

Stochastic RWA and Lightpath Rerouting in WDM Networks

Maryam Daryalal, Merve Bodur

Department of Mechanical and Industrial Engineering, University of Toronto
m.daryalal@mail.utoronto.ca, bodur@mie.utoronto.ca

In a telecommunication network, Routing and Wavelength Assignment (RWA) is the problem of finding lightpaths for incoming connection requests. When facing a dynamic traffic, greedy assignment of lightpaths to incoming requests based on predefined deterministic policies leads to a fragmented network that cannot make use of its full capacity due to stranded bandwidth. At this point service providers try to recover the capacity via a defragmentation process. We study this setting from two perspectives: (i) while granting the connection requests via the RWA problem and (ii) during the defragmentation process by lightpath rerouting. For both problems, we present the first two-stage stochastic integer programming model incorporating incoming request uncertainty to maximize the expected grade of service. We develop a decomposition-based solution approach, which uses various relaxations of the problem and a newly developed problem-specific cut family. Simulation of two-stage policies for a variety of instances in a rolling-horizon framework of 52 stages shows that our stochastic models provide high-quality solutions compared to traditionally used deterministic ones. Specifically, the proposed provisioning policies yield improvements of up to 19% in overall grade of service and 20% in spectrum saving, while the stochastic lightpath rerouting policies grant up to 36% more requests using up to just 4% more bandwidth spectrum.

Key words: Two-stage Stochastic Programming, Routing and Wavelength Assignment, Lightpath Rerouting, Optical Networks

1. Introduction

The global internet traffic is steadily growing both in the number of users and the amount of transmitted data: it is predicted that by 2023, the total number of connected devices can reach as high as three times the world population ([Cisco 2020](#)) and the amount of Internet traffic is multiplied by a factor of one thousand every twenty years ([Waldman 2018](#)). This ever increasing traffic is in constant combat with the capacity of the information carriers, specifically in the context of *optical networks*, the backbone of today's telecommunication systems ([Majumdar 2018](#)). While the cable operators have so far enjoyed the vast

capacity provided by upgrading their existing submarine cables over the years, the industry is approaching the theoretical capacity of the optical fibers (Lavallée 2017), called the Shannon limit (Shannon 1948, Essiambre et al. 2008). Combined with the increase in the number of users and the high bit-rate traffic requests such as online streaming services, this leads to the predicted *capacity crunch* problem in telecommunication networks (Ellis and Doran 2013, Waldman 2018, Jara et al. 2020, Zhou et al. 2020). Although capacity expansion by means of deploying fiber cables and extra equipment is a possible solution, ample investment and considerable amount of time required for the network expansion make it a strategic decision (Keiser 1999). The second and short-term option is efficient use of existing capacity in optical networks that can postpone reaching the Shannon limit. Traditional studies aiming at enhancing this usage for a deployed network often assume a deterministic setting, which is in stark contrast with practice. In this paper, we illustrate the importance of incorporating traffic uncertainty into decision making, and introduce the first two-stage stochastic integer programming (2SP) model for the classical resource allocation problem in optical networks.

As the dominant transmission system in the telecommunications industry, *wavelength division multiplexing* (WDM) networks are the subject of numerous studies aiming at improving the efficiency of optical networks. Therein, as the name of optical networks suggests, the transmission medium is an optical fiber carrying light. The key to their success is the ability to transmit multiple signals on different wavelengths through a single optical fiber, i.e., a WDM technology multiplexes and demultiplexes (joins and splits) these signals at the nodes along a connection. In WDM networks, the pair of an optical fiber link and a wavelength is called a *wavelink*. A connection is then set up via a *lightpath*, which is defined as a specific frequency occupied on a sequence of optical fibers, i.e., a path in the network together with a wavelength. The *routing and wavelength assignment* (RWA) problem is the provisioning problem in optical networks. Given a deployed WDM network and a set of connection requests as inputs, the RWA problem seeks to optimally assign one lightpath to each request considering two main conditions: (i) every lightpath needs to use the same wavelength throughout its path, and (ii) two lightpaths with the same wavelength cannot share a link. WDM networks make it possible to efficiently scale the network (Ciena 2020), thus the RWA problem plays a vital role in the telecommunications industry that relies mostly on such transmission systems.

The RWA problem is quite challenging with a rich literature that has mainly focused on heuristic solutions. Exact solution approaches that can live up to the scale of demand the industry is facing today have only recently been proposed ([Jaumard and Daryalal 2017](#)). The objective function of the RWA problem can be defined based on different perspectives ([Krishnaswamy and Sivarajan 2001](#)). One might consider minimizing the resource usage while serving a set of connections that have to be granted, which corresponds to minimization of the number of used wavelinks (the minRWA problem). Another possible objective function is to maximize the throughput measured as the number of granted requests (the maxRWA problem), which is equivalent to maximization of the grade of service (GoS), i.e., minimizing the blocking rate. For a brief overview of WDM networks, the RWA problem and the related literature, we refer the readers to [Kuri \(2003\)](#) and [Daryalal \(2016\)](#).

Although in many studies on the RWA problem, traffic request is assumed to be static (and given in advance), practical assumption in telecommunications industry, e.g. for a service provider, is a dynamic uncertain traffic. Mainly due to the immense size of traffic requests in optical networks, stochasticity of the future demand has not been well leveraged at operational levels for deployed networks. Existing works that consider stochasticity of traffic are mostly in the area of capacity planning and upgrade. [Kennington et al. \(2003\)](#) develop a robust optimization model to determine the number of necessary equipment for a dense WDM network considering the future uncertain demand. For a survivable WDM network, [Leung and Grover \(2005\)](#) propose a 2SP model that minimizes its deployment cost and possible augmentations in the future. More closely, [Lodha et al. \(2007\)](#) study the light-trail design problem and formulate it as a 2SP model to decide on the set of light-trails to be configured. Light-trails are a different technology in WDM networks, where between every pair of nodes supporting the light-trail architecture, connections can share a wavelength. Considering a scenario tree with three scenarios and with the goal of minimizing the pre-planning costs, [Kronberger et al. \(2011\)](#) propose a three-stage stochastic programming model for dealing with the cases where future demand is strongly overestimated in the first stage. To determine the capital expenditure needed for capacity upgrade of a network, [Aparicio-Pardo et al. \(2012\)](#) present a robust optimization approach that considers stochastic demand with normal distribution.

In the context of lightpath assignment and the RWA problem, although the stochasticity of traffic has been a discussion topic since the early studies ([Ozdaglar and Bertsekas](#)

2003), explicitly incorporating the uncertainty of the future traffic in granting decisions has been mostly avoided. Instead, the focus has been on the design of quasi-static/heuristic methods that tend to result in fewer congested links in the future. To address this shortcoming, we propose a stochastic maxRWA problem that, while assigning lightpaths, takes the future incoming traffic into account as a value function. To the best of our knowledge, this is the first study that proposes a stochastic model for lightpath assignment in a WDM network. Through numerical experiments, we show that the use of the deterministic maxRWA problem in a stochastic setting results in sub-optimal solutions in terms of the number of granted requests. In all studied instances designed based on standard backbone networks, the stochastic model can admit considerably more requests, up to 31 more connections, a 19% improvement over the deterministic model. We note that as a result of traffic grooming and SONET/SDH technology, a single wavelength can carry multiple low bit-rate traffic connections with transmission rates as low as 51.84Mbps (Zhu and Mukherjee 2003); and today’s metro and long-haul networks commonly operate at 100Gbps per wavelength (Inniss and Rubenstein 2016), and very recently at a remarkable bandwidth of 800Gbps (Alliance 2020, Xenos 2020). Therefore, even one more accepted request (assigned lightpath) is quite significant as it is equivalent to granting thousands of low bit-rate connections.

In a dynamic uncertain setting that traffic arrives and leaves in multiple time periods, a service provider sets up and tears down the lightpaths as needed. This type of traffic leads to network fragmentation (Ji et al. 2014), meaning that there exists capacity stranded over the network, but it cannot be retrieved to build a lightpath for certain requests, resulting in high blocking rates. After it is identified that a network is highly fragmented, a *network defragmentation* event is triggered. The state of the art in network defragmentation is to recover the stranded bandwidth by minimizing resource requirements modeled as a deterministic minRWA problem (Jaumard et al. 2019). It obtains a set of lightpaths assigned to the granted connections on the network (target provisioning). A defragmentation process then involves rerouting of the lightpaths in a seamless migration fashion, i.e., in a way that a rerouted lightpath is set up before the original one is torn down. Noting that minRWA ignores the future behavior of the traffic, as the second objective of this study, we question the effectiveness of minRWA solution as the target provisioning in a defragmentation process. It is presumed that a target with fewer used resources performs better in the future

for the incoming requests since available wavelinks are maximized (Jaumard and Daryalal 2017). In this paper, we propose to incorporate the natural uncertainty of incoming traffic in the design of the defragmentation target. We formulate the problem of finding this provisioning as a 2SP model and show its effectiveness in improving the performance of the network after multiple stages.

Contributions. Main contributions of this paper are summarized as follows:

- We introduce the *stochastic maxRWA* (SmaxRWA) problem, a stochastic variant of the maxRWA problem incorporating traffic uncertainty, and formulate it as a 2SP model.
- Rather than using minRWA as the target provisioning of defragmentation problem, we propose the *stochastic lightpath rerouting* (SmaxLR) problem, for obtaining a provisioning that maximizes the expected number of granted requests in the future. We formulate the SmaxLR problem as a 2SP model that shares many similarities with the SmaxRWA model.
- Relying on the empirical evidence of its strength, we propose to solve a certain relaxation of the stochastic models to derive high-quality solutions efficiently. For solving the proposed relaxation, we design a decomposition method enhanced by a new problem-specific family of cuts that significantly reduces the number of optimality cuts needed for its convergence.
- Through extensive numerical experiments, we show the notable value of considering stochasticity in provisioning and defragmentation problems in WDM networks, as well as the strength of our new family of cuts.

The rest of the paper is organized as follows. In Section 2 we briefly review the preliminaries and then introduce two novel stochastic problems to be used in provisioning and defragmentation processes. In Section 3 we propose 2SP models for these two problems. In Section 4 we design a decomposition framework for solving the proposed models and enhance it by introducing a new family of cuts. In Section 5 we provide a thorough numerical study that evaluates the value of incorporating stochasticity in provisioning and defragmentation, as well as the performance of our solution method. Finally, in Section 6 we provide concluding remarks and future research directions.

2. Background and Motivation

Consider a WDM network with a given physical topology and a set of available wavelengths. The traffic on this network is assumed to be asymmetric, and is characterized by the number of connection requests between two nodes of the physical network. A connection request between two nodes is served by a lightpath, consisting of a path between those nodes together with a certain wavelength. Note that, in the absence of wavelength conversion, shown to be practically ineffective in increasing GoS (Zhang et al. 2013), the same wavelength is used throughout the entire path associated with a lightpath. This is known as *wavelength continuity* property. As another important property, two lightpaths cannot share links with the same wavelength, otherwise they incur a *wavelength conflict*.




