A Fast Path-Based ILP Formulation for Offline RWA in Mesh Optical Networks

Zeyu Liu, George N. Rouskas
Department of Computer Science, North Carolina State University, Raleigh, NC 27695-8206, USA

***Abstract*—RWA is a fundamental problem in the design and control of optical networks. We introduce the concept of symmetric RWA solutions and present a new ILP formulation to construct optimally such solutions. The formulation scales to mesh topologies representative of backbone and regional networks. Numerical results demonstrate that the new formulation achieves a decrease of up to two orders of magnitude in running time compared to existing formulations. In particular, optimal solutions for topologies up to 20 nodes can be obtained within minutes using commodity CPUs, and larger networks can be solved in reasonable time. Our approach significantly lowers the barrier to entry in fully exploring the solution space of optical network design and in investigating the sensitivity of design decisions to forecast demands via extensive "what-if" analysis. Such analysis cannot be carried out currently without large investments in computational resources and time.**

I. Introduction

The global network infrustructure is built on a foundation of optical networking technologies, first deployed in the backbone and regional parts of the network but now also reaching into the access part in the form of PON architectures. Therefore, the planning and design of optical networks is crucial to the operation and economics of the Internet and its ability to support critical and reliable communication services.

In optical networks, traffic is carried over *lightpaths* that are optically switched at intermediate nodes. The routing and wavelength assignment (RWA) problem is one of selecting a path and wavelength for each connection demand, subject to certain constraints. RWA is a fundamental problem in the engineering, control, and design of optical networks, and arises in most design applications, including traffic grooming, survivability design, and traffic scheduling.

Offline RWA is a network design problem in which the input typically consists of a set of traffic demands. This problem is NP-hard, and several integer linear program (ILP) formulations have been proposed to solve it. Recently, we developed an exact decomposition approach for an ILP formulation based on maximal independent sets that makes it possible to obtain optimal solutions to the RWA problem for maximum size (i.e., 16-node) SONET rings in only a few seconds using commodity CPUs [9]. This new, fast technique achieves several orders of magnitude decrease in running time.

As backbone and regional networks evolve from ring to mesh, optimal RWA solutions for general topologies are becoming important to network designers and operators. Unfortunately, current optimization methods cannot be used to solve optimally mesh network instances arising in practice. Consequently, many heuristic solution methods have been developed under various assumptions and network settings (refer to the survey in [4]). Nevertheless, the lack of scalability of optimal methods makes it difficult to characterize the performance of heuristic algorithms, and severely limits the application of "what-if" analysis to explore the sensitivity of network design decisions to forecast traffic demands, capital and operational cost assumptions, and service price structures.

In this paper, we present a new path-based ILP formulation for tackling the RWA efficiently on mesh network topologies encountered in practice. In Section II, we review existing ILP formulations for the RWA problem and discuss in depth a path-based formulation that is the starting point for our work. In Section III we introduce the concept of a symmetric RWA solution and develop a new, fast formulation for obtaining such solutions, while in Section IV we show how to further improve this formulation. In Section V, we present numerical results that demonstrate the effectiveness of our formulation in solving problem instances representative of existing backbone networks, and we conclude the paper in Section VI.

II. Path ILP Formulation of the RWA Problem

Consider a connected graph $G=(V,A)$, where V denotes the set of nodes and A denotes the set of directed links (arcs) in the network. We define $N=|V|$ and $L=|A|$ as the number of nodes and links, respectively. Each directed link l consists of an optical fiber that may support W distinct wavelengths. Let $T=[t_{sd}]$ denote the traffic demand matrix, where t_{sd} is a non-negative integer representing the number of lightpaths to be established from source node s to destination node d. In general, traffic demands may be asymmetric, i.e., $t_{sd} \neq t_{ds}$. We also make the assumption that $t_{ss}=0, \forall s$.

There are three classes of ILP formulations for the RWA problem depending on the types of variables used: (1) link-based, (2) path-based, or (3) maximal independent set (MIS)-based. A comparison of link and path based formulations was carried out in [5], while several RWA algorithms based on LP relaxations were designed and studied in [3].

In this work, we focus on the path ILP formulation, whereby the entities of interest (i.e., decision variables) are path related. Specifically, a set of K paths are generated in advance for each source-destination pair (s,d), and all lightpaths from s to d are constrained to follow one of these K physical paths. We denote the k-th path, $k=1,\ldots,K$, from node s to node d as $p_{sd,k}$, and

This work was supported in part by the National Science Foundation under Grant CNS-1113191.

use binary parameters $X^l_{sd,k}$ to indicate whether path $p_{sd,k}$ uses link l. We also let $P = |\mathscr{P}|$ denote the total number of paths for all source-destination pairs.

