Stable Routing with Virtual Topology Capacity Adjustment: A novel paradigm for operating optical networks

Pablo Pavon-Marino, Nina Skorin-Kapov

***Abstract*— Advances in optical equipment permit network carriers to offer lightpath-on-demand services to Internet Service Providers (ISPs). These services are the key for constructing cost-efficient Virtual Topology Capacity Adjustment (VTCA) schemes, where ISPs dynamically adapt the number of lightpaths contracted between their IP routers, according to time-varying traffic volumes. An existing technique called lightpath bundling offers ISPs the possibility of grouping lightpaths between the same pair of routers into so-called bundles, perceived by the IP layer as single virtual links of aggregated capacity. Consequently, new lightpaths added or removed from bundles are not seen by the IP layer as new links necessary to advertise, but simply as capacity adjustments of already existing links. Adding lightpath bundling to the picture opens the door for developing VTCA schemes which maintain stability in the IP routing tables, a major requirement for ISPs. In this paper, we present and evaluate the merits of the proposed Stable Routing VTCA (SR-VTCA) paradigm and present algorithms for designing such a scheme. Results clearly show that SR-VTCA gives an advantageous trade-off between the fully static network with no capacity adjustment and the fully dynamic VTCA scheme where both the IP routing and the virtual topology are reconfigured over time.**

***Index Terms*— Optical fiber communications, virtual topology design, stable routing**

I. INTRODUCTION

Wavelength Division Multiplexing (WDM) technology and advances in optical switching are making it possible to build transparent WDM networks expected to satisfy the rapid growth of today's capacity demands. In such networks, transparent all-optical channels, called lightpaths, are established over the physical optical infrastructure. The set of established lightpaths comprises a so-called Virtual Topology (VT) where each virtual link corresponds to one lightpath. Packet-switched (IP) traffic is then routed over the virtual topology independently of the underlying physical optical network. In practice, the owners of the physical optical infrastructure are network operators or carriers who lease bandwidth, i.e. lightpaths, to Internet Service Providers (ISPs) or other institutional users of bandwidth [1]. The leased bandwidth is in the form of a set of limited-capacity lightpaths, i.e. a virtual topology, which can support the maximal expected traffic of the service provider. Generally, the ISP is not the owner of the physical optical network [2] and the network operator sets the leasing prices in the form of a constant bandwidth rate for the leased bandwidth over time. Naturally, ISPs and other bandwidth users prefer to minimize their leased capacity while supporting their traffic needs.

A. *Background and motivation*

The process of finding a set of lightpaths to support a given traffic demand is commonly referred to as Virtual Topology Design (VTD), and has been widely studied in the literature [3]-[6]. Common objectives include minimizing congestion, the number of lightpaths established and the average packet hop distance. These static planning approaches assume long-term traffic estimates which may not be ultimately precise. Short-term fluctuations of traffic, or long-term growth with unexpected shifts, can occur and lead to blocking. To overcome this, over-provisioning is employed to achieve robustness to traffic changes, but at a higher cost and degraded utilization, as well as higher energy consumption.

Instead of 'blind' over-provisioning of the calculated static topology, a method of 'smarter' over-provisioning for improved utilization can be achieved by planning according to known or expected traffic variations. Namely, real traffic traces, such as those collected for the Abilene network and the GEANT network (both available at [7]), indicate that traffic often follows expected periodic patterns, such as daily or weekly variations caused by peak hours. Knowledge of expected fluctuations allows the network planner to benefit from the sharing of lightpaths between time-disjoint traffic demands. This type of traffic is often modeled as a temporal

Manuscript received July 31st, 2011. Pablo Pavon-Marino is with the Technical University of Cartagena, Pza. Hospital 1, 30202, Cartagena, Spain (e-mail: pablo.pavon@upct.es).

Nina Skorin-Kapov is with Faculty of Electrical Engineering and Computing (FER), University of Zagreb, Unska 3, 1000, Zagreb, Croatia (email: nina.skorin-kapov@fer.hr).

sequence of traffic matrices each estimating the demand in a discrete time slot, referred as multi-hour traffic. The problem of finding a static virtual topology to support known multi-hour traffic has been investigated in [8] and [9].

To further reduce or even eliminate over-provisioning, lightpath reconfiguration approaches can be applied for better adaptation to traffic changes in a cost-effective manner. Initial work in this area was initiated with one-step virtual-topology reconfiguration approaches which migrate a current virtual topology to a target one which can support a new traffic demand [10]-[13]. The most common design objective is to minimize the associated reconfigurations to avoid traffic disruptions. However, such large-scale reconfigurations could still cause significant traffic disruption and are designed for long-term changes. As real traffic has been shown to exhibit daily or weekly variations, multi-hour traffic models have been applied to plan periodic virtual topology reconfigurations in advance, i.e. determine a periodic evolution of virtual topologies, in a cost effective manner [14], [15]. The problem of multi-period VTD considering not only periodic short term-fluctuations, but expected long-term changes in both the traffic demand and the physical topology, has been investigated in [16]. For cases when the traffic patterns are not known in advance, virtual-topology adaptation approaches have been developed which investigate lightpath on-demand scenarios based on real-time calculations of current traffic [1], [17].

While all the reconfigurable optical-layer approaches described above, such as VTD adaptation and reconfigurable multi-hour planning, can achieve efficient low-cost solutions to handle traffic variations, the most common approach employed by ISPs and other bandwidth users today is static planning with high over-provisioning. The reason for this is that approaches which assume a reconfigurable optical layer carry with them the assumption of variable flow routing, i.e. changing the routing in the upper (IP) layer as changes are made in the optical layer. However, this is something service providers strongly try to avoid. Namely, traffic flow rerouting complicates not only the control plane, but requires reconfiguration of monitoring services and the event correlation database in the management plane, as identified in [18]. Consequently, service providers prefer IP routing stability over efficient resource utilization and continue to aggressively over-provision, i.e. lease more lightpaths.

In an attempt to maintain routing stability while dynamically adding lightpaths to overcome short-lived traffic bursts, the authors of [18] propose an approach to mask the upper-layer flow rerouting from the IP routing protocols under the name of optical bypass. This method utilizes optical reconfiguration by adding temporary bypasses to overcome congestion peaks but without advertising the changes in the IP routing protocols. However, the router at ingress node of each bypass must be configured accordingly. This approach is well motivated but only provides additional temporary lightpaths for short-lived traffic bursts without VTD re-optimization and does not maintain truly fixed IP routing. Herein, we propose a new paradigm for operating optical networks, which enables efficient resource utilization via optical reconfiguration while maintaining true IP routing stability.

