Reach Adapting Algorithms for Mixed Line Rate WDM Transport Networks

K. Christodoulopoulos, K. Manousakis, E. Varvarigos
Department of Computer Engineering and Informatics, University of Patras, Greece, and
Research Academic Computer Technology Institute, Patra, Greece
{kchristodou, manousak, manos}@ceid.upatras.gr

Abstract—**We consider the problem of planning a mixed line rate (MLR) wavelength division multiplexing (WDM) transport optical network. In such networks, different modulation formats are usually employed to support transmission at different line rates. Previously proposed planning algorithms have used a transmission reach bound for each modulation format/line rate, mainly driven by single line rate systems. However, transmission experiments in MLR networks have shown that physical layer interference phenomena are more severe among transmissions that utilize different modulation formats. Thus, the transmission reach of a connection with a specific modulation format/line rate depends also on the other connections that co-propagate with it in the network. To plan a MLR WDM network, we present routing and wavelength assignment (RWA) algorithms that adapt the transmission reach of each connection according to the use of the modulation formats/line rates in the network. The proposed algorithms are able to plan the network so as to alleviate cross-rate interference effects, enabling the establishment of connections of acceptable quality over paths that would otherwise be prohibited.**

Keywords- Wavelenth Division Multiplexing (WDM), Mixed line rate (MLR) optical network, Routing and Wavelenth Assignment (RWA), planning (offline) phase, transmission reach, cross-rate interference.

I. INTRODUCTION

Optical networks using Wavelength Division Multiplexing (WDM) technology modulate multiple channels over a single fiber. The most common architecture utilized for establishing communication in WDM optical networks is *wavelength routing* [1], where the communication between a source and a destination node is performed by setting up optical channels between them, called lightpaths. From the network perspective, establishing a lightpath for a new connection requires the selection of a route (path) and a free wavelength on the links that comprise the path. The problem of selecting appropriate paths and wavelengths for a set of requested connections is called Routing and Wavelength Assignment (RWA), and its objective is to minimize the network resources used, or the network cost, or to maximize the traffic served for a given set of resources.

Given the rapid increase of traffic demand, the available bandwidth of many core networks has to be continuously upgraded. While the industry wants to move quickly to higher capacity optical transport networks and enhance the 10-Gbps systems currently employed, there are a number of technology issues that need to be addressed. Transmission performance, price, space and power dissipation per bit have to be improved to justify the use of 40- and 100-Gbps WDM transport as a more effective solution than 10-Gbps. As the technology matures, higher rate connections will be incorporated in existing 10-Gbps systems [2]-[4]. Thus, a transport network will end up managing a variety of line rates, what is usually referred to as a mixed line rate (MLR) WDM system (Fig. 1). Currently, 40-Gbps connections are deployed and we expect that in the near future even 100-Gbps transponders will reach production level.

Signal transmission is significantly affected by physical limitations of fibers and optical components [5]. Transmission reach is the distance an optical signal can travel before its quality and the bit-error-ratio (BER) degrade to an unacceptable level. Many factors affect the transmission reach: the launched power of the signal, the modulation format, the bit rate, the type of the amplification, the dispersion map, the interference from other signals, etc. To plan a single line rate (SLR) WDM system, the transmission reach can be used as a constraint in a coarse RWA planning algorithm without considering the utilization state of the network. More accurate physical layer models [6] hat take into account interference effects among the lightpaths can give better and more sophisticated algorithmic solutions [7].

For a given modulation format, higher rate transmissions have a shorter reach than lower rate transmissions, due to higher impairments. After a point, increasing the rate of a transmission becomes impractical, and is the main reason that we have to consider different and improved modulation techniques with a better reach-rate product. Note that 10-Gbps systems typically utilize ON/OFF keying (OOK) modulation. To move to higher rates more advanced modulation formats, such as duobinary or phase shift keying (PSK) modulation techniques, with higher spectral efficiency and more tolerance to impairments have to be employed [2]-[4]. Even with these advanced modulation techniques, transmission reach is expected to decrease as we move from 10 to 40-Gbps transmission and from 40 to 100-Gbps transmission while the relative cost of the transponders is expected to increase.

Planning a MLR network to support, e.g., 10/40/100-Gbps over the same system, can reduce the total cost of the transponders by exploiting the heterogeneity and flexibility that is provided by MLR transmissions. The total cost of the transponders is the sum of the products of the number of transponders of each type multiplied by their corresponding cost. To reduce the total transponder cost, some long-distance low-bit-rate connections could be served with inexpensive low-rate and long reach 10-Gbps transponders, while short-distance high-bit-rate connections could be served with more expensive, but fewer in number, high-rate connections using improved modulation format 40- or 100-Gbps transponders, so as to have the lowest possible total transponder cost.

Recently, routing and wavelength assignment (RWA) algorithms for MLR systems have been proposed [8]-[11]. The

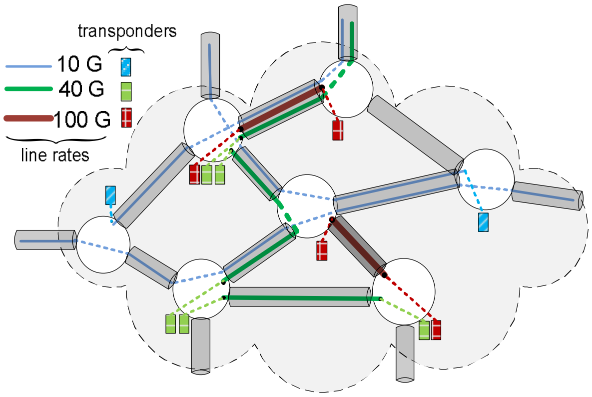

Figure 1: A part of a network that supports mixed line rates (MLR).

authors in [8] investigate the bit-rate migration from a networking point of view, by providing insight into the optimization of routing and aggregation in terms of overall capital expenditures (CapEx). For long-term migration, optimal network cost is achieved by early investments in 40-Gbps-only transmission systems. The authors in [9] formulate as an ILP the planning problem of a transparent MLR network under transmission-reach constraints for different modulations formats. Extending the work of [9], [10] proposes an algorithm for planning translucent MLR networks that consists of two phases. In the first phase the algorithm identifies candidate regenerators and then in the second phase solves the MLR cost optimization problem using the regeneration choices provided by the first phase. Reference [11] considers the logical topology planning problem of carrier Ethernet connections over a MLR network with transmission reach constraints. Both optimal ILP and heuristic algorithms are proposed and evaluated. Taking a different approach, optical Orthogonal Frequency Division Multiplexing (OFDM) can be used as a new networking solution that provides flexible bandwidth allocation to connections. A comparison study of the cost of a WDM and an OFDM-based network is presented in [12].

Multiplexing wavelength channels with different modulation format/line rates in a MLR system, introduces a number of additional technical issues. A field trial has been conducted to demonstrate the feasibility of accommodating 10-, 40- and 100-Gbps transmissions over a typical 50-GHz grid [3]. Depending on the signal power and other physical characteristics, the interference among simultaneously transmitted optical signals with different modulation formats or different rates can lead to considerable degradations in signal quality [3][4],[13][14], and consequent reductions in the transmission reach [15]. The authors in [15] used numerical simulations to examine the transmission reach by accounting for the nonlinear interaction between channels in a mixed-format system. They observed decreases in the transmission reach of up to 25% compared to the single line rate (SLR) system, depending on the transmission power of the connections. They then proposed a heuristic algorithm to plan a MLR system following a worst case approach, where decreased transmission reaches, calculated assuming worst case interference, are used for all supported rates, without considering the actual utilization state of the network. In [14] the authors reviewed analytical models that evaluate the quality of transmission of the lightpaths in a MLR system, taking also into account the cross-phase modulation (XPM) interference among the different formats/line rates. They then continued and proposed a number of solutions for establishing lightpaths in MLR systems for online traffic, that is serving a single connection at a time. The algorithms separate interfering rate connections using guardband or free wavelengths. A similar approach is adopted in [16]. The authors in [16] considered a MLR network with 10-Gbps OOK and 40-Gbps DQPSK connections and present algorithms that avoid interference between these two types of connections by leaving appropriate guardband wavelength space in-between interfering connections. Our approach is quite more sophisticated and explores a wider solution space than the aforementioned cases [14]-[16]. This is because we are able to adapt the transmission reaches of the connections according to the utilization state of the network and thus we are able to control and leave wavelength space between connections only when needed.

In this paper we present RWA algorithms for planning mixed line rate (MLR) optical transport networks. However, the presented model is general and can be used for dynamic (online) traffic problems, as well. The proposed heuristic algorithms serve sequentially the connections, which means that they are essentially online algorithms and can be used with small changes to serve dynamic traffic. In MLR networks, as discussed in [3]-[4], [13]-[16], the transmission reach of a lightpath at a given modulation format/rate, changes depending on the modulation format/rates of the connections that co-propagate with it along the path. For this reason, in MLR networks, it is not enough to consider a specific transmission reach for each modulation format/rate, but also the interactions among the connections for the specific modulation formats/rates they use, which we will call *cross-rate interference*. The proposed algorithms adapt the transmission reach of the connections according to the utilization state of the network, and plan the network so as to avoid cross-rate interference effects. We initially present optimal ILP algorithms for the MLR planning problem of both transparent and translucent networks, that is, without and with the use of regenerators. We also give sequential heuristic algorithms that serve the connections in a particular order, propose a specific ordering policy and also use simulated annealing to find even better orderings. Our results indicated that the proposed algorithms can efficiently utilize the wavelength domain to absorb cross-rate interference effects, enabling the establishment of connections with acceptable quality over paths that would otherwise be prohibited.

In our previous work [17] we also examined the problem of planning a MLR network, but we followed a per link worst case cross-rate interference assumption. In this paper, we extend our work and formulate the cross-rate interference based on the actual utilization of the wavelengths of the network. This work is more general, so that the problem considered in [17] is a special case of the one considered here. The new algorithms presented here are able to utilize the wavelength domain in order to avoid cross-rate interference effects. Moreover, in [17] we only provided algorithms for transparent networks, that is, networks that do not utilize regenerators, while here we provide algorithms for both transparent and translucent networks and perform a large number of simulation experiments to evaluate the performance of the algorithms in both network settings.

