(Non)-Reconfigurable Virtual Topology Design under Multi-hour Traffic in Optical Networks

Ramon Aparicio-Pardo, Nina Skorin-Kapov, Pablo Pavon-Marino and Belen Garcia-Manrubia.

Abstract—**This paper investigates offline virtual topology design in transparent optical networks under a multi-hour traffic demand. The main problem variant addressed here designs a *reconfigurable* virtual topology which evolves over time to more efficiently utilize network resources (the MH-VTD-R problem). The case of designing a static *non-reconfigurable* virtual topology which can accommodate the time-varying traffic (the MH-VTD-NR problem) is also considered. The objectives are to minimize (i) the number of transceivers, which make up for the main network cost; and, (ii) the frequency of reconfiguration (for MH-VTD-R), which incurs additional overhead and potential service disruption. We formulate this multi-objective problem as an exact MILP (Mixed Integer Linear Program). Due to its high complexity, we propose a very efficient heuristic algorithm called GARF (Greedy Approach with Reconfiguration Flattening). GARF not only solves both (non)reconfigurable problem variants, but it allows for tuning of the relative importance of the two objectives. Exhaustive experiments on real and synthetic traffic, and comparison with previous proposals and bounds, reveal the merits of GARF with respect to both solution quality and execution time. Furthermore, the obtained results indicate that the maximal transceiver cost savings achieved by the fully reconfigurable case may not be enough to justify the associated increase in reconfiguration cost. However, results show that an advantageous trade-off between transceiver cost savings and reconfiguration cost can be achieved by a allowing a small number of virtual topology reconfigurations over time.**

Index Terms— **Transparent optical networks, reconfigurable virtual topology design, multihour traffic.**

I. INTRODUCTION

Transparent optical networks based on wavelength division multiplexing (WDM) have been accepted as the enabling technology for high-speed backbone networks [1][2]. In transparent networks, traffic is carried over all-optical connections, called lightpaths. A lightpath originates at a transmitter and terminates at a receiver (together referred to as transceivers), and occupies a single wavelength channel in each traversed link. Since the traffic carried over a lightpath is not processed electronically at intermediate nodes, savings with respect to electronic switching equipment is achieved.

A set of lightpaths established over the physical topology creates a so-called virtual topology of all-optical connections. Virtual topology design (VTD) implies planning a set of lightpaths to carry a given set of electronic traffic demands or electronic traffic flows (i.e. measured in Gbps). In the upper layer, the traffic flows are routed on top of the virtual topology. In the lower layer, each lightpath in the virtual topology has to be routed over the physical topology and assigned a wavelength. This implies solving the Routing and Wavelength Assignment (RWA) problem [3].

In this paper, we consider a variant of the VTD problem, referred to as Multi-Hour Virtual Topology Design (MH-VTD), which addresses the planning of transparent optical networks under multi-hour or periodic traffic. In this case, the traffic demand takes the form of a temporal sequence of traffic matrices, reflecting the traffic variation along a given period of time (typically days or weeks). The periodic nature of the traffic has been confirmed with real traffic traces in numerous experiments, such as the Abilene backbone network [4], making the traffic load in the network fairly predictable.

This paper investigates offline planning of transparent optical networks for a given multihour traffic demand, comparing two scenarios: (i) planning a fixed virtual topology which is sufficient to carry the traffic demand at any time, (ii) planning a virtual topology which can change along time to more efficiently adapt to known traffic variations. For both approaches, we assume the general case where the routing of the electronic flows can be bifurcated (splittable) and time-varying. These two MH-VTD problem variants are denoted as MH-VTD-NR (non-reconfigurable) and MH-VTD-R (reconfigurable) and are illustrated in Fig. 1 and Fig. 2, respectively. For both problems, we consider the number of transceivers in the network to be the measure of the network cost to optimize. This is a common assumption for VTD problems [5]-[7]. A second criterion used to evaluate the solutions obtained for the MH-VTD-R problem, is the number of reconfigurations associated with the evolution of the virtual topology design.

Investigating the multi-hour problem variants allows us to explore the following trade-off. On the one hand, being able to change the virtual topology design along time may involve

The work described in this paper was carried out with the support of the BONE project ("Building the Future Optical Network in Europe"); a Network of Excellence funded by the European Commission through the 7th ICT-Framework Program. This research was partially supported by the Spanish project grants TEC2010-21405-C02-02/TCM (CALM) and TEC2010-12250-E (FIERRO) and the Croatian research projects UKF SAFE and MZOS 036-0362027-1641. It was also developed in the framework of "Programa de Ayudas a Grupos de Excelencia de la R. de Murcia, F. Séneca".

Ramon Aparicio-Pardo, Pablo Pavon-Marino and Belen Garcia-Manrubia are with Univ. Polit. de Cartagena (UPCT), Pza. Hospital 1, 30202, Cartagena, Spain (e-mail: {ramon.aparicio, pablo.pavon, belen.garcia}@upct.es).

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

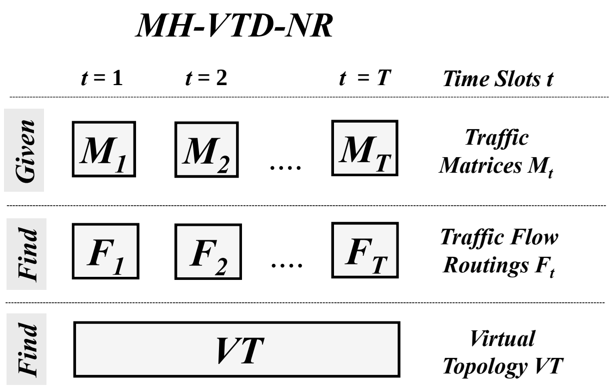

Fig. 1. Non-Reconfigurable Multi-Hour Virtual Topology Design

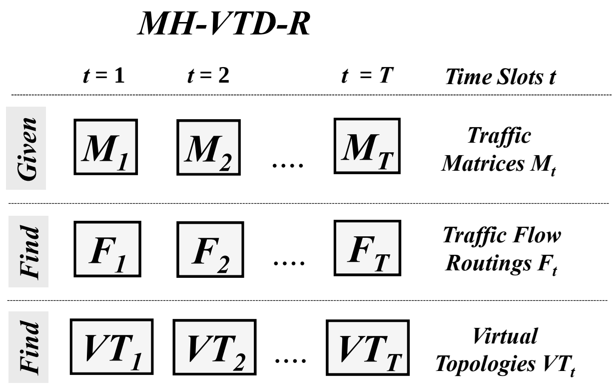

Fig. 2. Reconfigurable Multi-Hour Virtual Topology Design

savings in optical resources. This occurs when a node sends (receives) traffic to (from) different nodes at different peak hours. In such cases, part of the transceivers used by the nodes at certain time slots could be re-used for other lightpaths established at a different time. However, the ability to adapt the virtual topology comes at a cost. First, frequent reconfigurations of the virtual topology to adapt to traffic changes imply a significant increase in signaling complexity. Second, each reconfiguration interferes with existing traffic and disrupts network performance degrading the quality of service of the associated connections [8]. In summary, we are interested in assessing the reduction in the number of transceivers in the network that can be achieved with virtual topology reconfiguration, depending on the reconfiguration frequency.

It is important to remark that by solving the MH-VTD problem we obtain a set of lightpaths over time, together with the corresponding flow routing, but do not solve the RWA for each lightpath. This is because defining the evolution of lightpaths to be established between node pairs totally determines the number of transceivers and reconfigurations in the network, and thus determines the network cost differences and trade-off we intend to evaluate. Of course, solving the RWA problem for the determined set of lightpaths certifies the feasibility of the network plan with respect to wavelength availability and physical impairments. However, here we assume network scenarios with an over-provisioned fiber plant where a feasible RWA solution for almost any VTD exists, focusing the planning solely on virtual topology design and flow routing.

In this paper, we give a mixed integer-linear formulation (MILP) for the MH-VTD-R problem with the objective to minimize both transceivers and reconfiguration. Previously, we proposed MILP formulations for the MH-VTD-NR problem in [9]. Clearly, all the MH-VTD problem variants are NP-hard, as the reduced problem with constant traffic is known to be NP-hard (integer capacity planning) [10]. Our tests indicate that solving the problem optimally is feasible only for small network sizes and for a moderate number of time intervals. Consequently, we propose a very efficient heuristic approach, referred to as the GARF (Greedy Approach with Reconfiguration Flattening) algorithm. This heuristic can be tuned to allow various degrees of reconfiguration, from the extreme case where no reconfiguration is allowed (i.e., the MH- VTD-NR problem), to the case where reconfiguration is for 'free', i.e. the number of reconfigurations is not penalized at all. The paper provides an exhaustive evaluation of the merits of the GARF algorithm, including a comparison between GARF and (i) exact optimal solutions (in small networks) and cost lower bounds (in large networks), and (ii) other algorithms previously proposed for solving the MH-VTD-R/NR problems. Despite its simplicity, GARF has clearly shown its benefits in terms of the quality of the obtained solutions and scalability. Finally, the exhaustiveness of the tests included in this paper motivates and supports an analysis on the cost-efficiency of VTD reconfiguration as an adaptation mechanism for traffic changes in transparent optical networks.

The rest of the paper is organized as follows. Section II presents the state-of-the-art on multi-hour virtual topology design. In Section III we provide an optimal MILP formulation for the MH-VTD-R problem considering both transceivers and reconfiguration frequency. Section IV presents the GARF algorithm, whose merits are assessed and exhaustively evaluated in Section V. Finally, Section VI concludes the paper.

II. Related Work

Dynamic non-hierarchical routing (DNHR) [11] for telephone networks, which can be considered the first case of success of multi-hour network design, appeared in the 1980s. After this, network planning taking into account time-varying traffic demands has been researched for multiple network technologies [8],[9],[11]-[26] (see [12] for a comprehensive historical survey in the topic).

