Real-Time Power Control for Dynamic Optical Networks – Algorithms and Experimentation

Berk Birand, *Student Member*, Howard Wang, *Member*, Keren Bergman, *Fellow*, Dan Kilper, *Senior Member*, Thyaga Nandagopal, *Senior Member*, and Gil Zussman, *Senior Member*

***Abstract*—Core and aggregation optical networks are remarkably static, despite the emerging dynamic capabilities of the individual optical devices. This stems from the inability to address optical impairments in real-time. As a result, tasks such as adding and removing wavelengths take a substantial amount of time, and therefore, optical networks are over-provisioned and inefficient in terms of capacity and energy. Optical Performance Monitors (OPMs) that assess the Quality of Transmission (QoT) in real-time can be used to overcome these inefficiencies. However, prior work mostly focused on the single link level. In this paper, we present a *network-wide* optimization algorithm that leverages OPM measurements to dynamically control the wavelengths' power levels. Hence, it allows adding and dropping wavelengths quickly while mitigating the impacts of impairments caused by these actions, thereby facilitating efficient operation of higher layer protocols. We evaluate the algorithm's performance using a *network-scale optical simulator* under real-world scenarios and show that the ability to add and drop wavelengths dynamically can lead to significant power savings. Moreover, we experimentally evaluate the algorithm in an *optical testbed* and discuss the practical implementation issues. To the best of our knowledge, this paper is the first attempt at providing a global power control algorithm that uses live OPM measurements to enable dynamic optical networking.**

***Index Terms*—Optical networks, network management, power control algorithms, performance evaluation, energy efficiency.**

I. Introduction

Optical networks are the underlying infrastructure of core and aggregation networks [29]. In order to handle peaks in traffic demand, these networks are usually *static and over-provisioned* [38]. This leads to inefficient use of capacity and energy (due to the need to keep inactive lightpaths available). The increase in traffic demand and heterogeneity as well as the need to support energy-efficient operation [25] already pose challenges that cannot be addressed by over-provisioning. Dynamic adaptation of the network will become more pronounced as the fiber capacity is approached [15] and the energy consumption of the IT sector becomes more prominent [25]. Requirements for dynamic adaptation stem not only from ongoing changes in Internet traffic patterns (e.g., due to diurnal cycles), but also from the emerging needs to quickly provision high-bandwidth inter-data-center links [21, 24].

B. Birand, H. Wang, K. Bergman, and G. Zussman are with the Department of Electrical Engineering, Columbia University, New York, NY (e-mail: {berk,howard,keren,gil}@ee.columbia.edu). D. Kilper is with the Center for Integrated Aggregation Networks, College of Optical Sciences, University of Arizona, Tucson, AZ (e-mail: dkilper@optics.arizona.edu). T. Nandagopal is with the National Science Foundation, Arlington, VA (e-mail: tnandago@nsf.gov). Partial and preliminary versions of this paper appeared in OSA OFC'13 [8] and IEEE ICNP'13[7].

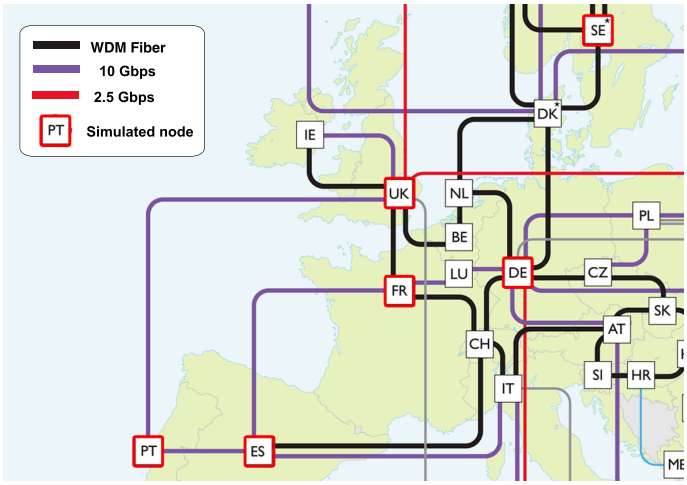

Fig. 1. The optical infrastructure of the Géant academic network[2], whose topology is used in order to evaluate the proposed algorithms. The highlighted nodes are used as part of the topology considered in the simulations in Section VI.

Wavelength-Switched Optical Networks (WSONs) (see e.g., Fig. 1) include various emerging dynamic optical devices which have the potential to address these challenges. Dynamic devices include, for example, Reconfigurable Add/Drop Multiplexers (ROADMs) that can transparently switch the transmissions from one lightpath to another [29], modulators that can adapt to the link state [16], and bandwidth variable transceivers that can modify band gaps between adjacent channels [17]. While the flexibility provided by such devices allows the network to adapt to the link conditions and traffic demands, optical networks are still mostly static, due to potential impairments that are hard to predict or model [6, 38]. Sources of these impairments are related to the optical transmission and fiber properties [29], and to factors such as temperature, component drift, component aging, and maintenance work [33]. Due to these impairments, lightpaths are rarely modified once assigned. This means that Routing and Wavelength Allocation (RWA) (e.g.,[6]) is done primarily at the planning phase, with significant over-provisioning. Any changes are executed manually which is both time-consuming and expensive [14].

Hence, our goal is to enable lightpath configuration, setup, and teardown with convergence times in the order of tens of seconds. This will allow higher layer protocols to adapt to traffic variations and customer demands, thereby leading to significant energy savings (e.g., [22, 35]). We build on the capabilities of the dynamic optical devices as well as various *Optical Performance Monitors (OPMs)* that have been recently

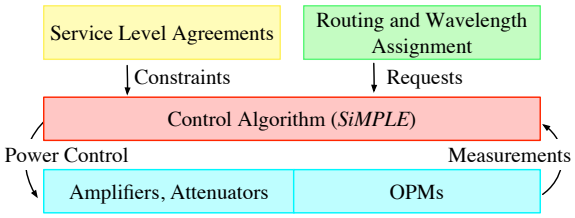

Fig. 2. Schematic view of the interaction of the control algorithm with the optical devices, and the higher layer algorithms and SLAs.

developed [33]. OPMs can measure Quality-of-Transmission (QoT) parameters such as the Optical Signal-to-Noise Ratio (OSNR), Bit Error Rate (BER), and chromatic dispersion in real-time. Yet, while OPM capabilities have improved, most control schemes that use them operate at the link-scale rather than at the network-scale [28]. Extensions of per-link policies to the entire network do not produce globally optimal results, and may not converge within the desired time [8].

We develop an impairment-aware, *network-wide power control algorithm*. As illustrated in Fig. 2, the algorithm will allow operators to control the dynamic devices such that the network will be maintained in a state that satisfies the QoT constraints and higher layer requirements. The algorithm would support quick reaction to changes (e.g., addition or removal of lightpaths), and can therefore facilitate the dynamic operation of higher layer RWA algorithms. We note that schemes that require close interaction between the layers are only starting to gain attention in the node/link-level of optical networks [36, 39]. *The development of network-scale schemes has rarely been addressed and is a challenging open problem*, due to the following reasons:

- **Continuous Operation** – Most optimization problems associated with optical networks are solved offline during the network planning phase or when lightpaths are added or removed. Dynamically solving these problems on a live production network requires always maintaining a feasible solution, which is challenging given the unpredictable and time-varying nature of the optical links.
- **Unknown Performance Functions** – The analytical expressions (and derivatives) for the BER and the OSNR as functions of the power levels in the network are intractable, and therefore, most optimization algorithms are inapplicable.
- **Limited Performance Evaluation Infrastructure** – Optical testbeds based on off-the-shelf networking equipment are limited in conducting dynamic experiments. The exposed functionality usually only allows higher-layer operations such as lightpath provisioning.

To overcome these limitations, we formulate the *Multi Link Optimization (MLO) problem* and present the *Simultaneous Multi-Path Lambda Enhancement (SiMPLE) algorithm* which controls the power levels of the wavelengths. Since the analytical models of BER and OSNR are intractable, the *SiMPLE* algorithm uses real-time OPM measurements. The functions are unknown and the measurements are noisy, and therefore, evaluating derivatives via finite-differences is unreliable. Moreover, the algorithm should operate on a live network (restricting the type of points that can be evaluated) and evaluations are costly in terms of time and energy. As a result, most convex solution methods cannot be used. Hence, the *SiMPLE* algorithm is based on derivative-free optimization (DFO) methods [12] and computes a live configuration of the wavelengths' power levels throughout the optical network.

We evaluate the performance of the *SiMPLE* algorithm in two ways. First, we use a *network-scale* optical simulator that provides fine-grained control over the optical physical layer. We then evaluate the algorithm's performance in a small-scale testbed that is built using commercial optical devices. We demonstrate that desirable convergence which does not impact the network reliability can be obtained through a proper choice of algorithm parameters. Furthermore, we show that the ability to add and drop wavelengths dynamically (as enabled by *SiMPLE*) can lead to significant power savings.

We note that our approach requires deploying OPMs at many locations throughout the network, and continuous communication between the OPMs and a central controller. While migrating to such a deployment may increase costs, new devices that have integrated OPMs may facilitate this transition and reduce the CAPEX [26]. Furthermore, transceivers, amplifiers, and ROADMs already consume a large share of the power in modern optical networks. Through a better use of resources, fewer such devices may need to be deployed, as the ones that are already in place can be configured to dynamically address the demands. As such, dynamic optical networks can have a positive impact both in terms of CAPEX and OPEX.

In summary, the main contributions of this paper are three-fold: (i) we develop a measurement-based power control algorithm that enables the dynamic addition and removal of lightpaths anywhere in the network at any time, (ii) we evaluate the performance of the algorithm using an optical simulator and in an optical testbed, and (iii) we estimate the possible energy savings based on the simulation results. To the best of our knowledge, this is the first attempt at providing a global power control algorithm that uses live OPM measurements to enable dynamic optical networking. The proposed algorithm can support optical network control in near real-time while allowing the higher layer protocols to dynamically adapt to traffic patterns. In other words, *we take one of the first steps towards designing software-defined optical networks that may be able to reduce the energy consumption of the network backbone via better resource allocation.*

The rest of the paper is organized as follows. We introduce related work in Section II and present the model that captures the dynamics of a single link in Section III. In Section IV, this link model is generalized to the entire network and the MLO problem is formulated. The *SiMPLE* algorithm is introduced in Section V, followed by extensive evaluations via simulation and experimentation in Sections VI and VII, as well as the energy savings analysis. We conclude and discuss future work in Section VIII.