Figure 1 A WDM network

Figure 1 provides an example of the physical topology of a WDM network, along with a provisioning for two requests from node 1 to 12, one request from node 12 to node 1, one request from node 6 to node 1, and one request from node 6 to node 2. In this figure, every link consists of bidirectional fibers. There are 3 wavelengths available on this network, represented by solid green, dashed blue and dotted red. While dashed and dotted lightpaths respect wavelength continuity with no wavelength conflicts, solid lightpaths share the same wavelength on a fiber link, hence they have a wavelength conflict. A provisioning is called valid if its lightpaths do not violate wavelength continuity and they do not have any wavelength conflicts. Thus, the provisioning illustrated in Figure 1(b) is not valid. If the request from node 6 to node 1 was granted on the same path with a different wavelength, then the provisioning would have been valid. For a set of incoming connection requests, the maxRWA problem finds a valid provisioning that maximizes the GoS, while minRWA minimizes the number of links used in the provisioning.

2.1. Dynamic Traffic

In a WDM network, dynamic traffic involves adds and drops of connection requests over time, through lightpath assignments and release of wavelinks, respectively. Figure 2 shows an example of a dynamic traffic in a WDM network with two available wavelengths. Incoming traffic is granted until period $t = 4$ (represented by solid green and dashed blue). At $t = 5$ a request from node 1 to 6 is blocked. At period $t = 6$, a connection is leaving the network, i.e., a connection drop occurs, then granting of the connections continues.




Figure 2 Dynamic traffic in a WDM network with two wavelengths

The consecutive adds and drops over time can result in a fragmented network that is probable to higher rate of request rejections, while there exists capacity over the network. The process of migrating a fragmented network to a provisioning with recovered bandwidth is called defragmentation. Figure 3 presents a fragmented provisioning in the WDM network of Figure 2 at $t = 7$, as well as a defragmented one that grants the same requests while using a fewer number of wavelinks. In a fragmented network, lightpaths with longer lengths are employed in the provisioning, leading to higher wavelink usage. For example, in Figures 3(a) and 3(b) the average lightpath length is 2.5 and 1.5 links, respectively. The provisioning in Figure 3(b) has recovered the stranded bandwidth by using shorter lightpaths. For instance, the defragmented network can grant four more requests from node 2 to 5, while the fragmented one has the capacity of only three more requests between the same nodes.

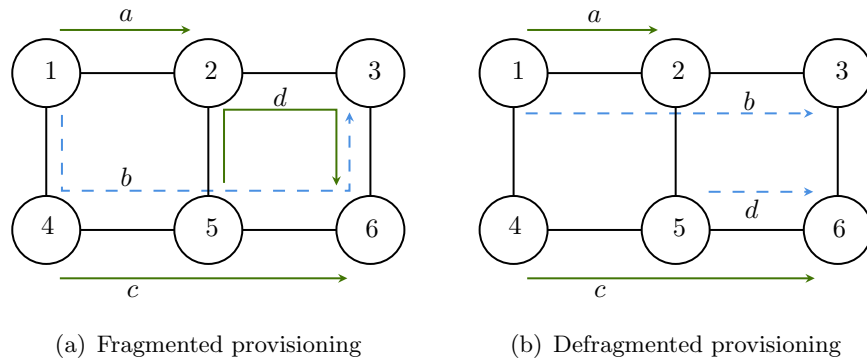


Figure 3 Defragmentation example

2.2. Motivation

A deployed and activated WDM network starts with full capacity. Over time, it grants the incoming connection requests by establishing lightpaths and frees the resources once a connection leaves the network. At certain points, we want to reroute the current existing lightpaths such that the network has a higher chance of granting the uncertain future demand. From this description, stochastic traffic impacts the performance of a network from two aspects: provisioning and defragmentation.

For the provisioning problem, disregarding the future fosters greedy assignment of lightpaths. It aims for locally maximum number of granted requests, even if it leads to selection of long lightpaths that take up considerable amount of resources just to accept one more connection at the current stage. By considering the future traffic one can make more informed lightpath assignment decisions, both in terms of acceptance or rejection of requests, and the characteristics of lightpaths that are assigned to the granted ones. Consequently, the goal is not to find locally optimal solutions at every stage, rather it is to optimize the overall performance in the longer term.

As for the defragmentation problem, there are two main steps in the process: (i) design of a defragmented provisioning target for the current set of connection requests, and (ii) conversion of the current provisioning to the target, namely performing a *network migration*. Common provisioning target in the network defragmentation problem is to minimize the usage of resources (e.g., minimization of the number of used links), with the assumption that it leads to a more efficient answer to traffic requests, leaving a larger number of wavelinks to serve the future traffic. Note that, over a multi-period planning horizon, merely minimizing the number of used resources as a local decision does not guarantee

the best performance in the future, and despite the momentary decrease in the wavelink usage, the network can quickly become fragmented once again after defragmentation. In fact, to the best of our knowledge, the impact of the defragmentation outcome on future GoS has not been studied before. We argue that a deterministic approach does not have the means to recognize and free up critical wavelinks that can contribute to granting more future requests, thus such a provisioning, although defragmented, might not be able to improve the GoS overall.

Following the discussions above, we suggest to look one step ahead while making the lightpath assignment/rerouting decisions, by solving two stochastic programming problems that correspond to stochastic versions of the maxRWA and lightpath rerouting problems, i.e., SmaxRWA and SmaxLR, respectively. Throughout the paper, we make the following set of assumptions:

- (A1) Lightpath assignment decisions are made periodically, specifically all the requests arrived in a time interval are granted/rejected at the end of the interval.
- (A2) The target of the defragmentation process is reachable, i.e., we can reroute the lightpaths without service interruption.

The first assumption is justified by the common practice of service providers, where the granting decisions are made for a batch of requests in order to maximize GoS and incremental planning has to be performed in given periods to determine the capacity upgrades needed for addressing the growth of the traffic. It is indeed a common practice in the context of the dynamic traffic literature (Wu et al. 2012, Chen et al. 2015, Jaumard and Daryalal 2017). The second assumption is nonrestrictive, since by leasing one unit of wavelength (Leung and Grover 2005), any provisioning can be migrated to a new one without service disruption.

In the following sections, we formally define SmaxRWA and SmaxLR and formulate them as 2SP models.

3. Problem Descriptions and Formulations

The stochastic maxRWA and lightpath rerouting problems, namely SmaxRWA and SmaxLR, both assign lightpaths to a known set of connection requests, such that the expected GoS is maximized. Therefore, their mathematical models are quite similar, with the exception of the definition of node pair set for which they are building a provisioning.

SmaxRWA attempts to choose from the set of new connection requests, while SmaxLR has to assign a (potentially new) lightpath to every connection already on the network. In other words, defragmentation (SmaxLR) is performed at the beginning of a time period before the arrival of a new batch of connection requests, while SmaxRWA is solved at the end of a period, when the new traffic has arrived, and future traffic is taken into consideration in both. In the following, we first present the 2SP formulation of SmaxRWA in Section 3.1, then in Section 3.2 we make the necessary modifications for SmaxLR.

3.1. The Stochastic maxRWA Problem

Consider a WDM network. Denote by $G = (\mathcal{V}, \mathcal{L})$ the multigraph that represents the physical topology of this WDM network, with $v \in \mathcal{V}$ and $\ell \in \mathcal{L}$ being a node and fiber link of the physical layer, respectively. $\omega \in \Omega$ is an available wavelength on the network. To grant a connection request between a node pair $(s, d) \in \mathcal{V} \times \mathcal{V}$, a lightpath $\lambda = (p, \omega)$ with path p between s and d over a wavelength ω is set up. In the following, the uncertainty in the future traffic is characterized by $r^{(s,d)}(\boldsymbol{\xi})$, denoting the number of requested lightpaths between node pair (s, d) , where $\boldsymbol{\xi}$ is the underlying random vector with known probability distribution \mathbb{P} and support Ξ . Suppose that a new batch of connection requests has arrived on a potentially non-empty network. Available wavelinks are denoted by $\mathcal{L}(\omega)$, the set of links without any granted requests on wavelength $\omega \in \Omega$.

The immediate (first-stage) decisions SmaxRWA has to make are the lightpath assignments for the new incoming traffic. Therefore the first-stage problem of SmaxRWA is a maxRWA problem with a slightly different objective function, given a set of available wavelinks. Let $\mathcal{SD}_N^1 \subseteq \{(s, d) \in \mathcal{V} \times \mathcal{V}\}$ be the set of node pairs with at least one new connection request in the most recent batch of arrivals, (N stands for ‘‘new’’). A node pair (s, d) can have multiple connection requests, denoted by $\bar{r}^{(s,d)} \in \mathbb{Z}_+$. $x_{\ell, \omega}^{(s,d)}$ is a binary decision variable equal to 1 if link ℓ is used in a lightpath with wavelength ω for a connection request between node pair (s, d) . Given \mathcal{SD}_N^1 , SmaxRWA is formulated as follows:

$$\max \quad \sum_{(s,d) \in \mathcal{SD}_N^1} \sum_{\omega \in \Omega} \sum_{\substack{\ell \in \\ \delta^+(s) \cap \mathcal{L}(\omega)}} x_{\ell, \omega}^{(s,d)} + \mathbb{E}_{\boldsymbol{\xi}} [\mathcal{Q}(x, \boldsymbol{\xi}, \mathcal{SD}_N^1)] \quad (1a)$$

$$\text{s.t. } x \in \mathcal{F}^{\text{RWA}}(\mathcal{SD}_N^1), \quad (1b)$$

where $\mathcal{Q}(\cdot)$ is the value function that measures GoS in the second stage and $\mathcal{F}^{\text{RWA}}(\mathcal{SD}_N^1)$ is the feasible set of the SmaxRWA problem. Objective function (1a) maximizes the expected

number of granted requests, considering both the new connection requests and future traffic. $\mathcal{Q}(\cdot)$ and $\mathcal{F}^{\text{RWA}}(\cdot)$ are also used in the modeling of SmaxLR, so in the following we define them for a generic set of node pairs to avoid repetition. For a given \mathcal{SD} , $\mathcal{F}^{\text{RWA}}(\mathcal{SD})$ is defined by the set of following constraints:

$$\sum_{(s,d) \in \mathcal{SD}} x_{\ell,\omega}^{(s,d)} \leq 1 \quad \omega \in \Omega, \ell \in \mathcal{L}(\omega) \quad (2a)$$

$$\sum_{\delta^+(v) \cap \mathcal{L}(\omega)}^{\ell \in} x_{\ell,\omega}^{(s,d)} = \sum_{\delta^-(v) \cap \mathcal{L}(\omega)}^{\ell \in} x_{\ell,\omega}^{(s,d)} \quad \omega \in \Omega, (s,d) \in \mathcal{SD}, v \in \mathcal{V} \setminus \{s,d\} \quad (2b)$$

$$\sum_{\omega \in \Omega} \sum_{\delta^-(s) \cap \mathcal{L}(\omega)}^{\ell \in} x_{\ell,\omega}^{(s,d)} = \sum_{\omega \in \Omega} \sum_{\delta^+(d) \cap \mathcal{L}(\omega)}^{\ell \in} x_{\ell,\omega}^{(s,d)} = 0 \quad (s,d) \in \mathcal{SD} \quad (2c)$$

$$\sum_{\omega \in \Omega} \sum_{\delta^+(s) \cap \mathcal{L}(\omega)}^{\ell \in} x_{\ell,\omega}^{(s,d)} \leq \bar{r}^{(s,d)} \quad (s,d) \in \mathcal{SD} \quad (2d)$$

$$x_{\ell,\omega}^{(s,d)} \in \{0, 1\} \quad (s,d) \in \mathcal{SD}, \omega \in \Omega, \ell \in \mathcal{L}(\omega), \quad (2e)$$

where $\delta^{+(-)}(v)$ is the set of links leaving (entering) node $v \in \mathcal{V}$. Constraints (2a) ensure that every link ℓ can be used for at most one path in any wavelength for the new traffic. Constraints (2b) establish the flow conservation in the chosen lightpaths. They make sure that for a given node pair (s,d) , the same number of wavelinks over the same wavelength is leaving and entering the nodes of the network for granting a request of (s,d) . Constraints (2c) eliminate the cycles by preventing the selection of incoming/outgoing link to/from source/destination of a request, and together with constraints (2b) set up a proper lightpath for a granted request. Constraints (2d) bound the number of lightpath assignments to a node pair by its demand. Constraints (2e) are the integrality and bound constraints.

