Quality of Transmission Estimation in WDM and Elastic Optical Networks Accounting for Space–Spectrum Dependencies

I. Sartzetakis, K. Christodoulopoulos, C. P. Tsekrekos, D. Syvridis, and E. Varvarigos

Abstract—**We develop a framework for estimating the quality of transmission (QoT) of a new lightpath before it is established, as well as for calculating the expected degradation it will cause to existing lightpaths. The framework correlates the QoT metrics of established lightpaths, which are readily available from coherent optical receivers that can be extended to serve as optical performance monitors. Past similar studies used only space (routing) information and thus neglected spectrum, while they focused on old-generation noncoherent networks. The proposed framework accounts for correlation in both the space and spectrum domains and can be applied to both fixed-grid wavelength division multiplexing (WDM) and elastic optical networks. It is based on a graph transformation that exposes and models the interference between spectrum-neighboring channels. Our results indicate that our QoT estimates are very close to the actual performance data, that is, to having perfect knowledge of the physical layer. The proposed estimation framework is shown to provide up to 4×10^{-2} lower pre-forward error correction bit error ratio (BER) compared to the worst-case interference scenario, which overestimates the BER. The higher accuracy can be harvested when lightpaths are provisioned with low margins; our results showed up to 47% reduction in required regenerators, a substantial savings in equipment cost.**

Index Terms—**Correlation; Fixed grid and elastic optical networks; Interference; Network kriging; Physical layer impairments; Quality of transmission (QoT) estimation.**

I. INTRODUCTION

Internet traffic has been growing continuously in recent years, with new applications, such as HD video on demand and cloud computing, requiring high capacity, which only optical networks can provide, and flexibility, which is the promise of elastic optical networks (EONs). In transport optical networks, the optical signals can transparently pass intermediate nodes (without undergoing optical–electrical–optical conversion) and traverse long links. The accumulated impairments may degrade the quality of transmission (QoT) of the signal to an unacceptable degree, necessitating the use of regenerators at certain intermediate hops.

The QoT of an established lightpath does not remain constant but decreases as time passes, due to increased interference from new lightpaths (the utilization of the network is light at the beginning of its life and increases as more connections are established), equipment aging, and maintenance operations (e.g., splices after fixing fiber cuts). Traditional provisioning of lightpaths makes use of tools (also referred to as Q-tools) that estimate the QoT of the lightpaths that use abundant margins. Margins are used to account for any inaccuracies in their models and to avoid subsequent interventions during the network's lifetime. High margins often force the deployment of regenerators or more robust transponders that are not strictly necessary during the initial setup. Clearly, provisioning with lower margins and placing equipment "just in time" (JIT) would be both desirable and beneficial [1–3] by avoiding or postponing the purchase of equipment by operators and reducing the overall cost, as equipment costs fall with time and saved capital translates to reduced loans or interest.

Lowering the margins used in provisioning, however, requires new feedback-based networking mechanisms. Such mechanisms can rely on the use of optical performance monitors (OPMs) to observe the state of the network and to (i) estimate in an accurate way the QoT before provisioning new lightpaths and (ii) anticipate, identify, and remedy the QoT problems that could occur at later times. In the past, hardware QoT monitors were expensive and, as a result, only power monitors were deployed. The work in [4] developed an algorithm for the placement of (expensive) hardware monitors in a few selected locations in order to supervise the QoT of lightpaths during network operation. During the past few years, optical coherent transceivers have started being installed in core networks, and the

I. Sartzetakis (e-mail: isartz@mail.ntua.gr), K. Christodoulopoulos, and E. Varvarigos are with the Computer Technology Institute and Press (CTI), Rion 26504, Greece.

I. Sartzetakis, K. Christodoulopoulos, and E. Varvarigos are also with the School of Electrical and Computer Engineering, National Technical University of Athens, Zografou, 15773 Athens, Greece.

C. P. Tsekrekos was with the Department of Informatics and Telecommunications, National and Kapodistrian University of Athens, Ilissia, 15784 Athens, Greece. He is now with Aston University, Birmingham B4 7ET, UK.

D. Syvridis is with the Department of Informatics and Telecommunications, National and Kapodistrian University of Athens, Ilissia, 15784 Athens, Greece.

general expectation is that we are moving toward all-coherent optical networks [5]. Such transceivers employ DSP processing at the receivers and are thus able to monitor and compensate certain impairments. The ORCHESTRA project [6] proposes to extend these coherent receivers to operate as software OPMs and develops a scalable and responsive monitoring and control plane to support and use such data, providing the solution for lowering the margins, as well as other benefits in dynamic network operation, as discussed in the following.

With the advent of flex-grid and tunable transceivers, optical networks are becoming more dynamic, typically referred to as flexible or elastic optical networks (EONs) [7,8]. In such a dynamic optical network environment, the reconfiguration actions would benefit from accurate QoT estimation. For example, assume a network that automatically adapts to traffic changes. Making dynamic decisions to change the spectrum used and/or the modulation format of existing lightpath(s) to cope with traffic changes requires the estimation of the QoT of the changed lightpath(s) and the effect the changes have on the other established ones. Therefore, accurate QoT estimation explored in the present paper conforms to a network that observes itself and uses such information in dynamic optimization decisions, closing the control cycle as envisioned in ORCHESTRA [6]. Note that this control cycle of observing, correlating, and re-adjusting the network is also very much in line with the current trends toward software defined networking (SDN).

Specifically, in this paper, we present a framework that correlates monitoring information from established lightpaths to estimate (a) the QoT of a new lightpath before it is established and (b) the degradation the new lightpath will cause to existing ones. This can be used to provision a lightpath with low margins and enable optimized and dynamic reconfiguration actions, which both increase network efficiency, as discussed above. The network studied is a single- or multi-rate traditional WDM or EON with coherent transmission. We assume that coherent receivers provide monitoring information on the electrical signal-to-noise ratio (SNR), which accounts for all optical layer impairments and is used to calculate the bit error ratio (BER), which is the ultimate QoT metric. Note that the SNR is already measured and reported by commercial coherent receivers. Alternative, individual impairments or other QoT metrics that can be measured by soft OPMs [6] can be estimated. The proposed framework correlates SNR measurements, taking into account the space (routing) and the spectrum domains (interference of spectrum neighboring channels), through the introduction of an auxiliary graph, which we call the interference-aware (IA) graph, on which the estimation algorithms run. As a result, the SNR and then the BER estimates provided by our framework are more accurate and realistic than those obtained by previous approaches (that only took the space domain into account). This in turn helps the network manager make more optimized decisions [9].

We use the Gaussian noise (GN) model [10,11] to approximate the behavior of the physical layer and conduct a set of simulation experiments to evaluate the accuracy of the estimation framework and the benefits it provides. The worst-case interference scenario that assumes all channels are simultaneously lighted, always overestimates the BER, as opposed to our estimation framework that takes into account the actual lightpath utilization. As a result, the estimated BER is much closer to reality and quite lower than the BER of the worst-case interference scenario (assuming that the actual network load is not very high and is therefore not close to the worst-case scenario). We observed that the estimate of the pre-forward error correction (pre-FEC) BER of a new lightpath that our framework provides is up to 4×10^{-2} lower than the corresponding worst-case estimate. This has profound implications on the required number of costly regenerators, as we also show in the simulations. In particular, our framework was shown to require up to 47% fewer regenerators than the worst-case scenario, while it required only 5% more than the perfect estimation scenario, where we accurately know the BER of the new lightpath.

The rest of the paper is organized as follows. In Section II we report on previous work. In Section III we present the model of the network under study and certain physical layer considerations. In Section IV we give the notation and briefly introduce the estimation techniques we use. Then, in Section V we present the proposed interference-aware (IA) estimation framework. In Section VI we present the simulation results obtained, showing the accuracy and the other benefits our scheme provides. Finally, in Section VII we conclude the paper and discuss future work.

II. Previous Work

Estimating the QoT of lightpaths is a key functionality that is typically performed by a "Q-tool," which is used when planning, upgrading, or operating an optical network. QoT estimation methods range from very complex solving of Schrödinger equations, to simulations (such as VPI), and to analytical models of lower complexity. Recently the GN model [10,11] has been introduced and shown to be quite accurate, while its approximated closed form analytical version combines reasonable accuracy and low computational complexity. Such models perform *forward* QoT estimation based on accurate knowledge of the network characteristics and physical layer parameters. Actually, since it is not possible to obtain accurate values for the parameters used by such tools, margins are used to account for inaccuracies (referred to as *design* margins in [1]). High margins are also used to account for equipment aging and interference (referred to as *system* margins in [1]). A second category for QoT estimation methods includes *backward* or *feedback-based* methods [12,13], where measurements and monitored information are correlated to estimate the QoT of new lightpaths, enabling the reduction of both design and system margins and a more dynamic network operation.

