

Offline Impairment-Aware Routing and Wavelength Assignment Algorithms in Translucent WDM Optical Networks

Konstantinos Manousakis, Konstantinos Christodoulopoulos, Euaggelos Kamitsas, Ioannis Tomkos, and Emmanouel A. Varvarigos

Abstract—Physical impairments in optical fiber transmission necessitate the use of regeneration at certain intermediate nodes, at least for certain lengthy lightpaths. We design and implement impairment-aware algorithms for routing and wavelength assignment (IA-RWA) in translucent optical networks. We focus on the offline version of the problem, where we are given a network topology, the number of available wavelengths and a traffic matrix. The proposed algorithm selects the 3R regeneration sites and the number of regenerators that need to be deployed on these sites, solving the regenerator placement problem for the given set of requested connections. The problem can be also posed in a slightly different setting, where a (sparse) placement of regenerators in the network is given as input and the algorithm selects which of the available regenerators to use, solving the regenerator assignment problem. We formulate the problem of regenerator placement and regenerator assignment, as a virtual topology design problem, and address it using various algorithms, ranging from a series of integer linear programming (ILP) formulations to simple greedy heuristic algorithms. Once the sequence of regenerators to be used by the non-transparent connections has been determined, we transform the initial traffic matrix by replacing non-transparent connections with a sequence of transparent connections that terminate and begin at the specified 3R intermediate nodes. Using the transformed matrix we then apply an IA-RWA algorithm designed for transparent (as opposed to translucent) networks to route the traffic. Blocked connections are re-routed using any remaining regenerator(s) in the last phase of the algorithm.

Index Terms—Physical layer impairments, regenerator assignment, regenerator placement, routing and wavelength assignment, translucent versus transparent networks, virtual topology.

I. INTRODUCTION

IN A Wavelength Division Multiplexing (WDM) network, connections are established concurrently through a common set of fibers, subject to the *distinct wavelength assignment* constraint, which states that connections sharing a

fiber must occupy different wavelengths. The most common architecture utilized for establishing communication in WDM optical networks is *wavelength routing* [1], where data are transmitted through all-optical WDM channels, called *lightpaths*, which may span multiple consecutive fibers. Given a set of connection requests, the routing and wavelength assignment (RWA) problem consists of selecting an appropriate path and a wavelength for each connection so as to minimize the use of network resources or maximize the traffic served [2].

The RWA problem is usually considered assuming two alternative traffic models. When the set of connections requests is known in advance the problem is referred to as *offline* or *static RWA*, while when the connections arrive at random times, over an infinite time horizon, and are served on a one-by-one basis, the problem is referred to as *online* or *dynamic RWA*. We will focus our study on offline RWA, which is known to be an NP-hard optimization problem [2]. The offline RWA problem is quite different from the online RWA problem. Although online algorithms can be used to solve the offline problem (by sequentially considering each connection in the given set of requests), this approach does not jointly optimize the solution for all the connections. Thus, combinatorial optimization algorithms, like the ones used in the current paper, are usually employed for offline problems.

The majority of offline RWA algorithms proposed in the literature assume an ideal physical layer where signal transmission is error free. However, signal transmission is significantly affected by physical limitations of fibers and optical components [3], including amplified spontaneous emission noise (ASE), chromatic dispersion (CD), polarization mode dispersion (PMD), filter concatenation (FC) and nonlinear effects such as self- and cross-phase modulation (SPM, XPM), and four-wave-mixing (FWM), etc. All these impairments degrade the received signal quality, so that the bit-error rate (BER) at the receiver may be too high for some lightpaths. For the remainder of this paper we will refer to such a phenomenon as physical-layer blocking, as opposed to the network-layer blocking that is due to the unavailability of an adequate number of wavelengths.

At present, and in the foreseeable future, the only satisfactory method to overcome these impairments is re-amplifying, re-shaping and re-timing of optical pulses (referred to as 3R regeneration). In opaque networks the signal is regenerated at every intermediate node along a lightpath via OEO conversion. The network cost could be reduced by employing regenerators

Manuscript received December 03, 2008; revised April 01, 2009. This work has been supported in part by the European Commission under the FP7-DI-COINET project. Current version published June 05, 2009.

K. Manousakis, K. Christodoulopoulos, E. Kamitsas, and E. A. Varvarigos are with the Computer Engineering and Informatics Department, University of Patras, Patras, 26500 Greece, and also with the Research Academic Computer Technology Institute, Patras, Greece 26500 (e-mail: manousak; kchristodou; kamitsas; manos@ceid.upatras.gr).

I. Tomkos is with the Athens Information Technology Center, Peania, Greece (e-mail: itom@ait.edu.gr).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JLT.2009.2021534

only at specific nodes instead of at all the nodes. When regenerators are available, a lengthy end-to-end connection that needs regeneration at some intermediate node(s) is set up in a multi-segment manner so that it is served by two or more consecutive lightpath-segments. The regenerator at the end of each lightpath-segment serves as a “refueling station” that restores signal quality. This type of optical networks, where some lightpaths are routed transparently, while others have to go through a sequence of 3R regenerators, are known as *translucent* optical networks [4]. The ultimate goal is the development of an all-optical *transparent* network, where the data signal remains in the optical domain for the entire path. Since this may not be feasible when the lengths of the links and the size of the network are large, translucent networks seem a more appropriate solution at present and in the near future.

In the current paper we study static IA-RWA algorithms for heterogeneous translucent optical networks. In such networks apart from the RWA problem there is also the additional problem of choosing the connections that have to be served using regenerators and the sequence of regenerators these connections are going to use. The network topology, the number of available wavelengths and a traffic matrix form the inputs to the algorithm. We identify the connections that are going to use regenerators and select the 3R regeneration sites and the number of regenerators that need to be deployed on these sites, a problem that is usually referred to as the *regenerator placement* problem. If the input is given in a slightly different setting where a sparse placement of regenerators is given and known, then the algorithm selects which of these regenerators to use, a problem that is usually referred to as the *regenerator assignment* problem. Note that the latter case is equivalent to the case where we are given the number of transmitter-receiver pairs (add-drop ports) of all the switches in the network, since unused transceivers can be employed back-to-back so as to function as regenerators.

We decompose the problem in the regenerator placement/assignment subproblem and the RWA subproblem. We propose an IA-RWA algorithm for translucent networks that consists of three phases. In the first phase, we find how to route non-transparent connections through 3R regenerators, so as to transform them into a sequence of transparent connections (that is, connections that do not use regenerators). To do so, we formulate a virtual topology problem and propose a series of algorithms for routing the non-transparent traffic demands. In particular, we propose: (i) some computationally expensive algorithms where the virtual topology problem is formulated as an *integer linear program* (ILP), and (ii) four *greedy heuristic* algorithms, corresponding to different choices of the link costs in the virtual topology. By the end of the first phase the initial traffic matrix has been transformed into a new traffic matrix whose source-destination pairs can, in principle, all be transparently connected. In the second phase, we apply an impairment-aware (IA-) RWA algorithm for transparent networks, with input the transformed transparent traffic matrix, in order to select the routes and wavelengths to be used. Finally, in the third phase of the algorithm, we reroute through the remaining (unused in the first phase) regenerators any connections that were rejected in the second phase due to physical-layer blocking. For assessing the Quality of Transmission (QoT) of lightpaths we use a Q

factor estimator that models through analytical formulas all the well known impairments. This estimator is used in the first phase of the algorithm to distinguish connection demands into transparent and non-transparent ones, and also at the end of the second phase of the algorithm to evaluate the feasibility of the IA-RWA solution and determine the paths that have to be handled again, in the third phase of the algorithm.

We evaluated the applicability of the proposed algorithms using simulation experiments for realistic topologies and traffic matrices. Our performance results indicate that the proposed IA-RWA algorithms for translucent networks can provide efficient solutions to medium- and even large-scale networks. The execution time of the algorithm is dominated by the IA-RWA algorithm applied in the second phase. Regarding the different algorithms investigated for the first phase, we observed that the ILP formulation that minimizes the maximum number of used regenerators among the regeneration nodes exhibits the best performance in terms of the number of wavelengths that it requires to reach zero blocking. The ILP formulations that minimize the total number of regenerators in the network or the number of regeneration sites have worse wavelength performance, but optimize regenerator usage characteristics, which, depending on the objective, can also yield cost efficient solutions. The proposed heuristic algorithms provide fast solutions to the virtual topology problem with wavelength and regenerator utilization performance that depends on the type of the problem considered and the cost of the virtual links utilized, but is in all cases satisfactory.

