

Evaluation of Elastic Modulation Gains in Microsoft's Optical Backbone in North America

Monia Ghobadi*, Jamie Gaudette*, Ratul Mahajan*, Amar Phanishayee*, Buddy Klinkers*, Daniel Kilper+

* Microsoft, Redmond, Washington, USA, 98052, + College of Optical Sciences, University of Arizona, Tucson, AZ, USA, 85721

{mgh, jagaudet, ratul, amar, buddyk}@microsoft.com, {dkilper}@optics.arizona.edu

Abstract: We study Q-factor and its variation in Microsoft's backbone network carrying live traffic.

We find that a capacity gain of at least 70% is achievable via elastic modulation.

OCIS codes: (060.4510) Optical communications; (060.4256) Networks, network optimization

1. Introduction

Wide-area backbone networks of Internet Service Providers (ISPs) and cloud providers, such as Microsoft, are the workhorses of Internet traffic delivery. To support the explosive demand, efficient use of the physical layer is imperative [1]. Bandwidth variable transceivers (BVTs) are an enabling technology for elastic optical networks. They can convert physical layer performance into network capacity using variable coded modulation, rate-adaptive FEC or parity coding, or other means [2-4]. For such techniques to find application, however, we need to understand the potential gains in practical field deployments. Signals require sufficient performance above transmission engineering margins (additional penalties that take into account uncontrolled system time and configuration dependent phenomena) in order to move to higher capacity modulation. In this paper, we examine both the measured and simulated optical performance of Microsoft's backbone network, and estimate the achievable benefits of BVTs.

Since the value of BVT is strongly dependent on the performance and available margins, we quantify the performance of a live backbone network and use that data to estimate the potential benefits. In particular, for a period of three months, we poll the signal Q-factor of all 100Gb/s PM-QPSK channels in the backbone network. From this, we can determine the performance of each channel relative to its minimum. We simulate the performance of 8-QAM and 16-QAM modulations in order to determine the propagation penalties associated with the fiber nonlinear response, which must be included due to its uncertainty. We then evaluate the number of paths that may be accessible using these formats and approximate further improvements from BVTs with a resolution of 25 Gb/s. Next, we consider the time and wavelength variations of Q-factor to determine whether fluctuations might be too rapid to adapt to as well as whether wavelength-dependent variation margins are accessible to BVT.

Our results show that substantial gain (up to 70%) can be achieved using a BVT capable of PM-QPSK, PM-8QAM, and PM-16QAM. A further 16% gain can be achieved using a BVT with 25 Gb/s resolution, increasing the total gain to 86%. We project another 13% gain using a BVT with 1Gb/s resolution. We also conclude that BVTs should have the ability to tune bandwidth on a channel-by-channel basis to optimize for wavelength-dependent performance. However, they can be relatively static in nature, commissioned at start-of-life or after major maintenances, and do not need the ability to tune bandwidth in real-time.

2. Cumulate Performance Analysis

We poll Microsoft's North American backbone network to collect Q-factor for all 100Gb/s PM-QPSK line-cards (thousands of line-cards) with an aggregate signal rate of 31.8 Gbaud. The optical fiber is primarily a mixture of LEAF and SSMF with a channel spacing of 50 GHz. Many segments on the network still have inline dispersion compensating modules (DCMs) to support legacy technology. The segments range from 5 km to 2600 km in length. The samples were taken as 15-minute bins, with the minimum, maximum and average extracted from each unique 15-minute sample. We focus on the average values since the minimum and maximum are similar to the average for a vast majority samples. Samples were collected continuously over a three-month period (Feb-April 2015).

To estimate the impact of BVTs, we first consider the network as a whole. We convert the Q-factor measurements to received electrical SNR [5] and plot the cumulative distribution function (CDF) for all samples in the three-month period (Fig. 1). The received SNR is defined as E_s / N_o , where E_s is the average symbol energy and N_o is the double-sided noise power spectral density [3,5]. The measured SNR values include all linear and nonlinear propagation penalties for PM-QPSK.

