Demand-Responsive Optical Networks

Weekly Meetings

Rafael Lourenço

April 26, 2018

UCDAVIS
UNIVERSITY OF CALIFORNIA

Outline

- Paper Review: *Monia Ghobadi and Ratul Mahajan. "Optical layer failures in a large backbone." In Proceedings of the 2016 Internet Measurement Conference. ACM, 2016.*

- Paper Review: *Rachee Singh, Monia Ghobadi, Klaus-Tycho Foerster, Mark Filer, and Phillipa Gill. "Run, Walk, Crawl: Towards Dynamic Link Capacities." In Proceedings of the 16th ACM Workshop on Hot Topics in Networks. ACM, 2017.*

- Demand-Responsive Optical Networks

UCDAVIS

Optical Layer Failures in a Large Backbone.

Monia Ghobadi (MSR) and Ratul Mahajan (MSR).
Proceedings of the 2016 Internet Measurement Conference. ACM, 2016.

Why Study Optical Links?

- WAN traffic is carried through optical networks
- Few reports on specific optical network availability
- Poor optical signal can corrupt or even lead to silent packet drops
- Main findings:
  1. Availability of different optical segments/channels differ by over three orders of magnitude
  2. Time to Repair (TTR) for segments and channels is similar, although segment faults are more damaging to network capacity
  3. Almost 80% of failures are unidirectional
  4. Q-factor is a good predictor of outages

Optics Beneath IP Layer

- *Segments*: sequence of fiber cables and amplifiers between two OXCs (5 to 2600km)
- OXC use WDM (each segment may have between 10 and 200 channels)
- Aggregation devices may "groom" traffic coming from multiple routers
- OXCs are transparent
- Different modulations yield from 10 to 250 Gbps throughput

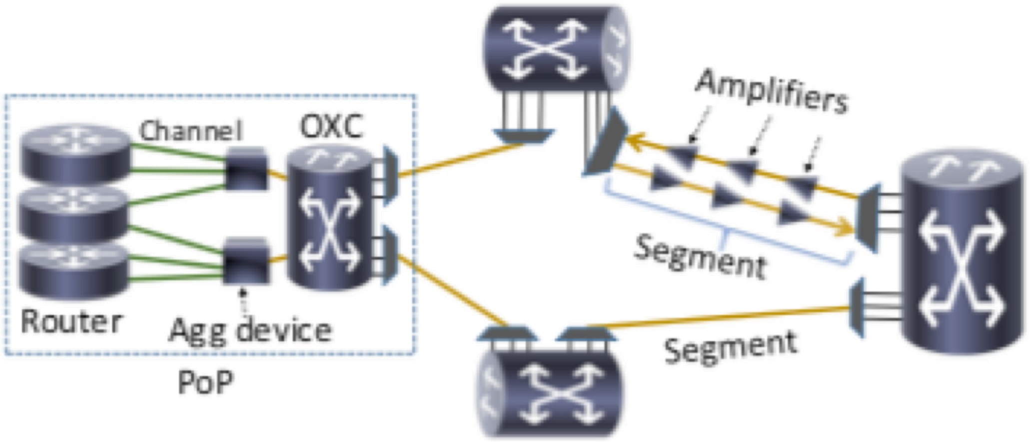

Figure 1: Overview of an IP over OTN wide are network.

UCDAVIS

Dataset

- 14 months of data (February 2015 to April 2016), polled every 15 min
- O(50) OXCs, O(100) segments, O(1000) channels

2015.03.18.21.30.00	14.5	-9.83	-34.75	8.79
2015.03.18.21.45.00	14.59	-9.83	-34.68	9.05 D
2015.03.18.22.00.00	14.71	-9.83	-33.47	8.92
2015.03.18.22.15.00	14.76	-9.84	-32.25	8.67
2015.03.18.22.30.00	14.76	-9.84	-31.26	8.68
2015.03.18.22.45.00	14.78	-9.84	-30.79	8.87
2015.03.18.23.00.00	14.79	-9.84	-31.02	8.98
2015.03.18.23.15.00	14.79	-9.84	-32.28	9.17
2015.03.18.23.30.00	14.62	-9.83	-32.25	8.75

- https://www.microsoft.com/en-us/research/project/microsofts-wide-area-optical-backbone/

