Active Monitoring of All-Optical Networks

Yvan Pointurier, Mark Coates, and Michael Rabbat

Department of Electrical and Computer Engineering, McGill University, Montreal, Canada
E-mail: $\{$*yvan.pointurier,mark.coates,michael.rabbat*$\}$*@mcgill.ca*

ABSTRACT

This paper shows that it possible to decrease the amount of hardware needed to monitor the Quality of Transmission (QoT) of lightpaths in all-optical networks by establishing carefully selected “active lightpaths”, without sacrificing estimation accuracy. The QoT of lightpaths that are not directly measured by monitoring equipment is estimated with the help of additional “active lightpaths”. We gain extra information about the unobserved lightpaths by measuring the QoT of the carefully chosen active lightpaths, which are activated solely for the purpose of monitoring. We demonstrate with simulations the possibility to trade-off the amount of costly hardware monitoring equipment with cheaper, temporary “active lightpaths”, while still achieving accurate monitoring.
Keywords: monitoring, QoT, physical impairments, linear estimation.

1 INTRODUCTION

In all-optical networks, clients are typically guaranteed a minimum QoT for their signals, measured through Bit-Error Rates (BER). Thus, network managers need to be able to know the quality of the signals that travel over their network at all times. Hardware monitors are expensive devices that need access to the signals to be monitored, which is difficult in all-optical networks since signals remain in the optical domain from end to end. The main contribution of this paper is a QoT monitoring scheme that reduces the amount of hardware needed to perform monitoring.

We call lightpaths that are established in the network and carry information “passive lightpaths”. In a recent paper [1], we presented a *passive* monitoring scheme for all-optical networks. To perform QoT estimation, we used the spatial correlation between the QoT metrics of the passive lightpaths to estimate the QoT for the lightpaths that are not directly observed. The spatial correlation is induced by the physical behavior of the network: physical impairments are caused at the link level and thus the BERs of two different lightpaths (on different wavelengths) sharing links are correlated. However, if few of the lightpaths established in the network terminate at hardware monitors, then estimation is based on very few observations, thereby leading to relatively high estimation errors. We propose a scheme to establish a limited number of lightpaths *in addition to* the lightpaths already established in the network. These new “active” lightpaths are established solely to gather QoT information about the existing lightpaths that carry traffic but are not directly observed by hardware monitors. There is a trade-off between the amount of monitoring hardware and the accuracy of the QoT estimates. We are able to essentially remove the trade-off, that is, to obtain accurate estimates with little monitoring equipment, at the cost of injecting more signals in the network. Note that these additional lightpaths do not carry any meaningful data (they could carry random bits) and are used only for a QoT monitoring purpose.

This paper is organized as follows. In Section 2, we state our assumptions concerning the network physical layer and hardware monitors. Section 3 is devoted to the presentation of our active monitoring scheme. We present simulation results in Section 4 and we conclude the paper in Section 5.

2 MODELING

Our monitoring technique exploits physical-layer properties and assumptions, which are presented here. We assume that all-optical regeneration, except for standard optical amplification, is not available; in particular, we assume that wavelength converters are not available and the *wavelength continuity constraint* holds: lightpaths are established over the same wavelength from end to end. In this paper, the impairments we consider are amplifier noise and intersymbol interference (ISI). We estimate the BER of a lightpath through its so-called Q-factor: $Q = \frac{\mu_0-\mu_1}{\sigma_0+\sigma_1}$ (where μ_0 and μ_1 are the means of the distributions of the “0” and “1” symbols at the photo-detection

This work was supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada and industrial and government partners, through the Agile All-Photonic Networks (AAPN) Research Network and the Mathematics of Information Technology and Complex Systems (MITACS) Network of Centres of Excellence (NCE).

stage, and σ_0 and σ_1 are their respective standard deviations) such that $\mathrm{BER} = \frac{1}{2}\mathrm{erfc}\big(\mathrm{Q}/\sqrt{2}\big)$. We model the impact of amplifier noise and ISI as additive variances (σ_n^2 and σ_{isi}^2, respectively) in σ_1^2 such that $\sigma_1^2 = \sigma_n^2 + \sigma_{isi}^2$ [2].