With the first-stage provisioning fixed, recourse decisions are made by solving another maxRWA problem for the future traffic $r^{(s,d)}(\xi)$ on the available resources. $y_{\ell,\omega}^{(s,d)}$ is a second-stage decision variable equal to 1 if for a future connection request between $(s,d) \in \mathcal{SD}^2(\xi) = \{(s,d) \in \mathcal{V} \times \mathcal{V} : r^{(s,d)}(\xi) > 0\}$, a wavelink with link ℓ and wavelength ω is used. $z^{(s,d)}$ is an integer (auxiliary) recourse decision variable measuring the number of granted requests between node pair (s,d) . For a first-stage solution \hat{x} and a future traffic scenario ξ , the recourse problem is defined as:

$$\begin{aligned} \mathcal{Q}(x, \xi, \mathcal{SD}) = \\ \max \sum_{(s,d) \in \mathcal{SD}^2(\xi)} z^{(s,d)} \end{aligned} \quad (3a)$$

$$\text{s.t.} \quad \sum_{(s,d) \in \mathcal{SD}^2(\xi)} y_{\ell,\omega}^{(s,d)} \leq 1 - \sum_{(s,d) \in \mathcal{SD}} x_{\ell,\omega}^{(s,d)} \quad \omega \in \Omega, \ell \in \mathcal{L}(\omega) \quad (3b)$$

$$\sum_{\delta^+(v) \cap \mathcal{L}(\omega)}^{\ell \in} y_{\ell,\omega}^{(s,d)} = \sum_{\delta^-(v) \cap \mathcal{L}(\omega)}^{\ell \in} y_{\ell,\omega}^{(s,d)} \quad \omega \in \Omega, (s,d) \in \mathcal{SD}^2(\xi), v \in \mathcal{V} \setminus \{s,d\} \quad (3c)$$

$$\sum_{\omega \in \Omega} \sum_{\delta^-(s) \cap \mathcal{L}(\omega)}^{\ell \in} y_{\ell,\omega}^{(s,d)} = \sum_{\omega \in \Omega} \sum_{\delta^+(d) \cap \mathcal{L}(\omega)}^{\ell \in} y_{\ell,\omega}^{(s,d)} = 0 \quad (s,d) \in \mathcal{SD}^2(\xi) \quad (3d)$$

$$z^{(s,d)} \leq r^{(s,d)}(\xi) \quad (s,d) \in \mathcal{SD}^2(\xi) \quad (3e)$$

$$z^{(s,d)} = \sum_{\omega \in \Omega} \sum_{\delta^+(s) \cap \mathcal{L}(\omega)}^{\ell \in} y_{\ell,\omega}^{(s,d)} \quad (s,d) \in \mathcal{SD}^2(\xi) \quad (3f)$$

$$y_{\ell,\omega}^{(s,d)} \in \{0,1\} \quad (s,d) \in \mathcal{SD}^2(\xi), \omega \in \Omega, \ell \in \mathcal{L}(\omega) \quad (3g)$$

$$z^{(s,d)} \in \mathbb{Z}_+ \quad (s,d) \in \mathcal{SD}^2(\xi). \quad (3h)$$

Objective function (3a) maximizes the number of granted requests. If a wavelink is used in the first-stage provisioning, Constraints (3b) prevent it to be used in future lightpath assignments. Constraints (3c)-(3e), are the counterparts of Constraints (2b)-(2d) for the second-stage provisioning, ensuring that a proper lightpath is selected for each granted request. Constraints (3f) calculate the number of granted connections for node pairs, for each realization of future traffic. Constraints (3g)-(3h) are the integrality and bound constraints in the second stage.

3.2. Stochastic Lightpath Rerouting Problem

In the defragmentation problem, when lightpath rerouting decisions are made, thanks to assumption (A2) it is irrelevant which wavelinks have been currently in use for the current provisioning. Therefore, the first-stage decisions simply choose the wavelinks in a new provisioning for the current set of granted requests on the network. In other words, in the first-stage problem we have $\mathcal{L}(\omega) = \mathcal{L}, \omega \in \Omega$. After the defragmentation procedure is complete and the uncertainty is revealed a maxRWA problem is solved to grant the new connection requests. Considering that, we propose a 2SP model for the SmaxLR problem, where in the first stage we make the rerouting decisions and in the second stage, with the occupied wavelinks now as given, we solve a maxRWA problem for the future traffic. Note that, the definition of the second-stage problem is the same as the recourse problem of SmaxRWA.

Let \mathcal{SD}_C^1 be the set of node pairs with an existing connection on the network (C stands for “current”). With the definitions and notations given in Section 3.1, the 2SP formulation for SmaxLR problem is:

$$\max \mathbb{E}_\xi[\mathcal{Q}(x, \xi, \mathcal{SD}_C^1)] \quad (4a)$$

$$\text{s.t. } x \in \mathcal{F}^{\text{RWA}}(\mathcal{SD}_C^1) \quad (4b)$$

$$\sum_{\omega \in \Omega} \sum_{\ell \in \delta^+(s)} x_{\ell, \omega}^{(s,d)} = \bar{r}^{(s,d)} \quad (s, d) \in \mathcal{SD}_C^1. \quad (4c)$$

As was mentioned earlier, SmaxRWA and SmaxLR formulations share many commonalities, such as constraints (4b) and their recourse problem. Differently in SmaxLR, every connection already on the network needs to be granted in the new defragmented provisioning, which is modeled by constraints (4c), and the objective function (4a) measures the expected GoS for the future traffic. Note that constraints (3e) are redundant for the SmaxLR model due to constraints (4c).

In the next section, we discuss how the 2SP formulation of SmaxRWA and SmaxLR problems can be solved efficiently.

4. Solution Methodology

In this section, we use *sample average approximation* (SAA), to obtain a deterministic equivalent (DE) form of models (1) and (4) for a given sample and discuss how to solve it by means of decomposition. Then we introduce a new family of cuts that enhance the performance of the decomposition algorithm.

4.1. Sample Average Approximation

The expectation term in the objective functions of SmaxRWA and SmaxLR models includes integration over a (high dimensional) random vector ξ . To overcome the difficulty of maximizing such a function, we use SAA, a technique common in solving such problems (see for e.g., [Shapiro et al. \(2014\)](#)). Given a sample of future traffic scenarios, SAA replaces the expectation in the objective function with the sample average. Assume that $\hat{\Xi} \subseteq \Xi$ is a set of independent and identically distributed (i.i.d.) observations drawn from \mathbb{P} using Monte Carlo sampling. For SmaxRWA, SAA gives the following approximation of (1):

$$\max \sum_{(s,d) \in \mathcal{SD}_N^1} \sum_{\omega \in \Omega} \sum_{\ell \in \delta^+(s) \cap \mathcal{L}(\omega)} x_{\ell, \omega}^{(s,d)} + \frac{1}{|\hat{\Xi}|} \sum_{\xi \in \hat{\Xi}} \eta_\xi \quad (5a)$$

$$\text{s.t. } x \in \mathcal{F}^{\text{RWA}}(\mathcal{SD}_N^1) \tag{5b}$$

$$\eta_\xi \leq \mathcal{Q}(x, \xi, \mathcal{SD}_N^1) \quad \xi \in \hat{\Xi} \tag{5c}$$

$$\eta_\xi \geq 0 \quad \xi \in \hat{\Xi}, \tag{5d}$$

where η_ξ is a continuous decision variable representing the recourse value function under scenario ξ . It is shown that as the sample size increases, the solution to (5) converges to that of the original problem (under certain mild conditions, as discussed in Shapiro (1996)). By creating copies of the second-stage decision variables and constraints for each scenario in the sample, we can reformulate the SAA problem (5) as a mixed-integer program, called the *extensive form* (given in Appendix, Section A). The extensive form often does not scale well with the number of scenarios, thus we try to solve it by a decomposition method.

The presence of integer decision variables in the second stage makes (5) quite challenging. The logic-based Benders decomposition (Hooker and Ottosson 2003), in particular the integer L-shaped method as its special case (Laporte and Louveaux 1993, Angulo et al. 2016), is a common solution method for solving the DE obtained from applying SAA on stochastic integer programs. However, such methods are often not efficient. As was mentioned in Section 3.1, our recourse function is itself a maxRWA problem on available wavelinks. In the literature of the RWA problem with asymmetrical traffic, it has been observed that the linear programming (LP) relaxation of its commonly used IP formulation tends to give a high-quality upper bound (Jaumard et al. 2007, Christodoulopoulos et al. 2010). Based on this observation (which is also confirmed by our extensive experiments in Section 5), we consider a relaxation of our DE where second-stage decisions are relaxed, i.e., assumed to be continuous. We call this relaxation IP-LP, in contrast to IP-IP as the original exact IP model. An important advantage of IP-LP is that, despite being a relaxation, it provides feasible first-stage decisions, thus can be directly implemented in practice. It can also provide a lower bound by an out-of-sample evaluation of its first-stage solution.