The first investigations of MH planning in optical networks were targeted towards the design of virtual topologies in multi-hop networks, based on passive stars [13]. In the last decade, the interest of the optical community has shifted to lightpath-based transparent optical networks. Initial works on virtual topology reconfiguration in lightpath-based networks focus on studying when and how to reconfigure an existing VTD, as a reaction to a variation in the traffic [8],[14]-[17]. The most common design objective considered is minimizing the associated reconfiguration to avoid traffic disruptions. Additional objective criteria include the minimization of (i) the maximum load in the lightpaths [8],[14], (ii) the average virtual hops in the network [15], or (iii) the average propagation delay [16]. A full multilayer approach is given in [17], where VTD reconfiguration is considered in conjunction with RWA reconfiguration. Note that all these approaches consider a one-time adaptation of a given virtual topology to a known change in traffic, and not periodic (multi-hour) traffic trends.

Incorporating the periodic nature of traffic in transparent optical networks planning has only more recently attracted the interest of the research community [9],[18]-[26]. Most research efforts have been focused on planning in the lower layer, i.e. the RWA of a given sequence of virtual topology designs corresponding to MH traffic. This is based on the Scheduled Lightpath Demand (SLD) model proposed in [18] where the evolution of individual lightpaths is known in advance. The planning problem then consists of finding a set of valid RWA solutions for the input VTDs, optimizing several network performances, such as the number of wavelengths used in the highest loaded fiber link. In [18], the authors proposed branch and bound and tabu search approaches, while in [19] a more scalable tabu-search based heuristic and greedy algorithms are presented. [20] and [21] 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.

While the aforementioned approaches deal with a known evolution of virtual topologies, we are interested in the problem problem of determining this evolution from the MH periodic traffic itself (i.e., the MH-VTD-R problem) to plan the upper layer. We also consider the problem of determining a single static virtual topology which can accommodate the traffic at any time (i.e., the MH-VTD-NR problem). To the best of our knowledge, the only works previous to our efforts in the literature on the MH-VTD problem variants are those in [22] and [23]. In [22], the authors propose a method, called Joint Configuration with Exact Traffic (JCET), based on a MILP formulation with a decomposition heuristic, for solving the MH-VTD-NR problem. The JCET method is compared with a simple MH-VTD-NR approach, denoted as Unique Configuration with Maximal Traffic (UCMT). In UCMT, the temporal sequence of traffic matrices is reduced to a single unique matrix where the traffic between node pairs is the maximum along time. The minimum-cost VT is then obtained from the maximal traffic matrix using the proposed static MILP formulation. Additionally, the solutions in [22] are also compared with a naive approach (Independent Configuration with Exact Traffic, ICET) for the MH-VTD-R problem, where the VTDs are independently planned in each time interval. The proposed MILP formulation is again used to solve a single traffic matrix problem for each time slot. The results comparing the MH-VTD-NR/R approaches indicate that establishing a non-reconfigurable virtual topology with fixed routing obtains solutions very close to those in which the virtual topology can be dynamically reconfigured over time. In [23] the authors address only the MH-VTD-NR problem, comparing the benefits of a variable and a static traffic routing on top of the (static) VTD, with several objective functions.

In the last years, the authors of this paper have intensively been investigating the different MH-VTD problem variants contributing with several works [9],[24]-[26]. In [24], we present a set of MILP formulations for several MH-VTD-NR and MH-VTD-R problem variants which minimize the number of transceivers. These formulations are extended in this paper to include penalization of VTD reconfiguration. In [25], a set of Tabu Search (TS)-based heuristic algorithms solving the previous problems is presented. These tabu search algorithms iteratively solve smaller single-time slot MILP formulations with constraints on transceivers for independent time intervals in order to jump between neighboring solutions, and thus explore the solution space in a directed manner. Depending on the problem variant, several fitness functions are proposed to assess the quality of neighboring solutions found iteratively.

In [9], we focus on the comparison of a set of variants of the static MH-VTD-NR problem, with and without flow routing reconfiguration. We propose a set of MILP formulations of the problem, together with a 3-step algorithm making use of the concept of traffic domination [27] for problem simplification. In this 3-step algorithm, as in the UCMT method, the sequence of matrices is reduced to unique single unique matrix. Here, however, traffic domination is employed to reduce the matrix sequence to one which mathematically dominates all of them.

Finally, in [26] we present two algorithms for the MH-VTD-R problem. The first one modifies the fitness function of the TS algorithms presented in [25] to include penalization of lightpath reconfiguration. The second algorithm is based on a classical heuristic method which iteratively solves the dual problem of the MH-VTD-R formulation built by Lagrangean relaxation of the primal problem. Once built, the dual problem is solved via a standard iterative subgradient optimization scheme.

The GARF algorithm presented in this paper strongly improves the performance of the previous proposals. One of the main advantages of GARF is its flexibility, enabling us to gradually tune the degree of reconfiguration allowed in the MH-VTD problem. In this sense, it is valid for both the MH-VTD-NR and MH-VTD-R variants. Despite its simplicity and its wider application, the GARF algorithm is more scalable and obtains solutions with a similar or better transceiver cost than previous algorithms, while strongly reducing the number of reconfigurations in the network for the reconfigurable case.

III. MH-VTD-R Problem Formulation

Let N be the number of nodes in the network, and T be the number of time intervals for which the traffic is defined. Let $i,j,s,d,n = \{1\ldots N\}$ be the indices for the nodes, and $t = \{1\ldots T\}$ be the index for the time intervals (since we are dealing with periodic traffic, we assume that the last time interval $t=T$ is followed by the first time interval $t=1$). Let M_t denote the traffic matrix at time slot t, and $M_t(i,j)$ 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. The cost of each transmitter and receiver is considered equal, and is represented by C_{TR}. An artificial cost of reconfiguring a lightpath is denoted as C_R.

The decision variables to the problem are:

- $p(i,j,t)=\{0,1,2,...\}$. Number of lightpaths from node i to node j, required during time interval t.
- $f(i,j,s,d,t)=[0,1]$. Fraction of the total traffic demand from node s to node d that is routed on the existing lightpaths from node i to node j, during time interval t.
- $T(i)=\{0,1,2,...\}$. Number of transmitters installed in node i.
- $R(i)=\{0,1,2,...\}$. Number of receivers installed in node i.
- $r^+(i,j,t)=\{0,1,2,...\}$. Number of new lightpaths set up at time t with respect to the number of existing lightpaths at time t-1 (or time T if t=1) between the nodes (i,j).
- $r^-(i,j,t)=\{0,1,2,...\}$. Number of lightpaths torn down at time t with respect to the number of existing lightpaths at time t-1 (or time T if t=1) between the nodes (i,j).

The problem formulation is given by (1).

$$\min c_{TR} \sum_{i=\{1...N\}} \big(T(i)+R(i)\big) + c_R \sum_{\substack{i,j=\{1...N\}\\ t=\{1...T\}}} r^+(i,j,t) \qquad (1a)$$

subject to:

$$\sum_{s,d=\{1...N\}} \{M_t(i,j)\cdot f(i,j,s,d,t)\} \le C\cdot p(i,j,t), \qquad (1b)$$
$$i,j=\{1,...,N\},\ t=\{1,...,T\}$$

$$\sum_{j=1}^{N} f(n,j,s,d,t) - \sum_{i=1}^{N} f(i,n,s,d,t) = \begin{cases} 1, if\ n=s \\ -1, if\ n=d, \\ 0\ otherwise \end{cases} \qquad (1c)$$
$$n,s,d=\{1,...,N\},\ t=\{1,...,T\}$$

$$T(n) \ge \sum_{j=1...N} p(n,j,t),\ n=\{1,...,N\},\ t=\{1,...,T\} \qquad (1d)$$

$$R(n) \ge \sum_{i=1...N} p(i,n,t),\ n=\{1,...,N\},\ t=\{1,...,T\} \qquad (1e)$$

$$p(i,j,t)-p(i,j,t-1) = r^+(i,j,t) - r^-(i,j,t), \qquad (1f)$$
$$i,j=\{1,...,N\},\ t=\{2...,T\}$$

$$p(i,j,1)-p(i,j,T) = r^+(i,j,1)-r^-(i,j,1),\ i,j=\{1,...,N\} \qquad (1g)$$

The objective function (1a) minimizes the total cost of the transmitters and receivers (C_{TR}), and the artificial lightpath reconfiguration cost (C_R). Recall that every reconfiguration of lightpaths causes expensive service disruption, since high data rates are involved. Consequently, we penalize such reconfigurations even though no extra cost is incurred by new equipment. This reconfiguration cost can be tuned to a desired value, reflecting its importance to the network planner. Note that only variable r^+ is used to compute the number of reconfigurations in (1a). Naturally, given the periodic nature of the planning, for every node pair, the sum along time of new lightpath establishments (r^+) is equal to the sum along time of lightpath disestablishments (r^-) in a full multi-hour period.

$$\sum_{t=1...T} r^+(i,j,t) = \sum_{t=1...T} r^-(i,j,t),\quad i,j=\{1,...,N\} \qquad (2)$$

Thus, the reconfiguration cost C_R represents the sum of two reconfiguration costs, i.e. the cost of setting up and then tearing down a lightpath. This rationale does not hold if these reconfiguration costs are non-linearly related. In that case, they should be considered separately in (1a). However, such scenarios are out of the scope of this paper.

Constraints (1b) force the traffic between two nodes at any time, to be limited by the number of lightpaths planned between those nodes. Constraints (1c) are the link-flow conservation constraints. Constraints (1d) and (1e) ensure that the number of lightpaths originating (terminating) at a given node at any time, must be below the number of transmitters (receivers) installed at that node. Constraints (1f) and (1g) are used to force the $r^+(i,j,t)$ and $r^-(i,j,t)$ variables to account for an absolute increase or decrease, respectively, in the number of lightpaths between nodes i,j, at time t.