B. Our approach: Stable Routing (SR) with Virtual Topology Capacity Adjustment (VTCA)

For better understanding of the approach, some preliminary terms are introduced. A *lightpath bundle* refers to a set of lightpaths established between the same set of nodes. The concept of lightpath bundling was first proposed in Internet Draft "Link bundling in optical networks" (2001) and included in RFC 3717 in 2004 as a mechanism to reduce the signaling burden in the control plane, e.g. by advertising bundles instead of individual lightpaths. This was further generalized with Link Bundling functionality in RFC 4201 and 4202 for any GMPLS-based network. The concept of lightpath bundling allows IP routers to perceive new lightpaths added/removed between the same pair of nodes as capacity expansions/reductions, i.e. capacity adjustments, of existing virtual links and not new virtual links which need to be advertised in the IP layer. The virtual topology is then composed of a set of lightpath bundles, where each lightpath bundle is a virtual link of variable capacity with lightpath granularity. Adjusting bandwidth capacity per virtual link in GMPLS-based networks was investigated in [19]. More recently, in [20] the lightpath bundling approach was studied in conjunction with a technique called anycast switching, to improve the statistical multiplexing gain in the network. Here we apply the lightpath bundling concept to virtual topology reconfiguration as a mechanism to maintain routing stability.

In order to validate our approach, we first investigate where the main savings lie in standard virtual topology reconfiguration approaches. Namely, since existing optical reconfiguration approaches also assume variable routing, their obtained savings with respect to the static non-reconfigurable case comes from the flexibility on both layers. Our objective is to assess how much of these savings can come from reconfiguring only the optical layer, assuming fixed IP routing. A comparison of fixed and variable flow routing over a static virtual topology for multi-hour traffic has been investigated in [8] and [9], the results of which are elaborated upon in Section II. However, to the best of our knowledge, the comparison of reconfigurable and non-reconfigurable virtual topologies assuming fixed routing has not been studied, as all reconfigurable approaches assume variable routing.

In an effort to ease realization of our SR-VTCA concept and assess its merits in the aforementioned context, we provide an efficient scheme for finding a fixed routing which implies the least-cost evolution of VTDs (i.e. the total lightpaths needed over time) for expected traffic variations. This reduces the lightpath leasing costs for the ISP while maintaining routing stability. If no information is available regarding future traffic variations, the routing can be planned according to standard static VTD approaches, while unexpected traffic variations over time are handled via straightforward Virtual Topology Capacity Adjustment (VTCA) as needed. This achieves savings with respect to the static over-provisioning approach by establishing lightpaths on-demand while keeping routing stability thanks to lightpath bundling. However, if periodic traffic variations estimation models or data is available, increased savings can be achieved by finding a routing which corresponds to lower-cost capacity adjustments, i.e. less lightpaths hired over time. To develop our scheme we apply multi-hour traffic models representing expected traffic fluctuations of so-called background traffic. These fluctuations make up for the bulk traffic whose variations are assumed to be known and used in the optimization to determine the best routing. Note that once we define a fixed routing, the decisions of which lightpaths to dynamically add or remove along time are straightforward, i.e. put the capacity (lightpaths) where the traffic is as it changes. Additional unexpected short-lived traffic bursts outside the multi-hour estimates can be handled through standard VTCA on-demand.

We formulate the problem of Stable Routing with Virtual Topology Capacity Adjustment (SR-VTCA) assuming multi-hour traffic as an ILP (Integer Linear Program). The objective is to minimize the total number of lightpaths over time representing the leasing costs of the service providers. The proposed scheme does not consider the Routing and Wavelength Assignment (RWA) of each lightpath. The RWA design is the task of the network carrier and does not affect the ISP's perspective as long as the optical infrastructure is dimensioned appropriately. We assume such a situation in the paper, where all the lightpaths requested by the ISP have a valid route and wavelength.

Due to the high intractability of the ILP, we propose an efficient heuristic algorithm, called Sensitivity-based Iterative Rounding Algorithm (SIRA), for larger instances. We enable an additional degree of flexibility which allows the service provider to decide on the desired frequency of VT reconfiguration, e.g. every 4 hours. Nowadays bandwidth markets have started to appear which will allow bandwidth users to employ on-demand bandwidth provisioning at granularities even as small as five-minute blocks [21] so assumptions of the desired frequency can be highly variable.

We perform tests on multi-hour traffic based on real traffic traces from [7] and compare with analytical lower bounds on the multi-layer fully (non)-reconfigurable cases. Results indicate that significant savings (in terms of lightpaths) can be achieved by reconfiguring the optical layer while keeping the flow routing fixed with SR-VTCA, in comparison to the "smart over-provisioning" approach of planning a fixed routing over a static non-reconfigurable optical layer based on MH traffic. In contrast, full reconfiguration of both the optical layer and IP flow rerouting, only achieves minor savings with respect to SR-VTCA indicating that not much benefit can be gained from IP rerouting. These findings, in conjunction with the increased network management complexity associated with variable routing, indicate that planning a fixed routing over a reconfigurable optical network utilizing virtual topology capacity adjustment offers and advantageous trade-off between cost and routing stability, which is very much in line with today's ISP demands.

C. *Outline of the paper*

The rest of this paper is organized as follows. Section II gives a state of the art on multi-hour planning in optical networks, followed by the formal problem definition and exact ILP formulation in Section III. Heuristic approaches are presented in Section V while Section VI describes the experimental method and obtained results. Finally, Section VII concludes the paper.

II. RELATED WORK ON OPTICAL NETWORKS MULTI-HOUR PLANNING

Network planning taking into account time-varying traffic demands has been researched for various network technologies [22]. In particular, the concept of adjusting the capacity of the Virtual Paths in ATM networks while keeping routing stability was investigated in ATM networking, in works such as [23][24]. Dimensioning lightpath-based optical networks under multi-hour traffic models has more recently attracted attention of the research community. The Non-Reconfigurable Multi-hour Virtual Topology Design (MH-VTD-NR) problem was investigated in [8] and [9] which find a static VTD to support multi-hour traffic. Both approaches compare fixed (stable) and variable flow routing on the fixed VTD. In [8], various objective functions are considered, such as the number of lightpaths, average packet hop distance and link utilization. Results indicate that allowing flow rerouting achieves some savings with respect to the static case, but maximally in the order of 6% for number of lightpaths and less than 4% for the other criteria. In [9], the savings was measured in the total number of transceivers needed over time. The migration from fixed to variable routing was found to be justified only if splittable (bifurcated) routing is also allowed in the network, and only for medium or high loads. These results indicate that routing variability may not be the best approach, which corresponds to the routing stability constraints of ISPs.

Assuming optical-layer reconfiguration, planning approaches can further exploit multi-hour traffic fluctuation by not only sharing link (lightpath) capacity over time between time-disjoint traffic, but also sharing node (transponder) capacity between lightpaths over time. Thus, a periodic evolution of virtual topologies can be determined, i.e. a planned set of temporary lightpaths, which make use of common transponders installed at network nodes at different moments in time, exploit time-disjoint wavelength reuse, as well as reduce the total number of lightpaths required over time.