To determine the potential capacity gain of PM-8QAM and PM-16QAM, we calculate the electrical SNR limits of these modulation formats. For this calculation, we assumed a 32 Gbaud transceiver with a FEC limit of 3.77e-2 (or a Q^2 [dB] of 5.0), and minimum OSNR at FEC limit of 10 dB/0.1nm [QPSK], 14 dB/0.1nm [8QAM], and 17 dB/0.1nm [16QAM]. For PM-8QAM and PM-16QAM, we estimate through simulations propagation penalties between 0-2.2 dB and 0-1.5 dB of SNR (depending on distance and fiber type), respectively. The lower penalty for PM-16QAM is

likely due to the shorter propagation distances supported by this modulation format. The vertical boxes in Fig. 1 mark the computed SNR limits and the propagation penalties. By applying these SNR limits, we estimate that 8QAM is applicable to between 78% (low) and 99% (high) samples on the network and 16QAM is applicable to between 12% and 43% of the samples, depending on propagation penalties (Fig. 2). The total capacity gain amounts to between 45% and 70% by using PM-8QAM and PM-16QAM where possible. We also use the simulation of representative segments to estimate the distribution of applicable modulation formats. As shown in Fig. 2, our simulations suggest that 95% of the channels on the network could improve to 8QAM and 45% could improve to 16QAM, leaving only 5% of the channels at QPSK, for a net capacity gain of 70%. The simulations match the high end of our measured prediction which included legacy 10G OOK interferers and DCMs on many segments, not in the simulations.




Fig. 1. CDF of converted electrical SNRs from measured Q-factors across Microsoft backbone. The vertical lines are modulation SNR limits.

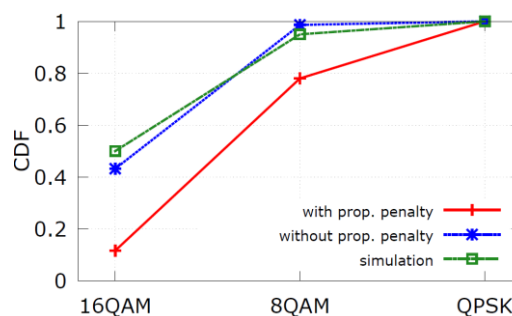


Fig. 2. Cumulative percentage of channels that can be upgraded to 8QAM/16QAM based on SNR estimation with and without propagation penalty (shown in Fig. 1). We also included simulation results for comparison.

Next, we analyze the benefits of a BVT with 25 Gb/s resolution operating between 100 Gb/s and 250 Gb/s. We use a simple approach to estimate the SNR limits of 125 Gb/s and 175 Gb/s by assuming they lie exactly between QPSK/8QAM/16QAM. In reality, the limits will vary with implementation. For 225 Gb/s and 250 Gb/s, we assume 512SP-QAM and PM-32QAM respectively [3]. The vertical lines in Fig. 3 show the approximated SNR limits with 25 Gb/s resolution. As shown in Fig. 4, by applying the SNR limits of 100 Gb/s up to 250 Gb/s with 25 Gb/s steps, we estimate that more than 90% of samples can increase their capacity to 175 Gb/s or higher. The average capacity gain using 25 Gb/s steps is 86%, compared to 70% best-case gain using coarse 8QAM/16QAM modulations. These results demonstrate that a 25 Gb/s resolution on optical modulation offers 16% additional capacity gain than a QPSK/8QAM/16QAM transponder (hatched area in Fig. 4).

To estimate the gain of even finer-grained modulation formats, we further divide the resolution to 1Gb/s steps and measure a total of 13% gain over 25Gb/s resolution. Given the disproportionately high complexity associated with 1Gb/s resolution, and the relative small gain over 25Gb/s resolution, we do not recommend such fine grained BVTs.




Fig. 3. CDF of converted electrical SNRs from measured Q-factors. The vertical lines are modulation SNR limits with 25 Gb/s resolution.

