Impact of Wavelength and Modulation Conversion on Translucent Elastic Optical Networks Using MILP

X. Wang, M. Brandt-Pearce, and S. Subramaniam

***Abstract*—Compared to legacy wavelength division multiplexing networks, elastic optical networks (EONs) have added flexibility to network deployment and management. EONs can include previously available functionality, such as signal regeneration and wavelength conversion, as well as new features such as finer-granularity spectrum assignment and modulation conversion. Yet each added feature adds to the cost of the network. In order to quantify the potential benefit of each functionality, we present a link-based mixed-integer linear programming (MILP) formulation to solve the optimal resource allocation problem. We then propose a recursive model in order to either augment existing network deployments (spectrum and regenerators) or speed up the resource allocation computation time for larger networks with higher traffic demand requirements than can be solved using an MILP. We show through simulation that systems equipped with signal regenerators or wavelength converters require a notably smaller total bandwidth, depending on the topology of the network. We also show that the suboptimal recursive solution speeds up the calculation and makes the running time more predictable, compared to the optimal MILP.**

***Index Terms*—Elastic optical networks; Mixed-integer linear programming; Modulation conversion; Regeneration placement; Spectral resource allocation.**

I. Introduction

Increasing traffic volume and growing heterogeneity of bandwidth requirements have pushed the development of optical transport networks. Using wavelength division multiplexing (WDM) technology, spectrum usage has greatly increased by allowing multiple-line-rates and traffic grooming [1]. Yet the current static spectrum allocation approach, which assigns a fixed 50 GHz (or even 12.5 GHz) channel to each user, is unable to handle increasing traffic heterogeneity. Elastic optical networks (EONs), on the other hand, provide flexibility in both bandwidth assignment (using subchannel granularity and superchannel assignments) and lightpath reconfigurability not available in fixed-grid WDM. As the technology matures, additional functionality such as modulation selection and conversion can be added, with the hope of further increasing the spectral efficiency. When major additions in physical layer features are being considered, the network design should be reexamined to determine the realized benefit gained by their implementation. This paper presents an optimal routing, regeneration, and spectrum allocation formulation that is then used to evaluate the merit of wavelength and modulation conversion on EONs affected by physical layer impairments.

The design of transport networks includes the placement and assignment of all physical resources, such as optical fiber and electronic devices (transponders, high-speed electrical signal regeneration circuits, etc.). The goal is usually to minimize the capital expenditure while fulfilling certain traffic accommodation expectations. One common way to solve this problem is to address it as a multi-commodity assignment by pairing the physical resources with traffic demands in order to minimize the resources used by each demand. For example, a traffic demand can be assigned the shortest route in order to reduce the cost. Such design principles have been used to develop many heuristic algorithms for network design [2–4]. Although these algorithms are computationally simple, they often yield poor performance when the problem becomes complex and consideration cannot be given to all influencing factors. Another approach is to formulate the resource allocation as an optimization problem with physical and network layer constraints and use linear programming (LP) to solve it. The available network resources become the LP design variables. Unlike arbitrary multi-commodity assignment problems, network design often requires its variables to be integer or Boolean, which leads to a mixed-integer linear programming (MILP) formulation. This significantly increases the computational complexity, not providing an approach that can scale to address larger networks. However, for small networks and few traffic demands, the MILP can be solved in reasonable time, and results in an optimal solution, unlike heuristic algorithms. It does this without requiring a complete understanding of the relationship between the multiple design factors, as heuristic algorithms often do [5].

In this paper we develop an MILP design method for EONs. Our formulation can implement modulation scheme selection, mid-lightpath modulation conversion (MC) and/or wavelength conversion (WC), and regeneration device

Manuscript received January 8, 2015; revised May 23, 2015; accepted May 24, 2015; published June 23, 2015 (Doc. ID 231575).

X. Wang and M. Brandt-Pearce (e-mail: mb-p@virginia.edu) are with the Charles L. Brown Department of Electrical and Computer Engineering, University of Virginia, Charlottesville, Virginia 22904, USA.

S. Subramaniam is with the Department of Electrical and Computer Engineering, The George Washington University, Washington DC 20052, USA.

http://dx.doi.org/10.1364/JOCN.7.000644

allocation (to satisfy either a quality of service constraint or a conversion function). MILP has previously been used to solve the resource assignment optimization problem in optical networks [6–8]. However, to the best of our knowledge, no published MILP solution has included these flexibilities in an optimal way for designing EONs.

Acknowledging the limitations of the MILP approach for solving realistically scaled problems due to its computational complexity, we envision the following two direct uses for our model. The MILP can be solved for a small network, which usually contains fewer than 10 nodes and 20 links, to quantify the potential benefit that can be obtained by implementing a particular feature, such as modulation conversion, without introducing artificial limitations imposed by a suboptimal resource allocation algorithm. This is quite common in recent literature. In [9] an ILP is used to solve a four- or five-node ring network and a six-node random network. An ILP is also used on ring networks in [10], with the number of nodes ranging from 6 to 24. In [7], their ILP is used on six-node, 16-link and nine-node, 26-link mesh networks. For larger topologies, solutions are obtained using heuristic algorithms. Our approach can also be used on realistic-size networks to solve for the optimal resource allocation of only a few traffic demands at a time. For the offline resource allocation problem (static network), we can partition the whole traffic matrix into small submatrices, and solve the assignment problem for the submatrices in a sequential manner. By doing this, we are able to greatly reduce the overall execution time, in exchange for obtaining a suboptimal solution. In the paper we discuss the trade-off between complexity and optimality for this approach that we call the recursive solution. For dynamic networks, we can use the recursive MILP to allocate resources for one or a few new connection requests given the current state of the deployed network, as we proposed for fixed-grid WDM systems in [11].

The rest of the paper is organized as follows: Section II introduces the network and node structure of our model and describes the advanced signal processing functionalities for EON that we consider. Section III explains how we implement the new functionalities with our MILP formulation. Section IV develops our recursive MILP implementation that balances optimality and complexity. Section V presents numerical simulation results collected by solving the design problem using our formulation. Finally, conclusions are given in Section VI.

II. Network Description

We are interested in long-haul transport optical networks such as the NSF network shown in Fig. 1 that covers the whole US mainland area. In order to investigate the effects of topology on the questions of interest, we also consider a symmetric network illustrated in Fig. 2 with the same number of nodes and network diameter as the NSF network (the link length is set to be 1330 km for this purpose) but with a different number of links and connectivity. The network must support a given traffic load as identified by a demand matrix specifying the source node, destination node, and throughput requested of each traffic demand.

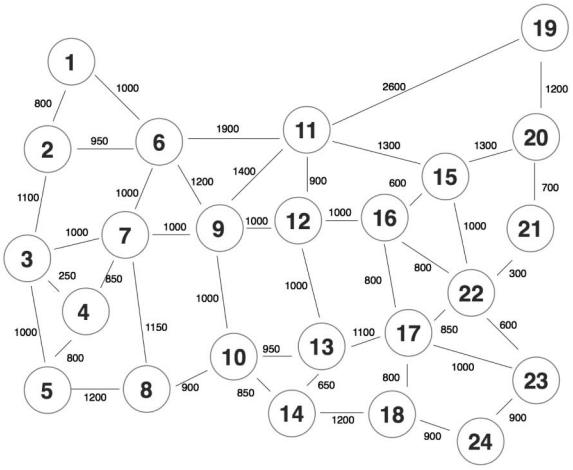

Fig. 1. NSF-24 network. The number on each link represents the physical length of the link in kilometers.