Previous studies [12,13] worked toward the estimation of only the QoT of new lightpaths, by applying network correlation techniques. In more detail, [12] aimed

at estimating end-to-end QoT metrics [optical SNR, polarization mode dispersion (PMD), chromatic dispersion (CD), and self-phase-modulation (SPM)] for a to-be-established lightpath, based on measured data from already established ones. Measurements of the means and standard deviations of the distribution of "0" and "1" symbols (μ_0, μ_1, σ_0, and σ_1) were used in [13]. Based on such estimation, the Q factor and then the BER of a new light-path were calculated. Network kriging (NK) [14] and norm l_2 minimization (NM) were used in [12,13] to exploit the correlation of information provided by already established lightpaths that share common links. However, [12,13] modeled and accounted only for routing (space) dependencies, neglecting the relative spectrum positioning of the established lightpaths and thus not accounting for interference. They focused only on linear effects (effects that are additive per link) and old generation 10 Gbps on–off keying WDM networks. Since interference was not included in the model, worst-case assumptions are needed, namely, that the network is fully loaded and all channels of a link are simultaneously lighted. The QoT estimates obtained in this way are pessimistic [15] and do not reflect the actual state of the network. Avoiding such a pessimistic worst-case interference approach and taking into account the load of the network can yield substantial regeneration savings, as shown in [9].

The novelties of the proposed framework are manifold. Compared to [12,13], the proposed QoT estimation models the dependencies not only in the space (routing) but also in the spectrum domain, accounting for the actual interference of the lightpaths, and thus yielding more accurate QoT estimates. By using this framework we can also assess the degradation that the new lightpath will cause to existing ones. Moreover, the proposed method is applicable to current and next-generation coherent optical networks that use either fixed- or flex-grid (elastic) technology, as opposed to [12,13], which considered previous generation networks. Note that the SNR that is used in the estimation is already measured in coherent receivers. Compared to forward techniques (such as the GN model), the proposed method is agnostic to the values of physical layer parameters, while the feedback from the physical layer eliminates the use of high margins for end-of-life aging and worst-case interference.

It is worth noting that this work is an extension of [16]. The extensions include (i) the generalization of the estimation framework so as to be applicable to EONs, (ii) detailed description of the framework with pseudo-code for all procedures, and (iii) improvements in estimation for cases with limited correlation information. The results are also enriched, evaluating the performance in EONs, examining the effect of the number of active neighbors that are accounted for in terms of interference, while results are presented for a second network topology (SPARKLE), with shorter links than the NSFNET topology studied in [16].

III. NETWORK MODELING

We consider a national or continental optical transport network that employs either a fixed- (WDM) or flex

(elastic)-grid [7]. The nodes are assumed to be equipped with reconfigurable optical add–drop multiplexers (ROADMs) so that the traffic can be switched to the desired direction while the signal remains in the optical domain according to the routing decisions. The nodes are connected through uncompensated fiber links, each consisting of a number of fiber spans terminating at an amplifier. The latter compensates exactly the loss of the span. We assume that there is no wavelength conversion, and, thus, the wavelength (or spectrum, in the case of a flex-grid network) continuity constraint holds for a lightpath that crosses several links. The lightpaths are provisioned using a routing and wavelength assignment (RWA) algorithm or a routing and spectrum assignment (RSA) algorithm for a WDM or flex-grid network, respectively. If the QoT of a certain connection is estimated to be unacceptable, the algorithm decides to place regenerators at certain locations, and then each segment between regenerators is considered a separate lightpath that can use a different wavelength or spectrum.

We assume that the network is enriched with OPM capabilities: certain (all or a subset of) receivers in the network are used as OPMs. Note that coherent receivers deployed today are packed with digital signal processing (DSP) capabilities, so they can be extended, almost for free, to function as OPMs [6]. The DSP at the receivers performs electronic dispersion compensation and multiple-input–multiple-output (MIMO) equalization. In particular we assume that an OPM (receiver) can provide information about the SNR of the lightpath. We also assume that, from the SNR and taking into account the modulation format, we can calculate the pre-FEC BER, and then, taking into account the FEC, we can calculate the BER, which is the ultimate QoT metric. Note that SNR is reported even today by commercial coherent transceivers. The SNR takes into account all impairments, amplified spontaneous emission (ASE), nonlinear interference (NLI), and residual dispersion. Alternatively, we can assume that an OPM, such as the one envisioned in the ORCHESTRA project, can provide separate information for several impairments, and that the SNR can be calculated based on the monitored parameters.

A key approximating assumption that we make is that the inverse of the SNR is additive per link, or that impairment parameters are additive per link if we estimate those separately. As we have mentioned in the introduction, in this study we use the GN model to approximate the behavior of the physical layer: to obtain QoT estimates of the established lightpaths and to evaluate the accuracy of our estimation framework. The impairments considered are ASE and NLI, which consist of self-channel interference (SCI) and cross-channel interference (XCI). We ignore multi-channel interference (MCI) because it is almost always negligible [10], while the GN model assumes also ideal dispersion compensation at the receivers. In particular, we used the analytical approximations of the GN model as described in [11]. The BER estimates of the GN model will serve as the ground truth for validating the accuracy of our correlation techniques. In practice, the proposed framework, when deployed in a real network, will use (instead of the GN model) the actual monitored values provided by the OPMs in the field, and validation of the estimation

accuracy will be done based on the actual monitored values, after the establishment of the new lightpath.

We should stress that our estimation framework does not depend on the GN (or any other physical layer) model. Actually, its use would be to replace or complement such a physical layer model in a Q-tool, when provisioning new or updating existing lightpaths. The GN model is used in this paper only to establish the ground truth for evaluating the results in Section VI. The proposed scheme estimates the QoT based solely on exploiting the space and spectrum correlations of existing lightpaths, using the assumption that certain parameters, and, in particular, the inverse of SNR (or ASE noise, polarization, and NLI) are additive per link. It could be used and validated with other physical layer models (e.g., VPI) with small modifications. It can also be verified with real data from monitors, which, however, requires setting up an extensive network-wide experimental setup, and is a topic of future work. The use of the GN model allows us to rapidly simulate different network scenarios and compare the results for different assumptions. Taking as the ground truth the GN model, our method provides very close to perfect QoT estimation, as shown in Section VI, significantly more accurate than that obtained using the worst-case interference assumption that all channels are simultaneously lighted.

IV. ESTIMATION TECHNIQUES

We now formally define the correlation techniques that are used in our QoT estimation framework. We go through the basic network algebraic representation and then describe how we can estimate an unknown end-to-end value of a new path from values of existing paths. Note that we follow a generic description, where we do not define from the beginning the end-to-end metric that we use, and discuss this issue at the end of the next subsection.

A. Network Notation

We consider a network $G = (V, E)$ where V denotes the set of nodes and E the set of unidirectional fiber links. We assume a set M of lightpaths already established in the network, which define what we call the *state* of the network. A connection needs a pair of two unidirectional lightpaths in the opposite direction that follow opposite routes and occupy the same wavelength on the opposite unidirectional links. The routing matrix of the established lightpaths is defined as the binary matrix $R_M \in \{0, 1\}^{|M| \times |E|}$, where $R_M[m, l] = 1$ when lightpath m contains link l, and is 0, otherwise. Consider the end-to-end vector of parameters $\mathbf{y}_M \in \mathbb{R}^{|M|}$, with the y_m member of $\mathbf{y}_M$ representing the value of lightpath m. Vector $\mathbf{y}_M$ can be written as a linear combination of link-level vector parameters $\mathbf{x} \in \mathbb{R}^{|E|}$ so that $\mathbf{y}_M = R_M \mathbf{x}$. We assume that we want to estimate the end-to-end parameters of a set N of new lightpaths, denoted by vector $\mathbf{y}_N \in \mathbb{R}^N$, assuming that we know their routing $R \in \{0, 1\}^{|N| \times |E|}$. Then, we have

$$\mathbf{y} = \begin{bmatrix} \mathbf{y}_M \\ \mathbf{y}_N \end{bmatrix} = \begin{bmatrix} R_M \\ R_N \end{bmatrix} \mathbf{x}. \tag{1}$$

Consider, for example, the network of Fig. 1, where a set of $M = \{p_1, p_2, p_3, p_4, p_5\}$ is already established, and it corresponds to submatrix R_M and the known end-to-end impairment parameter values $\mathbf{y}_M = \{y_1, y_2, y_3, y_4, y_5\}$ (we will discuss the nature of the end-to-end parameter at the end of this subsection). Note that, for illustration purposes, only one direction of a connection is depicted in Fig. 1. The same convention holds for the equations below. The lightpath to be established is denoted by $N = \{p_6\}$ whose end-to-end value y_6 we want to estimate. The routing state of the network of Fig. 1 can be described as

