

Virtual topology design and flow routing in optical networks under multi-hour traffic demand

Pablo Pavon-Marino · Ramon Aparicio-Pardo ·
Belen Garcia-Manrubia · Nina Skorin-Kapov

P. Pavon-Marino · R. Aparicio-Pardo · B. Garcia-Manrubia
Technical University of Cartagena (UPCT), Plaza Hospital 1, 30202, Cartagena, Spain

phone: +34 968325952

fax: +34 968325973

e-mail: pablo.pavon@upct.es

R. Aparicio-Pardo

e-mail: ramon.aparicio@upct.es

B. Garcia-Manrubia

e-mail: belen.garcia@upct.es

N. Skorin-Kapov

Faculty of Electrical Engineering and Computing (FER), University of Zagreb, Unska 3,1000, Zagreb, Croatia

phone: +385 16129731

fax: +385 1 6129832

e-mail: nina.skorin-kapov@fer.hr

Abstract This paper addresses the problem of finding a static virtual topology design and flow routing in transparent optical WDM networks under a time-varying (multi-hour) traffic demand. Four variants of the problem are considered, using fixed or dynamically adaptable (i.e., variable) flow routing, which can be splittable or unsplittable. Our main objective is to minimize the number of transceivers needed which make up for the main network cost. We formulate the problem variants as exact ILPs (Integer Linear Programs) and MILPs (Mixed ILPs). For larger problem instances, we also propose a family of heuristics based on the concept of domination between traffic matrices. This concept provides the theoretical foundations for a set of techniques proposed to reduce the problem complexity. We present a lower bound to the network cost for the case in which the virtual topology could be dynamically reconfigured along time. This allows us to assess the limit on the maximum possible benefit that could be achieved by using optical reconfigurable equipment. Extensive tests have been conducted, using both synthetically generated and real-traced traffic demands. In the cases studied, results show that combining variable routing with

splittable flows obtains a significant, although moderate, cost reduction. The maximum cost reduction achievable with reconfigurable virtual topologies was shown to be negligible compared to the static case in medium and high loads.

Keywords *Virtual Topology Design · Flow routing · Periodic traffic · Traffic domination*

1 Introduction

Transparent optical networks based on Wavelength Division Multiplexing (WDM) technology have been established as the enabling technology for today's high-speed backbone networks. In such networks, a set of all-optical connections, called lightpaths, are set up between pairs of nodes, forming a virtual topology on top of the physical interconnection of optical fibers. Each lightpath can

be routed over multiple physical links and occupies a single transmission wavelength on each traversed fiber. Traffic carried over the lightpaths is electronically processed only at the source and destination nodes, where it is said to be ‘added’ or ‘dropped’, while transparently traversing intermediate ones. Thus, each lightpath occupies one E/O transmitter at the ingress node, and one O/E receiver at the egress node.

The optical part of the switching nodes enabling WDM channels to be transparently switched from the input to the output ports is called the Wavelength Switch Fabric (WSF). The WSF can be realized as a fixed or a reconfigurable component. In the former, the interconnection of input and output ports is manually hard-wired and cannot change over time. In the latter, the lightpaths can be established, rerouted and torn down in tens of milliseconds, according to the commands received from management and/or the control plane. Naturally, reconfigurable equipment implies an increase in equipment costs and signaling complexity.

Given a set of nodes and a traffic demand, determining the set of lightpaths which are to be established between input-output node pairs is called virtual topology design (VTD). To establish such a set of lightpaths, it is necessary to find for each a sequence of traversed fibers and assign to it a wavelength subject to certain constraints. This is referred to as the Routing and Wavelength Assignment problem (RWA). Solving the RWA problem for the set of lightpaths which comprises the virtual topology certifies the feasibility of the network plan with respect to wavelength availability and physical impairments. However, the disparities in the network cost among different RWA solutions for the same virtual topology are not significant. Consequently, we assume that the links in the network are equipped with a sufficient number of wavelengths and that physical-layer constraints do not apply. Thus, the associated RWA constraints can be removed from the network planning optimization problem simplifying its complexity. Such assumptions can realistically depict several network scenarios, e.g. in metro-area

optical networks with an over-dimensioned fiber plant. In this paper we consider a network where the previous assumptions hold, and thus eliminate the RWA constraints from our analysis.

We can classify VTD planning problems according to the associated traffic demand variations along time, as static, scheduled or dynamic. The first two assume offline planning, while the third infers online provisioning. In the static traffic case, the traffic demand is constant along time. Therefore, the virtual topology obtained by solving the VTD problem is also constant where the lightpaths are established semi-permanently based on the estimated constant traffic. In the scheduled traffic case, the traffic demand varies along time, but its variation is assumed to be accurately known in advance. The problem can be solved either by finding a static virtual topology capable of handling all the given traffic variations over time, or by finding a succession of virtual topologies dynamically reconfigured in accordance with the traffic changes. The former can be realized using non-reconfigurable components, while the latter requires reconfigurable WSFs. Finally, in the dynamic traffic case, lightpaths are established as the connection requests arrive, with no *a priori* deterministic knowledge of connection request arrival time.

In this paper, we consider offline virtual topology planning in transparent optical networks given a multi-hour traffic demand, i.e. for known or estimated traffic variations. We denote this problem as MH-VTD (Multi-Hour Virtual Topology Design). The traffic demand over a certain period of time is modeled as a temporal series of traffic matrices, each estimating the demand in a discrete time slot. This can be used, e.g., to model periodic traffic as it changes over the course of a week due to peak hours [1]. We assume that the virtual topology can not be changed along time, i.e. we determine a static set of lightpaths capable of routing all traffic as it varies over all time slots. Thus, the network can be based on non-reconfigurable optical switching equipment. This assumption is in line with the current capabilities of the control plane in transparent optical networks

and avoids dynamic ‘hot’ reconfigurations of the virtual topology which can imply short but relevant traffic disruptions.

In summary, our target problem is finding an optimal fixed virtual topology, and the associated traffic flow routing on top of it, which is capable of fully carrying the given traffic in all time slots. The cost to minimize is measured in terms of the number of lightpaths in the network, which define the number of transceivers (transmitters and receivers) needed. This number is commonly considered as a network cost criterion in planning studies [2],[3],[4].

We are interested in investigating the trade-offs that arise in the MH-VTD planning problem with respect to two different criteria:

- Considering fixed or variable routing of higher layer flows on top of the static virtual topology over time. A fixed flow routing indicates that the traffic from a source node to a destination node is always transmitted via the same set of lightpaths. On the contrary, variable routing imposes no such constraints, but at a cost of higher signaling and network management complexity.
- Considering splittable or unsplittable flow routing. Unsplittable flow routing means that all the traffic between a given input-output pair at any given time is constrained to traverse the same sequence of nodes. If the traffic between a node pair exceeds the capacity of a single lightpath, the unsplittable traffic must be routed over a set of multiple lightpaths which are all-optically routed across the same sequence of nodes. However, if splittable flow routing is assumed, the traffic between two nodes can be split, where different fractions of the traffic can be routed over a different sequence of lightpaths not necessarily traversing the same nodes. Naturally, splittable routing allows for improved traffic balancing, but at a cost of increased signaling overhead.