UCDAVIS

Analysis of Optical Outages: Q-drop

- Q-drop: when the Q-factor drops but stays above the minimum required for the current modulation (i.e., for 100 Gbps PM-QPSK, at minimum 6.5)

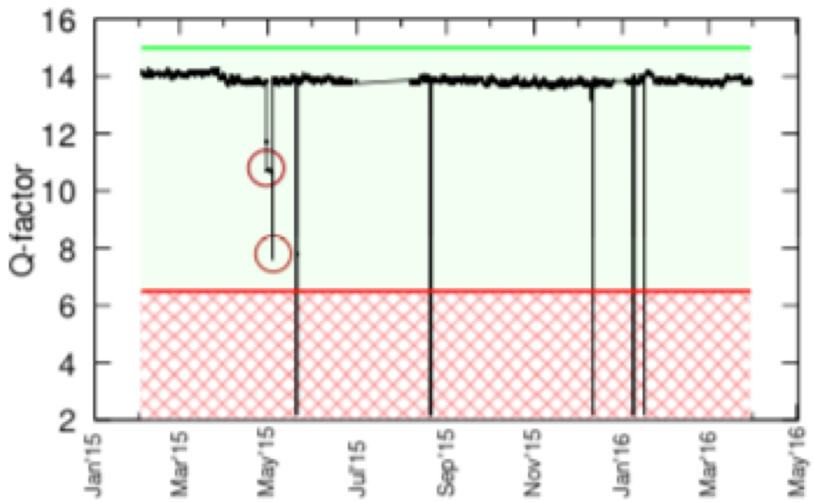

Figure 2: Q-factor variation of an optical channel over time. The graph is divided into healthy (solid green) and unhealthy (hashed red) areas. The circled areas are called Q-drops.

UCDAVIS

Analysis of Optical Outages: Availability and Time to Repair

- CDF of availability of all channels throughout the 14 months
- Different channels can differ by over three order of magnitude
- 40% of outages are full segment, 60% only channels
- Outages: planned or unplanned
- TTRs mostly change due to planned/unplanned (not segment/channel)

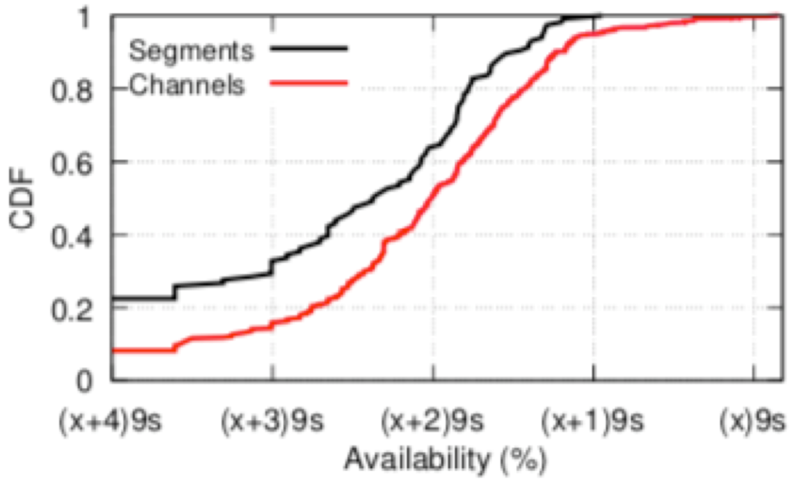

Figure 3: CDF of percentage of time individual channels are available (bottom curve) and CDF of percentage of time the entire segment (all channels traversing the segment) is available (top curve). The x-axis is labeled based on relative number of nines in the availability percentage; for example, (x)9s means the channel has x nines of availability.