In this paper, we study the following minimization problem.

Problem 2.1 (minRWA): Given graph G, number of wavelengths W, traffic demand matrix T, and path set $\mathscr{P} = \{p_{sd,k}\}$, assign each demand (i.e., lightpath) a path and wavelength so as to minimize the number of wavelength used in the network.

Let us define the following set of decision variables:

- $c^w_{sd,k} \in \{0,1\}$: binary variable indicating whether there is a lightpath from s to d assigned wavelength w on $p_{sd,k}$;
- $U^w \in \{0,1\}$: binary variable indicating whether wavelength w is used anywhere in the network; and
- W_{high}: the highest index of used wavelengths.

The path ILP formulation can then be expressed as:

$$\text{minimize:} \quad W_{high} \tag{1}$$

subject to:

- traffic constraints:
$$\sum_{k=1}^{K} \sum_{w=1}^{W} c^w_{sd,k} = t_{sd}, \quad \forall s,d \in V \tag{2}$$
- distinct wavelength constraints:
$$\sum_{s,d \in V} \sum_{k=1}^{K} c^w_{sd,k} X^l_{sd,k} \leq 1, \quad \forall l \in A, \forall w \tag{3}$$
- wavelength usage constraints:
$$\sum_{s,d \in V} \sum_{k=1}^{K} c^w_{sd,k} \leq U^w P, \quad \forall w \tag{4}$$
- highest wavelength index constraints:
$$W_{high} \geq w U^w, \quad \forall w \tag{5}$$
- integrality constraints:
$$U^w = 0,1, \quad \forall w; \qquad c^w_{sd,k} = 0,1, \quad \forall s,d,k,w \tag{6}$$

The traffic constraints (2) ensure that all the traffic demands are satisfied. The distinct wavelength constraints (3) guarantee that no two lightpaths sharing the same link are assigned to the same wavelength. The wavelength usage constraints (4) make sure that the decision variable U^w is set to 1 if wavelength w is used on any of the P paths. Constraints (5) count the number of distinct wavelengths used in the network.

The formulation has two main limitations: (1) its size increases rapidly with the size of the network and the number of wavelengths; and (2) it has a symmetry problem in that multiple solutions with the same objective value can be obtained by simply changing the order of wavelengths. Since the ILP solver has to evaluate all $W_{high}!$ distinct optimal solutions, the running time can be unnecessarily long. Hence, it has a severe scalability issue.

One approach to overcome these limitations has been to consider objectives that are piecewise linear functions of the link loads (i.e., the maximum number of wavelengths on each link), in place of (1), and relax the integrality constraints in the formulation [3]. However, there are no guarantees that the optimal solution returned by the LP solver will be integer. Moreover, this technique cannot be applied to other objectives, e.g., the one in (1) in this work.

Another approach has been to apply column generation. This technique has been applied to RWA [6], and does yield smaller problem sizes for each iteration. However, it may require a large number of iterations and a recent study specific to the RWA problem [6] reports low speed-up factors.

In the next two sections, we present a set of techniques to scale the path ILP formulation to problem instances encountered in practice. Our approach is general in the sense that it can be applied to *any* such formulation of the RWA problem that includes the traffic and distinct wavelength constraints (2) and (3), respectively, regardless of the exact form of the objective function or other constraints.

III. Optimal Symmetric RWA Solutions

In this work, we assume that routing is symmetric, i.e., for all source-destination pairs (s,d), path $p_{sd,k}$ from node s to node d consists of the same physical links as path $p_{ds,k}$ from node d to node s, but in the opposite direction. We now introduce the concept of a symmetric RWA solution.

Definition 3.1 (Symmetric RWA solution): Without loss of generality, assume that for a node pair (s,d) we have that $t_{sd} \leq t_{ds}$. For RWA problem instances with symmetric routing, we define a *symmetric solution* to be such that:

$$c^w_{sd,k} = c^w_{ds,k}, \ \forall \ w,k, \ \ \forall \ s,d : t_{sd} \leq t_{ds}. \tag{7}$$

In other words, for each demand between source s and destination d for which there is a symmetric demand from d to s, the two demands are assigned the *same* path (in opposite directions) **and** the *same* wavelength[1].