IV. Greedy Approach With Reconfiguration Flattening: The GARF Algorithm

Since the MH-VTD-R/NR problems are highly complex, solving the exact MILP is only feasible for smaller problem sizes. Herein we propose a fast and efficient heuristic approach to address both the MH-VTD-NR and MH-VTD-R problems. We refer to this algorithm as GARF (Greedy Approach with Reconfiguration Flattening). GARF is divided into three steps. In *Step 1*, a fast greedy approach is run which obtains a good feasible solution with respect to transceivers, without caring for the frequency of reconfiguration. *Step 2* applies a tabu search approach targeted to further reduce the transceiver cost and, finally, *Step 3* is aimed at reducing reconfiguration frequency according to a tunable factor which controls its tradeoff with transceiver cost. A flow chart illustrating the individual steps of GARF is given in Fig. 3.

A. The GARF algorithm: Step 1 - Greedy Approach

In this step, a greedy algorithm creates an initial solution with the objective to minimize the number of transceivers. First, as a preprocessing step, the lower bound shown in (3) and proposed in [9] is calculated. This is a bound on the number of transmitters (receivers) at a node based on the minimal number of lightpaths required to add (drop) the total traffic generated by (targeted to) the node in the maximal time slot.

$$LB_{TX}(n) = \max_{t=1...T} \left\{ \sum_{j=1}^{N} \left\lceil \frac{M_t(n,j)}{C} \right\rceil \right\} \qquad (3a)$$

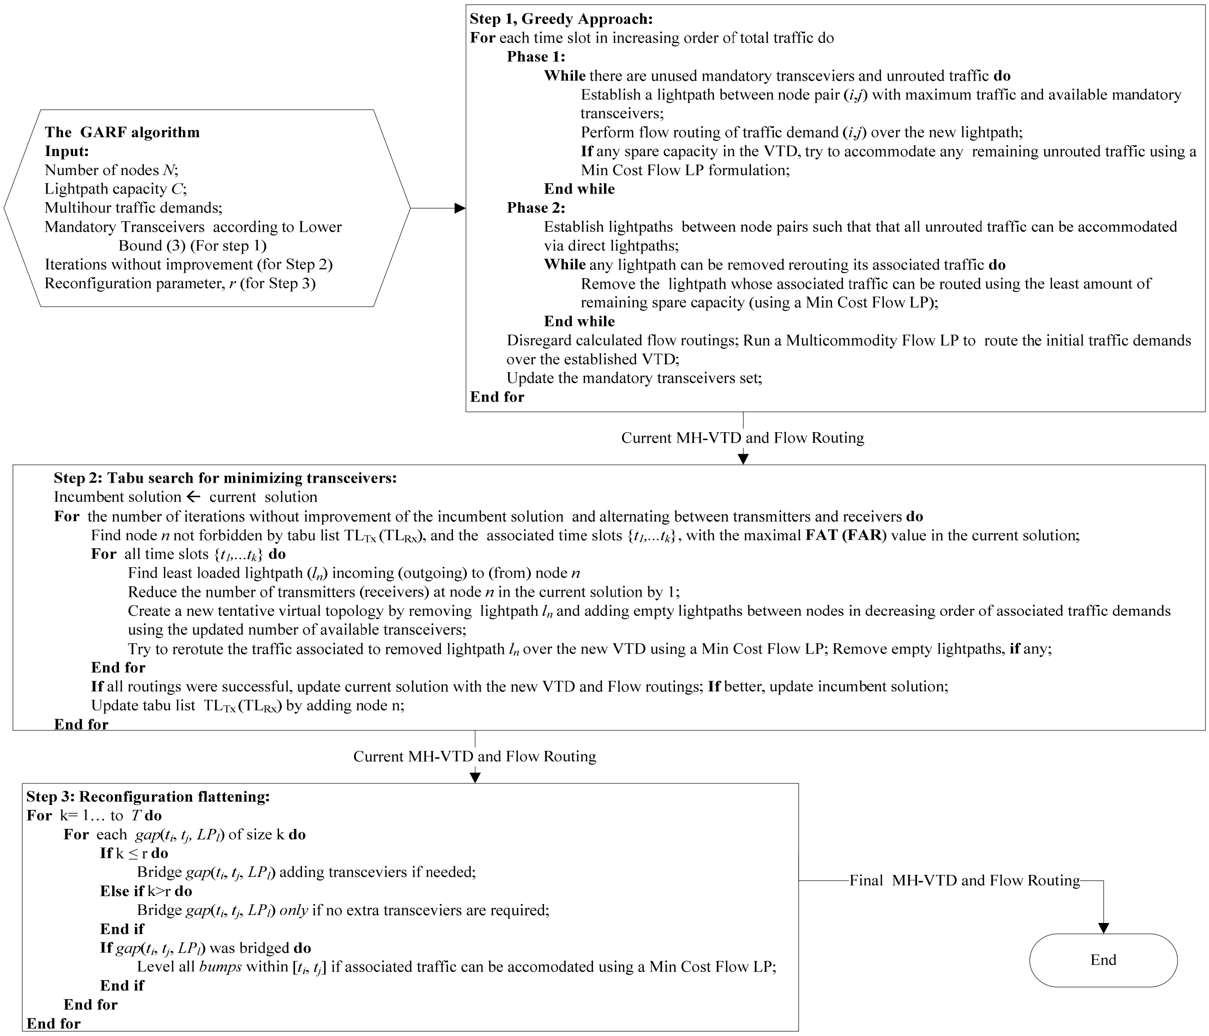

Fig. 3. A flow chart of the GARF algorithm

$$LB_{RX}(n) = \max_{t=1...T} \left\{ \sum_{i=1}^{N} \left\lceil \frac{M_t(i,n)}{C} \right\rceil \right\} \tag{3b}$$

The basic idea is as follows. Since the calculated transceivers will surely be needed at some point in time, they are the ones we want to use up first when planning each time slot. We call the set of these transceivers as *mandatory transceivers.*

For each time slot, in increasing order of its total traffic, two phases of the algorithm are performed. In phase 1, lightpaths are iteratively established using only transmitters/receivers from the *mandatory transceivers* set in decreasing order of their associated traffic demands, until the set is exhausted. Each time a new lightpath is established, as much of the associated direct traffic as possible is routed over the new lightpath. If this traffic demand exceeds the lightpath capacity, the excess traffic remains unrouted. Then, all remaining unrouted traffic demands between all node pairs are sorted in decreasing order and sequentially fed to a linear program (LP) formulation which models the Minimal Cost Flow Problem (Min Cost Flow Problem) [28]. This formulation is used to route each traffic demand over the available spare capacity. Where the LP is successful, the spare capacity is updated before the next demand is processed.

After the mandatory transceivers are exhausted in phase 1, the algorithm moves to the second phase which uses a modified version of the greedy algorithm proposed in [5]. Phase 2 begins by establishing as many lightpaths as necessary between node pairs such that all the remaining unrouted traffic can be routed via direct lightpaths. The algorithm then tries to iteratively remove the lightpath whose traffic can be accommodated by using the least amount of the remaining spare capacity. The algorithm proceeds for as long as such a lightpath removal is possible between any pair of nodes. After Phase 2 terminates, the currently calculated flow routings are disregarded and a Multicommodity flow LP formulation is run to re-route the original traffic demands in the considered time slot over the newly established virtual topology. The *mandatory transceivers* set is then updated to include all the transceivers associated with the new virtual topology to be used in subsequent time slots. This forces the algorithm to first use the transceivers that we have already decided to 'buy', before adding new ones. The algorithm terminates after all time slots are processed.

B. The GARF algorithm: Step 2 – Tabu Search for Minimizing Transceivers

In this step, we employ a tabu search to improve the solution generated in *Step 1*. In general, tabu search is an iterative meta-heuristic which guides the search through the solution space using a memory structure, called a tabu list, to avoid cycling between neighboring solutions around local optima. It does so by 'memorizing' a certain number of the most recently visited solutions, or some of their attributes, prohibiting the search to reconsider them for as long as they remain in the list.

First we introduce some concepts and notations required for a better understanding of Step 2 of the algorithm:

- *Active Transmitters matrix* (**AT**): an $N \times T$ matrix defined for a given solution to the *MH-VTD* problem, where a element $\mathbf{AT}(n,t)$ ={0, 1, 2,…} represents the number of transmitters that are active at node n in time slot t.
- *Fluctuation of Active Transmitters matrix* (**FAT**): an $N \times T$ matrix obtained from matrix **AT** by subtracting from each element in **AT**, the value of the *minimal* element in its row except itself. In other words:

$$\mathbf{FAT}(n_i, t_i) = \mathbf{AT}(n_i, t_i) - \min_{t | t \neq t_i}(\mathbf{AT}(n_i, t)) \qquad (4)$$

An example **AT,** and corresponding **FAT** matrix, are shown in Fig. 4 for a potential solution in a 3 node network with 4 time slots, i.e. N=3, T=4.

$$\mathbf{AT} = \begin{pmatrix} 3 & 1 & 1 & 2 \\ 2 & 2 & 1 & 1 \\ 1 & 4 & 2 & 4 \end{pmatrix} \quad \mathbf{FAT} = \begin{pmatrix} 2 & 0 & 0 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 3 & 1 & 3 \end{pmatrix}$$

Fig. 4. An example of an **AT** and corresponding **FAT** matrix.

AT(1,4)=2, for example, indicates that in the fourth time slot there are 2 transmitters active at node 1. The positive elements in the **FAT** matrix indicate the number of transmitters that are used in each time slot above the minimal value used in the least active time slot. The amplitude or maximal fluctuation of activity per node is the maximal element in each row of **FAT**.

