

Optical Layer Failures in a Large Backbone

Monia Ghobadi
mgh@microsoft.com
Microsoft Research

Ratul Mahajan
ratul@microsoft.com
Microsoft Research

ABSTRACT

We analyze optical layer outages in a large backbone, using data for over a year from thousands of optical channels carrying live IP layer traffic. Our analysis uncovers several findings that can help improve network management and routing. For instance, we find that optical links have a wide range of availabilities, which questions the common assumption in fault-tolerant routing designs that all links have equal failure probabilities. We also find that by monitoring changes in optical signal quality (not visible at IP layer), we can better predict (probabilistically) future outages. Our results suggest that backbone traffic engineering strategies should consider current and past optical layer performance and route computation should be based on the outage-risk profile of the underlying optical links.

Keywords

Wide-area backbone network; Optical layer; Q-factor; Availability; Outage

1. WHY STUDY OPTICAL LINKS?

Wide-area backbone networks (WANs) of Internet service providers and cloud providers are the workhorses of Internet traffic delivery. Providers spend millions of dollars building access points around the world and interconnecting them through optical links. Improving the availability and efficiency of the WAN is central to their ability to provide services in a reliable, cost-effective manner. Consequently, there has been significant research into measuring and characterizing various aspects of WANs, including topology, routing, traffic, and reliability [4, 9–11, 18].

However, prior studies tend to focus exclusively on the IP layer, and little is publicly known about the characteristics of the optical layer which forms the physical transmission

medium of WANs. There are studies that focus on dispersion and modulation [2, 5, 8, 14, 24, 25], but few reports exist on basic characteristics of operational optical backbones such as their availability and failures. While certain factors (e.g., expected signal quality) may be measured in the lab, only an “in the wild” characterization can shed light on the behavior and the combined impact of many others (e.g., equipment failures, maintenance events).

Studying optical layer characteristics of backbone networks is not simply a matter of curiosity. The health of this layer ultimately determines the network’s effectiveness at carrying traffic. For instance, poor optical signal quality can lead to corruption and even silent packet drops [22]. In addition, optical layer outages bring down IP routes and disrupt traffic. Thus, better awareness of this layer can improve WAN route selection and traffic engineering.

This paper presents preliminary results from our analysis of availability, outages, and signal quality of optical segments that form the backbone of a large cloud provider. An optical *segment* corresponds to a fiber and carries multiple *channels*, where each channel corresponds to a different wavelength and router port. IP layer links map to optical channels. Our study covers data from hundreds of optical segments and thousands of optical channels for over a year.

We uncover several notable characteristics of this backbone. First, the availability (i.e., uptime) of different optical segments and channels differs by over three orders of magnitude. Second, the distribution of time to repair of planned outages is similar for both optical segments and channels, even though a segment outage tends to represent an order of magnitude greater impairment in network capacity. Third, almost four in five optical segment outages are unidirectional; i.e., one direction is functional while the other is down. Finally, outages can be predicted (probabilistically) based on sudden drops in optical signal quality (which is not visible at the IP layer). There is a 50% chance of an outage within an hour of a drop event and a 70% chance of an outage within one day.

Our findings motivate smarter IP layer management and routing, one that is aware of optical layer characteristics. For instance, a common assumption in fault-tolerant routing schemes is that each IP layer link is equally likely to fail [18], but we find such links can have drastically different failure probabilities; moreover, many IP layer links

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IMC 2016, November 14–16, 2016, Santa Monica, CA, USA

© 2016 ACM. ISBN 978-1-4503-4526-2/16/11...\$15.00

DOI: <http://dx.doi.org/10.1145/2987443.2987483>




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

traversing the same segment often fail together. Because more accurate failure probability for IP layer links can be computed by monitoring optical signal quality, traffic engineering frameworks such as SWAN [10] and B4 [11] should monitor optical signal quality and compute IP layer routes based on physical link-failure probabilities.

Our results are preliminary and may not hold for other optical backbones. That being said, the management practices of the backbone we study are industry-standard, and it shares optical fiber conduits with many other backbones. In the future, we plan to expand our analysis to more backbone networks and fully develop our routing approach.