It has long been recognized [2] that symmetric routing offers many advantages, including better network capacity planning and utilization, faster problem resolution, and more consistent traffic flow characteristics in terms of delay, cost, and other metrics. Several Internet protocols, including RSVP, NTP, and multicast (e.g., reverse path forwarding) rely upon routing symmetry, despite the fact that symmetric routing of fine granularity flows is not guaranteed in today's multi-provider Internet [7]. On the other hand, optical transport networks are designed, deployed and engineered by a single provider, hence symmetric routing of lightpaths may be easily achieved within such a backbone or regional network. Our definition of a symmetric RWA solution goes one step further, requiring that symmetrically routed lightpaths also use the same wavelength. This additional requirement simplifies significantly the ILP formulation (and hence, the search for a solution) *with little, if any,* sacrifice in terms of optimality compared to general asymmetric solutions, as we explain and quantify in Section V.

It is possible to obtain an optimal symmetric solution by directly including constraints (7). However, the resulting formulation would be larger, and hence less scalable. In the following, we present a decomposition technique that obtains an optimal symmetric solution in times that are orders of magnitude faster: an intuitive decomposition that is applicable to instances with symmetric traffic demands is first described;

[1] Note that a symmetric solution does not constrain demands that do not have a symmetric counterpart.

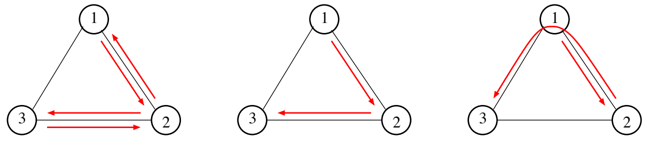

Fig. 1. RWA instance for the proof of Lemma 3.1: (a) optimal symmetric solution to the original RWA problem, (b) one optimal solution to RWA-1, (c) another optimal solution to RWA-1

then a generalized version of the decomposition for arbitrary traffic matrices is presented.

A. Symmetric Traffic Demands

In this subsection we make the additional assumption that traffic demands are symmetric, i.e., $t_{sd} = t_{ds}$, $\forall s, d \in V$; we will remove this assumption in the next subsection. Let us define two new traffic matrices $T^1 = [t^1_{sd}]$ and $T^2 = [t^2_{sd}]$ as:

$$t^1_{sd} = \begin{cases} t_{sd}, & s < d \\ 0, & \text{otherwise} \end{cases} ; \; t^2_{sd} = \begin{cases} t_{sd}, & s > d \\ 0, & \text{otherwise} \end{cases} \quad (8)$$

and the corresponding path sets: $\mathscr{P}^1 = \{p_{sd,k} \in \mathscr{P} | s < d\}$, $\mathscr{P}^2 = \{p_{sd,k} \in \mathscr{P} | s > d\}$. Note that T^1 and T^2 are lower and upper diagonal matrices, respectively, such that $T^1 + T^2 = T$ and $(T^1)^{\mathrm{T}} = T^2$, and that $\mathscr{P}^1 \cup \mathscr{P}^2 = \mathscr{P}$.

Consider the original RWA problem is decomposed into two subproblems, RWA-1 and RWA-2, such that input to RWA-1 is traffic matrix T^1 and path set $\mathscr{P}^1$, and input to RWA-2 is T^2 and $\mathscr{P}^2$. The two-step decomposition algorithm is as follows.

Algorithm SYM-RWA

1) Solve optimally subproblem RWA-1.
2) Construct a *symmetric solution* for RWA-2, as defined in expression (7), where $c^w_{sd,k}$ is the solution from Step 1.

Subproblem RWA-1 may be solved optimally in Step 1 using a *reduced* version of the ILP formulation (1)-(6) after making the following changes: (i) the new binary variables are $\{c^w_{sd,k} : s < d\}$; (ii) the new traffic constraints are those in (2) with $s < d$; and (iii) the new distinct wavelength constraints are those in (3) for links l used by paths in $\mathscr{P}^1$. This ILP formulation for RWA-1 has only one-half the decision variables $\{c^w_{sd,k}\}$ and one-half the traffic constraints (2) of the original formulation. The reduced formulation may also have fewer distinct wavelength constraints (3), if some of the links in the network are not used by paths of this subproblem. Hence, the decomposition algorithm can be applied to larger problem instances than is possible by tackling directly the original RWA problem through the ILP; we quantify the scalability of SYM-RWA in Section V.