The rest of this paper is organized as follows. In Section II we report on related work. In Section III we formally define the problem of offline RWA in translucent networks and describe our proposed algorithms. Performance results obtained under realistic traffic scenarios are presented in Section IV. Our conclusions follow in Section V.

II. PREVIOUS WORK

Previous studies address the RWA problem for translucent optical networks by dividing a large-scale optical network into several *islands of transparency* or *optically transparent domains*. Within the same island, a lightpath can transparently reach any other node without intermediate signal regeneration. For communication across islands, regeneration nodes at the island boundaries [5] perform 3R on connections crossing multiple island boundaries. An alternative approach, which is the one used in the present paper, is called *sparse placement of regenerators*, where some selected nodes in the network are used for regeneration and there is no explicit notion of transparent domains.

The majority of RWA algorithms proposed so far for translucent networks assume a dynamic traffic scenario, that is, they address the online version of the problem. [6] presents a two-dimensional Dijkstra RWA algorithm for translucent optical network, that takes into account the placement of regenerators and a constraint on the maximum transparent distance. When the length of a lightpath exceeds a maximum transparent distance bound, the lightpath is blocked. A different approach for dynamic resource allocation and routing is considered in [7] and [8], where spare transceivers (transmitter-receiver pairs

or add-drop ports) at the nodes are used to regenerate signals. In this kind of network, every node with a *spare transceiver* can become a *potential regenerator*. In networks where the lightpaths initiated and terminated at a node do not use up all its transmitters and receivers, many nodes will have spare transceivers available for regeneration purposes. A *Max-spare algorithm* for selecting the regeneration nodes for a lightpath is proposed in [9] and compared to a *Greedy* algorithm used in conjunction with a wavelength-weighted and a length-weighted RWA algorithm. In [10], two online RWA algorithms for translucent networks with sparse regenerator placement are presented. IA-RWA algorithms that assume worst-case physical transmission penalties corresponding to a fully loaded system or that take into account the current network status and the actual number of active channels are investigated.

In [11], the problem of maximizing the number of established connections, under a constraint on the maximum transparent length, is formulated as a *mixed-integer linear program* (MILP). Since MILP is NP-hard, the authors also propose a heuristic algorithm to establish connections. However, [11] does not consider impairment effects other than the transparent length. A carefully designed algorithm that uses the *sparse regeneration* technique to allocate regeneration resources in an efficient manner is given in [12]. In [13] the authors address the problem of translucent network design by proposing several regenerator placement algorithms based on different knowledge of future network traffic patterns. A quality of transmission based heuristic algorithm for IA-RWA in translucent optical networks is presented in [14]. In the first phase of that algorithm a random search heuristic RWA algorithm is used and in the second phase regeneration placement is performed after estimating the BER of the lightpaths comprising the solution of the first phase.

The main contribution of the current paper is a series of offline IA-RWA algorithms for heterogeneous translucent WDM networks. Our algorithms transform the problem of routing non-transparent connections (i.e., connections that require regenerators) into a virtual topology problem. The virtual topology problem is formulated as an ILP, which can be solved efficiently in most realistic cases considered, even for large optical networks, due to the small number of variables used (as compared to [11]). For instances where the ILP cannot be solved efficiently, an LP relaxation formulation with piecewise linear costs, together with iterative fixing and rounding techniques can be used [15]. Furthermore, we propose several heuristic polynomial-time online algorithms that yield fast solutions to the virtual topology problem.

The performance of our algorithms compares favourably to that of previously proposed offline algorithms with respect to accuracy and completeness of the solutions provided, as it uses the Q factor as a criterion of the received signal quality. It is not trustworthy for heterogeneous optical networks to decide on the feasibility (in terms of quality of transmission) of a lightpath based only on its length, which is the criterion used in most related works. Finally, the offline IA-RWA algorithm for transparent networks that we use in the second phase of the proposed algorithms [17] has been shown to exhibit superior performance to other approaches, being, to the best of our knowledge, the only offline scheme that accounts for the interference among

lightpaths [16]. Note that the algorithmic approach that is proposed in this paper does not depend on [17] and any other offline IA-RWA algorithm for transparent networks can be used in the second phase of the algorithm.

III. RWA IN TRANSLUCENT WDM NETWORKS

We consider a translucent WDM network with a known physical topology. The network supports a given number of wavelengths, and wavelength conversion is possible only at 3R nodes. We assume static traffic, given in the form of a traffic matrix that specifies the number of connections requested between any pair of nodes. Each connection requires transmission rate equal to that of one wavelength, but there may be multiple connection requests for a given source-destination pair.

The algorithms to be presented can work in two alternate settings. In the first setting, we assume that the algorithm's objective is to select the 3R regeneration sites and the number of regenerators to be deployed at these sites. In the second setting, we assume that the sites and number of regenerators are given in advance (this is usually referred to as a sparse regeneration placement), and the algorithm's objective is to select the regenerators to be used by the connections from the ones available. For the rest of this study we will refer to the former as the regenerator placement problem and to the latter as the regenerator assignment problem.

The proposed algorithm is given a specific RWA instance, described by the above parameters, and returns the instance solution, in the form of routed lightpaths and assigned wavelengths, including decisions about whether a connection will be served transparently or using regenerators. For the latter connections, the algorithm also decides the sequence of regenerators to be used. In case the algorithm cannot find a solution that serves all connections with the given number of available wavelengths, the algorithm returns the corresponding blocking ratio.

The proposed IA-RWA algorithms for translucent WDM networks consist of three phases. In phase 1, described in Section III-A, we initially decide which connections cannot be served transparently without the use of regeneration resources. For these non-transparent connection requests we formulate and solve a virtual topology design problem in order to decide the intermediate regeneration sites they should use. The virtual topology consists of the 3R sites with (virtual) links between any pair of transparently connected 3R sites. Each virtual link of the paths chosen in the virtual topology to serve a connection corresponds to a sub-connection (lightpath) in the physical topology (Fig. 1). In other words, a non-transparent connection is served using two or more consecutive segments (lightpaths), defined by the intermediate 3R sites it is going to utilize. Based on these decisions, we transform the *initial traffic matrix* into a *transparent traffic matrix*, where non-transparent connection requests are replaced by a sequence of transparent sub-connection requests.

In phase 2 of the algorithm, described in Section III-B, we apply an impairment-aware (IA-) RWA algorithm using as input the transformed transparent traffic matrix. The IA-RWA algorithm we used in our simulation experiments is the one proposed in [17], but other algorithms are also applicable. Note that

given an adequate number of wavelengths, the IA-RWA algorithm of phase 2 would solve the problem without any physical layer blocking, since, by construction, all the connections in the transformed traffic matrix can be served transparently if handled alone or with adequate wavelength spacing between them. However, if, due to the limited number of wavelengths, some connections cannot be served while there are still regenerators in the network that were not used in phase 1, in phase 2 of the algorithm, described in Section III-C, these physically-blocked connections are reattempted through the remaining regenerators, by repeating phase 2 of the algorithm. Fig. 2 provides an overview of the proposed IA-RWA algorithm for heterogeneous translucent networks.

A. Obtaining the Transparent Traffic Matrix (Phase 1)

Phase 1 of the algorithm aims at transforming the initial traffic matrix Λ into a *transparent traffic matrix* $\tilde{\Lambda}$ that consists of connection requests that can be served without regenerators. Towards this end, we formulate in Section III-A-I a virtual topology problem that considers only non-transparent connection requests. This problem is solved in Section III-A-II so as to obtain the intermediate regeneration sites these connections should use. The traffic matrix is then transformed into a transparent traffic matrix, as described in Section III-A-III.

The physical network topology is represented by a connected simple graph $G = (V, E)$, where V is the set of nodes with routing capabilities (switches) and $E \setminus \psi$ is the set of (point-to-point) single-fiber links. We assume that 3R regenerators can act as wavelength converters, and no other nodes are equipped with wavelength conversion capabilities. Each fiber supports a common set $C = \{1, 2, \dots, W\}$ of W distinct wavelengths. The static version of RWA assumes an a-priori known traffic scenario given in the form of a $|V| \times |V|$ traffic matrix Λ of non-negative integers, where $|\cdot|$ denotes the cardinality of a set. Then, Λ_{sd} denotes the number of requested connections from source s to destination d .