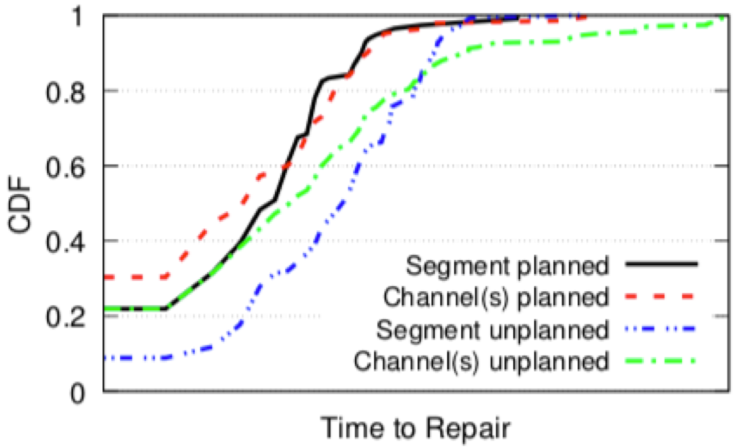

Figure 4: CDF of outage TTR for outages impacting all channels in a segment (black curve) and outages impacting some but not all channels (red curve). The x-axis is in log-scale.

UCDAVIS

Analysis of Optical Outages: Directional Symmetry

- Segment failures in both directions suggest unplanned failure or maintenance of shared component (e.g., fiber cut)
- 68% of unplanned and 58% of planned outages are one-sided
- Less than 30% of outages affect both directions: which suggests that failures of shared components are less frequent sources of outages, and failure of other equipment is more common

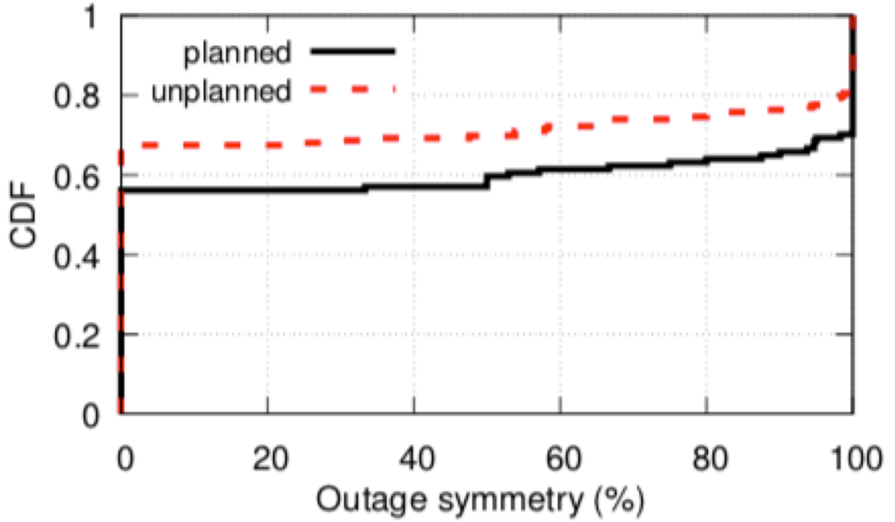

Figure 5: CDF of outage symmetry percentages across all segment outages. A completely symmetric outage is an outage impacting both directions of light completely.

UCDAVIS

Analysis of Optical Outages: Dependence on time of day

- Planned outages occur during maintenance windows
- Unplanned outages have diurnal pattern, peaking at 8 AM and 2 PM (times when people are working, e.g., construction workers accidentally cutting cables, etc.)
- Similar observation w.r.t. day of the week

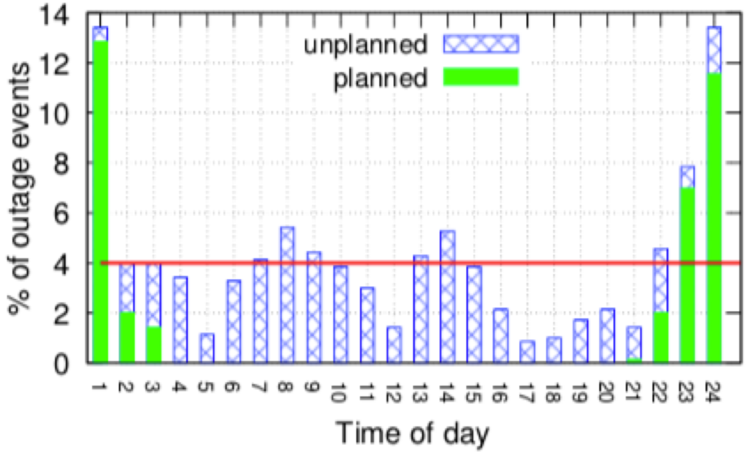

Figure 6: Probability of start of outage at different times of day. The horizontal red line indicates where the bars would have been if all hours had equal probabilities.