4.2. Decomposition Algorithm

IP-LP is a mixed-binary linear program with binary first-stage and continuous recourse decision variables, for which Bender decomposition (Benders 2005) is an established solution method. It decomposes the problem into a master problem and a subproblem that

decide on the first- and second-stage variables, respectively. For 2SPs, the latter further decomposes by scenario, thus leads to computational improvement. Recourse value function $\mathcal{Q}(\cdot)$ is iteratively approximated by solving the subproblems that return feasibility and optimality cuts, until no more cuts are found or a given optimality gap is achieved. Our second-stage feasibility space is nonempty for feasible first-stage decisions, since granting none of the future requests is always a feasible solution. Therefore, regardless of the \hat{x} values, as long as they are feasible, second stage remains feasible, i.e., we have relatively complete recourse, thus optimality cuts are solely needed. The Benders master problem (the so-called multi-cut version) for IP-LP is:

$$\text{MASTER}(\hat{\Xi}) : \max \sum_{(s,d) \in \mathcal{SD}} \sum_{\omega \in \Omega} \sum_{\substack{\ell \in \\ \delta^+(s) \cap \mathcal{L}(\omega)}} x_{\ell, \omega}^{(s,d)} + \frac{1}{|\hat{\Xi}|} \sum_{\xi \in \hat{\Xi}} \eta_{\xi} \quad (6a)$$

$$\text{s.t. } (2a) - (2e), (5d) \quad (6b)$$

$$(\eta_{\xi}, x) \in \Gamma(\xi) \quad \xi \in \hat{\Xi}, \quad (6c)$$

where $\Gamma(\xi)$ is the set described by optimality cuts, which we refer to as x -cuts, for scenario ξ . For a given scenario $\xi \in \hat{\Xi}$ and a first-stage solution \hat{x} , the Benders subproblem is:

$\text{SUB}(\hat{x}, \xi) :$

$$\max \sum_{(s,d) \in \mathcal{SD}^2(\xi)} z^{(s,d)} \quad (7a)$$

$$\text{s.t. } \sum_{(s,d) \in \mathcal{SD}^2(\xi)} y_{\ell, \omega}^{(s,d)} \leq 1 - \sum_{(s,d) \in \mathcal{SD}_N^1} \hat{x}_{\ell, \omega}^{(s,d)} \quad \omega \in \Omega, \ell \in \mathcal{L}(\omega) \quad (7b)$$

$$\sum_{\substack{\ell \in \\ \delta^+(v) \cap \mathcal{L}(\omega)}} y_{\ell, \omega}^{(s,d)} = \sum_{\substack{\ell \in \\ \delta^-(v) \cap \mathcal{L}(\omega)}} y_{\ell, \omega}^{(s,d)} \quad \omega \in \Omega, (s, d) \in \mathcal{SD}^2(\xi), v \in V \setminus \{s, d\} \quad (7c)$$

$$\sum_{\omega \in \Omega} \sum_{\substack{\ell \in \\ \delta^-(s) \cap \mathcal{L}(\omega)}} y_{\ell, \omega}^{(s,d)} = \sum_{\omega \in \Omega} \sum_{\substack{\ell \in \\ \delta^+(d) \cap \mathcal{L}(\omega)}} y_{\ell, \omega}^{(s,d)} = 0 \quad (s, d) \in \mathcal{SD}^2(\xi) \quad (7d)$$

$$\sum_{\omega \in \Omega} \sum_{\substack{\ell \in \\ \delta^+(s) \cap \mathcal{L}(\omega)}} y_{\ell, \omega}^{(s,d)} \leq r^{(s,d)}(\xi) \quad (s, d) \in \mathcal{SD}^2(\xi) \quad (7e)$$

$$z^{(s,d)} = \sum_{\omega \in \Omega} \sum_{\substack{\ell \in \\ \delta^+(s) \cap \mathcal{L}(\omega)}} y_{\ell, \omega}^{(s,d)} \quad (s, d) \in \mathcal{SD}^2(\xi) \quad (7f)$$

$$y_{\ell, \omega}^{(s,d)} \geq 0 \quad (s, d) \in \mathcal{SD}^2(\xi), \omega \in \Omega, \ell \in \mathcal{L}(\omega). \quad (7g)$$

Let $(\hat{x}, \hat{\eta})$ be an optimal solution of the master problem, and \hat{z} contain the optimal z variable values of the subproblem $\text{SUB}(\hat{x}, \xi)$ for scenario $\xi \in \hat{\Xi}$. If $\hat{\eta}_\xi > \sum_{(s,d) \in \mathcal{SD}^2(\xi)} \hat{z}^{(s,d)}$, an optimality cut (x -cut) in the form of

$$\eta_\xi \leq \sum_{\omega \in \Omega} \sum_{\ell \in \mathcal{L}(\omega)} \hat{\pi}_{\ell, \omega}^{(7b)} \left(1 - \sum_{(s,d) \in \mathcal{SD}} x_{\ell, \omega}^{(s,d)} \right) + \sum_{(s,d) \in \mathcal{SD}^2(\xi)} \hat{\pi}_{(s,d)}^{(7e)} r^{(s,d)}(\xi)$$

is added to the description of $\Gamma(\xi)$, where $\hat{\pi}^{(7b)}$ and $\hat{\pi}^{(7e)}$ are the optimal dual solutions associated with constraints (7b) and (7e), respectively.

SAA problem and optimality cuts for SmaxLR are very similar to the ones for SmaxRWA. The only modifications (as in their 2SP formulations) are the lack of first-stage costs, and respectively using node pair set \mathcal{SD}_C^1 and constraints (4c) instead of \mathcal{SD}_N^1 and (2d).

In Section 5, we evaluate the performance of the described Benders decomposition algorithm and show that, although the relaxation returns high-quality solutions, solving it with the above mentioned standard framework (even with some well-known enhancements such as implementing it in a branch-and-cut framework) is still computationally difficult. In the following we attempt to overcome this by introducing a new family of problem-specific cuts that improves the scalability of our decomposition algorithm.

4.3. A New Family of Cuts

The discussions in this section apply to SmaxRWA. For SmaxLR, the same results can be obtained via modifications discussed in the previous subsection. We define for each $\ell \in \mathcal{L}$ a new decision variable β_ℓ , measuring the total number of times link ℓ is used for serving traffic requests in the first stage, and add the constraints

$$\sum_{\omega \in \Omega} \sum_{(s,d) \in \mathcal{SD}_N^1} x_{\ell, \omega}^{(s,d)} = \beta_\ell \quad \ell \in \mathcal{L} \quad (8a)$$

$$(\eta_\xi, \beta) \in \Upsilon(\xi) \quad \xi \in \hat{\Xi} \quad (8b)$$

$$\beta_\ell \geq 0 \quad \ell \in \mathcal{L}, \quad (8c)$$

to the master problem, resulting in a problem denoted by $\overline{\text{MASTER}}(\hat{\Xi})$. In constraints (8b), $\Upsilon(\xi)$ is the set described by the inequalities called β -cuts of the following form:

$$\eta_\xi \leq \sum_{\ell \in \mathcal{L}} \hat{\pi}_\ell^{(9b)} (|\mathcal{L}(\omega)| - \beta_\ell) + \sum_{(s,d) \in \mathcal{SD}^2(\xi)} \hat{\pi}_{(s,d)}^{(9d)} r^{(s,d)}(\xi) + \sum_{\omega \in \Omega} \sum_{\ell \in \mathcal{L}(\omega)} \hat{\pi}_{\ell, \omega}^{(9e)},$$

where $\hat{\pi}^{(9b)}, \hat{\pi}^{(9d)}, \hat{\pi}^{(9e)}$ are the optimal dual solutions obtained by solving the following auxiliary subproblem:

$$\overline{\text{SUB}}(\hat{\beta}, \xi) : \max \sum_{(s,d) \in \mathcal{SD}^2(\xi)} z^{(s,d)} \quad (9a)$$

$$\text{s.t.} \sum_{\omega \in \Omega} \sum_{(s,d) \in \mathcal{SD}^2(\xi)} y_{\ell, \omega}^{(s,d)} \leq (|\mathcal{L}(\omega)| - \hat{\beta}_\ell) \quad \ell \in \mathcal{L} \quad (9b)$$

$$(7c) - (7d) \quad (9c)$$

$$\sum_{\omega \in \Omega} \sum_{\substack{\ell \in \\ \delta^+(s) \cap \mathcal{L}(\omega)}} y_{\ell, \omega}^{(s,d)} \leq r^{(s,d)}(\xi) \quad (s, d) \in \mathcal{SD}^2(\xi) \quad (9d)$$

$$\sum_{(s,d) \in \mathcal{SD}^2(\xi)} y_{\ell, \omega}^{(s,d)} \leq 1 \quad \omega \in \Omega, \ell \in \mathcal{L}(\omega) \quad (9e)$$

$$(7f) - (7g). \quad (9f)$$

Note that, $\overline{\text{SUB}}(\hat{\beta}, \xi)$ is a relaxation of $\text{SUB}(\hat{x}, \xi)$ where constraints (9b) are obtained by integrating over constraints (7b). They do not prevent the wavelength conflict, hence $\overline{\text{MASTER}}(\hat{\Xi})$ is only a relaxation of $\text{MASTER}(\hat{\Xi})$ and a complete set of β -cuts does not give an exact reformulation of (5). In the next section, we evaluate the performance of β -cuts and show that they are very effective in reducing the number of x -cuts that the algorithm needs to converge to optimality.

5. Numerical Results

In this section, we study the effectiveness of the proposed models and algorithms for dealing with uncertainty that arises in a WDM network with dynamic traffic. In Section 5.1 we present the characteristics of the studied networks and describe the data sets. Next, in Sections 5.2 and 5.3, we divide our analysis into two parts, SmaxRWA and SmaxLR, respectively. More specifically, for each problem we study the algorithmic performance of the decomposition framework given in Section 4, and through simulation, we compare the stochastic models with their traditional deterministic counterparts.

Implementation Details. Programs are written in Java and run on Niagara¹ servers (Ponce et al. 2019, Loken et al. 2010) using CPLEX (version V12.10.0) as the integer linear programming solver. In all experiments, a time limit of 600 seconds per IP-LP is imposed

¹https://docs.scinet.utoronto.ca/index.php/Main_Page

on the solver. In the implementation of the decomposition algorithms, we have benefited from the β -cuts in a pre-processing phase with first-stage decisions relaxed (LP-LP). The cuts are added to the LP-LP problem until no other β -cut is violated. Then the first-stage decisions are converted to integers and the IP-LP problem, together with the generated β -cuts are passed to CPLEX. In this phase, we have used callbacks for adding the cuts, where β -cuts are generated at both fractional and integral nodes, while x -cuts are added only if no violated β -cut is found at an integral node.

5.1. Data Sets

For the physical topology of our instances, we have considered six standard long-haul networks (Table 1) from the literature of WDM networks. These topologies represent different levels of connectivity degree (column “avg deg.”), ranging from ABILENE with a small number of links and paths, to the highly interconnected network BRAZIL. In Figure 4 we have reproduced the networks to give a visual representation of their meshness.

Table 1 Networks

Network	$ \mathcal{V} $	$ \mathcal{L} $	avg deg.	Source
ABILENE	12	30	2.5	Orlowski et al. (2010)
COST239	11	50	4.5	Tan and Sinclair (1996)
NSF	14	42	3.0	Ramaswami and Sivarajan (1996)
ATLANTA	15	44	2.9	Orlowski et al. (2010)
USA	24	88	3.7	Batayneh et al. (2011)
BRAZIL	27	140	5.2	Noronha and Ribeiro (2006)

Arrivals of the requests (size of the incoming batches) follow a Poisson distribution with parameter λ^R , while for service holding times (how long a connection remains on the network), we have considered an exponential distribution with mean $1/\lambda^L$ (as in Chen et al. (2015) and Xiong et al. (2018)). Distribution of lightpath requests among node pairs is uniform (the same as Giorgetti et al. (2015)). For each network, the arrival and service holding time parameters (provided for each analysis in their respective subsection) are selected via some preliminary experiments to avoid trivial solutions, i.e., such that the decisions are visibly affected by the uncertainty of incoming traffic. Note that, if the size of the batches with respect to the capacity of the network and the length of the planning horizon is too small, the assignment and routing decisions matter less, as the network would be free to accept all the requests regardless of the previous decisions. On the other hand, if the size is too large, the network becomes rapidly congested and soon there is no




Figure 4 WDM Networks

room to grant any new requests via deterministic or stochastic models. For this reason, for each experiment we have chosen the parameters λ^R and $|\Omega|$ based on the considered problem (SmaxRWA or SmaxLR), network and the number of requests on the network.