The rest of the paper is organized as follows. In Section II we formulate the adaptation of the transmission reach for a mixed line rate (MLR) optical transport network by introducing the effective length metric. Next, in Section III, we describe the proposed reach-adapting algorithms for planning MLR systems. Performance results are presented in Section IV. Finally, in Section V we give our concluding remarks.

II. NETWORK MODEL AND EFFECTIVE LENGTH

In a single line rate (SLR) system, given the modulation format and the rate that is going to be used, the network is designed to achieve long transmission reaches, using specifically designed amplification schemes, dispersion maps, etc. Typically, in an optical transport network that supports mixed line rates (MLR), different modulation formats are employed to support the transmissions at different rates. In such a MLR network the transmission reach of each modulation format/rate is not the same as the optimized reach in a corresponding SLR network, but is somewhat reduced [3],[4],[14],[15]. Due to interference effects between the different modulation formats/rates used, the transmission reach of each modulation format/rate is affected by the other transmissions. For example, intensity modulated connections (e.g. 10-Gbps OOK connections) induce significant cross-phase-modulation (XPM) on an xPSK modulated 40- or 100-Gbps connection [14]. However, according to [15], even different rate connections with the same modulaton format (xPSK connections) are affected by non-linear cross-rate interference. In particular, reductions of up to 25% for the cases of concurrent PDM-QPSK, PDM-BPSK and SP-BPSK are reported in [15]. Also the power budgeting and the dispersion maps employed play an important role and may deteriorate the transmission reach of the connections in a MLR as compared to a SLR system. Although in this paper we focus on cross-rate interference effects, some of the above parameters might be captured by our formulation as well. It is worth noting that the proposed model and algorithms are quite general, work for systems with any number of rates and can capture the interference effects between different rate connections in a non-uniform manner.

In what follows we present a way to formulate the variation of the transmission reach of a connection according to the utilization state of the network so as to capture cross-rate interference effects. In particular, depending on the modulation formats/rates transmitted over a link we calculate what we call the *effective length* metric of that link for a given connection. Instead of adapting-decreasing the transmission reach of the connection, we proportionally increase the effective lengths of the links that comprise its path in order to account for the cross-rate interference. For example, consider a connection that uses a specific modulation format/rate and shares a common link with another connection. Assume the second connection uses an interfering modulation format/rate and is within small enough spectrum/wavelength distance from the first to cause cross-rate interference. Instead of decreasing the transmission reach of the first connection, we increase by some amount the effective length of their common link so as to have exactly the same outcome as we would have if we decreased its transmission reach.

We consider a MLR network that supports a number of different rates R. For the sake of being specific, we will assume in this section and in the simulation results, to be presented in Section IV, that R={10,40,100}Gbps, and each link consists of a single fiber. However, the proposed model and the algorithms are quite general and also work for more and different rates.

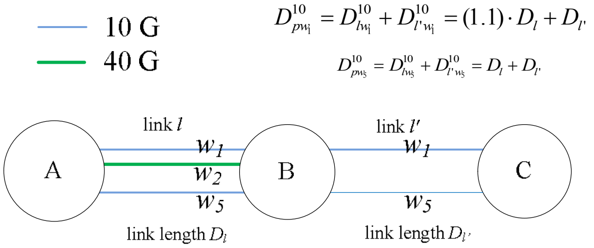

Figure 2: Calculation of the effective length of a lightpath, taking into account the interference of other lightpaths that utilize different rates/modulation formats.

We now formally define the adaptation of the link length and introduce the effective length metric. Assume a lightpath (p,w,r), that is, a lightpath utilizing path p and wavelength w using rate r. Assume a link l of length D_l crossed by path p ($l \in p$) and consider another lightpath (p',w',r') also crossing link l. We will say that lightpath (p,w,r) is subject to *cross-rate interfere* from lightpath (p',w',r'), if the lightpaths cross the same link l and their spectrum distance is within a given distance, $|w\text{-}w'| \leq I^{r,r'}$, where $I^{r,r'}$ is *the interfering distance threshold* in wavelengths. Lightpaths (p,w,r) and (p',w',r') sharing a link do not interfere if the wavelengths w and w' they use are more than $I^{r,r'}$ wavelengths apart from each other.

The *effective length* of the fiber link l of a lightpath (p,w,r) that is subject to interference from another lightpath of rate r' is calculated by $D^{r,r'}_{l,w} = m^{r,r'} \cdot D_l + D_l = (1+m^{r,r'}) \cdot D_l$, that is, it is equal to the physical length D_l of the link, increased by a proportional factor $m^{r,r'}$, due to cross-rate interference. We will refer to the parameters $m^{r,r'}$ as the *effective length factors*. In a similar manner, we can define the effective length of the link l for lightpath (p',w',r') that is subject to cross-rate interference from the first lightpath (p,w,r) to be $D^{r',r}_{l,w'} = (1+m^{r',r}) \cdot D_l$. In general, the effective length factors can be different for different directions of the interference, $m^{r',r} \neq m^{r,r'}$, and we can also have different wavelength interfering distance thresholds, $I^{r',r} \neq I^{r,r'}$. To have a consistent formulation we will assume that there is no cross-rate interference between lightpaths of the same rate r, and thus $m^{r,r}=0$. In other words, we assume in our formulation that the interference among lightpaths of the same rate is included in the calculation of the maximum transmission reach of each rate D^r.

We define the *effective length of link l for lightpath* (p,w,r) as

$$D^r_{lw} = D_l + \sum_{r':(p',w',r')\text{ is used, and } l \in p', \text{and} |w-w'| \leq I^{r,r'}} m^{r,r'} \cdot D_l \,. \quad (1)$$

Note that even if two or more lightpaths of rate r' are within interfering $I^{r,r'}$ distance, we increase the effective length of the link only once. Also note that the actual wavelength distance $|w\text{-}w'|$ is not taken into account as long as it less than $I^{r',r}$. These two assumptions mean that the effective length factor $m^{r,r'}$ accounts for the worst case interference effect that one or more connection(s) of rate r' within $I^{r,r'}$ have on a connection of rate r. Under this worst case assumption there is no need to consider how many r' connections are actually interfering, or their actual distance from the affected connection. More accurate models that would consider the exact number and distance of the cross-rate interfering lightpaths could be used. However, we argue that the used model is safe, since it captures the worst case assumption, but also gives us enough flexibility to use the wavelength domain to avoid cross-rate interference. Our performance results indicate that the proposed algorithms are able to utilize the wavelength domain to assign wavelengths to connections so as to avoid cross-rate interference. The proposed algorithms yield the same performance as if no cross-rate interference is present in the network, indicating that the used model, although coarse at certain points, is detailed enough to give the algorithms the required flexibility to avoid such effects.

Consider a lightpath (p,w,r) and assume that we know the rate utilization of the interfering wavelength channels on all the links $l_1,l_2, \ldots l_n$, comprising path p. We can use Eq. (1) to calculate $D^r_{l_i w}$, for all l_i, $i=1,2, \ldots,n$. The *effective length of lightpath* (p,w,r) is then given by

$$D^r_{pw} = D^r_{l_1 w} + D^r_{l_2 w} + \ldots + D^r_{l_n w}. \quad (2)$$

For the example of Fig. 2, the effective length for path p_{ABC} for wavelength w_1 at rate 10-Gbps is $D^{10}_{pw_1} = D^{10}_{lw_1} + D^{10}_{l'w_1} = 1.1 \cdot D_l + D_{l'}$, assuming that $m^{10,40} = 0.1$ and $I^{10,40}=2$. In comparison, the effective length of the same path and the same rate but for wavelength w_5 is $D^{10}_{pw_5} = D^{10}_{lw_5} + D^{10}_{l'w_5} = D_l + D_{l'}$. The difference in these two effective lengths is due to the 40-Gbps lightpath that utilizes link l and wavelength w_2, which interferes only with the 10-Gbps lightpath that uses wavelength w_1 and not with the one using w_5.

In our model, a lightpath of rate r has, in the absence of any cross-rate interference, maximum transmission reach D^r. As mentioned earlier, this transmission reach bound accounts for all other kinds of physical layer impairments a connection of rate r is subject to. We use this limit as an upper bound on the effective length (instead of the physical length) of all connections of rate r in the MLR system. If the effective length D^r_{pw} of lightpath (p,w,r) is higher than the given bound D^r then the lightpath is considered to have unacceptable quality of transmission (QoT) and cannot be used as part of the solution. As the effective lengths of the links change depending on the modulation formats/rates of the connections established in the network, the effective lengths of the lightpaths change accordingly. The effective lengths of the lightpaths are always higher than their corresponding real lengths, which correspond to the best possible case, that is, the case of zero-cross-rate interference.

Note that in the above model, we adapt the effective length of a link used by a given lightpath, based on the modulation formats/rates and wavelengths of the other lightpaths using it. The interference wavelength distance thresholds $I^{r,r'}$ constraints the number of adjacent wavelength channels that are considered for each lightpath. Threshold values in the order of 3 or 2 are logical, since interference effects degrade as we move away from the wavelength under examination. These values are inline with transmission experiments that have been conducted and corresponding analytical models that have been developed, which assumed 7 or 5 utilized wavelengths in total ($I^{r,r'}=3$ or 2, respectively), as, e.g., reported in [15]. Still the model used and the algorithms proposed are general and can be extended, if a higher number of adjacent wavelengths cause interference. The case $I^{r,r'}=W$ for all r,r', where W is the total number of wavelengths supported in the system, resembles the setting that we have previously examined in [17], where all the wavelengths of a link cause (substantial) interference to each other. Thus, the cross-rate interference model proposed in this paper is more general and includes [17] as a special case. Moreover, the proposed model is also flexible in the opposite direction, since it also includes the special case where cross-rate interference is not present, corresponding to $I^{r,r'}=0$ and/or $m^{r,r'}=0$, for all r,r'. Finally, note that in order to formulate link length adaptation as described above, we have to use link and wavelength related variables, similar to [7]. This more accurate and sophisticated model is also considerably more complicated than [17].