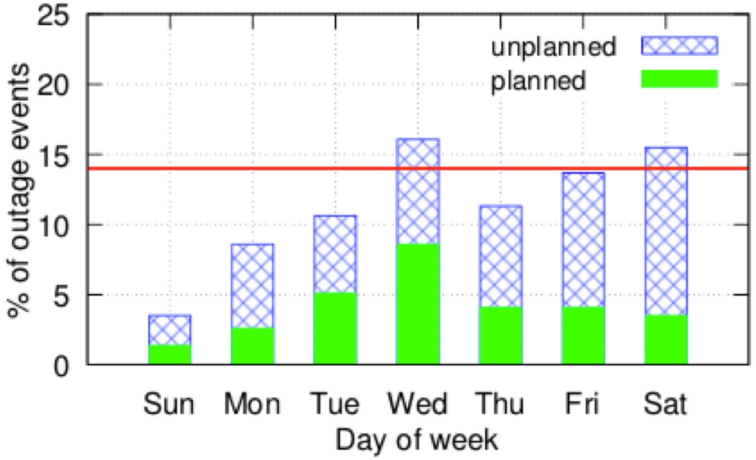

Figure 7: Probability of start of outage on different days of the week. The horizontal red line indicates where the bars would have been if all days had equal probabilities.

UCDAVIS

Outage Prediction

- P(outage) = probability of outage in some interval
- P(outage given Q-drop) = "probability of observing an outage in the channel given a prior Q-drop event within the same window"
- P(outage given outage) = "probability of an outage given the occurrence of a prior outage in the window"
- P(outage given outage) follows P(outage) somewhat closely, reinforcing the common assumption that failures are independent
- Thus, Q-drops are a good predictor of outages

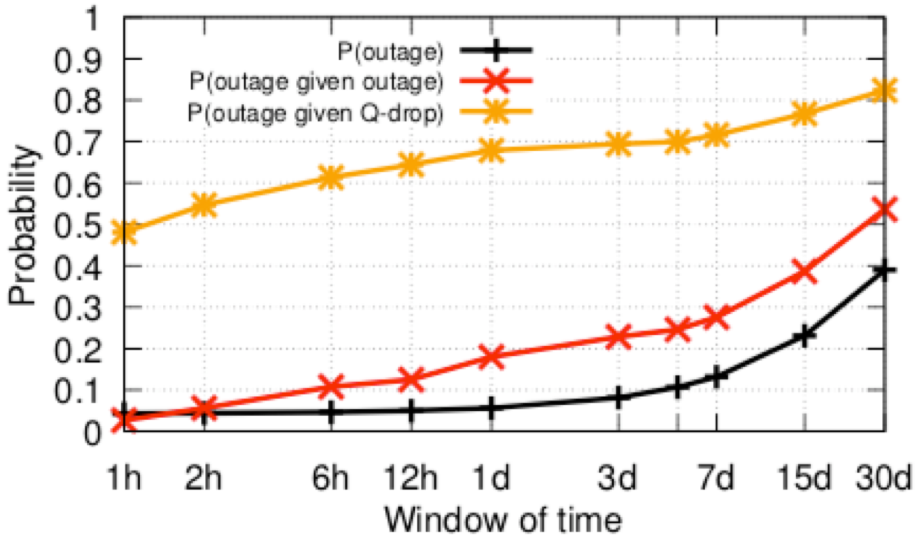

Figure 8: Probability of an outage in a given time window increases significantly after a Q-drop event. It increases only slightly after an outage.

UCDAVIS

Toward Optical-Layer-Aware Traffic Engineering

- Since the availability of optical layer directly affects the availability of upper layers, IP Traffic Engineering should take that into consideration
- Monitor Q-Drops to decide how to route traffic

Run, Walk, Crawl: Towards Dynamic Link Capacities

Rachee Singh (UMass), Monia Ghobadi (MSR), Klaus-Tycho Foerster (Aalborg Univ.), Mark Filer (MSR), Phillipa Gill (Umass).
Proceedings of the 16th ACM Workshop on Hot Topics in Networks. ACM, 2017.