II. Related Work

Recent years have seen all-optical networks garner increased attention as a viable option for reducing the power consumption of data-transport networks [25]. A recent special issue of

the *Proceedings to the IEEE* [5] addresses several important problems in all-optical networks, such as reconfigurablity, optical flow switching, optical network control, and cross-layer impairment-aware optical networks [36]. The end goal is to realize an all-optical network, and the problem posed in this paper is one of the building blocks needed to achieve this goal.

Modifying traffic patterns in an operational optical network requires one to be aware of the physical network constraints, the QoT requirements, and Physical Layer Impairments (PLI). This is true whether impairment-aware RWA algorithms are used [6, 13] or decisions are made in an optical switching fabric [20, 28]. A lot of algorithmic research of optical network control has looked into efficient Routing and Wavelength Allocation (RWA) problems. Recent developments focus on finding routes by considering impairments [6], and on the reduction of overall energy consumption [39]. Our work is independent of the type of RWA algorithm used, and will take the network-wide output of such an algorithm to determine the appropriate per-wavelength power assignments, if it is feasible, on a per-link basis, in near real-time.

The energy consumption of the ITU sector recently received a lot of attention [25, 30]. For instance, the GreenTouch Consortium aims to reduce the carbon footprint of telecommunications equipments and networks by several orders of magnitude [3]. Research in this area considered energy-reduction strategies for switching off optical devices at times of low demand[22]. In [35], the authors propose RWA algorithms that minimize the energy consumption based on the current load in the network. In both approaches, the algorithms use analytical models rather than measurements, and therefore, are unlikely to be implemented as-is in real systems given the unpredictable nature of optical transmissions.

OPMs offer real-time inspection of transmissions [33] by measuring OSNR and BER. OPMs have been successfully applied for dynamic optical network control, e.g., switching in an optical switch using BER as a metric [28] and changing the modulation format in real-time according to link OSNR [16]. Another important work considers minimizing the sum of convex cost functions (e.g., wavelength powers) of a *single* link based on the OSNR constraints on individual wavelengths [32]. In this paper, we go beyond a single link, and discuss the solution of this problem for a network of optical links.

III. Optical Link Model

We now focus on a single optical link of a network. Such a link consists of several *spans* of fiber connected by various optical devices such as amplifiers. The signal originates at a node with a transponder and is amplified at intermediate nodes. The receiver at the destination decodes the signal. Source and destination nodes can be, for example, ROADMs or Optical Cross Connects (OXCs) that connect several links.

On a single fiber of a Wavelength-Division Multiplexed (WDM) network, several transmissions can take place on different wavelengths, as illustrated in Fig. 3. We denote by E the set of spans. Each span $u \in E$ supports a set of wavelengths denoted by $\Lambda(u)$. The following definition will be useful in refering to individual wavelengths of each span.

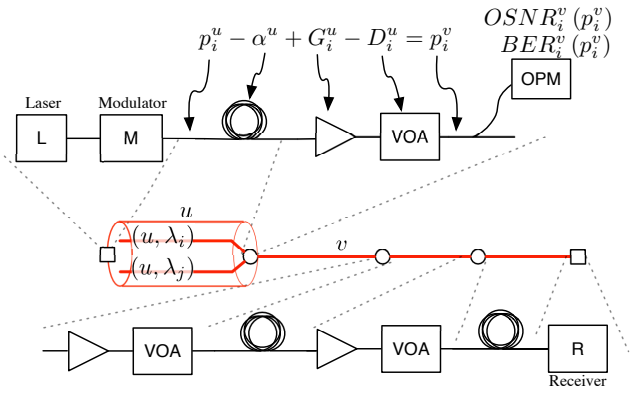

Fig. 3. The correspondence between a physical link and its spans, and our optical link model. The leftmost node includes an optical source composed of a laser (L) and a modulator (M). The intermediate nodes consist of amplifiers and Variable Optical Attenuators (VOAs). The rightmost span ends the link with a Receiver (R). OPMs can be located at any node.

Definition 1 (λ-span): A λ-span (u, λ_i) represents the transmission on fiber span $u \in E$ and wavelength $\lambda_i \in \Lambda(u)$.

Controllable parameters of the λ-spans include launch power, amplification, bandwidth, and modulation format. In this paper, *we focus on power control*[1]. Properties of a span, such as BER and OSNR, can be measured using an OPM.

A. Optical Power Dynamics

Each λ-span (u, λ_i) has an associated optical power-level p_i^u. All power levels are expressed in dBm. If the head of a span is a transmitter (laser), p_i^u is the power of the signal as it leaves the transmitter. If the head of the span is an amplifier, the power is the amplified signal power. During the transmission through the span, the signal power is first attenuated by a distance-dependent fiber loss α^u which is around $0.2\,\mathrm{dB/km}$ for single-mode fiber. The power at the receiving end of span is therefore $p_i^u - \alpha^u$ (in dBm).

At an intermediate node, the power can be modified in several stages, as shown in Fig. 3. The received signal $p_i^u - \alpha^u$ is first amplified by an amount G_i^u. The power can then be reduced using a variable optical attenuator (VOA) by a specified amount D_i^u. The launch power p_j^v of the signal at an intermediate node is

$$p_i^v = p_i^u - \alpha^u + G_i^u - D_i^u. \tag{1}$$

Depending on the network, different values of this expression will be the control variables. If a constant gain amplifier is used at λ-span (u, λ_i), the corresponding variable G_i^u will be a constant whose value is set to the gain of the amplifier. If the amplifier gain is configurable, G_i^u will be a decision variable that can be modified while adjusting the optical power levels. While these decision variables are real numbers, they can be rounded to the nearest value supported by the amplifier. Most optical amplifiers are ideally designed to amplify the entire spectrum by the same gain factor, i.e., $G_i^u = G_j^u$, for any two wavelengths i and j in the same span. If the power can be

[1]Extensions to other parameters such as modulation format and transmission wavelength will be considered in future work.

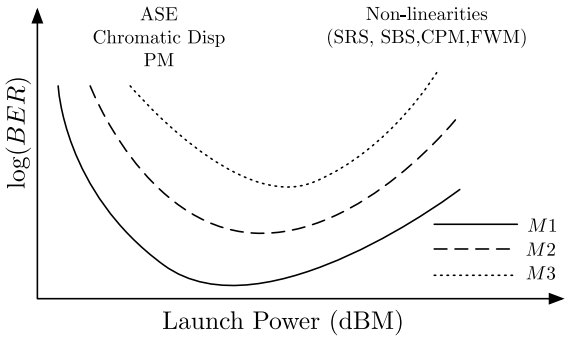

Fig. 4. Illustration of the relationship of BER and the optical launch power for a few modulation formats denoted $M1$, $M2$, $M3$ [23]. Prominent impairments in each region are marked (see [29] for detailed descriptions).

controlled at the launch of the λ-span, then the initial power p_i^u is a decision variable. Otherwise, it is a constant.

Regardless of the choice of parameters G_i^u and D_i^u, it is possible to express the power dynamics of the network as a function of the power variables, p_i^u. We will therefore write all future equations with respect to the λ-span power levels p_j^u, and use the notation $\mathbf{p}$ for the *power vector* of all power levels in the network.

B. Performance Measurements

There are direct relationships between power levels, BER and OSNR values, and these originate from the physical interactions of the optical transmission with the fiber. *As such, they are difficult to characterize analytically, but can be measured experimentally.* Fig. 4 provides an illustration of the relationship between the BER value and the launch power for a specific λ-span. The prominent impairments for different power levels are noted on the figure, and can be found in [29]. For instance, at low power levels, increasing the power improves the BER by mitigating the effects of Amplified Spontaneous Emission (ASE) noise. However, at higher power levels, increasing the power may negatively impact the BER, due to other, non-linear impairments such as Cross-Phase Modulation (CPM). The exact shape of Fig. 4 may depend on the characteristics of the fiber, amplifiers, and other equipment. Other factors, such as the used modulation format, temperature, component drift, aging, and fiber plant maintenance [33] affect the specifics of the curve, but the overall nature of the relationship remains the same [23].

In our setup, OPMs are used at the receiving end of a span to measure the quality of the transmission, including the BER and the OSNR. The BER and the OSNR metrics depend on the power p_i^u on λ-span (u, λ_i), and are denoted by $\mathrm{BER}_i^u(\mathbf{p})$ and $\mathrm{OSNR}_i^u(\mathbf{p})$, respectively, for a λ-span u.

There are no analytical expressions for BER and OSNR functions, due to the presence of many impairment factors. However, $\mathrm{BER}_i^u(\mathbf{p})$ is convex, while $\mathrm{OSNR}_i^u(\mathbf{p})$ is concave [18, 23]. We use our experimental testbed described in Section VII to show this fact. Fig. 5 shows a logarithmic plot of the $\mathrm{BER}_i^u(\mathbf{p})$ function when two lightpaths λ_1 and λ_2 on the same λ-span are gradually attenuated. The cuts along the x- and y-axes of this function yield convex curves similar to Fig. 4. However, this does not imply that the *multidimensional* function is convex over *both* variables. We

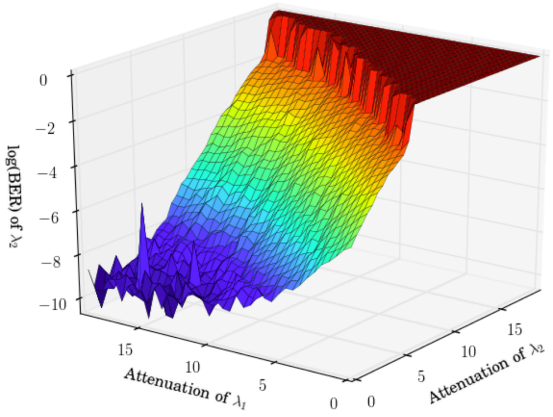

Fig. 5. The experimentally obtained 3-dimensional function that demonstrates the convexity of the BER curve of λ_2 with respect to the attenuation levels of λ_1 and λ_2.

verify this numerically by calculating the Hessian of this curve at all points, and making sure that it is positive semidefinite. This convexity property is leveraged in the next section to develop a network-wide power control algorithm.

We note that the problem formulations in Section IV and corresponding convergence results also hold under more general assumptions. Namely, when the BER and OSNR functions are not convex, but are *quasi-convex*[2]. Quasi-convex functions form a very large family that contains convex functions.