The idea of scheduled temporary lightpaths was proposed in [25] under the name of Scheduled Lightpaths Demands (SLDs) which assumes a known evolution of lightpaths. Work in this area was initially focused on planning the lower layer, i.e. routing and assigning wavelengths to a given time sequence of SLDs. In [25], the authors proposed branch and bound and tabu search approaches, while in [26] a more scalable tabu-search based heuristic and greedy algorithms are presented. [27] and [28] study a more general model, called the sliding scheduled traffic model where the known starting and holding times of lightpath demands are allowed to slide within a predefined window.

More recently, approaches which consider planning in the upper layer, i.e. determining the evolution or set of lightpaths from the periodic multi-hour traffic, called Reconfigurable Multi-hour Virtual Topology Design (MH-VTD-R) have been proposed. In [15], the authors present a MILP formulation and a decomposition heuristic for solving the non-reconfigurable case and compare its merits with a naïve approach for the reconfigurable variant where the VTDs are independently planned in each time interval. In [14], an efficient heuristic is proposed for MH-VTD-R with the objective to minimize the number of reconfigurations while minimizing the transceivers needed over time. However, both these approaches assume variable traffic flow routing in conjunction with the planned lightpath reconfiguration. Our approach of stable routing with virtual topology capacity adjustment also exploits multi-hour traffic and lightpath reconfiguration to minimize the cost (in terms of lightpaths), as do the MH-VTD-R approaches, but assumes fixed flow routing. Our results indicate that the main savings obtained by reconfiguration in fact lies in the dynamicity of the optical layer and not the variability of the routing.

III. AN ILP MODEL FOR THE SR-VTCA PROBLEM

In this section, we present an ILP (Integer Linear Programming) model for the Stable Routing with Virtual Topology Capacity Adjustment (SR-VTCA) problem. The target of the model is to find a stable flow routing and associated time-varying virtual topology which can support known time-varying traffic demands. The optimization target is minimizing the estimated cost of the capacity leased from the network carrier, assumed to be proportional to the number of lightpaths required per time interval.

Let N be the number of nodes in the network. The letters s,d,i,j,n are reserved as indices for nodes. We use the indices $s,d=\{1,...,N\}$ to denote the ingress and egress nodes of electronic traffic demands. In its turn, we use the indices $i,j=\{1,...,N\}$ for denoting the initiating and ending nodes of lightpath bundles. Let T be the number of time intervals for which the traffic is defined, and $t=\{1...T\}$ be the index indicating the time interval. Let M^t denote the traffic matrix at time slot t, and M_{ij}^t denote the traffic demand (measured in Gbps) from node i to node j, during time interval t. Let C denote the lightpath capacity in Gbps (e.g. 10 Gbps).

We assume that VT adjustments in the network can be made only at specific instants in the day/week. We call the time between two consecutive VT adjustments a non-reconfigurable epoch. Recall that the flow routing is kept unchanged along all the epochs. The number of non-reconfigurable epochs is denoted by R, and we use the letter $r=\{1,\ldots,R\}$ to index the epochs. We assume that all the non-reconfigurable epochs have the same duration, which is a multiple of the duration of the time intervals. Then, given an epoch $r=\{1,...,R\}$, the number of traffic matrices associated with it is given by $T/R = \{1,2,\ldots,T\}$. To simplify the notation, we denote as $r(t)$ the non-reconfigurable epoch index to which a time interval t belongs to, and we denote as $t(r)$ the *set* of consecutive time intervals associated to the epoch r. As an example, a plan based on T=168 traffic matrices corresponding to one traffic matrix per hour over the course of a week (24 hours x 7 days = 168), and the assumption of two reconfigurations per day, i.e. R=2x7=14 non-reconfigurable epochs per week, would make each epoch associated to T/R=12 matrices (twelve hours of real time).

The decision variables to the SR-VTCA problem are:

- $p_{ijr}\in\{0,1,2,...\}$. Number of lightpaths in the lightpath bundle from node i to node j, during non-reconfigurable epoch r.
- $x_{sdij}\in[0,1]$. Fraction of the total traffic demand from node s to node d that is routed over existing lightpaths in the lightpath bundle from node i to node j.

The problem formulation is given by (1).

$$\min \frac{1}{R}\sum_{i,j,r} p_{ijr} \tag{1a}$$

subject to:

$$\sum_j x_{sdnj} - \sum_i x_{sdin} = \begin{cases} 1, \text{if } n=s \\ -1, \text{if } n=d, n,s,d=1,\ldots,N \\ 0 \text{ otherwise} \end{cases} \tag{1b}$$

$$\sum_{s,d} M_{sd}^t x_{sdij} \le Cp_{ijr(t)}, i,j=1,\ldots,N,\ t=1,\ldots,T \tag{1c}$$

The objective function (1a) minimizes the average number of lightpaths required to support the traffic. This is an estimation of the total cost of leasing lightpaths for the ISP, assumed to be proportional to the number of lightpaths in the bundles and the amount of time each lightpath is active. Constraints (1b) are the routing conservation constraints for the stable routing x. Note that the traffic is routed on the lightpath bundles: a bundle is considered a single virtual link. Constraints (1c) are the capacity constraints in the lightpath bundles: the size (number of lightpaths) of the lightpath bundle at any moment in time should be enough to carry the traffic traversing it. In this formulation, it is assumed that all pairs of nodes can be interconnected with lightpaths. Consequently, the resulting virtual topology can be up to a full mesh with time-varying virtual link capacities defined by the number of parallel lightpaths between node pairs over time. Naturally, additional constraints can be included to limit the lightpath topology according to physical limitations in the network, as defined by the RWA problem. However, we do not consider this case here.

IV. THE SIRA ALGORITHM

In this section, we present the Sensitivity-based Iterative Rounding Algorithm (SIRA), for solving the SR-VTCA problem. Fig. 1 is a diagram illustrating the main steps of the SIRA algorithm. SIRA starts by solving the LP-relaxation of original problem (1) by dropping the integer constraints. If the solution of the relaxed problem, represented by a vector p of lightpath bundle sizes, is integral then the solution is optimal and the algorithm ends. If not, we round down the values in vector p and use it as a starting parameter for an iterative rounding-up stage of the algorithm. Vector p becomes the *tentative* lightpath bundling size vector.