UCDAVIS

The Case for Dynamic WANs

- Research community actively working on strategies to introduce programmability at higher layers of the networking stack:
- Distributed traffic engineering in WANs replaced by programmable centralized controllers (SDN)
- Commodity hardware load-balancers replaced by software
- Switch management APIs replaced by fully programmable stacks
- However, physical layer is generally considered static
- Programmability in lower layers may allow:

➢Higher throughput

➢Better availability

- From the perspective of the L3 Traffic Engineering systems, the extra complexity of a programmable optical layer must be justified (and simplified, if possible)

UCDAVIS

Quantifying the Opportunity: Improving Link Throughput

**High Density Region (HDR): the smallest region of the observed SNR in which 95% or more of the SNR values are concentrated.*

- SNR of 40 lighpaths that traverse the same fiber cable at 100 Gbps over 2.5 years (extended dataset)
- Avg. 12 dB, with occasional big drops: thus, operators use big margins to avoid outages
- 80% of SNR changes are within less than 2dB
- Fig.2b: percentage of lighpaths that could use higher modulations if the margins used were lower (2dB)

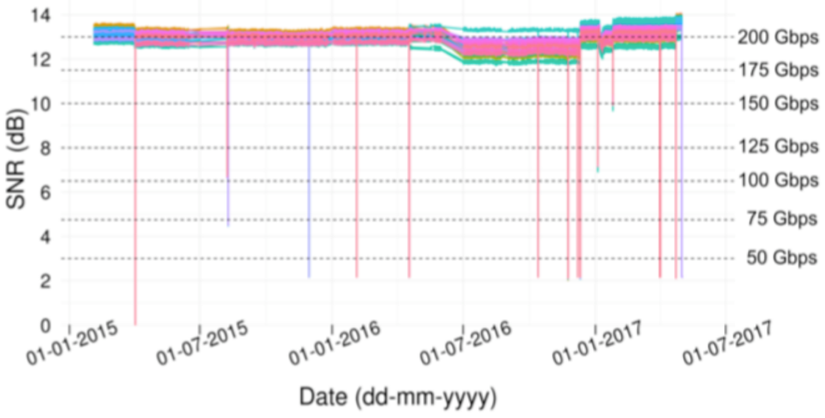

Figure 1: SNR changes of 40 optical wavelengths (i.e., IP links) on a wide area fiber cable. Dotted lines represent the feasible link capacity at and above a particular SNR.

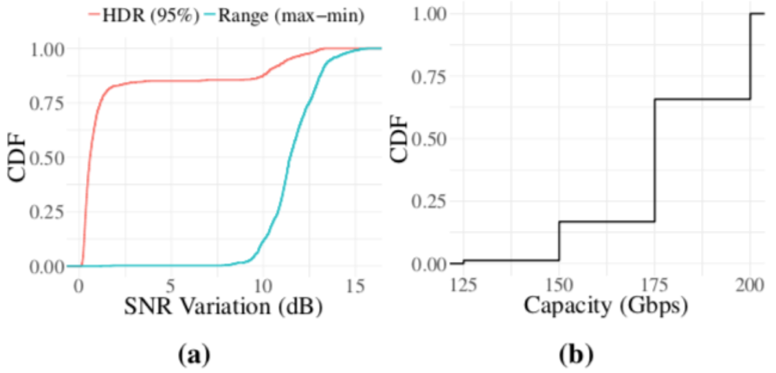

Figure 2: (a) shows the CDF of two metrics of optical SNR variation: the size of high density region (95%) and the range of SNR. Observe that SNR stays within a narrow band of less than 2 dB, 83% of the time. But the range of SNR is much larger, suggesting dramatic but infrequent changes. (b) shows the capacities of WAN links if they were to be utilized according to their signal quality. Over 80% of links can gain 75 Gbps or more capacity over their existing static configuration of 100 Gbps.