2. OPTICS BENEATH THE IP LAYER

This section provides a brief background of optical backbone networks with a focus on aspects relevant to our work. While there are multiple optical network architectures [3, 13, 16, 26], we focus on the most common one: IP-over-OTN (optical transport network). Figure 1 shows a simplified view of such a network.

The WAN is composed of optical links (*segments*) and optical cross-connects (OXCs) in different PoPs (points-of-presence). OXCs use wavelength division multiplexing (WDM) to combine multiple wavelengths onto a single fiber and steer different wavelength combinations to different neighbors.¹ Each wavelength is called an optical *channel* and is connected to a router port using transponders that convert electrical signals to and from modulated light. The maximum number of channels that a fiber can carry depends on the WDM technology. The actual number of channels carried over a deployed fiber depends on the provisioned hardware (e.g., transponders and router ports).

IP routers connect to aggregation devices that multiplex channels and isolate routers from the optical WAN, allowing routers to evolve independently. For instance, when router ports have lower capacity than an optical channel, the aggregation device can time-multiplex (or “groom”) multiple router ports onto the channel. Aggregation devices generate and terminate optical signals carried over the WAN. The connectivity between routers and aggregation devices may be optical as well, but the signal coming from the router

¹Some vendors market devices with electronic backplanes as “OXCs”. For our purposes, OXCs are devices with purely optical multiplexing and steering.

is decoded (i.e., converted to electrical domain) by aggregation devices. The aggregation devices will then generate a fresh optical signal, possibly after combining information from multiple router ports.

The data-rate of an optical channel depends on its modulation. Earlier systems used simple modulation, such as On-Off Keying (OOK) with 10 Gbps per channel. Today, 100 Gbps per channel is widely available through more efficient modulation such as Polarization-Multiplexed Quadrature Phase-Shift Keying (PM-QPSK). Our backbone has several modulations, ranging from 10 to 250 Gbps. We focus on 100 Gbps channels, which dominate deployment.

The top right link in Figure 1 zooms into an example of inter-OXC connectivity with two optical segments, one for each direction of traffic. To enable long reach, segments have amplifiers that amplify the signal in the optical domain (i.e., without decoding). The two (directional) segments share a fiber conduit (which is a raceway for enclosing optical fibers), but their fiber and amplifiers are separate.

3. DATASET

Our analysis is based on 14 months of data, from February 2015 to April 2016, taken from Microsoft’s optical backbone in North America. This backbone has $O(50)$ OXCs, $O(100)$ WAN segments, and $O(1000)$ optical channels. For confidentiality, we do not report the exact numbers.

We poll the aggregation devices for the optical signal Quality factor (*Q-factor*) for all 100Gbps channels. The Q-factor measures the quality of an analog signal in terms of its signal-to-noise ratio (SNR). It takes into account physical impairments to the signal (e.g., noise, chromatic dispersion) which can degrade the signal and ultimately cause bit errors. It is defined as $\frac{|\mu_1 - \mu_0|}{\sigma_1 + \sigma_0}$, where μ_1 and μ_0 correspond to the power levels of the transmitted ‘1’s and ‘0’s, and σ_1 and σ_0 correspond to the standard deviation of the noise on ‘1’s and ‘0’s, respectively. Higher values of the Q-factor imply better SNR and, thus, lower probability of bit errors.

The devices sample the Q-factor once per second, and at 15-minute intervals, they report the minimum, maximum and average across all samples in the last interval. In this work, we analyze the average values. Because our data have a 15-minute granularity, we do not capture events lasting much less than 15 minutes. The minimum and maximum values are same as the average for the vast majority of time.

The segments in our data range from 5 to 2600 Km in length. Each segment has multiple channels, and each channel traverses a fixed segment. There is no dynamic re-routing of wavelengths in our backbone. Segments carry between 10 and 200 channels.

