

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

Berk Birand*, Howard Wang*, Keren Bergman*, Dan Kilper†, Thyaga Nandagopal‡, Gil Zussman*

*{berk, howard, keren, gil}@ee.columbia.edu,

Department of Electrical Engineering, Columbia University, New York, NY

†dkilper@optics.arizona.edu,

Center for Integrated Access Networks, College of Optical Sciences, University of Arizona, Tucson, AZ

‡tnandago@nsf.gov, National Science Foundation, Arlington, VA

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.

I. INTRODUCTION

Optical networks are the underlying infrastructure of core and aggregation networks [18]. In order to handle peaks in traffic demand, these networks are usually *static* and *over-provisioned* [25], which 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 for energy efficient operation [13] already pose challenges that cannot be addressed by over-provisioning.

Wavelength-Switched Optical Networks (WSOs) (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 [18], modulators that can adapt to the link state [10], and bandwidth variable transceivers that can modify band gaps between adjacent

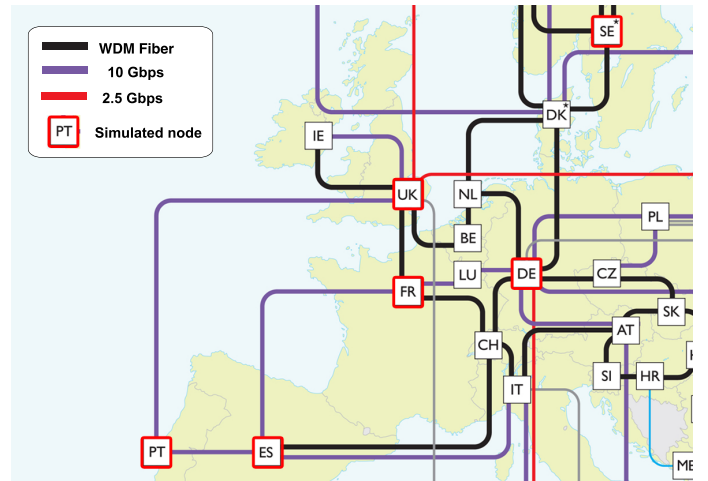


Fig. 1. The optical infrastructure of the Géant academic network [2], whose topology is used in order to evaluate the proposed algorithms. The highlighted nodes are used as part of the topology simulated in Section VI.

channels [11]. While the flexibility provided by such devices allows the network to adapt to the link conditions and traffic demands, the static nature of optical networks is mostly due to potential impairments that are hard to predict or model [5], [25]. Sources of these impairments are related to the optical transmission and fiber properties [18], and to factors such as temperature, component drift, component aging, and maintenance work [21]. Due to these impairments, lightpaths are rarely modified once assigned. This means that Routing and Wavelength Allocation (RWA) (e.g., [5]) is done primarily at the planning phase, with significant over-provisioning. Any changes are executed manually which is both time-consuming and expensive [9].

Hence, our goal is to enable lightpath configuration, setup, and teardown with convergence times in the order of tens of seconds. This will allow the network to efficiently react to traffic variations and customer demands. We build on the capabilities of the dynamic optical devices as well as various Optical Performance Monitors (OPMs) that have been recently developed [21]. 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

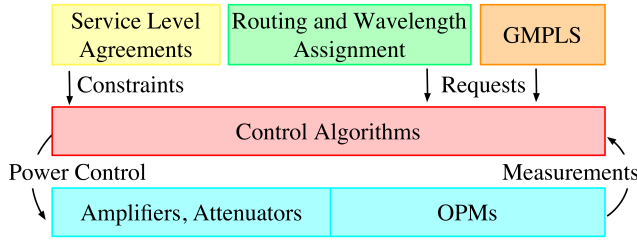


Fig. 2. Schematic view of the interaction of the control algorithm with the optical devices, and the higher layer algorithms and SLAs.

control schemes that use them operate at the link-scale rather than at the network-scale [14]. Extensions of per-link policies to the entire network do not produce globally optimal results, and may not converge within the desired time [6]. Moreover, although protocols based on Generalized Multi-Protocol Label Switching (GMPLS) can leverage OPM measurements [1], [15], [16], they mostly provide the infrastructure for using the measurements but still lack the protocol and network optimization aspects (for more details, see Section II).

