Hierarchical Traffic Grooming Formulations

Hui Wang, George N. Rouskas
Operations Research and Department of Computer Science, North Carolina State University, Raleigh, NC 27695-8206 USA

***Abstract*—Hierarchical traffic grooming facilitates the control and management of multigranular WDM networks. We define the hierarchical virtual topology and traffic routing (H-VTTR) problem, the grooming-specific subproblem of traffic grooming, and we present a suite of ILP formulations to solve it. The formulations represent various tradeoffs between solution quality and running time.**

I. Introduction

Traffic grooming is the field of study that is concerned with the development of algorithms and protocols for the design, operation, and control of networks with multigranular bandwidth demands [5]. Several variants of the traffic grooming problem have been studied in the literature under a range of assumptions regarding the network topology, the nature of traffic, and the optical and electronic switching model [6], [9]–[11], [16], [20]. Typically, an integer linear programming (ILP) formulation serves as the basis for reasoning about and tackling the offline problem. Most studies and formulations regard the network as a flat entity in the sense that grooming of traffic may take place at any node. Unfortunately, solving the ILP directly does not scale to instances with more than a handful of nodes, and cannot be applied to networks of practical size.

As the number of logical entities (including sub-wavelength channels, wavelengths, wavebands, and fibers) that need to be controlled in a multigranular network increases rapidly with the network size, wavelength capacity, and load, a scalable framework for managing these entities becomes essential for wide area WDM networks. In fact, network resources are typically managed and controlled in a hierarchical manner. The levels of the hierarchy either reflect the underlying organizational structure of the network or are designed in order to ensure scalability of the control and management functions. Accordingly, several studies have adopted a variety of hierarchical approaches to traffic grooming that, by virtue of decomposing the network, scale well and are more compatible with the manner in which networks operate in practice.

The study in [6] was the first to present several hierarchical ring architectures and to evaluate them under a model of dynamic traffic. Specifically, single-hub and double-hub ring structures were considered, as well as a more general hierarchical architecture in which ring nodes are partitioned into two types: *access* and *backbone*. A similar hierarchical ring structure was considered in [4] that used local (access) and bypass (backbone) wavebands to route traffic. A different hierarchical approach for grooming sub-wavelength traffic in ring networks was introduced in [14], in which the nodes are grouped into *super-nodes*, where each super-node consists of several consecutive ring nodes. The idea behind this partitioning is to pack (groom) all traffic from one super-node to another onto lightpaths that are routed directly between the two super-nodes. Finally, [6] also proposes the decomposition of a ring into contiguous segments; these are similar to the super-nodes of [14] but are referred to as *subnets*. With this decomposition, the ring network is organized in a hierarchical manner as a tree of subnets.

A hierarchical approach for networks with a torus or tree topology was presented in [4], and is based on embedding rings on the underlying topology and then selecting hub nodes along each ring and using bypass wavelengths to interconnect the hubs. Finally, a hierarchical grooming algorithm for networks with a star topology was developed in [2].

A framework for hierarchical traffic grooming that is applicable to networks with a general topology was presented in [3], and emulates the hub-and-spoke model used by the airline industry to “groom” passenger traffic onto connecting flights. Specifically, the network is first partitioned into clusters of nodes that form the first level of the hierarchy. Within each cluster, one node is designated as the *hub*, and is responsible for grooming intra-cluster traffic as well as inter-cluster traffic originating or terminating locally. Hub nodes collectively form the second level of the hierarchy, and are expected to be provisioned with more resources (e.g., larger number of switching ports and higher capacity for grooming traffic) than non-hub nodes. Returning to the airline analogy, a hub node is similar in function to airports that serve as major hubs. The hierarchical grooming algorithm of [3] takes less than a second to construct the virtual topology for networks with fifty or more nodes. However, the algorithm considers each cluster in isolation as a *virtual star*, and applies the grooming method for star networks in [2], regardless of the actual physical topology of the cluster. Consequently, the algorithm examines only a subset of the hierarchical traffic grooming solution space.

In this paper, we define several variants of the hierarchical grooming problem so as to explore the spectrum of solutions between (1) the flat grooming approach that is the subject of most studies, and (2) the hierarchical grooming algorithm of [3]. The rest of this paper is organized as follows. In Section II, we review the virtual topology and traffic routing (VTTR) subproblem of traffic grooming. In Section III, we define the hierarchical VTTR (H-VTTR) problem, and present several variants that arise naturally. We present a performance study of the problem variants in Section IV, and we conclude in Section V.