We assume that 3R regenerators will be sparsely placed in the network, forming *pools* of regenerators at some nodes. We let $R \subseteq V$ be the set of nodes that can be equipped with at least one 3R regenerator and r_n be the number of regenerators available at node $n \in R$. In the version of the problem where the algorithm is free to select the regeneration sites and the number of regenerators to deploy (regenerator placement problem), we assume that $R = V$ and that r_n is unbounded. In the version of the problem where the placement of regenerators is given (regenerator assignment problem) the set R and the parameters r_n are part of the input to the algorithm. The same problem can be given in the slightly different setting where, instead of the sparse regenerator placement, we are given the number of available transceivers at each node. Given the total number of transceivers at a node, and subtracting those that are used by originating or terminating traffic (described in matrix Λ), we can find the number of spare transceivers at each node. These spare transceivers can be connected back-to-back so as to function as 3R regenerators, and, thus, we can transform this problem to the typical regeneration assignment problem.

1) *Constructing the Virtual Topology*: In order to formulate the virtual topology problem, we first distinguish between

transparent and non-transparent connection requests in the given traffic matrix. Specifically, for each source-destination (s, d) pair we examine if the quality of transmission of at least one of its k -shortest length paths has acceptable performance, as indicated by its corresponding Q-factor, in an otherwise empty network. More formally, we let $P_k(s, d)$ be the set of k -shortest length paths between nodes s and d . We will say that a source-destination pair (s, d) is transparently connected, and will denote that by $(s, d) \in T$, when the following condition holds:

$$(s, d) \in T \text{ if } Q_p > Q_{\min} + Q_{\text{margin}} \\ \text{for some path } p \in P_k(s, d),$$

where Q_{\min} is the Q-factor value that corresponds to the minimal transmission quality requirement, and Q_{margin} is a constant safety margin. We use k -shortest paths instead of a single path in the above definition, since in a heterogeneous network the shortest length path does not always yield the best quality of transmission. The set of connections that do not satisfy the above constraint will be referred as the set of non-transparent connections and will be denoted by \bar{T} .

Note that for a source destination pair $(s, d) \in T$, at least one of its k -shortest paths has acceptable Q performance, and it is possible *in principle* to serve a connection between s and d without the use of regenerators. However, this possibility depends on the IA-RWA solution, since the Q-factor value calculated and used in distinguishing between transparent and non-transparent connections assumes an empty network. When all connections are present, interference among lightpaths may make some of the assumed transparent lightpaths infeasible in reality. The IA-RWA algorithm applied in the phase 2 of the proposed algorithms, aims at avoiding such problems by finding lightpaths that are feasible even when all impairments are taken into account, and for any connections for which this is not possible, the re-routing performed in phase 3 attempts to find a new sequence of (unused) regenerators to serve them.

Given the set T of transparently connected (s, d) pairs and the set R of nodes with regeneration capabilities, we define the virtual topology as the graph $\tilde{G} = (R, \tilde{E})$, where $\tilde{E} = \{(u, v) | u \in R, v \in R, (u, v) \in T\}$, denotes the set of virtual links between regeneration nodes that are transparently connected. Each virtual link corresponds to a transparent lightpath between two regeneration nodes, spanning one or several physical links. We also denote by R^s and R^d the set of regeneration sites source s and destination d are transparently connected to, respectively.

2) *Choosing the Regenerators for the Non-Transparent Connections*: In this step of phase 1 we consider only the set \bar{T} of source-destination pairs that are *not* transparently connected and that have to be routed through regenerators. Therefore, for a source-destination pair $(s, d) \in \bar{T}$, we have to choose the sequence of regeneration site(s) they are going to use. This is equivalent to finding a path from s to d in the virtual topology $\tilde{G}(s, d)$, obtained by adding to graph \tilde{G} the source node s and the virtual links connecting it to the elements of R^s and, also, the destination node d and the virtual links connecting it to the elements of R^d . In order to select the virtual path from s to d (equivalently, the sequence of regenerators to be used), we

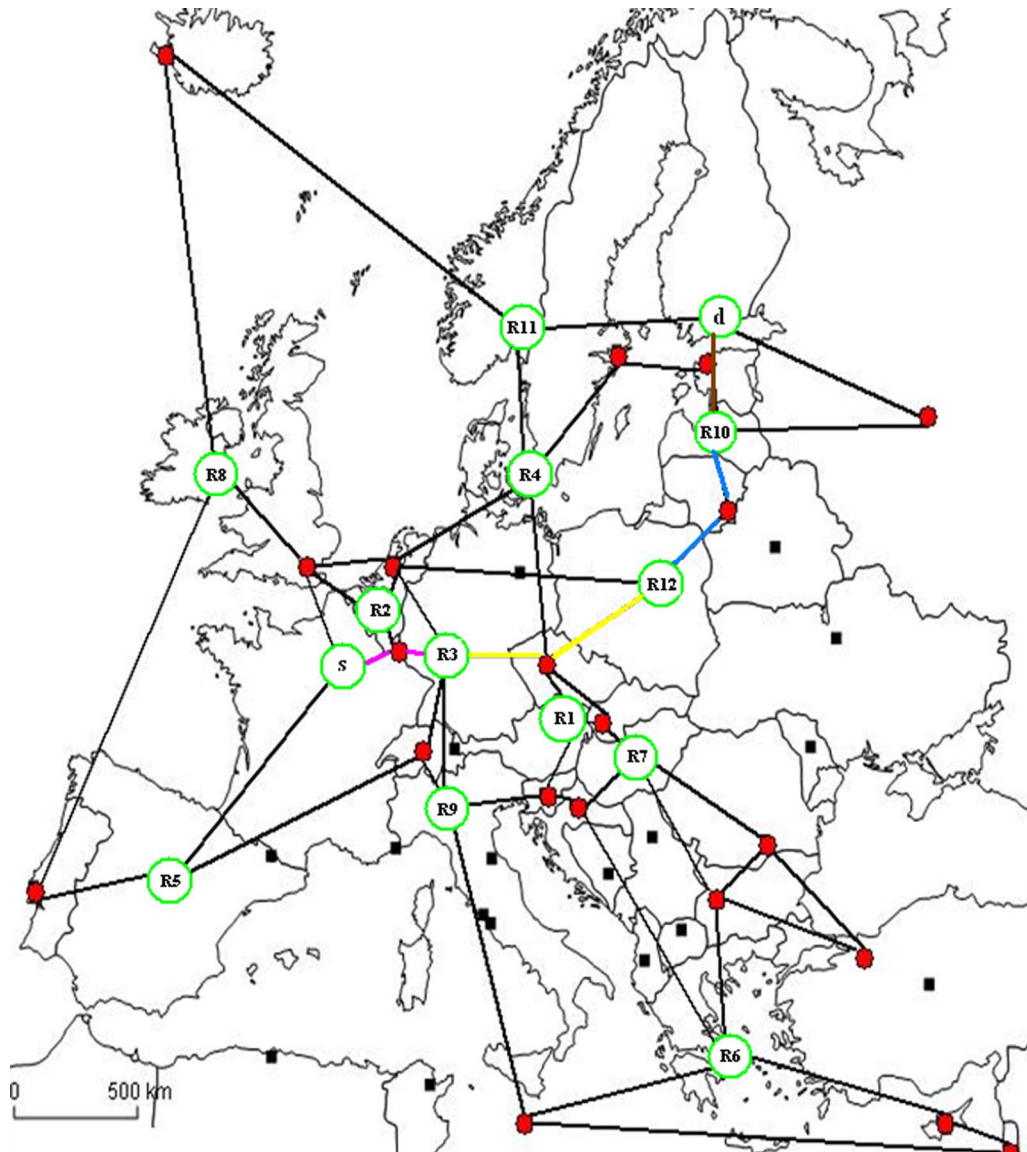


Fig. 1. Non-transparent connection request between source-destination pair (s, d) can be broken into four transparent sub-connection requests: s - $R3$, $R3$ - $R12$, $R12$ - $R10$ and $R10$ - d . Each of the three sub-connection requests can be served using a different wavelength.

investigated two classes of algorithms, which are presented in Sections III-A-II-A and B. In the example of Fig. 3 we focus on six regenerator sites (assuming 12 regeneration sites as presented in Fig. 1) and present the related virtual topology $\tilde{G}(s, d)$ for the source-destination pair (s, d) .