Here, we make no assumption about what hardware monitoring technique is deployed in the network. We only assume that the electrical power and noise can be measured directly or indirectly by the hardware monitors. In particular, we model a hardware monitor as a device located at the extremity of a link, after photo-detection. Hardware monitors can determine μ_0, μ_1, σ_0 and σ_1. Because of the high costs related to deploying hardware monitors, it is desirable to limit their number in the network. Note that placing a hardware monitor directly on a transmission line would imply diverting some signal power to measure the statistics of interest. For this reason, the hardware monitors are assumed to be located inside the receiver modules. A consequence is that only the lightpaths that terminate at a link where a monitor is located can be observed. Hardware monitor placement is out of the scope of this paper and we assume that monitors are already placed at fixed locations (chosen offline) when our monitoring technique is applied.

3 ACTIVE MONITORING

The active monitoring technique consists of two components. First, we select the set of active lightpaths that will allow for accurate estimation (Section 3.2). Then, once the active lightpaths are selected, the BERs of the unobserved lightpaths are estimated from passive and active observations (Section 3.3). To perform this estimation, we adapt the network kriging procedure [3]. Network kriging estimates the values of end-to-end metrics from a limited set of end-to-end observations; this is made possible by the existence of the aforementioned correlation between the end-to-end metrics (lightpaths' QoT). Network kriging also exploits the existence of a linear relationship (e.g., additivity over a path) between the link-level metrics and the end-to-end metrics to be estimated. In all-optical networks, BERs or Q factors are not additive. However, each of the four quantities μ_0, μ_1, σ_0, and σ_1^2 is approximately linear with respect to the number of spans over which a signal is transmitted (see Fig. 1), making network kriging suitable for lightpaths' BER estimation through the estimation of μ_0, μ_1, σ_0, and σ_1^2.

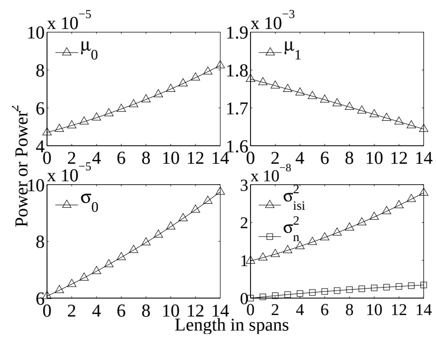

Figure 1: Linearity of the physical impairments metrics.

3.1 Notations and problem statement

Consider an all-optical network of n_ℓ links where n_p lightpaths are established. We denote by $\mathbf{y}$ a column vector containing path-level QoT metrics (μ_0, μ_1, σ_0, or σ_1^2 for some lightpaths) and by $\mathbf{x}$ the corresponding link-level metrics. Let n_s be the number of lightpaths observed through hardware devices. We denote by $G_m \in \{0,1\}^{n_s \times n_\ell}$ the routing matrix that describes the passive lightpaths that are observed, that is, the element on row i and column j of G_m is 1 when lightpath i traverses link j. Similarly, we denote by $G_n \in \{0,1\}^{(n_p - n_s)\times n_\ell}$ the routing matrix corresponding to the passive lightpaths that are not observed. We denote by $\mathbf{y}_m$ the column vector containing the QoT metrics for the observed passive lightpaths and by $\mathbf{y}_n$ the column vector containing the metrics for the passive lightpaths that are not observed, such that $\mathbf{y}_m = G_m\mathbf{x}$ and $\mathbf{y}_n = G_n\mathbf{x}$. .

We propose to establish a limited number n_a of observed, "active" lightpaths. Our goal is to design an "active routing matrix" $G_a \in \{0,1\}^{n_a \times n_\ell}$, where each row corresponds to an active lightpath. Let $\mathbf{y}_a$ be the column vector for metrics corresponding to these active lightpaths ($\mathbf{y}_a = G_a\mathbf{x}$). Call $G_A = \left[G_m^T, G_a^T\right]^T$ the routing matrix corresponding to all observed lightpaths and $\mathbf{y}_A = \left[\mathbf{y}_m^T, \mathbf{y}_a^T\right]^T$ the column vector containing the metrics for all observed paths.

Using the notations above, the best linear estimate $\hat{\mathbf{y}}_n$ for the metrics corresponding to non-observed lightpaths can be shown to be [3]:

$$\hat{\mathbf{y}}_n = G_n G_A (G_A G_A^T)^+ \mathbf{y}_A, \tag{1}$$

where $(\cdot)^+$ denotes the Moore-Penrose inverse. Active monitoring hence consists in finding the matrix G_a^* to minimize the relative mean square error (RMSE) for the non-observed metrics $\mathbf{y}_n$, given the observations y_A and the routing matrices G_n and G_A:

$$G_a^* = \arg\min_{G_a} \|\hat{\mathbf{y}}_n - \mathbf{y}_n\|_2 / \|\mathbf{y}_n\|_2, \tag{2}$$

where $\|\hat{\mathbf{y}}_n - \mathbf{y}_n\|_2/\|\mathbf{y}_n\|_2$ is a function of G_A and hence of the active routing matrix G_a via $\hat{\mathbf{y}}_n$.