Analogously for receivers, we define the *Active Receivers* (**AR**) matrix and the *Fluctuation of Active Receivers* (**FAR**) matrix.

Step 2 of GARF runs as follows. After obtaining a feasible solution in *Step 1*, tabu search iterations are run alternating between transmitters and receivers until the maximal number of iterations without improvement of the best solution is met. Upon termination, the incumbent (i.e., best visited) solution is deemed the final result. During these iterations, two tabu lists are maintained, realized as FIFO (First In First Out) queues of finite size corresponding to either transmitters (tabu list TL_{TX}) or receivers (tabu list TL_{RX}).

In each iteration, the **FAT (FAR)** matrix is computed to show the fluctuation of active transmitters (receivers) for that solution. A node n with the maximal amplitude of fluctuation, i.e. a row containing the maximal element in **FAT (FAR),** which is not forbidden by tabu list TL_{TX} (TL_{RX}) is chosen. In our example, this would correspond to node 3, where for time slots 2 and 4 the corresponding value **FAT**(3,2) **=** **FAT**(3,4) = 3 is the maximal value in the **FAT** matrix. For the chosen node n, we then attempt to remove one transmitter (receiver) from the current solution. This implies checking whether a lightpath initiating (terminating) at node n could be removed in each of the time slots t_n where $\mathbf{FAT}(n, t_n)$ is maximal (in our example, time slots 2 and 4).

The check associated with each time slot t_n consists of three steps. First we choose the lightpath l_n originating (terminating) at node n which is least loaded during that time slot. Second, we construct the so-called *tentative virtual topology* and, finally, we attempt to reallocate the traffic carried by lightpath l_n onto the *tentative VT* for time slot t_n. The *tentative VT* is constructed from the original one by (i) removing the aforementioned least loaded lightpath l_n and (ii) adding as many new empty lightpaths as the available transceivers in the network permit. Available transceivers assume all those which are allocated to the nodes over time but not used by the current VT in the considered time slot, with the number of transmitters (receivers) at node n reduced by 1. The new empty lightpaths are established between node pairs with available transceivers, in decreasing order of their original traffic demands. This traffic is not routed over the new lightpaths, i.e. the lightpaths remain empty, but is only used as a sorting criterion. After adding a lightpath to the *tentative VT*, the value of the corresponding traffic is decreased by the value of the lightpath capacity and the traffic demands are re-sorted.

After constructing a *tentative VT*, the check procedure attempts to reallocate the traffic associated to the removed lightpath on the spare capacities of the *tentative VT*, solving a Min Cost Flow problem. If this rerouting of the traffic in the least loaded lightpath is successful for all the time slots t_n, all empty left-over lightpaths are removed and the new set of virtual topologies and corresponding flow routings become the new current solution in the next iteration. Otherwise, the old current solution remains unchanged and the algorithm proceeds to the next iteration. If the solution in the current iteration is better than the incumbent solution, it is updated accordingly. In any case, tabu list TL_{TX} (TL_{RX}) is updated to include the considered node n.

C. The GARF algorithm: Step 3 – Reconfiguration flattening

Step 3 of GARF is aimed at minimizing the frequency of reconfiguration of the lightpaths obtained in *Step 2.* A tunable parameter controls the penalization of reconfiguration, allowing the solution to range from the static case (MH-VTD-NR) to various levels of reconfiguration for the MH-VTD-R problem variant.

First, we introduce some notation and basic concepts for better understanding of Step 3 of the algorithm:

- *Unique lightpath identifier* ($LP_{i,}$). A feasible solution to the MH-VTD-R (or MH-VTD-NR) problem consists of a set of lightpaths which are active at a certain interval(s) in time (for the MH-VTD-NR all established lightpaths are always active). We assign to each lightpath a unique identifier in the form of LP_i, i={1, 2, 3, …}. Note: we differentiate between multiple lightpaths established between the same pairs of nodes, but the same unique lightpath can be established and torn down over multiple time intervals.
- *Lightpath activity function* (*LAF*). Each unique lightpath corresponds to one lightpath activity function (LAF) which indicates its activity over time, i.e. is a graphic representation of its schedule. Each transition in a LAF indicates one reconfiguration: A drop in the LAF from 1 to 0 indicates that the associated lightpath was torn down, while a jump from 0 to 1 indicates that the lightpath was established.
- *gap*(t_i, t_j, LP_l). We call a gap any part of a LAF which starts with a 1 to 0 transition and ends with a 0 to 1 transition, while remaining at 0 in-between. We denote a specific gap as *gap*(t_i, t_j, LP_l), where t_i and t_j denote the first and last time slots of the gap and LP_l denotes the associated lightpath LAF. We define the size of a gap be the number of time slots it covers.
- *bump*(t_i, t_j, LP_l). A bump refers any part of a LAF starting with a 0 to 1 transition and ending with a 1 to 0 transition, while remaining at 1 in-between. Its notation and size is analogous to that of *gaps.*

Fig. 5 shows an example of a simple MH-VTD-R solution for a 3 node network with 7 time slots. Fig. 5(a) shows the set of unique lightpath identifiers while Fig. 5(b) indicates at which times the individual lightpaths are active (i.e. their schedule). Note that the time slots are periodic, i.e., after time slot 7, time slot 0 is active again. The LAFs corresponding to each unique lightpath are shown in Fig. 5(c), along with some examples of bumps and gaps and their corresponding sizes. Also note that each transition in a LAF indicates a reconfiguration, incurring service disruption and signaling overhead. Consequently, the flatter the LAFs, the better. Given an algorithm for the MH-VTD-R problem which does not consider the number of reconfigurations, it is possible that a lightpath may be torn down during a few time slots even though the corresponding transceivers are still available, simply because it is not needed to route traffic in those time slots. One of the basic ideas in *Step 3* of GARF is to eliminate such unnecessary reconfigurations.

Note that no matter its size, each gap and bump is associated with exactly two reconfigurations: one to tear down the lightpath and one to set it up (or vice versa). Intuitively, it seems reasonable to try to eliminate smaller gaps or bumps since they require the same number of reconfigurations as larger ones, but require planning a change with respect to flow routing and transceiver utilization in fewer time slots. 'Bridging' a gap implies letting an established lightpath remain active. 'Leveling' a bump implies trying to reroute traffic so that the associated lightpath can remain inactive over that time period.

The main objective of Step 3 is to maximally flatten (bridge/level) the LAF functions in order to reduce the total number of reconfigurations. Herein, we introduce a tunable parameter r={0, 1, ..., T-1}, which allows us to control the increase in the overall number of required transceivers as a tradeoff with minimizing reconfiguration frequency. GARF Step 3 runs as follows. For each time slot, all *gaps* of size r or less in the current solution are bridged, even if it is at the expense of additional transceivers. On the other hand, all *gaps* of size greater than r are bridged only if no additional

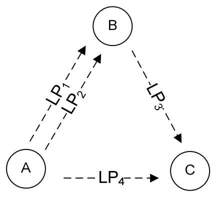

(a)

LP_l	Source	Dest.	Activity	No. of Reconfigs
LP_1	A	B	1:00-3:00; 5:00-6:00	4
LP_2	A	B	2:00-4:00; 5:00-100	4
LP_3	B	C	3:00 – 4:00	2
LP_4	A	C	0:00-0:00	0

(b)

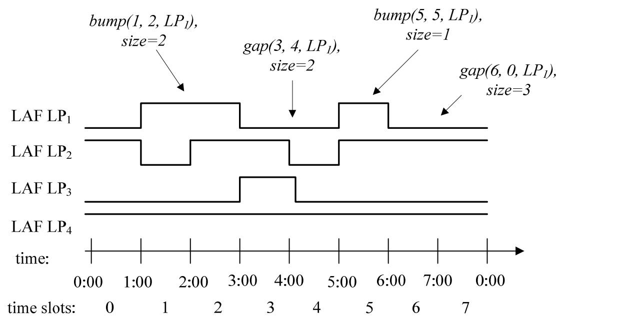

(c)

Fig. 5. An example a solution to the MH-VTD-R problem for a 3 node network: (a) the set of unique lightpath identifiers, (b) their schedule and (c) the corresponding lightpath activity functions.

transceivers are required. In the extreme case of r=0, minimizing transceiver cost is the primary objective, while only those *gaps* which correspond to inactivity of lightpaths due to lack of traffic and not lack of transceivers are bridged. For r=T, all LAFs will be completely flattened no matter the associated transceiver cost, eliminating reconfiguration all together and yielding a static virtual topology, i.e. a solution to the MH-VTD-NR problem. By tuning parameter r, we can tune this trade-off to a desired value.

During execution of Step 3, each time we bridge a gap of size s (with or without adding transceivers), we try to 'level' all bumps of size s or less within the time slots that have just been bridged. Leveling a bump is feasible only if the associated traffic can be routed over the remaining established lightpaths and the newly bridged one. This not only reduces additional reconfigurations, but in some cases reduces the total number of transceivers if an entire LAF is ultimately flattened to zero. After considering all gaps and bumps of size 1 to T for all time slots, the algorithm terminates.

V. RESULTS

This section presents the results of extensive tests conducted to assess the performance of the GARF algorithm, compared with those obtained by exact MILP formulations (in small networks), lower bounds on the network cost, and other heuristic algorithms. The algorithms are implemented in Matlab using the MatplanWDM tool [29] which links to the TOMLAB/CPLEX library [30] used to solve the MILP formulations.