Each iteration of the round-up stage of SIRA starts with Procedure P1, which involves solving linear program (2) shown below. An execution of LP (2) takes the tentative vector $p=(p_{ijr},\ i,j=1,\ldots,N,\ r=1,..,R)$ as an input parameter. It tries to find a stable routing $x=(x_{sdij},\ s,d,i,j=1,\ldots,N)$, and a *deficit* lightpath bundling size vector $s=(s_{ijr},\ i,j=1,\ldots,N,\ r=1,\ldots,R)$, so that a network with bundle sizes $(p+s)$ is able to support the traffic at every time. The name *deficit bundle size* is used to reflect that the bundle sizes given by p are not enough to carry the traffic, and thus have a deficit of capacity represented by s. Formulation (2) is as follows:

Find:

- $s_{ijr}\geq 0$. Size of the deficit in the lightpath bundle from node i to node j, during non-reconfigurable epoch r.
- $x_{sdij}\in[0,1]$. Fraction of the total traffic demand from node s to node d that is routed on the existing lightpaths in the lightpath bundle from node i to node j.

$$\min \sum_{i,j,r} s_{ijr} \tag{2a}$$

subject to:

$$\sum_j x_{sdnj} - \sum_i x_{sdin} = \begin{cases} 1, \text{if } n=s \\ -1, \text{if } n=d, n,s,d=1,\ldots,N \\ 0 \text{ otherwise} \end{cases} \tag{2b}$$

$$\sum_{s,d} M_{sd}^t x_{sdij} \leq C(p_{ijr(t)} + s_{ijr(t)}), i,j=1,\ldots,N,\ r=1,\ldots,R \tag{2c}$$

Equalities (2b) are the routing conservation constraints. Constraints (2c) ensure that bundle sizes $(p+s)$ are able to support the traffic at each moment. The objective function (2a) minimizes the deficit capacity that needs to be allocated. Note that since the deficit variables s are not restricted to be integer, Procedure P1 involves a Linear Program (LP), and can be completed in polynomial time.

If the deficit vector s resulting from Procedure P1 is zero in all its coordinates, then the routing x from (2) and the sizes p are a valid solution to the problem, the iterative section of the algorithm ends, and the algorithm continues with Procedure P3. If not, Procedure P2 is responsible for the rounding process. The new tentative vector p' is computed as follows. First, p' is set to $p'=p+s$. Then, vector p', which may have fractional coordinates because of the fractional nature of the deficit vector s, is *selectively* rounded, according to the current value of the utilization threshold U_{th}. The word selective in SIRA means that the coordinates (i,j,r) of p' with a fractional value equal or greater than U_{th} are rounded up, while the remaining coordinates are rounded down. Note that every coordinate of p' is always equal to or greater than the same coordinate in p, since we round vector $(p+s)$. In other words, the tentative capacity vector p is always non-decreasing along SIRA iterations.

The reasoning behind SIRA is that we iteratively increase

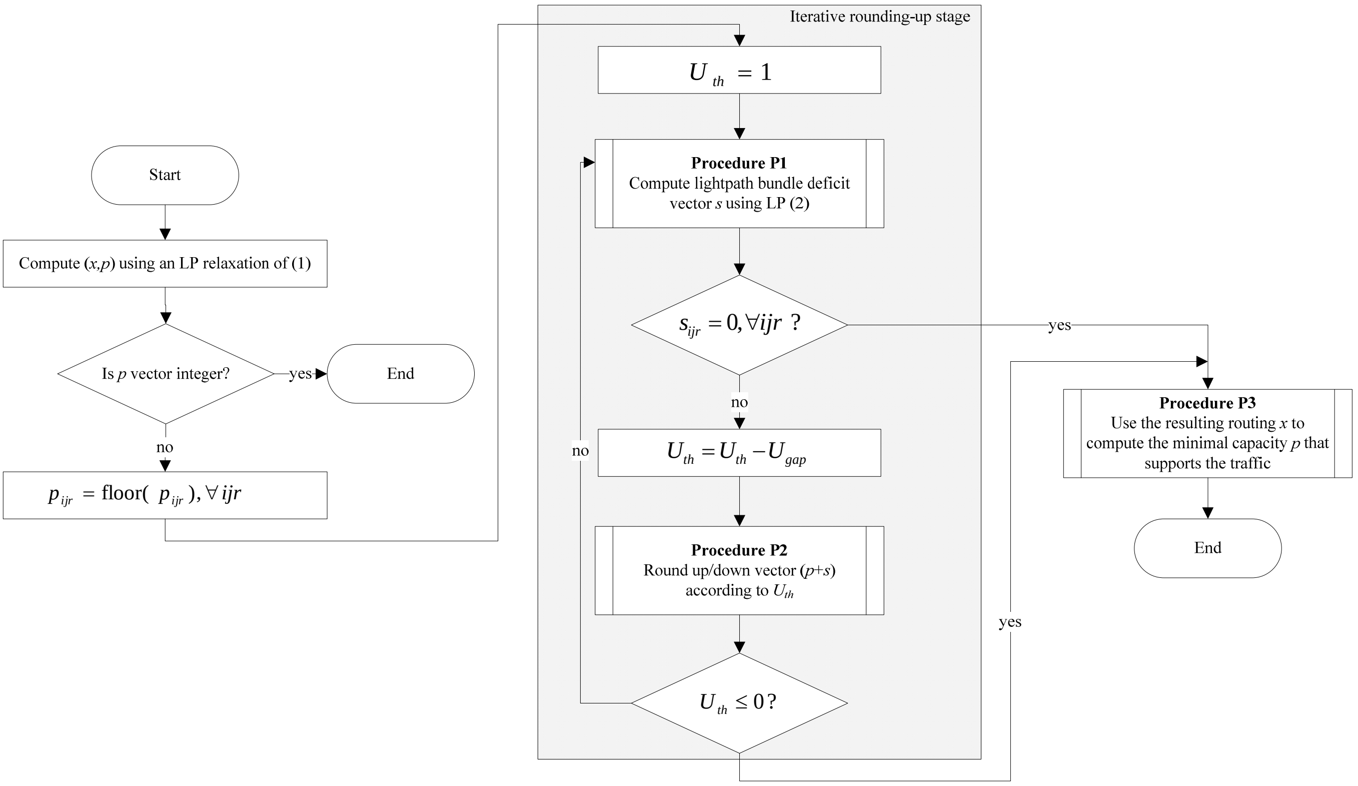

Fig. 1. SIRA algorithm

the tentative capacity vector p, and make the routing adapt to it. The extra capacity granted is accumulative, i.e., the capacity added to p in an iteration is used to route traffic in subsequent iterations. Then, given a tentative capacity p, formulation (2) finds a new routing minimizing the extra capacity needed *on top of p*. As the algorithm advances, the routing is shaped by (2) to make use of the allocated capacity, and consequently less and less excess capacity is needed.

Parameter U_{gap} is used to control the complexity vs. accuracy trade-off in SIRA. In the first execution of Procedure P2, the utilization threshold U_{th} has a value of $1-U_{gap}$. In subsequent iterations, U_{th} is decreased by parameter U_{gap}, and the algorithm ends when the utilization is zero. Then, the number of iterations of the algorithm is limited to a maximum of $1/U_{gap}$. As an example, if $U_{gap}=0.1$, ten iterations of SIRA corresponding to $U_{th}=\{0.9, 0.8, \ldots, 0.1, 0\}$ are completed. Note that in the iteration where $U_{th}=0$, *all* (i,j,r) triples which have a deficit of capacity ($s_{ijr}>0$) are rounded up, and the algorithm proceeds to Procedure P3. Also note that smaller U_{gap} parameters mean more SIRA iterations, and a more gradual adaptation of the routing to the discrete nature of the bundling sizes. Logically, this comes at a cost of a higher algorithm running time.