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix}. \tag{2}$$

The impairment parameter values in vector $\mathbf{y}$ can be different for different applications and use cases, and they also depend on the model used for the approximation of the physical layer. For example, in this paper we assume $\mathbf{y}_M$ to be the inverse of the SNR of the established paths, and we want to estimate the inverse of the SNR of the new lightpath $\mathbf{y}_N$. From the SNR, and assuming known modulation format and FEC, we can calculate the pre-FEC BER and the BER. Assuming an OPM that can report on impairments, such as ASE, dispersion, or NLIs, we can formulate an estimation problem for each metric, solve it, and then examine the feasibility of each metric separately or combine the values to obtain a BER estimate. In other cases, such as hard- or soft-failure localization (which require the estimation of link vector $\mathbf{x}$ that includes the
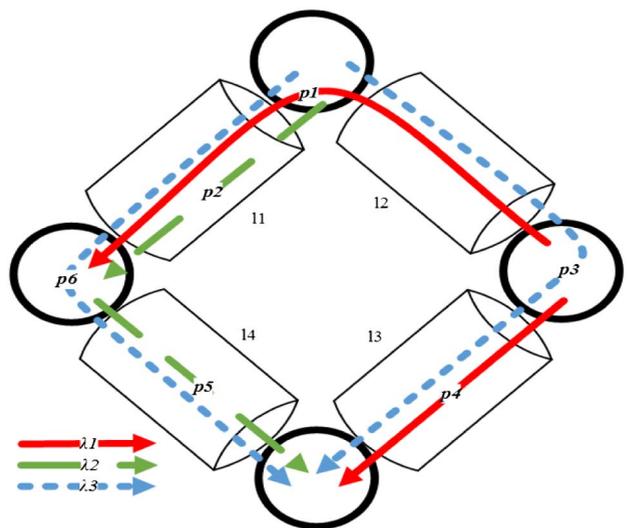

Fig. 1. Network with six lightpaths on three different wavelengths.

values per link), other parameters need to be considered. Soft failures are defined as events that progressively degrade the QoT. They cause subtle changes in performance and are often hard to detect. Most soft failures require complicated computations and information provided by OPMs to localize. In this case, the estimation framework can be used for specific impairment metrics that the OPMs provide or specifically tailored parameters, given that the parameters considered are additive per link.

Since the new lightpath p_6 contains links that are already in use (at a different spectrum) by other paths, it is possible to estimate its end-to-end value. To express this idea systematically, using the previous definitions, the end-to-end observations are expressed as $\mathbf{y}_M = R_M\mathbf{x}$, where the link parameters $\mathbf{x}$ are unknown, and it is not possible to exactly calculate them due to limited information from the available observations ($\mathbf{y}_m$). The objective is to estimate the end-to-end parameters $\mathbf{y}_N$, where $\mathbf{y}_N = R_N\mathbf{x}$. This was achieved in [12,13] using NK or NM and will be discussed in Subsections IV.B and IV.C.

Note that if we make the worst-case interference assumption (as in [12,13]) that all channels are simultaneously lighted, then all impairment parameters of a given link will have equal values regardless of the current wavelength utilization. Therefore, all established and new lightpaths crossing that link will exhibit equal QoT deterioration. In Section V we will describe our framework that accounts for the interference of the spectrum neighboring channels.

B. Network Kriging

NK [14] is a scheme comprised of linear prediction methods aiming to estimate network path characteristics based on a sample. It was applied to optical networks in [12,13]. It finds the best linear estimate (in terms of mean-square error) of $\mathbf{y}_N$, which is

$$\hat{\mathbf{y}}_N = R_N R_M^T (R_M R_M^T)^+ \mathbf{y}_M. \tag{3}$$

The symbol $(.)^+$ denotes a pseudo-inverse such as the Moore–Penrose inverse. The complexity of the algorithm is $O(|M|^3)$ [12], where $|M|$ is the number of established lightpaths.

C. Norm l_2 Minimization

This method can calculate nonlinear estimates of $\mathbf{y}_N$. It also has the advantage of being able to define constraints on the calculated solution. The respective problem is defined as in [12,13] as follows:

$$\min_{x,u}(\|\mathbf{u}\|_2^2 + \|\mathbf{x}\|_2^2)$$
$$\text{subject to } R_M\mathbf{x} + D_2\mathbf{u} = \mathbf{y}_M, \quad x \geq 0. \tag{4}$$

The symbol $\| \cdot \|_2^2$ denotes the Euclidean (l_2) norm. In Eq. (4), $\mathbf{u}$ is a regularization parameter and D_2 is a positive–definite diagonal matrix. D_2 determines whether each

row of $R_M\mathbf{x} \approx \mathbf{y}_M$ should be satisfied accurately or in the least-squares sense. Its diagonals should be typically small (10^{-4}) if the constraints should be satisfied reasonably accurately. This problem can be tackled using software libraries such as PDCO [17] that solve Eq. (4). The complexity of the algorithm was found in [12] to be $O(\alpha L^3)$, where α depends on the requested accuracy of the estimation and $L = |E|$ is the number of network links. The output of the algorithm is the estimate of the link level parameters $\mathbf{x}$. In the present paper we are interested in the estimation of the end-to-end parameter of the new lightpaths, which can be easily obtained from the relationship $\mathbf{y}_N = R_N\mathbf{x}$.

V. IA ESTIMATION

A. Modeling Neighboring Channel Interference

The definition of matrix R in the previous section depends only on the routing of the lightpaths and is thus able to convey information only on space dependencies, while it ignores dependencies in the spectrum domain. It records which lightpaths share links, but does not record, e.g., which ones use adjacent or distant channels and how much spectrum each of them utilizes. Consequently, the problem formulation in the previous section cannot accurately account for the interference neighboring channels (wavelength in WDM or spectrum bands in EONs) cause to each other, and it has to rely on a worst-case interference assumption as is done in [12,13]. It models lightpaths sharing the same link (of course at different channels) as having the same QoT value (the inverse SNR in our study), without taking into account the actual wavelength utilization, the position, and distance of the wavelengths used or the spectrum bands. To see this differently, in Eq. (1) and problem formulation (4) of Section IV, wavelength information plays no role, and the notion of a lightpath is indistinguishable from that of a path in the sense that, if a lightpath changes wavelength, the equations remain the same.

The model that we develop can be used for both WDM and EONs. It considers the spectrum that each lightpath occupies and this information is used to account for the interference between spectrum neighboring lightpaths. This is important because, by doing so, we can obtain a more realistic QoT estimation, avoiding the (worst-case) assumption that all channels are simultaneously lighted. This can result in significant regenerator savings, as we will show in the simulations. We will also show later that the model that we will develop can also be used to calculate the degradation that the new path will cause to the already established lightpaths.

To estimate the interference from neighboring channels, the initial problem statement has to be modified. We are given a network $G = (V, E)$ that is currently in a specific state. The network state is described by the routing matrix $A \in \{0,1\}^{|M| \times |E|}$ and the vector $\mathbf{w}$ of $|M|$ elements describing the wavelength used by each lightpath. To obtain a model that accounts for neighboring channel interference, we define an auxiliary *interference-aware* (IA) graph $G' = (V', E')$
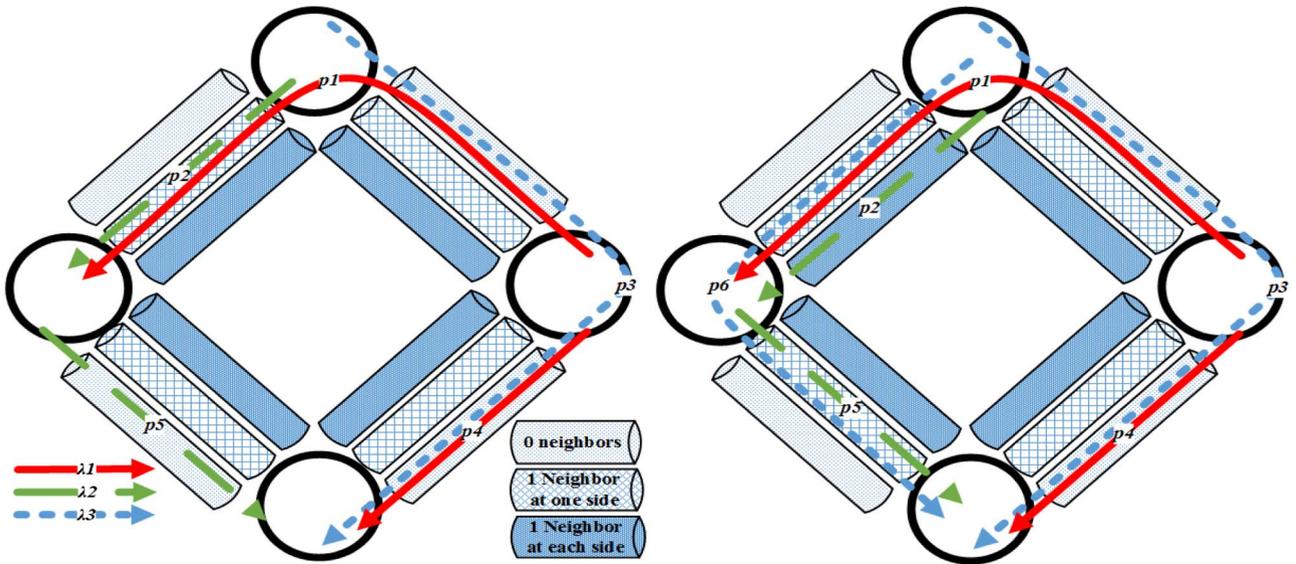
Fig. 2.   Auxiliary IA graph and its IA links and lightpaths (a) before and (b) after the establishment of lightpath p6. Notice lightpaths p2 and p5.