This optimization problem is subject to the following constraints: each row of G_a^* corresponds to a lightpath that meets the wavelength continuity constraint, and the last link of each active lightpath is equipped with a monitor. The problem, as formulated here, is difficult because the constraints on G_a^* are binary, making a full search computationally prohibitive; thus, heuristics are needed. In the following section, we present such a heuristic.

3.2 Active lightpaths selection

We propose the following heuristic to establish n_a active lightpaths:

Step 1: We compute, for each unobserved lightpath, the shortest path from the last node of the considered unobserved lightpath to any hardware monitor. This shortest path is appended to the route used by the considered unobserved lightpath, to form a new route comprehending all the links of the unobserved lightpaths. Lighting the active lightpaths uses extra resources and may disturb the passive lightpaths, for instance, through crosstalk injection. Using shortest paths mitigates this problem. Then, a free continuous wavelength is tentatively found over this route. Selecting candidates for active monitoring in this fashion ensures that the active lightpaths to be selected in a further step are indeed observed, that they contain many links from lightpaths that were previously not observed, and that the quantity of resources used by the active lightpaths remains low. The set of candidate lightpaths for active monitoring is called ACTIVECANDIDATES and the corresponding routing matrix is G'_a. Let n'_a be the number of rows of G'_a. If $n'_a > n_a$ then we need to choose n_a active lightpaths among the n'_a candidates (Step 2), otherwise the algorithm stops.

Step 2: The problem of choosing n_a lightpaths to minimize the estimation error for the linear estimation procedure amounts to an NP-complete row selection problem [3], but a good heuristic consists in finding the n_a rows of G'_a that approximate well the space spanned by the first n_a singular vectors of G'_a [4]. We use this technique to iteratively select an increasing number of rows of G_a^*, until the number of active candidates within the set of selected rows is equal to n_a. When the algorithm returns, at most n_a lightpaths of ACTIVECANDIDATES are selected. Because we used the aforementioned span maximization algorithm to select those candidates, the selection procedure returns lightpaths which are expected to yield a good accuracy in the estimation step.

3.3 Estimation

The goal of the estimation step is to estimate the QoT for unobserved lightpaths given the QoT of all observed lightpaths, and information about the (spatial) correlation structure between the QoT for different lightpaths. The existence of the linear relationships $\mathbf{y}_m = G_m\mathbf{x}$, $\mathbf{y}_n = G_n\mathbf{x}$ and $\mathbf{y}_a = G_a\mathbf{x}$ ensures that we can use the network kriging procedure to estimate $\mathbf{y}_n$ given G_n, G_m, G_a and $\mathbf{y}_n$. This estimation procedure is run in turn for $\mathbf{y}_m, \mathbf{y}_n, \mathbf{y}_a \in \{\mu_0, \mu_1, \sigma_0, \sigma_1^2\}$ to return estimates for each of these four quantities for the unobserved lightpaths.

4 SIMULATIONS

We simulate the operation of a scaled-down version of the NSFNET topology with standard physical parameters (10 Gbps NRZ signals traveling over 70 km spans of 100% post-dispersion compensated SMF, noise factor=2, 8 wavelengths). The results presented here are averaged over the establishment or termination of 100 lightpaths, totalling 100 different network states. In each state, around 50 lightpaths are simultaneously present in the network. The network contains 42 unidirectional links and hence would require 42 monitors to gather a complete observation set about the QoT in the network. We are interested in cases where far fewer devices are deployed.

Our main result is the evaluation of the active monitoring technique through the RMSE for $\log(\mathrm{BER})$. In the left panel of Fig. 2, we show the RMSE for our technique. Values for $n_a = 0$ active lightpaths denote passive monitoring results. Establishing a relatively small number (e.g., 15) of active lightpaths is shown to decrease sharply the estimation error, especially when fewer number of monitoring devices are installed. There exists a trade-off between the number of physical devices and the number of active lightpaths: for instance, the estimation error is the same, around 6%, whether 30 monitors are deployed and passive monitoring is used, or only 5 monitors are deployed in conjunction with active monitoring (establishing an additional 20 lightpaths in the network).