In the next section we propose algorithms that use the effective length model presented above to plan a MLR network.

III. REACH-ADAPTING MLR ALGORITHMS

We are given a network $G=(V,E)$, where V denotes the set of nodes and E denotes the set of (point-to-point) single-fiber links. We are also given the actual (physical) lengths D_l of all links $l \in E$. Each fiber is able to support a set $C=\{1,2,\ldots,W\}$ of W distinct wavelengths, and a set $R=\{r_1,r_2,\ldots,r_M\}$ of M different bit rates. Each rate is associated with a certain modulation format. Moreover, each rate r has an interfering wavelength distance threshold $I^{r,r'}$ and an effective length factor $m^{r,r'}$, for all $r' \in R$. The length of link l, normally D_l, is adapted to effective length D^r_{lw} for lightpath (p,w,r), depending on the other lightpaths that cross link l, according to Eq. (1). We are also given transmission reach bounds D^r and the corresponding transponder costs C^r for all the rates $r \in R$ supported in the network. It is natural to assume that the cost C^r of a transponder is higher for higher transmission rates r. Since the transmission reach of a modulation format/rate decreases as we move from lower to higher rates, e.g., from 10 to 40-Gbps and from 40 to 100-Gbps transmissions, there should be a cost benefit for using higher rates. Thus, the ratio of the transmission rate over the cost of the transponder (which is the per bit transmission cost) r/C^r should be higher for higher rates, or otherwise there would be no cost benefit of using higher rates. We assume an a-priori known traffic scenario given in the form of a matrix of aggregated demands Λ in Gbps, called the *traffic matrix*. Then, Λ_{sd} denotes the requested bandwidth from source s to destination d, that is Λ_{sd} is the end-to-end demand of commodity (s,d).

The objective of the RWA algorithm for planning a MLR system is to serve all traffic, described in Λ, and minimize the total cost of the transponders, related to the number and type of the transponders of different line rates used. Moreover, each lightpath selected in the solution has to satisfy an adaptive transmission reach constraint, modelled through the use of the effective lengths of the links that vary according to the utilization state of the network and the modulation formats/rates used, as described in Eq. (1) and (2).

In the following we present two sets of reach adapting algorithms to solve the planning problem of transport MLR systems. We describe algorithms to plan *transparent* networks, that is, networks with short lengths where lightpaths do not require regeneration, and also algorithms to plan *translucent* networks where regeneration may have to be performed at certain nodes in the network. We start by describing combinatorial optimization algorithms based on Integer Linear Programming (ILP) formulations to plan transparent and translucent MLR networks. Since these formulations cannot be solved efficiently for large input instances, we also propose heuristic algorithms that solve the planning MLR problems sub-optimally, but in polynomial time, by sequentially serving one-by-one the demands. The order in which demands are considered plays an important role in the performance of the heuristic algorithms. We propose and evaluate one ordering policy and also use a simulated annealing meta-heuristic to find good orderings that yield near-optimal performance.

A. ILP algorithms

1) Transparent MLR networks

In this section we focus on transparent MLR networks, which do not support regeneration, so that all connections are established end-to-end through transparent lightpaths. The proposed algorithm pre-calculates in a pre-processing phase for each source-destination pair (*s*,*d*) a set of *k* candidate paths P_{sd}, using a variation of the *k*-shortest path algorithm: at each step, a shortest path is selected and the costs of its links are increased (doubled in our experiments) so as to be avoided by the paths found in subsequent steps. The paths obtained in this way tend to use different edges so that they are more representative of the path solution space. Other *k*-shortest path algorithms are also applicable. We denote by $P = \bigcup_{sd} P_{sd}$ the set of all pre-calculated paths.

Variables

x^{r}_{pw}: Boolean variable. It is equal to 1 if transparent lightpath (*p*,*w*,*r*), that is, wavelength *w* with rate *r* over path $p \in P_{sd}$, is used to serve the commodity (*s*,*d*), and is equal to 0, otherwise.

$u^{r,r'}_{lw}$:Boolean variable. It is equal to 1 if at least one connection of rate *r’* is transmitted over link *l* using a wavelength in the range [max(0,w-$I^{r,r'}$), min(w+ $I^{r,r'}$,W)], and is equal to 0, otherwise. Thus, $u^{r,r'}_{lw}$ is equal to 1 if at least one lightpath of rate *r’* causes cross-rate interference to a lightpath (*p*,*w*,*r*) crossing link $l \in p$.

Objective

$$\text{minimize: } \sum_{p}\sum_{w}\sum_{r} C^{r} \cdot x^{r}_{pw}$$

subject to the following constraints:

- Capacity constraints

$$\text{For all } s,d \in V, \quad \sum_{p \in P_{sd}}\sum_{w}\sum_{r} r \cdot x^{r}_{pw} \geq \Lambda_{sd} \qquad \text{(C1)}$$

- Single wavelength assignment constraints

$$\text{For all } l \in E, \text{ for all } w \in W, \quad \sum_{p:l \in p}\sum_{r} x^{r}_{pw} \leq 1 \qquad \text{(C2)}$$

- Link-wavelength-rate utilization constraints

$$\text{For all } l \in E, \text{ for all } w \in W, \text{ for all } r,r' \in R,$$

$$\sum_{\max\left(0,w-I^{r,r'}\right) \leq w' \leq \min\left(w+I^{r,r'},W\right)} \; \sum_{p:l \in p} x^{r'}_{pw'} \leq B \cdot u^{r,r'}_{lw} \; , \qquad \text{(C3)}$$

where *B* is a large constant.

- Effective length constraints

$$\text{For all } p \in P, \text{ for all } w \in W, \text{ for all } r \in R,$$

$$\sum_{l \in p} D_{l} \cdot x^{r}_{pw} + \sum_{r' \in R}\sum_{l \in p} m^{r,r'} \cdot D_{l} \cdot u^{r,r'}_{lw} < D^{r} , \qquad \text{(C4)}$$

Constraints (C1) ensure that the lightpaths chosen to serve an end-to-end demand should have total capacity at least equal to the requested demand. Constraints (C2) prohibit the assignment of a wavelength to more than one lightpaths crossing the same link. Constraints (C3) identify cross-rate interfere among lightpaths so as to set accordingly the corresponding $u^{r,r'}_{lw}$ variables. To do so, constraints (C3) take into account the utilization of the lightpaths of the network. If at least one lightpath or rate *r’* crosses link *l* using wavelength *w*’ within interfering distance $I^{r,r'}$ from the examined wavelength *w*, then $u^{r,r'}_{lw}$ is forced to take the value of one. Then, variables $u^{r,r'}_{lw}$ are used in constraints (C4) to calculate the effective lengths of the lightpaths, based on the effective lengths of their links. The left hand side of (C4) calculates the effective length of a lightpath (*p*,*w*,*r*) according to Eq. (2). Then the lightpath’s effective length is constrained to be less than the accepted transmission reach at that rate [right hand side of (C4)]. Thus, constraint (C4) enables or disables the use of the specific lightpath: if the effective length of lightpath (*p*,*w*,*r*) is higher than the threshold, variable x^{r}_{pw} is forced to take the zero value, in which case lightpath (*p*,*w*,*r*) cannot be used in the solution.

The constant *B* used in constraints (C3) has to take values larger than $2 \cdot I^{r,r'}$. This is the highest value that the left hand side of (C3) can take, which corresponds to the case that all adjacent wavelengths within $I^{r,r'}$ distance from each side of the examined wavelength *w* are all utilized by lightpaths of rate *r’* that cross link *l*.

2) Tranlucent MLR networks

In this section we consider the planning of translucent MLR networks in which signal regeneration can be performed at intermediated nodes of an end-to-end connection. We assume that a regenerator is implemented by a transponder (transmitter-receiver connected back-to-back) and thus its cost is the same as the cost of the transponder of the same rate. However, we also comment on how to extend the proposed formulation to capture the case that the costs of the regenerators are different than the corresponding costs of the transponders. Following the above specification, when a lightpath is regenerated, it can also change its wavelength. Thus, a regenerator functions also as a wavelength converter. We assume that all nodes can be equipped with regenerators and there is no constraint on the number of regenerators that can be installed on each node.

The algorithm again pre-calculates a set P_{ij} of candidate paths between all pairs of nodes *i* and *j*. Note that in this case the nodes *i* and *j* can be intermediate nodes of a translucent end-to-end connection, instead of the actual source and destination nodes of the end-to-end connection, which was the case in transparent networks.

Variables

x^{r}_{pw}: Boolean variable. It is equal to 1 if lightpath (*p*,*w*,*r*), that is, wavelength *w* with rate *r* over path $p \in P_{ij}$, is utilized to connect (*i*,*j*), and is equal to 0, otherwise.

$u_{lw}^{r,r'}$: Boolean variable. It is equal to 1 if at least one connection with rate *r'* is transmitted over link *l* using a wavelength in the range [max(0,w-$I^{r,r'}$), min(w+ $I^{r,r'}$,W)], and is equal to 0, otherwise.

$f_{sd,ij}^{r}$: Integer variable. Equals to the number of lightpaths of rate *r* between nodes *i* and *j* that are used to serve commodity (*s*,*d*).

Note that in this formulation, indicator variable x_{pw}^{r} may correspond to a lightpath (*p*,*w*,*r*) that serves *transparently* an end-to-end demand between the given source and destination pair (*s*,*d*), or to an intermediate lightpath of a *translucent* connection that is realized by a series of lightpaths. In the latter case, the start and/or the end of the lightpath (*p*,*w*,*r*) are intermediate regeneration node(s) for the translucent connection. Variables $f_{sd,ij}^{r}$ are used as flow variables that identify the lightpaths used to serve the traffic of commodity (*s*,*d*). The lightpaths identified by the $f_{sd,ij}^{r}$ variables are realized though specific paths and wavelengths by the corresponding x_{pw}^{r} variables.