The two categorizations described above jointly define four different scenarios under study: fixed or variable routing (FR and VR, respectively), with splittable or unsplittable flows (s and u respectively). We denote these four variants of the problem as MH-VTD-FRs, MH-VTD-FRu, MH-VTD-VRs, MH-VTD-VRu.

In this paper, two approaches are proposed to address the MH-VTD problem and its variants. The first approach is based on exact ILP (Integer Linear Programming) and MILP (Mixed ILP) formulations, yielding optimal solutions to the problem. Clearly, all variants of the MH-VTD problem are NP-hard, as the static virtual topology design problem is [5]. For larger problem instances we propose a family of heuristic methods based on the concept of traffic domination [6], in conjunction with a heuristic for the static virtual topology design problem. Also, we propose a lower bound on the number of transceivers in the network for the case in which the virtual topology can be changed along time (and, thus, reconfigurable WSFs must be used). This lower bound is valid to evaluate the maximum cost saving in number of transceivers that could be achieved if the virtual topology could be reconfigured along time.

We present extensive results using synthesized traffic variations with various traffic loads and randomness factors, for networks of up to 8 nodes. In addition, we conduct tests for larger problem instances considering a time series of 672 traffic matrices obtained from a real traffic trace in the Abilene network **¡Error! No se encuentra el origen de la referencia.** The results are used to validate the heuristic methods, as well as extract interesting conclusions regarding the usefulness of dynamic flow routing, splittable flow routing, and the maximum benefits that can be obtained from allowing a dynamic reconfiguration of the virtual topology.

The rest of the paper is organized as follows. Section 2 presents the state-of-the-art for virtual topology design under a time-varying traffic demand. In Section 3 we provide optimal MILP formulations for the ST-VDT problem with

fixed/variable routing and (un)splittable flows. Section 4 presents our heuristic approach, which includes a reduction technique based on the concept of domination among traffic matrices. Section 5 describes the conducted experiments and obtained results. Finally, Section 6 concludes the paper.

2 State of the art

The periodic nature of traffic has been confirmed with real traffic traces, such as the Abilene backbone network trace [7], making the expected traffic load in a network fairly predictable. While multi-hour planning has been studied in the past for different communication networks (see [8] for details), to the best of our knowledge, very few works exist which study periodic or known traffic variations in the context of transparent optical networks planning.

In [9] the authors consider the problem of finding a static VTD based on scheduled traffic demands with fixed and variable flow routing assuming splittable flows. They analyze various objectives, such as the average packet hop distance, congestion, and total number of lightpaths, and propose a hybrid one combining all three. The objective to minimize lightpaths is equivalent to minimizing transceivers since each lightpath requires one transmitter at the source node, and one receiver at the destination node. However, the work in [9] assumes a given number of transmitters and receivers per node and, thus, does not investigate the maximal possible savings in this respect. Furthermore, they assume at most one lightpath between each pair of nodes, which is very restrictive. They provide formulations using the various objectives and test on 7 and 10 node networks for 3 time slots. In their results, flow rerouting achieved a 7% savings in the number of lightpaths. We extend these findings by proposing a heuristic approach which can handle a large number of time instances, and give MILP formulations with no constraints on the number transceivers and number of lightpaths between node pairs. As such, we can assess the maximal benefit which can be achieved with respect to the number of lightpaths needed for fixed vs. variable routing. Furthermore,

we investigate the results of a real traffic trace and compare with problem variants assuming unsplittable routing.

In [10], the multilayer problem of mapping a set of lightpaths onto the physical topology and mapping a set of LSPs (Label Switched Paths) for packet switched traffic over the set of lightpaths is considered. They assume known traffic variations and consider that one LSP is required for each flow. They study the case of designing a static virtual topology with fixed flow routing and unsplittable flows, assuming that a single lightpath capacity is sufficient for each LSP. They give a MILP formulation, as well as a decomposition heuristic method. They compare the results with those obtained for a reconfigurable approach, i.e. finding a succession of virtual topologies by solving the multilayer problem for each time slot separately. Their objective is a combination of the amount of average electronic packet processing along time, the number of optically switched wavelengths and the number of lightpaths in the network.

Their results indicate that establishing a static virtual topology with fixed routing obtains solutions very close to those in which the virtual topology can be dynamically reconfigured over time. They also compare their results with a method in which the temporal sequence of traffic matrices is reduced to one matrix: the traffic between nodes (i,j) is assumed to be the maximum along time of the traffic between those nodes. The static VTD problem is then solved for this matrix using a MILP formulation.

The approach that we follow in this paper is similar in the sense that an effort is first made to reduce the temporal sequence of traffic matrices to one unique matrix. Then, we solve the problem for that matrix. However, our approach applies the concept of traffic domination [6] to give a theoretical foundation to this process. It allows us to apply separately the concepts of so-called weak and total domination in order to solve the specific cases of variable and fixed flow routing, respectively. From this, we design a family of heuristics which reduce the temporal sequence of traffic matrices to one, and then apply a heuristic

for solving the VTD problem which is more scalable than solving the exact MILP.

In [11] and [12], we investigated the Multi-Hour VTD problem for the cases of variable and splittable flow routing for networks with reconfigurable and non-reconfigurable equipment. With reconfigurable equipment, the virtual topology can be dynamically changed over time, while the case with non-reconfigurable components maps to the MH-VTD-VRs problem variant described in this paper. MILP formulations [11] and a tabu search heuristic approach [12] are proposed but further testing is needed to more completely assess the benefits and trade-offs associated with using reconfigurable components. In this paper, we consider only the case of establishing a static non-reconfigurable virtual topology, but compare various possibilities for flow routing (i.e., fixed/variable; (un)splittable) in order to assess their cost-effectiveness. We extend the formulations given in [11] for the MH-VTD-VRs problem to describe all the mentioned variations of flow routing constraints and propose new heuristics based on traffic matrix domination to solve larger problem instances.

3 ILP/MILP-Based Approach

In this section, we propose exact ILP/MILP formulations for the four problem variants investigated. Let N be the number of nodes in the network, and T the number of time intervals for which the traffic is defined. 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 d_{sd}^t , $s,d=\{1,\dots,N\}$, $t=\{1,\dots,T\}$ denote the traffic demand (measured in Gbps) from node s to node d , during time interval t . Let C denote the lightpath capacity in Gbps.

3.1 MH-VTD-FR-s/u problems

The decision variables of the problem are:

- $u(i, j) = \{0, 1, 2, \dots\}$, $i, j = \{1, \dots, N\}$. The number of lightpaths from node i to node j .
- $f(i, j, s, d) \in [0, 1]$, $i, j, s, d = \{1, \dots, N\}$. The fraction of the total traffic flow from node s to node

d that is routed on the existing lightpaths from node i to node j .