Now, in Lemma 3.1, we show that if the ILP in (1)-(6) is used in Step 1 of the SYM-RWA algorithm (after removing the unnecessary variables and constraints, as discussed above), then the resulting symmetric solution may not be optimal (or even feasible) for the original problem. We then show how to modify the ILP so that SYM-RWA will always yield an optimal symmetric solution for the original problem.

Lemma 3.1: Consider an RWA problem instance with symmetric routing and symmetric traffic demands, RWA-1 and RWA-2 as defined above. If, in Step 1 of SYM-RWA, the formulation in (1)-(6) is used to solve RWA-1, then the solution obtained by SYM-RWA may not be an optimal (or feasible) symmetric solution for the original RWA problem.

Proof. We will construct a simple instance to show that the SYM-RWA algorithm may not always find an optimal symmetric solution. Consider an RWA instance with $N = 3$ nodes arranged in a ring topology, as shown in Figure 1, and 4 traffic demands are between node 1 and 2 (bidirectional), 2 and 3 (bidirectional), respectively. The optimal symmetric solution is shown in Figure 1(a) and requires only one wavelength.

For subproblem RWA-1 (with half of the traffic demands, i.e., from node 1 to 2, from node 2 to 3), there exist two optimal solutions, shown in Figures 1(b) and 1(c), respectively, that require a single wavelength. If Step 1 of the SYM-RWA algorithm returns the first solution, then in Step 2 the symmetric solution to the RWA-2 subproblem will yield the overall optimal symmetric solution in Figure 1(a). However, if the solution in Figure 1(c) is returned by Step 1 of the algorithm, then Step 2 of the algorithm cannot construct a feasible symmetric solution. In fact, a solution to RWA-2 that uses symmetric paths (but not wavelengths) necessarily requires one additional wavelength, such that the combined solution to the original problem is suboptimal. ■

To understand the reason underlying the negative result expressed by Lemma 3.1, observe that in the ILP formulation, subproblems RWA-1 and RWA-2 are coupled only through the distinct wavelength constraints (3). Once the decision variables $\{c^w_{sd,k}\}$ corresponding to RWA-2 are removed from the formulation used to solve RWA-1, the coupling between the two subproblems is also removed.

Fortunately, there is an easy way to account for this coupling by modifying slightly the distinct wavelength constraints (3) in the ILP formulation used to solve subproblem RWA-1. The key observation is that, although half of the paths (i.e., those that appear only in RWA-2) and the corresponding decision variables are no longer in the formulation used to solve RWA-1, due to symmetric routing, these removed paths are symmetric to the paths of RWA-1. Based on the above observation, we introduce a new set of binary parameters $Z^l_{sd,k} \in \{0,1\}$ to indicate whether link l is used by either path $p_{sd,k}$ or its symmetric path $p_{ds,k}$. We also introduce a new set of *bidirectional distinct wavelength constraints* to replace the ones in (3):

$$\sum_{s,d \in V : s<d} \sum_{k=1}^{K} c^w_{sd,k} Z^l_{sd,k} \le 1, \quad \forall l \in A, \forall w \quad (9)$$

These bidirectional constraints restore the coupling between the two subproblems in the ILP formulation used to solve subproblem RWA-1. Returning to the instance in Lemma 3.1, under these contraints, the solution shown in Figure 1(c) would require two wavelengths, making it suboptimal. Therefore, Step 1 of SYM-RWA would return the optimal solution shown in Figure 1(b), ensuring the symmetric solution constructed in Step 2 is the optimal symmetric solution of Figure 1(a).

We now have the following result.

Lemma 3.2: Consider an RWA problem instance with symmetric routing and symmetric demands, RWA-1 and RWA-2 as defined above. If, in Step 1 of the SYM-RWA algorithm, RWA-1 is solved using the ILP formulation in (1)-(6) after replacing the distinct wavelength constraints (3) with the bidirectional constraints (9), then the solution obtained by the algorithm is an optimal symmetric solution for the original RWA problem.