A. Description of the testing scenarios

The merits of the GARF algorithm are evaluated for three testing scenarios. As a consequence of the assumption established in Section I where any VTD has a feasible RWA solution, the optimization problem is independent of the physical network topology, save for the number of nodes which can be extracted from any traffic matrix. Therefore, the multi-hour-demand, i.e. a sequence of traffic matrices, is the only input data to the planning problem. Under these assumptions, the first scenario considers synthetically generated multi-hour traffic for a small 6-node network. Five sequences of T = 12 traffic matrices were generated randomly following the MH traffic generation method used in [9]. Equations (4)-(5) describe this model:

$$M_t(i,j) = b(i,j) \cdot \text{activity}(t) \cdot rf(R), \quad \forall i,j,t. \tag{4}$$

According to the model, the traffic $M_t(i,j)$ between two nodes (i,j) at a given time t is calculated as the product of three factors. First, factor $b(i,j)$ gives the (i,j) coordinate of a base traffic matrix computed for the sequence as follows. 80% of the values in matrix b (chosen randomly) were set to one, while the remaining 20% were set to two. This is meant to capture the effect of non-uniformities in the generated traffic matrices. Secondly, activity factor *activity*(t) in equation (4) intends to capture the effect of traffic intensity periodic variation along the day. Our intensity variation scheme is described by equation (5), based on the intensity model presented in [31].

$$\text{activity}(t) = \begin{cases} 0.1 \text{ if } t \in [1,6] \\ 1 - 0.9 \cdot \left(\cos\left(\dfrac{\text{mod}(t,T) - 6}{18} \cdot \pi \right) \right)^{10} & \text{otherwise} \end{cases} \tag{5}$$
$$\text{where } t = 1,\ldots,T$$

Finally, factor $rf(R)$ in (4) is an independent random sample, uniformly distributed over interval [1-R, 1+R]. The purpose of the rf factor is to include a randomness effect in the traffic intensity. In our tests R was set to 0.2.

In the remaining two scenarios, data from real traffic traces from the 11-node Abilene and the 23-node GÉANT networks were used. This data, publicly available at [4], consists of traffic matrices spanning several weeks. For each of the two networks, the *average week* was considered and calculated in the following manner. All values of the trace taken at the same time in the same day were averaged to obtain a sequence of 24x7=168 matrices representing the hourly traffic variations along one full week.

All the sequences of traffic matrices in all three scenarios are normalized to three different traffic intensities or traffic loads ρ={0.1, 1, 10}. The 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 ρ=0.1 corresponds to the case when the highest average traffic equals 10% of a single lightpath capacity. On the contrary, a value of ρ=10 captures cases in which the average traffic between two nodes in the highest loaded time slot fills on average 10 full lightpaths.

In Fig. 6, the temporal evolution of the average week of the traffic pattern is shown in the curve denoted as "offered traffic" for the medium load case, ρ=1. This curve represents the total volume of offered traffic for each time slot in terms of the *number of lightpaths*, i.e., the sequence of the sum of all the elements in each matrix M_t, t = 1…T divided by the lightpath capacity, calculated according to ρ. Of course, for other ρ values, the curve would change solely in the magnitude of the number of lightpaths, without modifications in the temporal shape. We can see that the traffic in an *average week* periodically experiences daily peaks corresponding to busy working hours and daily valleys corresponding to night hours. The amplitudes between peaks and valleys are much larger during the five working days, diminishing in the weekend.

The sequences of traffic matrices go through an additional pre-processing step to reduce the complexity of the problem before being applied as an input to the planning algorithms. For each sequence, we make use of the concept of traffic domination [27] to filter out those matrices which are

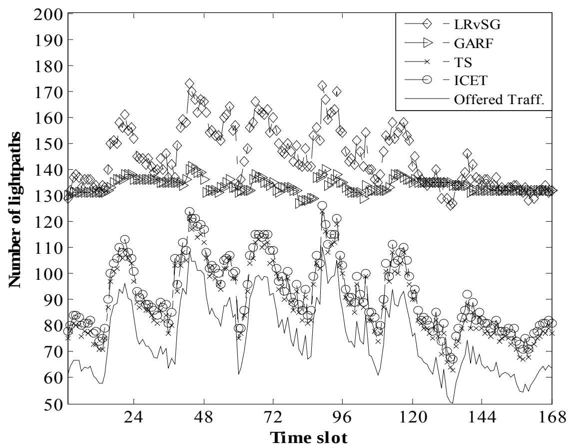

Fig. 6. Time evolution of the no. of lightpaths of each heuristic approach and the offered traffic for the average week of the Abilene nwk. with ρ=0.1

TABLE I
REDUCTION IN NUMBER OF MATRICES IN PREPROCESSING STEP

Problem	N = 6 T = 12	Abilene T = 168	GÉANT T = 168
MH-VTD-NR	5	52	84
MH-VTD-R	7.8	167	168

redundant to the planning problem, and thus can safely be eliminated. Informally speaking, if a traffic matrix M_1 dominates a traffic matrix M_2, every network which is able to carry traffic M_1, can also carry traffic M_2. Given a sequence of multihour traffic matrices M_t. we perform two types of traffic demand simplifications, depending on whether the sequence is to be applied to the non-reconfigurable or reconfigurable MH-VTD problem.

The traffic demand preprocessing applied in the non-reconfigurable case is the one described in [9]. The traffic matrices in the sequence which are dominated by *any other traffic matrix* in the sequence are filtered out, since they are redundant to the planning problem, i.e. every VTD which is able to carry the original traffic sequence is also able to carry the reduced sequence. In the reconfigurable case, a traffic matrix M_t is filtered out only if it is dominated by *the previous traffic matrix* M_{t-1} or by *the subsequent matrix* M_{t+1}. Since the matrix M_{t-1} (or M_{t+1}) dominates M_t, the virtual topology designed for M_{t-1} (or M_{t+1}) can be used without any reconfiguration in the next (or previous) time slot t. The reductions in the number of matrices obtained by these preprocessing steps are plotted in Table I. Note that the reductions are very significant in the static case (between 50% and 70%), but negligible in the Abilene and GÉANT networks for the reconfigurable case.

B. The MH-VTD-R problem

In this sub-section we focus on assessing the merits of the GARF algorithm for planning networks which permit VT reconfiguration with the main objective to minimize transceivers. Table II collects the results obtained for the MH-VTD-R problem using the following approaches:

- LB: Lower bounds on the number of transceivers calculated as shown in Section IV.A
- MILP: The exact solution of the MILP formulation for the MH-VTD-R problem presented in Section III.
- GARF: The GARF algorithm with reconfiguration parameter $r = 0$. In Step 2 of GARF, the number of iterations without improvement was set to 20 and the tabu list size to 4, 7 and 13 for the 6-node, Abilene and GÉANT networks, respectively, determined experimentally.
- TS: The reconfigurable Tabu Search heuristic from [26]. The number of iterations without improvement and the tabu list size were set to the same values as GARF Step 2 for the network scenarios it could solve.
- LRvSG: The Lagrangean Relaxation-based approach from [26]. The total number of iterations was set to 50, initial and minimal values of step-size parameter p were to 2 and 0.005, respectively; and the maximum number of iterations without improvement was set to 10.
- ICET: The Independent Configuration with Exact Traffic (ICET) strategy from [22] which minimizes only transceivers for each time slot separately. The algorithm in [5] is used to minimize transceivers for a single traffic matrix.
- Gap%: This column shows a bound on the GARF sub-optimality gap, i.e., the relative difference between the GARF solution cost, and the cost of the MILP solution (6-node network) or the lower bound cost (Abilene and GÉANT networks).

The MILP, GARF, TS, and LGvSG approaches all minimize both the number of transceivers and the reconfiguration frequency in their objective functions. Here, we focus on the case where the number of transceivers is the main objective to minimize, while the reconfiguration frequency is a secondary objective. Consequently, parameter r in GARF is set to r=0, while in the MILP, TS and LRvSG algorithms, the cost of reconfiguring a lightpath is set to a sufficiently small fraction ($\sim 10^{-7}$) of the transceiver cost C_{TR} to ensure that the transceivers completely dominate the optimization. The included reconfiguration cost simply prevents unnecessary reconfigurations. In the ICET algorithm, the objective function minimizes only the number of transceivers, not considering the number of reconfigurations.

Each cell in Table II shows the number of transceivers of each solution. In parenthesis, we include the number of reconfigurations averaged *per time slot* associated to the plan. This is an intuitive measure of the frequency of the reconfigurations, and thus the relative amount of extra signaling and traffic disruption caused by the VT changes. The heuristic solutions with the best results for each objective criterion are printed in bold. The cells marked with "- (-)" represent cases where the associated approach was not able to find a solution (MILP for Abilene and GÉANT) or even an initial solution (TS for GÉANT) in 3 days of computational time. Note that the TS algorithm contains single-slot MILP formulations which are less complex than the full multi-hour MILP formulation, but still pose severe scalability issues.

Results reveal that GARF and TS are the best algorithms in terms of minimization of the number of transceivers for the 6-node and Abilene networks. The difference between them is not significant in most cases, except for ρ=0.1 in the synthetic traffic scenario where TS showed 25% better results. However, TS fails with respect to reconfiguration performance. It requires a much higher amount of lightpath reconfiguration than GARF in all cases: specifically, {9%, 86%} more in the low load cases, and between 4 and 20 times more in the rest of the scenarios solved by TS. Note that for the GÉANT network, TS could not produce a solution in reasonable time making it unusable for larger network instances.

The high TS performance for low loads where the problem is "more integral", can be explained in the more thorough search of the solution space carried out by the algorithm, which: (i) considers several neighboring solutions in each iteration, and (ii) employs a single-time slot exact MILP formulation to check these neighbors. Conversely, the GARF is is a greedy approach whose tabu search performed by the Step 2 is less intense since (i) it considers only one neighbor in each iteration, and (ii) uses a simpler LP Min Cost Flow formulation to check it.