IV. Multi-Link Optimization Problem

In this section, the optical model for a single link introduced in Section III is generalized to the network setting and an optimization problem is formulated.

A. Network Model

We model the network as a directed graph (V, E). The nodes $v \in V$ represent ROADMs, OXCs, and amplifiers in which, it is possible to control the power and to perform measurements using OPMs.[3] The edges $u \in E$ are fiber spans between devices. In a WDM network, each fiber span can support several wavelengths which correspond to several λ-spans.

A lightpath P is a single optical stream of data that traverses several spans. Most lightpaths maintain the same wavelength throughout their route, although converters can be used to modify their wavelength along the route[29]. Lightpaths are represented as sequences of λ-spans $P = \{(u, \lambda_i), (v, \lambda_i), (w, \lambda_i), \ldots\}$.

As shown in Fig. 1, nodes can have several incoming links. At these locations, cross-connect devices such as ROADMs bridge the lightpaths from one span to another[34]. The assignment of routes and wavelengths to links is out of scope, as these are assumed to be handled by an RWA algorithm [6].

Not all λ-spans have all the capabilities introduced in Section III. We denote by Λ_{BER} and Λ_{OSNR} the sets of λ-spans that are equipped with the OPMs that measure BER and OSNR, respectively. Similarly, the set Λ_p correspond to the sets of λ-spans that have the ability to control the power.

[2]A function is said to be quasi-convex if its sublevel sets are convex.

[3]An optical link between two regional offices that includes several amplifiers is modeled as a path of several nodes.

B. Optimization Problem

The key requirement of network operators, as specified by their Service Level Agreements (SLAs), is to maintain the BER within a certain threshold value. Any network changes that are performed should also satisfy this requirement. Since network operators are unable to continuously adjust the power levels of the lightpaths in response to impairments, they typically compute an offline solution with added margins to the BER requirements, which leads to over-provisioning. While this approach works when network demands are largely static, with traffic variations seen in today's networks (e.g., diurnal patterns for video consumption, inter-data center traffic [21, 24]), a dynamic approach that can continuously guarantee BER requirements while adjusting to traffic demands is needed.

The *Multi-Link Optimization (MLO)* problem represents this requirement as a relationship between the desired threshold levels and the current outputs of the OPMs, provided by the BER and OSNR functions. The control variables are the power levels that need to be adjusted to change the OSNR or BER values. There can be several possible configurations that provide this guarantee, and the one that consumes the least amount of optical power is considered through the objective function $h_{\mathrm{MLO}}(\mathbf{p}, \mathbf{D})$. Based on the linear relationship of eq. (1), the power levels are adjusted by either minimizing the power levels $\mathbf{p}$ or by maximizing the attenuations $\mathbf{D}$. The formulation for this optimization problem is as follows.

Problem 1 (Multi-Link Optimization - MLO):

$$\underset{\mathbf{p},\mathbf{D}}{\text{minimize}} \quad \sum_{(u,\lambda_i)} p_i^u - \sum_{(u,\lambda_i)} D_i^u =: h_{\mathrm{MLO}}(\mathbf{p}, \mathbf{D}) \tag{O1}$$

$$\text{subject to } \mathrm{BER}_i^u(\mathbf{p}) \leq \overline{\mathrm{BER}_i^u}, \quad \forall (u,\lambda_i) \in \Lambda_{\mathrm{BER}} \tag{C1}$$

$$\mathrm{OSNR}_i^u(\mathbf{p}) \geq \overline{\mathrm{OSNR}_i^u}, \quad \forall (u,\lambda_i) \in \Lambda_{\mathrm{OSNR}} \tag{C2}$$

$$0 \leq \mathbf{p} \leq \mathrm{SAF}, \tag{C3}$$

where $\overline{\mathrm{BER}_i^u}$ and $\overline{\mathrm{OSNR}_i^u}$ are the respective performance thresholds on λ-span (u, λ_i), and SAF is the limit on the link power due to safety restrictions. Note that in this formulation, the objective function $h_{\mathrm{MLO}}(\mathbf{p}, \mathbf{D})$ effectively minimizes the difference of the power and the attenuation throughout the network. Bounding the power levels by the safety parameter SAF in constraint (C3) assures that the power levels are always reasonable. This effectively eliminates corner cases in which the difference can be small, while both the power levels $\mathbf{p}$ and the attenuation $\mathbf{D}$ are very large. We refer to a power vector $\mathbf{p}$ as *feasible*, if constraints (C1)-(C3) are all satisfied.

When an RWA algorithm needs to add a lightpath, the *MLO* formulation can be modified by adding constraints for the new λ-spans. To remove a lightpath, constraints involving the affected lightpath can be removed progressively. In the same manner, modifications in the threshold values for some lightpaths can be executed by changing the $\overline{\mathrm{BER}_i^u}$ and $\overline{\mathrm{OSNR}_i^u}$ parameters.

The *MLO* problem is convex due to the nature of the OSNR and BER functions, similar to the single-link case. They are also zero order oracle problems [12] because their analytical functions and first-order derivatives are unavailable (see Section III-B).

V. Power Control Algorithm

In this section, we present the Simultaneous Multi-Path Lambda Enhancement (*SiMPLE*) Algorithm that uses the characteristics of the *MLO* problem to solve it efficiently.

Computing an optimal solution for the convex *MLO* problem is not straightforward. The functions $\mathrm{BER}(\mathbf{p})$ and $\mathrm{OSNR}(\mathbf{p})$ can be evaluated for given points but their overall curves are unknown. Each evaluation of a performance function requires using an OPM device which is expensive both in terms of time and energy. The measurement process can be disruptive to existing traffic in the network, and the OPM readings may contain noise. Therefore, a general-purpose convex solver cannot be used for the solution of the *MLO* problem, and these special characteristics must be accounted for during the design of a novel algorithm.

We denote by $\mathbf{p}(k)$ the value of the power vector at iteration k (contrasted with p_i^u which is the power of λ-span (u, λ_i)). Similarly, the measurements from all the OPMs at iteration k are captured by vectors $\mathrm{BER}(\mathbf{p}(k))$ and $\mathrm{OSNR}(\mathbf{p}(k))$.

A. Design Considerations (DCs)

The requirement for the *SiMPLE* algorithm to operate in a live production network has several important implications:

(DC1) An algorithm that solves the *MLO* problem needs to evaluate the BER/OSNR functions at intermediate power levels while converging to a solution. To obtain these intermediate power levels, the amplifiers and attenuators need to be tuned throughout the network. A proper algorithm therefore must guarantee that these intermediate evaluations do not cause any disruptions to active lightpaths. Specifically, the algorithm needs to perform small perturbations and ideally the result of the most recent iteration needs to be the best one.

(DC2) For the network operator, it is more important to adhere to the SLA requirements than to find a setup with the optimal optical power allocation. Since satisfying the constraints (C1)-(C2) is the priority, the main aim of the algorithm is to obtain a feasible solution as quickly as possible. Once a feasible solution is reached, the algorithm must guarantee that the subsequent steps do not cause any of the feasible constraints to be violated by a large amount.

(DC3) Most convex optimization solvers need the first or second derivatives of the functions used in the problem formulation[31]. However, these methods are not appropriate for cases where the functions are not known and can be noisy. Therefore, derivative-free optimization (DFO) algorithms [12] are the most suitable solution methods.

B. SiMPLE Algorithm

We begin with a high-level overview of the *SiMPLE* algorithm. This algorithm is based on a constrained direct-search algorithm [12], and incorporates the design considerations discussed in Section V-A. Starting from a power vector $\mathbf{p}(0)$, the algorithm makes small changes to the power levels according

to a set of computed direction vectors, and evaluates the BER and OSNR functions. If there are two devices whose power can be modified, this set could be as simple as

$$\mathcal{D}_k = \left\{ \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix} \right\}, \forall k.$$

To improve convergence, search directions are generated dynamically as the set $\mathcal{H}_k$ according to a number of heuristic rules, denoted by H1-H3. When a direction that improves the current point is found, the next iteration begins. If an improvement direction is not found, the search starts over from the same point, with a smaller step size. The value of this step size variable α_k is changed throughout the run of the algorithm according to the parameters $(\theta-, \theta^+)$. These parameters have a large effect on the convergence properties, as shown in Section VI-D. An illustration of the high-level operations of *SiMPLE* for a one-dimensional case is shown in Fig. 6.

The pseudocode for the *SiMPLE* algorithm is shown above. It takes as input an instance $\mathcal{I}$ of the *MLO* problem and an initial power assignment $\mathbf{p}(0)$. The problem instance $\mathcal{I}$ corresponds to a set of constraints to *MLO*, as determined by the higher-layer algorithms and SLAs.

The first step is to create an augmented objective function $f(\mathbf{p}(k))$ by incorporating the constraints (line 6). Depending on the feasibility of the current point $\mathbf{p}(k)$, one of functions AUGMENTLOG or AUGMENTQUAD, defined below, is used.

If the initial point is feasible, the subsequent power levels must stay feasible for the remaining iterations (this is due to DC1). This is guaranteed by AUGMENTLOG, which returns the following log-barrier function:

$$\begin{aligned} f(\mathbf{p}(k);\mu) = h_{\mathrm{MLO}}(\mathbf{p}(k)) \\ - \frac{1}{\mu} \sum_{(u,\lambda_i)\in\Lambda_{\mathrm{BER}}} \log(\overline{\mathrm{BER}}_i^u - \mathrm{BER}_i^u(\mathbf{p}(k))) \\ - \frac{1}{\mu} \sum_{(u,\lambda_i)\in\Lambda_{\mathrm{OSNR}}} \log(\mathrm{OSNR}_i^u(\mathbf{p}(k)) - \overline{\mathrm{OSNR}}_i^u), \end{aligned}$$

where μ is parameter of the augmentation function[31]. Since the performance functions $\mathrm{BER}(\cdot)$ and $\mathrm{OSNR}(\cdot)$ are embedded in $f(\cdot)$, each evaluation of this function causes the OPMs to make a measurement. With this augmented function, *SiMPLE* can try power levels that violate the thresholds, as there is no knowledge of the feasible region boundary. However, evaluations outside the feasible region yield infinite values under the logarithm, and such points will not be accepted for the next iteration.