5.2. Analysis of the SmaxRWA Problem

In this section, we study the model and algorithms proposed for SmaxRWA for the networks given in Table 1.

5.2.1. Algorithmic Performance We have evaluated the performance of the decomposition method in comparison to solving the extensive form (see Section 4.1) with CPLEX, as well as the strength of the β -cut family. For this analysis, the smallest network ABILENE is selected (for other parameter settings, see Section 5.2.2). The evaluation is in terms of the solution time in seconds, optimality gap at the end of the time limit of 10 minutes, and in case of the decomposition methods, the number of added cuts. In Table 2, the average results of 30 repetitions for each scenario level is reported. In this table BENDERS- x stands for the Benders decomposition method when it purely relies on the x -cuts, while BENDERS- $x\beta$ embeds both family of cuts. If an instance does not converge within the time limit, it is indicated by “TL”, and if it does not return a gap, it is marked by “NA”. To make the comparison more meaningful, in this experiment, we omitted the

LP-LP phase and cut generation at fractional solutions. The reason for the latter is that, as demonstrated by the results in Table 2, BENDERS- x adds a large number of x -cuts, and increasing them by allowing cut generation at fractional solutions heavily impacts its performance, e.g., the algorithm is not able to converge even in the smallest scenario level.

Table 2 Average algorithmic performance of various methods for solving SmaxRWA problem over 30 repetitions, when ABILENE network is considered.

$ \hat{\mathcal{E}} $	CPLEX		BENDERS- x			BENDERS- $x\beta$			
	time (s)	gap (%)	time (s)	gap (%)	$\#x$ -cuts	time (s)	gap (%)	$\#\beta$ -cuts	$\#x$ -cuts
10	200	0	187	3	1124	17	0	8	15
20	TL	NA	465	2	1370	42	0	19	61
30	TL	NA	536	3	1538	157	0	31	204
40	TL	NA	TL	3	1785	164	0	34	203
50	TL	NA	TL	3	1980	206	0	58	209
100	TL	NA	TL	3	2164	305	0	99	400

The results of Table 2 clearly show the superiority of BENDERS- $x\beta$ compared to the other two methods. As expected, CPLEX is unable to solve the extensive form with more than a few scenarios, and after the first scenario level, it even fails to provide a nonzero feasible solution with an optimality gap. BENDERS- $x\beta$ solves every instance in the given time limit. It is able to significantly reduce the number of x -cuts by adding a small number of β -cuts, demonstrating the strength of this new family of cuts.

5.2.2. Comparison with the maxRWA Problem Over a planning horizon of 52 stages (representing weeks of a year), we have simulated the outcomes of SmaxRWA and maxRWA in granting new requests on the networks of Table 1, when we have batch arrivals and predetermined decision-making times. Networks are initially empty and the simulation is repeated for 50 times. Thus, for each instance, 50 sample paths are generated, which are characterized by the incoming requests at every stage, along with their service holding time. For the latter, $\lambda^L \in \{13, 26, 52\}$ are used in the experiments for each network. At every stage, SmaxRWA and maxRWA grant new requests by assigning lightpaths. As the stages move forward, some requests are dropped according to their service holding times, releasing resources that the models can use. This process is outlined in Figure 5.

Our decisions matter the most at certain ratios of arrival rate to the number of wavelengths, $\lambda^R/|\Omega|$, as a measure of the growth rate of the GoS, specifically in those where the network neither remains empty enough to accept all the requests over the planning




Figure 5 Simulation framework for SmaxRWA

horizon, nor quickly gets congested without any bandwidth room to grant new requests. In either of these two states, both models perform practically the same. By experimenting with various parameter settings, an interstate is achieved using the values in Table 3. In this table, we have focused on the ratio of $\lambda^R/|\Omega|$ and kept the size of Ω small enough such that the simulation framework is able to proceed for a long planning horizon of 52 stages.

Table 3 Parameter Setting for SmaxRWA vs. maxRWA Simulation

	ABILENE	COST239	NSF	ATLANTA	USA	BRAZIL
λ^R	40	10	50	60	40	50
$ \Omega $	10	10	10	10	5	5

In order to decide on the number of scenarios to be used in the SAA problems in the remaining experiments, an SAA analysis is performed for the two smallest networks, ABILENE and COST239. The full SAA analysis is summarized in Appendix (Section B). Based on the results, sample size of $|\hat{\Xi}| = 500$ is selected for ABILENE and COST239 instances. For the rest of the networks, the largest computationally affordable sample size in the given time limit is employed: $|\hat{\Xi}| = 250$ for NSF, ATLANTA and USA, $|\hat{\Xi}| = 50$ for BRAZIL.

Figures 6 and 7 present the relative performance of SmaxRWA in terms of overall GoS and the number of granted requests at each stage, compared to the deterministic maxRWA problem which does not take potential future arrivals into consideration. These figures show the percentage of improvement or loss obtained from incorporating SmaxRWA in comparison to maxRWA. Results for medium-sized networks, namely NSF and ATLANTA,




Figure 6 Relative performance of SmaxRWA for small-sized networks over different simulation horizons

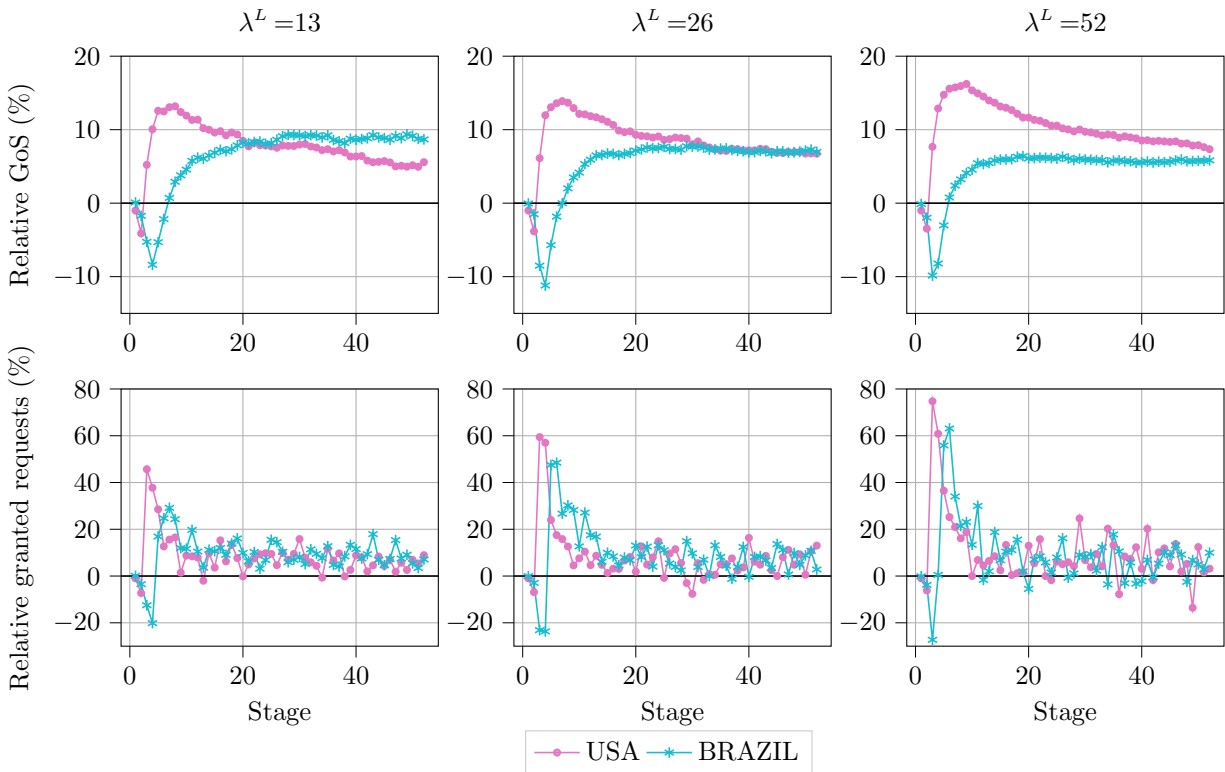



Figure 7 Relative performance of SmaxRWA for large-sized networks over different simulation horizons


are given in Appendix (Section C). In all instances, by using SmaxRWA for lightpath assignment decisions, the number of granted requests and GoS have improved over the planning horizon, with improvements up to 73% in granted requests (USA with $\lambda^L = 52$) and 19% for the GoS (COST239 with $\lambda^L = 26$). The results show that SmaxRWA opts for less request grants at earlier stages to reserve the resources for future arrivals. After few stages, it catches up with and surpasses the deterministic model.

In terms of granted requests, whose maximization is the actual objective function of SmaxRWA and maxRWA, the improvement is more significant when the service holding time of a request is longer and errors in allocation of resources by the decision maker have a longer term impact. For GoS though, COST239 and BRAZIL exhibit dissimilar patterns. As suggested by the shape of their relative GoS curve (which, unlike the other instances, appear not to have peaked), our planning horizon does not give a full picture of their behavior. These two networks are highly interconnected, therefore they have a large number of candidate lightpaths, i.e., high capacity for granting the requests. We conjecture that the peak of the relative GoS improvement should appear over a longer planning horizon, giving a full picture of their behavior.


Improvements in overall GoS and granted requests are more valuable when an important aspect of solution characteristics of both models is examined: bandwidth spectrum usage. We define the spectrum usage as the number of wavelinks utilized in granting the requests. From the results in Figures 6 and 7, we might wonder whether maxRWA is able to catch up to SmaxRWA in the long run, considering the higher number of granted requests on the network if we use SmaxRWA. The answer to this question depends on the resource utilization by SmaxRWA, i.e., spectrum usage. Figure 8 reveals that, despite the higher number of granted requests by SmaxRWA, it has a lower spectrum usage, with savings of up to 20% for the smaller networks in the earlier stages. In other words, SmaxRWA achieves higher GoS by refusing some requests, as well as making informed path assignments that utilize less resources. SmaxRWA effectively decreases the number of used wavelinks, an objective sought out by the minRWA problem. By doing so, it postpones the fragmentation of the network and maintains a lower level of resource consumption. Appendix (Section C) includes the set of results for the three other networks not presented in Figure 8, where the same trend emerges in all of our experiments.