UCDAVIS

Quantifying the Opportunity: Improving Link Throughput – Downside

- On the other hand, increasing lightpath modulation schemes increases the number of times the SNR falls bellow the minimum required for that scheme: i.e., an outage occurs
- Fig 3a: a good lightpath that generally has high SNR and would allow all different modulations: ***as throughput increases, so does number of failures***
- Fig 3b: all lighpaths at highest modulation possible, duration of failures per lighpath capacity: ***failures last for several hours***

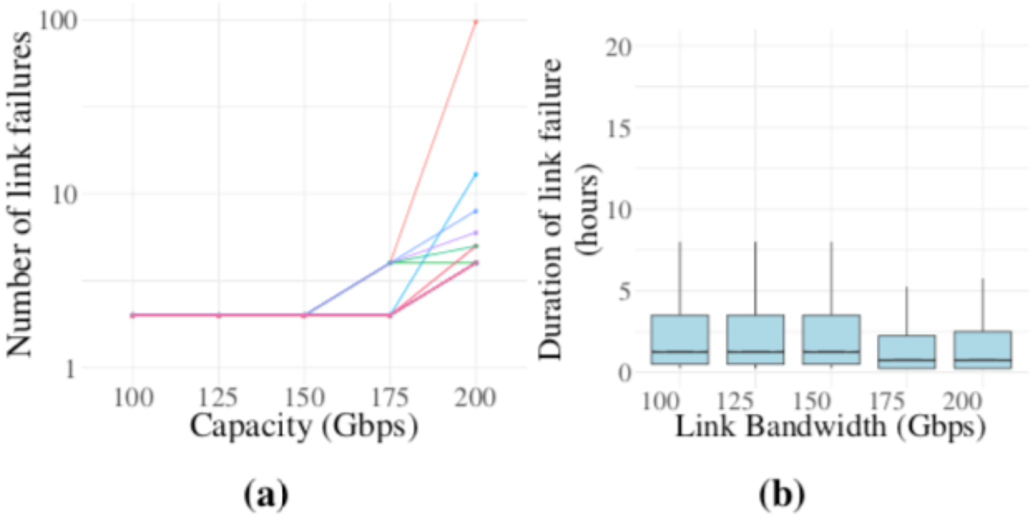

Figure 3: (a) the number of link failures for 40 different links (one color per link) for a given capacity. While increasing capacity up to 175 Gbps does not increase link failure events, achieving 200 Gbps capacity comes at the cost of increased link failures. (b) shows the duration of link failures if WAN links operate at a given capacity. Link failures for all capacities last for several hours on average.

Quantifying the Opportunity: Improving Link Availability

- Today, if SNR of lighpath falls bellow threshold, the link is *down*
- *However,* if a lower modulation was used when the SNR went down, service could still be provided (at lower rate)
- Authors classified seven months of maintenance tickets and classified unplanned failures into:
- Human: unplanned outages due to scheduled maintenance
- Fiber cuts: accidental cuts
- Hardware failure: equipment failure (amplifier, transponders, etc.)
- Undocumented
- Fiber cuts are only 10% of events
- Lowest SNR during failure is above 3dB in 25% of the cases (3dB SNR is sufficient for 50 Gbps)

