

Flow Isolation in Optical Networks

Hui Wang, George N. Rouskas
Operations Research / Department of Computer Science
North Carolina State University
Raleigh, NC 27695-8206 USA
Email: {hwang4,rouskas}@ncsu.edu

Abstract—We address the issue of ensuring the integrity and privacy of communications in optical networks. To this end, we extend the traffic grooming concept to encompass flow isolation considerations through a multi-class traffic model. We develop, solve, and compare ILP formulations that ensure that only traffic components within the same class are groomed onto (share) the same wavelength. Our approach provides flow isolation guarantees with an increase in overall cost as a tradeoff.

I. INTRODUCTION

Innovation and technology advances in computer networking have spawned new and exciting telecommunications applications and services that have transformed every aspect of business and commerce, education and scientific exploration, entertainment and social interaction, and government and defense services. In particular, the Internet has changed the way people and corporations interact, access information, or conduct business, and has evolved into a ubiquitous global communication infrastructure as vital to society as the power grid and the transportation system.

As many critical and sensitive transactions are conducted over the Internet, it is important for network providers to ensure the integrity and privacy of the information they transport on behalf of their customers. For instance, users will be reluctant to entrust their data to cloud computing operators unless the latter provide a solid guarantee that this data, and the resulting computations, will not be intercepted, or tampered with, on their way from the user premises to the cloud servers and back. Similar guarantees are important to transactions involving the health care industry, financial institutions, the defense and intelligence community, and a multitude of business undertakings.

In this work, we address the issue of ensuring the integrity and privacy of communications in optical networks. Our objective is to carefully isolate flows from each other so as to prevent unauthorized eavesdropping and/or tampering of data. We note that (electronically) routed networks suffer from several challenges as electrical signals may be intercepted, routers may be compromised, and even encrypted transmissions are subject to traffic analysis. Optical networks [7]–[9], on the other hand, offer several advantages with respect to data integrity and privacy. Optical signals are difficult to tap, since tampering with an optical fiber will result in severe power loss that is certain to be detected immediately (i.e., within a

few milliseconds) by downstream nodes. Also, switching at optical nodes is transparent, preventing a malicious user from accessing the data passing through even if the switch itself is compromised. Finally, wavelength division multiplexing (WDM) provides inherent traffic isolation as flows traveling on different wavelengths cannot interfere with each other.

The rest of the paper is organized as follows. In Section II, we discuss traffic grooming and define flow isolation as a variant of this well-known problem. Based on this model, in Section III, we develop algorithms to obtain grooming solutions that guarantee flow isolation. In Section IV, we compare these algorithms on a sample network topology, and we conclude the paper in Section V.

II. FLOW ISOLATION AS A TRAFFIC GROOMING PROBLEM

The physical topology of a WDM network consists of a set of nodes interconnected by fiber links. As shown in Figure 1, each node of the network is equipped with both an electronic switch, or digital cross-connect (DXC) device, and an optical cross-connect (OXC) that switches signals optically. A virtual topology is defined on the physical topology when wavelength paths, referred to as *lightpaths*, that may span several fiber links, are created to connect various pairs of nodes. User traffic is routed over the virtual topology, possibly traversing a sequence of lightpaths from source to destination. For instance, the traffic from node 1 to node 4 in Figure 1 is carried over the connection shown as a dashed line. The connection uses two lightpaths: one from node 1 to node 3, and the other one from node 3 to node 4. Lightpath 1 is optically switched at node 2, whereas the traffic is electronically switched between the two lightpaths at node 3.

A. Traffic Grooming

Fiber links currently support multiple wavelengths, each operating at data rates in the order of 2.5-10 Gbps, with rates of 40-100 Gbps likely to be commercially available in the near future. Although certain applications, including some related to large e-Science experiments or peta-scale data mining and analytics, can make efficient use of the large capacity of each wavelength channel, the vast majority of connections has bandwidth requirements that are small compared to the available data rates. Hence, there is a need for packing the sub-wavelength traffic demands into lightpaths for efficient transport across the optical network.