(see Fig. 2) and a corresponding routing matrix R_M, in the following way. The set V of nodes in G is the same as the set of nodes V' in the IA graph G', but each link l_i in E is replaced by a set of IA links in G', denoted by $l'_{i,j}$, with $j = 1, ..., q$, $q = 0.5(2^\gamma + 2^{\gamma/2})$, where γ denotes the *total* number of the closest neighbors that we assume contribute to interference, that is, $\gamma/2$ neighbors from each side. To capture the effect of the more distant neighbors, we can use an appropriate margin. Each IA link $l'_{i,j}$ represents the number and position of the active spectrum neighboring channels. Assuming equal launch power and baud rates, if two channels have the same neighbors, they will have almost equal interference values regardless of which side the neighbors are on. This assumption makes certain combinations of active neighboring channels be equivalent. Therefore, the number q of IA links, presented above, is calculated by subtracting from the total number of neighboring link combinations 2^γ, the number of equivalent IA links.

A lightpath in G' containing certain IA links $l'_{i,j}$ is called an IA lightpath. To define the routing matrix R_M of G' from the routing matrix A and the wavelength vector $\mathbf{w}$ of G, we take each lightpath (p, w) in G (each line of A and $\mathbf{w}$) and route it as a IA lightpath in G': if p passes through link l_i in G and has no occupied spectrum neighbors on l_i (i.e., no other active lightpath uses a wavelength on l_i that is γ-adjacent to w), the related IA lightpath is routed over link $l'_{i,0}$ in G'. If p has one active adjacent neighbor (on either side) on l_i in G, the IA lightpath passes through $l'_{i,1}$ in G', and so on, to finally capture the case where p has $\gamma/2$ active neighbors from each side, where the IA lightpath is routed over link $l'_{i,q-1}$ in G'. The respective pseudo-code is described in Algorithm 1.

To illustrate the graph and state transformation through an example, consider the optical network of Fig. 1 and assume that only the direct neighbor (from each side)

contributes significantly to the interference. In this case $\gamma = 2$, and there are $q = 3$ IA links in network G' for each link in Fig. 1. The five lightpaths $M = \{p_i\}$, $i = 1, 2, ..., 5$, are considered active on the network, using a set of three wavelengths λ_1, λ_2, and λ_3. We will describe the routing procedure for lightpath p_1. This lightpath contains links l_1 and l_2. At link l_1, lightpath p_1 has one neighbor from one side (namely, lightpath p_2 on wavelength λ_2), so the related IA lightpath is routed through $l'_{1,1}$ in G'[as seen in Fig. 2(a)]. At link l_2, lightpath p_1 has no neighbors from either side, so the IA lightpath is routed over link $l'_{2,0}$ of G'. If it had one neighbor from each side, it would have been routed over link $l'_{2,2}$. We follow the same procedure to route all the established lightpaths in G'.

The accuracy of the QoT estimation increases as the number of interfering sources (the number of neighbors γ) that we take into account increases. On the other hand, the number of IA links q increases exponentially with the number of neighbors considered. Since the accuracy of the estimation method depends on the number of lightpaths crossing a link, and the complexity of the estimation techniques (see Subsections IV.B and IV.C) depends on the number of links in the network, we would prefer to keep the number of additional IA links low. In practice, the assumption that the two closest neighbors from each side contribute most of the interference (i.e., $\gamma = 4$) is considered a sweet spot, as we demonstrate in the simulations. It requires only $q = 10$ IA links between two nodes, provides very good QoT estimation accuracy with relatively little information, and requires a very low margin to account for more distant neighbors. In our simulation experiments we also examine the performance for different values of γ, and observe the trade-offs between estimation accuracy and the required margin for distant neighbors, but the majority of simulations were performed using $\gamma = 4$.

Algorithm 1 Construction of the Auxiliary IA Graph

Input: Original routing matrix $A \in \{0,1\}^{|M| \times |E|}$, wavelengths per lightpath **w**, number of interfering neighbors γ

Output: IA lightpaths with their respective IA links (matrix R_M)

Procedure:

1: Initialize $q = 0.5(2^\gamma + 2^{\gamma/2})$ IA links to replace each link $\in E$
2: **for** lightpath $m \in M$
3:   **for** $l_i \in m$ {the links of lightpath m}
4:     Search in A and **w** for active $\gamma/2$ neighbors from each side
5:     Depending on number and type of active neighbors on link l_i, route m through $l'_{i,z}$ ($R_M[m, l'_{i,z}] = 1$), $z \in \{0, ..., q-1\}$
6:   **end for**
7: **end for**

By following the same methodology, it is also possible to consider a number of different baud rates or launch powers by inserting additional IA links in graph G' to represent the baud rate/launch power of each channel and that of its neighboring channels. In this case, the number of IA links is $q = (br + lp - 1)\{0.5((br + lp)^\gamma + (br + lp)^{\frac{\gamma}{2}})\}$, where br denotes the number of different baud rates and lp the number of different launch powers in the network. Note that for $br = 1$ and $lp = 1$, this equation is simplified to the one reported previously. As discussed above, the number of required IA links increases exponentially with γ and thus we need to keep the number of cases considered in our model low. This can be done by using relaxed grouping criteria. For example, assuming a network that has connections using two baud rates/powers, if two IA links differ in only the baud rate/power of the most distant neighbor, then we can assume that they exhibit similar impairment values. Baud rates that are close to each other (e.g., 28 and 32 Gbaud), can be grouped together for all their links. However, groupings will have an impact on the estimation error of the framework; a margin can be added so that we work on the safe side and always overestimate the QoT value. Taking a slightly different approach, calculations can be performed to infer the impairments of an equivalent IA link that uses another baud rate/launch power (e.g., use the Taylor representation of the GN model as a function of the baud rate).

The aforementioned graph transformation results in the conversion of the initial routing matrix G to a new IA matrix G', which keeps information not only for the routing of the lightpaths, but also for the spectrum neighboring lightpaths for each link of a connection and, thus, for the corresponding interference. As more IA links are added to the graph, the interference information is enriched with interference values that correspond to different neighbor attributes depending on the parameters considered (e.g., baud rate and/or power).

1. QoT Estimation Framework for EONs: The proposed IA estimation framework can also be used in EONs by changing the definition of the neighbor. In an EON, the spectrum is quantized in spectrum slots of F GHz ($F = 12.5$ GHz), and each lightpath is assigned a contiguous number of slots. Lightpaths whose spectra consist of variable numbers of slots are concurrently active in the network. Considering all possible combinations of spectrum utilization results in a high number (exponential) of additional IA links, and so we need to use some relaxed grouping criteria to limit that. In the following, we describe a set of such relaxed grouping criteria.