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 and GMPLS protocols.

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 [23], [26]. *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 a function 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 [8] and computes a live configuration of the wavelengths' power levels throughout the optical network.

In summary, the main contributions of this paper are two-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, and (ii) we evaluate the performance of the algorithm using a realistic optical simulator and in an optical testbed. 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 a software-defined optical network where a separate control plane running the SiMPLE algorithm determines the optical data plane behavior*.

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. 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 [13]. A recent special issue of the *Proceedings to the IEEE* [4] addresses several important problems in all-optical networks, such as reconfigurability, optical flow switching, optical network control, and cross-layer impairment-aware optical networks [23]. 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 we apply an impairment-aware RWA algorithms [5] or make local decisions in an optical switching fabric [14]. 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 [5], and on the reduction of overall energy consumption [26]. 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.

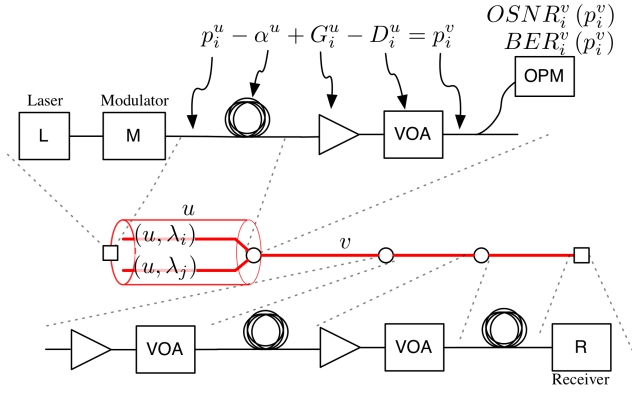


Fig. 3. The correspondence between a physical link and its spans, and our optical link model. The leftmost node includes an optical source composed of a laser and a modulator. The intermediate nodes consist of amplifiers and Variable Optical Attenuators (VOAs). The rightmost span ends the link with a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

GMPLS [16] is used widely to enable control of WSONs. Recent standardization efforts [15], [17] leverage RSVP-TE extensions to allow the network operator to collect (possibly imperfect) impairment parameters along a path and use the collected data in the computation of an impairment-aware RWA solution. However, even if such a solution can be computed, methods to switch the network from the old RWA to a new RWA solution in a stable, real-time fashion are still unknown. The *SiMPLE* algorithm in this paper offers a way to successfully solve this problem.

OPMs offer real-time inspection of transmissions [21] 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 [14] and changing the modulation format in real-time according to link OSNR [10]. 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 [20]. 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 (OXC)s 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 referring to individual wavelengths of each span.

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

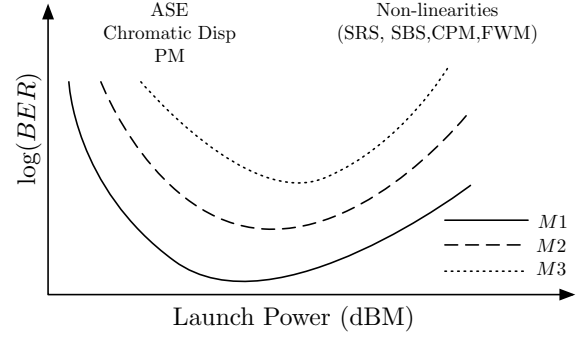


Fig. 4. Illustration of the relationship of BER and the optical launch power for a few modulation formats denoted $M1$, $M2$, $M3$ [12]. Prominent impairments in each region are marked (see [18] for detailed descriptions).