Objective

$$\text{minimize: } \sum_{p}\sum_{w}\sum_{r} C^{r}\cdot x_{pw}^{r}$$

subject to the following constraints:

- Capacity constraints – source node

$$\text{For all } s,d \in V,\ \sum_{r}\sum_{j} r\cdot f_{sd,sj}^{r} \geq \Lambda_{sd} \qquad \text{(C5)}$$

- Capacity constraints – destination node

$$\text{For all } s,d \in V,\ \sum_{r}\sum_{i} r\cdot f_{sd,id}^{r} \geq \Lambda_{sd} \qquad \text{(C6)}$$

- Flow constraints

$$\text{For all } s,d,n\in V,\ n\neq s,d, \text{ for all } r\in R,\ \sum_{i} f_{sd,in}^{r} = \sum_{j} f_{sd,nj}^{r} \qquad \text{(C7)}$$

- Lightpath assignment constraints

$$\text{For all } i,j \in V,\ \text{ for all } r\in R,\ \sum_{s}\sum_{d} f_{sd,ij}^{r} = \sum_{p\in P_{ij}}\sum_{w} x_{pw}^{r} \qquad \text{(C8)}$$

- Single wavelength assignment constraints

$$\text{For all } l, \text{ for all } w,\ \sum_{p:l\in p}\sum_{r} x_{pw}^{r} \leq 1 \qquad \text{(C9)}$$

- Link-wavelength- rate utilization constraints

For all *l*, for all *w*, for all *r*, for all *r'*

$$\sum_{\max\left(0,w-I^{r,r'}\right)\leq w'\leq\min\left(w+I^{r,r'},W\right)}\ \sum_{p:l\in p} x_{pw'}^{r'} \leq B\cdot u_{lw}^{r,r'}\ , \qquad \text{(C10)}$$

where *B* is a large constant, as before.

- Effective length constraints

For all *p*, for all $r\in R$,

$$\sum_{l\in p} D_{l}\cdot x_{pw}^{r} + \sum_{r'\in R}\sum_{l\in p} m^{r,r'}\cdot D_{l}\cdot u_{lw}^{r,r'} < D^{r}\ , \qquad \text{(C11)}$$

Constraints (C5) ensure that the lightpaths that start from the source node of an end-to-end demand have total capacity higher than the requested demand. Constraints (C6) function in a similar way at the destination node, while constraints (C7) ensure the flow conservation of lightpaths at intermediate regeneration nodes. Actually, constraints (C6) can be omitted, since the flow conservation constraints (C7) are applied to all nodes except for the destination, and thus constraints (C7) indirectly enforce the destination to act as the sink of each flow. Constraints (C8) assign paths and wavelengths to the required lightpaths between all node pairs of the network. Finally, constraints (C9),(C10) and (C11) are exactly the same as constraints (C2), (C3), (C4) of the transparent formulation.

To capture the case where regenerators have different costs than the corresponding transponders, we have to change the minimization objective and defined it as a function of the $f_{sd,ij}^{r}$ variables and not the x_{pw}^{r} variables. The $f_{sd,ij}^{r}$ variables can be used to distinguish between the first (source-initiated) and the intermediate (regenerated) lightpaths of a translucent connection. Thus, using this distinction, the source-initiated lightpaths would contribute the corresponding transponder cost, while the intermediate lightpaths would contribute the corresponding regenerator cost.

Note that the proposed formulation for translucent MLR networks is an extension of the formulation for transparent MLR networks, presented in the previous section. It actually solves a virtual topology problem on top of the transparent planning problem. Other approaches could be followed to formulate the translucent ILP problem, which could be even more efficient, but we have intentionally chosen to extend the transparent formulation so as to have a consistent approach to the whole MLR problem.

Table I presents the number of variables and constraints required in the above ILP formulations. In this table we denote by N =|V| the number of nodes, by L=|E| the number of links, by W=|C| the number of wavelengths, by M=|R| the number of different rates, and by k the number of pre-calculated candidate paths per connection.

TABLE I. NUMBER OF VARIABLES AND CONSTRAINTS

<table>
<tr><td rowspan="2">Transparent</td><td>Variables</td><td colspan="2"><i>x</i> (Boolean): $k\cdot N^{2}\cdot W\cdot M$</td><td colspan="2"><i>u</i> (Boolean): $L\cdot W\cdot M^{2}$</td></tr>
<tr><td>Constraints</td><td>C1: N^{2}</td><td>C2: $L\cdot W$</td><td>C3: $L\cdot W\cdot M^{2}$</td><td>C4: $k\cdot N^{2}\cdot W\cdot M$</td></tr>
<tr><td rowspan="2">Translucent</td><td>Variables</td><td><i>x</i> (Boolean): $k\cdot N^{2}\cdot W\cdot M$</td><td colspan="2"><i>u</i> (Boolean): $L\cdot W\cdot M^{2}$</td><td><i>f</i> (integer): $N^{4}\cdot M$</td></tr>
<tr><td>Constraints</td><td>C5 and C6: N^{2}</td><td>C7: $N^{3}\cdot M$</td><td>C8: $N^{2}\cdot M$</td><td>C9, C10 and C11
as
C2, C3 and C4</td></tr>
</table>

B. Heuristic algorithms

Since the above ILP formulations cannot be solved efficiently for large networks, it is desirable to obtain efficient heuristic algorithms. The heuristic approach we will propose consists of three phases. In the first phase, the algorithm breaks the demands into end-to-end connections of specific rates. In the second phase, the demands are ordered according to some criterion. Then, in the third phase, a heuristic algorithm designed to sequentially establish connections is used. The algorithm serves the connections of the same rate for all commodities of the network one after another, in the ordering identified in the second phase, and then moves to serve the connections for the next rate. In this way, the same rate/format connections are established in closely adjacent wavelengths, reducing cross-rate interference effects. The algorithm works for both transparent and translucent networks with small differences in the first and third phases, which will be indicated in the following paragraphs.

1) Breaking the demands into supported rates

To serve the demand of commodity (s,d), the algorithm first splits its requested capacity Λ_{sd} into the bit rates supported by the network, while minimizing the cost of the used transponders. We denote by f_{sd}^{R} the set of connections for all rates that are used for commodity (s,d), and by f_{sd}^{r} the number of connections of a specific rate $r \in R$. To find f_{sd}^{R} we use the following algorithm.

i) In the case of a transparent network, we pre-calculate for commodity (s,d) a set of k candidate paths P_{sd} (see the discussion in Section III.A.1 regarding the algorithm used). The lengths of the paths define the highest rates that can be used for transmission over these paths. For example, a path $p \in P_{sd}$ with length l_p can use all rates $r \in R$ for which $l_p < D^r$, where D^r is the transmission reach for rate r. For commodity (s,d), we denote by R_{sd} the set of rates that can be supported by all pre-calculated paths P_{sd}. The problem of minimizing the cost of the transponders for (s,d) can be formulated as follows:

$$\text{Minimize } \sum_{r \in R_{sd}} C^r \cdot f_{sd}^{r} \text{ , s.t. } \sum_{r \in R_{sd}} r \cdot f_{sd}^{r} \geq \Lambda_{sd} \, .$$

This problem can be solved optimally for a network that supports three rates, e.g., 10/40/100Gbps, but may be difficult in the general case where the network supports many rates. So, to be more general, instead of solving this problem for the given rates of interest, we use a heuristic algorithm that employs recursion. The recursive algorithm starts by the highest transmission rate, going downwards. At each examined rate r, the algorithm either covers completely the requested capacity with connections of rate r, utilizing $\lceil \Lambda_{sd} / r \rceil$ transponders of rate r, or it uses $\lfloor \Lambda_{sd} / r \rfloor$ transponders of rate r and the remainder $\Lambda_{sd} - \lfloor \Lambda_{sd} / r \rfloor \cdot r$ capacity is covered by transponders of lower rates, using recursion to calculate the cost of the lower rate transponders. The costs of these two options are calculated and the algorithm selects and returns the one with the smaller cost. The pseudocode of this algorithm is presented in Figure 3. The recursive algorithm examines $M(M+1)/2$ breaking options, irrespective of the value of Λ_{sd}, which is polynomial in M.

ii) In the case of a translucent network, we again pre-calculate for commodity (s,d) a set P_{sd} of k candidate paths. However, in this case, the network can utilize regenerators to support rates over paths that are longer than the corresponding transmission reach thresholds. Given a path p, rate r cannot be used over it, if there is a link on p with length more than D^r (we permit regeneration only at node locations); otherwise, rate r can in principle be used on p. Let, a_p^r be the number of regenerators required for rate r, over path p. To identify the minimum number and the placement of regenerators for a given rate over a path, we traverse the links of the path starting from the source. We keep the length of the path in a temporary variable that is initialized to zero. For each link we traverse, we add its length, until the temporary length of the path surpasses D^r. At that point we add a regenerator at the starting node of the last added link and re-initialize the temporary length of the path to be equal to the length of that link. After calculating a_p^r for all pre-calculated paths, we set, for each rate r, $a^r = \min_{p \in P_{sd}} a_p^r$. The problem of minimizing the cost of the transponders for (s,d) can be formulated as follows:

$$\text{Minimize } \sum_{r \in R} [C^r \cdot (1 + a^r)] \cdot f_{sd}^{r} \text{ , s.t. } \sum_{r \in R} r \cdot f_{sd}^{r} \geq \Lambda_{sd} \, ,$$

assuming that the cost of a regenerator is the same as that of a single transponder. If this is not the case, we can modify the above definition to use different cost values for the transponders and the regenerators. Note that this is the same as the related problem for transparent networks, but having transponder cost equal to $C^r \cdot (1+a^r)$.