(a) ABILENE



(b) COST239



(c) BRAZIL

Figure 8 Relative bandwidth spectrum usage

5.2.3. Quality of Relaxation To assess the quality of IP-LP bound, we have considered the network COST239 that has the smallest λ^R among the instances and CPLEX is able to solve its IP-IP extensive form with up 50 scenarios in reasonable time. Consider the simulation framework of Figure 5. For values of $\lambda^L \in \{13, 26, 52\}$, we compare the optimal objective values of IP-IP and IP-LP models at different points over the planning horizon. Let \hat{t} be this particular stage. For providing a comparable setting, at each experiment, up

to stage \hat{t} we grant the incoming traffic with maxRWA. At stage \hat{t} , we solve both IP-IP and IP-LP problems with the same sample $\hat{\Xi}$ of size 50, and repeat this 30 times. Table 4 presents the comparison results. The column labeled “ \neq ” gives the number of samples (out of 30) with different IP-IP and IP-LP optimal values. For these particular samples, the column “gap” measures the average relative gap between the two optimal values. The results show that in all cases the bound obtained by IP-LP problem is either exact or very close to the IP-IP optimal value. It is worth noting that, in every case, the first-stage cost (not necessarily the solution itself) of both models is the same, even if their optimal objective values slightly differ.

Table 4 Comparison of IP-LP bound and the optimal value of IP-IP, over 30 repetitions

λ^L	\hat{t}	\neq	gap (%)
	1	0/30	NA
	13	1/30	0.11
13	26	0/30	NA
	52	0/30	NA
	13	0/30	NA
26	26	4/30	0.09
	52	5/30	0.05
	13	1/30	0.10
52	26	3/30	0.09
	52	2/30	0.08

5.3. Analysis of the SmaxLR Problem

As a defragmentation method, SmaxLR is suitable for situations where there is a large number of requests on the network, but it is not at full capacity and stranded bandwidth cannot be used in granting new requests. As such, SmaxLR involves rerouting decisions for a large number of existing requests on the network, making it computationally more expensive than SmaxRWA.

5.3.1. Algorithmic Performance Table 5 presents the performance comparison of CPLEX, BENDERS- x and BENDERS- $x\beta$ (see Section 5.2.1) in solving instances of ABILENE with 100 existing requests on the network, $|\Omega| = 10$ and $\lambda^R = 10$. When they converge, CPLEX and BENDERS- x take on the average 2.5 and 2.3 times longer than BENDERS- $x\beta$, respectively. Unlike SmaxRWA, when BENDERS- x does not converge within the time limit, it yields a large optimality gap, making BENDERS- $x\beta$ the sole option for solving SmaxLR.

Table 5 Average algorithmic performance of various methods for solving SmaxLR problem over 30 repetitions, when ABILENE network is considered.

$ \hat{\Xi} $	CPLEX		BENDERS- x			BENDERS- $x\beta$			
	time (s)	gap (%)	time (s)	gap (%)	# x -cuts	time (s)	gap (%)	# β -cuts	# x -cuts
10	11	0	8	0	336	13	0	2	101
20	120	0	82	0	847	67	0	7	426
30	333	0	384	0	1545	91	0	6	418
40	528	0	439	0	4315	139	0	7	877
50	TL	NA	TL	36	4541	227	0	13	1388
100	TL	NA	TL	57	7286	338	0	21	1887

5.3.2. Comparison with the minRWA Problem In this section, we evaluate the impact of using SmaxLR for the defragmentation of small networks ABILENE and COST239. We simulate a planning horizon of 10 stages for 50 times, starting with a set of initial requests on each network. Simulation framework is depicted in Figure 9, where $\mathcal{N}_t^{\text{FRAG}}$ and $\mathcal{N}_t^{\text{DEFRAG}}$ are respectively the fragmented and defragmented provisioning at stage t . Over the course of n stages, at the beginning of each stage $t \in \{0, \dots, n-1, n\}$, a defragmentation event is triggered for $\mathcal{N}_t^{\text{FRAG}}$. Based on the result of the defragmentation target, $\mathcal{N}_t^{\text{DEFRAG}}$ (obtained via SmaxLR or minRWA), lightpaths are rerouted. At the end of stage t , the lightpaths associated with the connections leaving the network are dismantled, the uncertainty in the traffic is revealed, and new connection requests are granted using spare resources in the form of available lightpaths via the deterministic maxRWA problem. This process leads to a new fragmented provisioning. Note that for simulation purposes, our planning horizon is finite, and defragmentation event times are fixed and given.




Figure 9 Simulation framework for SmaxLR

By performing an SAA analysis, we choose $|\hat{\Xi}| = 1000$ scenarios for ABILENE and $|\hat{\Xi}| = 500$ for COST239, each resulting in worst-case optimality gaps of less than 5%. Appendix (Section B) provides the full SAA results for SmaxLR. Other parameters are selected such that non-trivial instances (amenable to defragmentation) are generated, while their size remains small and computationally affordable in our simulation framework. We consider two initial states (denoted by “Init.”) of 90 and 100 existing requests for ABILENE, when the network has 6 and 7 wavelengths, respectively. For COST239, we use 250 and 260 existing requests over 10 wavelengths in instance generation. Finally, for both networks, we have $\lambda^R = 10$ and $\lambda^L \in \{15, 20\}$.

Figures 10 and 11 show that for both networks, irrespective of the requests’ service holding time rate, defragmentation by SmaxLR grants more requests over the planning horizon, compared to minRWA. The improvements are up to 36% for ABILENE and 29% for COST239, each having 90 and 250 initial requests, respectively, with $\lambda^L = 15$. In order to achieve this, SmaxLR uses slightly more resources: up to 4% and 2.8% more for ABILENE




Figure 10 Relative performance of SmaxLR for defragmentation of ABILENE compared to minRWA

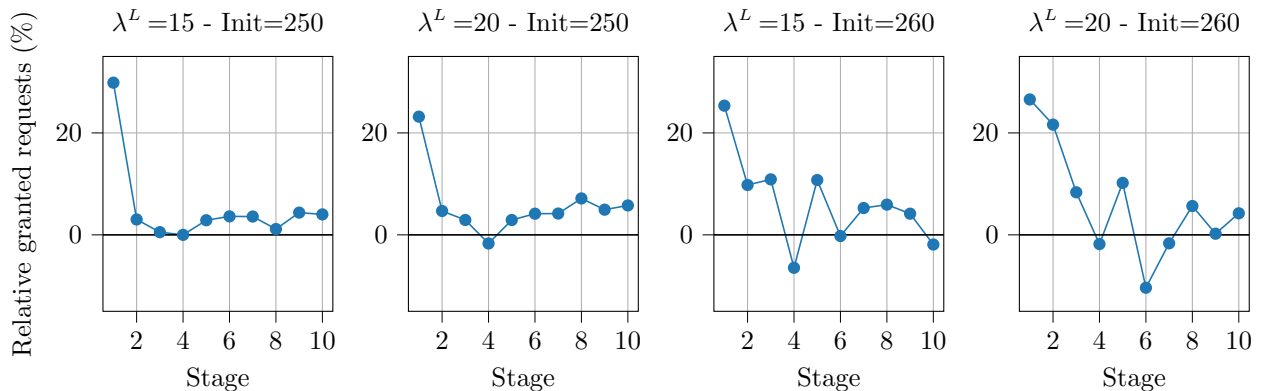


Figure 11 Relative performance of SmaxLR for defragmentation of COST239 compared to minRWA

and COST239, respectively (spectrum usage results are provided in Appendix, Section C). This shows that merely minimizing the number of links in defragmentation without considering the future traffic does not necessarily give the best GoS in longer term.

Since the simulation results for SmaxLR is not as compelling as SmaxRWA simulations, we further investigate the importance of considering the stochasticity of the traffic in our rerouting decisions. This is often achieved by computing the Value of Stochastic Solution (VSS). Let the *mean-value scenario* be the scenario built by using the expected value of random parameters, and a *mean-value solution* be an optimal solution to the deterministic model employing only the mean-value scenario. VSS is defined as the difference between the optimal value of the stochastic model and the expected long-run objective if the mean-value solution is implemented. Due to the specific discrete nature of the probability distribution of request arrivals, the mean-value scenario is not meaningful in our case. Therefore we use the notion of Expected VSS (EVSS) as defined by Mousavi et al. (2020) for the IP-IP model. To compute EVSS, we first solve the individual deterministic models associated with every scenario in our sample $\hat{\Xi}$, then evaluate the first-stage solutions of each scenario-specific problem using the scenarios in $\hat{\Xi}$. EVSS is the difference between the optimal value of the stochastic model and the average of these $|\hat{\Xi}|$ scenario objective values. Our experiments show that, for the instances where the extensive form of IP-IP model is solvable by CPLEX, there is no gap between the optimal values of IP-IP and IP-LP models (while their solutions can be different). Therefore we consider the optimal objective value of the IP-LP model as an approximation of the IP-IP model. The EVSS results are presented in Table 6, where EVSS numbers (in terms of the number of requests) are reported for different levels of initial requests on the network. σ^{det} is the standard deviation of the evaluation over $|\hat{\Xi}|$ of deterministic subproblems associated with each scenario in $\hat{\Xi}$. The results show that the stochastic model attains significant gains of up to 3 and 4 more requests out of an average of 10 incoming arrivals in ABILENE and COST239, respectively, showcasing the value of the stochastic model over a deterministic approach.

Considerable benefit from solving the stochastic model suggested by EVSS analysis is not reflected in the moderate performance of SmaxLR in simulation results. One reason could be the IP-LP approximation. Although in all the tested instances that we were able to solve IP-IP models (either using the extensive form or the integer L-shaped method), there is no gap between IP-IP and IP-LP optimal values, but their solutions are quite

Table 6 Expected value of stochastic solution for SmaxLR

Network	$\lambda^R = 10$			$\lambda^R = 20$		
	Init.	EVSS	σ^{det}	Init.	EVSS	σ^{det}
ABILENE	80	0.71	0.44	80	1.35	0.65
	90	0.77	0.65	90	1.60	1.1
	100	1.48	0.64	100	2.52	0.97
	110	2.92	0.67	110	3.86	0.84
	120	1.40	0.38	120	2.03	0.57
	130	1.19	0.24	130	1.52	0.45
	140	0.21	0.09	140	0.40	0.19
	150	0.06	0.06	150	0.11	0.11
COST239	200	1.57	0.60	200	2.85	1.06
	210	1.48	0.67	210	2.93	1.24
	220	2.66	0.84	220	4.72	1.37
	230	3.06	0.77	230	5.21	1.46
	240	4.04	0.79	240	6.66	1.48
	250	3.04	0.46	250	4.43	0.74

different. The first-stage solutions from the IP-LP model are not necessarily optimal for the IP-IP model. Another reason can be the intrinsically multistage nature of the problem that is approximated by SmaxLR, a two-stage model. SmaxLR looks one stage ahead into the future in its rerouting decisions, while the traffic arrives and leaves the network over the entire planning horizon. Figures 10 and 11 support this claim by showing that the most significant improvements in terms of granted requests are obtained at the early stages. Finally, note that, as a major source of network fragmentation, drops in the simulation framework are not modeled by SmaxLR, which could lead to losing more benefits.