Finally, Procedure P3 takes as an input parameter the routing vector x computed by the SIRA iterations. Routing x is set as the final stable routing which determines the traffic at each time instant that is carried over each lightpath bundle (i,j). Then, it is trivial to compute, in every non-reconfigurable epoch r, the minimum cost VT to support the traffic in the lightpath bundles. This is given by (3).

$$p_{ijr} = \max_{t \in t(r)} \left\lceil \frac{\sum_{sd} M^{t}_{sd} x_{sdij}}{C} \right\rceil, i, j = 1, \ldots, N, r = 1, \ldots, R \qquad (3)$$

V. RESULTS

This section presents results of the tests conducted to assess the potential benefits of the SR-VTCA paradigm in terms of network cost reduction, as well as the efficiency of the SIRA algorithm as a planning approach within this paradigm. We compare the benefits of SR-VTCA with respect to (i) a network where both the routing and the virtual topology are stable (we refer to this as a *static network*), and (ii) a network where both the virtual topology and the flow routing can be changed to adapt to traffic variations (we refer to this as a *fully-reconfigurable network*).

A. Testing scenario and traffic normalization

We use data from real traffic traces available for two different topologies: the 11-node Abilene network and the 23-node GEANT network, both of them publicly available at [7]. This data consists of traffic matrices spanning several weeks obtained from monitoring campaigns. For both cases, we investigate weekly variations of the traffic by partitioning a week into $T=42$ consecutive time intervals of four hours each. For each time interval $t=1,\ldots,42$ in the week, we compute a representative traffic matrix M^t, averaging the corresponding values from the original trace, taken at the same time interval within the week for all the weeks of the trace. As a result, a sequence of $T=42$ matrices representing the average periodic weekly variations in the traffic is obtained for each network topology (Abilene and GEANT).

Given such a sequence of traffic matrices M^t, we normalize its load to five different traffic load factors $\rho=\{0.5, 1, 2, 5, 10\}$ in an effort to assess the merits of the various approaches for different traffic load conditions. Value ρ represents the average amount of traffic between two nodes during the highest-loaded time slot measured in the number of lightpaths. Consequently, a value of $\rho=0.5$ corresponds to the case when the average traffic between two nodes in the peak traffic time slot of the week equals 50% of a single lightpath capacity. On the contrary, a value of $\rho=10$ captures cases in which the average traffic between two nodes in the highest loaded time slot fills on average 10 full lightpaths. We denote as $M^t(\rho)$ the normalized sequence of $T=42$ traffic matrices corresponding to load factor ρ.

In addition, we investigate the evolution of the network cost with respect to different durations of the non-reconfigurable epochs. We denote the duration of the epochs as D_e, and test for cases of $D_e=\{4h, 8h, 12h\}$. In the first case with $D_e=4h$, the network is reconfigured every four hours, while in the last case with $D_e=12h$, the network is reconfigured twice a day. Note that these cases correspond to $R=7\times6=42$ and $R=7\times2=14$ non-reconfigurable epochs in a week, respectively.

Finally, we assume that lightpaths between two nodes (i,j) in the network can be established only if at least one of the following two conditions is met: (i) there is a direct fiber link connecting nodes i and j, (ii) the length of the shortest path (in kilometers) between both nodes is below or equal to 2000 km. Namely, in lightpath-based networks, optical signal degradation constrains the maximum reach of lightpaths. The aforementioned conditions intend to capture limitations in the VT design that the ISP may experience due to the impossibility of leasing lightpaths between excessively distant nodes. The length 2000 km is consistent with the practical limits of currently available long-haul 10 Gbps equipment [29]. Incorporating such impairment-based constraints into the SR-VTCA problem formulation, as well as the SIRA algorithm, can easily be achieved by adding constraint (4) to formulations (1) and (2), respectively.

$$x_{sdij} = 0, s, d = 1, \ldots, N, (i, j) \text{ excessively distant nodes} \qquad (4)$$