4. ANALYSIS OF OPTICAL OUTAGES

Figure 2 illustrates the Q-factor of a sample 100 Gbps channel over time. In this channel, the Q-factor is mostly stable with mean 13.8 and variance 0.3. Occasionally, the Q-factor drops to smaller values, indicating a complete loss of light or low SNR on the channel. Low SNR can lead to

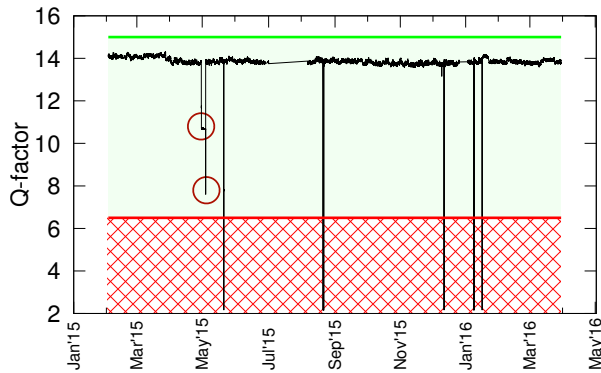


Figure 2: Q-factor variation of an optical channel over time. The graph is divided into healthy (solid green) and unhealthy (hashed red) areas. The circled areas are called *Q-drops*.

a complete inability to decode incoming transmissions or to an erroneous decoding process (i.e., packet corruption).

The Q-factor threshold for error-free signal recovery depends on several factors, including the optical gear, modulation format, and the forward error correction (FEC) [21] logic. In our backbone, the Q-factor threshold for 100Gbps QPSK channels is 6.5; below this threshold, the signal is unrecoverable (hashed red area), and above this region, it is recoverable (solid green area). When an optical channel is in the red area, it is considered “unavailable”, and the corresponding IP layer link is down.

The two circled areas in the figure mark when the Q-factor has dropped from its stable value but is still above the recovery threshold. In Section 5, we will demonstrate the power of such Q-factor drop events, which we call *Q-drops* from here on, to predict outages.

4.1 Availability and Time to Repair

We analyze the Q-factor time series for thousands of 100 Gbps channels and compute the percentage of time each channel is available (i.e., values above 6.5). The red (lower) curve in Figure 3 plots the CDF of all channels’ availability percentage. Due to confidentiality concerns, we do not report the exact availability numbers; instead, we report the relative number of nines in the availability percentage. For example, (x)9s means the channel is available with x number of nines; 99.9% has three nines.

For our purposes, the exact value of x is not as important as the difference in availability across channels. The graph shows that the availability of different channels can differ by over three orders of magnitude. We observe similar distributions for time between failures as well. Thus, unlike common practice, IP layer route computation systems should not assume equal failure probability for different IP links.

Some outages, like fiber cuts and amplifier failures, impact an entire segment, causing all channels in the segment to become unavailable. Other outages, like hardware failures, impact some channels but not all. A segment-level out-




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.

age happens when all channels traversing that segment are unavailable.

Of the hundreds of outage events in our data, 40% are segment-level, with the remaining 60% impacting only a subset of channels traversing the segment. This suggests that optical-link components in the middle of the link (e.g., amplifiers and fibers) fail with roughly the same probability as the components on the edges, though the impact of the former is higher because multiple Tbps of capacity are lost when a segment fails.

The black (upper) curve in Figure 3 shows the CDF of segment availability across all segments. Segments have higher availability in general; this is expected, as segment outages imply channel outages but not the other way around. As in channel availability, however, there is a wide difference in the availability of different segments.

Outages can be caused by planned maintenance (e.g., a line card replacement) or unplanned events (e.g., fiber cut, hardware failure, power failure). The impact of an unplanned outage on traffic is more drastic than that of a planned one. The Q-factor data do not distinguish between planned or unplanned outages, but based on our conversations with operations personnel and our investigation of maintenance records, we infer that planned maintenance normally starts at 10 PM Pacific time, regardless of location. In our analysis below, to study planned and unplanned separately, we assume that the planned outages are those that start between 10 PM and 12 AM Pacific time.

Time to repair (TTR) is an important factor in maintaining highly-available systems. Figure 4 plots the CDF of TTR for planned and unplanned channel and segment outages. The x -axis is in log-scale. For confidentiality, we do not report the values on the x -axis. As the figure shows, the TTR distributions for planned segment and channel outages are roughly similar, even though segment outages represent a larger im-




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.

pairment and loss in network capacity and repair staff must often travel to remote sites to repair them. On the other hand, the TTR distributions for unplanned segment and channel outages are different. Except for the tail, which may be explained by the difficulty to acquiring the failed component in a timely manner, channel outages tend to be repaired faster.