If the initial power levels are not feasible, the priority is to find a feasible point (due to DC2). The AUGMENTQUAD function returns an augmented function that is finite for infeasible points, and forces the points $\mathbf{p}(k)$ to feasibility:

$$\begin{aligned} f(\mathbf{p}(k)) = \sum_{(u,\lambda_i)\in\Lambda_{\mathrm{BER}}} \left(\left[\mathrm{BER}_i^u(\mathbf{p}(k)) - \overline{\mathrm{BER}_i^u} \right]^+ \right)^2 \\ + \sum_{(u,\lambda_i)\in\Lambda_{\mathrm{OSNR}}} \left(\left[\overline{\mathrm{OSNR}_i^u} - \mathrm{OSNR}_i^u(\mathbf{p}(k)) \right]^+ \right)^2, \end{aligned}$$

Algorithm 1 Pseudocode of *Simultaneous Multi-Path Lambda Enhancement (SiMPLE)*

1: **Input:** Problem instance $\mathcal{I}$ and initial power levels $\mathbf{p}(0)$
2: **Parameters:** θ^+, θ^-, and α_{tol}
3: **loop**
4: $\alpha_k \leftarrow 1$
5: **repeat**
6: **if** $\mathbf{p}(k)$ is feasible **then**
7: $f(\mathbf{p}(k)) \leftarrow$ AUGMENTLOG($\mathcal{I}$)
8: **else** $f(\mathbf{p}(k)) \leftarrow$ AUGMENTQUAD($\mathcal{I}$)
9: **end if**
10: $\mathcal{H}_k \leftarrow$ GENERATE($\mathbf{p}(k)$) ; $\mathcal{D}_k \leftarrow \mathcal{H}_k \cup \mathcal{G}$
11: // Try directions in $\mathcal{D}_k$
12: **if** $\exists d_i \in \mathcal{D}_k$ with $f(\mathbf{p}(k) + \alpha_k d_i) < f(\mathbf{p}(k))$ **then**
13: $\mathbf{p}(k+1) \leftarrow \mathbf{p}(k) + \alpha_k d_i$; $\alpha_{k+1} \leftarrow \theta^+ \alpha_k$
14: **else** $\mathbf{p}(k+1) \leftarrow \mathbf{p}(k)$; $\alpha_{k+1} \leftarrow \theta^- \alpha_k$
15: **end if**
16: **until** $\alpha_k \leq \alpha_{\mathrm{tol}}$
17: **end loop**

where $[x]^+ = \max(x, 0)$ is the positive projection of x. Under this function, infeasible power levels will evaluate to finite values. Yet, reducing the infeasibility decreases the function value. If the thresholds are attainable, the power levels are forced within the feasible region. Note that unlike other qudratic augmenting functions (e.g., [31]), the priority is to reach a feasible solution, and therefore, the objective function is not captured in this augmentation procedure.

It is also possible to selectively use AUGMENTQUAD and AUGMENTLOG for even more fine-grained control over feasibility of individual constraints. For instance, when a lightpath is first added, the corresponding new constraints will be infeasible by design. The AUGMENTQUAD function can be applied to the new constraints, and *SiMPLE* will progressively try to make that constraint feasible. Meanwhile, the AUGMENTLOG function can be applied to the constraints that are already feasible. This way, the lightpaths that are already provisioned are guaranteed to remain feasible. If the threshold of these constraints are violated even by a small margin in an intermediate iteration, the logarithm function will evaluate to infinity, thereby disallowing the acceptance of this point.

The direct search step in line 11 tries several directions d_k from a search set $\mathcal{D}_k$ to improve the objective function value. Each of these directions is tried with step size α_k. For each d_k of this set, it changes the power levels, and collects the OPM measurements. This search set consists of the union of two sets. $\mathcal{G}$ is a positive spanning set of the entire search dimension space, which means that for all $v \in \mathbb{R}^n$, there exists $\eta_k \geq 0$ such that $v = \sum \eta_k g_k$ with $g_k \in \mathcal{G}$. We use the columns of the block matrix $\mathcal{G} = \begin{bmatrix} I; -I \end{bmatrix}$ where I is the identity matrix. This condition guarantees that all points in the search space are reachable through a positive linear combination of these vectors and is crucial for the proof of convergence of direct search methods. Note that in this approach, neither the full BER curves, nor their derivatives are used, satisfying DC3.

The function GENERATE returns a set $\mathcal{H}_k$ of additional search directions based on the current and previous iterations. These directions are checked first, since they are more likely to be descent directions. There are several ways to obtain the

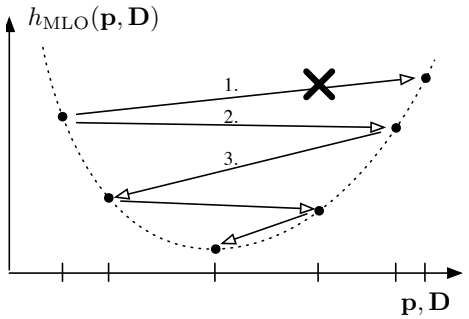

Fig. 6. Illustration of *SiMPLE*'s convergence in one dimension. The objective function is dashed to indicate that its realization is unknown. The first step (marked "1.") achieves a larger objective value and is discarded. The step size α_1 is reduced by θ^-, and the second iteration is successful. For third step, the "rightward" direction $d_i \in \mathcal{H}_i$ is attempted first (not shown), but the "leftward" direction is picked instead.

set of search directions $\mathcal{H}_k$, and we consider *three heuristics*:

- H1: $\mathcal{H}_k = \varnothing$, for comparison to the other methods,
- H2: $\mathcal{H}_k = d_{k-1}$, the last successful search direction,
- H3: $\mathcal{H}_k = d_{k-1}$ and a set of points around d_{k-1}. This corresponds to searching around $\mathbf{p}(k) + d_{k-1}$ in addition to searching around $\mathbf{p}(k)$.

These heuristics vary in the size of the set $\mathcal{H}_k$ they produce. If the search set contains many direction vectors, the likelihood of one of them being a descent direction is higher. However, larger sets will also result in wasted OPM evaluations, if none of the directions are viable, and the right strategy is to reduce the step size. The effects of the choice of heuristics are explored further in Section VI-D.

Once the directions are exhausted, there are two possible outcomes. If a descent direction is found, the step size parameter α_k is multiplied by $\theta^+ \geq 1$ to try a larger step in the next iteration. If no successful search direction is found, the step size is multiplied by $0 < \theta^- < 1$, and a smaller one is tried.

The inner loop exits when the step size α_k is reduced below a tolerance value α_{tol}. If AUGMENTQUAD was used as an augmentation function, a feasible power level is found at the end of the run (if such a value exists), and all the performance thresholds are met. The algorithm then restarts, using AUGMENTLOG to further optimize this new point.

The SiMPLE algorithm runs continuously in an outer loop and constantly optimizes the solution. If the MLO Problem constraints change (e.g., due to the addition or removal of a lightpath), these changes are reflected to the problem instance $\mathcal{I}$. A new penalty function $f(\cdot)$ is constructed by the appropriate augmentation function, and the algorithm is initiated again from its last successful point $\mathbf{p}(k)$.

As mentioned in Section III-B, the *MLO* formulation can be used under more general quasi-convexity assumptions regarding the $\mathrm{BER}_i^u(\mathbf{p})$ and $\mathrm{OSNR}_i^u(\mathbf{p})$ functions. In that case, the *SiMPLE* algorithm can be modified by using a standard method for solving quasi-convex optimization problems given in [9, Section 4.2.5]. This framework effectively transforms the quasi-convex optimization problem into a sequence of convex problems that are solved repeatedly. The downside of this method is that it may require a large number of evaluations.

The *SiMPLE* algorithm is unique in that it uses two different augmentation procedures to satisfy its design considerations.

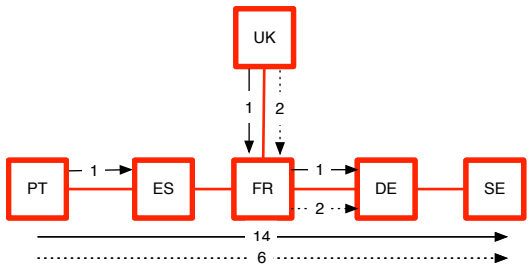

Fig. 7. The topology of a subnetwork of the Géant network (see Fig. 1) that was used for the evaluation of the *SiMPLE* algorithm. The solid lines indicates fibers between nodes. The solid arrows indicate the location and number of lightpaths that are present in the low-capacity assignment. The dotted arrows indicate the location and number of additional lightpaths that are provisioned in the high-capacity assignment, when extra capacity is needed.

Arguments for its convergence can be found in Appendix A.

VI. SIMULATION EVALUATION

In this section, we evaluate the performance of the *SiMPLE* algorithm and demonstrate the benefits of dynamic optical networks with regards to energy efficiency.

A. Evaluation Metrics

In order to evaluate the performance of the heuristics H1-H3 under different parameters and noise scenarios, we introduce several metrics. These metrics correspond to the objective of minimizing the disruptions and power fluctuations and reaching the target power levels as quickly as possible. We let $\mathcal{P} = \{\mathbf{p}(k)\}_{k\geq 0}$ denote the set of all power vectors throughout the run of the algorithm.

The running standard deviation (RStd) corresponds to the variability of the power levels. It is obtained by first finding the running average of the last 20 evaluations of the power vectors $\mathbf{p}(k)$. The standard deviation from this running average is then computed as follows:

$$\mathrm{RStd}_{\mathbf{p}}(k) = \mathrm{Std}\Big(\mathbf{p}(k) - \sum_{j=k-20}^{k} \mathbf{p}(j)/20\Big),$$

where Std is the standard deviation operator. The FeasTime metric measures the time until all the constraints are satisfied and the solution is feasible. It is measured in the number of iterations, and is defined as follows:

$$\mathrm{FeasTime}(\mathcal{P}) := \min_{\mathbf{p}(i)\in\mathcal{P}} \left\{ i : ||\,\mathrm{BER}(\mathbf{p}(i)) \leq \overline{\mathrm{BER}} \right\}.$$

The duration of a single iteration can be analyzed by separating it into three phases: measuring the BER/OSNR, computing the result, and making the appropriate modification. With current technology, it is possible to have an OSNR measurement in the order of microseconds (a photodiode takes the measurement in nanoseconds, and an FPGA processes in microseconds). Using a photonic integrated device, it is possible to have 80 such measurement on a single chip for obtaining a reading across all the wavelengths[11].

The processing time would currently be in the order of a hundred milliseconds due to the centralized nature of our architecture. Through the use of distributed algorithms, we believe it is possible to reduce this to the order of milliseconds.