The network we investigate is fully elastic in that the spectrum assignment is flexible over a very fine grid so that the system adapts to very heterogeneous traffic. The central frequency and the bandwidth of each channel can both be flexibly tuned. For mathematical simplicity, we model each channel as having a continuous-valued spectrum (i.e., not slotted) as an approximation to a fine grid. For our approach to be applied to a currently realistic implementation (e.g., the spectrum assignment follows a 6.25 GHz or smaller fixed grid), our result can be readily used to generate integer solutions by implementing methods such as integer rounding or local optimization to search for integer solutions [12]. In practice flexibility is achieved by using flexible-bandwidth and modulation-adaptive transponders, which are the fundamental building blocks of EONs. They transmit the signal on the fiber, performing all electrical–optical and optical–electrical conversions. Without loss of generality, we assume a node structure with transponders based on orthogonal frequency division multiplexing (OFDM) technology. In our formulation we do not

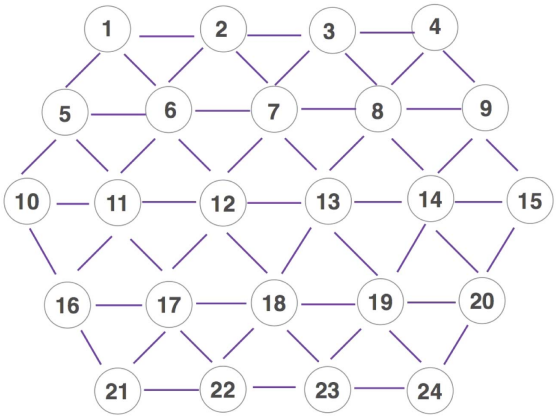

Fig. 2. Symmetric-24 network.

explicitly include guard bands, assuming they are subsumed within the spectrum assigned; they could easily be added to the model. The configuration of each transponder can be managed adaptively by the control plane as described in, for example, [13].

In this paper the resources being allocated to each demand include connected fiber links forming a route, spectrum on each link, signal regeneration devices on connecting nodes, and wavelength and/or modulation conversion at those nodes. (The MILP formulation given below could be easily modified to include other network attributes and constraints.) The functionality that these resources provide is described in the following sections.

A. Signal Regeneration

When optical signals traverse the network, they often suffer from physical impairments such as signal loss, noise, dispersion, and nonlinear effects (ignoring filtering impairments when guard bands are implemented). For networks with small physical dimensions (such as local or metro area networks), the impairments can be ignored. In long-haul transport-scale networks, the impairments need to be accounted for in order to maintain acceptable signal quality [often referred to as the quality of transmission (QoT)] at the receiving end. Due to the fact that impairments originate from many physical phenomena that accumulate over distance, and given that they usually depend on the network state, for simplicity a conservative constraint on the length of fiber a signal can traverse before regeneration, called the transmission reach (TR), is often used to guarantee the QoT. If the source–destination distance on a route exceeds the TR, then regeneration is needed to reduce the impact of physical impairments. The optical signal undergoes optical–electrical–optical (OEO) conversion at an intermediate node, and the regeneration [including reamplifying, retiming, reshaping (3R)] is performed in the electrical domain. In our work, we assume that 3R regeneration only occurs at intermediate nodes, not along the fiber links. To ensure proper QoT, the length of each transparent segment (part of the lightpath that has no intermediate regeneration) must be upper-bounded by the TR. As physical impairments depend on the bit rate and modulation scheme (spectral efficiency) used for a demand, so does the TR for the route used by that demand.

In a typical optical network, not all nodes are equipped with regenerators to save on capital expenditures as well as maintenance and other operational expenditures. OEO conversion is expensive since it requires high-speed electronic equipment, and therefore regeneration needs to be carefully and conservatively planned. The regeneration at intermediate nodes requires the optical signal to undergo OEO conversion through transponders and electrical regeneration (baseband processing). Reference [14] shows the fundamental structure of a regeneration node. The main cost of the infrastructure comes from the transponders, yet regeneration devices are also necessary, which implies that not all nodes with free transponders can function as regeneration nodes.

B. Wavelength Conversion

At a regeneration point, WC is often possible. We assume that the transponder can modulate the signal to an arbitrary new wavelength when converting the signal back from the electrical domain to the optical domain.

At nodes where no regenerators are available and therefore no WC is possible, the same wavelength must be used on both sides of the node on the lightpath. This so-called wavelength continuity constraint causes spectrum fragmentation as small portions of the spectral resources become trapped between other connections with rigid wavelength assignments. Due to asymmetry in topology and traffic, as network utilization increases, certain links and nodes (shared by the most demands) may become bottlenecks. WC can be beneficial as a way of reducing fragmentation of the network by filling in gaps in the spectrum. Defragmentation can lead to a significant reduction in total spectrum required by the system. In [15], defragmentation is shown to reduce the blocking probability for flexible-bandwidth networks by up to 71% under high traffic loads. A few defragmentation approaches are compared in [16] to show their blocking reduction performance and number of operations. WC can help reduce fragmentation without introducing the latency and transient effects that real-time defragmentation produces.

C. Flexible Modulation

Some networks may opt to invest in bandwidth-variable transponders that, in addition to being able to use a flexible number of OFDM subcarriers, can also modulate each subcarrier using a variety of different modulation schemes, such as binary phase-shift keying (BPSK), quadrature phase-shift keying (QPSK), and high-order quadrature-amplitude modulation (QAM). As the order of the modulation increases, so does the spectral efficiency of the transmission, requiring less bandwidth to transmit the same data-rate.

The length of a transparent lightpath being used limits the spectral efficiency on that lightpath because signals with higher spectral efficiency are more susceptible to physical impairments. By choosing an underlying modulation for a lightpath, one trades off data-rate and/or spectral efficiency for transmission distance. In this paper we approximate the TR for a particular bit rate b and spectral efficiency η by using published experimental data from [17]. A linear regression model for the TR, denoted as R, results in the approximation shown in Fig. 3, corresponding to

$$R(b,\eta) = \alpha b^{-1} + \beta\eta^{-1} + \gamma, \tag{1}$$

where α, β, and γ are coefficients optimized to fit the results in [17]. $R(b,\eta)$ is in units of kilometers, b is in units of Gbps, and η is in units of bits/symbol. η represents the efficiency of standard modulation schemes, e.g., $\eta = 2$ bits/symbol for QPSK and $\eta = 4$ bits/symbol for dual-polarization QPSK. The regression yielded values of $\alpha = 18,600$, $\beta = 8360$, and $\gamma = -250$.

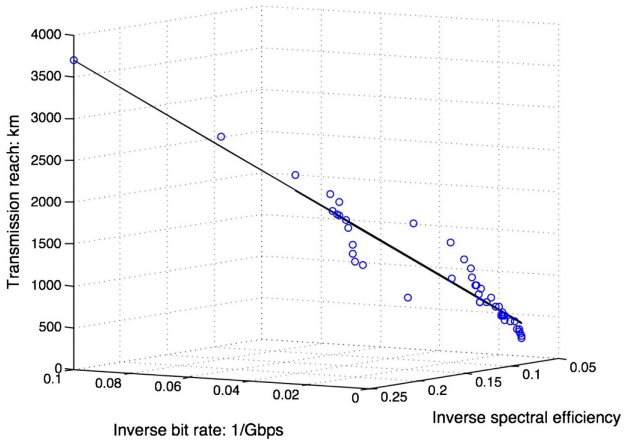

Fig. 3. Transmission reach based on bit rate and spectral efficiency using polynomial fitting over experimental result data from [17].

D. Modulation Conversion

In EONs that use bandwidth-variable transponders that transmit demands at different data-rates and with multiple modulation schemes, the regenerators can also be allowed to modify the underlying modulation for each transparent segment. We refer to this functionality as modulation conversion (MC), which can occur only at a regeneration node. If this functionality is not used, then only one modulation scheme is allowed for any single demand, and the spectral efficiency is then limited by the longest transparent segment. Since some transparent segments may be considerably shorter than others, those could have supported higher efficiencies. By allowing MC at regeneration nodes, the overall spectral efficiency can be improved and the required spectrum can be reduced.

III. MILP

In this section, we first introduce a basic link-based MILP formulation that solves a simple routing and spectrum assignment (RSA) problem with the same modulation scheme for all traffic and without signal regeneration, wavelength conversion, or modulation conversion functionalities, and then extend it to implement signal regeneration and multiple modulation schemes. Last we implement wavelength and modulation scheme conversion. Our general objective is to minimize the spectrum required by the system, as measured by the maximum frequency allocated over all links. We also examine the impact of simultaneously optimizing the spectral use and regeneration resources using a multi-objective function.

The network is modeled as a graph $G(\mathcal{N},\mathcal{L})$ with N nodes and L unidirectional links. We summarize the set notation used by our model in Table I. The model also depends on parameters specific to the network configuration and the traffic demands. We optimize resource assignment to traffic requests for node pairs in the demand set assuming the same flow is needed in each direction. The notation for the independent parameters needed is given in Table II.