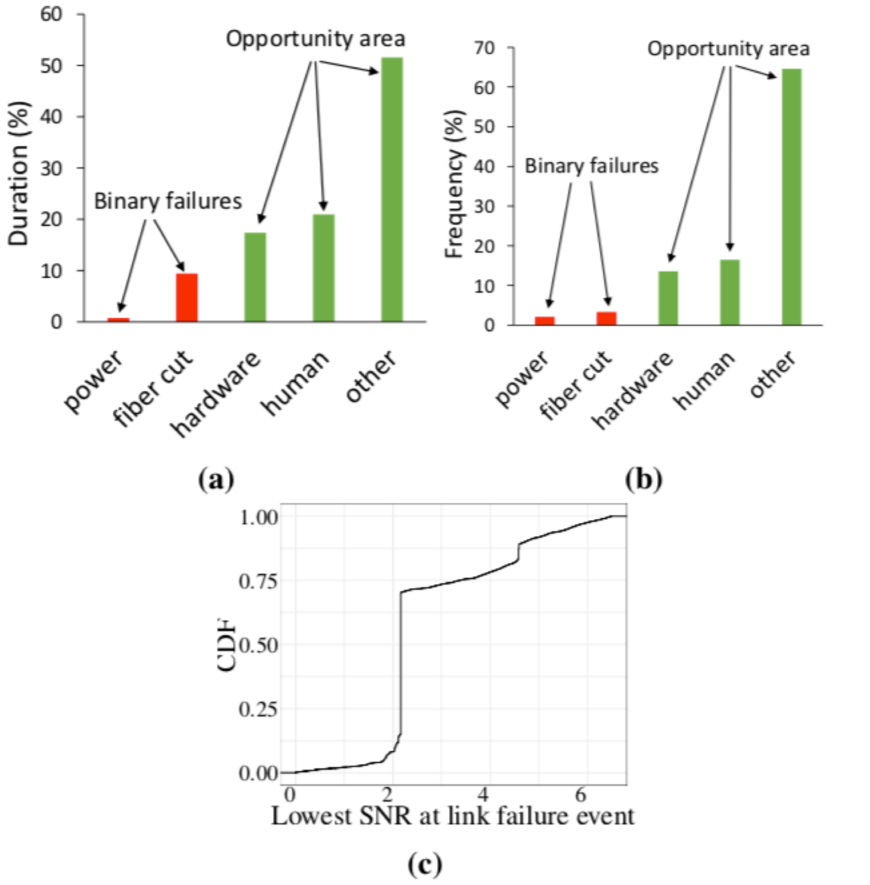

Figure 4: Categorization of major failure root causes in terms of duration of events (a) and frequency of events (b). Contrary to common belief, fiber cuts are not the major root cause of failures in WANs. Unplanned events during planned maintenance or hardware failures are more probable; they contribute to more outage duration than fiber cuts. Figure (c) shows the distribution of the SNR values at link failure events.

Deployment Constraints: Towards Hitless Capacity Change & Separating Layer 1 from TE

- Authors used an Acacia testbed to check the time it takes to change modulation formats
- Current technology avg: 68 s
- Mostly because laser is turned off and then on again
- Without turning off the laser: 35 ms
- Introducing complexity in already bulky Traffic Engineering systems is not desired by operators: the authors propose a graph abstraction to represent dynamic link capacities

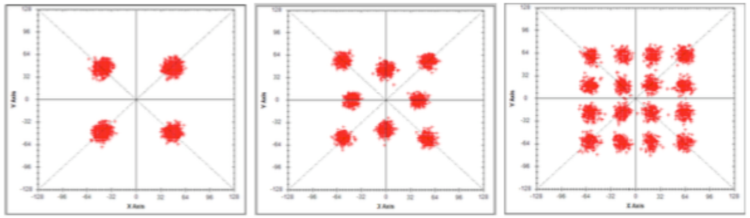

(a) 100 Gbps (b) 150 Gbps (c) 200 Gbps

Figure 5: Small testbed evaluation of dynamic capacity adjustment of two optical links.

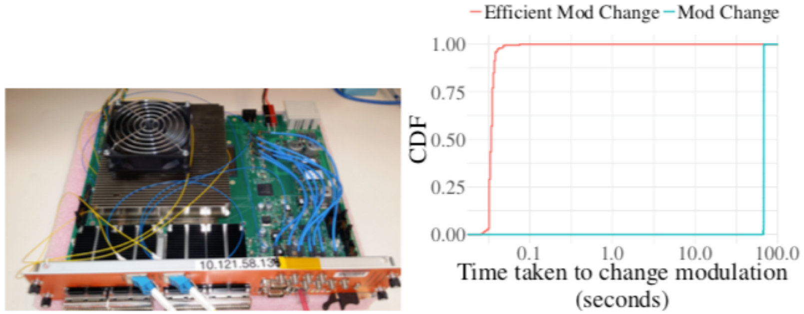

(a) Evaluation board (b)

Figure 6: (a) shows the testbed we build for evaluating the feasibility of capacity variable links. (b) is a CDF of the time taken to change modulation (capacity) of a fiber link in our testbed. Link capacity changes take 68 seconds on an average. But we demonstrate ways to change the modulation efficiently such that it takes only 35 milliseconds on average.