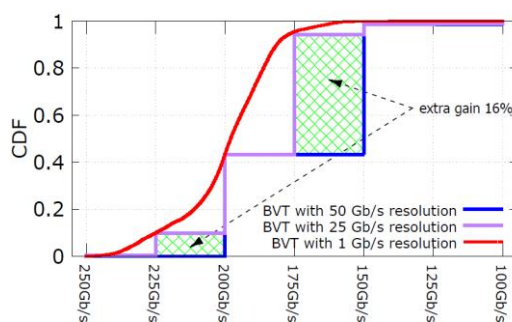


Fig. 4. Cumulative percentage of channels that can be upgraded to faster modulations using 50, 25, and 1 Gb/s resolution.

3. Q-factor Across Wavelength and Over Time

To understand the benefits of per-channel-varying BVTs or time-varying BVTs, we converted the Q-factor samples to SNR and analyzed them across wavelengths traversing the same optical path and across time.

Fig. 5 shows average SNR as a function of wavelength for a representative sample of network segments. Each line represents a different segment and each data point on a line represents a wavelength traversing the segment. The plot

shows that SNR values are different within and across segments as expected. Fig. 6 shows the CDF of the largest SNR value minus the smallest SNR value within the same segment. We see that the median SNR difference within a segment is 0.8 dB and 99th percentile is 3.8 dB. That the SNR of wavelengths within a segment can differ significantly implies that different wavelengths might benefit from different modulation formats even though the path is shared.




Fig. 5. SNR variation across wavelengths and segments. Each line represents a segment. Each point is a wavelength.




Fig. 6. Cumulative percentage of segments and their SNR variation across all wavelengths traversing the same segment.

Fig. 7 illustrates SNR over time at mid-band (~1550nm) across a representative sample of network segments. As shown, for each channel (each line in Fig. 7) SNR remains mostly stationary over time. For each channel, the performance range and standard deviation was calculated over the entire 3-months and their CDFs are shown in Fig. 8. The standard deviation was less than 0.2 dB for 95% of the data. However, (max – min) reveals many channels with a large change over time. The large range values were found to be due to occasional network changes (e.g. removal of legacy 10G OOK and/or removal of DCMs). We conclude that a majority of the BVT gains can be obtained by adapting over long time periods and to network operational changes or maintenance.




Fig. 7. SNR vs. time for some 100Gb/s PM-QPSK channels.




Fig. 8. CDF of SNR variation over time for each channel.

4. Conclusions

We study Q-factor in a large optical backbone network and estimate that substantial gain (70%) can be achieved using a BVT capable of PM-QPSK, PM-8QAM, and PM-16QAM. A further 16% gain, for a total of 86%, can be achieved using a BVT with 25 Gb/s resolution. We observe significant wavelength-dependent performance that might be accommodated by per-channel BVT to realize the estimated performance gains. The performance over time did not vary significantly under a static network configuration, suggesting that the BVT can be largely static in nature, configured at start-of-life or after major network maintenance (e.g., new channel provisioning). Therefore, at 25 Gb/s resolution, real-time BVTs do not appear necessary as long as the network configuration remains static.

5. Acknowledgement

This work was supported in part by NSF Center for Integrated Access Networks (CIAN) EEC-0812072.

6. References

- [1] X. Zhao et al., “The prospect of inter-data-center optical networks”, IEEE Communication Magazine, Sept. 2013
- [2] Y. Yoshida et al., “SDN-based network orchestration of variable-capacity optical packet switching network over programmable flexi-grid elastic optical path network”, Journal of Lightwave Technology, vol. 33, pp. 609–617, Feb. 2015.
- [3] J. K. Fischer, “Bandwidth-variable transceivers based on four-dimensional modulation formats”, Journal of Lightwave Technology, 2014.
- [4] B. Teipen et al., “Adaptive data rates for flexible transceivers in optical networks”, Journal of Networks, vol. 7, no. 5, 2012.
- [5] A. Carena et al., “Modeling of the impact of nonlinear propagation effects in uncompensated coherent transmission links”, Journal of Lightwave Technology, vol. 30, no. 10, May 15, 2012.