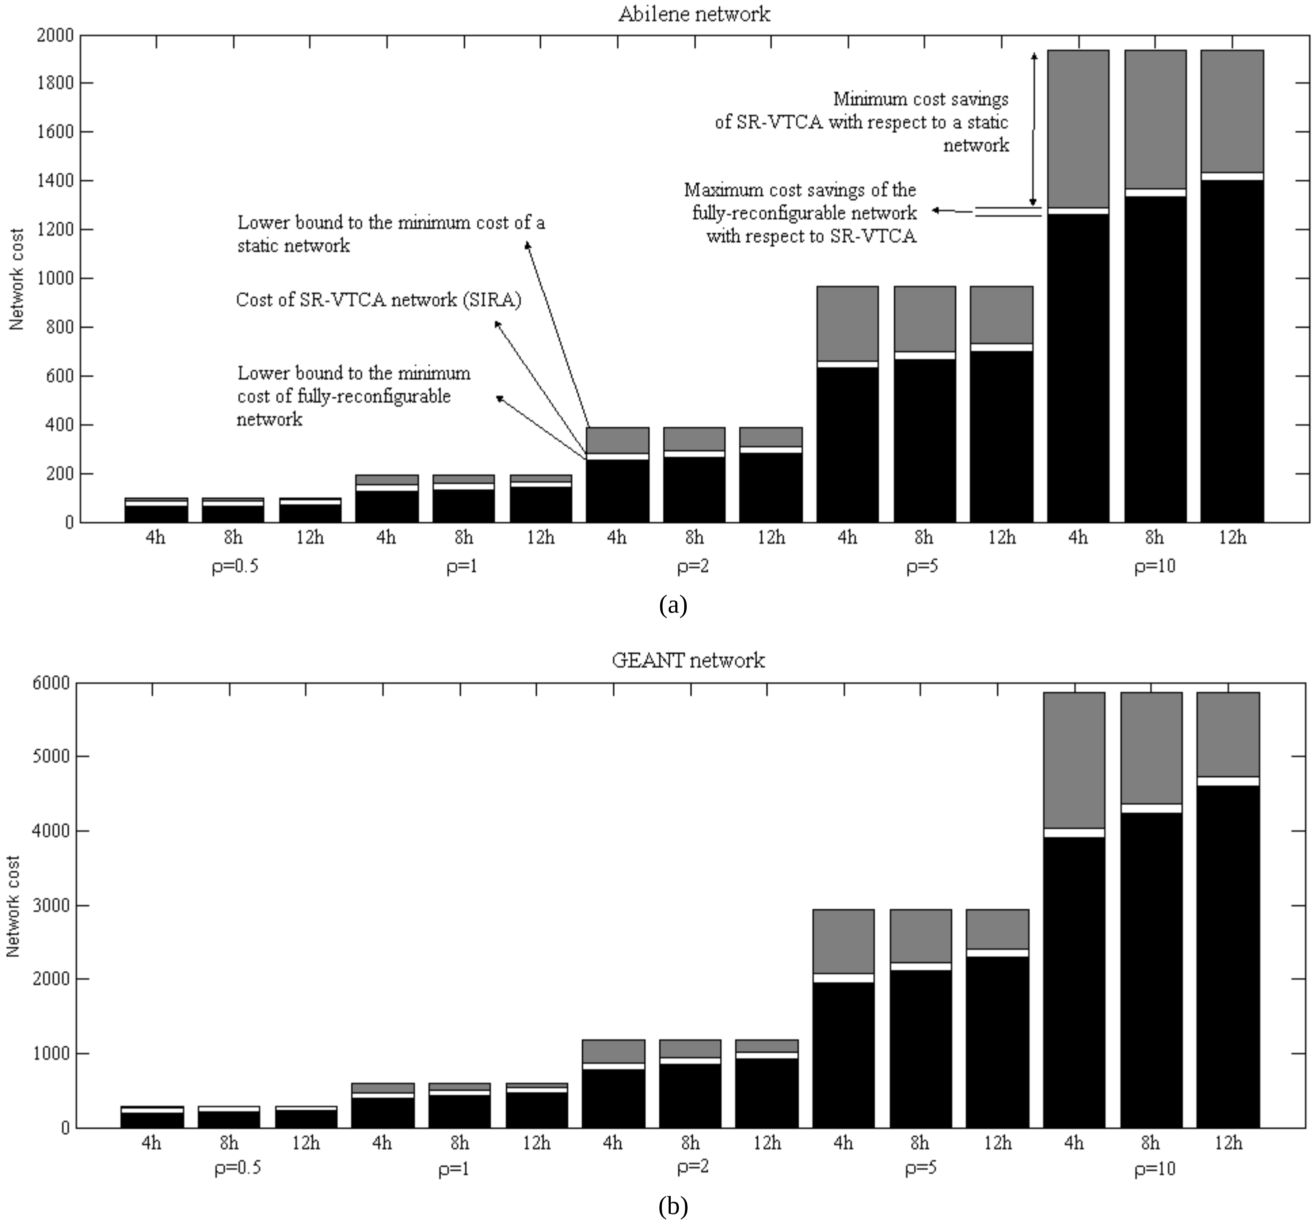

(a)

(b)

Fig. 2. Cost comparison static network (lower bound) / SR-VTCA (SIRA) / fully-reconfigurable network (lower bound), (a) Abilene network, (b) GEANT network. The x-axis ranges different load factors ρ={0.5, 1, 2, 5, 10} and reconfiguration epoch durations D_e={4h, 8h, 12h}.

B. Benefits of the SR-VTCA approach

In this sub-section, we evaluate the merits of the SR-VTCA approach, as an advantageous compromise between a static network and a fully-reconfigurable network. To do so, we compute the costs of the Abilene and GEANT networks applying the SR-VTCA paradigm using the SIRA algorithm. SIRA was implemented in the MatPlanWDM tool [30], which interfaces to the TOMLAB/CPLEX solver [31] for obtaining the solutions to the LP formulations involved. The results shown correspond to those obtained for parameter value U_{gap}=0.05 which limits the number of SIRA iterations to 1/0.05=20. The value of this parameter was determined experimentally. Namely, for U_{gap} values lower than 0.05, which implies more rounding iterations, the algorithm produced no improvement in the final solutions at the expense of larger execution times. Alternatively, setting the parameter U_{gap} to values higher than 0.05 to reducing the number of iterations ran faster, but produced a constant, albeit slight, increase in the cost of the obtained solutions.

As already mentioned, we ran SIRA for load conditions ρ={0.5, 1, 2, 5, 10} and for durations of the non-reconfigurable epochs set to D_e={4h, 8h, 12h}. In order to assess the guaranteed savings of our approach with respect to the static case, as well as the maximal possible savings of the fully reconfigurable case in comparison to SR-VTCA, we compare with associated lower bounds. Thus, we calculate the following two bounds:

- LB_{fr}: A *Lower bound* on the network cost of a fully-reconfigurable network. In a fully-reconfigurable network, the planning problem can be decoupled into R independent planning problems, one per non-reconfigurable epoch. Then, a lower bound to the optimal cost can be obtained by computing the optimal cost of the LP relaxations of formulation (1), applied to each epoch.
- LB_s: A *Lower bound* on the network cost of a static network. This lower bound is given by the optimal cost of the LP relaxation of problem (1), for the case where R=1 (i.e., there is only one non-reconfigurable epoch which implies no network reconfiguration at all).

Figs. 2-(a) and 2-(b) report the results obtained for the Abilene and GEANT cases respectively. The top of the black bars, represent the LB_{fr} cost, i.e., the minimum possible cost of the network for the fully-reconfigurable case. The top of the white bars show the SR-VTCA cost, while the top of the grey bars mark the lower bound on the cost of the static network (LB_s). This configuration of the bar plot helps us to illustrate the advantages of the SR-VTCA approach in a visual way. In particular:

- The size of the grey bar represents the *minimum (guaranteed)* cost saving obtained if we migrate from a static network to a network planned and operated following the SR-VTCA paradigm.
- The size of the white bar represents the *maximum possible* extra savings that can be achieved if we migrate from SR-VTCA to a fully-reconfigurable network plan.
- The size of the black bar is the *unavoidable* network cost: a lower bound to the cost of a fully-reconfigurable network.

We observe the following. First, for all the network conditions, except low loads (ρ=0.5), the white bars have a much smaller size with respect to the grey bars. In other words, the main reduction in the ISP network cost with respect to the static case comes from reconfiguring the virtual topology (the SR-VTCA approach), while only marginal extra cost reductions can be achieved if the routing is also allowed to be variable. The advantages of the SR-VTCA approach are increasingly evident for higher loads. As an example, if the network (Abilene-GEANT) is reconfigured every four hours, the SR-VTCA approach yields a *minimum* saving of 24%-28% (ρ=1), to 45%-50% (ρ=10) with respect to the static approach. In its turn, the maximum extra savings if also the routing reconfigurations are permitted, is in the order of 17% (ρ=1) to 2%-3% (ρ=10).

Naturally, shorter non-reconfigurable epochs implying more frequent reconfigurations give better results. However, VT reconfiguration twice daily (D_e=12h) already seems enough to obtain very relevant cost savings with respect to static networks, especially for high loads. Moreover, note that the results shown reflect a comparison between *guaranteed* savings achieved by SR-VTCA with respect to the *maximum achievable* extra savings of the fully-reconfigurable case, i.e., the worst cases for SR-VTCA. In other words, the relative merits of the SR-VTCA approach with respect to the other options may even be higher in reality. In particular, for the low load case (ρ=0.5) where the guaranteed savings of SR-VTCA seem to shrink, it cannot be deduced whether this is caused by a true lack of benefits of the SR-VTCA case or a lack of accuracy of the lower bound LB_s for these scenarios.

In summary, for the cases tested, results indicate that the SR-VTCA paradigm gives a clearly advantageous trade-off between the static and fully-reconfigurable approaches obtaining most of the potential cost savings derived from network reconfiguration, while keeping routing stability in accordance with ISP preferred requirements. The guarantee of these benefits increases with the network load.