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

II. The Virtual Topology and Traffic Routing (VTTR) Problem

Consider a connected graph $G = (\mathcal{V}, \mathcal{L})$, where $\mathcal{V}$ denotes the set of nodes and $\mathcal{L}$ denotes the set of directed links (arcs) in the network. We define $L = |\mathcal{L}|$ as the number of links. Each directed link $l \in \mathcal{L}$ consists of an optical fiber that may support W distinct wavelengths indexed as $1, 2, \ldots, W$. Let $T = [t^{sd}]$ denote the traffic demand matrix, where t^{sd} is a non-negative integer representing the traffic demand units 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 $t^{ss} = 0, \forall s$. Finally, we denote C as the capacity of a single wavelength channel in terms of traffic units.

We are interested in designing the network so as to carry all the traffic demands with the minimum total number of lightpaths; such an objective minimizes the use of critical resources and provides ample flexibility for future expansion of the network. This traffic grooming problem involves the following conceptual subproblems [17]:

1) *virtual topology and traffic routing (VTTR):* find a set of lightpaths to carry the offered traffic and route the traffic components over the lightpaths; and
2) *lightpath routing and wavelength assignment (RWA):* assign a wavelength and path over the physical topology to each lightpath.

The VTTR subproblem constitutes the grooming aspect of the problem and is defined formally as follows:

Definition 2.1 (VTTR): Given the traffic demand matrix T and the wavelength capacity C, establish the minimum number of lightpaths to carry all traffic demands.

In [17], we proposed a decomposition of the traffic grooming problem, where the objective is to minimize the number of lightpaths, into the VTTR and RWA subproblems that are then solved sequentially. We have shown in [17] that, whenever the network is not wavelength (bandwidth) limited, this sequential solution yields an optimal solution to the original traffic grooming problem. In [18], we developed partial LP relaxation techniques to solve the VTTR problem efficiently. We have also developed scalable optimal or near-optimal RWA algorithms for ring and mesh topologies in [12], [13], [19].

Note that the *VTTR* problem does not take as input the network graph G, only the traffic demand matrix T (and, hence, the number of nodes, $|\mathcal{V}|$). Consequently, the output of the problem is simply the set of lightpaths to be established but *not* the (physical) paths that these lightpaths take in the network. The physical path and wavelength for each lightpath included in the solution to VTTR must be determined in a second step by running an RWA algorithm on the given network graph G.

III. The Hierarchical VTTR Problem and Variants

In this work, we focus on hierarchical solutions to the VTTR problem. To define the *hierarchical VTTR (H-VTTR)* problem, we assume that a set $\mathcal{H} \subset \mathcal{V}$ of hub nodes in the network is given. Hub nodes are nodes with traffic grooming capabilities. However, in contrast with the work in [3] (and the problem variant we discuss in the following subsection), no clusters are defined in the network; in other words, non-hub nodes are *not* assumed to be assigned to clusters and associated with a "local" hub. We also let $\mathcal{N} = \mathcal{V} \backslash \mathcal{H}$ be the set of non-hub nodes, and $K = |\mathcal{H}|$ be the number of hubs.

Definition 3.1 (H-VTTR): Given the set $\mathcal{V}$ of nodes in the graph G, the set of hubs $\mathcal{H}$, the wavelength capacity C, and the traffic demand matrix T, establish the minimum number of lightpaths to carry all traffic demands, under two constraints: (1) only hub nodes may groom traffic that they do not themselves originate or terminate, and (2) no direct lightpaths between two non-hub nodes (i.e., nodes in $\mathcal{N}$) are allowed.

H-VTTR is a generalization of the VTTR problem we defined in the previous section and studied in [18]. Specifically, VTTR allows grooming of traffic to take place at any node in the network, as well as lightpaths to exist between any pair of nodes in the network. Therefore, if we let $\mathcal{H} = \mathcal{V}$ and $\mathcal{N} = \emptyset$, i.e., each node to be a hub node, H-VTTR reduces to VTTR. Note also that, because of the constraint on direct lightapths, traffic between two non-hub nodes has to be carried on at least two lightpaths via at least one hub node. The ILP formulation of the H-VTTR problem is provided in Appendix A.

In the following, we introduce several variants of the basic H-VTTR problem above.

