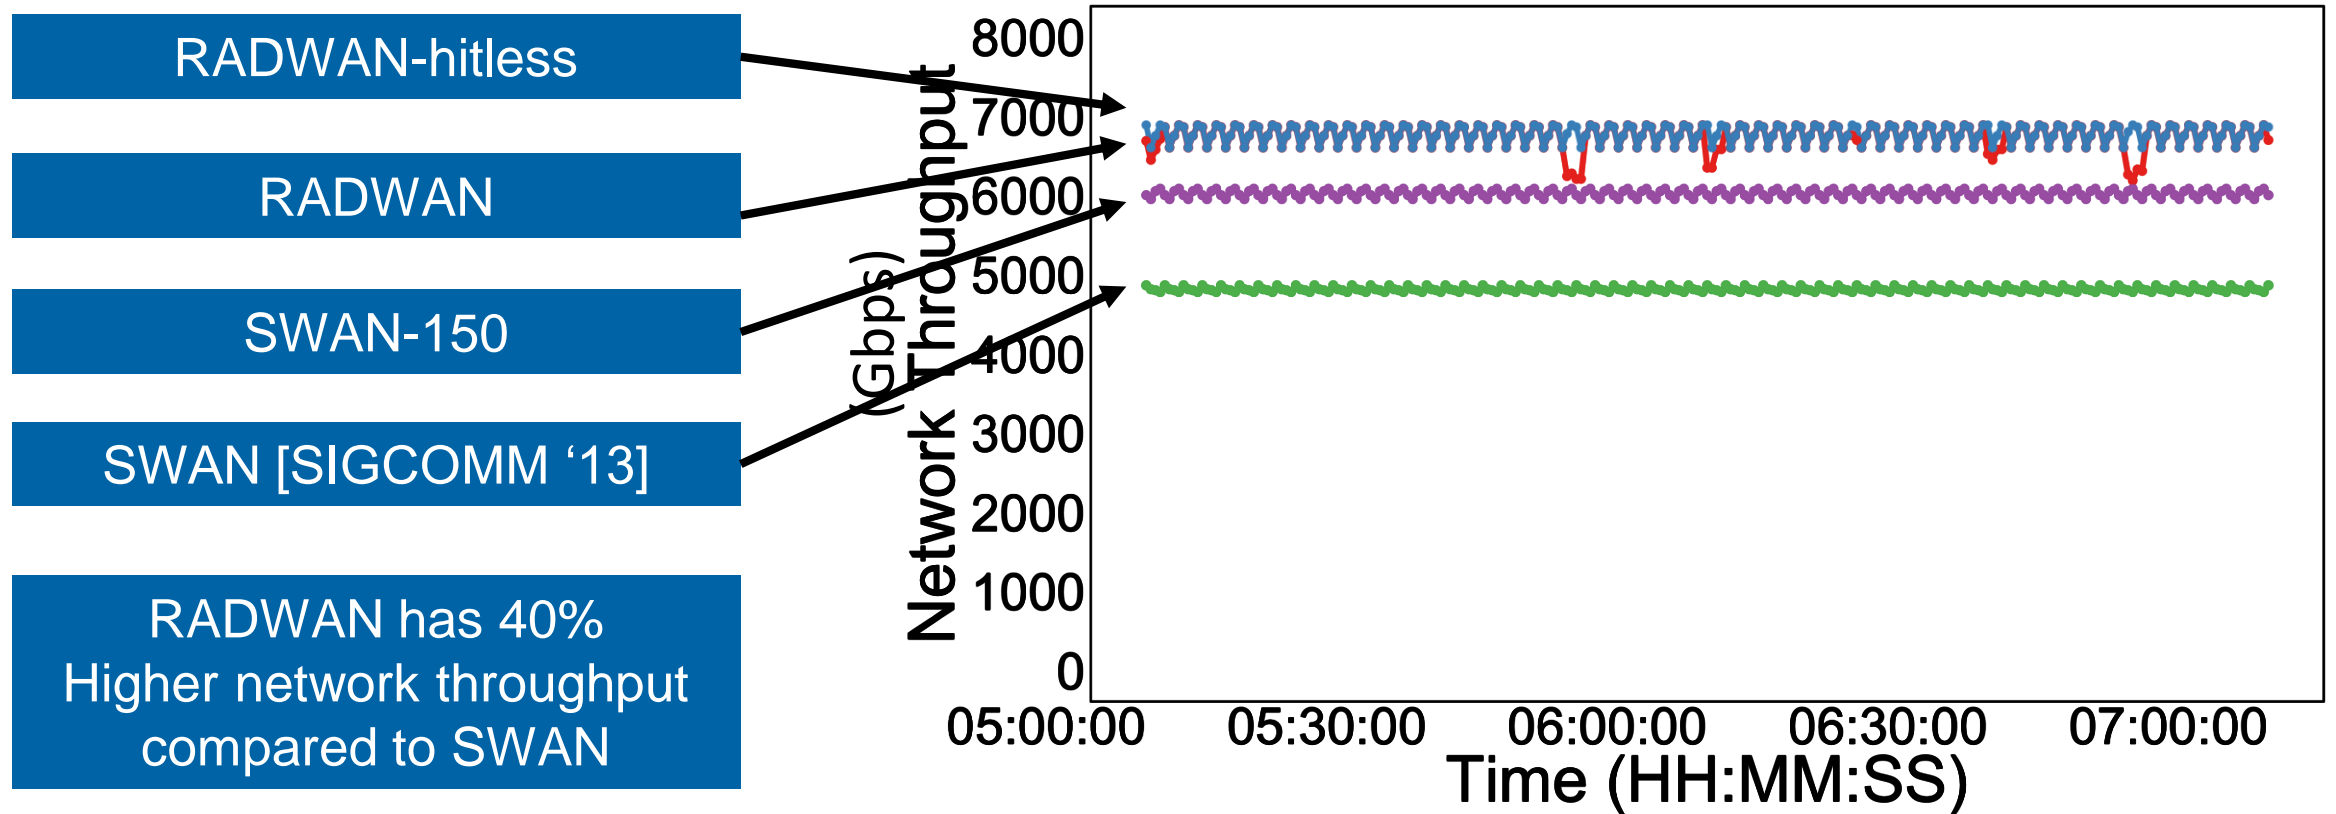
Outputs



## Proof of concept: RADWAN



## Throughput Gains with RADWAN



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%**

# RADWAN

Rate Adaptive Wide Area Networks

Rachee Singh / U. Massachusetts Amherst

Manya Ghobadi / Microsoft Research

**Klaus-T. Foerster** / University of Vienna

Mark Filer / Microsoft Research

Phillipa Gill / U. Massachusetts Amherst

---

# Interested in an Overview on Algorithmic Problems in Reconfigurable Networks?

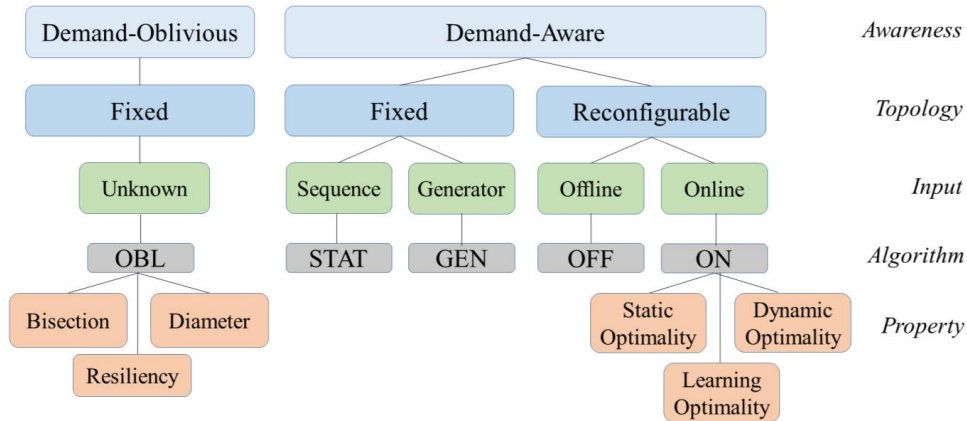


TABLE 1 Selected timeline of reconfigurable data centers

|      |  |
|------|--|
| 2009 | – <i>Flyways</i> [51]: Steerable antennas (narrow beamwidth at 60 GHz [78]) to serve hotspots  |
| 2010 | – <i>Helios</i> [33]/ <i>c-Through</i> [98, 99]: Hybrid switch architecture, maximum matching (Edmond’s algorithm [30]), single-hop reconfigurable connections ( $O(10)ms$ reconfiguration time).<br>– <i>Proteus</i> [21, 89]: $k$ reconfigurable connections per ToR, multi-hop path stitching, multi-hop reconfigurable connections (weighted $b$ -matching [69]), edge-exchanges for connectivity [72], wavelength assignment via edge-coloring [67] on multigraphs)   |
| 2011 | – Extension of <i>Flyways</i> [51] to better handle practical concerns such as stability and interference for 60GHz links, along with greedy heuristics for dynamic link placement [45]  |