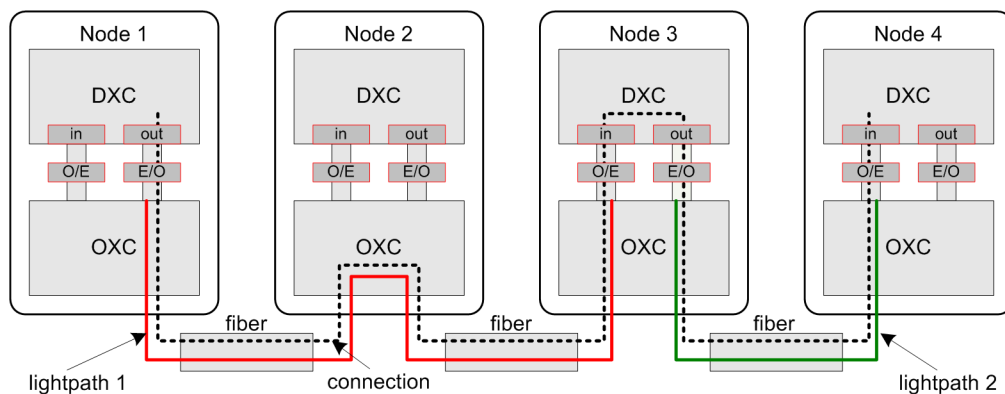


Fig. 1. Illustration of an optical WDM network and grooming

Traffic grooming [2] refers to a class of optimization problems developed to address the gap between the wavelength capacity and the traffic demands of individual connections in optical WDM networks. Returning to Figure 1, terminating lightpath 1 at node 3 makes it possible to groom traffic from node 1 with other traffic also destined to node 4 so as to better utilize the capacity of lightpath 2; this other traffic may originate at node 3, or it may have arrived there from other nodes via lightpaths not shown in the figure. Similarly, lightpath 1 may carry traffic not only to node 4, but to other nodes as well (e.g., node 3) that was appropriately groomed at node 1. As a result of such grooming, the overall utilization of the network may be significantly higher than it would be if each traffic demand were carried on its own dedicated lightpath directly from source to destination.

The grooming problem has been extensively studied in the literature, and basic integer linear programming (ILP) formulations are available in [3], [11]. The objective of most grooming studies has been to either minimize the total network cost (e.g., [1]), or to maximize the total revenue by satisfying as many demands as possible (e.g., [11]). A typical metric for representing the network cost is the number of lightpaths that need to be established to carry the traffic demands [6]. This metric captures both the capital expenses (i.e., the number of transceivers and hence the switch size) and the operating expenses (i.e., the total network capacity) associated with a given demand matrix. For a comprehensive treatment of the traffic grooming problem, the reader is referred to [2].

B. Flow Isolation

A fundamental assumption underlying most traffic grooming studies has been that a lightpath may carry *any* traffic demands, without any constraints on which traffic components can be groomed on a particular lightpath. This is a reasonable assumption if the only objective is to minimize network cost or maximize revenue: more opportunities for grooming traffic lead to better utilization of lightpaths, which in turn results in either lower cost (for a fixed set of demands) or in more accepted connections (for a given level of network resources).

We note, however, that the grooming of traffic entails terminating optical connections at a certain node and electronically

switching traffic from multiple incoming lightpaths to one or more outgoing lightpaths (refer to node 3 in Figure 1). As we discussed earlier, traffic at such electronic switching points may be vulnerable to interception and other malicious attacks. Therefore, users concerned with such attacks may require additional assurances from the network provider that their traffic be appropriately isolated to protect the privacy and integrity of their communications.

To address such concerns, we extend the concept of traffic grooming to encompass *flow isolation* considerations. Specifically, we assume that traffic carried by the network belongs to one of K distinct classes. A traffic class may represent a given level of security requirement as requested by the user; for instance, the network provider may classify traffic into $K = 3$ classes corresponding to low, medium, or high security, respectively. Alternatively, a class may represent a user or group of users; for example, if an organization does not wish its traffic to be groomed with traffic belonging to their competitor, the two organizations will be assigned to different classes.

Given such a classification of user traffic, grooming is permitted if and only if the traffic to be groomed belongs to the *same class*. As a result, each lightpath created in the virtual topology will carry traffic within a single class. Consequently, the virtual topology can be thought of as a collection of K disjoint sub-topologies, where the k -th sub-topology carries only traffic in class k , $k = 1, \dots, K$. This arrangement naturally supports flow isolation: each traffic class is carried optically on its own sub-topology and hence, it does not interfere with traffic on other sub-topologies.

This new variant of the traffic grooming problem balances the objective of minimizing network cost against the requirement to isolate flows. Clearly, the additional constraint of only grooming traffic within the same class is expected to limit the opportunities for grooming and drive the cost higher; we investigate this issue in Section IV. Nevertheless, such a cost increase may be acceptable given that the resulting solution provides higher security guarantees for which the users are likely to be willing to pay a premium.

Although the class-specific grooming constraints guarantee

flow isolation in the sense that each traffic class is carried on a dedicated set of lightpaths, they do not provide any guarantees with respect to isolating the *grooming operations per se*. In other words, it is possible that some node i will be responsible for grooming multiple classes of traffic. Although each traffic class arrives and departs on its own, dedicated, lightpaths, the actual grooming of the various classes takes place within the same DXC (refer to Figure 1) at node i . It is possible to extend the isolation guarantees to the grooming operations, if users deem such guarantees as important and are willing to incur the additional expenses. For instance, the network provider may install multiple distinct DXCs at each grooming node such that each DXC exclusively handles traffic from a single class. Another approach would be to designate certain nodes in the network as grooming nodes for a particular class of traffic, such that class- k traffic is only groomed at class- k grooming nodes. In this case, the network provider may dedicate resources to fortifying grooming nodes against malicious attacks based on the security class of the corresponding traffic. Both these approaches can be modeled by introducing additional constraints in the ILP formulation that we discuss in the following section. However, due to page constraints, we do not consider such extensions further in this paper.