TABLE I
SETS USED BY BASIC ILP

$\mathcal{N}$	Set of nodes in the network.
$\mathcal{L}$	Set of unidirectional links in the network. Each link L_{ij} is represented by its source and destination node, $L_{ij} \in \mathcal{L}$.
$\mathcal{D}$	Set of unidirectional traffic demands. Each demand D_{sd} is represented by its source node s and destination node d, $D_{sd} \in \mathcal{D}$.

The objective function of the MILP is to minimize the highest frequency required to support the network traffic:

$$\min_{F_{sd},V_{ij,sd},\delta_{sd,s'd'}} c, \tag{2}$$

where the optimization variables are defined in Table III. The optimization requires several constraints, listed below:

- Highest required spectrum:

$$c \geq F_{sd} + B_{sd} \quad \forall\, D_{sd} \in \mathcal{D}, \tag{3}$$

  where B_{sd} is the bandwidth required by D_{sd} for a given η_{sd}, $B_{sd} = b_{sd} \times \eta_{sd}^{-1}$.
- Flow conservation constraints:

$$\sum_{L_{ij}\in\mathcal{L},j=n} V_{ij,sd} - \sum_{L_{ij}\in\mathcal{L},i=n} V_{ij,sd} = S_{n,sd} \quad \forall\, n \in \mathcal{N}, D_{sd} \in \mathcal{D}. \tag{4}$$

TABLE II
PARAMETERS USED BY BASIC ILP

b_{sd}	Bit rate requested by demand D_{sd}.
η_{sd}	Spectrum efficiency according to particular modulation scheme (e.g., 2 for QPSK).
$S_{n,sd}$	Relationship between nodes and demands: $S_{n,sd} = -1$ if node n is the source node of demand D_{sd} (i.e., $n = s$); $S_{n,sd} = 1$ if node n is the destination node of demand D_{sd} (i.e., $n = d$); $S_{n,sd} = 0$ otherwise (i.e., $n \neq s$, $n \neq d$).
G	Guard band in gigahertz.

TABLE III
VARIABLES USED BY BASIC ILP

F_{sd}	Starting frequency index of demand D_{sd}.
$V_{ij,sd}$	Link assignment: $V_{ij,sd} = 1$ if link L_{ij} is assigned to demand D_{sd}; $V_{ij,sd} = 0$ otherwise.
$\delta_{sd,s'd'}$	Order of the starting frequency index of demand D_{sd} and $D_{s'd'}$.[a] $\delta_{sd,s'd'} = 1$ if $F_{sd} \leq F_{s'd'}$; $\delta_{sd,s'd'} = 0$ if $F_{sd} > F_{s'd'}$.
c	Highest frequency index required by the network traffic.

[a]This relationship between two demands is only of interest if they share a link. We use this relationship in the following constraints to guarantee no overlapping between spectra assigned to multiple demands.

- No spectrum overlap constraints, $\forall D_{sd}, D_{s'd'} \in \mathcal{D}$:

$$\delta_{sd,s'd'} + \delta_{s'd',sd} = 1, \tag{5}$$

$$F_{sd} - F_{s'd'} \le T(1 - \delta_{sd,s'd'} + 2 - V_{ij,sd} - V_{ij,s'd'}), \tag{6}$$

$$F_{sd} - F_{s'd'} + B_{sd} + G \le (T + G) \times (1 - \delta_{sd,s'd'} + 2 - V_{ij,sd} - V_{ij,s'd'}), \tag{7}$$

where T is the total spectrum required by the network traffic, $T = \sum_{D_{sd} \in \mathcal{D}} b_{sd} \times \eta_{sd}^{-1}$.

Equations (2)–(7) define a general link-based RSA formulation for EON. Together Eqs. (5)–(7) enforce a contiguous spectrum assignment to each demand. Equation (5) says that for any two demands sd and $s'd'$ that share a link, one demand has to have a starting frequency lower than the other, and therefore one of the ordering variables is 0 and the other is 1. Equation (6) enforces the necessary relationship between starting frequencies of the two demands based on the variable $\delta_{sd,s'd'}$. Then Eq. (7) forces the starting frequency of the demand with the higher starting frequency to be far enough away from the starting frequency of the lower adjacent channel, i.e., provides room for the signal bandwidth and guard band. These expressions can be modified to implement the more sophisticated signal processing we consider in this paper. Each functionality is discussed below, together with the additional variables and constraints needed.

A. Multiple Modulation Schemes

When each demand has different spectral efficiency, their transmission reach also varies. This is implemented by making η_{sd}^{-1}, the inverse spectral efficiency of demand D_{sd}, a variable instead of a constant parameter. In our model we relax this value from its normal discrete nature to be a real number bounded by the largest and smallest inverse spectral efficiencies allowed: $\eta_{sd,\mathrm{MIN}}^{-1} \le \eta_{sd}^{-1} \le \eta_{sd,\mathrm{MAX}}^{-1}$. For actual implementation, the noninteger modulation efficiency could be rounded to an integer value, or an OFDM modulation with nonuniform constellation per subcarrier could be used. When η_{sd}^{-1} is optimized as part of the solution, the trade-off becomes a higher spectral efficiency for higher η_{sd} in exchange for a higher susceptibility to signal degradation.

B. Signal Regeneration

Signal regeneration can be used to increase the length of transmission beyond the transmission reach. The following constraints, using additional parameters and variables defined in Tables IV and V, respectively, must be satisfied so that the QoT requirements are fulfilled for all demands.

We consider two cases. In the first case, the nodes in the network that are equipped with regeneration devices have been preselected. There has been quite a bit of research recently on how to select regeneration nodes, including [18]. The constraints that the MILP must satisfy for all $D_{sd} \in \mathcal{D}$ and $L_{ij} \in \mathcal{L}$ are as follows:

TABLE IV
PARAMETERS USED BY TRANSMISSION REACH CONSTRAINT

ℓ_{ij}	Length of link L_{ij} in kilometers.
R_{sd}	Transmission reach of demand D_{sd} according to particular spectral efficiency, e.g., in Eq. (1)
$\mathcal{N}^r$	Set of regeneration nodes.

TABLE V
VARIABLES USED BY TRANSMISSION REACH CONSTRAINT

$Y_{n,sd}$	$Y_{n,sd} = 0$ if node n is not on the lightpath assigned to demand D_{sd}. Otherwise, $Y_{n,sd}$ is the physical distance from node n on the lightpath to the beginning of that transparent segment for demand D_{sd}.
$U_{ij,sd}$	$U_{ij,sd} = 0$ if the entire link L_{ij} is not assigned to demand D_{sd} ($V_{ij,sd} = 0$). Otherwise, $U_{ij,sd}$ is the physical distance from node i to the beginning of the transparent segment for demand D_{sd}. Equivalently, if we were not restricted to linear functions, we could have defined $U_{ij,sd} = V_{ij,sd} Y_{i,sd}$.

$$U_{ij,sd} \le V_{ij,sd} R, \tag{8}$$

$$U_{ij,sd} \le Y_{i,sd}, \tag{9}$$

$$Y_{i,sd} - U_{ij,sd} \le R(1 - V_{ij,ld}), \tag{10}$$

$$Y_{n,sd} = \begin{cases} \sum_{L_{ij} \in \mathcal{L}; j=n} \ell_{ij} V_{ij,sd} & \text{if } \exists L_{ij}, i \in \mathcal{N}^r \\ & \text{and } V_{ij,sd} = 1, \\ \sum_{L_{ij} \in \mathcal{L}; j=n} U_{ij,sd} + \ell_{ij} V_{ij,sd} & \text{otherwise,} \end{cases} \tag{11}$$

where $R = R(b_{sd}, \eta_{sd})$ from Eq. (1). For cases in which node i is an intermediate node but link L_{ij} is not an intermediate link for demand D_{sd} (i.e., $V_{ij,sd} = 0$), the distance $U_{ij,sd} = 0$ but $Y_{i,sd}$ is not necessarily zero but has to be less than the transmission reach. In this case Eq. (10) reduces to $Y_{i,sd} \le R$.