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 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

$$p_j^v = p_i^u - \alpha^u + G_i^u - D_i^u. \quad (1)$$

Depending on the network, different values of this expression will be the control variables. If the amplification cannot be modified, G_i^u will be a constant. 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 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 looked up in [18]. For instance, at low powers, increasing the power levels improves the BER by mitigating the effects of Amplified Spontaneous Emission (ASE) noise. However, at higher power levels, increasing the powers 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 [21] affect the specifics of the curve, but the overall nature of the relationship remains the same [12].

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 $\text{BER}_i^u(\mathbf{p})$ and $\text{OSNR}_i^u(\mathbf{p})$, 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, $\text{BER}_i^u(\mathbf{p})$ is convex, while $\text{OSNR}_i^u(\mathbf{p})$ is concave. We used our experimental testbed described in Section VII to gradually attenuate two lightpaths λ_1 and λ_2 traversing a single fiber, and verified the convexity of the curves. We also numerically verified the convexity by computing the Hessian of this curve at all points. We leverage this fact in the next section to develop a network-wide power control algorithm.

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.² 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 [18].

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

²An optical link between two regional offices that includes several amplifiers is modeled as a path of several nodes.

Lightpaths are represented as sequences of λ -spans $P = \{(u, \lambda_i), (v, \lambda_i), (w, \lambda_i), \dots\}$.

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 [22]. The assignment of routes and wavelengths to links is out of scope, as these are assumed to be handled by an RWA algorithm [5].

All λ -spans may not 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.

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), 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, as measured 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. The formulation for this optimization problem is as follows.

Problem 1 (Multi-Link Optimization - MLO):

$$\begin{aligned} \underset{\mathbf{p}, D}{\text{minimize}} \quad & h_{\text{MLO}}(\mathbf{p}, D) = \sum_{(u, \lambda_i)} (p_i^u - D_i^u) \\ \text{subject to} \quad & \text{BER}_i^u(\mathbf{p}) \leq \overline{\text{BER}}_i^u, \quad \forall (u, \lambda_i) \in \Lambda_{\text{BER}} \quad (1) \\ & \text{OSNR}_i^u(\mathbf{p}) \geq \overline{\text{OSNR}}_i^u, \quad \forall (u, \lambda_i) \in \Lambda_{\text{OSNR}} \quad (2) \\ & 0 \leq \mathbf{p} \leq \text{SAF}, \quad (3) \end{aligned}$$

where $\overline{\text{BER}}_i^u$ and $\overline{\text{OSNR}}_i^u$ 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 the power can be minimized either by decreasing the power levels directly, or by increasing the attenuation, D .

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 $\overline{\text{BER}}_i^u$ 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 [8] 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 $\text{BER}(\mathbf{p})$ and $\text{OSNR}(\mathbf{p})$ 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 may introduce noise that needs to be accounted for during the computation of the optimal solution.

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 $\text{BER}(\mathbf{p}(k))$ and $\text{OSNR}(\mathbf{p}(k))$.

A. Design Considerations (DCs)

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

(DC1) To evaluate the BER and/or OSNR functions at a given power level, the attenuations or gains of the amplifiers must be modified throughout the network. This restricts the type of points that can be evaluated, since the process should cause as little disruption to active lightpaths as possible, and keep the changes in the network to a minimum.

(DC2) For the network operator, it is more important to adhere to the SLA requirements than to find the setup that consumes the least amount of power. Since constraint satisfaction 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 use the first or second derivatives of the functions used in the optimization to choose the next iteration [19]. However, these methods are not appropriate for cases where the functions to be optimized are not known and can be noisy. Therefore, derivative-free optimization (DFO) algorithms [8] 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 [8], and incorporates the design considerations discussed in Section V-A. Starting from a point $\mathbf{p}(0)$, the algorithm evaluates points along a set of computed directions. For a search on the plane, this set could be as simple as