Algorithm 2 Estimating the QoT of a New Lightpath

Input: Original routing matrix $A \in \{0,1\}^{|M| \times |E|}$, wavelengths per lightpath **w**, monitored values $\mathbf{y}_M$, new lightpath $n(N = \{n\})$, number of interfering neighbors γ

Output: The QoT estimation for lightpath n ($\mathbf{y}_N$)

Procedure:

1: Initialize $q = 0.5(2^\gamma + 2^{\gamma/2})$ IA links to replace each link $\in E$
2: Use algorithm 1 with input A, **w**, and γ to generate R_M
3: **for** $l_i \in n$ {the links of lightpath n}
4:   Search in A and **w** for active $\gamma/2$ neighbors from each side
5:   Depending on number and type of active neighbors on link l_i we find the link $l'_{i,z}$ to route n, $z \in \{0, ..., q-1\}$
6:   **If** ($l'_{i,z}$ is used in R_M) **then** set $R_N[n, l'_{i,z}] = 1$
7:   **else** use an IA link with more neighbors
8: **End for**
9: Use Eq. (3) or (4) with R_M, R_N, and $\mathbf{y}_M$ to estimate $\mathbf{y}_N$

We introduce a parameter z to represent a measure of the width of an *elastic channel*. In our experiments, we set z to be the minimum number of slots that a lightpath may occupy in the network, but many other options are available, such as the mean or the maximum number of slots. We consider for a given lightpath the z slots from each side as the first neighbor and the z to $2z$ slots from each side as the second neighbor, and so on, until we include $\gamma/2$ neighbors from each side. For example, consider an EON with two baud rates, 32 and 56 Gbaud, that a 32 Gbaud lightpath may occupy three or four slots, and that a 56 Gbaud channel may occupy five, six, or seven slots. Note the additional slots act as a guardband to reduce interference from neighboring channels and increase the transmission reach of a lightpath, if needed. The first neighbor is considered to be up to $z = 3$ slots away (the minimum number of slots among all options), and the second neighbor is considered between three and six slots. This means, for example, that a first neighbor might be one slot away, and another might be two slots away. Since they are both considered first neighbors, they are both mapped to the same IA link, and are assumed to have equal impairment values, while in reality they do not. Another issue that affects the accuracy in the elastic scenario is that two lightpaths that use the same baud rate may occupy different numbers of slots, as mentioned previously. This means that they produce slightly different interference effects because their central frequencies have different distances from the neighboring channels. However, since they have the same baud rate, we will map them to the same IA link. All these issues add "noise"

in the estimation, and the accuracy is expected to worsen for every additional baud-rate/slot-allocation option that is available. Thus, we expect the accuracy of the estimation framework for EONs to be lower than that for WDM networks, as it is shown in our performance studies. To account for such inaccuracy, we need to use larger margins, compared to the case of WDM networks. Still, the margins required are much lower than those required under worst-case assumptions.

B. Estimating the QoT of a New Lightpath Before It Is Established

In this section we describe how to combine the IA graph concept with the estimation techniques described in the previous Section IV [Eq. (3) or (4)]. We denote by A the original routing matrix and the end-to-end observations by $\mathbf{y}_M$ (1/SNR in this study from which we can derive the BER). We then construct submatrices R_M and R_N that correspond to the routing matrix R of the IA graph, using Algorithm 2. In order to construct matrix R_M we use Algorithm 1 (previously described in Subsection V.A) without taking into consideration the new lightpath (lines 1–2 of Algorithm 2). The columns of the matrix R_M are the IA links that represent the number and position of the neighboring lightpaths. Consider Fig. 2(a), for example. The IA graph is constructed without considering the new lightpath $N = \{p_6\}$. The IA lightpaths $M = \{p_1, \ldots, p_5\}$ correspond to matrix R_M. The matrix R_N is in fact a single vector that represents the to-be-established (IA) lightpath for which we would like to estimate the QoT. The columns of this vector are again the IA links of the IA graph, representing the kind of neighbors the links of the new lightpath would have if it was inserted into the current network. To find that, we search the original routing matrix A and take into account wavelength utilization $\mathbf{w}$ to identify the neighboring lightpaths of each link of the new lightpath and build the respective vector using again a procedure similar that of Subsection V.A (lines 3–8 of Algorithm 2). The correlation method makes use of the fact that the IA links of the new lightpath are used by the established ones. However, there might arise cases in which some IA links of the new lightpath are not used by any other existing lightpath in R_M and, therefore, there is no information about the impairments of this IA link. Note that such a case could not arise when creating R_M (it is not covered in Algorithm 1). In this case (lines 6–7 of Algorithm 2), we route the new lightpath over an IA link with more neighbors so that the estimation of the new lightpath is conservative and safe (overestimated QoT). If there is more than one IA link with more neighbors, there are several parameters to consider when picking the one to use, such as the number of lightpaths crossing those IA links and the lengths of those lightpaths. In the case that there is no IA link with more neighbors, we can fall back to a worst-case assumption (all channels active) for that link. In a multi-baud-rate scenario, if an IA link of a certain baud rate is not used by any lightpath, we can assign the equivalent IA link of a higher baud rate, expected to have higher interference. When the matrices

R_M and R_N are created, NK or NM (see Section IV) can be used to calculate the QoT of the new lightpath.

C. Estimating the Degradation a New Lightpath Will Cause to Existing Lightpaths

The estimation framework we presented can also be used to estimate the QoT degradation that a new lightpath would cause to the existing ones if it was established. We will use the problem statement of Subsection V.B. Again, the routing matrix R_M is constructed without taking into consideration the new lightpath (line 1 of Algorithm 3) and $\mathbf{y}_M$ includes the respective observed values. However, in this case R_N represents the lightpaths N that will be affected by the new lightpath n if it is established into the network. To find which lightpaths are affected, we insert the new lightpath into the original routing matrix A (line 3 of Algorithm 3) and construct a new temporary matrix A' and a new wavelength vector $\mathbf{w}'$. By using those with Algorithm 1 (line 3 of Algorithm 3), we find a new IA routing matrix T_M. Then we compare the differences between matrices R_M and T_M to find which existing (IA) lightpaths have different kinds of neighbors (this is the set N) and will thus use different IA links (line 4 of Algorithm 3). These lightpaths are the ones that will be mostly (if we limit the number of considered interfering channels to γ) affected by the establishment of the new lightpath. We create for these IA lightpaths the related R_N matrix (we copy the related entries of T_M into R_N) and run the estimation algorithm of Eqs. (3) or (4).

Algorithm 3 Estimating the degradation a new lightpath will cause to existing lightpaths

---

Input: Original routing matrix $A \in \{0, 1\}^{|M| \times |E|}$, wavelengths per lightpath $\mathbf{w}$, monitored values $\mathbf{y}_M$, new lightpath n, number of interfering neighbors γ
Output: The QoT estimation of the affected existing lightpaths ($\mathbf{y}_N$)
Procedure:
1: Use algorithm 1 with input A, $\mathbf{w}$, and γ, to generate R_M
2: Add lightpath n to routing matrix A and $\mathbf{w}$ to form temporal routing matrix A' and temporal wavelengths per lightpath $\mathbf{w}'$
3: Use algorithm 1 with input A', $\mathbf{w}'$, and γ to generate T_M
4: Compare matrices R_M and T_M to find the set N of existing IA lightpaths that have modified neighbors and set as R_N the related lines of T_M
5: Use Eq. (3) or (4) with R_M, R_N, and $\mathbf{y}_M$ to estimate $\mathbf{y}_N$, the QoT of affected lightpaths N

Consider, for example, Figs. 2(a) and 2(b). Figure 2(a) depicts the matrix R_M, which represents the routing of the IA lightpaths $M = \{p_1, \ldots, p_5\}$ that are computed before the establishment of the new lightpath p_6. Figure 2(b) depicts the matrix T_M, which represents the routing of the IA lightpaths $\{p_1, \ldots, p_5, p_6\}$ after the establishment of lightpath p_6. If we compare the two IA graphs (matrices R_M and T_M), we can see that only p_2 and p_5 contain different IA links. Therefore, we set $N = \{p_2, p_3\}$ and create the related R_N

matrix. Using such input, our estimation framework evaluates whether establishing the new lightpath makes infeasible some existing ones or how much it degrades their performance.

D. Measurement Database

The proposed estimation framework can be enhanced by adding a database (DB) that stores past end-to-end measurements along with the IA lightpaths that the measurements correspond to. The DB can be updated whenever a new lightpath is established or when an estimation of a new lightpath is required (since certain operations of the estimation framework and the DB are shared, namely, the operations of Algorithm 1). The DB can also have a field to store the age of each measurement, and remove measurements as time passes to account for aging and other time-varying effects. The technical details concerning the integration of the DB with the control plane of the network are outside the scope of this paper, and are currently under definition within the framework of ORCHESTRA [6]. Note that such a DB is filled up very quickly, since we store information for the IA lightpaths, and, therefore, a single lightpath may occupy multiple database entries. That is, whenever a new lightpath is established, it triggers multiple DB entries, since it affects all its close neighbors and results in the rerouting of the related IA lightpaths over G'. To see this differently, each affected lightpath for which we estimate its feasibility (see Subsection V.C and description of the R_N matrix), can give us a new measurement (a new entry in the DB) once the new lightpath is established. Note that some IA lightpaths in the DB may provide duplicate information for certain IA links. Such information can be used to reduce the uncertainty of the measurements, due, e.g., to low-accuracy DSP monitoring algorithms or aging effects. When the DB is sparse, the use of *probe* lightpaths, which do not carry useful data but are deliberately established for estimation purposes, can help improve estimation accuracy. Such lightpaths can be carefully selected to provide the DB with the most new information possible and/or reduce the number of IA lightpaths required in the DB for accurate estimation. Such extensions are left for future work.