Proof. By definition, we have that $Z^l_{sd,k} = X^l_{sd,k} + X^l_{ds,k}$, for all s,d,k,l. Therefore, we can rewrite (9), for all l and w, as:

$$\sum_{s,d \in V: s<d} \sum_{k=1}^{K} c^w_{sd,k} X^l_{sd,k} + \sum_{s,d \in V: s<d} \sum_{k=1}^{K} c^w_{sd,k} X^l_{ds,k} \leq 1. \quad (10)$$

Since at most one of $X^l_{sd,k}$ and $X^l_{ds,k}$ may be equal to one, the above set of constraints is equivalent to the two sets:

$$\sum_{s,d \in V: s<d} \sum_{k=1}^{K} c^w_{sd,k} X^l_{sd,k} \leq 1, \quad \forall l \in A, \forall w \quad (11)$$

$$\sum_{s,d \in V: s<d} \sum_{k=1}^{K} c^w_{sd,k} X^l_{ds,k} \leq 1, \quad \forall l \in A, \forall w \quad (12)$$

Therefore, solving subproblem RWA-1 and then constructing a symmetric solution for RWA-2, results in an optimal symmetric solution for the original RWA problem. ■

Note that the proof does not depend on the objective function, hence this decomposition is applicable to RWA variants with different objective functions.

B. Arbitrary Traffic Demands

Now let us assume that traffic demands are arbitrary and not symmetric, i.e., generally $t_{sd} \neq t_{ds}$. To accommodate such asymmetric demands, we generalize the decomposition approach such that each of the two subproblems includes only one-half of the source-destination pairs and paths of the original problem, as before, but RWA-1 contains more traffic demands than RWA-2.

Let O^1 and O^2 be sets of ordered pairs of nodes such that:

$$(s,d) \in \begin{cases} O^1, & t_{sd} > t_{ds} \vee (t_{sd} = t_{ds} \wedge s < d) \\ O^2, & \text{otherwise.} \end{cases} \quad (13)$$

Then the traffic matrices $T^1 = [t^1_{sd}]$ and $T^2 = [t^2_{sd}]$ of RWA-1 and RWA-2, respectively, are:

$$t^1_{sd} = \begin{cases} t_{sd}, & (s,d) \in O^1 \\ 0, & \text{otherwise} \end{cases} ; t^2_{sd} = \begin{cases} t_{sd}, & (s,d) \in O^2 \\ 0, & \text{otherwise} \end{cases} \quad (14)$$

and correspondingly $\mathscr{P}^1 = \{p_{sd,k} \in \mathscr{P} | (s,d) \in O^1\}; \mathscr{P}^2 = \{p_{sd,k} \in \mathscr{P} | (s,d) \in O^2\}$.

Let $t^{min}_{sd} = \min\{t_{sd}, t_{ds}\}$, and $t^+_{sd} = t_{sd} - t^{min}_{sd}$. We can rewrite the traffic demands of matrix T^1 as:

$$t^1_{sd} = \begin{cases} t^{min}_{sd} + t^+_{sd}, & (s,d) \in O^1 \\ 0, & \text{otherwise} \end{cases} \quad (15)$$

With this notation, traffic matrix T^1 can be written as $T^1 = T^{1,1} + T^{1,2}$, where $T^{1,1}$ is the transpose of T^2 and $T^{1,2}$ contains the additional demands t^+_{sd} that present in T^1.

We solve the original RWA problem using the same decomposition algorithm SYM-RWA algorithm. However, the subproblem RWA-1 has a "heavier" traffic matrix than subproblem RWA-2. Specifically, there are two types of demands in RWA-1, those that have a symmetric demand in RWA-2 and those that do not. Hence, the formulation is further modified to account for the two types of traffic differently: for traffic in $T^{1,1}$ for which a symmetric demand exists in T^2, use the parameter $Z^l_{sd,k}$; and for traffic in $T^{1,2}$ for which a symmetric demand does not exist in T^2, use the parameter $X^l_{sd,k}$.

We also define the following two sets of variables that replace variables $\{c^w_{sd,k}\}$:

- $d^w_{sd,k} \in \{0,1\}$: binary variable that indicates whether there exists a lightpath assigned wavelength w on path $p_{sd,k}$ to carry traffic in $T^{1,1}$.
- $e^w_{sd,k} \in \{0,1\}$: binary variable that indicates whether there exists a lightpath assigned wavelength w on path $p_{sd,k}$ to carry traffic in $T^{1,2}$.