4.2 Directional symmetry

We now ask whether segment outages are symmetric, i.e., both directions fail simultaneously. While separate fiber cores and amplifiers are used for each direction, fiber conduits and electrical power sources are shared. Hence, a segment outage in both directions suggests the unplanned failure or maintenance of a shared component (e.g. a fiber cut that damages the conduit), whereas a single-sided outage suggests hardware issues that are unique to each direction.

We measure the percentage of time overlap between the two directions of light in all segment-level outages. Figure 5 shows that 68% of the unplanned and 58% of the planned outages are completely one-sided, i.e., with 0% overlap between two directions. Less than 30% of outages impact both directions of a segment, i.e., 100% overlap. This result suggests that failures of shared components are a less frequent source of segment level outages, and failure of other equipment is a more frequent cause. Our analysis of trouble tickets from the backbone supports this conclusion.

4.3 Dependence on time of day

We next ask if outages are more likely to happen at a specific time of day or day of the week. Figure 6 shows the histogram of outages binned by time of day. The time reference is with respect to the location of the receiving end of the segment and the start of the outage. The horizontal line denotes the y-value if the outages were equally likely at all hours of the day.

As the figure indicates, most planned outages happen between 11 PM and 1 AM local time. As mentioned earlier, this is because of the scheduling of many maintenance activ-




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




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.

ities at 10 PM Pacific time. In contrast, unplanned outages have a diurnal pattern. They peak at 8 AM and 2 PM and dip around noon and 5 PM. Our trouble ticket logs show many outages are accidental damage of maintenance work that is unrelated to the backbone (e.g., road construction).

The day of week pattern in Figure 7 suggests unplanned outages are likely to happen any day of the week, with Sundays and Mondays the least probable and Wednesdays and Saturdays the most probable.

5. OUTAGE PREDICTION

In this section, we reveal the power of Q-drop events to predict channel-level outages. For a 100 Gbps QPSK channel, a Q-drop event is when the Q-factor value drops from its stable value but remains above the threshold of 6.5. In Figure 2, these events are marked with circles. During such events, the channel is still available, and the degradation is not visible at the IP layer.

In this analysis, for each channel, we first compute the probability of an outage within a window of time and call




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.

it $P(\text{outage})$. We then define $P(\text{outage given Q-drop})$ as the probability of observing an outage in the channel given a prior Q-drop event within the same window (in the same channel). Figure 8 plots the average probabilities across all channels as a function of window size, from 1 hour to 30 days. In this analysis, we do not distinguish between planned and unplanned outages. Expectedly, $P(\text{outage})$ increases as window size increases; the larger the window of time, the greater the likelihood of an outage happening within that window. In fact, for a window of 7 days, the probability of an outage occurrence is 0.12; on average there is a 12% chance of an outage happening each week. However, there is a significant jump in outage probability if there has been a Q-drop event in the past. For example, for a window of 7 days, the probability of outage occurrence increases to 70% if there has been a Q-drop event within that week. This means Q-drop events are strong predictors of future outages.

This result suggests integrating two new features for IP layer network management. First, we should monitor Q-drop events and raise an alarm when they occur. Investigating such events and addressing their root causes may prevent some outages. Second, once a Q-drop happens, until the event has been diagnosed and fixed, the failure probability of the IP layer link should be updated and appropriate actions taken. For example, high priority traffic should be moved away from the impacted link.

We also find that, compared to Q-drop events, past outages are less predictive of future outages. To arrive at this finding, we compute the probability of an outage given the occurrence of a prior outage in the window and plot the average among all channels labeled as $P(\text{outage given outage})$, as shown in Figure 8. The fact that $P(\text{outage given outage})$ closely follows $P(\text{outage})$ curve means outages are memoryless, and the occurrence of an outage does not necessarily mean an increased probability of another outage for the same channel.