Greedy First-Fit Algorithms for Routing in the Virtual Topology: In what follows we present several Greedy algorithms for routing non-transparent connections in the virtual topology. We denote by $\Lambda^{\bar{T}}$ the part of the traffic matrix Λ that corresponds to source-destination pairs in \bar{T} , that is

$$\Lambda_{sd}^{\bar{T}} = \Lambda_{sd} \text{ if } (s, d) \in \bar{T}; \text{ and } \Lambda_{sd}^{\bar{T}} = 0, \text{ otherwise.}$$

Non-transparent connection requests are treated by the Greedy routing algorithms one by one. In particular, for each pair (s, d) for which $\Lambda_{sd}^{\bar{T}} \neq 0$, we run Dijkstra's shortest path algorithm on graph $\tilde{G}(s, d)$ with appropriate link costs. The way the costs

of the virtual links in $\tilde{G}(s, d)$ are defined is important for the performance of the algorithms, and is described next.

A non-transparent connection may be blocked due to the unavailability of free regenerators, or due to the unavailability of free wavelengths, or due to significant physical impairments (including interference by other lightpaths). The latter factor is important for paths that are long or use a large number of hops. Given the above considerations that affect blocking probability we studied the following definitions for the cost of the virtual links in the virtual topology:

1. *Virtual-Hop (VH) shortest path algorithm.* In this algorithm all the links of the virtual graph have cost equal to 1, and the cost of a virtual path is equal to the number of regenerators it crosses. The optimal virtual path is the one consisting of the fewest regenerators (virtual hops).
2. *Physical-Hop (PH) shortest path algorithm.* Here the cost of a virtual link is taken to be equal to the number of physical links (physical hops) it consists of. With this definition,

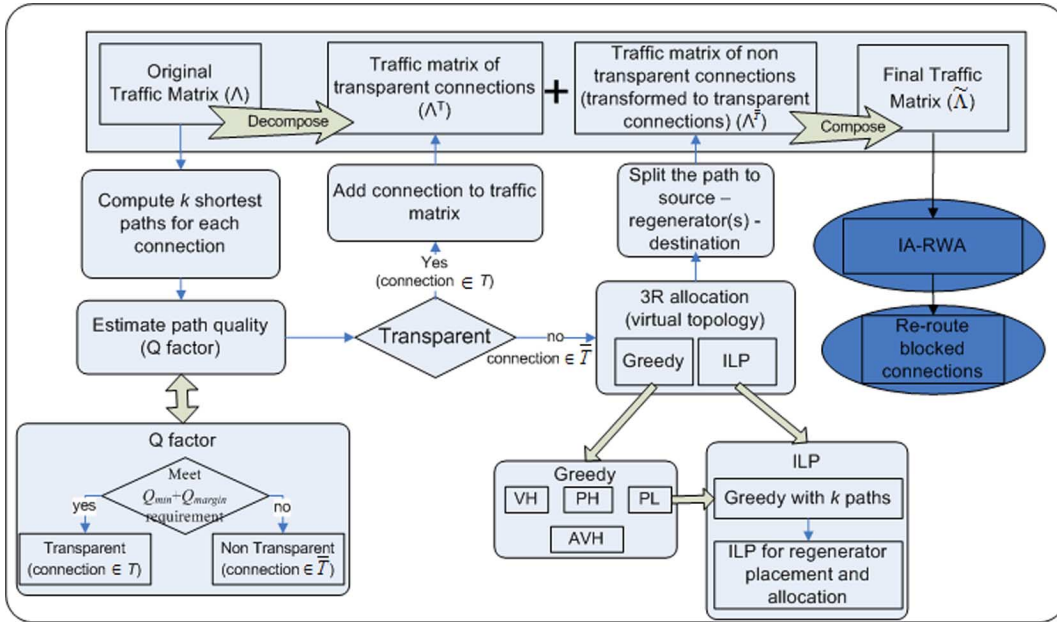


Fig. 2. Flow chart of the proposed algorithm.




Fig. 3. Assuming regeneration sites at nodes R2, R3, R4, R10, R11 and R12, we present the related virtual topology $\tilde{G}(s, d)$. The source node s and the destination node d are transparently connected to regeneration sites R2 and R3, and R10 and R11, respectively.

the optimal virtual path is the one that traverses the minimum number of physical nodes.

3. *Physical-Length (PL) shortest path algorithm.* In this case the cost of a virtual link is taken to be equal to the sum of the lengths of the physical links that comprise it. Thus, the PL algorithm selects the virtual path that has the shortest physical length.
4. *Adjusted-Virtual-Hop (AVH) shortest path algorithm.* Since in the three algorithms described above, Dijkstra’s algorithm is executed for each non-transparent (s, d) pair separately and successively, many connections may try to use certain regeneration sites, and, if the number of regen-

erators per site are limited, some of these connections may be blocked. To avoid this and to also distribute the load more equitably among the regeneration sites, the AVH algorithm adjusts the cost of using the regenerators of a specific site according to a properly increasing function of the number of regenerators assigned up to that point. This cost function is chosen to be concave, so as to imply a greater amount of ‘undesirability’ when a regenerator site becomes highly congested. In our simulation experiments, the cost of a virtual link l between virtual nodes i and j equipped with r_i and r_j regenerators, respectively, from which u_i and u_j regenerators have already been used by previous non-transparent connections, is taken to be

$$w_l = \left[\frac{u_i}{r_i} + \frac{u_j}{r_j} \right] \cdot \max_{n \in R} (r_n)$$

The performance of the VH, PH, PL and AVH heuristic algorithms will be given in Section IV. In Section III-A-II-B we present several alternative algorithms for selecting the sequence of regenerators to be used by the non-transparent connections based on combinatorial optimization approaches.

ILP Routing Algorithm for Routing in the Virtual Topology:

The ILP formulation of the virtual topology problem presented in this section result in an algorithm that is more computationally demanding than the heuristic algorithms presented in the previous section, but it is still practical for networks and traffic matrices of realistic size, as the results in Section IV will indicate.

The algorithm in this section uses a variation of the k -shortest path algorithm, with the link costs defined as in the VH, PH or PL algorithms of the previous section, in order to find for each non-transparent source-destination pair $(s, d) \in \bar{T}$ a set of k candidate virtual paths $\tilde{P}_k(s, d)$ on \tilde{G} . Clearly, by using only k alternative paths we may lose the optimal solution, but we decrease dramatically the problem size. In any case, the optimal solution can

be found by selecting the parameter k to be large enough. Then we select the routes to be followed by non-transparent connections by minimizing one of the following: (i) the maximum number of regenerators used among all network nodes, or (ii) the total number of regenerators used in the network, or (iii) the number of regeneration sites.

Minimizing the maximum number of regenerators used among the sites yields solutions that distribute the load more evenly in the network and thus also require fewer wavelengths to serve the connections. Minimizing the total number of regenerators used or minimizing the number of regeneration sites is important when the cost of deploying the regenerators is considered. Each of these objectives can be used to obtain good solutions, depending on the criterion that we want to optimize. If we are given cost parameters for the expenses of utilizing a wavelength, a regenerator, and a regenerator site, a weighted combination of these objectives can be used so as to minimize the CAPEX (Capital expenditure) and OPEX (Operational expenditure) of the translucent network, which depends not only on the number of wavelengths but also on the number of regenerator sites.

In what follows we formulate these algorithms as multicommodity flow problems.

Parameters:

- $s, d \in V$: source—destination physical network nodes
- $P_k(s, d)$: set of alternative paths that source-destination pair (s, d) may use in virtual topology $\tilde{G}(s, d)$
- $p \in P_k(s, d)$: a candidate virtual path (source, sequence of regenerators, destination) for source-destination pair (s, d)

Constants:

- Λ_{sd}^T : the number of requested connections for a non-transparent source-destination pair $(s, d) \in T$.
- r_n : the number of available regenerators on virtual node n .