The second case we consider is one in which the regeneration nodes are not preselected. We then use the MILP to optimize the placement of regeneration equipment on the network. We treat the regeneration node assignments as binary variables I_n, and, with the help of additional variables defined in Table VI, we use Eqs. (8)–(10) from above and replace Eq. (11) with

$$Y_{n,sd} = \sum_{L_{ij} \in \mathcal{L}; j=n} X_{ij,sd} + \ell_{ij} V_{ij,sd}. \tag{12}$$

Consider Fig. 4 to understand the TR constraints. There is a demand D_{sd} and a path from node s to node d, while link L_{ia} does not belong to the path of demand D_{sd}.

TABLE VI
VARIABLES USED FOR REGENERATOR CIRCUIT ASSIGNMENT

I_n	Regeneration nodes: $I_n = 1$ if node n is used as a regeneration node; $I_n = 0$ otherwise.
$N_{n,c}$	Number of regeneration devices used on node n.
$I_{n,sd}$	Regeneration at node n: $I_{n,sd} = 1$ if demand D_{sd} is regenerated at node n, $I_{n,sd} = 0$ otherwise. For $I_{n,sd} = 1$, node n has to be a regeneration node, i.e., $I_n = 1$, but also its regeneration circuit has to be used by demand D_{sd}.
$X_{ij,sd}$	Distance used to calculate $Y_{n,sd}$ based on whether regeneration occurs at node i: $X_{ij,sd} = U_{ij,sd}$ if $I_{i,sd} = 0$, $X_{ij,sd} = 0$ otherwise.

Assume that node i is the only regeneration node on the path. Inequality (8) says that for any link on the path (e.g., L_{ij}), $U_{ij,sd} \le R$. Since node i is the only regeneration node, U_{ij} for links L_{ij} and $Y_{i,sd}$ both represent the distance from node s to node i. Since link L_{ia} is not on the path for demand D_{sd} (i.e., $V_{ia,sd} = 0$), $U_{ia,sd} = 0$. Equation (11) says that, for node i, since no link that leads to it starts with a regeneration node, $Y_{i,sd}$ is the sum of U's for all links that lead to node i plus the link length of the link that is on the path. Since all links except for link L_{hi} have U equal to zero, $Y_{i,sd}$ is the sum of $U_{hi,sd}$ and the link length of link L_{hi}. But, for node j, since node i is a regeneration node, $Y_{j,sd}$ is just the link length of link L_{ij}. Inequalities (9) and (10) say that $Y_{i,sd} < R$, $U_{ia,sd} < Y_{i,sd}$, and $U_{ij,sd} = Y_{i,sd}$. When the allocation of regeneration nodes is unknown, we use the variable $X_{ij,sd}$ to differentiate the cases in which node i is a regeneration node or not, as in Eq. (12).

Constraints that limit the number of OEO devices per regeneration node can also be included using

$$N_{n,c} = \sum_{D_{sd} \in \mathcal{D}} I_{n,sd}, \tag{13}$$

$$I_n N_{n,c\,\mathrm{MAX}} \ge N_{n,c}, \tag{14}$$

where $N_{n,c\,\mathrm{MAX}}$ is the largest number of regeneration devices that can be equipped on a regeneration node.

When the cost of regeneration resources is a concern, we can build a multi-objective function to balance the cost of regeneration and spectrum resources:

$$\min \left\{ ac + (1-a) \sum_{n \in \mathcal{N}} I_n \right\}, \tag{15}$$

where the coefficient $a \in [0, 1]$ represents the cost relationship between using the two resources. This objective function minimizes the total cost of all resources together, according to their relative costs. While we do not presume to know the exact cost relationship among the two, a network designer can base their objective function on a realistic requirement, and use Eq. (15) to determine what resources are needed and where.

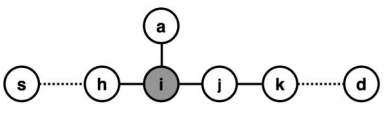

Fig. 4. Illustration used to explain the variables defined for constraining the transmission reach.

C. Wavelength and Modulation Conversion

When WC is available, the frequencies used for a demand can be different on the links entering a regeneration node and exiting it. In order to represent this flexibility we define the starting frequency on a link-by-link basis, as shown in Table VII, and rewrite Eqs. (6) and (7), the constraint that guarantees no spectrum overlap, for all $n \in \mathcal{N}$ as

$$\sum_{L_{ij} \in \mathcal{L}: j=n} F_{ij,sd} - \sum_{L_{ij} \in \mathcal{L}: i=n} F_{ij,sd} \ge -T \times (I_{n,sd} + |S_{n,sd}|),$$
$$\sum_{L_{ij} \in \mathcal{L}: j=n} F_{ij,sd} - \sum_{L_{ij} \in \mathcal{L}: i=n} F_{ij,sd} \le T \times (I_{n,sd} + |S_{n,sd}|). \tag{16}$$

This constraint requires that if node n is an intermediate node for demand D_{sd}, i.e., $n \ne s, n \ne d$, and n is not used as a regeneration node for demand D_{sd}, then the starting frequency assignments entering node n equal the starting frequency assignments exiting node n. For other cases, this constraint does not apply.

Similar to WC, when MC is available the spectral efficiency of each demand on each link can be different than its immediate uplink or downlink if the joining node is used as a regeneration node. We must define the spectral efficiency on a link-by-link basis, as listed in Table VII. The MC constraint can be written as

$$\sum_{L_{ij} \in \mathcal{L}: j=n} \eta^{-1}_{sd,ij} - \sum_{L_{ij} \in \mathcal{L}: i=n} \eta^{-1}_{sd,ij} \ge -\eta^{-1}_{sd,\mathrm{MAX}} \times (I_{n,sd} + |S_{n,sd}|)$$
$$\sum_{L_{ij} \in \mathcal{L}: j=n} \eta^{-1}_{sd,ij} - \sum_{L_{ij} \in \mathcal{L}: i=n} \eta^{-1}_{sd,ij} \le \eta^{-1}_{sd,\mathrm{MAX}} \times (I_{n,sd} + |S_{n,sd}|). \tag{17}$$

Similar to the WC constraint, this constraint requires that modulation (equivalently, spectral efficiency) only be converted at nodes n where the demand is regenerated, i.e., where $I_{n,sd} = 1$.

IV. RECURSIVE MILP

The computational complexity of the MILP formulation for EONs restricts its implementation to offline calculation only. And even then, it does not scale well as the size of the network or the number of traffic demands increase. An effective alternative would be to reduce the problem size to such an extent that the results can be found in acceptable

TABLE VII
VARIABLES USED BY WAVELENGTH AND MODULATION CONVERSION

$F_{ij,sd}$	Starting frequency index of demand D_{sd} on link L_{ij}.
$\eta^{-1}_{sd,ij}$	Inverse spectral efficiency of demand D_{sd} on link L_{ij}.

time and with reasonable computational resources. Many works have attempted to break the whole RSA problem into subproblems with or without losing some degree of optimality [19,20]. In this section, we propose a different way to reduce the problem size by splitting the traffic matrix into submatrices and solving them sequentially, a technique we call *recursive MILP*.

The recursive MILP approach is motivated by the understanding that the complexity of the problem is greatly affected by the number of traffic demands that need to be accommodated at once. By separating them into subsets and allocating those in sequential iterations, the overall runtime as well as other computational resources, such as memory, can be reduced. The solution from the previous iteration forms new MILP constraints for the new iteration. In particular, the first iteration can be viewed as a subproblem with the same constraints but with fewer demands. In the new iterations, the constraints (e.g., nonoverlapped spectrum assignment) apply to both the assigned and unassigned resources. In this manner, the original problem can be solved, albeit not optimally, after all iterations are done.

Another advantage of using recursive MILP is that the complexity is easy to estimate. For example, in the aforementioned MILP formulation, the complexity-dominating variable is $\delta_{sd,s'd'}$, which grows with the number of demands $|\mathcal{D}|$ squared, i.e., $O(|\mathcal{D}|^2)$. By solving the same problem recursively, the number of variables of each calculation is reduced. If the number of subsets is S, the last subproblem (which has the highest number of variables) will have about $\frac{|\mathcal{D}|^2}{S^2}$ many $\delta_{sd,s'd'}$. Running the MILP in recursive mode does not require reformulating the problem: the constraints remain the same, but the variables that represent demands from previous iterations become constants.