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:  $\theta^+$ ,  $\theta^-$ , and  $\alpha_{\text{tol}}$ 
3: loop
4:    $\alpha_k \leftarrow 1$ 
5:   repeat
6:     if  $\mathbf{p}(k)$  is feasible then
7:        $f(\mathbf{p}(k)) \leftarrow \text{AUGMENTLOG}(\mathcal{I})$ 
8:     else  $f(\mathbf{p}(k)) \leftarrow \text{AUGMENTQUAD}(\mathcal{I})$ 
9:     end if
10:     $\mathcal{H}_k \leftarrow \text{GENERATE}(\mathbf{p}(k))$ ;  $\mathcal{D}_k \leftarrow \mathcal{H}_k \cup \mathcal{G}$ 
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  $\alpha_k \leq \alpha_{\text{tol}}$ 
17: end loop

```

$\mathcal{D}_k = \left\{ \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix} \right\}, \forall k$. 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 (θ^-, θ^+) . These parameters have a large effect on the convergence properties, as shown in Section VI-D.

The pseudocode for the *SiMPLE* algorithm is shown above. It takes as input an instance \mathcal{I} of the *MLO* problem and a initial power assignment $\mathbf{p}(0)$. The problem instance \mathcal{I} corresponds to a set of constraints to the problem 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 textscAugmentLog , which returns the following log-barrier function:

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

where μ is parameter of the augmentation function [19]. Since the performance functions $\text{BER}(\cdot)$ and $\text{OSNR}(\cdot)$ 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:

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

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 quadratic augmenting functions (e.g., [19]), the objective function of minimizing the total power consumption is captured in this function, as the priority is to reach feasibility.

The direct search step in line 11 tries several directions d_k from a search set \mathcal{D}_k to improve the objective value. Each of these directions is tried with step size α_k . For each d_k of this set, it changes the power levels of this network, 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 $\mathcal{G} = [I; -I]$ 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 strategy, 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 search set of the algorithm, and we consider *three heuristics*:

- H1: $\mathcal{H} = \emptyset$, for comparison to the other methods,
- H2: $\mathcal{H} = d_{k-1}$, the last successful search direction,
- H3: $\mathcal{H} = d_{k-1}$ 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 search set that 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 $\mathbf{p}(k)$.

VI. SIMULATION EVALUATION

In this section, we evaluate the performance of the SiMPLE algorithm and demonstrate the benefits of dynamic optical network 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 as quickly as possible. We let $\mathcal{P} = \{\mathbf{p}(k)\}$ denote the set of power vectors over iteration numbers k .

The running standard deviation (RStd) measures the variability of the power levels. This value 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:

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

where Std is the standard deviation operator.

The FeasTime metric measures the time until all the constraints are satisfied and the problem is feasible:

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

Finally, the feasibility probability FeasProb is defined as the probability that SiMPLE finds a feasible solution to the given problem.

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 problem, the measurement and actuation overheads of each OPM dominate the running time compared to the operations of the algorithm.

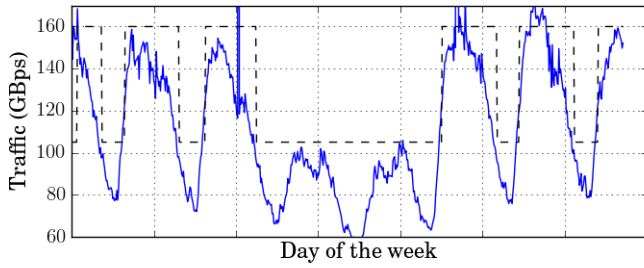


Fig. 5. Sample traffic pattern between London and Stockholm in the Géant network over a period of one week. The dashed line corresponds to the required capacity to satisfy the demand. In peak times, additional lightpaths are needed to support a higher capacity.

B. Simulation Setup

We developed the simulator by using a detailed physical model of an optical amplifier developed at Bell Labs [7]. 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, as well as 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.

We use the Géant network topology (Fig. 1) consisting of 25 lightpaths that follow four routes. The endpoints for the lightpaths used in the simulations are highlighted in red in Fig. 1 (e.g., PT to SE). These 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 received power levels are measured at each destination. Gaussian noise of different variances was added to the OPM evaluations to mimic measurement noise.

C. Traffic Data

Traffic data between each pair of cities in the Géant network obtained at 15 minute intervals for a four month period in 2005 is available in [24]. This data was averaged over a week-long period to get traffic variations for each weekday, for all the simulated nodes in Fig. 1. A sample of the traffic variations over a week is shown in Fig. 5. The data was scaled tenfold to accommodate the traffic growth statistics based on [3]. Furthermore, since the values provided in [24] correspond to the average data rates, we provisioned five times as much capacity in order to account for bursts in traffic.

Many approaches can be used to generate the optimal topologies that satisfy these traffic demands. In the most complex case, a new topology can be computed in real-time using the live traffic matrices. Since we do not focus on such algorithms, two topologies are designed to satisfy the traffic demands. While this approach seems simple, it already provides a vast improvement over current optical networks, in

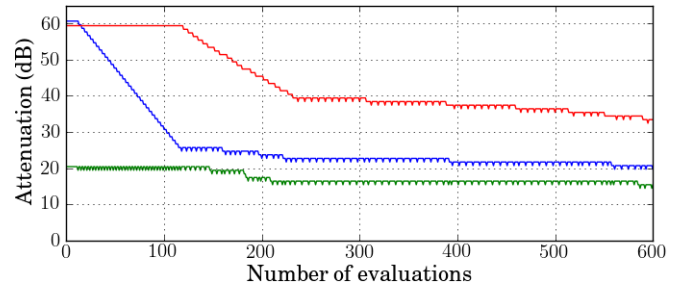


Fig. 6. Evolution of the attenuation of three lightpaths in our simulator while the *SiMPLE* algorithm transitions from the low-capacity topology to the high-capacity one in order to satisfy the extra demand. Two lightpaths are progressively added by decreasing their attenuation.

which such drastic changes rarely occur over timescales less than the order of months.

In the two considered topologies, the high-capacity one is used during the day, while the low-capacity one is used when the demand is low on nights and weekends. The capacities for these two topologies are illustrated on Fig. 5 as the envelope that covers the traffic demands. The objective of the *SiMPLE* algorithm is to switch between such two topologies to optimize resource usage, as desired in dynamic optical networks.

D. Simulation Results