The problem formulation is given by (1).

$$\min \sum_{i,j=\{1,\dots,N\}} u(i, j) \quad (1a)$$

subject to:

$$\sum_{s,d=\{1,\dots,N\}} d_{sd}^t \cdot f(i, j, s, d) \leq C \cdot u(i, j), \quad (1b)$$

$$\forall i, j = \{1, \dots, N\}, \forall t = \{1, \dots, T\}$$

$$\sum_{j=1,\dots,N} f(n, j, s, d) - \sum_{i=1,\dots,N} f(i, n, s, d) = \begin{cases} 1, & \text{if } n = s \\ -1, & \text{if } n = d \\ 0, & \text{otherwise} \end{cases} \quad (1c)$$

$$\forall n, s, d = \{1, \dots, N\}, \forall t = \{1, \dots, T\}$$

MH-VTD-FRu:

$$f(i, j, s, d) = \{0, 1\}, \forall i, j, s, d = \{1, \dots, N\} \quad (1d)$$

The objective function (1a) searches for a virtual topology that minimizes the number of lightpaths, which is equivalent to minimizing the number of transceivers in the network. Constraints (1b) force the traffic between two nodes at any time, to be limited by the number of lightpaths planned in the virtual topology between those nodes. Constraints (1c) are the link-flow conservation constraints. Integrality constraints (1d) appear only in the unsplittable case (MH-VTD-FRu problem), preventing the flow between two nodes to be fractioned among different paths.

3.2 MH-VTD-VR-s/u problems

The decision variables of the problem are:

- $u(i, j) = \{0, 1, 2, \dots\}$, $i, j = \{1, \dots, N\}$. The number of lightpaths from node i to node j .
- $f(i, j, s, d, t) \in [0, 1]$, $i, j, s, d = \{1, \dots, N\}$, $t = \{1, \dots, N\}$. The fraction of the total traffic flow from node s to node d that is routed on the existing lightpaths from node i to node j during time interval t .

The problem formulation is given by (2).

$$\min \sum_{i,j=\{1,\dots,N\}} u(i,j) \quad (2a)$$

subject to:

$$\sum_{s,d=\{1,\dots,N\}} \{d_{sd}^t \cdot f(i,j,s,d,t)\} \leq C \cdot u(i,j), \quad (2b)$$

$$\forall i,j = \{1,\dots,N\}, \forall t = \{1,\dots,T\}$$

$$\sum_{j=1,\dots,N} f(n,j,s,d,t) - \sum_{i=1,\dots,N} f(i,n,s,d,t) = \begin{cases} 1, & \text{if } n=s \\ -1, & \text{if } n=d \\ 0, & \text{otherwise} \end{cases} \quad (2c)$$

$$\forall n,s,d = \{1,\dots,N\}, \forall t = \{1,\dots,T\}$$

MH-VTD-VRu:

$$f(i,j,s,d,t) = \{0,1\}, \forall i,j,s,d = \{1,\dots,N\}, \forall t = 1,\dots,T \quad (2d)$$

As in the formulation (1), the objective function (2a) searches for a virtual topology that minimizes the number of lightpaths. Both capacity constraints (2b) and link-flow conservation constraints (2c) have been trivially adapted to consider a routing which can change along time, by considering independent constraints in each time slot. Again, integrality constraints (2d) appear only in the unsplittable case MH-VTD-VRu.

4 Heuristic approach

Herein, we propose a family of heuristics to solve the four variants of the MH-VTD problem described in this paper. We make use of the concept of domination between traffic matrices introduced in [6], as a foundation for our approach. For the sake of completeness, the next subsection briefly presents the relevant definitions and properties of the traffic matrix domination concept used in the proposed heuristics.

4.1 Traffic Domination

Let N be the number of nodes in the network. A traffic matrix D is a non-negative $N \times N$ matrix, where each element d_{ij} represents the amount of traffic demand from node i to node j . A $N \times N$ capacity matrix U is a non-negative matrix where

each element u_{ij} represents the amount of capacity installed on the link between nodes i and j . A routing $F = \{f_{ijhk}\}$, $i,j,h,k=1,\dots,n$ is a function which assigns to each element f_{ijhk} , the fraction (in the interval $[0,1]$) of traffic demand d_{ij} that is routed on the link between nodes h and k . A routing can be splittable or unsplittable. In this model, a routing F is unsplittable if it is integer, i.e., if $f_{ijhk} = \{0,1\}$, for all $i,j,h,k=1,\dots,n$. Finally, we say that a capacity matrix U supports a traffic matrix D if there exists a feasible routing F which enables all the traffic in D to be carried over U .

Formal definitions of the two types of domination applied in this paper, called *weak* and *total domination*, follow.

Definition 1: Let D^1 and D^2 be two traffic matrices. We say that traffic matrix D^1 *weakly dominates* D^2 if for any capacity matrix U supporting D^1 , then there exists at least one routing F such that (U,F) also supports D^2 .

Definition 2: Let D^1 and D^2 be two traffic matrices. We say that traffic matrix D^1 *totally dominates* D^2 if any capacity and routing pair (U,F) supporting D^1 , also supports D^2 .

The difference between total and weak domination comes from the flow routing perspective. If a traffic matrix D^1 totally dominates another traffic matrix D^2 , then any capacity and routing pair (U,F) that supports D^1 also supports D^2 . If the domination is weak, then finding a capacity U and a routing F which supports D^1 guarantees that capacity U is enough to carry the traffic D^2 , but does not guarantee that it can be done with the same routing F . Naturally, total domination is a stronger condition, which implies weak domination.

A second variation of weak domination was also considered in [6], *weak domination with respect to unsplittable flows*, which applies to our planning scenario analysis.

Definition 3: Let D^1 and D^2 be two traffic matrices. We say that traffic matrix D^1 *weakly dominates* D^2 *with respect to unsplittable flows* if for any capacity matrix U supporting D^1 with unsplittable flows, then there exists at least one unsplittable routing F such that (U,F) also supports

D^2 .

Note that the general weak domination case applies to general (and, thus, splittable) routing. Also, note that if a traffic matrix D^1 totally dominates another matrix D^2 , the capacity and routing which is feasible for D^1 is also feasible for D^2 , whether the routing is splittable or not. Thus, definitions 1-3 sufficiently cover all variants considered in the paper.

The following properties are extracted from [6] and are the instrument we use to assess domination between traffic matrices. We refer to [6] for the proofs.

Property 1: D^1 weakly dominates D^2 with respect to (un)splittable flows if and only if D^1 (considered as capacity matrix) supports D^2 by (un)splittable flows.

Property 2: D^1 totally dominates D^2 if and only if $d_{ij}^1 \geq d_{ij}^2$, for any coordinate of the traffic matrix $i, j=1, \dots, n$.

4.2 The general scheme of the proposed 3-step heuristic approach

The four MH-VTD problem variants considered in this paper have as an input parameter a sequence of T traffic matrices D^1, \dots, D^T , each of them of size $N \times N$, being N the number of nodes in the network. The heuristics proposed for each of the 4 problem variants follow a common scheme, composed of three basic steps.

Step 1: Replace-with-1 matrix. The initial sequence of traffic matrices D^1, \dots, D^T is replaced with one unique matrix D , such that matrix D dominates each of the matrices D^1, \dots, D^T . The type of domination depends on the problem variant, and will be described in detail in the subsection below.

Step 2: Static Virtual Topology Design (VTD). Dominating matrix D is used to feed a heuristic algorithm for calculating a static virtual topology U (i.e. a capacity matrix) and a routing F (unsplittable for problems MH-VTD-FRu and MH-VTD-VRu), which minimizes the number of transceivers, our primary cost measure of interest.