Variables:

- x_p : represents the flow on path $p \in \tilde{P}_k(s, d)$ for a non-transparent pair $(s, d) \in T$, and takes integer positive values according to the number of non-transparent (s, d) connections that are routed over p . Note that x_p can take values larger than one, since more than one wavelengths can be requested between (s, d) .
- u_n : represents the utilization of regenerator site $n \in R$. It takes the value equal to 1, if n utilizes at least one regenerator, and 0, otherwise.

The virtual topology problem is then formulated as follows:

$$\begin{aligned} \text{Minimize : } R_{\max} &= \max_n \left(\sum_{\{p|n \in p\}} x_p \right) \\ \text{or } R_{\text{total}} &= \sum_n \sum_{\{p|n \in p\}} x_p \\ \text{or } R_{\text{sites}} &= \sum_n u_n \end{aligned}$$

where R_{\max} corresponds to the maximum number of regenerators among all nodes in the network needed to




Fig. 4. Variable u_n takes the value 1, if $\sum_{\{p|n \in p\}} x_p > 0$, and 0, otherwise.

serve the non-transparent connections, R_{total} to the total number of regenerators used in the network, and R_{sites} to the number of regeneration sites. We name the algorithm that minimizes R_{\max} as ILPmax, the one that minimizes R_{total} as ILPsum and the one that minimizes R_{sites} as ILPsites. Also, the name of the algorithm is extended based on the greedy algorithm used to calculate the k candidate virtual paths. For example, PH-ILPsum refers to the ILPsum algorithm where we have used PH algorithm to find the set of k candidate virtual paths.

Subject to the following constraints:

- Incoming Traffic constraints:

$$\sum_{p \in P(s,d)} x_p = \Lambda_{sd}^T, \text{ for all } (s, d) \in T.$$

- Regenerator constraints:

$$\sum_{\{p|n \in p\}} x_p \leq r_n, \text{ for all } n \in R.$$

- Regeneration sites utilization constraints:

$$\frac{1}{B} \cdot \sum_{\{p|n \in p\}} x_p \leq u_n \text{ for all } n \in R.$$

The regenerator constraints limit the number of connections that are routed through the regeneration sites, if the number of regenerators available per node is limited. The incoming traffic constraint represents the requirement that the connections of the non-transparent pair (s, d) should be satisfied by the chosen flows. Note that if more than one connections are requested for a given source—destination pair, they can be routed via different routes and regeneration nodes. For a regenerator site $n \in R$, the regeneration sites utilization constraint forces the site utilization variable u_n to take value equal to 1, if at least one virtual path crosses n is selected. In the case no virtual path that crosses n is selected, the site utilization variable u_n takes value equal to 0 (Fig. 4).

The number of x_p variables in the ILP problem is $k \cdot |\bar{T}|$, where k is the number of alternative paths calculated for each (s, d) pair and $|\bar{T}|$ is the number of active non-transparent source-destination pairs. The number of site utilization variables u_n is equal to the number of regeneration sites $|R|$. Finally, the number of constraints is $k \cdot |\bar{T}| + |R|$. This ILP can be solved even for large-sized networks, since the number of variables and constraints is small and depends only on the number of non-transparent connections and the number of regeneration sites.

When the ILP is intractable (which may be the case for very large networks), the LP relaxation of the problem combined with piecewise linear costs for the links and iterative fixing and rounding techniques, as presented in [15], can be used and will also yield good solutions. Note that in the simulation experiments we performed with realistic network and traffic input, we were always able to solve the corresponding ILP problems in a few seconds time.

3) *Transforming the Traffic Matrix*: Having found the sequence of regenerators to be used by non-transparent connections, the IA-RWA problem in the translucent network is transformed into a corresponding IA-RWA problem in a transparent network that has a different traffic matrix. In particular, after obtaining the solution to the virtual topology problem, which effectively breaks each non-transparent connection into multiple transparent segments, we transform appropriately the original traffic matrix, as follows. From the part of the initial traffic matrix that corresponds to the non-transparent connections, Λ^T , we use the virtual paths chosen for the non-transparent connections, to obtain the traffic matrix $\tilde{\Lambda}^T$ that corresponds to these virtual paths. By adding this to the matrix $\Lambda^T = \Lambda - \Lambda^T$ that represents the initial transparent connections, we obtain the transparent traffic matrix $\tilde{\Lambda} = \Lambda^T + \tilde{\Lambda}^T$ that if routed, will serve all requested (transparent and non-transparent) connections.

B. IA-RWA Using the Transformed Traffic Matrix (Phase 2)

Phase 2 of the algorithm takes as input the transparent traffic matrix $\tilde{\Lambda}$, calculated in the previous phase, which only contains connections that can be served transparently, and applies an IA-RWA algorithm designed for transparent networks to route and assign wavelengths to these connections. Given the traffic matrix $\tilde{\Lambda}$, the physical network topology G , the number of supported wavelengths W and impairment related parameters, the IA-RWA algorithm returns the solution to the specific instance in the form of routed lightpaths and the corresponding blocking probability.

Blocking in phase 2 may occur for two reasons: (i) if the transformed traffic matrix $\tilde{\Lambda}$ cannot be served with the given number of wavelengths some connections will have to be dropped, or (ii) the interference among lightpaths may make the quality of transmission of some lightpaths (which would be feasible in an empty network) unacceptable. Blocking of category (i) is usually not encountered, since for the given placement and number of regenerators, enough wavelengths W should be available in the incoming and outgoing links of the regeneration nodes to support the traffic that makes use of these regenerators. Otherwise, some of the available regenerators would be unusable, and the network would have been badly designed. Blockings of category (ii) can be avoided by a well-designed IA-RWA algorithm that takes into account the interference among lightpaths. In any case, connections that are blocked in phase 2 are re-attempted (re-routed) in phase 3 of the algorithm, as described in Section III-C.

The IA-RWA algorithm we used in our simulation experiments is the one proposed in [17], which is based on a ILP formulation that is solved in its LP relaxation form. Both linear and non-linear impairments are handled by this algorithm. More

specifically, the algorithm pre-calculates a set of k -candidate paths for each (s, d) pair of the given traffic matrix (matrix $\tilde{\Lambda}$ in this case). For each lightpath inserted in the LP formulation, we calculate a noise variance bound after accounting for the impairments that do not depend on the other lightpaths. We then express the interference among lightpaths by noise variance parameters. Using the noise variance bounds and these parameters we formulate new constraints and insert them in the ILP formulation. Solutions that satisfy the noise variance bounds are expected to yield acceptable transmission performance. The algorithm is executed in its LP relaxation form, using the Simplex algorithm that is generally considered efficient for the great majority of possible inputs. We use piecewise linear costs for the links, and iterative fixing and rounding techniques in order to obtain an integer solution [15].

We note that other IA-RWA algorithmic approaches can also be used in this second phase. The advantage of the algorithm in [17] is that it incorporates all the key requirements of the problem under consideration, and is, to the best of our knowledge, the only offline IA-RWA algorithm that accounts for the interference among lightpaths.

C. Rerouting the Blocked Connections (Phase 3)

In the final phase 3 of the algorithm we try to reroute any connections that were blocked in phase 2 of the algorithm.

Given the RWA solution provided by the IA-RWA algorithm of phase 2 and the set B of blocked (s, d) pairs, we want to serve the connections in B without altering the previously accepted lightpaths. Note that the set B of blocked connections refers to the initial traffic matrix Λ and not to the transformed traffic matrix $\tilde{\Lambda}$. For each $(s, d) \in B$, we check if it belongs to the set \bar{T} of transparent connections or to the set \bar{T} of non-transparent connections. In the first case [that is, when $(s, d) \in B$ and $(s, d) \in \bar{T}$], the transparent k -shortest paths that were used for the connection (s, d) as input in the IA-RWA algorithm were not appropriate. We now treat this (s, d) connection as if it were a non-transparent connection, and try to find a sequence of regenerators that it could use. In the second case [that is, when $(s, d) \in B$ and $(s, d) \in \bar{T}$], the chosen set of regenerator(s) was not appropriate and we also have to route this connection through a different set of regenerators.