A. H-VTTR with Clustering (HC-VTTR)

The hierarchical VTTR with clustering (HC-VTTR) problem is a variant of H-VTTR that adopts the concept of clustering considered in [3]. Specifically, we assume that the set $\mathcal{V}$ of network nodes is partitioned into $K = |\mathcal{H}|$ clusters, $v_1, \ldots, v_K$, and that node $h_i \in \mathcal{H}$ is the hub node of cluster v_i. In HC-VTTR, traffic originating from, or terminating at, a non-hub node in cluster v_i may only be groomed with other traffic at the local hub h_i. More formally, we have the following definition.

Definition 3.2 (HC-VTTR): Given the set $\mathcal{V}$ of nodes in the graph G, the set of hubs $\mathcal{H}$, a set of $K = |\mathcal{H}|$ clusters $\{v_1, \ldots, v_H\}$ such that each node $h_i \in \mathcal{H}$ is the hub of cluster v_i, the wavelength capacity C, and the traffic demand matrix T, establish the minimum number of lightpaths to carry all traffic demands, under three constraints: (1) only hub nodes may groom traffic that they do not originate or terminate, (2) traffic originating from, or terminating at, a non-hub node in cluster v_i may only be groomed with other traffic at the local hub h_i, and (3) no direct lightpaths between two non-hub nodes (i.e., nodes in $\mathcal{N}$) are allowed.

The key idea in HC-VTTR is to ensure that grooming of traffic takes place "near" non-hub nodes (i.e., at their local hub). Local grooming handles small traffic demands efficiently, and it prevents solutions with long underutilized lightpaths. On the other hand, traffic between two non-hub nodes in different clusters must be carried on at least three lightpaths: from the source node to its local hub, then to the remote hub, and finally to the destination node. Although we omit the ILP formulation of the HC-VTTR problem, it is similar to that of the H-VTTR problem with additional constraints to prevent the establishment of lightpaths between a non-hub node and hubs other than the one in its own cluster.

B. Hierarchical Grooming with Direct Lightpaths

The H-VTTR and HC-VTTR problems explicitly prevent direct lightpaths between non-hub nodes. Note, however, that if there is sufficient traffic between two non-hub nodes to fill a lightpath, forcing this traffic to travel via a hub node results in more lightpaths: sending the traffic directly to its destination requires only one lightpath, whereas sending it through one or more hubs requires at least two lightpaths without improving the grooming of other traffic (since this traffic takes up the whole capacity of these lightpaths). Our experience [3] also indicates that it is often cost-effective to establish partially filled direct lightpaths as long as these lightpaths have high utilization (i.e., the traffic between the two non-hub nodes is close to the capacity of a lightpath). Such high direct traffic demands may not present effective opportunities to groom other traffic on the same lightpaths; furthermore, including partially filled lightpaths in the solution makes it possible to accommodate future increases in traffic demands without the need to establish new lightpaths, an important consideration for long-term network planning.

We now formally define the H-VTTR problem with direct lightpaths (H-VTTR/DL):

Definition 3.3 (H-VTTR/DL): Given the set $\mathcal{V}$ of nodes in the graph G, the set of hubs $\mathcal{H}$, the wavelength capacity C, the traffic demand matrix T, and a threshold $\theta, 0 < \theta \leq 1$, establish the minimum number of lightpaths to carry all traffic demands, under two constraints: (1) only hub nodes may groom traffic that they do not originate or terminate, and (2) direct lightpaths between two non-hub nodes (i.e., nodes in $\mathcal{N}$) are allowed only if the traffic between these nodes is at least equal to θC.

The ILP formulation of H-VTTR/DL is presented in Appendix A.

A similar HC-VTTR/DL problem with clustering can be defined, in which direct lightpaths between non-hub nodes, or a non-hub node and a remote hub, are allowed as long as the traffic between these nodes is at least equal to θC. The ILP formulation of the problem is omitted, but it is similar to the formulation of H-VTTR/DL with additional constraints.

The HC-VTTR/DL problem is identical to the one studied in [3]. But whereas the virtual topology algorithm developed in [3] treated each cluster in isolation as a virtual star and used a heuristic to determine the lightpaths, the ILP formulation we developed in this work considers the clusters in an integrated manner and solves the HC-VTTR/DL problem optimally.

IV. Numerical Results

In this section we evaluate the performance of hierarchical solutions to the VTTR problem in terms of two metrics: quality of solution (i.e., the number of lightpaths produced by the solution) and running time. Specifically, we compare the following five ILP formulations:

1) H-VTTR (the problem is defined in Section III and the ILP formulation is shown in Appendix A);
2) H-VTTR/DL (the problem is defined in Section III-B and the formulation is shown in Appendix A);
3) HC-VTTR (the problem is defined in Section III-A);
4) HC-VTTR/DL (the problem is defined in Section III-B); and
5) VTTR (the problem is defined in Section II).