Step 3: (MH-VTD-VR problems) Per-time-slot flow routing calculation. For the MH-VTD-VR

problem variations, the flow routing calculated in the previous step is not guaranteed to be a feasible routing over all time slots. Consequently, T instances of the multicommodity flow routing problem are solved to individually route the traffic flows in matrices D^t , $t=1, \dots, T$ on top of virtual topology U for each time slot separately.

The crucial point in the method is that the traffic dominance concept guarantees that a virtual topology U (a capacity matrix) calculated in Step 2 which supports the traffic of D , will also be able to support the traffic in all the time slots D^1, \dots, D^T . As such, virtual topology U is a suitable virtual topology for our MH-VTD planning problem.

A more detailed explanation of each individual step in the proposed 3-step heuristic follows in the next subsections.

4.3 Step 1: Replace-with-1 methods

Step 1 is a procedure to replace a sequence of traffic matrices D^1, \dots, D^T with one dominating matrix D , such that D dominates all matrices in the sequence, i.e. matrix D dominates each matrix D^t , $t=1, \dots, T$. Depending on the problem variant considered, different types of domination between matrix D and matrices D^1, \dots, D^T are applied. For each case, we aim to find that matrix D among all possible matrices D' which dominate the sequence (which includes, e.g., a matrix with all the coordinates set to infinity), which yields the most cost-effective virtual topology U . This aspect is described for each problem variant separately.

4.3.1 Total domination for problems MH-VTD-FRs/u

In the MH-VTD-FRs and MH-VTD-FRu problems, the flow routing must remain constant along time. Consequently, we must guarantee that a common routing F exists, such that (U, F) supports traffic D^i , for all $i=1, \dots, T$. The following lemma proves the validity of the total domination relation as a strategy to solve this problem.

Lemma 1: Let D be a matrix which totally

dominates a sequence of traffic matrices D^t , $t=1,\dots,T$. Let U_s be a virtual topology and F_s a flow routing calculated such that (U_s, F_s) supports matrix D with splittable flows. Then, (U_s, F_s) is a solution for the MH-VTD-FRs problem. For the same matrix D , let U_u be a virtual topology and F_u a flow routing calculated such that (U_u, F_u) supports matrix D with unsplitable flows. Then, (U_u, F_u) is a solution for the MH-VTD-FRu problem.

Proof: The proof comes immediately from applying the concept of total domination: as D totally dominates the sequence, every capacity-routing pair (U, F) that support D , also supports every matrix in the sequence D^t , $t=1,\dots,T$. If the routing F is unsplitable, (U, F) is a solution to the MH-VTD-FRu problem.

The following lemma provides a method to calculate the total dominating traffic matrix D we are searching for.

Lemma 2: Given a sequence of traffic matrices D^1, \dots, D^T , the matrix $D = \{d_{ij}\}$, $i, j = 1, \dots, n$, such that $d_{ij} = \max\{d_{ij}^1, \dots, d_{ij}^T\}$, totally dominates the sequence.

Proof: As $d_{ij} = \max\{d_{ij}^1, \dots, d_{ij}^T\}$, then for any $i, j = 1, \dots, N$, it holds that $d_{ij} \geq d_{ij}^t$ for all $t = 1, \dots, T$. Then, matrix D totally dominates the sequence.

The next lemma states that *any* other matrix which totally dominates the sequence, other than the one above, yields a less cost-effective solution for any cost measure we devise.

Lemma 3: Let D^a and D^b be two traffic matrices which totally dominate the same traffic matrix sequence D^1, \dots, D^T . Let D^a be calculated as shown in the previous lemma, while D^b is calculated differently. Let $c(U, F)$ be a cost function which associates a cost to a virtual topology U and a flow routing F . Let (U^a, F^a) be the virtual topology and routing pair that supports D^a with the minimum cost attending to the cost function c . Let (U^b, F^b) be the virtual topology and routing pair that supports D^b with the minimum cost attending to the same cost function c . Then, the cost obtained with D^a is always equal to or lower than the cost obtained with D^b : $c(U^a, F^a) \leq c(U^b, F^b)$.

Proof: The proof is based on showing that matrix

D^b must totally dominate matrix D^a . This is proved, by reduction to the absurd. Let us assume that D^b does not dominate D^a . Then, there exists a node pair $i, j = 1, \dots, N$ such that $d_{ij}^b < d_{ij}^a$. Let D^t be the matrix in the sequence such that $d_{ij}^a = d_{ij}^t$, that is, the matrix associated to the time slot with the maximum value in the position (i, j) . Then, it holds that $d_{ij}^b < d_{ij}^t$. Therefore, a contradiction is found, as if matrix D^b does not dominate D^a , then matrix D^b can not dominate matrix D^t , and thus can not dominate the sequence.

As matrix D^b totally dominates matrix D^a , every (U, F) pair that supports D^b , also supports D^a . It follows that the (U^b, F^b) pair that minimizes the cost figure supports both D^a and D^b . Then, $c(U^a, F^a) \leq c(U^b, F^b)$, as we wanted to prove.

Note that the method to calculate the totally dominating matrix D is applied to both the MH-VTD-FRs and MH-VTD-FRu problems. As will be shown later, the heuristics proposed for these two problems differ only in Step 2 of the approach.

4.3.2 Weak domination for problems MH-VTD-VRs/u

In the MH-VTD-VRs and MH-VTD-VRu problems, we are interested in finding a common virtual topology U and a time varying flow routing for the traffic sequence D^1, \dots, D^T . The following lemmas prove the validity of the weak domination relation as a strategy to solve this problem.

Lemma 4: Let D be a matrix which weakly dominates a sequence of traffic matrices D^t , $t=1,\dots,T$. Let U be a virtual topology and F a flow routing calculated such that (U, F) supports matrix D with splittable flows. Then, there is at least one solution to the MH-VTD-VRs problem which has U as its common virtual topology, but potentially with different flow routings in different time slots.

Proof: The proof is based on the definition of weak domination: as matrix D weakly dominates the sequence, and virtual topology U supports D , then the same virtual topology U supports any matrix in the sequence. However, note that the weak domination relation does not guarantee a common flow routing along time.

Lemma 5: Let D be a matrix which weakly dominates a sequence of traffic matrices $D^t, t=1, \dots, T$ with respect to unsplittable flows. Let U be a virtual topology and F an unsplittable flow routing calculated such that (U, F) supports matrix D with unsplittable flows. Then, there is at least one solution to the MH-VTD-VRu problem which has U as its common virtual topology, but potentially with different flow routings in different time slots, all of them unsplittable.

Proof: The proof is analogous to the one for the previous lemma, but replaces the “weak domination” relation with the “weak domination with respect to unsplittable flows” relation.