Based on the above, we formulate a new virtual topology problem similar to that of phase 1 (Section III-A). The input is the set of remaining regenerators in the network and the set B of blocked connections. For each connection $(s, d) \in B$ we want to find a virtual path through the remaining regenerators, with the additional constraint that the connections that were considered earlier in phase 1 (i.e., connections $(s, d) \in \bar{T}$) are not served again by the same set of regenerator(s). Then we re-execute the IA-RWA algorithm of phase 2 (Section III-B), assuming that the lightpaths calculated in the previous solution are established (i.e., we set equal to 1 the corresponding variables). The output of this reduced IA-RWA problem will indicate if we can route the connections in B with acceptable transmission quality without affecting the other connections.

The proposed algorithm is terminated at this point and outputs the routed lightpaths from phase 2 in addition to any new lightpaths that were re-routed in phase 3. Connections that have




Fig. 5. Link model assumed in the experiments.

unacceptable Q performance, even after the re-routing attempt made in phase 3, are blocked.

IV. PERFORMANCE RESULTS

In order to evaluate the performance of the proposed impairment-aware RWA algorithm for translucent networks we carried out a number of simulation experiments. We implemented the virtual topology algorithms of Section III-A and the IA-RWA algorithm of Section III-B in Matlab, and we used LINDO API [20] to solve the related ILP problems.

To evaluate the feasibility of the lightpaths we used a Q-factor estimator that relies on analytical models to account for the most important impairments. The link model of the reference network is presented in Fig. 5. We assumed NRZ-OOK modulation format, 10 Gbps transmission rates and 50 GHz channel spacing. The span length in each link was set to 80 km. Each link was assumed to consist exclusively of SSMF fibers with dispersion parameter $D = 17$ ps/nm/km and attenuation parameter $a = 0.25$ dB/km. For the DCF we assumed parameters $a = 0.5$ dB/km and $D = -80$ ps/nm/km. The launch power was 3 dBm/ch for every SMF span and -4 dBm/ch for the DCF modules. The EDFAs' noise figure was 6 dB with small variations (± 0.5 dB) and each EDFA exactly compensates for the losses of the preceding fiber span. We assumed a switch architecture similar to [18] and a switch-crosstalk ratio $X_{sw} = 32$ dBs with small variations per node (± 1 dB). Regarding the dispersion management scheme, a pre-compensation module was used to achieve better transmission reach: initially the dispersion was set to -400 ps/nm/km, every span was under-compensated by a value of 30 ps/nm/km to alleviate non-linear effects, and the accumulated dispersion at each switch input was fully compensated to zero using an appropriate post-compensation module at the end of the link. The acceptable Q-factor limit was taken equal to $Q_{\min} = 15.5$ dB and the transparency margin equal to $Q_{\text{margin}} = 0.5$ dB.

The network topology used in our simulations was the Geant-2 network, shown in Fig. 1, which is a candidate translucent network, as identified by the DICONET project [19] with 34 nodes and 54 bidirectional links (for our simulations we assumed 108 directional links). All single-hop connections were able to be served transparently, but some multi-hop connections were not, making the use of regenerators necessary for some connections.

We performed two sets of simulation experiments. In the first set, reported in Section IV-A, we assume that all nodes can accommodate regenerators and the algorithm decides the regeneration sites and the number of regenerators to deploy at each

site (regenerator placement problem). In the second set of experiments, reported in Section IV-B, we assume that the regeneration sites are given, that is, we are given a sparse regenerator placement, and the algorithm decides which regenerators to use from the ones available (regenerator assignment problem). In Fig. 1 we indicate the nodes where the regenerators were placed for the second set of experiments.

The traffic matrix used in our simulation experiments corresponds to a realistic traffic load where there are 826 connection demands. For the given topology, the described link models, the given traffic matrix, and using $k = 3$ alternative candidate paths between any source-destination pair, the set of non-transparent connections consists of $|\bar{T}| = 373$ connections, which have to be routed through regenerators.

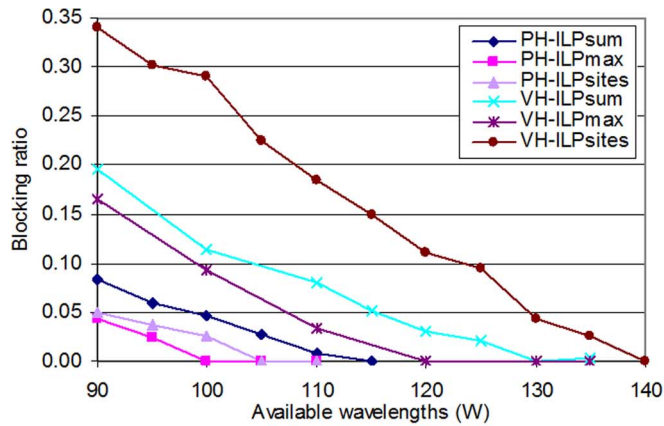
A. Regenerator Placement Simulation Experiments

In this set of simulation experiments we assume that the number of regeneration sites is not restricted; that is, every node in the network is capable of accommodating regenerators. In this version of the problem it is up to the proposed algorithms to solve the regeneration placement problem, in order to decide the regeneration sites and the number of regenerators to deploy on each site.

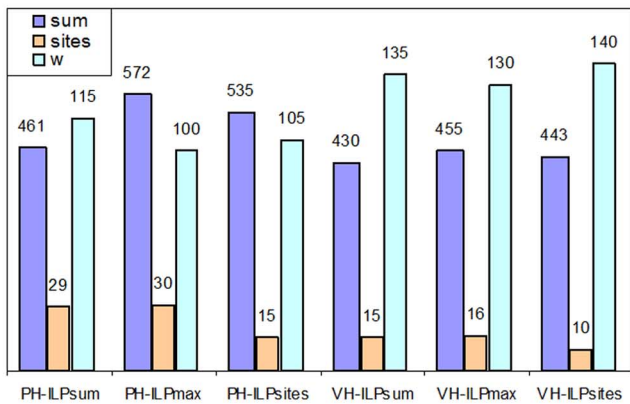
We evaluated the performance of the ILP virtual topology algorithms presented in Section III-A-II combined with VH and PH heuristics to pre-calculate their paths. Thus, the algorithms that are evaluated are: VH-ILPsum, VH-ILPmax, VH-ILPsites, PH-ILPsum, PH-ILPmax, and PH-ILPsites.

Note that we were able to obtain results for all these ILP algorithms with running times in the order of a few seconds. The related problems consist of $k \cdot |\bar{T}| + |R| = 1153$ variables and $k \cdot |\bar{T}| + |R| + |R| = 1187$ constraints (Section III-A-II-B), which are rather easily tractable using a typical personal computer.

In Fig. 6(a) we graph the blocking ratio as a function of the number of available wavelengths. From Fig. 6(a) it is obvious that PH-based algorithms have better wavelength performance, since they have lower blocking probability and require fewer wavelengths to reach zero blocking. Among these algorithms, the PH-ILPmax algorithm outperforms all the other examined algorithms, requiring the fewest number of wavelengths, in particular $W = 100$, in order to serve all the connections with zero blocking. This was expected, since the PH-ILPmax algorithm distributes the regenerators in the network so as to minimize the number of connections crossing a specific regeneration node, preventing regeneration nodes from becoming 'bottlenecks'. In this way, PH-ILPmax algorithm distributes the load and ends up



(a)

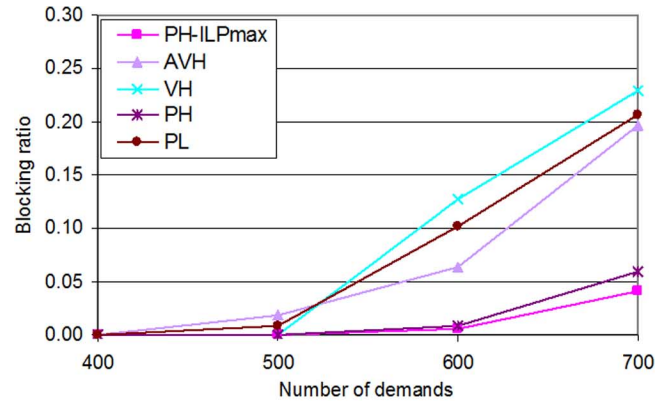


(b)

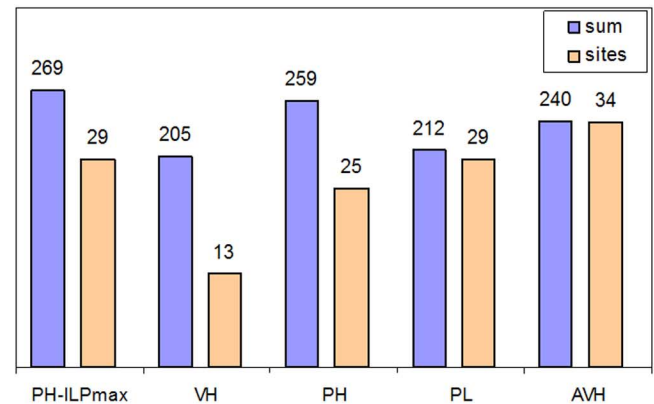
Fig. 6. (a) Blocking ratio for realistic traffic load versus number of available wavelengths and (b) total number of regenerators, total number of regeneration sites, and total number of wavelengths in the network required to obtain blocking equal to zero.

using fewer different wavelengths to route all the connections with zero blocking.