Again the above problem can be solved easily for a network that supports a small number (e.g., 3) rates, but we also developed a recursive heuristic, similar to the one presented in Fig. 3, to solve it in a quick and efficient way for a larger number of rates.

```
$f_{sd}^{R}$ ← Break_transmission_in_rates ($\Lambda_{sd}$, $R_{sd}$, $C^R$)
    For all $r_i \in R_{sd}$ - starting from higher rates and going downwards
        Let $R_i = \{r_1, r_2 \ldots, r_i\}$ be the set of rates with $r_i$ being the highest rate
        ($f_{sd}^{R_i}$, cost_up_to_$r_i$) = Calculate_recursive_cost($\Lambda_{sd}$, $R_i$, $C^R$)
    EndFor
Return $f_{sd}^{R_i}$ that yielded the smallest cost_up_to_$r_i$

($f_{sd}^{R}$, cost) ← Calculate_recursive_cost ($\Lambda$, $R$, $C^R$)
    Let $r$ be the highest rate in $R$
    Option 1:
        $f_{sd}^{R}$: $f_{sd}^{r}$ = ceiling($\Lambda_{sd}/r$), for all $r' \in R$, $r' \neq r$: $f_{sd}^{r'} = 0$
        cost_r_ceil = $C^r \cdot f_{sd}^{r}$
    Option 2:
        $R' = R - \{r\}$
        $f_{sd}^{r}$ = floor($\Lambda_{sd}/r$), $\Lambda'$ = remainder($\Lambda_{sd}/r$)
        ($f_{sd}^{R'}$, cost_of_remainder) = Calculate_recursive_cost($\Lambda'$, $R'$, $C^R$)
        $f_{sd}^{R} = f_{sd}^{r} \cup f_{sd}^{R'}$
        cost_r_floor = $C^r \cdot f_{sd}^{r}$ + cost_of_remainder
Return the option that yields the smaller total cost (minimum of cost_r_ceil and cost_r_floor)
```

Figure 3: Pseudo-code for breaking a requested demand to the available rates

2) *Ordering the demands and Simulated Annealing*

The heuristic algorithm that will be described in the following paragraph establishes connections, one-by-one, in some particular order. The ordering in which the commodities are served is quite important in this process, and different orderings result in planning solutions of different costs. We implement the following ordering policy:

Highest Demand First (*HDF*) ordering: We order the demands according to their requested rate, and serve first the demand that requires the highest rate.

A number of other policies can be easily defined, based on the length and hop count of the paths used by the demands, and/or other network and traffic parameters. However, since the performance of specific policies depends on many parameters, it is quite difficult to come up with a good ordering policy that would yield good performance for diverse inputs. Thus, to find good orderings, we use a simulated annealing (SimAn) meta-heuristic, which works as follows. We start with the HDF ordering and calculate its cost (viewed as "energy" in the SimAn terminology) by sequentially serving the connections, using the heuristic algorithm described in the following subsection III.B.3 (this is the "fitness function" in the SimAn terminology). For a particular ordering $((s_1,d_1),(s_2,d_2),\ldots,(s_n,d_n))$ of n demands, we define its neighbor as the ordering where (s_i,d_i) is interchanged with (s_j,d_j) for some i and j. To generate a random neighbor we choose pivots (s_i,d_i) and (s_j,d_j) uniformly among the n demands. We use this random neighbor creation procedure and the single demand heuristic as the fitness function in a typical simulated annealing iteration.

3) *Sequential Heuristic Algorithm*

For each link l, we define a Boolean wavelength-rate availability vector

$$\overline{w_l^r} = [\, w_{li}^r \,] = (\, w_{l1}^r,\, w_{l2}^r, \ldots,\, w_{lW}^r \,),$$

whose i^{th} element w_{li}^r is equal to 0 if the i^{th} wavelength of link l is utilized by a connection of rate r, and equal to 1, otherwise. Then the wavelength availability vector $\overline{w_l}$ of link l is given by

$$\overline{w_l} = [w_{li}] = \underset{r\in R}{\&}\, \overline{w_l^r} = [\underset{r\in R}{\&}\, w_{li}^r], \qquad (3)$$

where "&" denotes the Boolean AND operation. Note that the wavelength availability vector $\overline{w_l}$ does not distinguish among different rates, as wavelength-rate availability vector $\overline{w_l^r}$ does.

The wavelength availability vector of a path p consisting of links $l \in p$ can be computed as follows:

$$\overline{W_p} = [W_{pi}] = \underset{l\in p}{\&}\, \overline{w_l} = [\underset{l\in p}{\&}\, w_{li}]. \qquad (4)$$

Thus, the element W_{pi} is equal to 1 if wavelength i is available for transmission over path p. Note that the above equation enforces the wavelength continuity constraint among the links comprising a path.

We start with an "all ones" links wavelength-rate availability vectors, to map an initially completely empty network. We pre-calculate k candidate paths P_{sd}, for each commodity (s,d). We denote by U the set of established lightpaths in the network. Initially, $U=\mathbf{O}$.

We sequentially establish the connections of a specific rate for all commodities and then move to serve the next rate connections. We start from the connections of the highest rate, and then continue to lower rates. For a given rate the commodities are served according to the ordering defined in the second phase of the algorithm. When establishing a lightpath we take into account the lightpaths established up to that point for the previous connections. So for each $r\in R$ and for each commodity (s,d) we establish the corresponding f_{sd}^r calculated in the first phase of the algorithm. After establishing a connection, we update the wavelength-rate availability vectors $\overline{w_l^r}$ for the links l that comprise the chosen path and also update the set of established lightpaths U. Thus, at each step the choices made are stored so as to affect the following connections. Note the algorithm serves the connections of the same rate one after another and assigns wavelengths to them that are quite close to each other. In this way, cross-rate interference is reduced, since the connections of the same rate are not affected by such effects ($m^{r,r}=0$).

We now describe the single demand heuristic algorithm for the case of a transparent network. We want to establish f_{sd}^r lightpaths for (s,d) under the current utilization state of the network, given in the form of the wavelength-rate availability vectors $\overline{w_l^r}$, for all l and r, and the established lightpaths U up to that point. We calculate the wavelength utilization $\overline{W_p}$ of the pre-calculated paths $p \in P_{sd}$, using Eq (3) and (4). For the given rate r we examine only the paths that can support the specific rate, starting from the shortest path. We order the available wavelengths over these paths according to the most used wavelength (MUW) policy. Then, for each available wavelength we check if the corresponding lightpath (identified by the path, wavelength and rate tuple) has acceptable total effective length to support the transmission of the specific rate. To evaluate this we use the wavelength-rate availability vectors $\overline{w_l^r}$ to identify the interfering established lightpaths, and then use Eq. (1) to calculate the effective lengths of the links. Then we use Eq. (2) to calculate the effective length of the lightpath and compare it to the transmission reach threshold D^r. We also check the effect that establishing this new lightpath would have on the already established connections. In particular, we calculate again the effective lengths of the already established connections in U that are affected by the new lightpath and check if the acceptance of the new lightpath will violate their transmission thresholds. This second set of checks is very important, since inserting a new lightpath might turn infeasible some of the already established lightpaths, canceling the previous correct choices made by the algorithm. If all checks are passed then the lightpath is established. Thus, we update U

and $\overline{w_l^r}$ and we also decrease f_{sd}^r so as to know at each point the number of lightpaths of rate *r* that remain to be established for (*s*,*d*). For the given rate *r*, we continue to check the available wavelengths over all paths until either $f_{sd}^r = 0$ or there are no remaining available wavelengths to check. In the latter case, the remaining unserved connections are *blocked*. We continue with establishing lightpaths for the next commodity, that is, the next source-destination pair, in the ordering defined in the second phase of the algorithm. After all commodities are served we move to the next rate and start from the first commodity of the ordering, and so on, until all rates are examined. For a given rate and a given commodity, the single demand heuristic algorithm returns the number of blocked lightpaths and also the updated wavelength-rate availability vectors and the updated set of established lightpaths. Figure 4 presents the pseudocode of the heuristic algorithm for establishing f_{sd}^r connections of rate *r* for commodity (*s*,*d*).

```
(Blocked, $\overline{w_l^r}$ ,U) ← establish_connections ($P_{sd}$, $f_{sd}^r$ , $\overline{w_l^r}$ , U)
For all paths $p \in P_{sd}$,
    calculate utilization $\overline{W_p}$ using Eq. (3) and (4)
EndFor
Blocked=0;
For all paths $p \in P_{sd}$, starting from the shortest path
    If rate r is supported by path p ($l_p<D^r$)
        Order the available wavelengths according to the MUW policy
        For each available wavelength i ($W_{pi}$=0), starting from the most used wavelength (MUW) wavelength
            Check_lightpath_effective_lenght (p,w,r, $\overline{w_l^r}$ ) using Eq (2)
            Define temporary wavelength-rate vector $\overline{w'^r_l}$ equal to $\overline{w_l^r}$ with the addition of the candidate lightpath (p,w,r)
            For all established lightpaths (p',w',r') in U
                Check_lightpath_effective_length (p',w',r', $\overline{w'^r_l}$ )
            Endfor
            If all checks have passed
                Establish lightpath (p,w,r), $f_{sd}^r = f_{sd}^r - 1$
                Insert lightpath (p,w,r) in U and update $\overline{w_l^r}$
            Endif
        EndFor
    EndIf
    Blocked= $f_{sd}^r$ (remaining/unserved lightpaths of rate r)
EndFor
```

Figure 4: Pseudo-code of the algorithm for establishing f_{sd}^r connections of rate *r* for commodity (*s*,*d*).

The above described algorithm is a quick and efficient greedy algorithm that establishes for each demand the lightpaths defined in the first phase of the algorithm. Pre-calculation of paths is used for speeding up the procedure, especially in the simulated annealing variation of the algorithm, where the algorithm is executed multiple times for the different orderings. The algorithm returns the total number of blocked connections for all (*s*,*d*) pairs, for the given number of available wavelengths. Since we are considering the planning problem of a MLR network, we are interested in finding the minimum number of wavelengths that can satisfy the demands with zero blocking, what we call a *zero-blocking solution*. To find zero blocking solutions, we iteratively increase the number of available wavelengths until we can serve all demands without blocking.

In a similar manner, we develop a heuristic algorithm for the case of a translucent network. The difference in the translucent network case is that for each path that we pre-calculate, we also identify the regeneration points for each rate (see the discussion in phase 1 about finding the number of regenerators a_p^r). Thus, an end-to-end connection can be served by a single transparent lightpath, or broken down into a tandem of transparent lightpaths to form a translucent connection. When establishing a transparent lightpath the process is exactly as previously described. When establishing a translucent connection, we establish the series of lightpaths that comprise it. Each lightpath in this series is established as a separate connection, by using Eq. (4) to compute the wavelength availability of the corresponding path, thus enforcing the wavelength continuity constraint along its links. The wavelength continuity constraint is not enforced among the different lightpaths comprising a translucent connection, since the regenerators that are allocated can also perform wavelength conversion. This is why the lightpaths of the series that define the translucent connection are considered as separate and individual demands.