It can be expected that the recursive MILP suffers loss of optimality compared to the nonrecursive counterpart. The gap between the suboptimal solution from the recursive MILP and the optimal solution depends on the size of the subset and the grouping and ordering of traffic demands. Since the complexity is easy to estimate, network designers can base the implementation of the formulation on the complexity they can accept. In Table VIII we show a comparison between the order of computational complexity for nonrecursive (single-run) and recursive MILP solutions.

The complexity and optimality are affected not just by the size of the demand subsets, but also by the selection of demands in each iteration. The grouping of demands can be done in many ways such as sorting them randomly or based on their characteristics such as volume and locality. In Section V we show a comparison of the required spectrum using different ordering schemes.

Another use for the recursive MILP implementation is to help accommodate network expansion with existing infrastructure. The existing assignments of physical resources (links, regeneration devices, spectrum) can be input into the MILP as constraints, in the same way as was done for results from earlier iterations of the recursive MILP. In our simulation results below we show a progressively increasing number of traffic demands, each demand set including the same demands from previous sets. The spectrum predicted shows the resources required for network expansion. This problem is similar to the so-called *dynamic* resource allocation problem [21,22], the difference being whether connections are also torn down. In that scenario, the dynamic resource assignment problem can also be solved using the above MILP in a similar recursive way as was shown for fixed-grid WDM systems in [11].

V. Numerical Results

We test our MILP formulations on the NSF-24 mesh network, shown in Fig. 1, which is often used as a benchmarking topology in the literature. In addition, to gauge the sensitivity of our results on network topology we simulate a symmetric network illustrated in Fig. 2, with 24 edge nodes that also serve as intermediate nodes for routing traffic. The NSF-24 network has 43 bidirectional links (or unidirectional link pairs of 86 links), while the symmetric-24 network has 55 bidirectional links (or 110 unidirectional links). In our simulation, we assume that traffic demands are generated between randomly selected node pairs so that the results are not biased by the traffic distribution; we expect that this will lead to a worst case compared with realistic traffic. The demands have random bit-rate requests ranging uniformly from 1 to 100 Gbps in order to represent the highest expected heterogeneity of Internet traffic. This assumption is to show the full flexibility that EON introduces in spectrum assignment versus fixed-grid WDM (or even coarse-grid EON) where the network operator is forced to overprovision spectrum for odd-size traffic to fit standard rates (e.g., 10, 40, or 100 Gbps). The demands are assumed to be unidirectional and are between different node pairs. We assume that the traffic in the opposite direction is assigned over the same route and with the same physical resources in the opposite links. We thus reduce the complexity of our problem by halving the number traffic demands in the search for a solution. The traffic is assumed static, and no traffic grooming or reverse grooming (i.e., traffic splitting) is considered. In all cases we collect simulation results over 20 independent random demand sets and report average results. To the best of our knowledge, there is no other published work on the effects of wavelength and modulation conversion

TABLE VIII
Complexity of One Iteration of Recursive and Nonrecursive Basic MILP

Number of Variables		Number of Constraints	
Nonrecursive	Recursive	Nonrecursive	Recursive
$(1+L+\lvert\mathcal{D}\rvert)\times\lvert\mathcal{D}\rvert+1$	$(1+L+\lvert\mathcal{D}\rvert/S)\times\lvert\mathcal{D}\rvert/S$	$\lvert\mathcal{D}\rvert+N\lvert\mathcal{D}\rvert+(1+2L)\times\lvert\mathcal{D}\rvert^2$	$\lvert\mathcal{D}\rvert/S+N\lvert\mathcal{D}\rvert/S+(1+2L)\times(\lvert\mathcal{D}\rvert/S)^2$

on EONs, so we can only compare our algorithms with each other.

In the numerical results, we investigate one or two features at a time in each section. Unless otherwise stated, the algorithm tested is an optimum MILP (nonrecursive) assuming multiple modulations are available with $\eta_{sd} \in [1, 10]$ and signal regeneration capability at all nodes, but no modulation or wavelength conversion is used so that each demand uses the same modulation and spectrum from source to destination.

A. Recursive and Nonrecursive MILP

We first verify the applicability of the MILP solution to the network sizes we have chosen, and compare the optimality and computational complexity between the recursive and nonrecursive approaches. In Fig. 5 we show the required spectrum for both single and multiple modulation schemes. We also plot the standard deviation for our results. As the number of demands allocated increases, the spectral usage increases approximately linearly. For the single modulation case (marked with Δ), we choose (arbitrarily) $\eta_{sd} = 2\ \forall D_{sd}$ (QPSK), while for the multiple modulation case (marked with ∇) η_{sd} is optimized. When the MILP assigns resources to all demands together [the "single solve" approach (dashed lines)], the performance is notably better than the recursive approach [assuming a random partition of the demands into sets of five (solid lines)], but too computationally burdensome for more than 30–40 simultaneous traffic demands. We also conclude that the added flexibility of optimizing the spectral efficiency for each demand more than halves the required spectrum. Both networks show similar results.

In Fig. 6 we compare the complexity of the four approaches from Fig. 5. The running time for each computation of 25 demands is shown in a histogram for 10 trials using different demand sets. For example, for the recursive solution using a single modulation, the solution for eight of the 10 trials converged in less than 1 s, while the other two finished between 1 and 10 s. The running times for the optimal MILP vary considerably between trials (we set a time limit of 3000 s), while the running times for the recursive approach are uniformly short.

In Fig. 7 we show that the partitioning and ordering of traffic demands in the recursive solution has a small but nonnegligible impact on the required spectrum. When only a few demands have been assigned resources, accommodating high data-rate (b_{sd}) demands first (the "Descending b_{sd}" case) leads to a lower required spectrum. Also, assigning demands that have the shortest paths (labeled "SP") first typically requires less spectrum than assigning the longer-distance connections first. The differences are slightly more pronounced on the symmetric network. Note that this relationship may also depend on the traffic distribution and the node degrees of the network. Nevertheless, we do expect that the results presented are typical of real networks.

We then investigate the effect of the demand subset size $|\mathcal{D}|/S$ on the required spectrum obtained by the recursive MILP. The single solve approach, which finds the globally optimal result for all the traffic demands together, must always yield the smallest required spectrum. In Fig. 8, we see that as the subset size decreases, the required spectrum increases. When few demands are assigned per iteration, the required spectrum appears to be stair-stepped, since new connections can often use gaps in the allocated spectrum left by fragmentation induced by the suboptimal resource allocation. In Fig. 9, we show the cumulative running time that each case needs to assign a certain

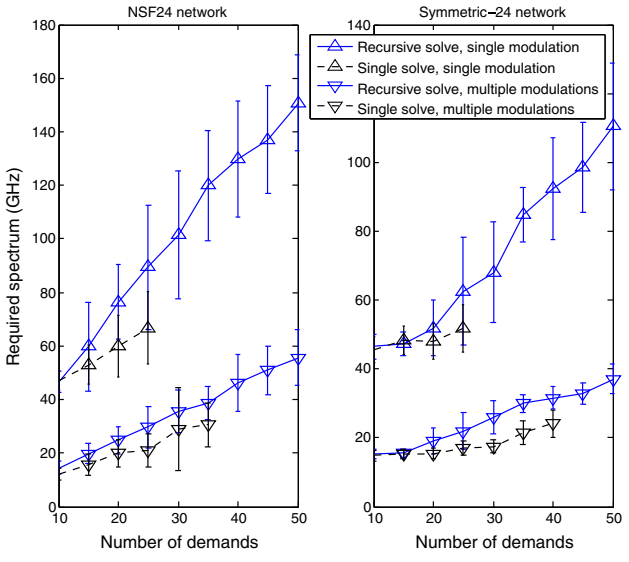

Fig. 5. Required spectrum using the recursive MILP and single solve MILP for a single modulation scheme ($\eta = 2$) and multiple modulation schemes ($1 \le \eta \le 10$).