With respect to reconfiguration frequency, GARF outperforms all the algorithms in all cases except the Abilene network test with low load where LRvSG is best. Nevertheless, this should not be seen as a true merit of LRvSG since the large number of extra transceivers planned by the LRvSG algorithm in that case (69% more) permits solutions with fewer reconfigurations. Generally, TS, LRvSG and ICET all require mutually similar reconfiguration frequencies which are much higher (between 7% and 170% for low loads in 6-node and Abilene networks cases, and between 3 and 21 times more for the remaining cases) than GARF in all cases but the one mentioned above. For ICET this is expected as this approach does not consider reconfiguration frequency. The fact that the reconfiguration results of TS and LRvSG (which explicitly consider reconfiguration minimization in their objective functions) are very similar to those of ICET (which one ignores reconfiguration) suggests that TS and LRvSG do not effectively minimize reconfiguration frequency as their secondary objective.

The Gap% row evaluates the quality of the solutions of the GARF algorithm against the optimal MILP solution (in the 6-node network) and the lower bounds (in the Abilene and GÉANT networks). The optimality gap results are fairly good: below 1% in the high load scenarios, and below 6% in medium load scenarios. The sub-optimality gap of the GARF algorithm for the low load case, on the other hand, is in the order of 30-40% for the worst cases. However, it is important to note that this gap is based on the lower bound for the Abilene and GÉANT networks. Considering there are significant disparities between the optimal MILP solution and the lower bound for the 6-node network at load ρ=0.1, for the Abilene and GÉANT networks at low load, it cannot be concluded whether the lower bounds are weak or the solutions obtained by GARF are far from optimal. In any case, GARF gives the best solutions among the heuristics for these scenarios.

For further comparison of the performance of the algorithms, in Fig. 6 we show the temporal evolution of the *number of lightpaths* of the heuristic solutions for the Abilene scenario at medium load (ρ=1). It is clear that the LRvSG, ICET and TS designs follow the changes of the offered traffic pattern, implying frequent reconfigurations, while the GARF design is practically flat, i.e. the number of lightpaths remains almost constant with only small peaks covering the highest traffic points. Note that the *number of lightpaths active* over time in GARF is higher than in the solutions obtained by the TS and ICET approaches, and yet the *number of transceivers needed* is similar. This strongly supports the efficiency of the temporal transceiver reuse achieved by GARF which is able to maintain a larger number of lightpaths over time using the same equipment as the other approaches. This is mainly realized in Step 3 of GARF where lightpaths are maintained in gaps where they are not strictly necessary (“bridging gaps”), but only between nodes with already installed transceivers, supporting both the primary objective of transceiver minimization and the secondary objective of reducing reconfiguration frequency.

To achieve further insight into the performance of GARF, we observe the joint evolution of the number of transceivers and the frequency of the reconfigurations in the GARF designs. Namely, there exists a trade-off between the increase in the number of transceivers and the obtained decrease in the number of reconfigurations, depending on parameter r. To illustrate this relation, we ran the GARF algorithm with parameter r ranging from 0 to T for the Abilene scenarios. Figs. 7(a), 7(b) and 7(c) collect the results for loads ρ=0.1, ρ=1 and ρ=10, respectively.

We can see that at medium and especially high loads, the number of transceivers slowly, and almost linearly, increases with parameter r. The drop in reconfiguration frequency, on the other hand, is much steeper for small values of r. This suggests that most of the gaps/bumps are of small size, and are removed with the addition of a few transceivers. Consequently, setting r to a small value for medium and high loads can create an advantageous trade-off between transceivers and reconfiguration.

Finally, at low loads, an anomaly is observed where the number of transceivers initially grows with r as expected, but is smoothly reduced after a certain r. This anomaly is presumably due to the existence of some large gaps, whose “bridging” allows us to permanently tear down some lightpaths along time, reducing the number of transceivers. Namely, at low loads, a gap “bridged” provides bigger spare capacities to re-accommodate traffic from candidate bumps to “flatten”.

C. The MH-VTD-NR problem

In this sub-section we are interested on assessing the merits of the GARF algorithm for planning those networks in which the VT is constrained to be static along time. We compare the designs obtained by GARF with those obtained by other relevant algorithms previously proposed for the MH-VTD-NR problem. For all the tests, the transceiver cost C_{TR} is fixed to 1, while reconfiguration cost is infinite since reconfiguration is prohibited. The parameters of GARF, TS and LRvSG are the same as those used in Subsection V.B., except those explicitly stated below. Table III collects the obtained results where the columns correspond to:

- LB: Equal to that of the MH-VTD-R case.
- MILP: Exact solution to the MILP formulation to the MH-VTD-NR problem presented in [9].
- GARF: The GARF algorithm with parameter $r = T$
- TS: The static version of the Tabu Search (TS) algorithm from [25].
- LRvSG: The Lagrangean Relaxation-based algorithm, following the same technique as in [32], applied to the MILP formulation presented in [9] for the MH-VTD-NR problem.
- 3-Step-H: The three-step heuristic approach from [9].
- UCMT: The UCMT strategy from [22], using algorithm in [5] to minimize transceivers for the maximal traffic matrix.

- Gap%: A bound on the GARF sub-optimality gap as, in Table II.

As before, the best heuristic solutions are printed in bold, and the cells corresponding to tests whose initial solution running times exceed 3 days are dashed. We can see that GARF and 3-Step-H generally provide the best heuristic solutions, where GARF gives the best solution or a solution within 2% of the best in all medium and high load cases. These results certify the efficiency of GARF for these cases considering that its closest competitor (3-Step-H) calls many exact LP formulations to reduce the multihour sequence of matrices to a single "dominating" matrix.

For low loads, the UCMT and 3-Step-H approaches outperform GARF, where GARF is 23%, 17%, and 5% worse for the 6-node, Abilene, and GÉANT networks, respectively. This can be explained by Step 3 of the GARF which sets up more lightpaths than necessary to convert the reconfigurable design from Step 2 into the final non-reconfigurable one. Note that the UCMT and 3-Step-H approaches are both based on reducing the multihour traffic to a single traffic matrix (the maximal or "dominating" one, respectively). For low loads, were the lightpath capacity is in the order of ten times larger than the traffic demands, their success seems logical as smaller differences in the traffic matrices are not relevant.

The solutions obtained by LRvSG and TS are significantly worse in all cases, even though the approaches are analogous to their reconfigurable variants indicating that their heuristic philosophies are not efficient for the MH-VTD-NR problem.

The GARF sub-optimality gap for the MH-VTD-NR problem is in the order of 1% for the 6-node network in medium and high load cases, being the best among the heuristic algorithms. For the Abilene and GÉANT networks, the bound on the sub-optimality gap is between 8% and 11% for ρ=1 and, between 3% and 5% for ρ=10. The true gaps should be presumably better since the lower bounds used are the same as those calculated for the reconfigurable case, i.e., they may be somewhat weaker for the static problem. The accuracy of GARF in the low load experiments cannot be easily assessed given the presumable weakness of the lower bound for those cases, as indicated in the previous subsection for the MH-VTD-R problem.

D. Trade-offs between transceiver and reconfiguration costs

In this subsection we study the tradeoffs associated with tuning the mutual relationship of the transceiver cost C_{TR} and the reconfiguration cost C_R. GARF, TS and LRvSG are used to perform the study, as they are the unique heuristic proposals which consider both objectives. In contrast to the previous subsections, where either transceiver cost (Subsection V.B) or reconfiguration cost (Subsection V.C) completely dominated the optimization, we consider various combinations of the multi-objective problem driving to intermediate solutions. We fix the transceiver cost C_{TR} to 1, while varying the reconfiguration cost C_R from a very small fraction (10^{-6}) of the transceiver cost to a value of 0.3. C_{TR} was increased logarithmically from 10^{-6} to 0.01, and then linearly in increments of 0.05 from values 0.05 to 0.3.

Given a solution to the MH-VTD problem and a cost pair

TABLE IV
MAXIMAL COST SAVINGS – ABILENE, N=11, T=168

		Reconfiguration Cost (C_R)				
ρ		**0.001**	**0.01**	**0.1**	**0.2**	**0.3**
0.1	**GARF**	**64.73**	**68**	**68**	**68**	**68**
	TS	65.01	73.62	151.2	216	216
	LRvSG	108.40	111.98	147.8	168	168
1	**GARF**	301.16	**302.61**	**310.6**	**312**	**312**
	TS	**291.94**	309.4	406	406	406
	LRvSG	396.68	402.76	430	430	430
10	**GARF**	**2671.44**	**2675.36**	**2714.6**	**2746.8**	**2764**
	TS	2678.04	2759.37	3150	3150	3150
	LRvSG	2778.59	2855.9	3240	3240	3240

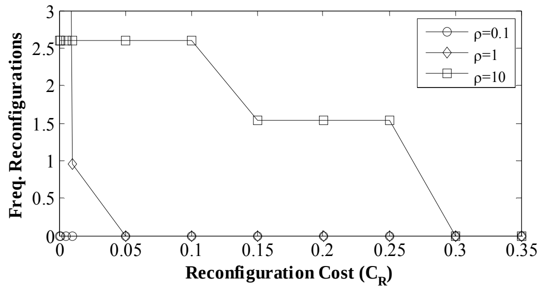

Fig 8. Frequency of Reconfigurations vs. C_R. Abilene Network.

(C_{TR}, C_R), the cost of the solution can be calculated according to the objective function of MILP formulation (1). Note that the tradeoff between transceivers and reconfigurations in GARF is controlled by parameter *r,* and not the values of C_{TR} and C_R. Consequently, we calculate the cost of all the solutions obtained by GARF in the previous experiments (In Tables II and III, and Fig. 7) using each cost pair (C_{TR}, C_R) and take the best result for each cost pair. For the TS and LRvSG algorithms we run additional experiments explicitly using the above mentioned values of C_R and C_{TR}. However, for some cost pairs, the fully dynamic or static solutions obtained by TS and LRvSG in Tables II and III may be better. Therefore, to make a fair comparison with GARF, we choose their best solution for each pair (C_{TR}, C_R) among Tables II and III and the newly run experiments.