Given a traffic matrix sequence D^1, \dots, D^T , the method to calculate a matrix D which weakly dominates the sequence is based on the application of the Property 1 from Section 4.1. This property states that the traffic matrix D we are searching for, seen as a non-integral capacity matrix, should be able to support all the traffic matrices in the sequence. Thus, finding a matrix D for the replace-with-1 step in the MH-VTD-VR case, implies finding a common non-integral capacity matrix D which is capable of carrying the traffic over all the time slots in the sequence. In other words, finding a matrix D which weakly dominates the sequence, and finding a virtual topology U which solves the MH-VTD-VR problem, are *almost* equal problems. They can be solved with *almost* the same formulation. The difference between the formulation (2), which optimally solves the MH-VTD-VRs/u problems, and the formulation for Step 1 of our heuristic is that in Step 1, the integral constraints for the capacity matrix are removed. Therefore:

- Calculating a matrix D which weakly dominates the sequence can be done by solving formulation (2) for the MH-VTD-VRs problem, by relaxing the integrality constraints for capacity matrix U . The capacity matrix U obtained by the relaxed problem is the weakly dominating matrix searched. Note that by relaxing the integrality constraints in variables U , an LP

(Linear Programming) problem is obtained, which can be solved in polynomial time.

- Analogously, calculating a matrix D which weakly dominates the sequence with respect to unsplittable flows can be done by solving formulation (2) for the MH-VTD-VRu problem, by relaxing the integrality constraints for capacity matrix U . However, for this case, a MILP problem must still be solved since the routing variables remain binary as a consequence of the unsplittable flow constraint.

In summary, the application of the weak domination concept in the Replace-with-1 step allows us to reduce the complexity of the problem by relaxing the integrality constraints in the capacity variables. Furthermore, the domination concept guarantees that any virtual topology supporting the resulting matrix will also support the traffic along all the time slots.

Note that the method proposed above for finding matrix D , based on relaxing formulation (2), implies finding the weakly dominating matrix which minimizes the sum of its coordinates. This is a reasonable way to produce cost-effective virtual topologies and routings. Unfortunately, the method does not guarantee that any other matrices D' which weakly dominate the same sequence, would yield virtual topologies of equal or higher cost. In contrast to the case for the total dominating relation described in Lemma 3, it is not true that the matrix D calculated by solving the relaxed version of formulation (2) is weakly dominated by all other matrices D' which weakly dominate the given sequence. Consequently, finding the matrix which provides the most cost-effective virtual topology, is presumably as difficult as optimally solving the MH-VTD-VR problem.

While the Replace-with-1 step described in this paper, implies a reduction in the complexity with respect to the original problem, calculating the weakly dominating matrix of a long traffic matrix sequence can still be too complex (or intractable in the MH-VTD-VRu case, which involves solving a MILP). As such, we propose two additional complexity-reduction methods which make use of the domination concept and the procedure described

in this section. They can be combined and used together as Step 1 in the heuristic approach for the variable routing problem variants.

Method 1: Reduce to the maximal ‘non-dominated set’, i.e. Reduce the original set of matrices to the maximal subset in which no matrix weakly dominates any other.

This method eliminates a subset of matrices from the original sequence of traffic matrices $S=\{D^1, \dots, D^T\}$ in such a way that the reduced sequence yields an *exactly* equivalent solution to the MH-VTD-VR problem, i.e. the optimal solution remains the same. A matrix D^i is eliminated from the sequence, if it is weakly dominated by other matrix $D^{i'}$ in the sequence. The weak domination relation states that any virtual topology that supports $D^{i'}$ also supports D^i , and thus the set of feasible solutions to the MH-VTD-VR problem is not modified by its removal. Reducing to a maximal non-dominated set can be done by solving a series of small linear formulations where the domination between each pair of matrices is checked. This involves solving a flow routing problem in the MH-VTD-VRs case, and its unsplittable version in the MH-VTD-VRu case.

Method 2: Hierarchical calculation of the weakly dominating matrix.

This method finds a matrix D which weakly dominates sequence $S^1=\{D^1, \dots, D^T\}$ as follows. First, the matrices in S^1 are partitioned into groups of two matrices consecutively. Then, for each group of two matrices, a matrix that weakly dominates both is calculated. This is done by using the relaxed version of formulation (2). For the MH-VTD-VRu problem, weak domination with respect to unsplittable flows is used. The calculated matrices form a sequence S^2 comprised of $\text{ceil}(T/2)$ matrices (if T is odd, one matrix from S^1 is directly mapped to S^2). The process is repeated, reducing the sequence size by half in each repetition. In the last repetition, a single matrix is obtained. Every virtual topology U that supports this matrix, also supports the matrices in the original sequence S^1 , and in all intermediate sequences. The method is suboptimal in the sense that it can not guarantee that other

matrices D^i which dominate the original sequence cannot be found, which lead to more cost-effective virtual topologies. However, since there is no guarantee that the weakly dominating matrix solved using formulation (2) for the original sequence yields the optimal solution to the MH-VTD-VR, the solution obtained with this approximation technique is not necessarily worse. The complexity reduction obtained by applying this method is quite relevant, as the relaxed version of formulation (2) is executed for instances of only 2 matrices. The number of instances solved for a sequence of $T=2^k$ matrices, equals $2^{k-1}+2^{k-2}+\dots+2^0=2^k-1=T-1$.

4.4 Step 2: Virtual topology design calculation

In Step 2, the dominating matrix D calculated in the previous step is used to feed an algorithm that calculates a virtual topology which supports traffic D . In this paper, we have chosen the heuristic algorithm presented in [9] which sub-optimally minimizes the number of transceivers. It starts with a fully connected virtual topology where all the traffic is carried between node pairs in one direct virtual hop. If the traffic between a node pair exceeds the capacity of a lightpath, multiple lightpaths between the same pair of nodes are assumed. The algorithm iteratively tries to reroute all the traffic on the least loaded lightpath onto its shortest path using the spare bandwidth available in the current virtual topology. If it succeeds, the lightpath under consideration is eliminated from the virtual topology and the corresponding traffic is rerouted. The process is repeated until the traffic in the least loaded lightpath of some iteration can not be rerouted using the available capacity. While the algorithm described in [9] considers only unsplittable flows, we have made a modification which allows us to consider also a splittable routing version. The modification consists of solving a fractional min-cost flow problem for rerouting the traffic on the least loaded lightpath, instead of a shortest-path approach. The algorithm from [9] is used for the unsplittable problem variants MH-VTD-FRu and MH-VTD-VRu, while the adaptation to fractional routing of [9] is used in the splittable problem variants.

In the MH-VTD-FR fixed routing problems, the method guarantees that the common virtual topology U and the common flow routing F obtained in this Step, are a feasible solution for the planning problem. In the MH-VTD-VR variable routing problem variants, a further step is needed.

4.5 Step 3. Variable flow routing calculation

In the MH-VTD-VR variable routing problems, the weak domination relation applied in the Replace-with-1 step, does not guarantee that the (U, F) pair calculated in Step 2, supports the traffic in all time slots. However, the method guarantees that for the obtained virtual topology U , there exists at least one routing F^t for each time slot $t=1, \dots, T$, which supports the traffic in that time slot.

Step 3 is responsible of finding such a set of T flow routings F^1, \dots, F^T of interest, where each pair (U, F^t) supports traffic matrix D^t , $t=1, \dots, T$. If changes in the flow routing along time are not penalized, this can be done by solving T independent conventional multicommodity flow problems. If not, the flow routing problems in different time slots become dependent, yielding to a more complex problem. Some suitable heuristics have been proposed in the literature which could be then used in this Step [13],[14].