In the developed QoT estimation scheme, every time the algorithms of Subsections V.B and V.C are used, the matrix R_M is concatenated with the (disjoint) rows of the DB that is a superset of the established lightpaths, thus improving estimation accuracy.

VI. PERFORMANCE EVALUATION

To evaluate the performance of the QoT estimation scheme, we carried out simulation experiments. In particular, we evaluated the framework under two network scenarios, namely, for (i) a WDM fixed-grid network with single and dual baud-rate transmissions, and (ii) an EON using two baud rates and several spectrum occupation options. We used two network topologies: (i) NSFNET

(Fig. 3), with 14 nodes and 22 bidirectional links, with link lengths taken to be the Euclidean distances multiplied by 1.2, and (ii) a modified European SPARKLE topology (Fig. 4) with 49 nodes and 63 bidirectional links. The modifications introduced to SPARKLE were to set the minimum link length to 80 km, to satisfy the requirement of the GN model for having span loss greater than 7 dB, and to neglect double links and take only the shortest link between certain nodes.

Links were assumed to consist of single-mode fiber (SMF) with attenuation coefficient 0.25 dB/km, dispersion parameter 16.7 ps/nm/km, and nonlinear coefficient 1.3 1/W/km. The span length in both network topologies was set to at most 100 km, and erbium-doped fiber amplifiers that fully compensate span losses with noise figure of 6 dB were assumed. We assumed that there are no dispersion compensation modules, and the DSP at the receivers performs ideal electronic dispersion compensation and MIMO equalization.

For the WDM network simulations, we consider a 50 GHz grid, 80 wavelengths, and two different transmission scenarios: the "WDM-1 baud rate" scenario assumes 100 Gbps polarization multiplexed quadrature phase shift keying (PM-QPSK) with 28 Gbaud, while the "WDM-2 baud rates" scenario assumes two different baud rates existing at the same time in the network: 28 and 32 Gbaud, which are represented by different IA links in the graph transformation. For the EON, we assumed
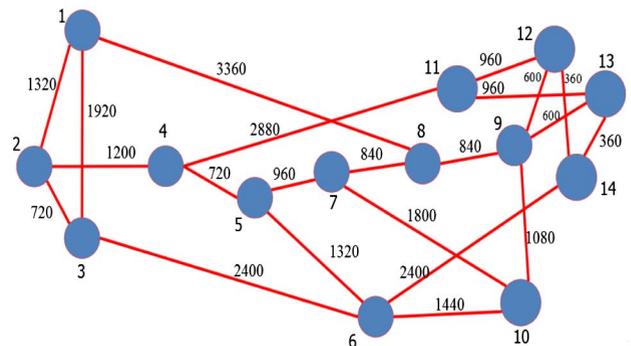

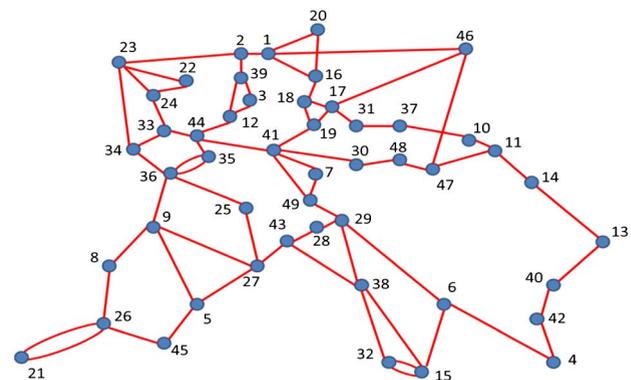Fig. 3.   NSFNET network with the link lengths in kilometers.

Fig. 4.   Extended European SPARKLE Network.

PM-QPSK in 32 and 56 Gbaud (100 and 200 Gbps, respectively, both of them single carrier). We assumed two transmission scenarios: "EON-ScenA" assumes 32 Gbaud connections served in three slots (37.5 GHz), and 56 Gbaud served in five and six slots (62.5 and 75 GHz). "EON-ScenB" assumes three and four slots for the 32 Gbaud and five, six, and seven slots for the 56 GBaud. The total number of neighbors considered was $\gamma = 4$, unless otherwise stated, for both the WDM and EON. For the EON scenarios we used as the width of an elastic channel z and the minimum number of required slots by all lightpath transmission options (so, $z = 3$ slots in both ScenA and ScenB).

Requests for the establishment of lightpaths arrive according to a Poisson process with rate λ. Destinations are uniformly distributed over all nodes, and connection durations are exponentially distributed with mean $1/\mu$. Upon the arrival of a request for a lightpath from a specific source to a specific destination, a shortest path routing with a first-fit wavelength assignment algorithm is used for RWA. In the EON scenarios, a RSA algorithm is used with the same principles as the aforementioned RWA. Note that the proposed estimation mechanism would also work for any other RWA/RSA algorithm, since any such algorithm would benefit from a more accurate estimation of the physical conditions. Actually, the benefits are more pronounced for longer paths, and, thus, any RWA/RSA algorithm that uses alternative paths to the shortest ones considered here will exhibit even higher benefits.

The time required for our algorithms to construct the IA graph and to estimate the pre-FEC BER from the SNR (note that the SNR in our case considers the ASE and NLIs) in MATLAB was on average 0.9, 1.8, or 8.2 s for 300, 600, of 1000 known IA lightpaths (including the ones in the database), respectively, on an Intel Core-i5 4210U (2 cores at 2.7 GHz).

A. Estimating the QoT of a Lightpath Before It Is Established

We examine the accuracy of our estimation framework using the mean squared error (MSE) for the pre-FEC BER estimation as a function of the number of IA lightpaths available in the DB. In particular, we calculate the log(pre-FEC BER) estimation of a new lightpath and then compare it with the log(pre-FEC BER) that is obtained using the GN model. NK and NM were found to provide almost identical results, with NM providing slightly more accurate estimations, and, thus, it is used in all the following results. Figure 5 shows the MSE of our estimation framework for the WDM scenario in the NSFNET network. Note that the accuracy is better than the one presented in our previous work [16] because we improved the way that we handle cases where the correlation information is limited. We observe that when the number of IA lightpaths in the DB is low, the MSE is high. This was expected, since in this case there is not enough information for most of the IA links. In most cases the estimation accuracy is between 0.01 and 0.5 dB. The difference was larger in single-link
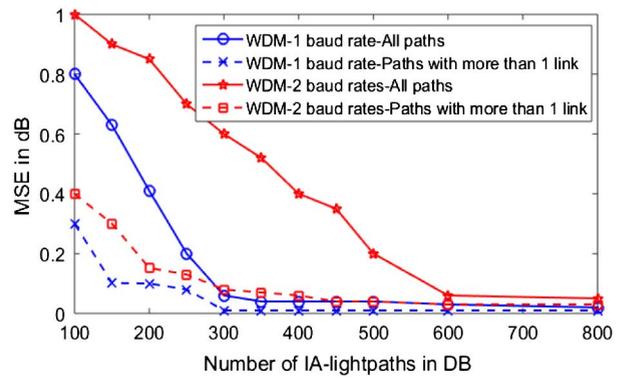

Fig. 5.   Pre-FEC BER estimation MSE for two WDM scenarios.

lightpaths, which usually have very low pre-FEC BER, making such inaccuracy insignificant in practice. To show this, in Fig. 5 we also plot the MSE for lightpaths consisting of at least two links, which is observed to be much lower. To achieve a negligible MSE (less than 0.05) in the WDM NSFNET network, the DB must have around 400 IA lightpaths (which translates to approximately 160 lightpaths in the original network) for the single baud rate, and around 700 (180 lightpaths in the original network) for the dual-baud-rate transmission scenarios. Note that the database is filled up quickly, since establishing a new lightpath creates interference and thus reroutes several IA lightpaths, which in turn generates new entries in the DB (see the discussion in Subsection V.D).

Another interesting metric, apart from the MSE graphed in Fig. 5, is the maximum underestimation (MU) of the pre-FEC BER, since this gives us the QoT margin that has to be used to work on the safe side (never underestimate the QoT). For the single-baud-rate scenario, the MU was 0.1 dB for 1000 IA lightpaths, while for the dual-baud-rate scenario, the same MU required around 1800 IA lightpaths.