The actuation phase of changing the power levels can be done almost instantaneously, within picoseconds. Under the current architecture, we therefore estimate the duration of a single iteration to be in the order of 100s of milliseconds, with the prospect of a dramatic decrease through distributed computation.

Finally, the feasibility probability $\mathrm{FeasProb}$ is defined as the probability that *SiMPLE* finds a feasible solution to the given problem. It is computed as the ratio between the number of simulation runs that result in a feasible assignment and the total number of simulations runs.

For each of these metrics, we collect the result of *every measurement*, even if these measurements are not selected as the optimal point of an iteration. This is in contrast to most evaluations of convex algorithms where the number of *iterations* until convergence is used as a benchmark for computational complexity. In our setting, the measurement and actuation overheads of each OPM dominate the running time rather than the operations of the algorithm.

B. Simulation Setup

We developed a simulator which is based on a detailed physical model of an optical amplifier developed at Bell Labs [10]. The network level functionality was written in Python, and the code was designed to run in a parallelized manner on a computing cluster. The simulations were executed on an 8-core virtual machine running on the Amazon EC2 system.

This optical network simulator models a large-scale WSON that contains lasers, receivers, and ROADMs in a mesh topology. Many concurrent transmissions can take place across several lightpaths, and the optical power levels can be measured at every span of the lightpaths. The OSNR is estimated at the receiver by comparing the received signal power with the noise floor.

In the simulations, a subset of the Géant network topology was considered. It consists of 6 nodes and 27 lightpaths that follow four routes as shown in Fig. 7. Between the nodes, lightpaths go through several spans separated by ROADMs as shown in Fig. 3. The optical power of each lightpath can be modified at ROADM nodes on their path using an attenuator (VOA). The same attenuation level is used for all the ROADMs during the entire lightpath. The dimension of the optimization problem is therefore 7, one for each lightpath group in Fig. 7. The received power levels are measured at each destination node, amounting to a total of 4 OPMs. The *MLO* formulation has 7 OSNR constraints, one for each lightpath group. Gaussian noise with different standard deviation values was added to the OPM evaluations to emulate measurement noise. In terms of the parameters, we set the launch power of the transceivers to 20 dBm and the gain of the amplifiers to 15 dB. The OSNR thresholds are 20 dB. A tolerance value of $\alpha_{\mathrm{tol}} = 0.5$ was used to terminate the simulation runs.

C. Traffic Data

Traffic demand between each pair of cities in the Géant network obtained at 15 minute intervals for a four month period in 2005 is available in [37]. This data was averaged over a week-long period to obtain traffic variations for each weekday, for all the nodes that are part of the subtopology. A sample of the traffic variations for two specific nodes over a week is shown in Fig. 8. The data was scaled tenfold to accommodate the traffic growth statistics based on [4]. Optical networks are also typically overprovisioned to account for bursts in traffic, so the data was further scaled by a factor of five.

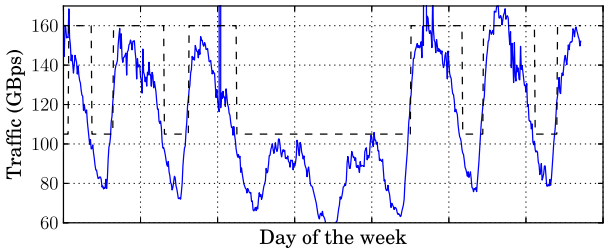

Fig. 8. Sample traffic pattern between two cities in the Géant network over a period of one week. The dashed line corresponds to the required capacity to satisfy the demand. In peak times, additional lightpaths are needed to provide higher capacity.

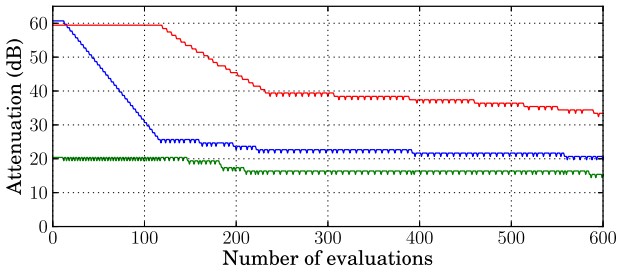

Fig. 9. Evolution of the attenuation of three lightpaths in our simulation while the *SiMPLE* algorithm transitions from the low-capacity assignment to the high-capacity one in order to satisfy the extra demand. Two lightpaths are progressively added by decreasing their attenuation.

Many approaches can be used to generate the optimal wavelength assignment that satisfy these traffic demands. In the most complex case, a new assignment can be computed in real-time using the live traffic matrices. Since the focus of this paper is not RWA algorithms, two specific assignments are designed to satisfy the traffic demands. In particular, based on the physical topology of Fig. 1, we define two *virtual* topologies referred to as the *high-capacity* and *low-capacity* assignments. The high-capacity assignment has enough capacity to satisfy traffic demands at peak times, while the low-capacity one is used when the demand is low on nights and weekends. The capacities for two specific source and destination nodes are illustrated in Fig. 8 as the dashed line that covers the traffic demands. The objective of the *SiMPLE* algorithm is to switch between such two assignments to optimize resource usage, as desired in dynamic optical networks. While this approach seems simple, it already provides a vast improvement over current optical networks, in which such drastic changes rarely occur over timescales less than the order of weeks.

D. Simulation Results

Parameters and Heuristics: For a given network deployment, two types of parameters should be considered: the (θ^-, θ^+) parameters for adjusting the step size (line 2 of

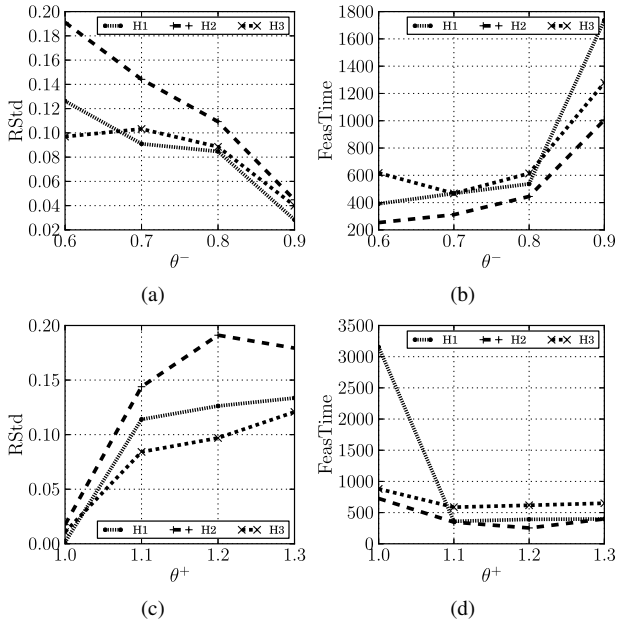

Fig. 10. Values of the RStd and FeasTime metrics that demonstrate the fluctuations and convergence time of the *SiMPLE* algorithm when switching from the low capacity to the high capacity assignment. Two sets of parameters are plotted: (a)-(b) θ^+ is fixed and θ^- is modified, and (c)-(d) θ^- is fixed and θ^+ is modified.

Algorithm 1), as well as the search direction heuristics H1-H3. We ran extensive simulations on the Géant subtopology to evaluate the effects of these parameters and heuristics on the convergence of *SiMPLE*. Specifically, we simulated the transition from the low capacity assignment used during nights and weekends, to the high capacity assignment used during peak times. In this setup, only the minimally necessary lightpaths are initially turned on. Additional lightpaths that can support the peak traffic are off, and therefore, have a very low OSNR. A problem instance is created for this scenario that requires all the lightpaths to have high OSNR. This problem is used as an input to *SiMPLE* which instructs the attenuators to bring up the additional lightpaths, while monitoring the other lightpaths.

Fig. 9 shows the attenuation evolution over time for a single run of simulation. In this setup, heuristic H1 was used with $\theta^- = 0.6$ and $\theta^+ = 1.2$. The green lightpath is initially active, while the red and blue lightpaths are being provisioned. The *SiMPLE* algorithm progressively decreases the attenuation of these lightpaths until the OSNR constraints are satisfied. The process of adding these lightpaths takes around 400 OPM evaluations. However, it can be seen that this process causes fluctuations in the power levels.

To understand the fundamental tradeoff between fluctuations and convergence speed, we repeated this experiment and *averaged the results over 250 runs for each parameter and heuristic combination*. The results are shown in Fig. 10.

Figs. 10(a)-10(b) show the effect of the θ^- parameter when $\theta^+ = 1.2$. One can notice that smaller values of θ^- cause larger fluctuations, as measured by the running standard deviation metric (RStd). However, these small values also decrease the convergence time dramatically by almost 80% compared to larger values. These observations are explained by recalling the definition of θ^- which affects the amount by which the search directions are reduced in unsuccessful iterations. Large reductions (e.g., when $\theta^- = 0.6$) cause large fluctuations. However, the algorithm can also adapt faster to the topology of the feasible region to reduce convergence times.

The results for the three heuristics are also plotted on Figs. 10(a)-10(b). The tradeoff between fluctuations and convergence time is also present among the heuristics. Namely, H3 achieves convergence at the expense of larger fluctuations. Recall that the H3 heuristic has a large set of candidate directions that is considered at every iteration. These trials create larger variations, but improve the algorithm capability to find the correct direction. Heuristic H2 adds an extra direction to H1, but its benefit over H1 is not evident in this figure.

Figs. 10(c)-10(d) illustrate the metrics for different values of θ^+, when $\theta^- = 0.6$. The parameter θ^+ governs the amount by which the search directions are increased after a successful iteration. Therefore, larger values of θ^+ do increase the fluctuations, as the step sizes become larger. However, we surprisingly find that there is no significant improvement in the convergence time with larger values. This is because the inability to take large steps is not the main bottleneck. Increasing the step size therefore plays a smaller role, since *SiMPLE* operates close to the boundary of the feasible region.

The previous observations on the heuristics also hold when observing the effects of θ^+. H3 achieves better convergence times, compared to the other two. One interesting point to note is the very large convergence time in Fig. 10(d) for H1. If the step size is not allowed to increase (when $\theta^+ = 1$), H1 takes an unusually long time to converge, as it takes many small steps. The other two heuristics avoid this by choosing their search directions more intelligently, thereby making better progress.

Noise: OPMs deployed in a real network will suffer from measurement noise as part of their operation. Furthermore, faster OPMs will have larger measurement noise. To operate in realistic scenarios, the *SiMPLE* algorithm needs to perform well in a noisy environment. It is also important to understand the magnitude of the tolerable noise. We applied our evaluation metrics to different noise conditions, and plotted them across the three heuristics in Fig. 11. Each point on this plot is averaged over 250 iterations as in Section VI-D.