C. *Assessing the quality of the SIRA algorithm*

This section is devoted to give an insight on the merits of the SIRA algorithm proposed in this paper, as a method for solving the SR-VTCA problem.

Table I reports the optimality gap of the SIRA algorithm (run with parameter value U_{gap}=0.05), i.e., the gap between the cost of the solution provided by SIRA and a lower bound on the optimal cost of the SR-VTCA problem. This lower bound is obtained by solving the LP-relaxation of (1). The optimality gap values are given as a percentage of the obtained SIRA cost.

Results show that optimality gap decreases approximately according to sequence 20%, 15%, 10%, 5% and 2.5% for increasing values of the load factor ρ={0.5, 1, 2, 5, 10}. Results are consistent for both topologies and various durations of the non-reconfigurable epochs D_e. Consequently, results support the efficiency of the SIRA algorithm for solving the SR-VTCA problem. In particular, the narrow optimality gaps at medium and high loads are a valuable proof of the success of SIRA in these scenarios.

TABLE I
OPTIMALITY GAPS

	Abilene			GEANT		
ρ	D_e=4h	D_e=8h	D_e=12h	D_e=4h	D_e=8h	D_e=12h
0.5	25%	23%	22%	26%	23%	21%
1	17%	15%	14%	17%	15%	14%
2	10%	9%	9%	11%	10%	9%
5	4%	4%	4%	6%	5%	5%
10	2%	2%	2%	3%	3%	3%

D. *SIRA running complexity*

SIRA algorithm has a polynomial complexity. The maximum number of iterations executed is given by ($1/U_{gap}$), and can thus be determined in advance. In each rounding iteration, an LP program is solved. In practice, the complexity of these LP programs determines the running complexity of SIRA.

Table II reports the execution times of the SIRA algorithm with parameter U_{gap}=0.05 (which produced the results shown in Figs. 2-(a) and 2-(b)) run on a laptop with an Intel(R) Core(TM) i5 CPU, 2.27 GHz and 8 GB of RAM. The running times of SIRA are always below 30 seconds for the Abilene network. The running time increases somewhat with network loads, ranging from about 10 seconds for ρ=0.5 and about 25 seconds for ρ=10. In turn, the running times of SIRA for the GEANT network range from ~8 minutes to a maximum of 4.25 hours. The significant increase in the running times associated with the network size (Abilene vs GEANT) is explained by an increase in the number of decision variables and constraints of the problem. In addition, running times for the larger GEANT case tend to increase for higher values of D_e. Intuitively, higher values of D_e create a more coupled problem since the virtual topology at each non-reconfigurable epoch supports more traffic matrices. Since the sophisticated CPLEX solver achieves it efficiency by making various problem reductions, this coupling can make it more difficult to achieve such reductions resulting in longer execution times.

TABLE II
SIRA EXECUTION TIME (MINUTES). U_{GAP}=0.05

	Abilene			GEANT		
ρ	D_e=4h	D_e=8h	D_e=12h	D_e=4h	D_e=8h	D_e=12h
0.5	25%	23%	22%	26%	23%	21%
1	17%	15%	14%	17%	15%	14%
2	10%	9%	9%	11%	10%	9%
5	4%	4%	4%	6%	5%	5%
10	2%	2%	2%	3%	3%	3%

VI. CONCLUSIONS

In this paper, we present a novel paradigm for operating reconfigurable optical networks while maintaining IP routing stability, a major concern for ISPs. Namely, although planning fully reconfigurable networks with changing lightpaths and variable traffic routing can achieve significant bandwidth-leasing cost savings for ISPs, in practice, ISPs still apply static over-provisioning in order to keep their IP routing stable. Intuitively, if we reconfigure the lower layer (i.e., the lightpaths), it seems logical that adapting the upper layer (i.e., the traffic flows) which are routed over these lightpaths is also needed. However, by applying lightpath bundling to group parallel lightpaths into single virtual links of variable capacity, this assumption need not be so. Our approach, called Stable Routing with Virtual Topology Capacity Adjustment (SR-VTCA) is based on enabling efficient lightpath (virtual topology) reconfiguration in the form of capacity adjustment, while maintaining higher layer IP routing stability. By applying traffic prediction models, such as the multi-hour model, fixed routing schemes can be found which allow for virtual topology adjustments which can achieve almost as much savings as the fully reconfigurable case. For the multi-hour model, we give an exact problem formulation, along with an efficient heuristic algorithm called SIRA (Sensitivity-based Iterative Rounding Algorithm) to obtain suboptimal solutions for larger instances. To assess the guaranteed benefits of the SR-VTCA paradigm with respect to the fully static and reconfigurable cases, we compare with associated lower bounds. Results for real traffic traces show that the main savings achieved from network reconfiguration lies in reconfiguring the virtual topology, while only marginal benefits come from routing variability. Consequently, the SR-VTCA concept of maintaining a static IP-layer on top of a reconfigurable optical layer provides an excellent trade-off between cost savings and network stability. These findings support the validity of this paradigm as a promising solution able to cope with the growing dynamicity of next-generation optical networks.

ACKNOWLEDGMENT

This research was partially supported by the MICINN/FEDER project grant TEC2010-21405-C02-02/TCM (CALM) and MICINN project grant TEC2010-12250-E (FIERRO). It was also developed in the framework of project "Programa de Ayudas a Grupos de Excelencia de la Región de Murcia” funded by F. Séneca (Plan Regional de Ciencia y Tecnología 2007/2010), project “A Security Planning Framework for Optical Networks (SAFE),” funded by the Unity Through Knowledge Fund (UKF) in Croatia and project 036-0362027-1641, funded by the Ministry of Science, Education and Sports, Croatia.

REFERENCES

[1] S. F. Gieselman and N. K. Singhal, B. Mukherjee, "Minimum-cost virtual-topology adaptation for optical WDM mesh networks," in *Proc. of IEEE International Conference on Communications* (IEEE, 2005) pp. 1787-1791.

[2] E. L Varma, S. Sankaranarayanan, G. Newsome, Z. Lin and H. Epstein., “Architecting the services optical network,” *IEEE Commun. Magazine*, vol. 39, no. 9, pp. 80-87, Sept 2001.

[3] R. Dutta and G. N. Rouskas, “A survey of virtual topology design algorithms for wavelength routed optical networks,” *Optical Networks Magazine*, vol. 1, no.1, pp. 73-89, Jan. 2000.

[4] R. Ramaswami and K. Sivarajan, “Design of logical topologies for wavelength-routed optical network*”, IEEE Journal on Selected Areas in Communications,* vol. 14, no. 5, pp. 840-851, June 1996.

[5] R. M. Krishnaswamy and K. N. Sivarajan, “Design of logical topologies: a linear formulation for wavelength-routed optical networks with no wavelength changers”, *IEEE/ACM Transactions on Networking*, vol. 9, no. 2, pp. 186-198, April 2001.