In Fig. 6(b) we graph (i) the total number of regenerators used in the network, (ii) the number of regeneration sites, and (iii) the number of wavelengths required in order to reach zero blocking. In this figure we can observe the difference between the VH- and PH-based algorithms. PH-based algorithms need a smaller number of wavelengths to reach zero blocking, but utilize more regenerators and regeneration sites. On the other hand VH-based algorithms need more wavelengths to reach zero blocking, but have better regenerator utilization performance. For example, if we consider PH-ILPsum and VH-ILPsum algorithms, they both minimize the total number of regenerators used in the network. PH-ILPsum needs 115 wavelengths while VH-ILPsum needs 135 to reach zero blocking. To achieve this performance PH-ILPsum needs 461 regenerators while VH-ILPsum needs 430. This difference can be explained by the fact that PH-based algorithms are fed with good candidate paths that have been calculated in the physical topology, while VH-based algorithms are fed with good candidate paths that have been calculated in the virtual topology. PH algorithm calculates paths that minimize the number of physical hops utilized, which tends to give good wavelength utilization performance, since shorter physical hop



(a)



(b)

Fig. 7. (a) Blocking ratio versus number of demands, assuming 80 available wavelengths and unbounded regeneration sites. (b) Total number of regenerators and total number of regeneration sites for 400 connection demands, 80 available wavelengths and unrestricted regeneration sites.

paths use less wavelengths than longer physical hop paths. This is the reason shorter physical hop paths are widely used in pure (without impairments) RWA problems. On the other hand, these paths are not directly related to the virtual topology and thus the ILP algorithm that runs over the virtual topology cannot find a good solution that minimizes its objective. The VH-based algorithms use as input good virtual paths that are not related to the physical topology. So, although these paths can yield good solutions to the regenerator placement problem that is performed over the virtual topology, they may be long physical hop paths that waste wavelength resources.

From Fig. 6(b) we can see that there is a trade-off between the number of wavelengths required to reach zero blocking and the number of regenerators used to reach this goal. All examined algorithms provide good solutions depending on the objective that has to be minimized and there is no single algorithm that outperforms the others with respect to all criteria. For example, if the cost of the regenerators is the dominant expense, one has to use VH-base algorithms and probably VH-ILPsum; if the objective is to minimize the number of wavelengths, PH-ILPmax seems to perform better.

In Fig. 7 we compare the performance of the PH-ILPmax algorithm that requires the smallest number of wavelengths to reach zero blocking (see Fig. 6(a)) with the proposed heuristic

algorithms, presented in Section III-A-I. Fig. 7(a) plots the blocking ratio as a function of the number of requested connections (network load) with constant number of available wavelengths equal to 80. Fig. 6(b) plots the total number of regenerators required in the network and the total number of regeneration sites, assuming a traffic matrix consisting of 400 connections. The performance of each algorithm is closely related to the metric it minimizes.

From Fig. 7(a) and (b) we can observe that PH-ILPmax and PH algorithms exhibit similar performance. They both have good blocking performance but not so good regenerator placement performance. These algorithms have similar objectives, since PH-ILPmax is a combinatorial extension of PH heuristic that uses as input a set of paths calculated by the PH heuristic (in particular, $k = 3$ paths per connection). The VH heuristic has the worst blocking performance but the best regenerator utilization performance, since it focuses on the virtual topology and the minimization of the number of regenerators used. The AVH heuristic distributes the load on all 34 nodes of the network, since the number of regenerators and the number of regeneration sites is not limited in these simulation experiments. By distributing the load its blocking performance is better than the corresponding unadjusted VH algorithm. Finally, PL algorithm selects virtual paths that have small physical lengths in order to obtain good impairment performance. It seems that this objective yields medium blocking and regenerator utilization performance.

The execution of all the phases of the IA-RWA algorithm for translucent networks requires about 3 hours and is dominated by phase 2 of the algorithm, that is, by the execution of the IA-RWA algorithm for the transparent transformed traffic matrix [17]. Phases 1 and 3 of the algorithm have execution times in the order of tens of seconds, respectively. Given the size of the network and the traffic matrix, both of which correspond to realistic scenarios, we consider these execution times acceptable, and it is evident that the proposed algorithms can scale to large networks.

B. Regenerator Assignment Simulation Experiments

In this set of simulation experiments we consider the case where the set of regeneration sites available in the network is given. In particular, we assume that there are 12 regeneration sites as shown in Fig. 1, equipped with limited number of regenerators and in particular equal to 100. This corresponds to the regeneration assignment problem.

In Fig. 8 we graph the blocking ratio obtained for the PH-ILPmax algorithm and the heuristic algorithms used in phase 1, as a function of the number of connection demands, for 80 available wavelengths. Comparing the performance of the algorithms that is reported in Figs. 7(a) and 8 we can see the significant improvements obtained by the AVH algorithm. This is because the AVH algorithm adjusts the cost of the virtual links by taking into account the remaining number of regenerators on the sites. PH-ILPmax and PH algorithms also have good performance, since they minimize the maximum number of regenerators used in the network and are not affected by the restriction on this parameter. On the other hand, the performance of VH and PL algorithms deteriorate when

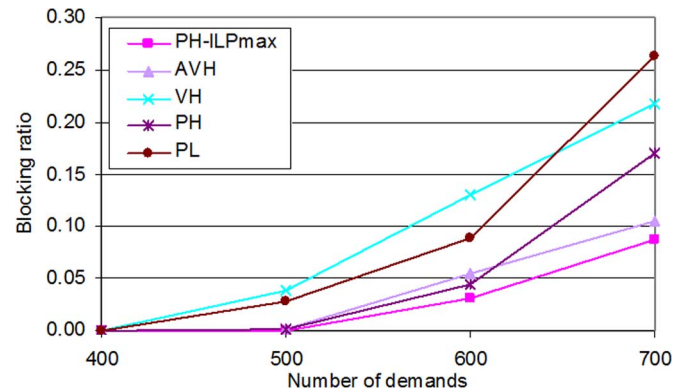


Fig. 8. Blocking ratio versus number of demands, assuming 80 available wavelengths and bounded regeneration sites and regenerators.

compared to the unrestricted case (Fig. 7), since they are not designed to take into account these kinds of restrictions. These results indicate that the performance of some of the heuristic algorithms depends significantly on the placement and the number of regenerators available in the network and this issue has to be considered when solving the regeneration assignment problem.

V. CONCLUSIONS

We considered the offline routing and wavelength assignment problem in translucent WDM networks in the presence of physical impairments. The proposed algorithms start by transforming the initial traffic matrix into a transparent traffic matrix, by choosing the regenerators that should be used by non-transparent connections. This is done by transforming the problem into a virtual topology design problem that is solved by a series of algorithms, ranging from ILP to greedy heuristics. Then a transparent impairment-aware RWA algorithm is applied to serve the transformed transparent traffic matrix, with any blocked connections being re-routed through the remaining regenerators in a third phase of the algorithm. The use of the Q factor as the quality of transmission metric in order to distinguish between the transparent and non-transparent connections and the low complexity of the virtual topology design problem make the proposed algorithms efficient for medium and large scale heterogeneous mesh networks. With respect to the different algorithms examined for solving the virtual topology problem, we observed that the ILP formulation that minimizes the maximum number of regenerators among the regeneration nodes exhibits the best wavelength performance. We also proposed heuristic algorithms that provide fast solutions to the virtual topology problem with acceptable performance, which depends on the type of the problem (regenerator assignment or regenerator placement) and the version of the heuristic algorithm (choice of the link cost function) used. In particular, the PH heuristic was shown to exhibit good blocking performance and stable function under all examined cases.