This result confirms the independence of outages assumption in previous traffic engineering work using Forward Fault Correction [18]. However, FFC assumes constant and equal




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.

outage probability over time and across links, while we show that the Q-drops indicate the outage probabilities change over time and differ across links (Section 4). In the next section we explain the implications of our analysis on IP layer route computation and traffic engineering.

6. TOWARD OPTICAL LAYER AWARE TRAFFIC ENGINEERING

Our findings suggest several ways IP layer network management and routing can be improved by factoring in optical layer characteristics. In this section, we focus on a particular case we are currently investigating in detail: improving traffic engineering so that traffic is less likely to traverse the paths more likely to fail. This change would improve application performance, as IP layer failures result in heavy packet losses, at least in the time between the occurrence and the detection of failures.

Improving the efficiency of routing in the WAN has already been the focus of several studies. Recently, SWAN [10] and B4 [11] introduced traffic engineering (TE) systems with software control planes that perform periodic and centralized route computation for different classes of service in the WAN. However, neither work considers the availability of the physical layer.

The results in Sections 4 and 5 have direct implications for TE decisions. First, the wide range of availabilities and time between failures among optical links implies the common assumption in fault-tolerant routing designs that all links have equal failure probabilities [18] does not hold, and new failure models based on measured optical link outages should be adopted. Second, the predictive power of Q-drop events on outages implies that route selection can greatly benefit from periodically updating risk estimates of optical links and making it less likely for traffic to traverse high-risk paths.

Hence, we suggest using centralized TE controllers to periodically compute routes based on dynamic risk profiles of optical links, along with IP layer metrics. The controller

monitors the optical layer Q-factor of all links in real-time, and computes the outage probability for each link based on its baseline outage probability and the occurrence of Q-drop events in the past month.

Specifically, if there were no Q-drop events in the past month, the outage probability is the baseline outage probability computed using past outages. Otherwise, the outage probability is based on how far in the past the Q-drop event occurred. This probability is computed using an analysis similar to that shown in Figure 8 but is done for individual channels (not aggregated across channels). We leave details of a route computation formulation that maximizes allocated bandwidth over most available paths to future work.

7. RELATED WORK

Our work builds on three lines of related work.

Optical layer characterization: Several studies have characterized dispersion in deployed fiber. For example, Feuerstein [5] considered polarization mode dispersion (PMD), chromatic dispersion (CD), and dispersion slope in 70 spans (i.e., portions of segments between two amplifiers); Woodward et al. [24, 25] studied the state-of-polarization (SOP) and real-time PMD on six fiber paths; Bulow et al. [2] examined PMD fluctuations of an installed fiber; finally, Karlsson et al. [14] compared the temporal drift properties of two similar fibers. In contrast to these previous studies, our work studies a much larger deployment and focuses on signal quality, availability, and outages.

Ghobadi et al. [8] reported on a three-month study of Q-factor data from Microsoft’s optical backbone. They computed the SNR and evaluated whether fiber segments can support higher-order modulations to increase the amount of traffic carried. In this study, we use 14 months of data to report on channel and segment outages.

Filer et al. [6] studied the deployed optical infrastructure of Microsoft’s backbone and discussed the benefits of optical elasticity. They expressed a long-term goal of unifying the optical control plane with routers under a single Software Defined Network controller. They recognized YANG [1] and SNMP as potential starting points for a standard data model and control interface between the optical layer and the extend WAN traffic controller. In this work, we do not focus on the control interface, rather we explore how the optical data can be used by the WAN controller.

Ji et al. [12] evaluated the outage probability of optical networks due to dispersion variations from the seasonal and regional temperature variations; Li et al. [17] modeled the probability of an optical link outage due to optical impairments such as seasonal soil temperature variations. Unlike our work, which is based on measurements in the wild, their research was based on laboratory studies.

Durairajan et al. [4] constructed a map of US optical backbone infrastructure using public data. They observed a correspondence between long-haul fiber-optic, roadway, and railway infrastructures and used traceroute to identify parts of the backbone with high levels of infrastructure sharing and high volumes of traffic. Our work is complementary; we

have access to the internal topology of our optical backbone, so our focus is on measuring and predicting outages.