With above notations, the following modified formulation is used to solve RWA-1:

$$\text{minimize: } W_{high} \quad (16)$$

subject to:

- traffic constraints:

$$\sum_{k=1}^{K} \sum_{w=1}^{W} d^w_{sd,k} = t^{min}_{sd}, \quad \sum_{k=1}^{K} \sum_{w=1}^{W} e^w_{sd,k} = t^+_{sd}, \quad (s,d) \in O^1 \quad (17)$$

- distinct wavelength constraints:

$$\sum_{(s,d) \in O^1} \sum_{k=1}^{K} (d^w_{sd,k} Z^l_{sd,k} + e^w_{sd,k} X^l_{sd,k}) \leq 1, \ \forall l \in A, \forall w \quad (18)$$

- wavelength usage constraints:

$$\sum_{(s,d) \in O^1} \sum_{k=1}^{K} (d^w_{sd,k} + e^w_{sd,k}) \leq U^w P, \quad \forall w \quad (19)$$

- highest wavelength index constraints:

$$W_{high} \geq w U^w, \quad \forall w \quad (20)$$

- integrality constraints:

$$U^w = 0,1, \ \forall w; \ d^w_{sd,k} = 0,1, \ e^w_{sd,k} = 0,1, \ \forall s,d,k,w \quad (21)$$

Finally, we have the following result.

Lemma 3.3: Consider an RWA problem instance with symmetric routing and arbitrary demands, RWA-1 and RWA-2 as defined above. If, in Step 1 of the SYM-RWA algorithm, RWA-1 is solved using the ILP formulation in (16)-(21), then the solution obtained by the algorithm is an optimal symmetric solution for the original RWA problem.

Proof. The proof is similar to the proof of Lemma 3.2, and is omitted. ■

IV. Further Improvements to the ILP Formulation

We may further improve the scalability of the new ILP formulation by incorporating information from fast, high-quality heuristic algorithms for the RWA problem. We adopt the LFAP algorithm [8], a fast and conceptually simple heuristic that we have found to perform consistently well across a range of topologies [1]. For each problem instance to solve, we first run the LFAP algorithm to obtain a feasible solution that uses W_{LFAP} distinct wavelengths. This provides an upper bound on the optimal solution, that can be used to effect a further reduction in the size of the ILP formulation.

Referring to the general formulation for symmetric RWA solutions in (16)-(21), we observe that the number of variables and constraints is a function of the number of wavelengths W. Without any upper bound on the number of wavelengths for a particular instance, one might initialize W to a value (e.g., the number supported by DWDM technology) that may be much larger than necessary. Doing so would have a negative impact on the ILP solver due to: (1) a large increase in the size of the formulation; and (2) the symmetry problem, since the number of equivalent solutions obtained by wavelength permutations is proportional to $W!$. As we explain in the next section, we have found that LFAP in many cases constructs solutions that use 25% more wavelengths than the optimal one. Therefore, we set the number of wavelengths in the formulation (16)-(21) to $W = \lceil 0.8 \times W_{LFAP} \rceil$. (If this value turns out to be too low, CPLEX can determines quickly the problem is infeasible, and can then be invoked again with a higher value.)

V. Numerical Results

In this section, we present the results of an experimental study we conducted to investigate the performance of the optimal symmetric RWA formulation in terms of scalability (running time) and quality of solution. All results were obtained by running CPLEX 11 on a cluster of identical compute nodes with Woodcrest Xeon CPU at 2.33GHz with 1333MHz memory bus, 4GB of memory and 4MB L2 cache.

Our study involves a large set of problem instances defined on several network topologies with random traffic matrices. In particular, we consider the following topologies: (1) the 14-node, 42-(directed) link NSFNet; (2) the 17-node, 52-link German network; (3) the 20-node, 78-link EON network; and (4) a 32-node, 106-link USA topology. These networks have irregular topologies of increasing size that are representative of existing backbone networks, and have been used extensively in networking research. For each network topology, we consider several problem instances. For each problem instance, the traffic demand matrix $T = [t_{sd}]$ is generated by drawing the (integer) traffic demands (in units of lightpaths) uniformly at random in the interval $[0, T_{max}]$. We generate both symmetric (i.e., $t_{sd} = t_{ds}, \forall s, d$) and asymmetric traffic matrices. Each data point in the figures we present in this section represents the average of 10 random problem instances for the same settings.