REFERENCES

- [1] T. E. Stern and K. Bala, *Multiwavelength Optical Networks: A Layered Approach*. Englewood Cliffs, NJ: Prentice Hall, 1999.

- [2] B. Mukherjee, *Optical WDM Networks*. Berlin, Germany: Springer-Verlag, 2006.
- [3] G. P. Agawal, *Fiber-Optic Communication Systems*, 3rd ed. Hoboken, NJ: Wiley, 2002.
- [4] B. Ramamurthy, D. Datta, H. Feng, J. P. Heritage, and B. Mukherjee, "Transparent vs. opaque vs. translucent wavelength-routed optical networks," in *Proc. OFC*, 1999, pp. 59–61.
- [5] A. M. Saleh, "Transparent optical networking in backbone networks," in *Proc. OFC*, 2000, pp. 62–64.
- [6] G. Shen, W. D. Grover, T. H. Cheng, and S. K. Bose, "Sparse placement of electronic switching nodes for low-blocking in translucent optical networks," *OSA J. Opt. Netw.*, vol. 1, pp. 424–441, 2002.
- [7] Y. Ye, T. Y. Chai, T. H. Cheng, and C. Lu, "Novel algorithm for upgrading of translucent optical networks," *Opt. Express*, vol. 11, no. 23, pp. 3022–3033, 2003.
- [8] X. Yang and B. Ramamurthy, "Dynamic routing in translucent WDM optical networks: The intra-domain case," *IEEE/OSA J. Lightw. Technol.*, vol. 23, no. 3, pp. 955–971, 2005.
- [9] Y. Ye, T. Y. Chai, and C. Lu, "Routing and wavelength assignment algorithms for translucent optical networks," *Opt. Commun.*, pp. 233–239, 2004.
- [10] S. Pachnicke, T. Paschenda, and P. Krummrich, "Assessment of a constraint-based routing algorithm for translucent 10 Gbits/s DWDM networks considering fiber nonlinearities," *OSA J. Opt. Netw.*, vol. 7, no. 4, pp. 365–377, 2008.
- [11] Y. Ye, T. Y. Chai, T. H. Cheng, and C. Lu, "Algorithms for the design of WDM translucent optical networks," *Opt. Express*, vol. 11, no. 22, pp. 2917–2926, 2003.
- [12] B. Ramamurthy, S. Yaragorla, and X. Yang, "Translucent optical WDM networks for the next-generation backbone networks," in *Proc. Globecom*, 2001, pp. 60–64.
- [13] X. Yang and B. Ramamurthy, "Sparse regeneration in translucent wavelength-routed optical networks: Architecture, network design and wavelength routing," *Photon. Netw. Commun.*, vol. 10, no. 1, 2005.
- [14] M. A. Ezzahdi, S. A. Zahr, M. Koubaa, N. Puech, and M. Gagnaire, "LERP: A Quality of transmission dependent heuristic for routing and wavelength assignment in hybrid WDM networks," in *Proc. ICCCN*, 2006, pp. 125–136.
- [15] K. Christodouloupoloulos, K. Manousakis, and E. Varvarigos, "Comparison of routing and wavelength assignment algorithms in WDM networks," in *Proc. Globecom 2008*, pp. 1–6.
- [16] K. Manousakis, K. Christodouloupoloulos, and E. Varvarigos, "Impairment-Aware offline RWA for transparent optical networks," in *Proc. Infocom 2009*, to be published.
- [17] K. Christodouloupoloulos, K. Manousakis, and E. A. Varvarigos, "Cross layer optimization of static lightpath demands in transparent WDM optical networks," in *Proc. IEEE ITW Netw. Inf. Theory*, 2009, to be published.
- [18] B. Ramamurthy, D. Datta, H. Feng, J. P. Heritage, and B. Mukherjee, "Impact of transmission impairments on the teletraffic performance of wavelength-routed networks," *IEEE/OSA J. Lightw. Technol.*, vol. 17, no. 10, pp. 1713–1723, 1999.
- [19] Dynamic Impairment Constraint Network for Transparent Mesh Optical Networks (DICONET) [Online]. Available: <http://www.diconet.eu/>
- [20] LINDO API [Online]. Available: <http://www.lindo.com/>

Konstantinos Manousakis received the Diploma degree in 2004 and the M.Sc. degree in Computer Science and Engineering in 2007, both from the Computer Engineering and Informatics Department, the University of Patras, Greece. He is currently a Ph.D. candidate there. His research activities focus on optimization algorithms for high-speed and optical networks.

Konstantinos Christodouloupoloulos received a Diploma in Electrical and Computer Engineering from the National Technical University of Athens, Athens, Greece, in 2002, the M.Sc. degree in Advanced Computing from the Imperial College London, London, U.K., in 2004, and the Ph.D. degree from the Computer Engineering and Informatics Department, University of Patras, Patras, Greece, in 2009.

His research interests are in the areas of protocols and algorithms for optical networks and grid computing.

Euaggelos Kamitsas received a Diploma in Electrical and Computer Engineering in 2005 and the M.Sc. degree in Computer Science and Engineering from the Computer Engineering and Informatics Department in 2008, both from the University of Patras, Patras, Greece.

His research interests are in the areas of physical layer and network layer algorithms for optical networks.

Ioannis Tomkos received the B.Sc. degree in Physics from the University of Patras, Greece, the M.Sc. degree in Telecommunications Engineering and the Ph.D. degree in Optical Telecommunications from the University of Athens, Greece.

He is a Full Professor at the Athens Information Technology (AIT) Center, Athens, Greece, and serves as its serves as the Associate Dean since 2004. He is an Adjunct Faculty at Carnegie-Mellon University, Pittsburgh, PA. He was a Senior Scientist at Corning Inc., from 1999 to 2002. He joined AIT in 2002, where he founded and serves as the Head of the High-Speed Networks and Optical Communications (NOC) Group that has been involved in many EU funded research projects (including five running FP7 projects), within which he has represented AIT as Principal Investigator, and has a consortium-wide leading role. These have included Project Leader of the EU ICT STREP Project DICONET, Technical Manager for the EU IST STREP Project TRIUMPH, Chairman of the EU COST 291 Project, and WP leader of many other projects.

Dr. Tomkos has received the prestigious title of "Distinguished Lecturer" of the IEEE Communications Society on the topic of transparent optical networking. He serves as the Chair of the International Optical Networking Technical Committee of the IEEE Communications Society and the Chairman of the IFIP working group on photonic networking. He is currently the Chairman of the OSA Technical Group on Optical Communications. He has been the General Chair, Technical Program Chair, Subcommittee Chair, Symposium Chair and/or member of the steering/organizing committees for several major conferences (e.g., OFC, ECOC, IEEE Globecom, IEEE ICC, ONDM, etc.) in the area of telecommunications/networking, as well as more than 50 conferences and workshops. He is a member of the editorial boards of the IEEE/OSA JOURNAL ON OPTOELECTRONICS, the *IET Journal on Optoelectronics*, and the *International Journal on Telecommunications Management*.

Emmanuel A. Varvarigos received a Diploma in Electrical and Computer Engineering from the National Technical University of Athens, Athens, Greece, in 1988, and the M.S. and Ph.D. degrees in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology, Cambridge, in 1990 and 1992, respectively.

He has held faculty positions at the University of California, Santa Barbara from 1992–1998, as an Assistant and later an Associate Professor, and Delft University of Technology, the Netherlands from 1998–2000, as an Associate Professor. In 2000, he became a Professor at the Department of Computer Engineering and Informatics at the University of Patras, Patras, Greece, where he heads the Communication Networks Lab. He is also the Director of the Network Technologies Sector (NTS) at the Research Academic Computer Technology Institute (RA-CTI), which, through its involvement in pioneering research and development projects, has a major role in the development of network technologies and telematic services in Greece. His research activities are in the areas of high-speed network, protocols and architectures, distributed computation and grid computing.