The associated costs for the Abilene network are shown in Table IV with the best results in bold, while Fig. 8 plots the reconfiguration frequencies of the best solution for all tested C_R values. Table IV shows only an illustrative subset of the C_R values due to lack of space. We can see that GARF found the best solution for all cases, except for the smallest reconfiguration costs where transceivers dominate the optimization. This table indicates the superiority of GARF to find reconfigurable VTs where tradeoffs between transceiver and reconfiguration reduction are present.

Interestingly enough, from Fig. 8 we can see that even very small reconfiguration costs dominate the optimization. Values of the reconfiguration cost of 5% of the transceiver cost already drive the low and medium load cases to create static solutions. For high loads, this occurs at 30% of the transceiver cost.

E. Scalability

This section comments on the scalability properties of the approaches according to tests run on an Intel Core2 Duo CPU P8400 2.26 GHz processor. Table V gives the execution times for the Abilene and GÉANT networks with the fastest algorithm's times in bold. For the small 6-node networks, all algorithms ran in under 15 seconds (except TS which ranged from 10 to 600 s), so this is omitted for lack of space. Excluding the naive UCMT approach, GARF generally is the fastest heuristic for both reconfigurable and non-reconfigurable cases, while TS is by far the slowest due to intractable MILP formulations used in its iterations. GARF's execution time naturally grows with the network size *N*, and also with the increasing number of lightpaths (based on network load ρ), reaching 9 hours for the GÉANT network at high load. In spite of this significant computational effort, considering that the growth is not exponential, this is an acceptable time for a large instance (23 nodes and 168 time slots) of an offline planning problem.

F. Reconfigurable vs Non-reconfigurable MH-VTD

Based on the experiments presented in this paper, we investigate the benefits of periodic virtual topology reconfiguration as a valid method to reduce transceiver costs. Table VI summarizes some results for this discussion. The *Upper bound* column contains the maximum theoretical cost saving that can be achieved, i.e. the maximum possible savings that periodic virtual topology reconfiguration could bring. For the 6-node networks, the upper bound shown is in fact a comparison of the optimal solutions of the exact MILPs for the two problem variants. For the Abilene and GÉANT network scenarios, the upper bound is calculated by comparing the best non-reconfigurable solution found to the lower bound on the transceivers cost (valid for the reconfigurable case). Note that this bound is calculated using the lower bound on transceiver cost which is presumably weak for low loads. Consequently, the upper bound for the Abilene and GÉANT scenarios shown in the table may also be weak at low loads. The *Achieved savings* column shows the maximal savings actually achieved in the number of transceivers comparing the best solution obtained for the non-reconfigurable case and the best solution obtained for the reconfigurable case. These solutions come from the exact MILP in the 6-node network, and the best solutions among the heuristics in the Abilene and GÉANT experiments. Naturally, given the exact solutions of the MILP, both columns are equal for the 6-node case.

Optimal results for the 6-node networks show that the transceiver cost benefits of reconfiguring increase with the network load. However, the upper bounds for the Abilene and GÉANT instances indicate the opposite trend. Although it is inconclusive whether the bound or the solutions are weak for low loads, the fact that the achieved savings at medium loads are close to or even greater than the upper bound at high loads for both Abilene and GÉANT, seems to suggest that the possible savings does not increase dramatically from medium to high loads. For both these networks, the savings achievable are generally greater than for the 6-node network for all the loads, with a maximum achieved value of 5.8 %. However the average achieved reduction over all the scenarios seems marginal (2 %) which may not be significant enough to justify the trade-off in the extreme case where transceiver cost is the primary objective allowing unlimited reconfiguration. The results shown in Fig. 7 indicate that a more advantageous trade-off may be achieved by allowing for some (limited) reconfiguration while still comparably reducing transceivers cost. In these cases, a narrow but significant reduction in total transceivers cost (e.g. 5%) could be achieved in some scenarios with a minimal number of reconfigurations (on average, less than one lightpath reconfiguration in the whole network per hour).

TABLE VI
MAXIMAL COST SAVINGS ACHIEVED WITH RECONFIGURATION (MH-VTD-R) WITH RESPECT TO THE STATIC CASE (MH-VTD-NR)

Network	*T*	*ρ*	Upper Bound (%)	Achieved savings (%)
***N*=6**	**12**	**0.1 / 1 / 10**	0 / 0.6 / 2.3	0 / 0.6 / 2.3
Abilene *N*=11	**168**	**0.1 / 1 / 10**	32.8 / 9.7 / 4.1	0 / 5.8 / 3.8
GÉANT *N*=23	**168**	**0.1 / 1 / 10**	32.0 / 8.4 / 3.7	0 / 3.0 / 3.0

VI. CONCLUSION

This paper investigates virtual topology design (VTD) under multi-hour (MH) traffic, assuming non-reconfigurable (NR) and reconfigurable (R) lightpaths. Although the reconfigurable solution can achieve a reduction in the number of transceivers needed to establish the VTD in relation to the non-reconfigurable solution, this comes at a cost. Namely, each reconfiguration implies increased signaling complexity and network service disruption. Consequently, we investigate the associated trade-offs between transceiver reduction and increased reconfiguration frequency for the MH-VTD-R problem variant.

We propose an exact MILP formulation for small instances, as well as an efficient heuristic, called GARF (Greedy Approach with Reconfiguration Flattening), for larger problem cases which allows for tuning of the relative importance of the considered objectives. Furthermore, GARF can be used to efficiently solve the MH-VTD-NR problem variant. The efficiency and scalability of the proposed algorithm was studied through exhaustive tests conducted on both synthetic and real traffic traces, comparing GARF with analytical lower bounds, exact solutions computed by MILP formulations and sub-optimal solutions generated by other heuristic algorithms from the literature. Despite its simplicity, GARF obtains solutions which improve or match the network transceiver cost with respect to other algorithms, while strongly reducing the number of reconfigurations and with better scalability than its closer competitors. For the non-reconfigurable case, GARF provides designs with a similar cost to the best algorithms previously proposed in the literature, with similar or better scalability.

The obtained results also motivate a discussion on the benefits of virtual topology reconfiguration as a mechanism to successfully adapt to traffic variations and reduce the number of transceivers needed. Results suggest two main conclusions:

- The maximal transceiver cost reduction achieved by *unlimited* lightpath reconfiguration may not be dramatic enough to justify the trade-off with an increased reconfiguration cost;
- However, allowing a small *(limited)* amount of reconfiguration can achieve an advantageous trade-off between the two costs.

The GARF algorithm proposed in this paper provides us with an efficient mechanism to tune this trade-off to a desired value in multi-hour virtual topology design for optical networks planning.

REFERENCES

[1] J. Berthold, A. A. M. Saleh, L. Blair, and J. M. Simmons, "Optical Networking: Past, Present, and Future," *Journal of Lightwave Technology*, vol. 26, no. 9, pp.1104-1118, May 2008.

[2] S. Sygletos, I. Tomkos, and J. Leuthold, "Technological challenges on the road toward transparent networking," *Journal of Optical Networking*, vol. 7, no. 4, pp. 321-350, Apr. 2008.

[3] H. Zang, J. P. Jue, and B. Mukherjee, "A Review of Routing and Wavelength Assignment Approaches for Wavelength-Routed Optical WDM Networks," *Optical Network Magazine*, pp. 47-59, Jan. 2000.

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

[5] V.R. Konda et al., "Algorithm for traffic grooming in optical networks to minimize the number of transceivers," in *Proc. IEEE Workshop on High Performance Switching and Routing*, Dallas, TX, 2001, pp. 218-221.

[6] K. Zhu, H. Zhu, B. Mukherjee, "Traffic engineering in multigranularity heterogeneous optical WDM mesh networks through dynamic traffic grooming," *IEEE Network.*, vol. 12, no. 2, pp. 8-15, 2003.

[7] N. Srinivas, C. Siva Ram Murthy, "Design and dimensioning of a WDM mesh network to groom dynamically varying traffic," *Photonic Network Communications*, vol 7, no. 2, pp. 179-191, 2004.

[8] I. Baldine and G. N. Rouskas, "Traffic adaptive WDMnetworks:A study of reconfiguration issues," *Journal of Lightwave Technology*, vol. 19, pp. 433–455, Apr. 2001.

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

[10] B. Mukherjee, D. Banerjee, S. Ramamurthy, A. Mukherjee, "Some Principles for Designing a Wide-Area WDM Optical Network," *IEEE/ACM Transactions on Networking*, vol. 4, vo. 5, pp. 684-696, Oct. 1996.

[11] G. R. Ash, R. H. Cardwell, and R. P. Murray, "Design and optimization of networks with dynamic routing," *Bell Systems Technical Journal*, vol. 60, pp. 1787-1820, Oct. 1981.

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

[13] G. Rouskas, M. H. Ammar, "Dynamic Reconfiguration in Multihop WDM Networks," *Journal of High Speed Networks*, vol. 4, no. 3, pp. 221-238, 1995.

[14] A. Gencata, B. Mukherjee, "Virtual-Topology Adaptation for WDM Mesh Networks under Dynamic Traffic," *IEEE/ACM Transactions on Networking*, vol. 11, no. 2, pp. 236-247, Apr. 2003.

[15] D. Banerjee, B. Mukerjee, "Wavelength-Routed Optical Networks: Linear Formulation, Resource Budgeting Tradeoffs, and a Reconfiguration Study," *IEEE/ACM Transactions on Networking*, vol. 8, pp. 598-607, 2000.