Note that the VTTR ILP formulation is similar to the one for H-VTTR shown in Appendix A, but takes a flat view of the network such that grooming may take place at any node, not just hubs, and lightpaths are allowed between any pairs of nodes without any threshold constraints on the traffic demands. Since the four hierarchical formulations we presented in this paper are derived from the VTTR formulation by adding appropriate constraints, the solution to the VTTR formulation provides a lower bound for the solutions to the hierarchical formulations. (In fact, as we showed in [17], the solution to the VTTR formulation is a lower bound to the solution of the original traffic grooming problem, and it is optimal whenever the network is not wavelength limited.) Hence, we are interested in characterizing the performance of the hierarchical solutions relative to the baseline VTTR formulation.

In our study we consider four network topologies (link counts refer to directed links): the 14-node, 42-link NSFNet [15]; the 17-node, 52-link German network [8]; the 32-node, 106-link network we studied in [3]; and the 47-node, 192-link network from [1]. For each problem instance, we generate the traffic matrix $T = [t^{sd}]$ by drawing each traffic demand T^{sd} uniformly and randomly in the interval $[0, t_{max}]$. Each data point in the following figures is the average of ten problem instances. For the experiments, we fix the wavelength capacity $C = 16$, and we vary the value of parameter $t_{max} = 10, 20, 30, 40, 50, 60$, to investigate various traffic loads. For the two formulations that allow direct lightpaths, we fixed the threshold value to $\theta = 0.6$, as our experiments indicate that this value represents the best tradeoff between running time and solution quality; results that support this finding are omitted due to space constraints but are available in the dissertation of the first author. The results we present were obtained by running the IBM CPLEX 12 optimization tool on a cluster of identical compute nodes with dual Woodcrest Xeon CPU at 2.33GHz with 1333MHz memory bus, 4GB of memory and 4MB L2 cache. We imposed a 3% relative optimality gap in solving the optimization problems with CPLEX.

Figures 1 and 2 compare the five formulations above across the four network topologies, in terms of the objective value and the CPU time it takes CPLEX to solve them, respectively. For these experiments, we set $t_{max} = 40$, and we used the k-center algorithm [7] to determine the hubs for each topology. Specifically, we set the number of hubs to four for the 14-

and 17-node topologies, and eight for the 32- and 47-node networks. We also set a time limit of two hours. As we can see, CPLEX was able to solve all the formulations within the time limit, except for the VTTR formulation on the 47-node network; hence, the two figures do not present results for this formulation and topology.

Let us first refer to Figure 1 that compares the five formulations in terms of solution quality. We first note that the objective value increases with the size of the network topology, as expected: for a given value of t_{max}, a larger network has more traffic to carry than a smaller one, requiring a larger number of lightpaths. We also note that the objective value obtained by solving the HC-VTTR formulation is always higher than that obtained by H-VTTR. Recalling the problem definitions, HC-VTTR includes more constraints than H-VTTR: in the former, traffic from a non-hub node must be groomed at the local hub, whereas in the latter it may be groomed at any hub node. Therefore, the solution to HC-VTTR cannot be better than that to H-VTTR. Also, the variants that allow for direct lightpaths (H-VTTR/DL and HC-VTTR/DL) lead to solutions that are better than variants that do not allow direct lightpaths (H-VTTR and HC-VTTR, respectively). Again, this result can be explained by the fact that allowing direct lightpaths increases the space of candidate solutions. Finally, the original VTTR formulation produces the best solution, as expected, for the three topologies for which a solution to this formulation was obtained within the time limit. However, in all three cases, the solution to H-VTTR/DL is very close to that of VTTR. Overall, the relative performance of the five formulations is consistent across the four topologies: VTTR leads to the best solution, followed by H-VTTR/DL, HC-VTTR/DL, H-VTTR, and HC-VTTR, in this order.