The probability that *SiMPLE* will reach a feasible solution is plotted in Fig. 11(a). For reasonable noise levels where the variance of the Gaussian noise is less than 10^{-1}, the algorithm finds a feasible solution with high probability (greater than 90%). This behavior is independent of the choice of heuristics.

Fig. 11(b) shows the running standard deviation RStd of the power levels, with respect to the noise variance. It can be observed that large noise variations cause large power fluctuations. This suggests that when picking the θ^- and θ^+ parameters, it is also important to factor the noise performance of the OPMs. Also note that for very large noise values for which *SiMPLE* does not find a feasible solution, the fluctuations in the network are small.

E. Energy model

A major motivator for dynamic optical networks is better use of resources, specifically with respect to energy consump-

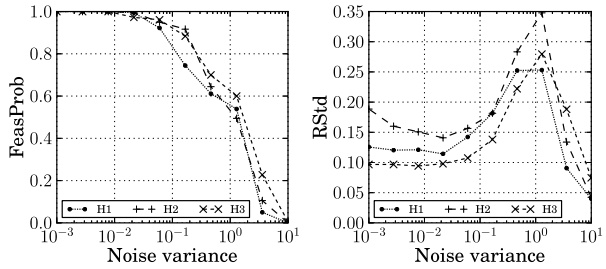

(a) (b)

Fig. 11. The performance of *SiMPLE* as a function of the measurement noise variance, for two metrics: (a) the probability that *SiMPLE* finds a feasible solution (FeasProb), and (b) the running standard deviation (RStd).

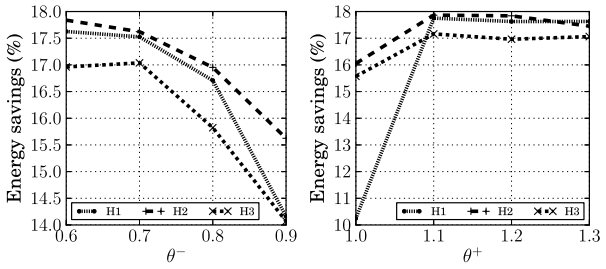

(a) (b)

Fig. 12. The potential energy savings provided by the *SiMPLE* algorithm for the heuristics H1-H3, when (a) $\theta^+ = 1$ and θ^- is varied, and (b) $\theta^- = 0.6$ and θ^+ is varied.

tion. In this section, we analyze the benefits of the *SiMPLE* algorithm in terms of the energy consumption of the network equipment. The topology outlined in Fig. 7 is considered and the traffic data that will satisfy the needs based on the time of day is used. We begin by analyzing the power consumption for the low-capacity and high-capacity assignments.

The power consumption of both the electrical and the optical components of a network stand to gain from dynamic behavior. However, the WDM portion of the transmission network typically consumes much less power compared to the electrical components. Indeed, the power density of a ROADM is about three orders of magnitude less than that of a core router[19]. Furthermore, the relationship between minimizing the optical power through the $h_{\mathrm{MLO}}(\mathbf{p}, \mathbf{D})$ objective function and the power consumption of optical devices is nontrivial. In general, it is not necessarily true that reducing the optical output power of a device translates into a power saving, and developing devices that exhibit this behavior is an active research area [3]. Due to these reasons, we focus our analysis on the benefits of the dynamic operation over the electrical components only, noting that this is effectively a lower bound analysis.

We assume a network based on the Cisco CRS-1 routers [1] . The optical transmission originates at WDM optical linecards each consuming 500 W. Depending on the number of wavelengths used at each node, we accounted for a 4-, 8-, or 16-shelf chassis, consuming roughly 2200 W each. For certain nodes, a multi-shelf system consisting of two 16-shelves were necessary to satisfy the high traffic demands.

Based on these estimates, the high capacity assignment consumes 47 kW when all the linecards and chassis are active. For a static network, this would be the permanent power consumption of this network. At times of low demand, a portion of this equipment is turned off, and the low capacity assignment consumes 32 kW.

Based on the traffic patterns described in Section VI-C, on average, the high capacity assignment is needed 41% of the time. A first order approximation would conclude that the average power consumption of a dynamic optical network in this scenario would be 38 kW, resulting in a saving of 19%.

However, the previous reasoning corresponds to the best case energy savings, in which *SiMPLE*'s overhead is not accounted for. In reality, each time the wavelength assignment needs to change, energy is spent on the convergence of the algorithm. Furthermore, the amount of energy spent depends on the parameters and heuristics. The worst-case energy overhead of the *SiMPLE* algorithm can be computed as follows:

$$E_{\mathrm{overhead}} = \frac{1}{p_{\mathrm{conv}}} \times t_{\mathrm{conv}} \times t_{\mathrm{round}} \times P_{\mathrm{diff}},$$

where p_{conv} is the probability of convergence measured by the metric FeasProb, introduced in Section VI-A. Therefore, the expressions $1/p_{\mathrm{conv}}$ correspond to the expected number of rounds the algorithm will take until convergence. Each such round of the algorithm takes on average t_{conv} evaluations, measured by the metric FeasTime and plotted in Fig 10. The duration for the OPM evaluation, the message transmission duration to the central server, and the actuation is taken to be $t_{\mathrm{round}} = 5\,\mathrm{s}$, which is fairly realistic by the standards of current equipment. Finally, $P_{\mathrm{diff}} = 15\,\mathrm{kW}$ is the difference in power consumed between the two assignments that is wasted during the convergence of the algorithm.

Fig. 12 shows the energy savings in the simulated topology due to the *SiMPLE* algorithm, as a percentage of the energy consumed by the static network. The savings are plotted with respect to different (θ^+, θ^-) parameters, as well as heuristics H1-H3. It can be observed that the effect of θ^+ on the energy savings is insignificant compared to the effect of θ^-. On the other hand, the right choice of θ^- can lead to greater energy savings (18%). Note that these large values come at the expense of large variations in the power levels, as indicated in Fig. 10(a) which may negatively affect the availability of the network by causing unpredicted downtime. Therefore, the choice of parameters plays a critical role in obtaining a desirable convergence behavior.

A key value that affects these savings is the measurement and actuation time t_{round}. While we assume that $t_{\mathrm{round}} = 5\,\mathrm{s}$, advances in OPMs and actuation technology is expected to bring this value further down, which will reduce the overall energy consumption. For example, decreasing t_{round} to 1 s will reduce the energy consumption by 18% for all (θ^-, θ^+) values. This is within 1% of the ideal savings of 19%.

VII. Testbed Evaluation

A small-scale testbed was built using commercial optical devices to evaluate *SiMPLE* in a real-life scenario.

A. Experimental Setup

The experimental setup is shown in Fig. 13. Eight closely-spaced Continuous-Wave (CW) laser sources in the ITU C-band are partitioned into two contiguous wavebands–Λ_{long}

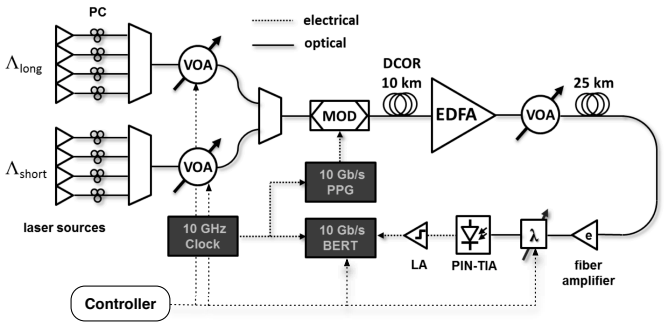

Fig. 13. Schematic of the experimental setup. Two wavebands each consisting of four wavelengths are modulated with a PRBS pattern and are transmitted via an amplified link segment. The optical powers of each waveband are controlled by a computer configured as a controller that runs the *SiMPLE* algorithm.

and Λ_{short}–of 4 wavelengths each. A $2^{15}-1$ Pseudo-Random Bit Sequence (PRBS) is inscribed using a single intensity modulator to produce a 10 Gb/s On-Off-Keyed (OOK) Non-Return-to-Zero (NRZ) pattern on each channel. In order to provide individual fine-tuned control, the injected power of each waveband is independently set via Variable Optical Attenuators (VOA) preceeding the optical modulator.

Inter-channel impairments are induced through a segment consisting of an Erbium-Doped Fiber Amplifier (EDFA), VOA, and 25 km of single-mode fiber. The EDFA is tuned to operate in saturation within a subrange of the operating powers of the incident wavebands. As a result, at significantly high powers, each waveband will "steal gain" from the other, resulting in mutual degradation.

Representative channels in each waveband are isolated for observation using a tunable grating filter (λ). Each data stream is recovered using a photoreceiver assembly with an inline digitizer (comparator) and subsequently fed into a Bit-Error-Rate Tester (BERT). Using the BERT, we can quantify the effect of the experimentally induced impairments on each waveband to not only characterize the parameter space of our system, but to serve as a real-time performance metric utilized by the experimental implementation of *SiMPLE*.

The *SiMPLE* algorithm runs in an automated fashion on a laptop configured as a central controller. This controller interfaces with the VOAs, tunable optical filter, and BERT using the IEEE-488 General Purpose Interface Bus (GPIB) interface. Through this interface, the algorithm collects measurements from each data stream and iteratively modifies the attenuation levels of the wavebands.