A. Scalability Comparison

Let us first investigate the scalability of the the new formulation. Figure 2 compares the original formulation (1)-(6) for the RWA problem to the new formulation (16)-(21) in terms of running time, for the four networks. Recall that the original formulation obtains the overall optimal solution given the input path set, whereas the new formulation obtains the optimal symmetric solution for the same path set. The figure plots, in logarithmic scale, the CPU time it takes for CPLEX

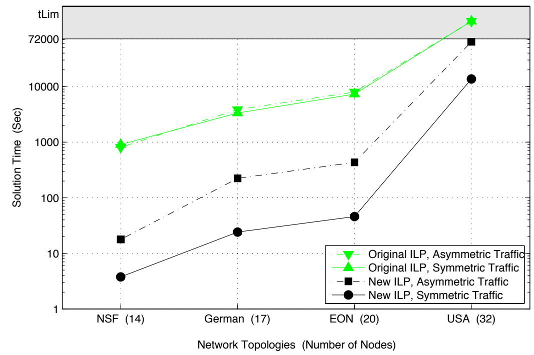

Fig. 2. CPU time comparison, $K = 2$, $T_{max} = 2$

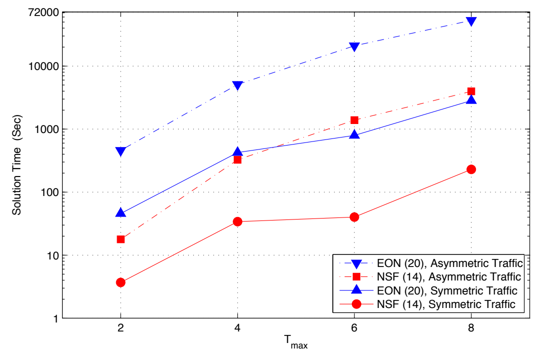

Fig. 3. CPU time against T_{max}, $K = 2$, NSF and EON networks

to find an optimal solution under symmetric and asymmetric traffic. We imposed a limit of 20 CPU hours for CPLEX to find a solution; if it failed to do so within the 20-hour limit, we terminated the execution run and plotted the data point in the light gray area of the figure labeled "tLim." For the results shown in Figure 2, there are $K = 2$ paths between each source-destination pair and the traffic matrix (symmetric or asymmetric) was generated by setting $T_{max} = 2$.

As we can see, the original path ILP formulation can find the optimal solution within the 20-hour limit for the NSF, German, and EON topologies, but not for the 32-node USA topology. Also, the running time for a given topology is the same regardless of whether traffic is symmetric or asymmetric, as the number of variables in the formulation depends on the size of the network and the traffic load (which depends on T_{max}), not on the form of the traffic matrix. The new formulation, on the other hand, performs much better than the original one. For symmetric traffic, we observe a reduction of more than two orders of magnitude in running time, and the 32-node topology can be solved in only about 3 hours. For asymmetric traffic, the new formulation achieves a reduction in running time of more than one order of magnitude (up to a factor of 25), solving the 32-node network in about 16 hours. The main reason for the higher running time in the case of asymmetric traffic is due to the additional variables that need to be included in the formulation, as we explained in Section III-B. Importantly, symmetric instances on the NSF, Germany, and EON topologies may be solved in about one minute and asymmetric instances on the same topologies take

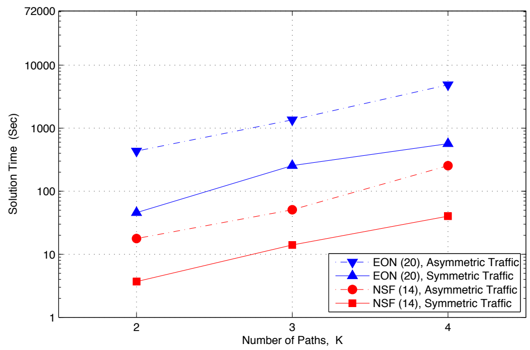

Fig. 4. CPU time against K, $T_{max} = 2$, NSF and EON networks

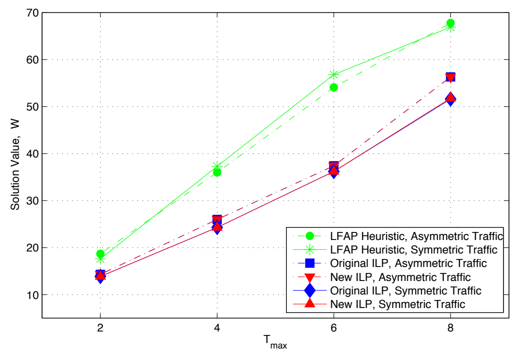

Fig. 5. Optimality comparison for NSFNet, $T_{max} = 2$, $K = 2$

about ten minutes to solve. As a result, the new formulation makes it possible to solve networks of practical size fast enough to allow network designers and operators to perform extensive "what-if" analysis so as to investigate large numbers of scenarios regarding forecast demands. We have obtained similar results for several other topologies, including regular (torus) topologies (omitted due to space constraints).