Let us now turn our attention to Figure 2 that compares the running time for the five formulations. We observe that solving the HC-VTTR formulation takes the least amount of time, less than a second, on average, even for the 47-node network. Among the hierarchical formulations, the next fastest solution time is achieved by HC-VTTR/DL, followed by H-VTTR and H-VTTR/DL. We also note that, for a given formulation, the running time is similar for the 14- and 17-node networks, and is also similar (but higher) for the 32- and 47-node networks. On the other hand, for the two small networks, the VTTR formulation that does not impose any hierarchical structure on the topology, takes about the same time as H-VTTR/DL, the hierarchical formulation with the worst running time. But whereas the running time of H-VTTR/DL increases by a small factor as we move from the 17- to the 32-node network, the running time of VTTR increases by almost three orders of magnitude; similarly, the running time of H-VTTR/DL increases slightly from the 32- to the 47-node network, but the running time of VTTR increases significantly and exceeds the two-hour limit we imposed. From these results, we conclude that imposing a hierarchical structure on the virtual topology is not beneficial in terms of running time when the size of the network is relatively small (in our study, up to 17 nodes). However, as the network size grows, flat solutions (i.e., VTTR)

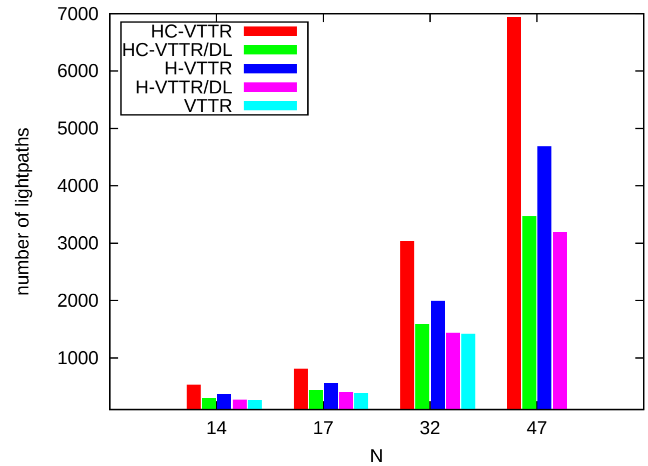

Fig. 1. Objective value comparison, $t_{max} = 40$

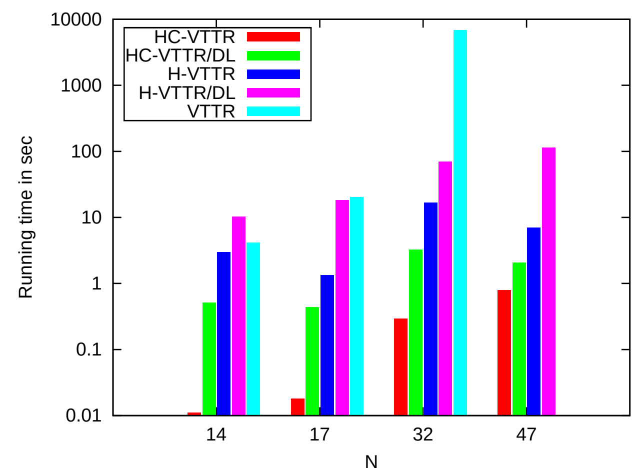

Fig. 2. CPU time comparison, $t_{max} = 40$

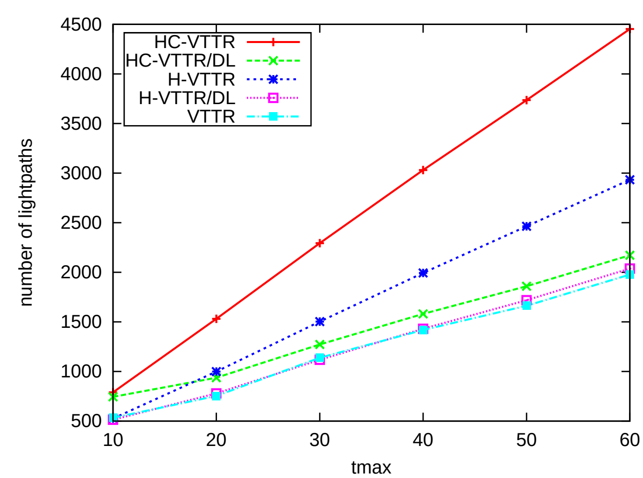

Fig. 3. Objective value comparison, 32-node network, $K = 8$ hubs

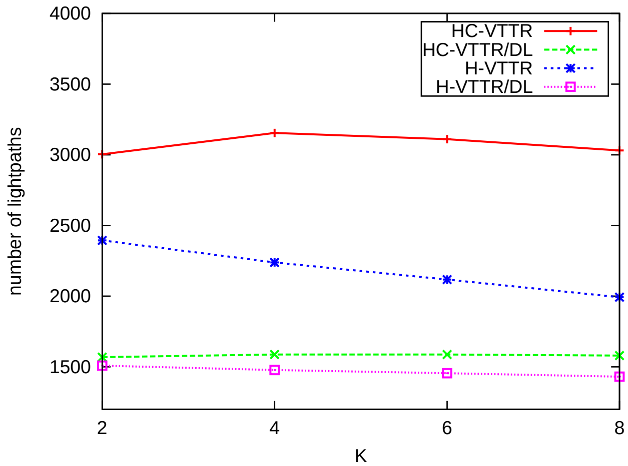

Fig. 4. Objective value comparison, 32-node network, $t_{max} = 40$

do not scale whereas hierarchical solutions scale quite well; indeed, it is at larger network sizes that one would expect the benefits of hierarchical structures to materialize. Overall, these results indicate that H-VTTR/DL represents the best tradeoff between running time and solution quality, as it takes, on average, about 100 seconds or less to obtain solutions close to the optimal (i.e., that obtained by VTTR).