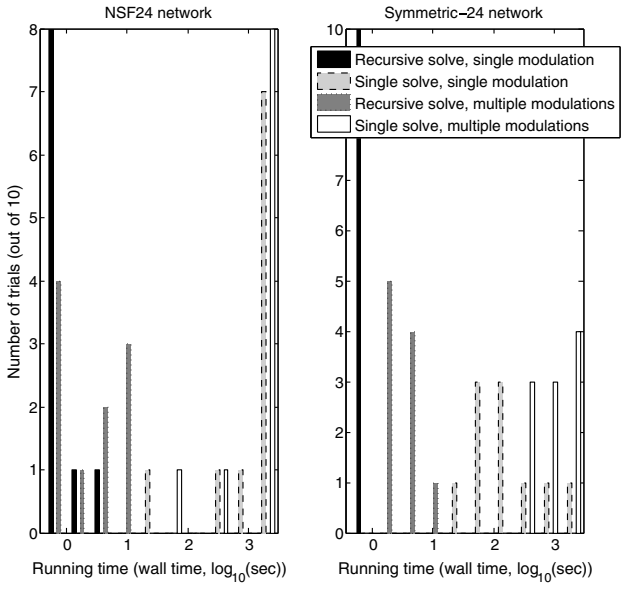

Fig. 6. Histogram of running times for the recursive MILP and single solve MILP for a single modulation scheme ($\eta = 2$) and multiple modulation schemes ($1 \le \eta \le 10$), for 10 different trials of 25 demands.

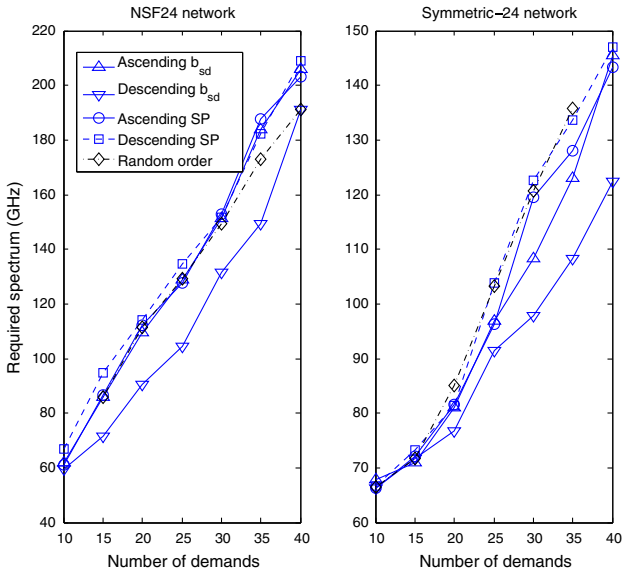

Fig. 7. Required spectrum for the recursive MILP with different ordering schemes for the same 40 demands, as the resource assignment for the demand subsets progresses (single modulation scheme with $\eta = 2$).

number of demands. Having a smaller subset size requires more time to accommodate traffic. As more demands are assigned, the new demands also require more time to be assigned because the solutions from previous subsets have to be used as constraints on the existing network state.

B. Multi-Objective Formulation

We investigate the ability of the MILP optimization to effectively trade off spectral usage with regenerator usage by including the cost of regeneration resources (namely the number of regeneration nodes) in our objective function according to Eq. (15). The resulting spectrum and regeneration node requirements based on different values of the cost coefficient a are shown in Figs. 10 and 11, respectively. We show the effect of the number of demands $|\mathcal{D}|$. Assuming a *priori* knowledge of the cost relationship between spectrum and regeneration resources, network designers can choose the cost coefficients accordingly. The results show that when the spectrum cost is not considered (i.e., $a = 0$), the number of regeneration nodes is minimized and the required spectrum is large. Since we have configured the demand volume and spectral efficiency so that no demand absolutely requires regeneration, in the case of $a = 0$ the number of regeneration nodes is always minimized to zero. The required spectrum in this case is also highly irregular, as it is entirely unconstrained. On the contrary, when the regeneration cost is not considered (i.e., $a = 1$) the required spectrum is minimized but the number of regeneration nodes is high. It is interesting to note that by assigning even a relatively small coefficient to the regeneration cost (i.e., $a = 0.99$), we are able to maintain a similar required spectrum but greatly reduce the number of regeneration nodes needed. Results also show that the coefficient

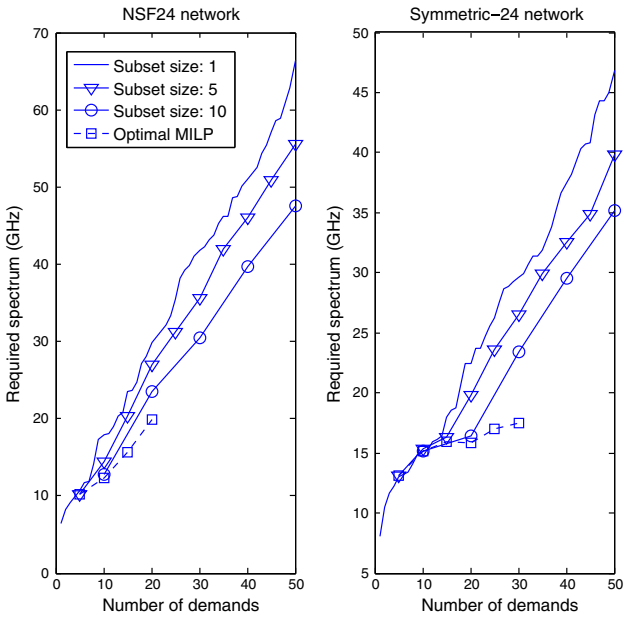

Fig. 8. Required spectrum by solving the same recursive formulation with different demand subset sizes (multiple modulation schemes with $1 \le \eta \le 10$). The "optimal MILP" results correspond to the "single solve" in Fig. 5.

$a = 0.5$ introduces a good trade-off between the required spectrum and the required number of regeneration nodes.

C. Wavelength and Modulation Scheme Conversion

Since regeneration involves OEO conversion, it can improve the signal quality, so as to extend the TR, and can

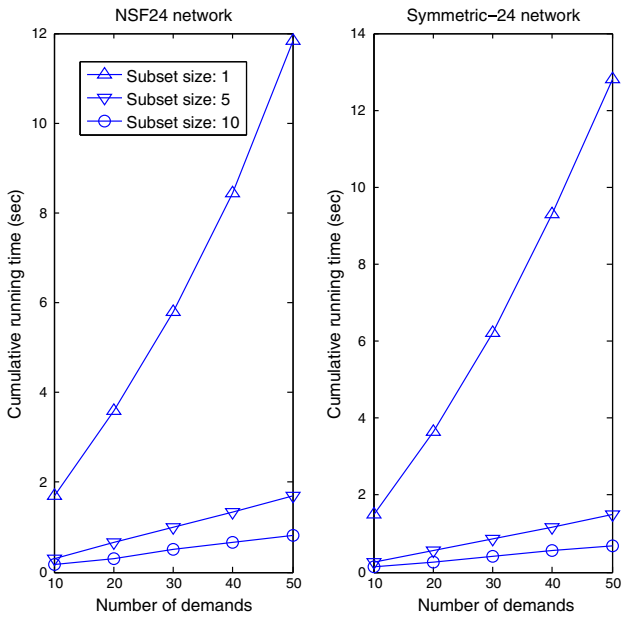

Fig. 9. Cumulative running time by solving the same recursive formulation with different demand subset sizes.