Parameters and Heuristics: For a given network deployment, there are two types of parameters that should be considered; the (θ^-, θ^+) parameters for adjusting the step size (line 2 of 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 topology used during nights and weekends, to the high capacity topology 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. 6 shows the attenuation evolution over time for a single run of simulation. In this setup, heuristic H1 was simulated 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 trade-off 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. 7.

Figs. 7(a)-7(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 devi-

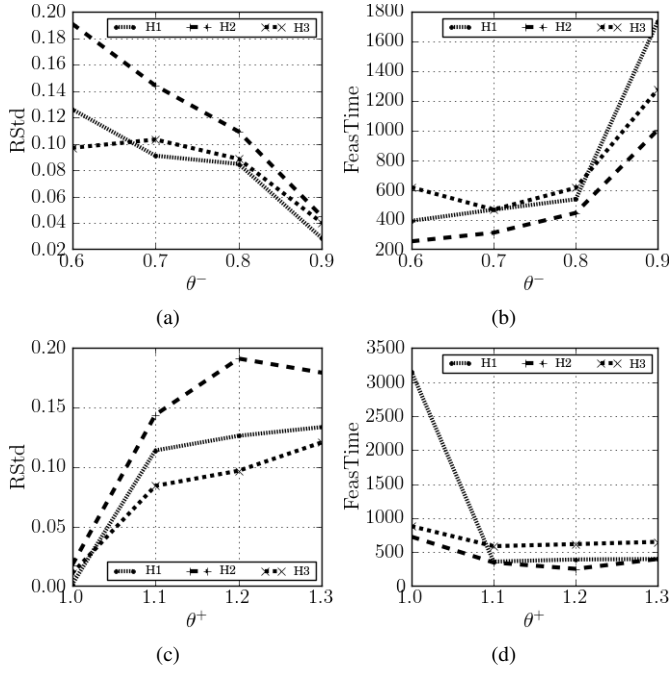


Fig. 7. Values of the RStd and FeasTime metrics that demonstrate the fluctuations and convergence time of the *SiMPLE* algorithm when switching from the low capacity to the high capacity topology. Two sets of parameters are plotted: (a)-(b) θ^+ is fixed and θ^- is modified, and (c)-(d) θ^- is fixed and θ^+ is modified.

ation 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. 7(a)-7(b). The trade-off 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. 7(c)-7(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. 7(d) for H1. If the

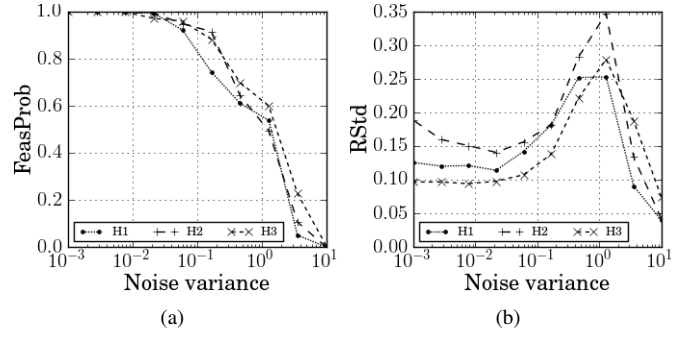


Fig. 8. The performance of *SiMPLE* as a function of the measurement noise variance, for two metrics: (a) the probability that the algorithm will yield a feasible solution (FeasProb), and (b) the running standard deviation (RStd) of the run.

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. 8.

The probability that *SiMPLE* will reach a feasible point is plotted in Fig. 8(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. 8(b) shows the variations in the power levels with respect to the noise variance. We can observe 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 consumption. To illustrate the benefits of dynamic optical networks, we perform a simple computation of the energy consumption of the low-capacity and high-capacity topologies. Our focus is on electrical components since they are more power-hungry than the optical components. We assume a network based on the Cisco CRS-1 router. 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 kW each. For certain nodes, a multi-shelf system consisting of two 16-shelf were necessary to satisfy the high traffic demands.

Based on our estimates, the high capacity topology consumes 47 kW when all the linecards and chassis are active. For a static network, this would be the permanent power

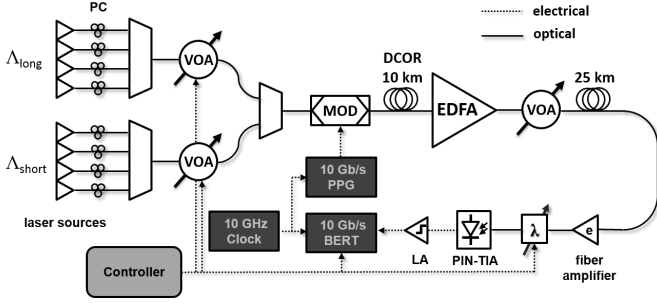


Fig. 9. Schematic of the experimental setup. Two wavebands each consisting of four wavelengths are modulated with a PRBS pattern and are transmitted via an amplified link segment. The optical powers of each waveband are controlled by a computer configured as a controller that runs the *SiMPLE* algorithm.

consumption of this network. The low capacity topology consumes 32 kW of power. Based on the traffic patterns described in Section VI-C, we estimate that the high capacity topology was needed 41% of the time. The average power consumption of a dynamic optical network in this instance would be 38 kW. By using a dynamic optical network enabled by the *SiMPLE* algorithm, and adapting the link capacities to the traffic demands, it is therefore possible to reduce the power consumption of this network by 20%.