Figure 6 depicts the accuracy in terms of MSE for a single-baud-rate WDM scenario by accounting for the interference of a different number of neighboring channels, and in particular when considering $\gamma = 2$, $\gamma = 4$, and
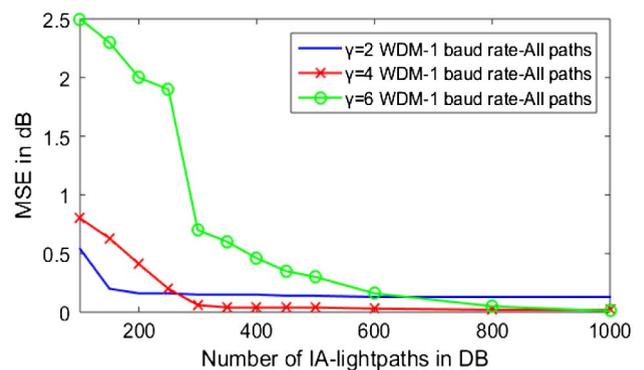

Fig. 6.   Pre-FEC BER estimation MSE for different number of neighbors γ.

$\gamma = 6$ neighboring channels in total from both sides. We notice that for $\gamma = 2$ (i.e., three IA links per original link), the MSE of the pre-FEC BER is relatively low even for a small number of IA lightpaths. However, after around 200 IA lightpaths, the MSE remains constant at 0.13 dB and the MU is 0.2 dB. The good accuracy for few IA lightpaths stems from the low number of IA links required in this scenario, making even a small number of DB entries enough to provide the required information for those links. However, this also causes the saturation of the estimation accuracy after 200 IA lightpaths. The absence of enough detail in modeling the interference limits the estimation framework from additional correlation sources and the estimation accuracy it can achieve. When $\gamma = 6$ (then 36 IA links are needed), we notice that the MSE is initially very high and 1000 IA lightpaths are required for the MSE to drop to 0.01 (MU of 0.08 dB), which is quite lower than the accuracy reached for $\gamma = 2$. This was expected, since the high number of IA links requires many IA lightpaths to provide the relevant correlated values. The low number of IA lightpaths contains a, respectively, low number of IA links, and, as we saw in Subsection IV.B when there is no information for an IA link, we use an equivalent IA link with more neighbors. This decision is taken to always be on the safe side, sacrificing some of the estimation accuracy. The computation time required for $\gamma = 2$ is almost the same with $\gamma = 4$, while the time required for $\gamma = 6$ is close to that required for $\gamma = 4$ (approximately 1 s slower for 1000 lightpaths). The good estimation accuracy and low number of IA lightpaths that are required for $\gamma = 4$ makes it the favorable choice for many applications, since the DB can be filled quickly and updated promptly in response to network changes such as aging and other time-varying effects.

Figure 7 depicts the accuracy of the estimation framework for the EON for two scenarios: EON-ScenA and EON-ScenB. We notice that the MSE is generally higher than the WDM scenario, due to the inaccuracies introduced by the use of relaxed neighboring grouping criteria, as discussed in Subsection V.A. Also, the MSE of Scenario A is slightly worse than Scenario B's, because of the extra set of slots for each baud rate that infer higher inaccuracies. To achieve an MSE of less than 0.05 for Scenarios A and B, 2000 IA lightpaths are required, indicating that the relaxed grouping criteria used (discussed in Subsection V.A)
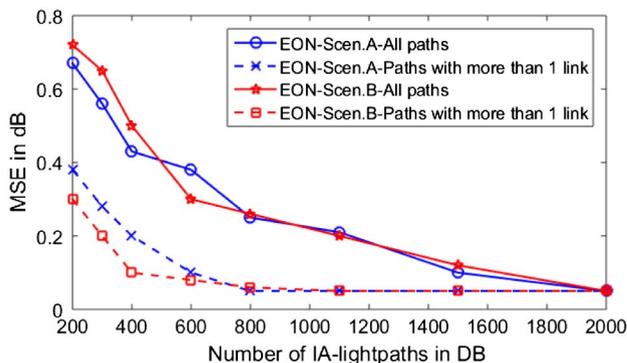

Fig. 7.   Pre-FEC BER estimation MSE for the EON scenario.

are successful. The MU of the BER for this number of IA lightpaths was 0.3 dB for lightpaths that contain more than one link and 1.2 dB for one-link lightpaths. If we take into account only the lightpaths whose QoT is close to the blocking threshold (in our case it is set at −2 dB), then the MU is 0.2 dB. That is because the slope of the erfc function that is used to calculate the BER [11] is small when the BER is high. As a result, small deviations in the SNR do not translate into respective BER deviations. This MU needs to be used as a margin so that our scheme always overestimates the QoT, and, as expected, the margins are higher than in the case of the WDM network.

B. Regeneration Savings

As discussed in the introduction, obtaining better QoT estimates translates into various benefits for the network operator. In this section, we calculate the regenerator savings that our estimation framework can achieve when compared to provisioning lightpaths using the worst-case interference assumption, that is, assuming that all channels are lighted simultaneously. We also calculate the number of the regenerators that would be required if the actual BER values of the lightpaths were known before we established them (referred to as *perfect estimation*). In our experiments, this is done by establishing the lightpath and using the GN model to calculate its pre-FEC BER. Regenerators are placed at a node when a lightpath has unacceptable pre-FEC BER. If the BER is unacceptable after only one link (e.g., 200 G transmissions over a long link), then we assume that a regenerator is placed in the middle of the link (this is actually required only in a single link in the NSFNET topology). The pre-FEC BER threshold was set at −2 dB for the 28 and 32 Gbaud, assuming 20% FEC, and at −1.88 dB for the 56 Gbaud, assuming 25% FEC. Driven by the accuracy results presented in the previous section (MU metric) we use a 0.1 dB margin to account for the estimation error in our framework for the WDM case. This means that if our estimation framework calculates the pre-FEC BER of a WDM lightpath to be equal to −2.1 dB, then a regenerator will be required even though the related threshold is set at −2 dB. Similarly, for the EON we use a 0.2 dB margin. Our proposed framework is also used to estimate the degradation a new lightpath causes to existing lightpaths: the RWA/RSA algorithm initially chooses the first available wavelength/set of slots and then picks the next in the case that the chosen one turns unacceptable an existing lightpath.

Figure 8(a) shows the savings in the maximum number of regenerators in the WDM network and one baud rate for the NSFNET topology. Savings were observed to be up to 47% (at the lowest load) for our proposed framework when compared to the worst-case interference assumption. This was achieved because we obtain up to 0.04 lower pre-FEC BER (1.4 dB) when compared to the worst-case assumption that overestimates the BER. Our framework uses less than 5% more regenerators than the perfect estimation case, which assumes that we accurately know the (pre-FEC) BER, proving the great accuracy of our estimation. As
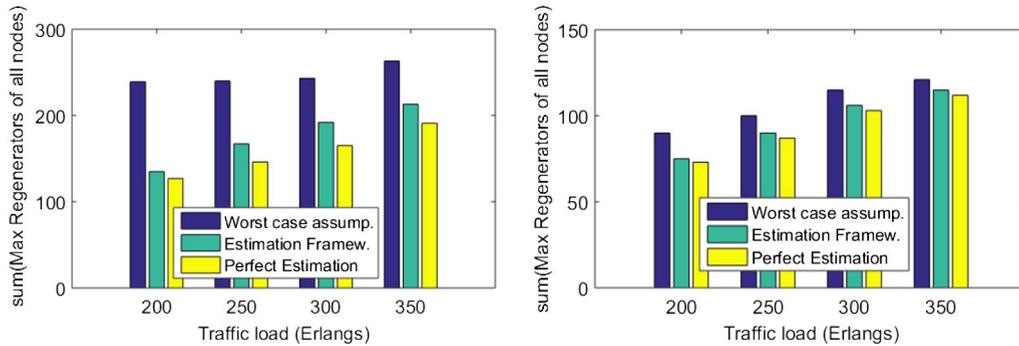
Fig. 8.   (a) Sum of maximum number of regenerators of all nodes for the WDM scenario in NSFNET and (b) in SPARKLE.

the network load increases, more lightpaths are concurrently active and thus the QoT of the lightpaths degrades and becomes close to that in the worst-case scenario. As a result, as the load increases, the required regenerators for our estimation framework and the perfect estimation case become close to those required for the worst case, and the gains are reduced. Similar results were observed for the maximum savings in the number of regenerators required in a single node (not shown in graphs due to space limitations). Our estimation framework required 21 regenerators in a single node, as opposed to 41 in the worst-case scenario.

In Fig. 8(b) we present the savings for the SPARKLE network for the same WDM scenario. The savings in the total number of regenerators required for all the nodes were up to 15%. The maximum number of regenerators in a single node was 11 for our estimation framework and 19 for the worst-case interference scenario. The reason for lower savings in the SPARKLE topology is the much shorter links of that topology when compared to NSFNET. As a result, the QoT of most lightpaths does not degrade up to the regeneration threshold that was set. Note that the savings achieved by the perfect estimation case are also limited, with our scheme being less than 5% worse. Therefore, the framework works quite well, even in a topology where the possible gains are limited.