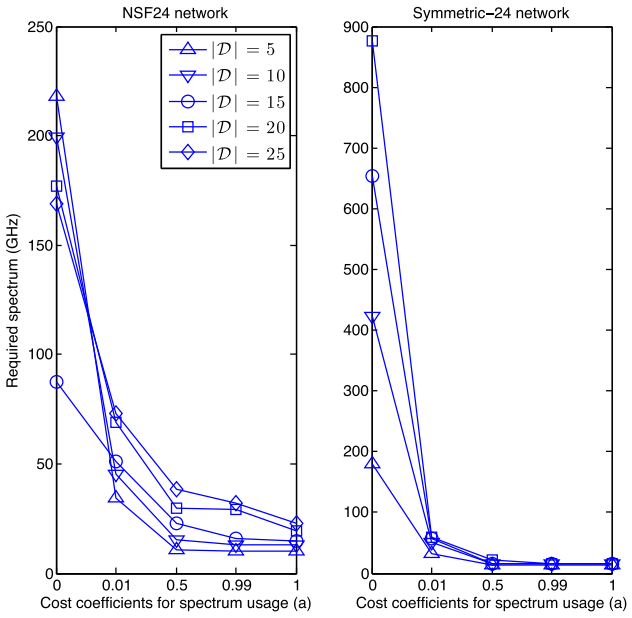

Fig. 10. Required spectrum as the coefficient a in the objective function (15) varies. Note that the horizontal axis is not drawn to scale.

also provide an opportunity to change the spectrum and modulation assigned starting from that node. We run simulations to show the impact on the spectrum requirements of using the capability to convert the wavelength and/or modulation at the regeneration nodes. In Fig. 12, we solve the resource allocation problem with our recursive MILP formulation, since the added flexibility of WC and MC increases the complexity of the problem considerably. Both WC and MC reduce the amount of spectrum required to support the traffic demand. WC allows signals to be reallocated on some links so as to fill gaps left by other traffic demands. The improvement made by wavelength conversion depends on the fragmentation condition of the network. When one link is heavily congested, it then becomes a bottleneck of the network, creating unnecessary fragmentation on other links. MC takes advantage of the difference in the length of transparent segments for each lightpath. If they are significantly different, the spectrum saved by optimizing the spectral usage based on each segment becomes significant. The NSF-24 network has well-documented bottleneck paths, and thus benefits more from WC than a more symmetric topology. The symmetric network also has equal link lengths, and can therefore not exploit MC as much as the more heterogeneous NSF-24.

In Fig. 12, when we compare the WC case with the case in which both wavelength and modulation scheme conversion are available, we expect the latter to always outperform the former, since it has more flexibility. However, the results (which are averaged over 20 trials) for the symmetric network show that this relationship is not guaranteed. Solving the MILP recursively optimizes the solution in each iteration according to its constraints, yet does not necessarily lead to the optimal solution for the whole traffic matrix when looking at multiple iterations together. Each iteration can only find the local optimum for its subproblem, and the local optimal solution of a subproblem may not be one part of the global optimal for the whole problem. To put it in terms of resource assignment, in the previous iterations resources may be assigned to demands according to the subproblem, but such assignment may be not optimal when one considers the whole traffic matrix, which leads to an increased spectrum requirement. The same limitation exists in dynamic real-time RSA, where no future traffic information is available. The RSA algorithm

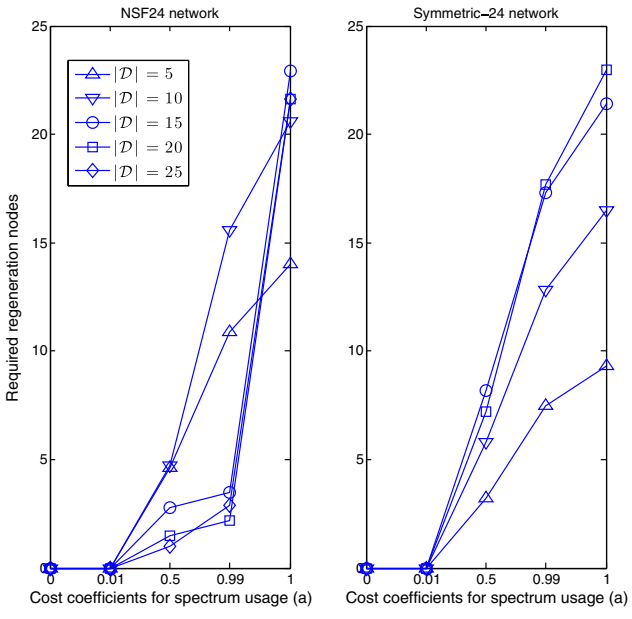

Fig. 11. Required number of regeneration nodes as the coefficient a in objective function (15) varies. Note that the horizontal axis is not drawn to scale.

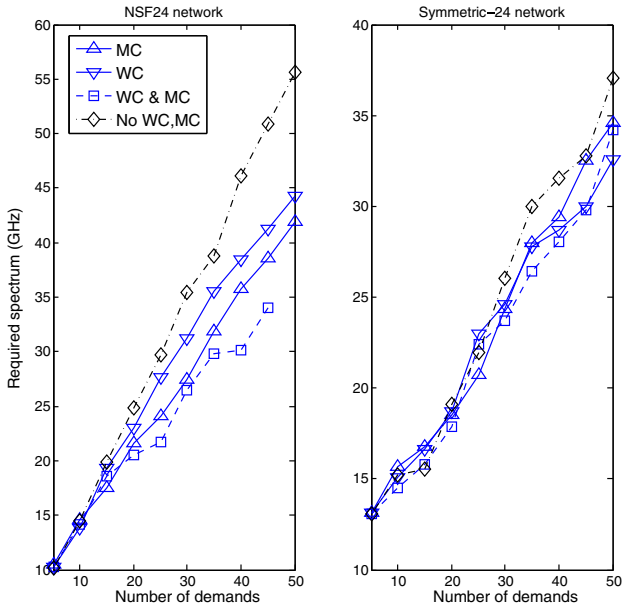

Fig. 12. Spectrum usage comparison using the recursive MILP with and without wavelength and/or modulation conversion.

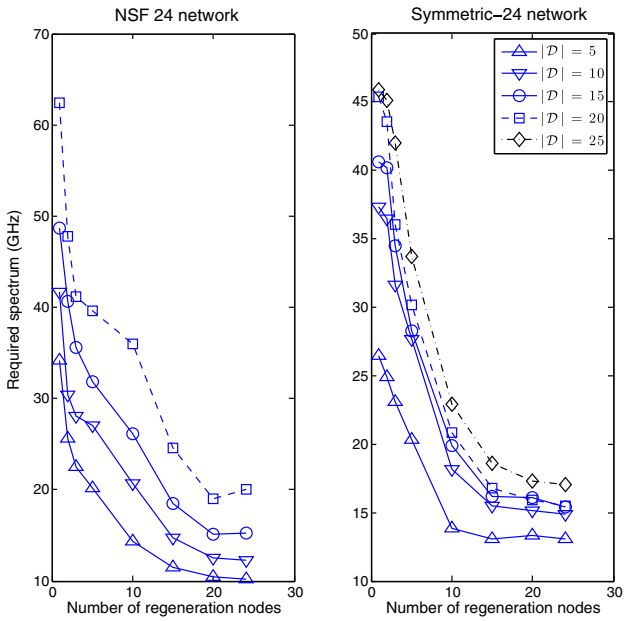

Fig. 13. Spectrum usage with a limited number of regeneration nodes.

cannot assign physical resources to current traffic demands taking unknown future demands into consideration. One could introduce a conservative rule based on long-term statistics that prevents overassigning physical resources to current demands.

D. Regeneration Node Placement

When regeneration resources are scarce, careful network planning is important to minimize capital expenditure. In order to show the trade-off between the number of regeneration nodes and the required spectrum, we simulate a case in which the network has a limited number of regeneration nodes. In order to show the change of performance by adding additional regeneration nodes we keep the existing regeneration nodes unchanged once assigned. We use our results in Fig. 10 for cost coefficient $a = 0.5$ to find and rank the most often used regeneration node locations on average (over 20 trials). After allocating a limited number of nodes as regeneration nodes according to this ranking, we minimize the required spectrum. Our results in Fig. 13 show that for the NSF-24 network, the required spectrum reaches its minimum when 20 nodes have been allocated as regeneration nodes; for the symmetric 24-node network, the required spectrum is close to minimum already when the number of regeneration nodes reaches 15. We believe this is due to the symmetric structure of the symmetric 24-node network, where most node pairs share joint intermediate nodes.

VI. Conclusion

In this paper we propose an MILP formulation to investigate the impact of technologies such as allowing multiple modulation schemes, signal regeneration, wavelength conversion, and modulation conversion on the required spectrum of the EON. In order to balance the optimality and complexity of the algorithm, we propose a recursive MILP formulation that yields a suboptimal yet comparable solution to the MILP by using a lower and more consistent running time. The performance of the recursive model depends on factors such as heuristic demand ordering and iteration size. We show through simulation that equipping systems with signal regenerating nodes that control physical impairments and allow for modulation and/or wavelength conversion reduces the amount of spectrum required. Such improvements depend on the topology of the network. We also show the impact of having a limited number of regeneration nodes and different topology structures.