Thus, our simulations show that the *SiMPLE* algorithm can adjust power levels to enable the dynamic operation of optical equipment. Using appropriate parameters and heuristics, it is possible to prioritize convergence speed or power fluctuations, while mitigating the effects of measurement noise.

VII. TESTBED EVALUATION

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

A. Experimental Setup

Our experimental setup is shown in Fig. 9. Eight closely-spaced Continuous-Wave (CW) laser sources in the ITU C-band are partitioned into two contiguous wavebands— Λ_{long} 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) preceding 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

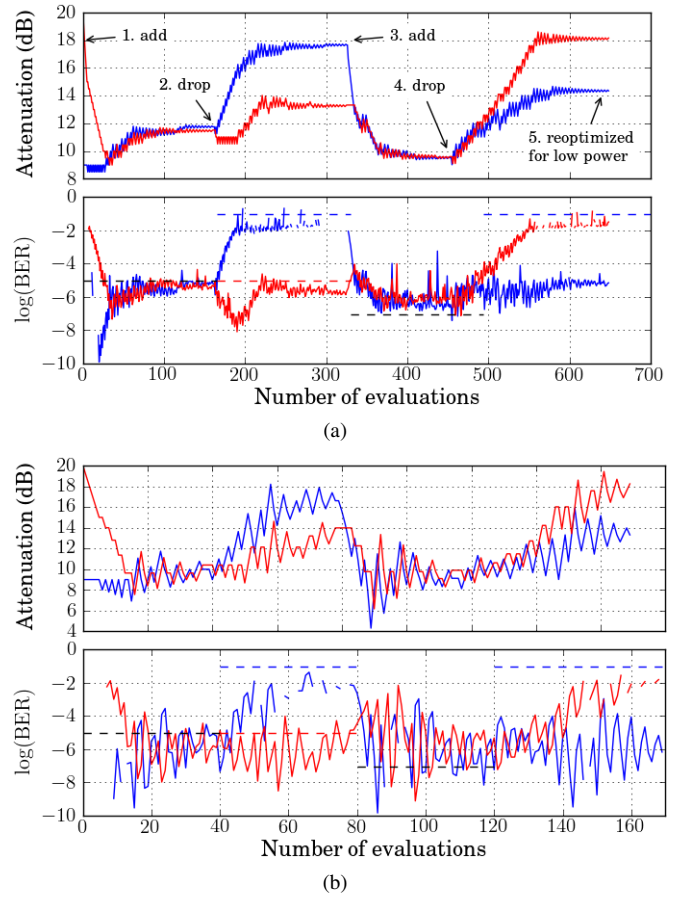


Fig. 10. Evolution of the attenuation and the BER of two lightpaths in our testbed for the lightpath lifecycle scenario. Two sets of (θ^+, θ^-) parameters are considered, resulting in (a) slower but more gradual convergence, or (b) high fluctuations but fast convergence.

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 a large-scale experiment in our testbed. This experiment has several parts, as labeled in Fig. 10(a). First, the red lightpath is added similarly to the simulation scenarios. Then, the threshold constraint $\overline{\text{BER}}$ for one of the wavelengths is modified to be very large (10^{-1}), and *SiMPLE* readjusts the power levels, effectively dropping the blue wavelength. Next, a second $\overline{\text{BER}}$ constraint is modified, and finally the red lightpath is dropped. This experiment is designed to capture the full scale of actions that future dynamic optical networks are expected to perform. Note that the attenuation level at step

5 on Fig.10(a) is larger compared to that just before step 2, even though their BER levels are the same. This shows that the same QoT constraint can be met using less optical power.

We ran this experiment over different values of θ^- and θ^+ . We show two sample outcomes in Fig. 10. Fig 10(b) shows the variations of the power level and the corresponding BER 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.

Fig 10(a) corresponds to the same scenario with $\theta^- = 0.9$ and $\theta^+ = 1$. It can be seen that the variations in power and BER are much lower, and the convergence is smoother. This smoothness comes at a penalty in time since the entire scenario takes about 650 OPM evaluations. 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.

To conclude, we tested our algorithm running on a computer that controlled on optical devices. We showed that the insights 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. CONCLUSION

In this paper, we formulated a global optimization problem, *MLO*, that captures the QoS guarantees of optical networks. This problem is unique since there do not exist analytical models that capture impairments in optical fibers, and measurements are the only practical way to characterize performance. We designed a global network management algorithm, *SiMPLE*, for solving this problem by using feedback from real-life OPMs. The convergence of *SiMPLE* to the optimal solution is demonstrated using extensive simulations on a network-wide optical network simulator, as well