IV. PERFORMANCE RESULTS

We carried out a number of simulation experiments to evaluate the performance of the proposed reach-adapting MLR algorithms. We implemented both the ILP and the heuristic algorithms in Matlab. We used ILOG CPLEX to solve the corresponding ILP problems and Matlab's built in simulated annealing tool. We performed two sets of experiments, so as to evaluate the proposed algorithms in transparent and translucent network settings.

We assumed that the network supports three transmission rates, and in particular 10-, 40- and 100-Gbps, using e.g. OOK and QPSK and DQPSK transmitters, respectively. The transmission reaches D^r were taken equal to 2500, 1500 and 800 km, and the relative costs of the transponders were set to 1, 2.5 and 5.5, respectively, driven from [9] and [10]. Note that, as previously discussed, as we move to higher rate transmitters the cost per bit decreases, but also the transmission reach decreases. Unless otherwise stated, in the simulations we have set the effective length factors $m^{r,r'}$=0.1, for all $r' \neq r$, and the wavelength interfering distance thresholds $I^{r',r}$=2, for all $r' \neq r$. For all the algorithms we used *k*=3 candidate paths.

A. Transparent network experiments

We performed experiments assuming two transparent network topologies: the simple 6 node topology shown in Fig.

5a, and the generic Deutsche Telekom network topology consisting of 14 nodes and 46 directed links shown in Fig 5b.

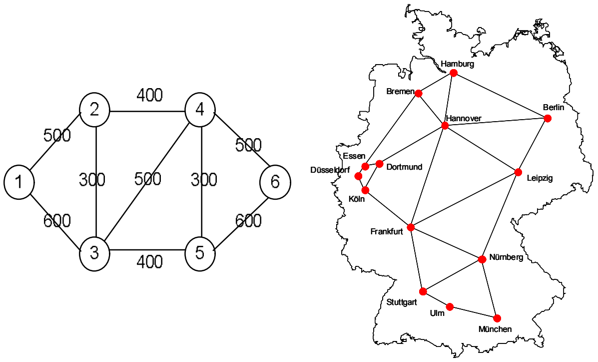

Fig 5: (a) 6-node network topology, (b) the generic DT network topology, with 14 nodes and 23 undirected links.

For the simple 6 node topology (Fig 5a) and for a given traffic load, we randomly created 10 traffic matrices, where the requested capacity for each (*s*,*d*) pair was an exponential random variable with average the given traffic load. We created matrices for loads ranging from 10 to 100-Gbps, with a 15 Gpbs step. Table II reports the average cost, the average number of wavelengths (*W*), and the average running time for the different values of the load and the different algorithms. In particular, we examined the performance of the reach-adapting ILP algorithm for transparent networks (Section III.A.1), the heuristic algorithm (Section III.B), using the HBF ordering policy and also using simulated annealing meta-heuristic (SimAn) with 10, 100 and 1000 iterations.

We also report what we call the "zero-cross-rate interference" and the "worst-cross-rate-interference" ILP cases for these experiments. In the zero-cross-rate-interference case, which corresponds to the best possible case, we assumed that the network is not subject to cross-rate interference, that is, reaches do not decrease by cross-rate interference effects and remain always equal to 2500, 1500 and 800 km, for the 10-, 40- and 100-Gbps transmissions, irrespectively of the utilization of the network. To obtain the results for the zero-cross-rate case we used the reach-adapting ILP algorithm of Section III.B and assumed that $m^{r,r'}=0$, and/or that $I^{r',r}=0$. In the worst-cross-rate-interference case, we assumed that the reaches of the connections are reduced due to always present cross-rate interference. To obtain the results for the worst case we used the reach-adapting ILP algorithm of Section III.B and divided the default transmission reach bounds D^r by 1.2 (remember that we have two interfering rates for each rate under examination), which is equal to the worst case increase of the effective length, and also set $m^{r,r'}=0$, and/or that $I^{r',r}=0$. In other words, the zero-cross-rate-interference and the worst-cross-rate-interference cases correspond to a typical MLR algorithm with different transmission reach bounds, in which the transmission reaches do not adapt. Note that, in the literature, heuristic algorithms assuming the worst-cross-rate-interference have been examined [8][15]. Thus, the comparison of the proposed algorithms with the worst case is indicative of the improvement of the proposed algorithm over previous works. On the other hand, the comparison of the proposed algorithms with the zero-cross-rate-interference case helps us quantify the degree to which the proposed algorithms can find solutions that avoid the cross-rate interference effects.

From Table II we can observe that the optimal reach-adapting ILP algorithm was able to track solutions with average times up to a few seconds (78 sec for load=100-Gbps). The performance of the proposed heuristic is quite good, and was able, in all cases, to find solutions with transponders cost equal to that reported by the optimal ILP. This shows that the 1st phase of the heuristic algorithm (Section III.B.1) succeeds in dividing the connections to the optimal number of lightpaths. These lightpaths are then established in the third phase of the algorithm (Section III.B.3), using the available wavelengths. We can see that the heuristic algorithm requires different number of wavelengths to find zero blocking solutions, depending on the ordering that is used (Section III.B.2). When using simulated annealing (SimAn) with 1000 iterations (1000 corresponds to the different orderings that are examined), the number of wavelengths required to find zero blocking solutions were equal to that of the ILP algorithm. The running time of SimAn with 1000 iterations is comparable to that of the ILP algorithm, while as stated above, the wavelength and transponder cost performance are the same. As the number of SimAn iterations decreases, the number of wavelengths required to find zero blocking solutions increases. The case where we use only one ordering, and in particular the HDF ordering, without employing simulated annealing, has obviously the worst performance in terms of the number of wavelengths required to serve the traffic. As expected, the running time of the heuristic algorithm decreases as the number of SimAn iterations decreases. Thus, using SimAn we obtain a tradeoff between the running time and the wavelengths performance. At least for this small network, the results show that even with few SimAn iterations (e.g., 100) we can have wavelength performance quite close to the optimal solution found by the ILP and very low average running times.

Also, from Table II we can observe that the cost and the number of wavelengths reported for the reach-adapting ILP algorithm are equal to those reported for the zero-cross-rate-interference case. This shows that the proposed reach-adapting ILP algorithm (and the heuristic) is able to assign wavelengths effectively to the connections so as to absorb cross-rate interference. On the other hand, the number of wavelengths required assuming worst-case cross-rate interference is higher, and so is the transponders cost. This is because under the worst-cross-rate interference scenario the effective lengths of the paths are larger, or equivalently the transmission reaches are shorter. This results in many paths in the network being considered infeasible (even though they are not), negatively impacting the wavelength and cost performance of the network. The running time of the reach-adapting ILP algorithm compared to the zero- and worst-cross-rate-interference ILP cases is higher, due to the additional active constraints (constraints (C3) and (C4)) that formulate the adaptation of the effective lengths.

TABLE II: PERFORMANCE OF THE TRANSPARENT ALGORITHMS FOR THE SMALL NETWORK (10 TRAFFIC MATRICES PER LOAD)

Load (Gbps)	**10**			**25**			**40**			**55**			**70**			**85**			**100**		
Algorithms	**Average cost**	**Average W**	**Average time (s)**	**Average cost**	**Average W**	**Average time (s)**	**Average cost**	**Average W**	**Average time (s)**	**Average cost**	**Average W**	**Average time (s)**	**Average cost**	**Average W**	**Average time (s)**	**Average cost**	**Average W**	**Average time (s)**	**Average cost**	**Average W**	**Average time (s)**
HDF	44.05	4.9	0.08	69.10	5.7	0.11	95.85	6.9	0.12	119.60	7.4	0.13	145.15	9.4	0.19	170.90	10.3	0.26	196.00	11.9	0.27
SimAn (10 iter)	44.05	4.7	0.16	69.10	5.1	0.42	95.85	6.3	0.27	119.60	7.2	0.40	145.15	8.6	0.91	170.90	9.7	0.99	196.00	11.1	1.13
SimAn (100 iter)	44.05	4.6	0.31	69.10	5.1	0.41	95.85	6.3	0.24	119.60	7.1	2.29	145.15	8.4	8.81	170.90	9.5	9.95	196.00	10.7	20.91
SimAn (1000 iter)	44.05	4.6	3.57	69.10	5.1	2.69	95.85	6.3	3.43	119.60	7.1	4.45	145.15	8.4	10.17	170.90	9.4	12.12	196.00	10.7	78.38
ILP	44.05	4.6	2.29	69.10	5.1	2.55	95.85	6.3	3.11	119.60	7.1	4.07	145.15	8.4	4.27	170.90	9.4	6.17	196.00	10.7	11.40
Zero cross-rate (ILP)	44.05	4.6	2.29	69.10	5.1	2.55	95.85	6.3	3.11	119.60	7.1	4.07	145.15	8.4	4.27	170.90	9.4	6.17	196.00	10.7	11.40
Worst cross-rate (ILP)	44.45	5.0	2.25	71.10	6.6	2.66	98.95	8.5	3.61	124.10	10.3	3.64	150.90	12.8	4.76	178.40	14.6	5.93	204.60	16.9	9.12