[16] B. Ramamurthy, A. Ramakrishnan, "Virtual Topology Reconfiguration of Wavelength-Routed Optical WDM Networks," *IEEE Globecom 2000*, 2000.

[17] K. Liu, C. Liu, J. Pastor, A. Roy, J. Wei, "Performance and Testbed Study of Topology Reconfiguration in IP over Optical Networks," *IEEE Transactions on communications*, vol. 50, no. 10, pp. 1662-1679, Oct. 2002.

[18] 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, pp. 1231–1240, Oct. 2003.

[19] 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, pp. 2–15, 2006.

[20] 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*, Boston, USA, pp 13-22, Oct. 2005.

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

[22] F. Ricciato, S. Salsano, A. Belmonte, and M. Listanti, "Off-line Configuration of a MPLS over WDM Network under Time-Varying Offered Traffic," *IEEE Infocom*, pp. 57-65, 2002.

[23] G. Agrawal and D. M edhi, "Lightpath Topology Configuration for Wavelength-Routed IP/MPLS Network for Time-Dependent Traffic", *Proc. of IEEE Globecom 2006*, San Francisco, Nov-Dec 2006.

[24] B. Garcia-Manrubia, R. Aparicio-Pardo, P. Pavon-Marino, N. Skorin-Kapov, and J. Garcia-Haro, "MILP Formulations for Scheduling Lightpaths under Periodic Traffic," in *Proc. 11th International Conference on Transparent Optical Networks*, *ICTON 2009*, Island of São Miguel, Azores, Portugal, June 2009.

[25] N. Skorin-Kapov, P. Pavon-Marino, B. Garcia-Manrubia, and R. Aparicio-Pardo, "Scheduled Virtual Topology Design Under Periodic Traffic in Transparent Optical Networks," in *Proc. 6th ICST International Conference on Broadband Communications, Networks and Systems*, *BroadNets 2009*, Madrid, Spain, Sept. 2009.

[26] R. Aparicio-Pardo, P. Pavon-Marino, N. Skorin-Kapov, B. Garcia-Manrubia, and J. Garcia-Haro, "Algorithms for virtual topology reconfiguration under multi-hour traffic," in *Proc. 12th International Conference on Transparent Optical Networks*, *ICTON 2010*, Munich, Germany, June 2010, p. We.B1.2.

[27] G. Oriolo, "Domination between traffic matrices," *Mathematics of Operations Research*, vol. 33, no. 1, pp. 91-96, Feb. 2008.

[28] R. K. Ahuja, T. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*, Prentice Hall, 1993, pp. 4–8

[29] 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. IFIP Int. Conf. Optical Network Design and Modelling*, Athens, Greece, pp. 58-67, 2007.

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

[31] J. Milbrandt, M. Menth, S. Kopf, "Adaptive Bandwidth Allocation: Impact of Traffic Demand Models for Wide Area Networks," in *19th International Teletraffic Congress (ITC19)*, Beijing, China, 2005.

[32] D. Medhi, "Multi-hour, multi-traffic class network design for virtual path-based dynamically reconfigurable wide-area ATM networks," *IEEE/ACM Transactions on Network*, vol. 3, no. 6, pp. 809-818, 1995.

Ramón Aparicio-Pardo received the Engineering in Telecommunications and the Ph.D. degree from the Universidad Politécnica de Cartagena (UPCT), Spain, in 2006 and 2011, respectively. His research interests include planning and performance evaluation of optical switching networks.

Nina Skorin-Kapov received the Dipl.-Ing. and Ph.D. degrees in electrical engineering from the University of Zagreb, Zagreb, Croatia, in 2003 and 2006, respectively. She has been working as 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.

Pablo Pavón Mariño received his Telecommunications Engineering degree from the University of Vigo (Spain) in 1999. In 2000 he joined the Universidad Politécnica de Cartagena (Spain), where he received his Ph.D. degree in Telecommunications in 2004. In 2010 he received his degree in Mathematics from the UNED University (Spain). He is Associate Professor in the Department of Information Technologies and Communications. His research interests include performance evaluation and optimization of communication networks.

Belen Garcia-Manrubia received the Engineering degree in Telecommunication and the M.Sc. degree from the Universidad Politécnica de Cartagena (UPCT), Spain, in 2007 and 2008, respectively. She is currently working toward the Ph.D. degree at the Department of Information Technologies and Communications at the same university. Her research interests include planning and performance evaluation of optical networks, especially cross-layer planning under wavelength-routing paradigm.

TABLE II
MH-VTD-R CASE: NUMBER OF TRANSCEIVERS (AND RECONFIGURATIONS) FOR EACH TESTING SCENARIO

Network	T	ρ	LB	MILP	GARF	TS	LRvSG	ICET	Gap(%)
N=6	12	0.1	12.0	16.0 (0)	21.6 (**2.2**)	**16.4** (2.4)	48.2 (4.8)	24.6 (2.6)	2.4%
		1	71.0	71.2 (0.2)	72.4 (**0.2**)	**71.4** (3.5)	89 (2.1)	73.0 (3.2)	0.3%
		10	647.0	647.2 (2.9)	651.4 (**6.1**)	**648.6** (33.6)	671.8 (34.7)	653.6 (33.5)	0.2%
Abilene *N*=11	168	0.1	39	- (-)	**64** (4.4)	66 (8.2)	108 (**2.4**)	88 (11.9)	39.1%
		1	278	- (-)	301 (**0.9**)	**290** (11.6)	396 (4.0)	303 (11.7)	4.1%
		10	2659	- (-)	2671 (**2.6**)	**2669** (53.8)	2770 (51.2)	2682 (54.1)	0.4%
GÉANT *N*=23	168	0.1	136	- (-)	**203 (7.5)**	- (-)	983 (116.6)	212 (16.8)	33.0%
		1	1139	- (-)	**1207 (2.7)**	- (-)	1895 (17.3)	1240 (36.7)	5.6%
		10	11190	- (-)	**11265 (9.8)**	- (-)	11768 (196.0)	11293 (215.8)	0.7%

TABLE III
MH-VTD-NR (STATIC CASE). NUMBER OF TRANSCEIVERS IN THE SOLUTIONS OBTAINED FOR EACH TESTING SCENARIO

Network	T	ρ	LB	MILP	GARF	TS	LRvSG	3-Step-H	UCMT	Gap(%)
N=6	12	0.1	12.0	**16.0**	23.2	25.6	23.6	19.2	**18.8**	31.0%
		1	71.0	**71.6**	**72.4**	74.0	100.8	76.4	80.4	1.1%
		10	647.0	**662.4**	**673.6**	724.0	747.2	686.4	728	1.7%
Abilene *N*=11	168	0.1	39	-	68	216	168	**58**	62	42.6%
		1	278	-	312	406	430	**308**	346	10.9%
		10	2659	-	2778	3150	3240	**2774**	3156	4.3%
GÉANT *N*=23	168	0.1	136	-	210	-	1034	**200**	222	35.2%
		1	1139	-	**1244**	-	2052	1306	1476	8.4%
		10	11190	-	**11620**	-	13940	12044	13664	3.7%

TABLE V
ALGORITHMS' EXECUTION TIMES (HOURS)

		MH-VTD-R				*MH-VTD-NR*				
Network	**ρ**	**GARF**	**TS**	**LRvSG**	**ICET**	**GARF**	**TS**	**LRvSG**	**3-Step-H**	**UCMT**
Abilene *N*=11, *T*=168	0.1	**0.05**	101.98	1.74	0.14	0.01	65.49	1.04	0.24	**<0.01**
	1	**0.06**	28.97	1.74	0.11	0.03	1.63	1.05	0.23	**<0.01**
	10	0.67	10.53	1.74	**0.17**	0.41	1.64	1.08	0.23	**0.01**
GÉANT *N*=23, *T*=168	0.1	**0.55**	-	43.71	6.71	0.32	-	16.64	39.00	**0.16**
	1	**1.17**	-	44.71	7.63	0.68	-	16.05	38.42	**0.17**
	10	**9.00**	-	44.18	16.55	6.12	-	15.64	38.52	**0.35**

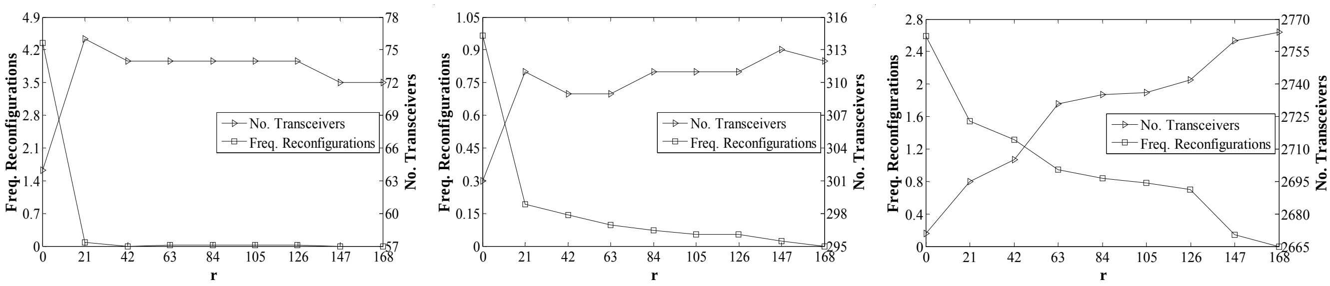

(a) (b) (c)

Fig 7. Frequency of Reconfigurations (i.e. the number of reconfigurations per time slot) and Number of Transceivers vs. parameter *r* for the Abilene average week (*T*=168) with (a) $\rho = 0.1$, (b) $\rho = 1$, (c) $\rho = 10$.