In Figure 3 we compare the solution quality of the five formulations as a function of t_{max}, i.e., the traffic load. The results shown are for the 32-node network with eight hubs; the relative behavior of the various curves is representative of that for the other topologies. As the traffic load increases, the number of lightpaths increases almost linearly, but the rate of increase depends on the particular formulation. We also observe that the relative performance across the various values of t_{max} is similar to that in Figure 1, i.e., HC-VTTR requires the largest number of lightpaths, followed by H-VTTR and HC-VTTR/DL, while H-VTTR/DL and VTTR have very similar objective values. These reults further support our earlier conclusion that H-VTTR/DL provides the best tradeoff between running time and solution quality.

Finally, Figure 4 shows the effect of the number K of hubs on the objective value for the four hierarchical grooming formulations. As we can see, the number of hubs has little effect on the number of lightpaths for formulations that allow direct lightpaths; this result is due to the fact that a good amount of traffic is sent over such direct lightpaths and hence the number of hubs is not very important. For H-VTTR, as the number of hubs increases, the solution improves; since H-VTTR allows a node to send traffic to any hub, increasing the number of hubs also increases the number of candidate solutions. HC-VTTR, on the other hand, requires each node to send its traffic to the local hub, hence increasing the number of hubs may also initially increase the overall number of lightpaths. The running time (not shown due to page constraints) increases by about two orders of magnitude from $K = 2$ to $K = 8$ across all formulations. Therefore, if one of the formulations that allow direct lightpaths is adopted, these results indicate that a smaller number of hubs should be used.

V. Concluding Remarks

Hierarchical traffic grooming is an efficient and scalable approach to grooming multigranular traffic in large-scale WDM networks with a general topology. We presented a number of ILP formulations for solving the virtual topology and traffic routing subproblem of traffic grooming in a hierarchical manner. The formulations have been shown to perform well over a range of network topologies and traffic patterns, and scale to networks of realistic size.