Figure 3 plots the running time of the new formulation for the NSF and EON networks as a function of the value of parameter T_{max}. We observe that the running time increases with the traffic load, as expected, due to the larger number of wavelengths needed to accommodate the traffic. Note that a value of $T_{max} = 8$, the highest one considered in this experiment, corresponds to an average of four lightpaths between a source-destination pair (in each direction); for the NSF and EON networks, the optimal solution for this value would require more than 50 and 70 wavelengths, respectively. Although such instances are well beyond what can be realized in deployed networks, the new formulation is capable of solving them efficiently.

Finally, Figure 4 plots the running time of the new formulation as a function of the number K of candidate paths, for the NSF and EON networks and $T_{max} = 2$. Again, the running time increases with K due to the larger number of variables in the formulation. Nevertheless, the new formulation makes it possible to solve realistic instances in a short amount of time. We also note that, of the 80 instances (i.e., ten instances per data point) we run with $K > 2$, only one instance resulted in an optimal value better (by a single wavelength) than the one obtained with $K = 2$.

B. Quality of Optimal Symmetric Solution

Let us now turn our attention to the quality of the optimal symmetric solution. Figure 5 plots the value of the solution returned by the original formulation, the new formulation, and the LFAP heuristic [8] for NSFNet, as a function of T_{max}. Recall that LFAP and the original formulation consider the whole solution space, while the new formulation only considers symmetric solutions to the RWA problem as defined in (7). The new formulation obtained the *same* optimal solution as the original one for 79 of 80 instances from which the results in the figure are generated. There was only one instance ($T_{max} = 8$, symmetric traffic) for which the optimal symmetric solution required two additional wavelengths (52 versus 50 for the overall optimal). We obtained similar results for EON network (omitted), in which the new approach resulted in the same optimal solution as the original one over all 80 instances. These findings confirm the intuitive view that, except in rare cases when the demand matrix creates severe bottlenecks, optimal symmetric solutions are also optimal overall.

We also observe that in NSF network, LFAP, one of the best heuristic algorithms for the RWA problem, constructs solutions that are between 25-60% higher than the optimal one. Hence, while LFAP is quite fast, relying on such a heuristic may result in significantly higher costs in deploying and operating the network. Our new formulation, on the other hand, achieves an excellent tradeoff between running time and optimality.

VI. Concluding Remarks

We have presented a new ILP formulation to construct optimal symmetric solutions to the RWA problem, that scales well to network topologies encountered in practice and enables network designers and operators to carry out extensive "what-if" analysis. We also demonstrated that optimal symmetric solutions, in addition to their practical advantages, often achieve the overall optimal or a value very close to it.

References

[1] B. Chen, G. Rouskas, and R. Dutta. On hierarchical traffic grooming in WDM networks. *IEEE/ACM Trans. Netw.*, 16(5):1226-1238, Oct. 2008.
[2] E. Chen and T. Bates. Current practice of implementing symmetric routing and load sharing in the multi-provider Internet. IETF Draft <draft-ietf-idr-symm-multi-prov-01.txt>, June 1995.
[3] K. Christodoulopoulos, K. Manousakis, and E. A. Varvarigos. Offline routing and wavelength assignment in transparent WDM networks. *IEEE/ACM Trans. Netw.*, 18(5):1557-1570, Oct. 2010.
[4] R. Dutta and G. N. Rouskas. A survey of virtual topology design algorithms for wavelength routed optical networks. *Optical Networks*, 1(1):73–89, January 2000.
[5] B. Jaumard and etc. Comparison of ILP formulations for the RWA problem. *Optical Switching and Networking*, 3-4:157–172, 2007.
[6] B. Jaumard and etc. On column generation formulations for the RWA problem. *Discrete Applied Mathematics*, 157(6):1291–1308, 2009.
[7] W. John, M. Dusi, and kc claffy. Estimating routing symmetry on single links by passive flow measurements. *IWCMC'10*, Caen, France, 2010.
[8] H. Siregar, H. Takagi, and Y. Zhang. Efficient routing and wavelength assignment in wavelength-routed optical networks. *Proc. 7th Asia-Pacific Network Oper. and Mgmt Symposium*, pages 116–127, Oct. 2003.
[9] E. Yetginer, Z. Liu, and G. N. Rouskas. Fast exact ILP decompositions for ring RWA. *Journal of Optical Communications and Networking*, vol. 3, no. 7, 577-586, July 2011.