4.6 Summary

To summarize, Table 1 collects the main aspects that define the family of heuristics proposed for the MH-VTD problem variants.

Table 1 Methodology Of The 3-Step Heuristic Approach According to the Planning Problem

Planning Problem	Step 1	Step 2	Step 3
MH-VTD-VRVRs	Weak domination	Splittable variant of [9]. Obtain U .	Required to obtain the flow routings F^t , $t=1, \dots, T$
MH-VTD-VRVRu	Weak domination (unsplittable)	[9]. Obtain U .	Required to obtain the flow routings F^t , $t=1, \dots, T$
MH-VTD-FRs	Total domination	Splittable variant of [9]. Obtain (U, F)	Not required
MH-VTD-FRu	Total domination	[2]. Obtain (U, F)	Not required

5 Results

This section presents exhaustive results obtained by applying the algorithms proposed in this paper to the four problem variants of interest. The algorithms were implemented in the MatPlanWDM tool [15] which links to the TOMLAB/CPLEX library [16] used to solve the MILP formulations.

5.1 Results for synthetic traffic

In a first step, both the MILP formulations and the heuristic approaches were tested for networks consisting of $N=\{4,6,8\}$ nodes, using traffic sequences generated synthetically. Each traffic sequence was composed of $T=12$ traffic matrices, each meant to describe the traffic of a 2-hour interval over the course of a day. Let $M(i, j, t)$ denote a traffic matrix sequence, where $M(i, j, t)$ is the traffic in Gbps from node i to node j during time interval t , $i, j=1, \dots, N$, $t=1, \dots, T$. Equation (3) describes the traffic synthesis model:

$$M_1(i, j, t) = M_{\text{Base}}(i, j) \cdot \text{activity}(t) \cdot \text{rf}(R), \forall i, j, t.$$

$$M(i, j, t) = M_1(i, j, t) \cdot \text{nf}(\rho), \forall i, j, t \quad (3)$$

where $\text{nf}(\rho) = \frac{N \cdot (N-1) \cdot \rho \cdot C}{\max_t \sum_{i,j} M(i, j, t)}$

Initially, the traffic between two nodes at a given time, was calculated as the product of 3 factors. Factor $M_{\text{Base}}(i, j)$ gives the (i, j) coordinate of a base traffic matrix computed for the sequence as follows. 80% of the values in M_{base} (randomly chosen) were set to 1, while the remaining 20% were set to 2. This is meant to capture the effect of non-uniformities in the generated traffic matrices. Activity factor $\text{activity}(t)$ in equation (3) intends to capture the effect of traffic intensity variation along the day. Our intensity variation scheme is described by equation (4), based on the intensity model presented in [17].

$$\text{activity}(t) = \begin{cases} 0.1 & \text{if } t \in [1, 6] \\ 1 - 0.9 \cdot \left(\cos\left(\frac{\text{mod}(t, T) - 6}{18} \cdot \pi\right) \right)^{10} & \text{otherwise} \end{cases} \quad (4)$$

where $t = 1, \dots, T$

Factor $rf(R)$ in (3) is a random value, uniformly distributed over interval $[1-R, 1+R]$. A new independent sample of the $rf(R)$ factor was used for each value $M(i, j, t)$. The purpose of the rf factor is to include a randomness effect in the traffic intensity. The random factors used were $R = \{0.1, 0.2, 0.5\}$, which correspond to low, medium and high random variation scenarios.

After the full sequence $M_1(i, j, t)$ was calculated, a traffic normalization stage followed. All the traffic values in the sequence were multiplied by a normalization factor (nf). Value nf was calculated such that the average traffic between two nodes which occurs in the highest loaded time slot equals ρC , where C is the lightpath capacity. Value ρ represents the traffic load of the sequence normalized with respect to the lightpath capacity. The values tested in our study were $\rho = \{0.1, 1, 10\}$. A value of $\rho = 0.1$ corresponds to the case when the average traffic between two nodes in the highest loaded time slot fills only 10% 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 10 whole lightpaths.

For each network size (N), randomness factor (R) and normalized load (ρ), we generated five independent sequences of traffic matrices, and solved all four variants of the MH-VTD problem. Table 2 shows the average number of transceivers used in the solutions obtained for each of the 5 sequences, using both the exact MILP formulation described in Section 3, and the heuristic methods presented in Section 4. For the exact MILP formulations, the solver was configured to find a solution within a 5% of optimality gap due to its high complexity. In some cases, this optimality gap has produced solutions in which the cost of the MILP was higher than the solution obtained by the heuristics, or solutions in which the cost of the more constrained case (e.g. unsplittable and/or fixed routing) was found to be better than the less constrained case. Note, however, that these situations are a symptom of the small cost differences between both solutions.

For the heuristic methods in the MH-VTD-VR variants, we calculated the matrix which weakly

dominates the sequence by applying the two complexity reduction methods described in Subsection (4.3.2: first we found the maximal non-dominated set, and then applied the proposed hierarchical scheme to find a matrix which weakly dominates all matrices in the reduced set.

From the obtained MILP solutions, we can extract some interesting conclusions regarding the effect of the different factors on the network cost. As a general rule, and not surprisingly, the number of transceivers used is higher for series of traffic matrices with higher randomness factor R , higher normalized loads ρ and a higher number of nodes N .

With respect to the reduction in the number of transceivers obtained by allowing dynamic reconfiguration of flows, in contrast to the fixed routing case, the following conclusions can be made: (i) the reduction is negligible in the unsplittable routing case, (ii) in the splittable routing case, this reduction is small (a maximum of ~13% at higher loads and higher randomness factors).

Analyzing the same data obtained from the MILP solutions, we can compare splittable vs. unsplittable flow routing. For fixed routing, the cost reduction obtained by allowing the routing to be splittable was only relevant at medium loads $\rho = 1$, increasing for higher randomness factors. The maximum reduction observed was of ~16%. For variable routing, the same trend was found, but with more intense cost reductions at medium loads (which reached up to ~22%) and non-negligible reductions at higher loads (up to 14%).

Comparison of the variable-splittable routing case, in contrast to the fixed-unsplittable routing case, indicates a significant cost reduction only at medium loads (up to ~25%) and high loads (up to ~17%).

When we compare the results obtained from the exact MILP formulations with the ones obtained by the heuristic algorithms, we can see that, in almost all the cases, the results are fairly similar. The cost gap between the heuristic and the MILP solutions has a significant value (up to 40%) only at low traffic loads. Note that in some cases, the heuristic method gave slightly better solutions for the fixed routing and/or unsplittable case than for the less

constrained variable and/or splittable routing. These situations are caused by the inherent inexactitudes in the heuristics

The LB column in Table 2 shows a lower bound on the number of transceivers in the optimal solution calculated as follows. Given a traffic sequence, the minimum number of transmitters (receivers), LB_{TR} (LB_{RE}), that a node n requires is at least the number of lightpaths needed to transmit (receive) the ingress (egress) traffic to the node, in its worst time slot.

$$LB_{TR}(n) = \max_{t=1,\dots,T} \left\{ \sum_{j=1}^N \left\lceil \frac{M(n,j,t)}{C} \right\rceil \right\} \quad (5a)$$

$$LB_{RE}(n) = \max_{t=1,\dots,T} \left\{ \sum_{i=1}^N \left\lceil \frac{M(i,n,t)}{C} \right\rceil \right\} \quad (5b)$$

It follows that a lower bound on the number of transceivers in the network is given by:

$$LB = \sum_{n=1}^N \{LB_{TR}(n) + LB_{RE}(n)\} \quad (6)$$

It is very important to note that the lower bound described above is a lower bound for the *fully reconfigurable* case in which (i) the virtual topology can be reconfigured along time, (ii) the flow routing is variable, and (iii) the flow routing is splittable. Therefore, for cases where the lower bounds are close to the MH-VTD solutions, we can claim that only a small reduction in the number of transceivers can be obtained by allowing the virtual topology to be dynamically reconfigured along time.