Inferring optical characteristics using IP layer data:

Markopoulou et al. [20] analyzed IS-IS routing updates from Sprint’s backbone and characterized failures that affected IP connectivity. They inferred the causes of failures based on shared patterns in the data and classified failures into maintenance, router-related and optical layer, and they found that 11% of failures can be related to the optical layer.

Marian et al. and Freedman et al. [7, 19] focused on IP and TCP layer measurements, for example, packet loss and packet inter-arrival times on fiber optics spans. In contrast, we capture outages in the optical layer directly by probing optical devices. We leave correlation of IP layer performance and optical outages to future work.

Kompella et al. [15] point out that IP routing and the underlying optical fiber plan are typically described by disparate data models and propose SCORE, a system that automatically identifies root causes across layers using only IP layer data. Our work is complementary and can enable such systems to root cause optical issues more accurately.

Risk-aware network management:

Govindan et al. [9] studied 100 failure events at Google WAN and data center networks, offering insights on why maintaining high levels of availability for content providers is challenging. Although they did not isolate optical layer failures, they have built a set of design principles for high availability, such as avoiding and preventing failures. Our work agrees with this principle, and we provide a concrete proposal to improve IP layer availability using optical layer signals.

Finally, Vidalenc et al. [23] built analytical models for the dynamic estimation of failure risks and proposed accounting for risk when determining routes in optical backbones so that service level agreement violations are minimized. We argue for measurement-based failure probabilities and the use of live data to update failure risks.

8. CONCLUSIONS

In this paper, we take a preliminary look at optical impairments in a large backbone by studying the quality of the signal of thousands of channels for over a year. We find optical channels have a wide range of availabilities, and failures can be predicted using drops in signal quality that are not visible at the IP layer. Our findings question the common assumption in fault-tolerant routing designs that all links are equally likely to fail. They also suggest that traffic engineering systems should factor in the dynamic failure probabilities of different links using optical layer data.

9. ACKNOWLEDGMENTS

We thank our colleagues Victor Bahl, Jamie Gaudette, Jeffrey Cox and Buddy Klinkers for their support and sharing their expertise. We also thank our shepherd Paul Barford and the IMC reviewers for feedback on the paper.