6. Conclusion

Stochastic nature of the traffic in WDM networks is a contributing factor in sub-optimal usage of the available bandwidth. To tackle this issue, in this paper we have introduced SmaxRWA, the first 2SP model for maxRWA as the provisioning problem of WDM networks, with the purpose of making more informed lightpath assignment decisions. Furthermore, for the costly operation of network defragmentation that aims to improve the spectrum usage in WDM networks, we have proposed SmaxLR, a 2SP lightpath rerouting model as the provisioning target of the procedure.

In order to solve the 2SP models, we have designed a decomposition framework based on two different relaxations that make use of our newly developed problem-specific family of cuts. Our numerical results show that these cuts are quite effective in reducing the

number of Benders optimality cuts needed for the algorithm to converge, in turn improving the scalability of the solution framework. Through a set of comprehensive numerical experiments over six standard long-haul networks in the literature, we have showed that SmaxRWA considerably improves the overall GoS in all instances compared to the results from deterministic maxRWA. At the same time, SmaxRWA reduces the spectrum usage despite having a higher number of connection requests on the network, leaving more room for future traffic. As for the SmaxLR problem, our numerical experiments are promising, showing moderate improvements in the number of granted requests and overall GoS.

Our numerical study indicates the merit of considering traffic uncertainty in the decision making involved in provisioning and defragmentation of WDM networks. This leads to several questions that open new areas of research. First, our 2SP models look one step ahead into the future traffic, while the multistage nature of the provisioning and defragmentation decisions suggests that a multistage model might better capture the process. Multistage problems involving integer decision variables are notoriously difficult to solve, therefore the practicality of them for SmaxRWA and SmaxLR problems remains open. Second, although our designed solution method is able to demonstrate the value of using a stochastic model for several networks, it is not able to solve very large instances where network size, arrival of the connection requests or the number of available wavelengths is considerably larger than the studied instances. In addition, as demonstrated by our evaluation of SmaxLR problem, solving the exact IP-IP model is also of interest. Therefore, new methodological contributions are needed. Finally, significant savings from SmaxRWA problem for WDM networks unfold new questions about the impact that ignoring the future traffic can have on other practical decisions concerning WDM and more generally other types of optical networks. For instance, an interesting question is whether a stochastic model can determine the best time for triggering a defragmentation process by taking the future traffic into account. As another example, flex-grid networks, the popular emerging transmission systems, face the same traffic uncertainty as WDM networks. Thus, their provisioning, a much more difficult problem than the RWA, might also benefit from stochastic models.

Acknowledgments

This work was supported by Natural Sciences and Engineering Research Council of Canada [Grant RGPIN-2018-04984]. Computations were performed on the Niagara supercomputer at the SciNet HPC Consortium. SciNet is funded by: the Canada Foundation for Innovation; the Government of Ontario; Ontario Research Fund - Research Excellence; and the University of Toronto.

References

- Alliance E (2020) 2020 Ethernet roadmap. <https://ethernetalliance.org/technology/2020-roadmap/>.
- Angulo G, Ahmed S, Dey SS (2016) Improving the integer L-shaped method. *INFORMS Journal on Computing* 28(3):483–499.
- Aparicio-Pardo R, Pavon-Marino P, Mukherjee B (2012) Robust upgrade in optical networks under traffic uncertainty. *16th International Conference on Optical Network Design and Modelling (ONDM)*, 1–6 (IEEE).
- Batayneh M, Schupke DA, Hoffmann M, Kirstaedter A, Mukherjee B (2011) On routing and transmission-range determination of multi-bit-rate signals over mixed-line-rate WDM optical networks for carrier Ethernet. *IEEE/ACM Transactions on Networking* 19(5):1304–1316.
- Benders JF (2005) Partitioning procedures for solving mixed-variables programming problems. *Computational Management Science* 2(1):3–19.
- Chen X, Zhu S, Jiang L, Zhu Z (2015) On spectrum efficient failure-independent path protection p-cycle design in elastic optical networks. *Journal of Lightwave Technology* 33(17):3719–3729.
- Christodoulopoulos K, Manousakis K, Varvarigos E (2010) Offline routing and wavelength assignment in transparent WDM networks. *IEEE/ACM Transactions on Networking* 18(5):1557–1570.
- Ciena (2020) What is WDM? <https://www.ciena.com/insights/what-is/what-is-wdm.html>.
- Cisco (2020) Cisco annual internet report, 2018-2023.
- Daryalal M (2016) Efficient spectrum utilization in large-scale RWA and RSA problems.
- Ellis A, Doran N (2013) Are few-mode fibres a practical solution to the capacity crunch? *15th International Conference on Transparent Optical Networks (ICTON)*, 1–4 (IEEE).
- Essiambre RJ, Foschini GJ, Kramer G, Winzer PJ (2008) Capacity limits of information transport in fiber-optic networks. *Physical Review Letters* 101(16):163901.
- Giorgetti A, Paolucci F, Cugini F, Castoldi P (2015) Dynamic restoration with GMPLS and SDN control plane in elastic optical networks. *Journal of Optical Communications and Networking* 7(2):A174–A182.
- Hooker JN, Ottosson G (2003) Logic-based Benders decomposition. *Mathematical Programming* 96(1):33–60.
- Inniss D, Rubenstein R (2016) *Silicon photonics: Fueling the next information revolution* (Morgan Kaufmann).
- Jara N, Pempelfort H, Rubino G, Vallejos R (2020) How much the wavelength dimensioning methods and a tightened QoS provision impact on the dynamic WDM optical networks capacity? *Optical Switching and Networking* 35:100540.
- Jaumard B, Daryalal M (2017) Efficient spectrum utilization in large scale RWA problems. *IEEE/ACM Transactions on Networking* 25(2):1263–1278.

- Jaumard B, Meyer C, Thiongane B (2007) Comparison of ILP formulations for the RWA problem. *Optical Switching and Networking* 4(3-4):157–172.
- Jaumard B, Pouya H, Coudert D (2019) Wavelength defragmentation for seamless migration. *Journal of Lightwave Technology* 37(17):4382–4393.
- Ji F, Chen X, Lu W, Rodrigues JJ, Zhu Z (2014) Dynamic p-cycle protection in spectrum-sliced elastic optical networks. *Journal of Lightwave Technology* 32(6):1190–1199.
- Keiser GE (1999) A review of WDM technology and applications. *Optical Fiber Technology* 5(1):3–39.
- Kennington J, Lewis K, Olinick E, Ortynski A, Spiride G (2003) Robust solutions for the DWDM routing and provisioning problem: Models and algorithms. *Optical Networks Magazine* 4(2):74–84.
- Krishnaswamy RM, Sivarajan KN (2001) Algorithms for routing and wavelength assignment based on solutions of LP-relaxations. *IEEE Communications Letters* 5(10):435–437.
- Kronberger C, Schondienst T, Schupke DA (2011) Impact and handling of demand uncertainty in multiperiod planned networks. *IEEE International Conference on Communications (ICC)*, 1–6 (IEEE).
- Kuri J (2003) *Optimization problems in WDM optical transport networks with scheduled lightpath demands*. Ph.D. thesis.
- Laporte G, Louveaux FV (1993) The integer L-shaped method for stochastic integer programs with complete recourse. *Operations Research Letters* 13(3):133–142.
- Lavallée B (2017) Shannon’s limit, or opportunity? <https://www.ciena.com/insights/articles/Shannons-Limit-or-Opportunity.html>.
- Leung D, Grover WD (2005) Capacity planning of survivable mesh-based transport networks under demand uncertainty. *Photonic Network Communications* 10(2):123–140.
- Lodha A, Gumaste A, Bafna P, Ghani N (2007) Stochastic optimization of Light-trail WDM ring networks using Bender’s decomposition. *2007 Workshop on High Performance Switching and Routing*, 1–7 (IEEE).
- Loken C, Gruner D, Groer L, Peltier R, Bunn N, Craig M, Henriques T, Dempsey J, Yu CH, Chen J, et al. (2010) SciNet: lessons learned from building a power-efficient top-20 system and data centre. *Journal of Physics-Conference Series*, volume 256, 012026.
- Majumdar AK (2018) *Optical wireless communications for broadband global internet connectivity: Fundamentals and potential applications* (Elsevier).
- Mousavi K, Bodur M, Roorda M (2020) Stochastic last-mile delivery with crowd-shipping and mobile depots https://www.optimization-online.org/DB_FILE/2020/02/7619.pdf.
- Noronha TF, Ribeiro CC (2006) Routing and wavelength assignment by partition colouring. *European Journal of Operational Research* 171(3):797–810.

-
- Orlowski S, Wessälly R, Pióro M, Tomaszewski A (2010) SNDlib 1.0—Survivable network design library. *Networks: An International Journal* 55(3):276–286.
- Ozdaglar AE, Bertsekas DP (2003) Routing and wavelength assignment in optical networks. *IEEE/ACM Transactions on Networking* 11(2):259–272.
- Ponce M, van Zon R, Northrup S, Gruner D, Chen J, Ertinaz F, Fedoseev A, Groer L, Mao F, Mundim BC, et al. (2019) Deploying a top-100 supercomputer for large parallel workloads: The Niagara supercomputer. *Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines (learning)*, 1–8.
- Ramaswami R, Sivaraajan KN (1996) Design of logical topologies for wavelength-routed optical networks. *IEEE Journal on Selected Areas in Communications* 14(5):840–851.
- Shannon CE (1948) A mathematical theory of communication. *The Bell System Technical Journal* 27(3):379–423.
- Shapiro A (1996) Simulation-based optimization—convergence analysis and statistical inference. *Stochastic Models* 12(3):425–454.
- Shapiro A, Dentcheva D, Ruszczyński A (2014) *Lectures on stochastic programming: Modeling and theory* (SIAM).
- Tan L, Sinclair M (1996) Wavelength assignment between the central nodes of the COST 239 European optical network. *Performance Engineering of Computer and Telecommunications Systems*, 235–247 (Springer).
- Waldman H (2018) The impending optical network capacity crunch. *International Optics and Photonics Conference*, 1–4 (IEEE).
- Wu J, Zhang J, von Bochmann G, Savoie M (2012) Forward-looking WDM network reconfiguration with per-link congestion control. *Journal of Network and Systems Management* 20(1):6–33.
- Xenos H (2020) 800G – Nothing but the facts. <https://www.ciena.com/insights/articles/800g-nothing-but-the-facts.html>.
- Xiong Y, Li Y, Zhou B, Wang R, Rouskas GN (2018) SDN enabled restoration with triggered precomputation in elastic optical inter-datacenter networks. *Journal of Optical Communications and Networking* 10(1):24–34.
- Zhang JY, Wu J, Bochmann GV (2013) A proof of wavelength conversion not improving Lagrangian bounds of the sliding scheduled RWA problem. *Computer Communications* 36(5):600–606.
- Zhou Q, Zhang R, Chen YW, Shen S, Su SJ, Finkelstein J, Chang GK (2020) Combining efficient probabilistic shaping and deep neural network to mitigate capacity crunch in 5G fronthaul. *2020 Optical Fiber Communications Conference and Exhibition (OFC)*, 1–3 (IEEE).
- Zhu K, Mukherjee B (2003) A review of traffic grooming in WDM optical networks: Architectures and challenges. *Optical Networks Magazine* 4(2):55–64.