Notice that the estimation error does not converge to 0 as n_a increases: rather, it exhibits a floor. The value for this floor depends on the number of monitors, and decreases as the number of monitors increase and permits

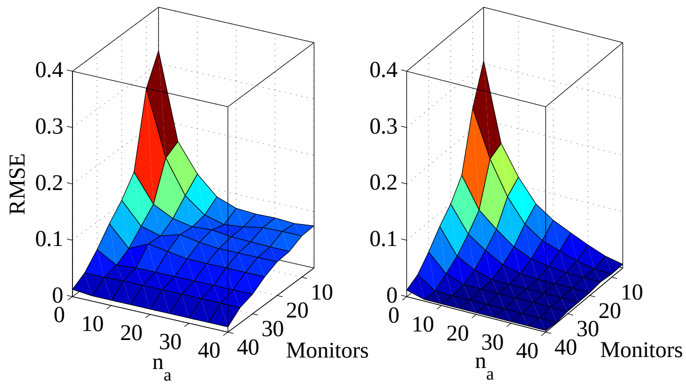

Figure 2: RMSE: without/with linearized metrics (left/right).

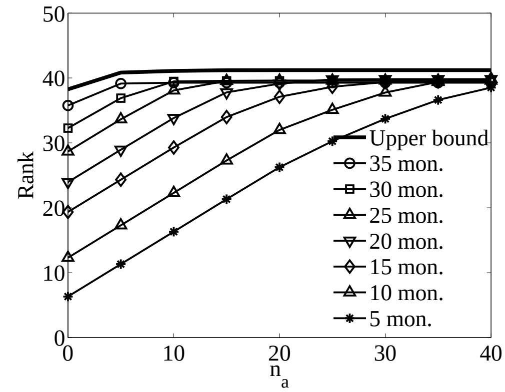

Figure 3: The rank of $\left[G_m^T, {G^*}_a^T\right]^T$ increases with n_a.

more accurate estimation. This can be attributed to the linearization noise, that is, the approximation incurred in the linearization step. Indeed, we show in the right panel of Fig. 2 the RMSE in the case where all four metrics μ_0, μ_1, σ_0, and σ_1^2 are artificially linearized: they are given values that are exactly linear with distance. It is seen that the estimation error now converges to 0 as the amount of monitoring hardware or n_a increase.

Insight as to why our technique indeed decreases RMSE for $\log(\text{BER})$ can be gained by observing the variations of the rank of $\left[G_m^T, G_a^{*T}\right]^T$ with the number of monitors and active lightpaths n_a (see Fig. 3). For low numbers of monitors and active lightpaths, this rank is well below the numbers of lit links, making BER estimation difficult or inaccurate. However, as n_a increases, this rank increases with a slope of approximately 1, converging to $\operatorname{rank}\left(\left[G_m^T, G_n^T, G_a^{*T}\right]^T\right)$ (bold curve). The fact that the latter rank (bold curve) increases with n_a very little (by 1) indicates that establishing active lightpaths introduces at most a single new unknown, making the estimation problem easier to solve.

5 CONCLUSIONS

We showed that it possible to decrease the amount of hardware needed to monitor the QoT of lightpaths in a all-optical networks by establishing carefully selected "active lightpaths", without sacrificing estimation accuracy. The estimation technique relies on a linear assumption for the quantities involved in the QoT computations. Generalizing our technique to more complex cases where linearity does not hold will be the subject of future work. Moreover, integrating the impact of establishing the additional lightpaths on the QoT of the passive lightpath should be further investigated.

ACKNOWLEDGMENTS

The authors would like to thanks Boris Oreshkin for his insightful comments.

REFERENCES

[1] M. Coates, Y. Pointurier, and M. Rabbat, "Compressed network monitoring for IP and all-optical networks," in *Proc. ACM/Usenix Internet Measurement Conf.*, San Diego, CA, USA, Oct. 2007.

[2] Y. Pointurier, M. Brandt-Pearce, T. Deng, and S. Subramaniam, "Fair QoS-aware adaptive Routing and Wavelength Assignment in all-optical networks," in *Proc. IEEE ICC*, Jun. 2006.

[3] D. B. Chua, E. D. Kolaczyk, and M. Crovella, "Network kriging," *IEEE J. Select. Areas Commun.*, vol. 24, no. 12, pp. 2263–2272, Dec. 2006.

[4] G. Golub and C. V. Loan, *Matrix Computations*. The Johns Hopkins University Press, 1996.