10. REFERENCES

- [1] M. Bjorklund. YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF). RFC 6020, Oct. 2010.
- [2] H. Bulow, W. Baumert, H. Schmuck, F. Mohr, T. Schulz, F. Kupperts, and W. Weiershausen. Measurement of the maximum speed of PMD fluctuation in installed field fiber. *Optical Fiber Communication Conference and the International Conference on Integrated Optics and Optical Fiber Communication. OFC/IOOC '99*, 2:83–85, Feb 1999.
- [3] A. L. Chiu, G. Choudhury, G. Clapp, R. Doverspike, J. W. Gannett, J. G. Klinecicz, G. Li, R. A. Skoog, J. Strand, A. V. Lehmen, and D. Xu. Network design and architectures for highly dynamic next-generation IP-Over-Optical long distance networks. *Journal of Lightwave Technology*, 27(12):1878–1890, June 2009.
- [4] R. Durairajan, P. Barford, J. Sommers, and W. Willinger. Intertubes: A study of the US long-haul fiber-optic infrastructure. *SIGCOMM'15*, 45(4):565–578, Aug. 2015.
- [5] R. J. Feuerstein. Field measurements of deployed fiber. *Optical Fiber Communication Conference and Exposition and The National Fiber Optic Engineers Conference*, page NThC4, 2005.
- [6] M. Filer, J. Gaudette, M. Ghobadi, R. Mahajan, T. Issenhuth, B. Klinkers, and J. Cox. Elastic optical networking in the microsoft cloud. *Journal of Optical Communications and Networking*, 8(7):A45–A54, Jul 2016.
- [7] D. A. Freedman, T. Marian, J. H. Lee, K. Birman, H. Weatherspoon, and C. Xu. Exact temporal characterization of 10 Gbps optical wide-area network. *IMC'10*, pages 342–355, 2010.
- [8] M. Ghobadi, J. Gaudette, R. Mahajan, A. Phanishayee, B. Klinkers, and D. Kilper. Evaluation of elastic modulation gains in Microsoft's optical backbone in North America. *Optical Fiber Communication Conference*, page M2J.2, 2016.
- [9] R. Govindan, I. Minei, M. Kallahalla, B. Koley, and A. Vahdat. Evolve or die: High-availability design principles drawn from Google's network infrastructure. *SIGCOMM'16*, pages 58–72, 2016.
- [10] C.-Y. Hong, S. Kandula, R. Mahajan, M. Zhang, V. Gill, M. Nanduri, and R. Wattenhofer. Achieving high utilization with Software-driven WAN. *SIGCOMM'13*, pages 15–26, 2013.
- [11] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, J. Zolla, U. Hölzle, S. Stuart, and A. Vahdat. B4: Experience with a globally-deployed software defined wan. *SIGCOMM'13*, pages 3–14, 2013.
- [12] H. Ji, J. H. Lee, and Y. C. Chung. System outage probability due to dispersion variation caused by seasonal and regional temperature variations. *Optical Fiber Communication Conference and Exposition and The National Fiber Optic Engineers Conference*, page OME79, 2005.
- [13] I. P. Kaminow, T. Li, and A. E. Willner. *Optical Fiber Telecommunications*. Academic Press, Burlington, fifth edition, 2008.
- [14] M. Karlsson, J. Brentel, and P. Andrekson. Long-term measurement of PMD and polarization drift in installed fibers. *Journal of Lightwave Technology*, 18(7):941–951, July 2000.
- [15] R. R. Kompella, J. Yates, A. Greenberg, and A. C. Snoeren. IP fault localization via risk modeling. *NSDI'05*, pages 57–70, 2005.
- [16] G. Li, D. Wang, R. Doverspike, C. Kalmanek, and J. Yates. Economic analysis of IP/optical network architectures. page FH5, 2004.
- [17] J. C. Li, K. Hinton, P. M. Farrell, and S. D. Dods. Optical impairment outage computation. *Opt. Express*, 16(14):10529–10534, Jul 2008.
- [18] H. H. Liu, S. Kandula, R. Mahajan, M. Zhang, and D. Gelernter. Traffic engineering with forward fault correction. *SIGCOMM'14*, 44(4):527–538, Aug. 2014.
- [19] T. Marian, D. Freedman, K. Birman, and H. Weatherspoon. Empirical characterization of uncongested optical lambda networks and 10GbE commodity endpoints. *IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, pages 575–584, June 2010.
- [20] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C.-N. Chuah, and C. Diot. Characterization of failures in an IP backbone. 4:2307–2317, March 2004.
- [21] T. Mizuochi. Recent progress in forward error correction and its interplay with transmission impairments. *IEEE Journal of Selected Topics in Quantum Electronics*, 12(4):544–554, July 2006.
- [22] Y. Tremblay. Circuit and method of testing for silent faults in a bi-directional optical communication system, 1998. US Patent 5,781,318.
- [23] B. Vidalenc, L. Ciavaglia, L. Noirie, and E. Renault. Dynamic risk-aware routing for OSPF networks. pages 226–234, May 2013.
- [24] S. Woodward, L. Nelson, M. Feuer, X. Zhou, P. Magill, S. Foo, D. Hanson, H. Sun, M. Moyer, and M. O'Sullivan. Characterization of real-time PMD and chromatic dispersion monitoring in a high-pmd 46-gb/s transmission system. *IEEE Photonics Technology Letters*, 20(24):2048–2050, Dec 2008.
- [25] S. Woodward, L. Nelson, C. Schneider, L. Knox, M. O'Sullivan, C. Laperle, M. Moyer, and S. Foo. Long-term observation of PMD and SOP on installed fiber routes. *IEEE Photonics Technology Letters*, 26(3):213–216, Feb 2014.
- [26] H. Zhu, K. Zhu, H. Zang, and B. Mukherjee. Cost-effective WDM backbone network design with OXCs of different bandwidth granularities. *IEEE Journal on Selected Areas in Communications*, 21(9):1452–1466, 2003.