TABLE III: PERFORMANCE OF THE TRANSPARENT ALGORITHMS FOR THE DT NETWORK

Load (times the reference matrix)	**1**			**2**			**3**			**4**			**5**			**6**			**7**			**8**		
Algorithms	**cost**	**W**	**time (sec)**	**cost**	**W**	**time (sec)**	**cost**	**W**	**time (sec)**	**cost**	**W**	**time (sec)**	**cost**	**W**	**time (sec)**	**cost**	**W**	**time (sec)**	**cost**	**W**	**time (sec)**	**cost**	**W**	**time (sec)**
HDF	342.0	24	0.90	501.0	25	1.10	640.0	26	1.09	797.0	29	2.93	955.0	32	3.02	1103.0	35	3.27	1252.0	40	7.29	1408.5	41	15.84
SimAn (10 iter)	342.0	22	0.92	501.0	25	1.35	640.0	26	1.10	797.0	29	3.59	955.0	32	4.59	1103.0	34	5.59	1252.0	37	20.76	1408.5	40	26.77
SimAn (100 iter)	342.0	22	2.15	501.0	24	3.40	640.0	25	10.39	797.0	27	16.68	955.0	30	17.08	1103.0	33	23.02	1252.0	37	32.17	1408.5	38	43.14
SimAn (1000 iter)	342.0	21	81.11	501.0	23	64.96	640.0	23	153.45	797.0	26	165.01	955.0	29	195.44	1103.0	31	164.23	1252.0	35	223.62	1408.5	36	328.52
ILP	342	20	280.2	501.0	22	2028	640.0	23	4900	797.0	26	7200	-	-	-	-	-	-	-	-	-	-	-	-
Zero cross-rate (ILP)	342	20	149.4	501.0	22	1646	640.0	23	2892	797.0	26	7200	-	-	-	-	-	-	-	-	-	-	-	-
Worst cross-rate (ILP)	342	20	164.5	501.0	24	1845	640.0	25	3127	797.0	30	7200	-	-	-	-	-	-	-	-	-	-	-	-

Figure 6 presents the average number of wavelengths required to find zero blocking solutions with the reach-adapting ILP algorithm for load equal to 70 Gbps and for different values of the interference distance $I^{r,r'}$ and effective length factors $m^{r,r'}$. From this graph we can observe that the proposed algorithm is able to exploit the wavelength domain and avoid cross-rate interference even for high values of the interference distance parameter $I^{r,r'}$. In particular, the average number of wavelengths does not change for values up to $I^{r,r'}=10$ and remains equal to 8.4 which is the average number of wavelengths for the best case ($I^{r,r'}=0$ – zero-cross-rate). Even for high values of $I^{r,r'}$, where it is no longer possible to use the wavelength space to absorb cross-rate interference, the increase in the average number of wavelengths required is not that significant. Note that the case that $I^{r,r'}$ is equal to the number resembles the problem setting previously examined in [17].

Also, from Fig. 6 we can observe that network performance deteriorates significantly even for small increases of the effective length factor $m^{r,r'}$. Thus, the dependence of the performance on $m^{r,r'}$ is more significant than the dependence on $I^{r,r'}$. The effective length factor $m^{r,r'}$ affects directly the decrease of the transmission reach (increase of the effective length), and large values of $m^{r,r'}$ turn many paths unusable, if they are subject to cross-rate interference. Thus, at high values of $m^{r,r'}$, the algorithm spreads the lightpaths, leaving wavelength space between them, to avoid the cross-rate interference effects, increasing in this way significantly the average number of wavelengths required to find the solution. It is worth noting that the running times of the reach-adapting ILP algorithm deteriorate as the parameters $I^{r,r}$ and $m^{r,r'}$ increase. High values of these parameters correspond to stronger cross-rate interference effects. The problem becomes more complicated and the algorithm has to search many more options to avoid these stronger interference effects, resulting in increased running time. Some instances of the problem became intractable and thus the results presented in Fig. 6 have been produced by stopping the ILP algorithm after running for 2 hours per instance.

Next we performed experiments for the DT network (Figure 5b). We used a realistic traffic matrix for year 2009 and traffic predictions for the following years (please refer to deliverable D2.1 in www.diconet.eu/deliverables.asp). In this traffic matrix the capacity requirements among the demands range from 4.5 up to 47 Gbps, with an average of 15 Gbps. We uniformly scaled up the reference traffic matrix to obtain traffic matrices up to 8 times larger than that, corresponding to the expected traffic growth in the following few years. Table III shows the corresponding results. In this set of experiments for the DT network, the optimal ILP algorithm could not track solutions for high loads, and in particular for loads higher than 4 times the reference traffic matrix, within reasonable time, i.e., within 2 hours. The same holds for the results obtained under the zero- and worst-cross-rate-interference assumptions, for which we also were not able to track solutions within 2 hours for loads higher than 4. Note that for loads lower than 4, the majority of traffic is served through 10- and 40-Gbps connections, while for higher loads 100-Gbps connections start to appear, complicating the problem to a greater extend. Even for the low loads for which we obtained optimal ILP solutions, we can see that the number of wavelengths required by SimAn is quite close to the optimal solution, at least when 1000 iterations were used. As expected, the wavelength performance of the heuristic algorithm deteriorates while the running time improves as the number of iterations decrease. Again the running time of the heuristic can be controlled by the number of SimAn iterations that are performed.

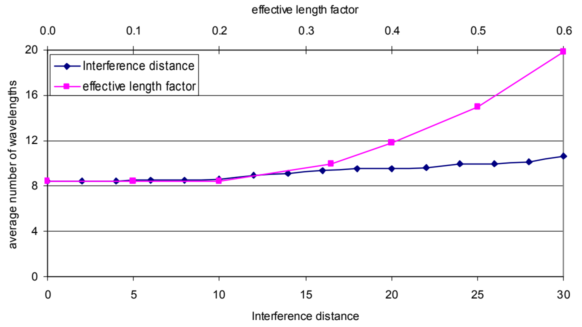

Fig. 6: Average number of wavelengths as a function of the interference distance $I^{r,r'}$ and the effective length factor $m^{r,r'}$.

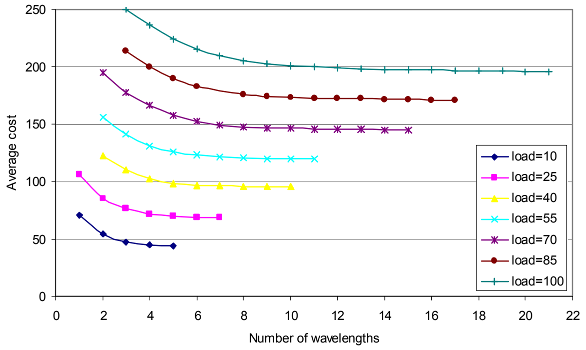

Fig. 7: Average transponders cost as a function of the number of available wavelengths for different loads found by the ILP translucent algorithm for the small network topology.

B. *Translucent network experiments*

We now turn our attention to the case of translucent networks and evaluate the performance of the corresponding ILP (Section III.A.2) and heuristic algorithms (Section III.B).

We start by reporting the results obtained for the small 6-node network of Fig. 5a. This is a network that has relatively small link distances that can be planned transparently quite efficiently, as done in the previous section (see Table II). We used the same network and traffic, in order to observe the difference between the functioning of the transparent and the translucent algorithms.

In Fig. 7 we report the average transponder cost as a function of the number W of available wavelengths in the network, found by the translucent reach-adapting ILP algorithm. When the number of available wavelengths is small for a given load, the translucent algorithm decides to utilize higher rate transponders and more regenerators, so as to save wavelengths that are the constraining and scarce network resource in this case. Thus, for low number of wavelengths, the algorithm utilizes a high number of expensive transponders (including the regenerators) and yields a high total transponder cost. As the number of wavelengths increase, the algorithm utilizes more efficiently the heterogeneous transponders/rates, reducing the cost, which converges to that of the transparent algorithm. Indeed, for the number of wavelengths reported for the transparent ILP algorithm in Table II, the translucent algorithm finds exactly the same solutions. Note that the transparent network is a special case of the translucent network in which no regenerators are employed. In a network that can be planned in a transparent way, such as the one that is considered in this set of experiments, the optimal cost solution would always be to plan the network transparently without the use of regenerators. When the available wavelengths are sufficient to accommodate transparently the traffic, the optimal translucent algorithm converges and finally produces a transparent solution that has the minimum cost. This behavior has been verified in this set of experiments. When the number of wavelengths is small, however, the translucent algorithm explores solutions that use more expensive higher rate transponders that have higher cost but save the scarce wavelength resources. Note that, as Fig. 7 shows, the average cost converges slowly and the values are quite close to the optimal –transparent- solution many wavelengths before the optimal solution is found. Thus, it seems very efficient to use the translucent algorithm to obtain e.g. 20% reduction in the required wavelengths with an increase of about 2% in the transponders cost. In any case the reductions depend on the traffic and network parameters so they have to be evaluated for each problem instance separately.

Next, we doubled the lengths of the links of the small network of Fig 5a so as to turn it to a translucent network. Table IV reports the average cost, the average number of wavelengths, and the average running times for the different load values and the different algorithms. We can again observe that the performance of the heuristic algorithm with SimAn is quite close to that of the optimal ILP algorithm. In all cases the heuristic algorithm is able to find the same total transponder cost (1st phase of the heuristic), while using a close to optimal number of wavelengths. However, compared to the transparent case, where we had a total match in terms of the number of wavelengths required between the heuristic and the optimal ILP algorithm, we can observe that in the translucent case the heuristic algorithm is not that close to the optimal solution. This is because the translucent problem is quite more complicated than the transparent one, since it also includes choices for the allocation of the regeneration points. The heuristic translucent algorithm assumes a specific regeneration placement for each connection (the placement that minimizes the number of regenerators for that connection– see discussion in 1st and 3rd phases of the heuristic), which might not always be optimal for the concurrent establishing of all the lightpaths in the network. Remember that regenerators function also as wavelength converters, so, at regeneration points the wavelength continuity constraint is relaxed, resulting to a smaller number of required wavelengths. The optimal reach-adapting ILP algorithm that searches among all possible regeneration options for all connections, can find better solutions, at least for this small network where we can track optimal solutions. Also, from Table IV we observe that the performance of the reach-adapting ILP algorithm is identical to that of the zero-cross-rate-interference ILP case, in terms of transponders cost and number of required wavelengths. Thus, the proposed reach-adapting translucent algorithms (ILP and heuristic) are able to absorb the cross-rate interference among the connections by intelligently assigning wavelengths to them. On the other hand, the performance of the ILP worst-cross-

TABLE IV: PERFORMANCE OF THE TRANSLUCENT ALGORITHMS FOR THE SMALL NETWORK WITH DOUBLE LENGTHS (10 TRAFFIC MATRICES PER LOAD)