The rightmost column in Table 2 shows the maximal reduction in the number of transceivers which could be obtained, in the variable and splittable flow routing case, if the virtual topology could change along time. In other words, this column shows the maximal transceiver reduction that could be obtained by using reconfigurable equipment over its non-reconfigurable variant. Interestingly enough, results show that only a minor reduction (a maximum of ~7%) can be achieved in all the scenarios tested, except for those with a low load and $N=\{6,8\}$ nodes. In these cases the maximal cost reduction have shown to be high

(~33% for 6 nodes, and 65% for 8 nodes) and independent of the traffic randomness factor.

5.2 A realistic case study: the Abilene network

In order to further validate the results and conclusions presented in this paper, we performed a case study on realistic (not synthetically generated) traffic. To do so, we chose the 11-node Abilene network for which a real traffic trace is publicly available. The traffic matrices used were taken from the measures carried out in the TOTEM project (Toolbox for Traffic Engineering Methods) **¡Error! No se encuentra el origen de la referencia.** The sequence of matrices available in **¡Error! No se encuentra el origen de la referencia.** spans several weeks. From this data, we averaged the values taken at the same time and day in the week to obtain a sequence representing the average week in 15-minute time intervals (i.e., 96 matrices per day, 672 matrices in total). In order to test different traffic intensities, the average week sequence was normalized for values $\rho=\{0.1, 1, 10\}$, in the same manner as in the study on synthetic traffic described above.

Naturally, the exact MILP formulations for such large problem sizes are intractable. Therefore, all tests were conducted using the 3-step heuristic approach proposed in this paper, along with the two complexity reduction stages described to simplify the first step of the algorithm. First, we applied the reduction to obtain the maximal non-dominated set. This reduced the number of matrices in the set from 672 to 172 in the splittable case, and from 672 to 479 in the unsplittable one. Then, the hierarchical version of the replace-to-1 step was applied to the maximal non-dominated sets.

Table 3 shows the obtained results for all problem variants. Unfortunately, some parts of the analysis made in the previous section can not be repeated here. Namely, the results obtained by the heuristic algorithm for the unsplittable and variable routing case show a higher cost than the ones for the fixed routing case. This comes from the accumulated suboptimality that appear in the hierarchical reduction method which become very significant with such a large number of matrices.

The analysis of the rest of the data confirms the main trends observed for synthetic traffic, although the cost reductions achieved by allowing variable routing and/or splittable routing are somewhat smaller. Firstly, we can see that for splittable routing, the reduction in the number of transceivers obtained by using variable routing, as opposed to fixed, is minor (a maximal reduction of $\sim 7\%$ was achieved for medium and high loads). In the fixed routing case, the cost reduction obtained by using splittable instead of unsplittable routing, is also minor (with a maximal reduction of $\sim 8\%$ at medium loads). Furthermore, if we compare the case of variable-splittable routing with the fixed-unsplittable case, the results show only a moderate cost reduction, with a maximum of $\sim 15\%$ achieved at medium loads.

As before, comparison of the MH-VTD results with the lower bounds indicates the maximal possible reduction in the number of transceivers that could be achieved by using reconfigurable optical equipment, i.e. allowing dynamic reconfiguration of the virtual topology. This maximal benefit was shown to be approximately 40%, 15% and 10% for low, medium and high loads, respectively.

6 Conclusions and future work

This paper studies the problem of finding a static virtual topology in transparent optical networks for a known time-varying traffic demand. Four problem variants are considered assuming variable/fixed flow routing with splittable/unsplittable flows. Exact MILP formulations, as well as a family of 3-step heuristics, are proposed for each problem variant. The heuristics apply the concept of traffic domination, which provides a theoretical basis allowing us to devise a set of complexity reduction techniques for the problem. Extensive results using both synthetic traffic and realistic traffic traces are presented.

Results support interesting conclusions regarding the cost reduction (in terms of the number of transceivers used) that network operators can achieve in different circumstances. For instance, in the scenarios studied the migration from fixed to variable routing is justified only if splittable routing

is also allowed in the network, and only for medium or high loads. However, the maximum reduction obtained is still moderate ($\sim 15\%$) and is a trade-off with increased overhead. Analogously, for network operators evaluating the advantages of splittable routing in contrast to unsplittable routing, our results indicate that the migration is cost-effective for medium load scenarios, with a larger cost reduction if variable routing is allowed. Still, the reductions are moderate (with a maximum of $\sim 20\%$).

A lower bound to the network cost has been provided for the case in which the network uses variable splittable routing, and the virtual topology can be dynamically reconfigured along time (i.e., uses reconfigurable optical equipment). The comparison of the lower bounds with our results obtained in the MH-VTD-VRs problem, allows us to measure the maximal benefit that can be achieved in the network with variable and splittable routing if it were to also permit dynamic reconfiguration of the virtual topology. The results show that the maximum cost reduction is not significant except in the low load scenario, where it can reach values of $\sim 40\%$ (in the Abilene case study) and up to 65% (in the studies with synthetic traffic). Nevertheless, it should be noted that in the low load scenario, as the number of lightpaths in the network is low, the reduction in the number of transmitters in absolute value is also small. Therefore, it is still a matter of study whether the increase in cost caused by using reconfigurable optical equipment, at all or some nodes in the network, as well as the extra signaling complexity, could be compensated for by the cost reduction achieved by using fewer transceivers. For future work, we will further investigate the trade-offs associated with replacing non-reconfigurable equipment with their reconfigurable counterparts, as well as continue our studies concerning the application of traffic domination to optical networks planning.

Table 3 Abilene network. Total Number of Transceivers

ρ	ST-VTD-VR		ST-VTD-FR		LB	Cost gap
	Split.	Unsplit.	Split.	Unsplit.		
0.1	62	68	64	68	38	38.7%
1	336	402	358	392	286	14.9%
10	3046	3556	3292	3360	2739	10.1%

Acknowledgments 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 has been partially supported by the projects from the Spanish Ministry Of Education TEC2007-67966-01/TCM CON-PARTE-1, and TEC2008-02552-E, and it is also developed in the framework of the projects from *Fundación Seneca* (Regional Agency of Science and Technology of Region of Murcia) 00002/CS/08 (FORMA) and "Programa de Ayudas a Grupos de Excelencia de la Región. de Murcia", F. Séneca (*Plan Regional de Ciencia y Tecnología 2007/2010*).".