B. Results

To demonstrate the effects of the (θ^-, θ^+) parameters, we ran an experiment in our testbed corresponding to a series of changes in a dynamic optical network. This experiment was repeated for various sets of θ^- and θ^+ parameters. Fig. 14 shows the result for two different sets of θ^- and θ^+ parameters. The experiment has several stages, as labeled in Fig. 14(a). First, the red lighpath is added similarly to the simulation scenarios (see Section VI-D). Then, after 160 evaluations (marked as step 2 in Fig. 14(a)), the threshold constraint $\overline{\text{BER}}$ for the blue lightpath is significantly increased ($\overline{\text{BER}} = 10^{-1}$). As a result,

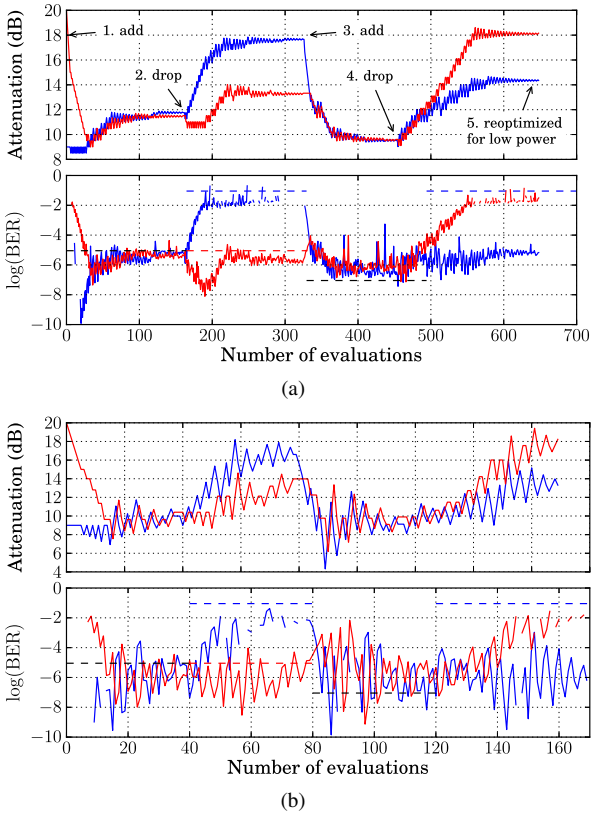

Fig. 14. Evolution of the attenuation and the BER of two lightpaths in our testbed for the lightpath lifecycle scenario. Two sets of (θ^+, θ^-) parameters are considered, resulting in (a) slower but more gradual convergence, or (b) high fluctuations but fast convergence.

SiMPLE readjusts the power levels, effectively dropping the blue wavelength. Next, a second $\overline{\text{BER}}$ constraint is modified in step 3, and finally the red lightpath is dropped in step 4. This experiment is designed to capture the full scale of actions that future dynamic optical networks are expected to perform, based on higher-layer requests. Note that the attenuation level at step 5 in Fig.14(a) is larger compared to that just before step 2, even though BER levels at these steps are the same. This illustrates that the same QoT constraint can be met using less optical power.

Fig. 14(a) shows the variations of the power level and the corresponding BER when $\theta^- = 0.9$ and $\theta^+ = 1$. It can be seen that the variations in power and BER are low, and the convergence is smooth. This smoothness comes at a penalty in time, since the entire scenario takes about 650 OPM evaluations. This may be long with current OPMs, if evaluations are in the order of minutes. However, as future OPMs are expected to perform evaluations in the order of milliseconds, the *SiMPLE* algorithm can complete this scenario in under a second, with reasonable convergence behavior.

Fig 14(b) shows the same scenario when $\theta^- = 0.6$ and $\theta^+ = 1.2$. Similar to the insight obtained from simulations, these parameters cause large variations in the step size, leading to large fluctuations in the power levels and the BER. However, these parameters also allow the entire test sequence to complete in about 170 OPM evaluations, substantially faster than in the scenario illustrated in Fig 14(a).

Similar to the simulations, the algorithm was evaluated

TABLE I
AVERAGE VALUES OVER 5 EXPERIMENTS OF THE RUNNING STANDARD DEVIATION (RStd) AND FEASIBLE TIME (FeasTime) METRICS FOR SEVERAL PARAMETER COMBINATIONS WHEN USING HEURISTIC H3.

(θ^-, θ^+)	RStd $(\times 10^{-3})$	FeasTime
$(0.9, 1)$	1.80	76
$(0.9, 1.1)$	9.16	30
$(0.6, 1.2)$	5.83	28
$(0.6, 1.3)$	66.2	31

on the experimental testbed for different set of parameters (θ^-, θ^+). Due to the longer duration of the experiments, the results were averaged over 5 runs, for each heuristic and a representative set of parameters. The averaged running standard deviation (RStd) and feasible time (FeasTime) for heuristic H3 are shown in Table I. The results are comparable to the simulation results obtained in Fig. 10. When the parameters are picked such that the step size does not vary greatly with $\theta^-, \theta^+ = (0.9, 1.1)$, there are less fluctuations, but the convergence time is larger. At the other extreme, when $\theta^-, \theta^+ = (0.6, 1.3)$, the power levels fluctuate a lot, but the convergence is faster. Unlike the simulation, there is an optimal set of parameters for these experiments. When $\theta^-, \theta^+ = (0.6, 1.2)$, the best possible convergence is obtained. It should be noted that noise is present naturally in the experiments and can be the cause for this somewhat different behavior.

To conclude, we tested the *SiMPLE* algorithm running on a computer that controlled optical devices. We showed that the insights obtained from simulations also hold with real equipment, and that proper selection of parameters for the *SiMPLE* algorithm can enable complex operations in future dynamic optical networks.

VIII. CONCLUSIONS

In this paper, we formulated a global optimization problem, *MLO*, that captures the QoT constraints in optical networks. This problem is unique since analytical models that capture impairments in optical fibers do not exist, and measurements are the only practical way to characterize performance. We designed a power control algorithm, *SiMPLE*, for solving this problem by using feedback from OPMs. The convergence of *SiMPLE* to the optimal solution was demonstrated using extensive simulations on a network-scale optical network simulator, as well as experiments with commercial optical network equipment. We evaluated the performance of *SiMPLE* and showed that even relatively simple dynamic policies can result in *substantial power savings through a better use of resources*.

The *SiMPLE* algorithm enables dynamic control of optical networks in near real-time, compared to the days-long setup times for lightpaths in the existing optical networks. Hence, *SiMPLE* is the first step towards allowing optical networks to react rapidly to user demands and traffic variations. Namely, it is a step towards software-defined optical network.

In future work, we will expand the model and consider other variables that can be controlled in dynamic optical networks. For example, through the use of adaptive modulation techniques, the OSNR of a lightpath can be improved without modifying its power level. Similarly, spectrum utilization can be improved by decreasing the guard bands between lightpaths. An extended version of the *SiMPLE* algorithm can be used to monitor the QoT in such transmission systems, and adapt the bandwidth and modulation schemes as necessary. Finally, we also intend to evaluate the performance of the *SiMPLE* algorithm by incorporating in the simulation a different constant-gain optical amplifier model, and by experimenting in a large-scale testbed.

APPENDIX A
CONVERGENCE OF *SiMPLE*

Being a general convex solution algorithm, *SiMPLE* can be used to solve a wide variety of problems. First the case of deterministic functions is discussed, followed by functions that have probabilistic noise.

The convergence of *SiMPLE* when the $\mathrm{BER}_i^u(\mathbf{p})$ and $\mathrm{OSNR}_i^u(\mathbf{p})$ functions are not affected by random noise follows from several well-known results. We outline these results for completeness. The proof of convergence for the family of direct search algorithms is given in [27]. In the simplest case, when the *MLO* problem does not have constraints (C1)-(C2) but only has an objective function, the proof follows from bounding the step size parameter α_k from above. The algorithm converges since the step size can be shown to get arbitrarily close to zero. The proof for the case in which the constraints (C1)-(C2) are present follows from similar arguments, as the augmentation procedures AUGMENTLOG or AUGMENTQUAD are used to transform the constrained problem to an unconstrained one.

To the best of our knowledge, it was not proved that direct search methods converge if measurement noise is present. However, our application of these methods differs from other typical scenarios in that *SiMPLE* runs the algorithm continuously. In our setting, if there is a positive probability that the direct search method converges, it can be shown that *SiMPLE* will also converge. While in practice this convergence may take a long time, it is not as crucial because the priority for *MLO* is to reach feasibility, and not necessarily optimality.

The use of two different augmenting functions in a single setting is unique to the formulation of the *MLO* problem. The fact that the algorithm converges when the log-barrier function AUGMENTLOG is used is well-known [31]. AUGMENTQUAD generates a function that is similar to the well-known quadratic penalty function, which converges when the parameter $\mu \to \infty$. The proof that *SiMPLE* yields a feasible point (not necessarily optimal) follows from the following argument. When a typical quadratic penalty function is used, the algorithm converges to an optimal solution when $\mu \to \infty$. It therefore follows that for μ large enough, an intermediate value of the augmenting quadratic penalty function is feasible. Using the AUGMENTQUAD function effectively corresponds to finding this intermediate value, thereby yielding feasibility.

ACKNOWLEDGEMENTS

This work was supported in part by CIAN NSF ERC under grant EEC-0812072, NSF grant CNS-1018379, and DTRA grant HDTRA1-13-1-0021. This material is based upon work supported by (while

one of the authors was serving at) the National Science Foundation. Any opinion, findings and conclusions or recommendations expressed here do not necessarily reflect the views of the National Science Foundation.