Load (Gbps)	**10**			**25**			**40**			**55**			**70**			**85**			**100**		
Algorithms	**Average cost**	**Average W**	**Average time (s)**	**Average cost**	**Average W**	**Average time (s)**	**Average cost**	**Average W**	**Average time (s)**	**Average cost**	**Average W**	**Average time (s)**	**Average cost**	**Average W**	**Average time (s)**	**Average cost**	**Average W**	**Average time (s)**	**Average cost**	**Average W**	**Average time (s)**
HDF	45.05	5.8	0.26	76.3	10.3	0.42	108.95	14.7	0.78	139.45	18.9	0.82	172.00	25.2	2.30	204.45	30.3	2.48	236.00	38.3	3.32
SimAn (10 iter)	45.05	5.7	0.32	76.3	10.3	0.61	108.95	13.8	1.61	139.45	17.6	2.26	172.00	22.9	3.23	204.45	28.4	3.67	236.00	37.8	4.29
SimAn (100 iter)	45.05	5.4	0.39	76.3	9.6	2.91	108.95	13.9	2.99	139.45	18.6	10.74	172.00	23.4	13.55	204.45	27.8	26.56	236.00	37.3	73.69
SimAn (1000 iter)	45.05	5.4	1.64	76.3	9.4	4.94	108.95	13.8	9.63	139.45	18.3	13.89	172.00	23.2	28.81	204.45	27.5	49.95	236.00	37.2	81.24
ILP	45.05	5.4	9.75	76.3	9.4	22.43	108.95	13.6	45.14	139.45	18.0	75.68	172.00	22.8	131.31	204.45	27.1	220.09	236.00	36.8	634.18
Zero cross-rate (ILP)	45.05	5.4	6.32	76.3	9.4	8.45	108.95	13.6	8.92	139.45	18.0	12.15	172.00	22.8	18.43	204.45	27.1	21.46	236.00	36.8	35.40
Worst cross-rate (ILP)	45.05	5.6	4.52	76.4	9.7	6.75	109.55	14.3	8.94	141.95	19.6	9.45	181.50	24.9	14.94	214.60	30.8	22.12	252.00	41.2	33.40

TABLE V: PERFORMANCE OF THE TRANSLUCNET HEURISTIC ALGORITHM FOR THE GEANT NETWORK

Load (times the reference matrix)	**1**			**2**			**3**			**4**			**5**			**6**			**7**			**8**		
Algorithms	**cost**	**W**	**time (sec)**	**cost**	**W**	**time (sec)**	**cost**	**W**	**time (sec)**	**cost**	**W**	**time (sec)**	**cost**	**W**	**time (sec)**	**cost**	**W**	**time (sec)**	**cost**	**W**	**time (sec)**	**cost**	**W**	**time (sec)**
HDF	387.5	31	30.20	435.0	35	31.32	508.0	36	95.26	574.0	42	56.49	617.0	44	132.82	713.0	48	141.72	780.5	51	100.64	817.5	53	113.81
SimAn (10 iter)	387.5	29	109.71	435.0	31	96.13	508.0	35	148.93	574.0	41	175.66	617.0	42	206.95	713.0	46	218.55	780.5	47	262.46	817.5	50	273.17
SimAn (100 iter)	387.5	24	214.45	435.0	28	579.10	508.0	34	793.64	574.0	37	939.72	617.0	40	1040.97	713.0	44	1463.67	780.5	47	1724.60	817.5	49	2245.92
SimAn (1000 iter)	387.5	22	3189	435.0	26	4819	508.0	32	8231	574.0	35	11440	617.0	38	16320	713.0	43	19322	780.5	47	22321	817.5	48	26418

rate-interference algorithm is inferior, since it results in increased transponders cost and more required wavelengths.

Finally, we also performed experiments assuming a realistic translucent network topology. In particular we used the GEANT-2 network topology and a realistic reference traffic matrix, which we again scaled it up to 8 times (please refer to deliverable D2.1 in www.diconet.eu/deliverables.asp for the topology and the reference traffic matrix). Table V reports the corresponding results only for the heuristic algorithm, since the ILP algorithm was not able to track solutions within a 2 hours limit. From this table we can verify that the proposed heuristic algorithms are able to find solutions for realistic networks and traffic matrices in reasonable time. Again using SimAn we are able to control the running time of the algorithm by trading-off wavelength performance.

V. CONCLUSIONS

We presented algorithms for planning mixed line rate (MLR) optical transport networks. In MLR systems the transmission reach can differ significantly from those typically used in single rate (SLR) systems. We modeled the cross-rate interference due to the different modulation formats/rates used in a MLR system, by defining an effective length metric that helps us adapt the transmission reach of the connections based on the utilization state of the network. We used the effective length metric to formulate the adaptive reach planning problem for transparent and translucent MLR optical networks. We initially presented optimal ILP algorithms for the MLR planning problem for both transparent and translucent networks. We also gave sequential heuristic algorithms, proposed a specific ordering policy and also used simulated annealing to find even better orderings. Our results indicated that the proposed algorithms can efficiently utilize the wavelength domain to absorb cross-rate interference effects. The algorithms assign wavelengths to the lightpaths so as to reduce or avoid cross-rate interference and yield solutions that have the same transponder cost and utilize the same number of wavelengths as if no cross-rate interference was present in the network. The performance of the proposed reach-adapting algorithms was shown to be superior to that of other planning algorithms that are based on the worst transmission reach assumption.

REFERENCES

[1] T. E. Stern, K. Bala, "Multiwavelength Optical Networks: A Layered Approach", Prentice Hall, 1999.

[2] X. Liu, S. Chandrasekhar, "High Spectral-Efficiency Mixed 10G/40G/100G Transmission", OFC 2008.

[3] T. J. Xia et. al., "Multi-Rate (111-Gb/s, 2x43-Gb/s, and 8x10.7-Gb/s) Transmission at 50 GHz Channel Spacing over 1040-km Field-Deployed Fiber", ECOC 2008.

[4] T. Wuth, M. W. Chbat, V. F. Kamalov, "Multi-rate (100G/40G/10G) Transport Over Deployed Optical Networks", OFC 2008.

[5] G. P. Agawal, "Fiber-Optic Communication Systems", 3rd ed, Wiley - Interscience, 2002.

[6] S. Azodolmolky, et al., "A Dynamic Impairment Aware Networking Solution for Transparent Mesh Optical Networks", IEEE Communications Magazine, 47 (5), pp. 38-47, 2009.

[7] K. Christodoulopoulos, K. Manousakis, E. Varvarigos, "Offline Routing and Wavelength Assignment in Transparent WDM Networks", IEEE/ACM Transactions of Networking, Vol. 18 , Is. 5, 2010.

[8] C. Meusburger, D. A. Schupke, A. Lord, "Optimizing the Migration of Channels with higher Bit-rates", IEEE/OSA Journal of Lightwave Technology, vol 28, no 5, pp. 608-615, 2010.

[9] A. Nag, M. Tornatore, B. Mukhergee, "Optical Network Design with Mixed Line Rates and Multiple Modulation Formats", IEEE/OSA Journal of Lightwave Technology, vol 28, no 4, 2010.

[10] A. Nag, M. Tornatore, "Transparent vs. Translucent Optical Network Design with Mixed Line Rates", OFC 2009.

[11] M. Batayneh, D. Schupke, M. Hoffmann, A. Kirstaedter, B. Mukherjee, “On Routing and Transmission-Range Determination of Multi-Bit-Rate Signals Over Mixed-Line-Rate WDM Optical Networks for Carrier Ethernet”, IEEE/OSA Transactions on Networking.

[12] A. Bocoi, M. Schuster, F. Rambach, D. A. Schupke, C. A Bunge, B. Spinnler, “Cost Comparison of Networks using Traditional 10 and 40 Gb/s Transponders versus OFDM Transponders”, OFC 2008.

[13] O. Rival, A. Morea, J.-C. Antona, “Optical Network planning with Rate-Tunable NRZ Transponders”, ECOC 2009.

[14] N. Sambo, M. Secondini, F. Cugini, G. Bottari, P. Iovanna, F. Cavaliere, P. Castoldi, “Modeling and Distributed Provisioning in 10–40–100-Gb/s Multirate Wavelength Switched Optical Networks”, IEEE/OSA Journal of Lightwave Technology, vol. 29 is. 9, pp. 1248 – 1257, 2011.

[15] O. Rival, G. Villares, A. Morea, “Impact of Inter-Channel Nonlinearities on the Planning of 25-100Gb/s Elastic Optical Networks”, IEEE/OSA Journal of Lightwave Technology, 2011.

[16] R. Aparicio-Pardo, P. Pavon-Marino, S. Zsigmond, "Mixed line rate virtual topology design considering non linear interferences between amplitude and phase modulated channels”, to be published in Photonic Network Communications.

[17] K. Christodoulopoulos, K. Manousakis, E. Varvarigos, “Adapting the Transmission Reach in Mixed Line Rates WDM Transport Networks”, ONDM 2011.

Konstantinos Christodoulopoulos (kchristodou@ceid.upatras.gr) received a Diploma in Electrical and Computer Engineering from the National Technical University of Athens, Greece, in 2002, the M.Sc. in Advanced Computing from Imperial College London, U.K., in 2004, and the Ph.D. degree from the Computer Engineering and Informatics Department, University of Patras, Greece, in 2009. His research interests are in the areas of protocols and algorithms for optical networks and grid computing.

Konstantinos Manousakis (manousak@ceid.upatras.gr) received the Diploma from the Computer Engineering and Informatics Department, University of Patras, Greece, in 2004, and the M.Sc. degree in Computer Science and Engineering from the Computer Engineering and Informatics Department in 2007. He is currently a Ph.D. candidate at the same department. His research activities focus on optimization algorithms for high speed and optical networks.

Emmanouel (Manos) Varvarigos (manos@ceid.upatras.gr) received a Diploma in Electrical and Computer Engineering from the National Technical University of Athens in 1988, and the M.S. and Ph.D. degrees in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology in 1990 and 1992, respectively. He has held faculty positions at the University of California, Santa Barbara (1992-1998, as an Assistant and later an Associate Professor) and Delft University of Technology, the Netherlands (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, 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.