[6] B. Mukherjee, *Optical WDM Networks*, Springer, 2006

[7] TOTEM Project: Toolbox for Traffic Engineering Methods, http://totem.run.montefiore.ulg.ac.be/datatools.html

[8] G. Agrawal and D. Medhi, “Lightpath Topology Configuration for Wavelength-Routed IP/MPLS Network for Time-Dependent Traffic”, in *Proc. of IEEE Global Communications Conference (GLOBECOM)* (IEEE 2006), pp. 1-5.

[9] P. Pavon-Marino, R. Aparicio-Pardo, B. Garcia-Manrubia and N. Skorin-Kapov, "Virtual topology design and flow routing in optical networks under multi-hour traffic demand", *Photonic Network Communications*, vol. 19, no. 1, pp.42-54, February 2010.

[10] D. Banerjee and B. Mukherjee, “Wavelength-routed optical networks: Linear formulation, resource budgeting tradeoffs, and a reconfiguration study,” *IEEE/ACM Transactions on Networking*, vol. 8, no. 5, pp. 598-607, Oct. 2000.

[11] E. Bouillet, J.-F. Labourdette, R. Ramamurthy, and S.Chaudhuri, "Lightpath re-optimization in mesh optical networks", *IEEE/ACM Transactions on Networking*, vol. 13, no. 2, April 2005, pp. 437 – 447.

[12] G. N. Roukas and M. H. Ammar, “Dynamic reconfiguration in multihop WDM networks,” *Journal of High Speed Networks*, vol. 4, no. 3, pp. 221-238, 1995.

[13] M. Saad and Z. Luo, “Reconfiguration with no service disruption in multifiber WDM networks”, *IEEE/OSA Journal*

of Lightwave Technology, vol. 23, no. 10, pp. 3092-3104, October 2005.

[14] R. Aparicio-Pardo, N. Skorin-Kapov, P. Pavon-Marino and B. Garcia-Manrubia, "(Non)-Reconfigurable Virtual Topology Design under Multi-hour Traffic in Optical Networks", submitted to IEEE/ACM Transactions on Networking.

[15] F. Ricciato, S. Salsano, A. Belmonte, and M. Listanti, "Off-line Configuration of a MPLS over WDM Network under Time-Varying Offered Traffic," in *Proc. IEEE International Conference on Computer Communications (Infocom)*, (IEEE 2001) pp. 57-65.

[16] P. Manohar, A. Padmanath, S. Singh, and D. Manjunath, "Multiperiod Virtual Topology Design in Wavelength Routed Optical Networks", *IEE Proc. - Circuits, Devices and Systems*, vol. 150, no. 6, pp. 516-520, Dec. 2003.

[17] A. Gencata and B. Mukherjee, "Virtual-topology adaptation for WDM mesh networks under dynamic traffic," *IEEE/ACM Transactions on Networking*, vol. 11, no. 2, pp. 236-247, April 2003.

[18] M. Chamania, M. Caria, A. Jukan, "Effective Usage of Dynamic Circuits for IP Routing", *Elsevier Journal on Optical Switching and Networking*, vol. 7, no. 4, pp. 173-184, December 2010.

[19] S. Gunreben, S. Spadaro, J. Sole-Pareta, "A unified model for bandwidth adaptation in next generation transport networks", in *Proc. IEEE International Workshop on Bandwidth on Demand* (IEEE , 2006), pp. 58-64.

[20] P. Pavon-Marino, "Lightpath bundling and anycast switching: a good team for multilayer optical networks", in *Proc. 15th Conference on Optical Network Design and Modelling (ONDM 2011)*, pp. 1-6.

[21] C. Cavdar, A. Gencata, B. Mukherjee, "CATZ: Time-zone-aware Bandwidth Allocation in Layer 1 VPNs", *IEEE Communications Magazine*, vol. 45, no. 4, pp. 60-66, April 2007

[22] M. Pioro and D. Medhi, *Routing, Flow and Capacity Design in Communication and Computer Networks*. San Francisco, CA: Morgan Kaufmann, 2004, pp. 455–474.

[23] M. Logothetis and S. Shioda, "Medium-term centralized virtual-path bandwidth control based on traffic measurements", *IEEE Transactions on Commununications*. vol. 43, no. 10, pp. 2630-2640, October 1995.

[24] S. Shioda and H. Saito, "Sizing and provisioning for physical and virtual path networks using self-sizing capability", vol. E80-B, no.2, pp. 252-262, February 1997.

[25] J. Kuri, N. Puech, M. Gagnaire, E. Dotaro, and R. Douville, "Routing and wavelength assignment of scheduled lightpath demands," *IEEE Journal on Selected Areas in Communications*, vol. 21, no. 8, pp. 1231–1240, October 2003.

[26] N. Skorin-Kapov, "Heuristic algorithms for the routing and wavelength assignment of scheduled lightpath demands in optical networks", *IEEE Journal on Selected Areas in Communication*, vol. 24, no. 8, pp. 2–15, August 2006.

[27] B.,Wang, T. Li, X. Luo, Y. Fan, and C. Xin, "On service provisioning under a scheduled traffic model in reconfigurable WDM optical networks" in *Proc. Of BroadNets 2005*, pp 13-22.

[28] C. V. Saradhi, M. Gurusamy, and R. Piesiewicz, "Routing fault-tolerant sliding scheduled traffic in WDM optical mesh networks" in *Proc. of BroadNets 2008,* pp.197-202.

[29] NOBEL phase 2 –Deliverable D 2.4. "Migration Guidelines with Economic Assessment and New Business Opportunities Generated by NOBEL phase 2".

[30] P. Pavon-Mariño, R. Aparicio-Pardo, G. Moreno-Muñoz, J. Garcia-Haro, J. Veiga-Gontan, "MatPlanWDM: An educational tool for network planning in wavelength-routing networks," in *Proc. 11th Conference on Optical Network Design and Modelling (ONDM 2007)*, pp. 58-67.

[31] TOMLAB Optimization. Available: http://tomopt.com/.

Pablo Pavon-Marino received the Telecommunications Engineering, and Ph.D. degrees from the University of Vigo (Spain) in 1999 and from the Technical University of Cartagena (Spain) in 2004 respectively. He is Associate Professor in the Department of Information Technologies and Communications of the Technical University of Cartagena. He has authored or coauthored over 60 papers published in international journals and conferences. His research interests include optimization and planning of optical networks.

Prof. Pavon-Marino is part of the Editorial Board of the Optical Switching and Networking Journal. He was General Co-Chair of the 2011 IEEE High Performance Switching and Routing conference (IEEE HPSR 2011), and was a member of the technical program committees of several other conferences.

Nina Skorin-Kapov received the Dipl.-Ing. and Ph.D. degrees in electrical engineering in the field of telecommunications from the University of Zagreb, Zagreb, Croatia, in 2003 and 2006, respectively. She has been an Assistant Professor at the Department of Telecommunications at the same university since 2008. She completed a Post-Doctoral Fellowship at Télécom ParisTech (formerly ENST) in Paris, France, in academic year 2006-7. Her main research interests include optimization of communication networks and network security.