REFERENCES

[1] "Cisco CRS-1 datasheet," http://www.cisco.com/en/US/products/ps5763/.
[2] "Géant project," http://www.geant.net.
[3] "Greentouch consortium," http://www.greentouch.org.
[4] "Cisco visual networking index: Forecast and methodology, 2011-2016," *White Paper*, 2012.
[5] *Proc. IEEE, Special Issue on The Evolution of Optical Networking*, vol. 100, no. 5, May 2012.
[6] S. Azodolmolky, M. Klinkowski, E. Marin, D. Careglio, J. S. Pareta, and I. Tomkos, "A survey on physical layer impairments aware routing and wavelength assignment algorithms in optical networks," *Comput. Netw.*, vol. 53, no. 7, pp. 926–944, 2009.
[7] B. Birand, H. Wang, K. Bergman, D. Kilper, T. Nandagopal, and G. Zussman, "Real-time power control for dynamic optical networks – algorithms and experimentation," in *Proc. IEEE ICNP'13*, Oct. 2013.
[8] B. Birand, H. Wang, K. Bergman, and G. Zussman, "Measurements-based power control – a cross-layered framework," in *Proc. OFC'13*, Mar. 2013.
[9] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
[10] C. Chekuri, P. Claisse, R.-J. Essiambre, S. Fortune, D. C. Kilper, W. Lee, N. K. Nithi, I. Saniee, B. Shepherd, C. A. White, G. Wilfong, and L. Zhang, "Design tools for transparent optical networks," *Bell Labs Tech. J.*, vol. 11, no. 2, pp. 129–243, 2006.
[11] M. R. Chitgarha, S. Khaleghi, W. Daab, M. Ziyadi, A. Mohajerin-Ariaei, D. Rogawski, M. Tur, V. Vusirikala, W. Zhao, J. Touch, and A. E. Willner, "Demonstration of WDM OSNR performance monitoring and operating guidelines for polmuxed 200-Gbit/s 16-QAM and 100-Gbit/s QPSK data channels," in *Proc. OFC'13*, Mar. 2013.
[12] A. Conn, K. Scheinberg, and L. Vicente, *Introduction to derivative-free optimization*. MPS/SIAM Series on Optimization, 2009, vol. 8.
[13] T. Deng, S. Subramaniam, and J. Xu, "Crosstalk-aware wavelength assignment in dynamic wavelength-routed optical networks," in *Proc. IEEE BroadNets'04*, Oct. 2004.
[14] R. Doverspike and J. Yates, "Optical network management and control," *Proc. IEEE*, vol. 100, no. 5, pp. 1092–1104, May 2012.
[15] R.-J. Essiambre, G. Kramer, P. Winzer, G. Foschini, and B. Goebel, "Capacity limits of optical fiber networks," *J. Lightw. Technol.*, vol. 28, no. 4, pp. 662–701, Feb. 2010.
[16] D. J. Geisler, R. Proietti, Y. Yin, R. P. Scott, X. Cai, N. K. Fontaine, L. Paraschis, O. Gerstel, and S. J. B. Yoo, "Experimental demonstration of flexible bandwidth networking with real-time impairment awareness," *Opt. Express*, no. 26, pp. B736–B745, Dec. 2011.
[17] O. Gerstel, M. Jinno, A. Lord, and S. Yoo, "Elastic optical networking: a new dawn for the optical layer?" *IEEE Commun. Mag.*, vol. 50, no. 2, pp. s12–s20, Feb. 2012.
[18] A. Gnauck, "Advanced amplitude- and phase coded formats for 40-gb/s fiber transmission," in *Proc. IEEE LEOS'04*, Nov. 2004.
[19] K. Guan, D. Kilper, and G. Atkinson, "Evaluating the energy benefit of dynamic optical bypass for content delivery," in *Proc. IEEE INFOCOM'11 Workshop*, Apr. 2011.
[20] P.-H. Ho and H. Mouftah, "Routing and wavelength assignment with multigranularity traffic in optical networks," *J. Lightw. Technol.*, vol. 20, no. 8, pp. 1292–1303, Aug. 2002.
[21] C.-Y. Hong, S. Kandula, R. Mahajan, M. Zhang, V. Gill, M. Nanduri, and R. Wattenhofer, "Achieving high utilization with software-driven WAN," in *Proc. ACM SIGCOMM'13*, Aug. 2013.
[22] F. Idzikowski, S. Orlowski, C. Raack, H. Woesner, and A. Wolisz, "Saving energy in IP-over-WDM networks by switching off line cards in low-demand scenarios," in *Proc. IEEE ONDM'10*, Feb. 2010.
[23] M. S. Islam and S. P. Majumder, "Bit error rate and cross talk performance in optical cross connect with wavelength converter," *J. Opt. Netw.*, vol. 6, no. 3, pp. 295–303, Mar. 2007.
[24] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, J. Zolla, U. Hölzle, S. Stuart, and A. Vahdat, "B4: experience with a globally-deployed software defined WAN," in *Proc. ACM SIGCOMM'13*, Aug. 2013.
[25] D. Kilper, K. Guan, K. Hinton, and R. Ayre, "Energy challenges in current and future optical transmission networks," *Proc. IEEE*, vol. 100, no. 5, pp. 1168–1187, May 2012.
[26] D. Kilper, R. Bach, D. Blumenthal, D. Einstein, T. Landolsi, L. Ostar, M. Preiss, and A. Willner, "Optical performance monitoring," *J. Lightw. Technol.*, vol. 22, no. 1, pp. 294–304, Jan. 2004.
[27] T. Kolda, R. Lewis, and V. Torczon, "Optimization by direct search: New perspectives on some classical and modern methods," *SIAM review*, vol. 45, no. 3, pp. 385–482, 2003.
[28] C. Lai, A. Fard, B. Buckley, B. Jalali, and K. Bergman, "Cross-layer signal monitoring in an optical packet-switching test-bed via real-time burst sampling," in *Proc. IPC'10*, Nov. 2010.
[29] B. Mukherjee, *Optical WDM networks*. Springer-Verlag New York Inc, 2006.
[30] T. Nandagopal, K. C. Guan, and D. Kilper, "Energy efficiency and delay performance of data transfer using dynamic optical switching," in *Proc. ICC'12*, June 2012.
[31] J. Nocedal and S. Wright, *Numerical optimization*. Springer-Verlag, 1999.
[32] Y. Pan, T. Alpcan, and L. Pavel, "A system performance approach to OSNR optimization in optical networks," *IEEE Trans. Commun.*, vol. 58, no. 4, pp. 1193–1200, Apr. 2010.
[33] Z. Pan, C. Yu, and A. Willner, "Optical performance monitoring for the next generation optical communication networks," *Opt. Fiber Technol.*, vol. 16, no. 1, pp. 20–45, 2010.
[34] R. Ramaswami, K. Sivarajan, and G. Sasaki, *Optical networks: a practical perspective*. Morgan Kaufmann, 2009.
[35] S. Ricciardi, F. Palmieri, U. Fiore, D. Careglio, G. Santos-Boada, and J. Solé-Pareta, "An energy-aware dynamic RWA framework for next-generation wavelength-routed networks," *Comput. Netw.*, vol. 56, no. 10, pp. 2420–2442, 2012.
[36] J. Sole-Pareta, S. Subramaniam, D. Careglio, and S. Spadaro, "Cross-layer approaches for planning and operating impairment-aware optical networks," *Proc. IEEE*, vol. 100, no. 5, pp. 1118–1129, May 2012.
[37] S. Uhlig, B. Quoitin, J. Lepropre, and S. Balon, "Providing public intradomain traffic matrices to the research community," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 36, no. 1, pp. 83–86, 2006.
[38] S. Woodward and M. Feuer, "Toward more dynamic optical networking," in *Proc. OECC'10*, July 2010.
[39] Y. Wu, L. Chiaraviglio, M. Mellia, and F. Neri, "Power-aware routing and wavelength assignment in optical networks," in *Proc. ECOC'09*, Sept. 2009.

Berk Birand (S'04) received his dual B.S. degree in Electrical & Computer Engineering and Computer Science with distinction from the Worcester Polytechnic Institute in 2008. He received his M.S. degree in Electrical Engineering from Columbia University in 2009, and is currently a Ph.D. candidate in the Electrical Engineering Department at Columbia University.

His research interests are in the area of wireless scheduling algorithms and energy-efficient optical networks. He was an IBM Ph.D. Fellow for the 2011-2012 academic year and was a research intern at IBM T.J. Watson Research Center in 2010 and 2011. He was awarded the "Best Theory Session Talk Award" in the ACM MobiHoc'10 S^3 workshop, and the "Millman Award for Outstanding Teaching Assistant" by Columbia University. He was a co-chair of the ACM S^3'11 workshop co-located with ACM MobiCom'11.

Howard Wang (S'06) received his B.S. in Electrical Engineering (minoring in computer science and economics) from Columbia University in 2006 and the M.S. and Ph.D. in Electrical Engineering in 2008 and 2013, respectively. As a graduate researcher in the Lightwave Research Laboratory at Columbia University, his doctoral research focused on optical networks and photonic switch architectures for high performance computing systems and data centers.

Keren Bergman is the Charles Batchelor Professor and Chair of Electrical Engineering at Columbia University where she also directs the Lightwave Research Laboratory (http://lightwave.ee.columbia.edu/). She leads multiple research programs on optical interconnection networks for advanced computing systems, data centers, optical packet switched routers, and chip scale nanophotonic networks. Dr. Bergman holds a Ph.D. from M.I.T. and is a Fellow of the IEEE and of the OSA.

Dan Kilper is a research professor at the College of Optical Sciences, University of Arizona and the administrative director of the Center for Integrated Access Networks (CIAN), an NSF engineering research center. Previously he was a member of technical staff at Bell Labs, Alcatel-Lucent. He served as the founding technical committee chair of the GreenTouch Consortium, a global organization with over 50 members, and was the Bell Labs Liaison Executive for the Center for Energy Efficient Telecommunications at the University of Melbourne, Australia. While at Bell Labs, he received the Presidents Gold Medal Award in 2004 and was a member of the Presidents Advisory Council on Research. He is an adjunct professor at Columbia University and is a senior member of IEEE. Currently he is serving as the general co-chair of IEEE GreenComm 2014 and as the TPC co-chair for Photonics in Switching 2014. Dr. Kilper has conducted research on optical performance monitoring and on transmission, architectures and control systems for transparent and energy-efficient optical networks. He holds eight patents and authored four book chapters and more than one hundred peer-reviewed publications.

Thyaga Nandagopal serves in the Directorate of Computer & Information Science and Engineering (CISE) of the National Science Foundation. He manages wireless networking and mobile computing research within the Networking Technologies and Systems (NeTS) program at NSF. He has been with the Foundation since February 2012.

Dr. Nandagopal received his Ph.D. in Electrical Engineering in 2002 from the University of Illinois at Urbana-Champaign. He was at Bell Labs from 2002 to 2012. His research interests have spanned several areas over these years: wireless ad hoc/mesh networks, RFID/sensor networks, internet routing architectures and protocols, cloud computing, and energy-efficient networks.

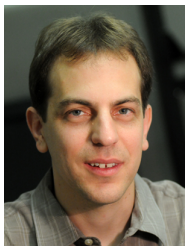

Gil Zussman (S'02-M'05-SM'07) received the Ph.D. degree in electrical engineering from the Technion in 2004 and was a postdoctoral associate at MIT in 2004–2007. In 2008 he joined the faculty of the Department of Electrical Engineering at Columbia University where he is currently an Associate Professor. His research interests are in the area of networking, and in particular in the areas of wireless, mobile, and resilient networks. He has been an editor of IEEE Transactions on Wireless Communications and Ad Hoc Networks, the Technical Program Committee (TPC) co-chair of IFIP Performance 2011, and a member of a number of TPCs (including the INFOCOM, MobiCom, SIGMETRICS, and MobiHoc committees). He is a co-recipient of 5 paper awards including the ACM SIGMETRICS'06 Best Paper Award and the 2011 IEEE Communications Society Award for Outstanding Paper on New Communication Topics. He received the Fulbright Fellowship, two Marie Curie Fellowships, the DTRA Young Investigator Award, and the NSF CAREER Award, and was a member of a team that won first place in the 2009 Vodafone Foundation Wireless Innovation Project competition.