| 2012 | – <i>Mirror Mirror on the ceiling</i> [106]: 3D-beamforming (60 Ghz wireless), signals bounce off the ceiling  |
| 2013 | – <i>Mordia</i> [31, 32, 77]: Traffic matrix scheduling, matrix decomposition (Birkhoff-von-Neumann (BvN) [18, 97]), fiber ring structure with wavelengths ( $O(10)\mu s$ reconfiguration time)<br>– <i>SplayNets</i> [6, 76, 82]: Fine-grained and online reconfigurations in the spirit of self-adjusting datastructures (all links are reconfigurable), aiming to strike a balance between short route lengths and reconfiguration costs  |
| 2014 | – <i>REACToR</i> [56]: Buffer burst of packets at end-hosts until circuit provisioned, employs [77]<br>– <i>Firefly</i> [14] Combination of Free Space Optics and Galvo/switchable mirrors (small fan-out)   |
| 2015 | – <i>Solstice</i> [57]: Greedy perfect matching based hybrid scheduling heuristic that outperforms BvN [77]<br>– Designs for optical switches with a reconfiguration latency of $O(10)ns$ [3]  |
| 2016 | – <i>ProjecToR</i> [39]: Distributed Free Space Optics with digital micromirrors (high fan-out) [38] (Stable Matching [26]), goal of (starvation-free) low latency<br>– <i>Eclipse</i> [95, 96]: $(1 - 1/e^{(1-\epsilon)})$ -approximation for throughput in traffic matrix scheduling (single-hop reconfigurable connections, hybrid switch architecture), outperforms heuristics in [57]   |