References

[1] P. Baran. On distributed communications networks. *IEEE Transactions on Communications*, 12(1):1–9, March 1964.
[2] B. Chen, R. Dutta, and G. N. Rouskas. Traffic grooming in star networks. In *Proceedings of the First Traffic Grooming Workshop*, October 2004.
[3] B. Chen, G. N. Rouskas, and R. Dutta. On hierarchical traffic grooming in WDM networks. *IEEE/ACM Transactions on Networking*, 16(5):1226–1238, October 2008.
[4] L.-W. Chen and E. Modiano. Dynamic routing and wavelength assignment with optical bypass using ring embeddings. *Optical Switching and Networking*, 1(1):35–49, 2005.
[5] R. Dutta and G. N. Rouskas. Traffic grooming in WDM networks: Past and future. *IEEE Network*, 16(6):46–56, November/December 2002.
[6] O. Gerstel, R. Ramaswami, and G. Sasaki. Cost-effective traffic grooming in WDM rings. *IEEE/ACM Transactions on Networking*, 8(5):618–630, October 2000.
[7] T. Gonzalez. Clustering to minimize the maximum inter-cluster distance. *Theoret. Comput. Sci.*, 38:293–306, 1985.
[8] W.D. Grover and D.P. Onguetou. A new approach to node-failure protection with span-protecting p-cycles. In *Proceedings of ICTON 2009*, June 2009.
[9] J.Q. Hu and B. Leida. Traffic grooming, routing, and wavelength assignment in optical WDM mesh networks. In *Proceedings of IEEE INFOCOM 2004*, pages 495–501, March 2004.
[10] V. R. Konda and T. Y. Chow. Algorithm for traffic grooming in optical networks to minimize the number of transceivers. In *IEEE Workshop on High Performance Switching and Routing*, pages 218–221, 2001.
[11] D. Li, Z. Sun, X. Jia, and K. Makki. Traffic grooming on general topology WDM networks. *IEE Proc. Commun.*, 150(3):197–201, June 2003.
[12] Z. Liu and G. N. Rouskas. A fast path-based ILP formulation for offline RWA in mesh optical networks. In *Proceedings of IEEE GLOBECOM 2012*, December 2012.
[13] Z. Liu and G. N. Rouskas. Link selection algorithms for link-based ILPs and applications to RWA in mesh networks. In *Proceedings of ONDM 2013*, April 2013.
[14] J. Simmons, E. Goldstein, and A. Saleh. Quantifying the benefit of wavelength add-drop in WDM rings with distance-independent and dependent traffic. *IEEE/OSA Journal of Lightwave Technology*, 17:48–57, Jan 1999.
[15] B. Mukherjee *et al.* Some principles for designing a wide-area WDM optical network. *IEEE/ACM Transactions on Networking*, 4(5):684–696, October 1996.
[16] P-J. Wan, G. Calinescu, L. Liu, and O. Frieder. Grooming of arbitrary traffic in SONET/WDM BLSRs. *IEEE Journal on Selected Areas in Communications*, 18(10):1995–2003, 2000.
[17] H. Wang, Z. Liu, and G. N. Rouskas. Scalable traffic grooming in optical networks. In *Proceedings of ACP 2012*, November 2013.
[18] H. Wang and G. N. Rouskas. An efficient algorithm for solving traffic grooming problems in optical networks. In *Proceedings of IEEE ICC 2013*, June 2013.
[19] E. Yetginer, Z. Liu, and G. N. Rouskas. Fast exact ILP decompositions for ring RWA. *Journal of Optical Communications and Networking*, 3(7):577–586, July 2011.
[20] K. Zhu and B. Mukherjee. Traffic grooming in an optical WDM mesh network. *IEEE Journal on Selected Areas in Communications*, 20(1):122–133, Jan 2002.

APPENDIX A
ILP FORMULATION OF H-VTTR AND H-VTTR/DL

We now present a formulation of the H-VTTR problem. We use the following notation:

- $\mathcal{H}$ denotes the set of hub nodes, and $K = |\mathcal{H}|$ denotes the number of hub nodes.
- $\mathcal{N}$ is the set of non-hub nodes, and $N = |\mathcal{N}|$ represents the number of non-hub nodes.
- $\mathcal{L}$ denotes the set of (directed) physical links, and $L = |\mathcal{L}|$ is the number of links.
- $T = \{t^{sd}\}$ is the traffic demand matrix representing demands from any source node s to any destination node d.
- $\mathcal{Z} = \{(i,j) | i \in \mathcal{H}$ or $j \in \mathcal{H}, i \neq j\}$ is the set of pairs of nodes such that at least one node in the pair is a hub node; in other words, $\mathcal{Z}$ is the set of pairs of nodes between which direct lightpaths are allowed.

We also define these decision variables:

- b_{ij}, $(i,j) \in \mathcal{Z}$: the number of lightpaths originating at node i and terminating at node j.
- $t^{s,d}_{h_i,h_j}$: the amount of traffic from any source node s to destination node d carried on lightpaths from hub node h_i to hub node h_j.
- $t^{n,d}_{n,h_i}$: the amount of traffic from non-hub node n to any destination node d carried on lightpaths from n to hub node h_i.
- $t^{s,n}_{h_i,n}$: the amount of traffic from any source node s to a non-hub node n carried on lightpaths from hub node h_i to n.

With these definitions, we have the following multi-commodity flow formulation for the H-VTTR problem:

Objective function: minimize the number of lightpaths

$$\min \sum_{(i,j) \in \mathcal{Z}} b_{ij} \tag{1}$$

Constraints:

Capacity Constraint

$$\sum_{s,d \in \mathcal{N} \cup \mathcal{H}, s \neq d} t^{sd}_{ij} \leq b_{ij} C, \quad (i,j) \in \mathcal{Z} \tag{2}$$