UCDAVIS

Graph Abstraction

- The authors propose a graph abstraction of the network topology that allows for minimal changes in current Traffic Engineering systems

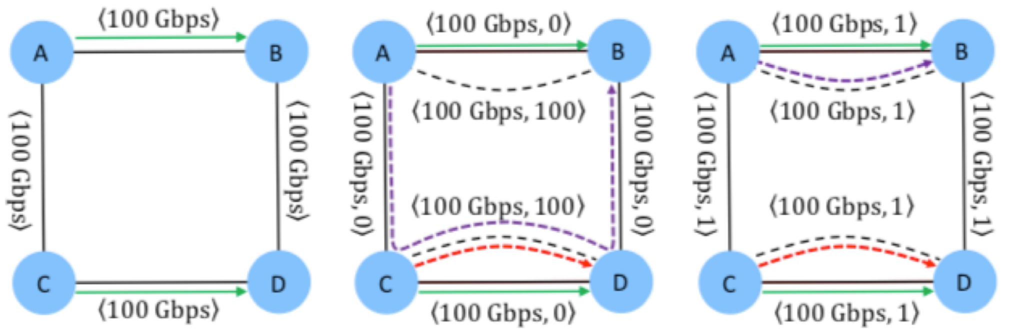

(a) Current flows **(b) Few increases** **(c) Short paths**

Figure 7: Intuition for graph abstraction: In 7a, we have the current network, with 100 Gbps flows (in green) between A, B and C, D, respectively. In 7b, we annotate each link with $< capacity, cost >$, where the cost reflects the penalty of using the (updated) link. When both A and C want to send an additional 25 Gbps flow (in dotted red and purple), updating one link's capacity suffices. If short paths are favored, all costs are set equal, as in 7c.

Demand-Responsive Optical Networks

UCDAVIS

OSNR margins

- Adding to what the authors of the previous paper propose, OSNR margins are used because:
- They help avoid unnecessary disruptions of optical lightpaths
- They allow for new lightpaths to be established in the future (long term), since:

1. Placing new lightpaths in adjacent lambdas might create crosstalk
2. Equipment and fibers may age

- Nevertheless, the idea of using tighter OSNR margins is interesting (and has also been investigated before, as in last OFC: "Capacity Enhancement in Optical Networks using Margin Extraction, Mohammad Sheikh Zefreh (Ciena, Canada). *In this paper, the value of margin extraction in optical networks is investigated. Simulation results for a sample network show that up to 64% of multi-rate transponders can run with higher rates using extracted margins."*)

Tighter Margins = Higher chance of disruption

- Dataset provided and the results of the previous papers clearly show that tighter margins lead to higher occurrence of outages
- Authors suggest that if fast modulation changes are available, the network would still benefit from tighter margins:
- Although the authors do not cover the extensive literature about hitless modulation changes in BVTs, if their approach is in fact state-of-the-art, a modulation change would take 35ms
- Since their data has 15 min granularity, we have little knowledge of instantaneous variations of Q-factor.

➢With a 5x9s requirement, service can be unavailable at most 315 seconds per year on average. If the modulation is changed 25 times per day (and absolutely no failures occur), this SLA is infringed… Thus, contrary to what the authors propose, always using a very tight margin seems very dangerous.

UCDAVIS

Q-factor Outage Prediction

- The fact that the Q-factor is good predictor of outages is very interesting…
- In fact, is there a ML model that may use the Q-factor information to predict with higher accuracy possible failures (several works focus on this, Prof. Massimo's OFC paper, Prof. Velasco's latest papers on JOCN presented in the last few weeks…)

➢We may consider using tighter margins only if the Q-factor-based prediction is that no failure will happen in the next few hours (or choosing to tighten margins only for those lighpaths with smallest risk of outages)

UCDAVIS

Our idea

- Lightpaths tend to use large OSNR margins, which would allow for higher modulation schemes and, thus, more throughput
- On the other hand, tighter margins incur in more unavailable time (even by simply switching the modulation to lower schemes)
- Not only that, but modulation changes should not be too frequent

➢Thus, during regular hours => use large margins, lower modulations (higher availability)

➢**During Resource Crunch** => use tighter margins, higher modulations (with possibly higher risk of outages)

➢If we have a good way to predict the risk of future outages based on the current Q-factor (and other information) => even better

UCDAVIS

Problem Statement

- Given: Optical network topology with some lightpaths with large OSNR margins, L3 topology, increase in offered load (configuring a Resource Crunch)
- Output: What lighpaths to increase modulation (and what modulation to use)
- Goal: Increase network capacity during Resource Crunch, but minimizing the chances of outage occurring due to tight OSNR margins

UCDAVIS