Appendix

A. Extensive Form for the SmaxRWA Problem

Denote by $\mathcal{F}_R^{\text{RWA}}(x, \xi, \mathcal{SD})$ the feasible set of $\mathcal{Q}(x, \xi, \mathcal{SD})$, defined in the paper by (3b)-(3h).

We can reformulate the SAA problem (5) as a mixed-integer programming model:

$$\max \sum_{(s,d) \in \mathcal{SD}_N^1} \sum_{\omega \in \Omega} \sum_{\substack{\ell \in \\ \delta^+(s) \cap \mathcal{L}(\omega)}} x_{\ell, \omega}^{(s,d)} + \frac{1}{|\hat{\Xi}|} \sum_{\xi \in \hat{\Xi}} \eta_\xi \quad (\text{A.1a})$$

$$\text{s.t. } x \in \mathcal{F}_R^{\text{RWA}}(\mathcal{SD}_N^1) \quad (\text{A.1b})$$

$$\eta_\xi \leq \sum_{(s,d) \in \mathcal{SD}^2(\xi)} z^{(s,d)}(\xi) \quad \xi \in \hat{\Xi} \quad (\text{A.1c})$$

$$(y(\xi), z(\xi)) \in \mathcal{F}_R^{\text{RWA}}(x, \xi, \mathcal{SD}_N^1) \quad \xi \in \hat{\Xi} \quad (\text{A.1d})$$

$$\eta_\xi \geq 0 \quad \xi \in \hat{\Xi}, \quad (\text{A.1e})$$

i.e., by creating copies of the second-stage decision variables and constraints for each scenario in the sample. This is the *extensive form* of the SAA problem.

B. SAA Analysis

In an SAA analysis, we decide on the sample size to be used in the SAA problem, in order to get statistically valid lower and upper bound on the optimal value of the 2SP model. Consider a maximization problem. In an SAA analysis with n repetitions, for each scenario level $|\hat{\Xi}|$, we generate n samples of size $|\hat{\Xi}|$ and solve the SAA problem for each sample. The average of the optimal values of these SAA problems provides a statistically valid upper bound. We then evaluate the solutions of each SAA problem with a much larger evaluation sample and pick the best one, which gives us a statistical lower bound. This process is repeated for different scenario levels, until a predefined condition is met. For the details of the steps involved in this analysis, interested reader can refer to [Shapiro et al. \(2014\)](#). In the following, for both SmaxRWA and SmaxLR, we have performed an SAA analysis considering the two smallest networks, ABILENE and COST239, with 30 repetitions and an evaluation sample of 5000 scenarios.

For SmaxRWA, we have considered two scenario levels of $|\hat{\Xi}| \in \{100, 500\}$, as the larger sample sizes failed to converge in the given time limit. A contributing factor in this analysis can be the state of the network when SmaxRWA is solved, i.e., which wavelinks are already

serving connections on the network. For generating such initial states, we have considered the simulation framework given in Section 5.2.2 of the paper, and following one sample path (with maxRWA used to grant new arrivals), we have taken three snapshots of the state of the network at stages 13, 26 and 52, to be used as the initial state. Then, for each combination of network, stage and scenario level, we have solved the instance we have performed an SAA analysis with 30 repetitions. Results are summarized in Table B.1, where the mean and width of the confidence intervals for upper bound (UB) and lower bound (LB) results are given, as well as the final worst-case optimality gap. Based on these results, for ABILENE and COST239 instances, we have chosen sample size of $|\hat{\Xi}| = 500$ where the optimality gap is less than or equal to 3%. Note that, our main goal is to retrieve feasible solutions to be used in the simulation framework, hence the small LB widths are adequate.

Table B.1 SmaxRWA SAA results

Network	Stage	$ \hat{\Xi} $	UB		LB		gap (%)
			Mean	Width	Mean	Width	
ABILENE	13	100	19.63	0.07	19.10	0.11	4.36
		500	19.70	0.04	19.23	0.02	2.18
	26	100	21.88	0.07	21.30	0.16	5.08
		500	21.89	0.05	21.47	0.18	2.35
	52	100	21.44	0.07	21.10	0.11	4.17
		500	21.45	0.07	21.19	0.07	2.06
COST239	13	100	21.02	0.09	20.40	0.18	3.71
		500	20.36	0.04	20.09	0.13	2.73
	26	100	19.03	0.09	18.37	0.17	3.83
		500	19.12	0.03	18.72	0.02	3.09
	52	100	21.96	0.09	21.33	0.17	2.44
		500	22.03	0.07	21.70	0.04	1.93

For each network considered in the SmaxLR problem, two levels of initial connection requests on the network (denoted by “Init.”) are considered. Requests are then randomly generated while making sure that they can result in a feasible provisioning. Table B.2 presents the SAA results for three scenario levels of $|\hat{\Xi}| \in \{10, 500, 1000\}$. Based on these experiments, for ABILENE and COST239 we choose $\hat{\Xi} = 1000$ and $\hat{\Xi} = 500$, respectively, where at this level the worst-case optimality gap is less than 5% for instances of both networks.

Table B.2 SmaxLR SAA results

Network	Init.	$ \hat{\mathcal{E}} $	UB CI		LB CI		gap (%)
			Mean	Width	Mean	Width	
ABILENE	100	10	9.97	0.31	9.04	0.05	14.36
		500	9.98	0.04	9.71	0.08	4.10
		1000	9.97	0.03	9.89	0.08	2.03
	120	10	4.59	0.24	4.09	0.03	19.00
		500	4.69	0.03	4.48	0.04	6.18
		1000	4.68	0.02	4.53	0.04	4.79
COST239	200	10	9.83	0.39	9.62	0.08	7.24
		500	9.97	0.05	9.85	0.08	2.57
		1000	9.99	0.03	9.99	0.09	1.24
	220	10	9.84	0.35	9.52	0.08	7.97
		500	10.03	0.04	9.77	0.08	3.87
		1000	10.02	0.03	10.01	0.09	1.30

C. Simulation Results

In this section, we present the simulation results that are removed from the main text for the sake of brevity. Figure C.1 illustrates the relative improvements obtained from employing SmaxRWA over the maxRWA problem for medium-sized networks NSF and ATLANTA. Together with Figures 6 and 7 given in the paper, they provide a general picture of the merit of incorporating stochasticity in our decision-making in terms of granting requests over different network.

In addition to the results presented for ABILENE, COST239 and BRAZIL in Figure 8 of the paper, we have analyzed the spectrum usage of NSF, ATLANTA and USA and the results are depicted in Figure C.2. The results are consistent with previously mentioned network.

Likewise, Figure C.3 depicts the relative spectrum usage by SmaxLR compared to the minRWA problem, which is slightly more than that of the latter.

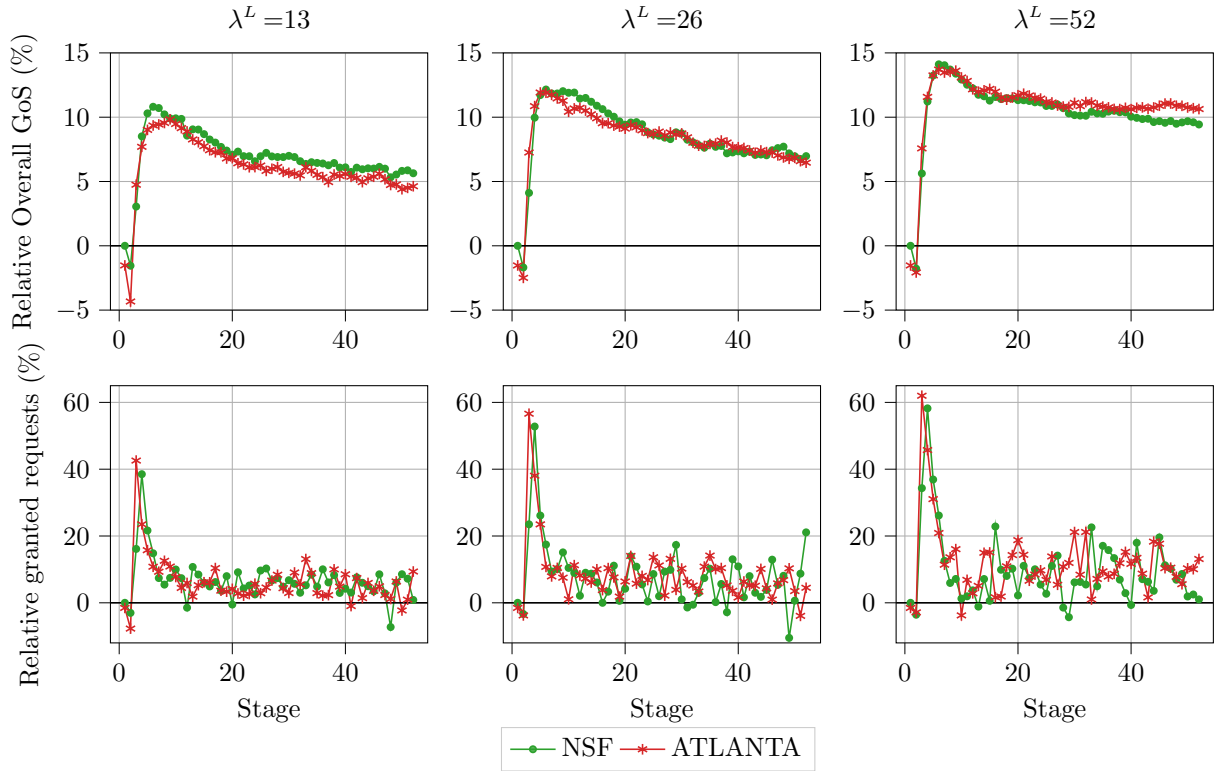





Figure C.1 Relative performance of SmaxRWA for medium-sized networks over different simulation horizons.



(a) NSF




(b) ATLANTA

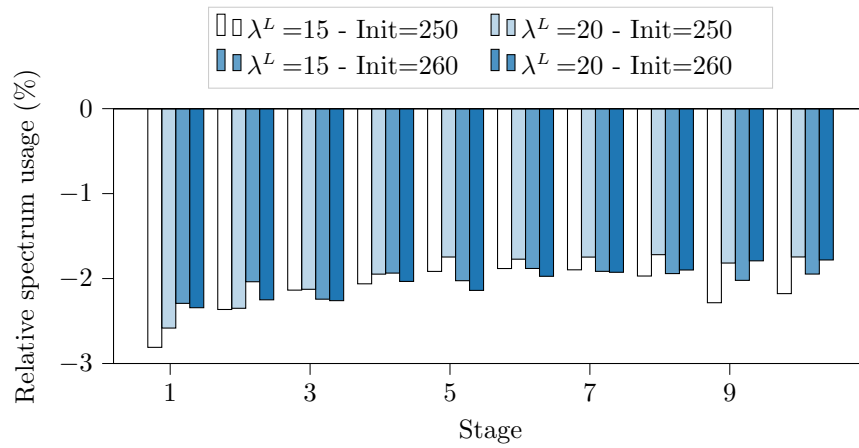


(c) USA

Figure C.2 Relative bandwidth spectrum usage



(a) ABILENE



(b) COST239

Figure C.3 Relative bandwidth spectrum usage after each rerouting