Flow Conservation Constraints at Intermediate Nodes

$$\sum_{h_j \in \mathcal{H}, h_j \neq h_i} t^{sd}_{h_i h_j} + t^{sd}_{h_i d} - \sum_{h_j \in \mathcal{H}, h_j \neq h_i} t^{sd}_{h_j h_i} - t^{sd}_{s h_i} = 0, \quad s,d \in \mathcal{N}, h_i \in \mathcal{H} \tag{3}$$

$$\sum_{h_j \in \mathcal{H}, h_j \neq h_i} t^{sd}_{h_i h_j} - \sum_{h_j \in \mathcal{H}, h_j \neq h_i} t^{sd}_{h_j h_i} = 0, \quad s,d \in \mathcal{H}, h_i \in \mathcal{H}, h_i \neq s, h_i \neq d \tag{4}$$

$$\sum_{h_j \in \mathcal{H}, h_j \neq h_i} t^{sd}_{h_i h_j} + t^{sd}_{h_i d} - \sum_{h_j \in \mathcal{H}, h_j \neq h_i} t^{sd}_{h_j h_i} = 0, \quad s \in \mathcal{H}, d \in \mathcal{N}, h_i \in \mathcal{H}, h_i \neq s \tag{5}$$

$$\sum_{h_j \in \mathcal{H}, h_j \neq h_i} t^{sd}_{h_i h_j} - \sum_{h_j \in \mathcal{H}, h_j \neq h_i} t^{sd}_{h_j h_i} - t^{sd}_{s h_i} = 0, \quad s \in \mathcal{N}, d \in \mathcal{H}, h_i \in \mathcal{H}, h_i \neq d \tag{6}$$

Flow Conservation Constraints at Source Nodes

$$\sum_{h_i \in \mathcal{H}, h_i \neq s} t^{sd}_{s h_i} = t^{sd}, \quad s \neq d, s,d \in \mathcal{H}, \text{ or } s,d \in \mathcal{N}, \text{ or } s \in \mathcal{N}, d \in \mathcal{H} \tag{7}$$

$$\sum_{h_i \in \mathcal{H}, h_i \neq s} t^{sd}_{s h_i} + t^{sd}_{sd} = t^{sd}, \quad s \neq d, s \in \mathcal{H}, d \in \mathcal{N} \tag{8}$$

$$\sum_{h_i \in \mathcal{H}, h_i \neq s} t^{sd}_{h_i s} = 0, \quad s \in \mathcal{H}, d \in \mathcal{H} \cup \mathcal{N} \tag{9}$$

Flow Conservation Constraints at Destination Nodes

$$\sum_{h_i \in \mathcal{H}, h_i \neq d} t^{sd}_{d h_i} = 0, \quad s \in \mathcal{H} \cup \mathcal{N}, d \in \mathcal{H} \tag{10}$$

$$\sum_{h_i \in \mathcal{H}, h_i \neq d} t^{sd}_{h_i d} = t^{sd}, \quad s \neq d, s,d \in \mathcal{H}, \text{ or } s,d \in \mathcal{N}, \text{ or } s \in \mathcal{H}, d \in \mathcal{N} \tag{11}$$

$$\sum_{h_i \in \mathcal{H}, h_i \neq d} t^{sd}_{h_i d} + t^{sd}_{sd} = t^{sd}, s \neq d, s \in \mathcal{N}, d \in \mathcal{H} \tag{12}$$

Constraint (2) ensures that enough lightpaths are established to satisfy the traffic demand between each pair of nodes. Constraints (3) to (12) are the flow conservation constraints at intermediate nodes ((3-(6)), source nodes ((7)-(9)), and destination nodes ((10)-(12)). There are different flow conservation constraints depending on whether the source and destination nodes are hub or non-hub nodes, as this determines whether a drect lightpaths can be established between the two.

A. ILP Formulation of H-VTTR/DL

The only difference between the H-VTTR/DL and H-VTTR problems is that in the former we allow direct lightpaths between non-hub nodes, whereas such lightpaths are not allowed in the latter. Therefore, the formulation of H-VTTR/DL is very similar to (1)-(12), with the following differences:

- The set $\mathcal{Z}$ is redefined as: $\mathcal{Z} = \{(i,j) | i,j \in \mathcal{H} \cup \mathcal{N}, i \neq j$ to allow direct lightpaths between any pair of nodes;
- Flow conservation constraints (7) and (11) are removed;
- Flow conservation constraints (8) and (12) are modified to apply to all $s,d \in \mathcal{H} \cup \mathcal{N}, s \neq d$; and
- The following constraints are added to ensure that direct lightpaths between a pair of non-hub nodes may be established only if the traffic between these nodes exceeds the given threshold θ:

$$b_{sd} = 0, \quad t^{sd} < \theta C, s,d \in \mathcal{N}. \tag{13}$$

Note that the fact that direct lightpaths may be established between non-hub nodes (something that is not allowed under H-VTTR) makes it possible to simplify the formulation by removing and modifying, respectively the above pairs of flow conservation constraints.