Acknowledgments

This work was supported in part by NSF grants CNS-0915795 and CNS-0916890.

References

[1] X. Wang, M. Brandt-Pearce, and S. Subramaniam, “Grooming and RWA in translucent dynamic mixed-line-rate WDM fiber optics networks subject to physical impairments,” in *Proc. OFC*, Mar. 2012.

[2] T. Erlebach and S. Stefanakos, “On shortest-path all-optical networks without wavelength conversion requirements,” in *STACS 2003*, H. Alt and M. Habib, Eds. (Lecture Notes in Computer Science, Vol. 2607). Springer, 2003, pp. 133–144.

[3] N. Garcia, P. Lenkiewicz, M. Freire, and P. Monteiro, “On the performance of shortest path routing algorithms for modeling and simulation of static source routed networks—An extension to the Dijkstra algorithm,” in *2nd Int. Conf. on Systems and Networks Communications (ICSNC)*, Aug. 2007, p. 60.

[4] A. Zyane, Z. Guennoun, and O. Taous, “Performance evaluation of shortest path routing algorithms in wide all-optical WDM networks,” in *Int. Conf. on Multimedia Computing and Systems (ICMCS)*, Apr. 2014, pp. 831–836.

[5] S. Talebi, F. Alam, I. Katib, M. Khamis, R. Salama, and G. N. Rouskas, “Spectrum management techniques for elastic optical networks: A survey,” *Opt. Switching Networking*, vol. 13, pp. 34–48, July 2014.

[6] M. Klinkowski and K. Walkowiak, “Routing and spectrum assignment in spectrum sliced elastic optical path network,” *IEEE Commun. Lett.*, vol. 15, no. 8, pp. 884–886, Aug. 2011.

[7] M. Klinkowski and K. Walkowiak, “Offline RSA algorithms for elastic optical networks with dedicated path protection consideration,” in *4th Int. Congr. on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT)*, Oct. 2012, pp. 670–676.

[8] A. Pages, J. Perello, S. Spadaro, and G. Junyent, “Split spectrum-enabled route and spectrum assignment in elastic optical networks,” *Opt. Switching Networking*, vol. 13, pp. 148–157, 2014.

[9] Y. Wang, X. Cao, and Y. Pan, “A study of the routing and spectrum allocation in spectrum-sliced elastic optical path

networks," in *Proc. IEEE INFOCOM*, Apr. 2011, pp. 1503–1511.

[10] E. Yetginer, Z. Liu, and G. Rouskas, "Fast exact ILP decompositions for ring RWA," *J. Opt. Commun. Netw.*, vol. 3, no. 7, pp. 577–586, July 2011.

[11] X. Wang, M. Brandt-Pearce, and S. Subramaniam, "Dynamic grooming and RWA in translucent optical networks using a time-slotted ILP," in *IEEE Global Communications Conf. (GLOBECOM)*, Dec. 2012, pp. 2996–3001.

[12] A. Ozdaglar and D. Bertsekas, "Optimal solution of integer multicommodity flow problems with application in optical networks," in *Frontiers in Global Optimization*, C. Floudas and P. Pardalos, Eds. (Nonconvex Optimization and Its Applications, Vol. 74). Springer, 2004, pp. 411–435.

[13] R. Casellas, R. Munoz, J. Fabrega, M. Moreolo, R. Martinez, L. Liu, T. Tsuritani, and I. Morita, "Design and experimental validation of a GMPLS/PCE control plane for elastic CO-OFDM optical networks," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 1, pp. 49–61, Jan. 2013.

[14] F. Kuipers, A. Beshir, A. Orda, and P. Van Mieghem, "Impairment-aware path selection and regenerator placement in translucent optical networks," in *18th IEEE Int. Conf. on Network Protocols (ICNP)*, Oct. 2010, pp. 11–20.

[15] D. Geisler, Y. Yin, K. Wen, N. Fontaine, R. Scott, S. Chang, and S. Yoo, "Demonstration of spectral defragmentation in flexible bandwidth optical networking by FWM," *IEEE Photon. Technol. Lett.*, vol. 23, no. 24, pp. 1893–1895, Dec. 2011.

[16] M. Zhang, Y. Yin, R. Proietti, Z. Zhu, and S. Yoo, "Spectrum defragmentation algorithms for elastic optical networks using hitless spectrum retuning techniques," in *Optical Fiber Communication Conf. and Expo. and the Nat. Fiber Optic Engineers Conf. (OFC/NFOEC)*, Mar. 2013, pp. 1–3.

[17] A. Klekamp, R. Dischler, and F. Buchali, "Limits of spectral efficiency and transmission reach of optical-OFDM superchannels for adaptive networks," *IEEE Photon. Technol. Lett.*, vol. 23, no. 20, pp. 1526–1528, Oct. 2011.

[18] Y. Chen, A. Bari, and A. Jaekel, "Optimal regenerator assignment and resource allocation strategies for translucent optical networks," *Photonic Netw. Commun.*, vol. 23, no. 1, pp. 16–24, 2012.

[19] X. Wan, N. Hua, and X. Zheng, "Dynamic routing and spectrum assignment in spectrum-flexible transparent optical networks," *J. Opt. Commun. Netw.*, vol. 4, no. 8, pp. 603–613, Aug. 2012.

[20] X. Liu, L. Gong, and Z. Zhu, "Design integrated RSA for multicast in elastic optical networks with a layered approach," in *IEEE Global Communications Conf. (GLOBECOM)*, Dec. 2013, pp. 2346–2351.

[21] Z. Zhu, W. Lu, L. Zhang, and N. Ansari, "Dynamic service provisioning in elastic optical networks with hybrid single-/multi-path routing," *J. Lightwave Technol.*, vol. 31, no. 1, pp. 15–22, Jan. 2013.

[22] K. Kuang, X. Wang, S. Wang, S. Xu, H. Liu, and G. Liu, "Dynamic routing and spectrum allocation in elastic optical networks with mixed line rates," in *IEEE 15th Int. Conf. on High Performance Switching and Routing (HPSR)*, July 2014, pp. 1–6.

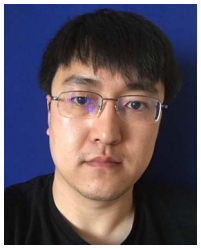

Xu Wang received his B.S. degree in electrical engineering from Beijing Technology and Business University, China, in 2006, his M.S. degree in electrical engineering from Polytechnic University, USA, in 2008, and his Ph.D. degree in electrical engineering from the University of Virginia, USA, in 2014. After graduating, he joined VLNComm Inc. working on free-space visible-light communication research projects. His research interests include fiber-optic and free-space optical communications.

Maïté Brandt-Pearce is a professor in the Charles L. Brown Department of Electrical and Computer Engineering at the University of Virginia. She received her Ph.D. in electrical engineering from Rice University in 1993. Her research interests include nonlinear effects in fiber optics, free-space optical communications, cross-layer design of optical networks subject to physical layer degradations, body area networks, and radar signal processing. Dr. Brandt-Pearce is the recipient of an NSF CAREER Award and an NSF RIA. She is a co-recipient of Best Paper Awards at ICC 2006 and GLOBECOM 2012. She was the General Chair of the Asilomar Conference on Signals, Systems & Computers in 2009 and the Chair of the Optical Networks and Systems Symposium at IEEE GLOBECOM 2015.

Suresh Subramaniam (S'95-M'97-SM'07-F'15) received the Ph.D. degree in electrical engineering from the University of Washington, Seattle, in 1997. He is a Professor in the Department of Electrical and Computer Engineering at the George Washington University, Washington, DC. His research interests are in the architectural, algorithmic, and performance aspects of communication networks, with current emphasis on optical networks, cloud computing, and data center networks. He has published over 150 peer-reviewed papers in these areas. Dr. Subramaniam is a co-editor of three books on optical networking. He is or has been on the editorial boards of 6 journals including *IEEE/ACM Transactions on Networking* and the IEEE/OSA *Journal of Optical Communications and Networking*, and has chaired several conferences including IEEE INFOCOM 2013. He is a co-recipient of Best Paper Awards at ICC 2006 and at the 1997 SPIE Conference on All-Optical Communication Systems. He is a Fellow of the IEEE.