| 2017 | – <i>DAN</i> [7, 8, 11, 12]: Demand-aware networks based on reconfigurable links only and optimized for a demand snapshot, to minimize average route length and/or minimize load<br>– <i>MegaSwitch</i> [23]: Non-blocking circuits over multiple fiber rings (stacking rings in [77] doesn’t suffice)<br>– <i>Rotornet</i> [63]: Oblivious cyclical reconfiguration w. selector switches [64] (Valiant load balancing [94])<br>– <i>Tale of Two Topologies</i> [105]: Convert locally between Clos [24] topology and random graphs [87, 88] |
| 2018 | – <i>DeepConf</i> [81]/ <i>xWeaver</i> [102]: Machine learning approaches for topology reconfiguration   |
| 2019 | – Complexity classifications for weighted average path lengths in reconfigurable topologies [34, 35, 36]<br>– <i>ReNet</i> [13] and <i>Push-Down-Trees</i> [9] providing statically and dynamically optimal reconfigurations<br>– <i>DisSplayNets</i> [75]: fully decentralized <i>SplayNets</i><br>– <i>Opera</i> [60]: Maintaining expander-based topologies under (oblivious) reconfiguration   |

*Toward Demand-Aware Networking: A Theory for Self-Adjusting Networks*  
C. Avin, S. Schmid  
ACM SIGCOMM Computer Communication Review, October 2018

*Survey of Reconfigurable Data Center Networks: Enablers, Algorithms, Complexity*  
K.-T. Foerster, S. Schmid  
ACM SIGACT News, June 2019