Figure 9(a) shows the regeneration savings for the EON-ScenA in the NSFNET network. In this case, the worst-case scenario is set as having all neighboring channels on, using the lowest baud rate and the minimum number

of slots for the grid. The savings observed for the total number of regenerators required were up to 36%. Also, the maximum regenerators in a single node were 47 with the worst-case scheme, as opposed to 34 for our estimation framework. Figure 9(b) shows the results for the EON for the SPARKLE network. The respective savings for the SPARKLE network were observed to be around 12%. The maximum number of regenerators in a single node was 12 for our estimation framework and 21 for the worst-case scenario. We notice that the percent savings are smaller in the EON case than in the WDM case in both network topologies. The reason is that the relaxed neighboring criteria that are used for the EON case result in higher MSE, as discussed in Subsection VI.A, and we use a higher margin (0.2 dB) to account for this. The increased margin is expected to require more regenerators, since it is added to the regeneration threshold of the ideal case. This results in at most 7% more regenerators than the perfect estimation case in which we accurately know the BER, which is slightly higher than the 5% that was observed in the WDM case. We also notice that the difference in the number of regenerators between the ideal and the worst case is less than that of the WDM scenario. This is attributed to the fact that the EON scenario employs 56 Gbaud connections, which are more susceptible to impairments than the 32 Gbaud, and, therefore, more regenerators are required even in the perfect estimation. So, in conclusion, despite the higher margins used for EONs, our framework's accuracy is still quite good, and yields significant savings in regenerators.
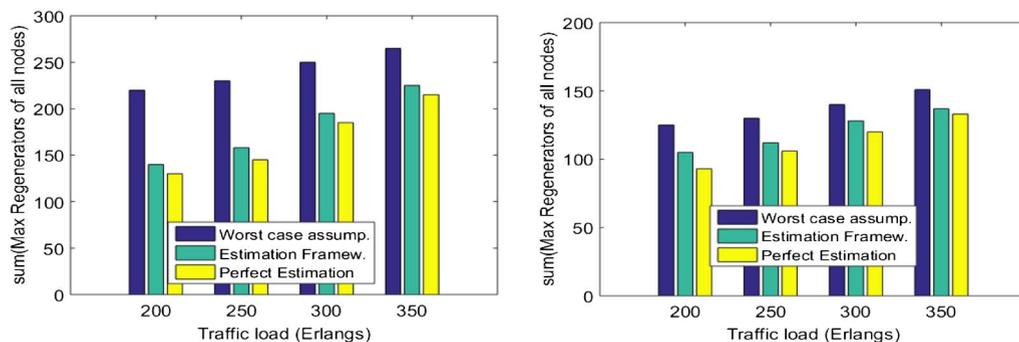

Fig. 9.   (a) Sum of maximum number of regenerators of all nodes for the elastic scenario in NSFNET and (b) in SPARKLE.

Note that in the above experiments we focused on cases of zero blocking. Since regenerators function as wavelength converters, a higher number of regenerators would improve the wavelength blocking performance when blocking would appear. However, when the network reaches such a limit, the operator would typically take more effective measures than adding regenerators to relax the wavelength continuity constraint, such as upgrading part of the network with higher rate transponders, or employing additional fibers. So our focus here is on cost savings, neglecting the blocking performance of the network close to its end of life. Also note that the comparison presented above focuses only on interference, assuming good knowledge of network parameters and constant performance of equipment. Provisioning the lightpaths for worst-case interference but also for end-of-life aging (system margins), and considering estimation inaccuracies (design margins), as is typically done (see discussion in Sections I and II), would result in even more waste of resources. So, the savings that the proposed QoT estimation framework can achieve are even higher, considering that it can reduce both design and system margins based on current network conditions feedback.

VII. Conclusion

We presented a novel QoT estimation framework for fixed-grid WDM and EONs that can be used when provisioning new lightpaths. Using a proposed graph transformation, we took into account both space and spectrum information when calculating the QoT of a new lightpath before it is established, and the degradation it causes to the existing ones if it was established. The estimations obtained were shown to be quite accurate, and result in lower margins in provisioning new lightpaths. The accuracy of the framework was slightly worse in EON as opposed to WDM networks, due to the higher number of possible interference states that were modeled through relaxed grouping criteria. The high accuracy of the proposed estimation framework leads to significant regeneration savings compared to provisioning based on the worst-case interference assumption that all channels are simultaneously on, as indicated by our performance studies. Future research efforts include the application of the proposed estimation framework under uncertainty, for soft-failure localization, and the use of probing lightpaths to improve estimation accuracy.

Acknowledgment

References

[1] J. L. Auge, "Can we use flexible transponders to reduce margins?" in *Optical Fiber Communication Conf.*, Anaheim, CA, Mar. 2013.

[2] A. Mitra, S. Kar, and A. Lord, "Effect of frequency granularity and link margin at 100G and beyond flexgrid optical networks," in *Nat. Conf. on Communications*, Kanpur, India, Feb. 2014.

[3] Y. Pointurier, "Design of low-margin optical networks," in *Optical Fiber Communication Conf.*, Anaheim, CA, Mar. 2016.

[4] M. Angelou, Y. Pointurier, D. Careglio, S. Spadaro, and I. Tomkos, "Optimized monitor placement for accurate QoT assessment in core optical networks," *J. Opt. Commun. Netw.*, vol. 4, no. 1, pp. 15–24, Jan. 2012.

[5] Infonetics Research, "In telecom optics market, 100G transceiver growth suppressed until 2016," May 2015 [Online]. Available: http://www.infonetics.com/pr/2015/2H14-Telecom-Optics-Mkt-Highlights.asp.

[6] K. Christodoulopoulos, P. Kokkinos, A. Di Giglio, A. Pagano, N. Argyris, C. Spatharakis, S. Dris, H. Avramopoulos, J. C. Antona, C. Delezoide, P. Jennevé, J. Pesic, Y. Pointurier, N. Sambo, F. Cugini, P. Castoldi, G. Bernini, G. Carrozzo, and E. Varvarigos, "ORCHESTRA—Optical performance monitoring enabling flexible networking," in *17th Int. Conf. on Transparent Optical Networks (ICTON)*, Budapest, Hungary, July 2015, pp. 1–4.

[7] O. Gerstel, M. Jinno, A. Lord, and S. J. B. Yoo, "Elastic optical networking: A new dawn for the optical layer?" *IEEE Commun. Mag.*, vol. 50, no. 2, pp. s12–s20, Feb. 2012.

[8] F. Cugini, F. Paolucci, F. Fresi, G. Meloni, N. Sambo, L. Potí, A. D'Errico, and P. Castoldi, "Towards plug-and-play software-defined elastic optical networks," *J. Lightwave Technol.*, vol. 34, no. 6, pp. 1494–1500, Mar. 2016.

[9] A. Bononi, P. Serena, A. Morea, and G. Picchi, "Regeneration savings in flexible optical networks with a new load-aware reach maximization," *Opt. Switching Netw.*, vol. 19, no. 3, pp. 123–134, Jan. 2016.

[10] P. Poggiolini, "The GN model of non-linear propagation in uncompensated coherent optical systems," *J. Lightwave Technol.*, vol. 30, no. 24, pp. 3857–3879, Dec. 2012.

[11] P. Poggiolini, G. Bosco, A. Carena, V. Curri, Y. Jiang, and F. Forghieri, "A detailed analytical derivation of the GN model of non-linear interference in coherent optical transmission systems," arXiv:1209.0394, 2012.

[12] N. Sambo, Y. Pointurier, F. Cugini, L. Valcarenghi, P. Castoldi, and I. Tomkos, "Lightpath establishment assisted by offline QoT estimation in transparent optical networks," *J. Opt. Commun. Netw.*, vol. 2, no. 11, pp. 928–937, Nov. 2010.

[13] Y. Pointurier, M. Coates, and M. Rabbat, "Cross-layer monitoring in transparent optical networks," *IEEE J. Opt. Commun. Netw.*, vol. 3, no. 3, pp. 189–198, Mar. 2011.

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

[15] K. Christodoulopoulos, K. Manousakis, M. Angelou, and E. Varvarigos, "Considering physical layer impairments in offline RWA," *IEEE Netw. Mag.*, vol. 23, no. 3, pp. 26–33, May–June 2009.

[16] I. Sartzetakis, K. Christodoulopoulos, C. Tsekrekos, D. Syvridis, and E. Varvarigos, "Estimating QoT of unestablished lightpaths," in *Optical Fiber Communication Conf.*, Anaheim, CA, Mar. 2016.

[17] M. Saunders, "Primal-dual interior method for convex objectives," [Online]. Available: http://web.stanford.edu/group/SOL/software/pdco/.