References

- [1] TOTEM Project: Toolbox for Traffic Engineering Methods. <http://totem.run.montefiore.ulg.ac.be/datatools.html> Accessed 29 June 2009
- [2] Konda, V. R., & Chow, T. Y. (2001): Algorithm for traffic grooming in optical networks to minimize the number of transceivers. *Proceeding of IEEE Workshop on High Performance Switching and Routing*, pp. 218-221, Dallas, TX, USA.
- [3] Zhu, K., Zhu, H., Mukherjee, B.(2003): Traffic engineering in multigranularity heterogeneous optical WDM mesh networks through dynamic traffic grooming. *IEEE Network*. **17**(2), 8 - 15.
- [4] Srinivas, N., & Siva Ram Murthy, C.,(2004): Design and dimensioning of a WDM mesh network to groom dynamically varying traffic. *Photonic Network Communications*. **7**(2), 179 - 9.
- [5] Mukherjee, B., Banerjee, D., Ramamurthy, S., Mukherjee, A. (1996): Some Principles for Designing a Wide-Area WDM Optical Network. *IEEE/ACM Transactions on Networking*. **4**(5), 684-696.
- [6] Oriolo, G. (2008): Domination Between Traffic Matrices. *Mathematics of Operations Research*. **33**(1), 91-96.
- [7] Internet 2 (Former Abilene) Global Research Network Operations Center. <http://www.abilene.iu.edu/> Accessed 29 June 2009
- [8] Pioro, M., Medhi, D. (2004): Multi-Hour and Multi-Time-Period Network Modelling and Design. In *Routing, Flow, and Capacity Design in Communication and Computer Networks* (pp. 455-474). San Francisco, CA, USA: Morgan Kaufmann.
- [9] Agrawal, G., Medhi, D. (2006): Lightpath Topology Configuration for Wavelength-routed IP/MPLS Networks with Time-Dependent Traffic, *Proceedings of IEEE Globecom'06*, pp. 1-5. San Francisco, CA.
- [10] Ricciato, F., Salsano, S., Belmonte, A., Listanti, M. (2002): Offline configuration of a MPLS over WDM network under Time-varying Offered Traffic. *Proceedings of Infocom'02*, pp.57-65, New York.
- [11] Garcia-Manrubia, B., Aparicio-Pardo, R., Pavon-Mariño, P., Skorin-Kapov, N., Garcia-Haro, J. (2009): MILP Formulations for Scheduling Lightpaths under Periodic Traffic. *Proceedings of ICTON'09* (invited paper), Azores, Portugal.
- [12] Skorin-Kapov, N., Pavon-Mariño, P., Garcia-Manrubia, B., Aparicio-Pardo, R., (2009): Scheduled Virtual Topology Design Under Periodic Traffic in Transparent Optical Networks, *Proceedings of Broadnets'09*, Madrid, Spain. (*accepted for publication*)
- [13] Wang, B., Li, T., Luo, X., Fan, Y., Xin, C. (2005): On Service Provisioning under a Scheduled Traffic Model in Reconfigurable WDM Optical Networks, *Proceedings. of IEEE Broadnets'05*, pp. 15 -24, Boston, MA, USA.
- [14] Aggarwal, A., Bar-Noy, A., Coppersmith, D., Ramaswami, R., Schieber, B., Sudan, M. (1994): Efficient routing and scheduling algorithms for optical networks, *Proceedings of 5th annual ACM-SIAM symposium on Discrete algorithms*, pp. 412-423, Arlington, VA, USA.
- [15] Pavon-Marino, P., Aparicio-Pardo, R., Moreno-Muñoz, G., Garcia-Haro, J., Veiga-Gontan, J., (2007): MatPlanWDM: An educational tool for network planning in wavelength-routing networks, *Proceeding of IFIP Int. Conf. Optical Network Design and Modeling*, pp. 58-67, Athens, Greece.
- [16] TOMLAB Optimization. <http://tomopt.com/> . Accessed 29 June 2009
- [17] Milbrandt, J., Menth, M., Kopf, S. (2005): Adaptive Bandwidth Allocation: Impact of Traffic Demand Models for Wide Area Networks, *Proceedings of 19th International Teletraffic Congress (ITC19)*, Beijing, China.

Table 2 Synthetic traffic. Number of transceivers in the network in the tested scenarios

N	ρ	R	MH-VTD-VR				MH-VTD-FR				LBs	Cost gap
			Split.		Unsplit.		Split.		Unsplit.			
			Heur.	Exact MILP	Heur.	Exact MILP	Heur.	Exact MILP	Heur.	Exact MILP		
4	0.1	0.1	10	8	9.6	8	9.6	8	9.6	8	8	0.0%
		0.2	9.6	8	10	8	10.8	8	10	8	8	0.0%
		0.5	10.4	8	9.6	8	9.2	8	9.2	8	8	0.0%
	1	0.1	28	28	30.4	30.4	28.4	28	30.4	30.4	28	0.0%
		0.2	30	28.4	40.4	34.8	34	30.8	40.4	34.8	28.2	0.7%
		0.5	34.8	32.8	44	38	36.8	33.6	46.8	38	31.4	4.5%
	10	0.1	258.4	255.2	270.4	270.4	264	262.8	270.4	270.4	253	0.9%
		0.2	272.8	264.8	294	294	288.4	285.2	294	294	260	1.9%
		0.5	295.2	283.6	319.2	317.2	319.6	309.6	326.8	317.2	274	3.5%
6	0.1	0.1	18.8	16	19.6	16	19.6	16	19.6	16	12	33.3%
		0.2	19.2	16	19.6	16	18.8	16	19.6	16	12	33.3%
		0.5	20.8	16	22	16	22	16	22	16	12	33.3%
	1	0.1	72	71.2	75.6	75.2	72	72	75.6	75.2	70.6	0.9%
		0.2	76.4	71.6	96.8	87.2	80.4	76	96.8	88	71	0.9%
		0.5	85.2	76	111.6	96.4	95.2	84.8	116.8	102	74.2	2.4%
	10	0.1	646.8	635.6	686	670	668.4	670.8	686	670.8	627	1.4%
		0.2	686.4	662.4	747.2	747.2	728	726	747.2	747.2	647	2.4%
		0.5	753.6	716.4	847.2	834	849.2	817.6	869.6	868	693.4	3.3%
8	0.1	0.1	32.8	26.4	34	26.4	34.8	26.4	34	26.8	16	65.0%
		0.2	33.2	26	34	26.4	34.8	26	34	26.4	16	62.5%
		0.5	34.8	26.8	38.4	27	38.4	26.8	40	27.6	16.6	61.5%
	1	0.1	131.6	127.2	145.2	140	134	132.8	145.2	142.4	123	3.4%
		0.2	136.4	130	177.6	158	148.4	141.6	177.6	160	127.6	1.9%
		0.5	155.6	140.4	200.8	181	179.2	159.2	225.2	188.8	134.8	4.2%
	10	0.1	1209.2	1224.8	1301.6	1310.2	1260.8	1261	1301.6	1319.6	1166.8	5.0%
		0.2	1264	1262.8	1397.6	1390	1359.2	1359	1396.8	1401.2	1196	5.6%
		0.5	1404	1368.8	1536.4	1540.4	1627.6	1586	1665.6	1651.2	1275.6	7.3%