as measurements with commercial optical network equipment. We showed that even simple dynamic policies in optical networks 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 current optical networks, using *SiMPLE* is the first step in allowing optical networks to react rapidly to user demands and traffic variations, and can lead to a software-defined optical network. In future work, we will investigate controlling modulation schemes and the optical bandwidth as part of this dynamic network control plane.

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 impairment-aware WSON control plane," http://www.cisco.com/en/US/prod/collateral/optical/ps5724/ps2006/data_sheet_c78-689160.html.
- [2] "Géant project," <http://www.geant.net>.
- [3] "Cisco visual networking index: Forecast and methodology, 2011-2016," *White Paper*, 2012.
- [4] *Proc. IEEE, Special Issue on The Evolution of Optical Networking*, vol. 100, no. 5, May 2012.
- [5] 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.
- [6] B. Birand, H. Wang, K. Bergman, and G. Zussman, "Measurements-based power control – a cross-layered framework," in *Proc. OFC'13*, Mar. 2013.
- [7] C. Chekuri, P. Claisse, R.-J. Essiambre, S. Fortune, D. C. Kilper, W. Lee, N. K. Nithi, I. Sanjeev, 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.
- [8] A. Conn, K. Scheinberg, and L. Vicente, *Introduction to derivative-free optimization*. Society for Industrial Mathematics, 2009, vol. 8.
- [9] R. Doverspike and J. Yates, "Optical network management and control," *Proc. IEEE*, vol. 100, no. 5, pp. 1092–1104, May 2012.
- [10] 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*, vol. 26, pp. B736–B745, Dec. 2011.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] 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.
- [15] Y. Lee, G. Bernstein, D. Li, and G. Martinelli, "A framework for the control of wavelength switched optical networks (WSONs) with impairments," Internet Engineering Task Force, RFC 6566, Mar. 2012.
- [16] E. Mannie, "Generalized Multi-Protocol label switching (GMPLS) architecture," Internet Engineering Task Force, RFC 3945, Oct. 2004.
- [17] G. Martinelli, "GMPLS Signaling Extensions for Optical Impairment Aware Lightpath Setup," Internet Engineering Task Force, Internet-Draft, 2010. [Online]. Available: <http://tools.ietf.org/html/draft-martinelli-ccamp-optical-imp-signaling-03.txt>
- [18] B. Mukherjee, *Optical WDM networks*. Springer-Verlag New York Inc, 2006.
- [19] J. Nocedal and S. Wright, *Numerical optimization*. Springer-Verlag, 1999.
- [20] 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.
- [21] Z. Pan, C. Yu, and A. Willner, "Optical performance monitoring for the next generation optical communication networks," *Optical Fiber Technology*, vol. 16, no. 1, pp. 20–45, 2010.
- [22] R. Ramaswami, K. Sivarajan, and G. Sasaki, *Optical networks: a practical perspective*. Morgan Kaufmann, 2009.
- [23] 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.
- [24] 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.
- [25] S. Woodward and M. Feuer, "Toward more dynamic optical networking," in *Proc. OECC'10*, July 2010.
- [26] Y. Wu, L. Chiaraviglio, M. Mellia, and F. Neri, "Power-aware routing and wavelength assignment in optical networks," in *Proc. ECOC'09*, Sept. 2009.