III. GROOMING ALGORITHMS FOR FLOW ISOLATION

We consider an optical network with a physical topology represented as a graph $G = (\mathcal{N}, \mathcal{A})$, where \mathcal{N} is the set of network nodes and \mathcal{A} is the set of physical links (arcs) interconnecting the nodes. We assume that each physical link (arc) is directed and composed of a single fiber supporting W wavelength channels. The nodes are connected to each other by a pair of links in opposite directions. The capacity of each wavelength channel is measured in multiples of a basic traffic rate (e.g., OC-3) and is denoted by integer C . There are K traffic classes such that only traffic demands within the same class are permitted to be groomed onto a given lightpath. The amount of class- k traffic originating at node s and terminating at node d , in terms of the same basic rate, is an integer and represented as t_k^{sd} , $k = 1, \dots, K$. We also let $T_k = [t_k^{sd}]$ denote the network traffic matrix for class k .

Our objective is to determine the virtual topology of lightpaths and the routing of traffic components $\{t_k^{sd}\}$ over these lightpaths, under the class-specific grooming constraints we mentioned above, so as to minimize the number of lightpaths. Note that the special case $K = 1$ corresponds to the classical traffic grooming problem that has been studied in the past. In particular, we presented a link-based [5] ILP formulation of this problem in [10]. In the following subsections, we present three approaches to tackling the flow isolation problem, and describe how the corresponding ILP formulations can be derived from the one in [10].

A. Integrated Formulation

The first formulation considers all traffic classes in an integrated manner, and hence produces a virtual topology that

is jointly optimal for the K classes. This formulation can be derived from the one in [10] through the following two modifications:

- Expand each decision variable in the original formulation to a set of K new decision variables, each corresponding to one of the traffic classes. Consider, for instance, variable b_{ij} in the original formulation, that is defined as the number of lightpaths from node i to node j . This variable is replaced by K variables $b_{ij,k}$, $k = 1, \dots, K$, defined as the number of lightpaths from node i to node j that are used for grooming class- k traffic.
- Expand accordingly the set of constraints (i.e., those related to lightpath routing, wavelength assignment, and traffic routing) in the original formulation in [10] to account for the class-specific decision variables.

Although this formulation provides the jointly optimal solution, its main drawback is lack of scalability. Specifically, the size of the formulation (in terms of number of variables and constraints) is larger by roughly a factor of K compared to the size of the original formulation in [10]. The formulation we present next trades optimality for better scalability.

B. Separate Formulation

As the name implies, in the separate formulation we consider each traffic class in isolation, rather than considering all classes jointly as in the integrated formulation. Let us assume that the traffic classes are labeled $1, 2, \dots, K$, in some arbitrary manner. This approach can be described by the following algorithmic steps:

- 1) Let $k \leftarrow 1$.
- 2) Solve the ILP formulation in [10] with the traffic matrix $T_k = [t_k^{sd}]$.
- 3) Update the number of wavelengths available on each link by subtracting the number of wavelengths used in the solution of Step 2.
- 4) Let $k \leftarrow k + 1$; if $k \leq K$ then repeat from Step 2, otherwise terminate.

During the k -th iteration, this algorithm constructs the virtual sub-topology for class- k traffic. This sub-topology is optimal for the resources (i.e., link wavelengths) left over after accounting for the resources taken up by the sub-topologies of traffic classes $1, \dots, k - 1$.

The main advantage of this approach is that the size of the formulation solved in Step 2 is independent of the number K of traffic classes. As we will see in the following section, solving the original ILP formulation in [10] K times can be substantially more efficient computationally than solving the integrated formulation of the previous subsection.

Importantly, the following lemma states that, if there are sufficient wavelength resources in the network, then the separate formulation is equivalent to the integrated one, i.e., it is jointly optimal for all traffic classes.

Lemma 3.1: Let L^{INT} denote the number of lightpaths obtained as the optimal solution of the integrated formulation. Let L_k^{SEP} , $k = 1, \dots, K$, denote the number of lightpaths in

the optimal solution for class- k traffic obtained through the separate formulation (i.e., when each traffic class is considered in isolation in Step 2 of the above algorithm). Define $L^{SEP} = \sum_{k=1}^K L_k^{SEP}$ as the total number of lightpaths in the virtual topology that represents the union of all the class- k sub-topologies. Suppose that the number W of wavelengths is sufficiently large such that $L_k^{SEP}, k = 1, \dots, K$, is optimal when each link has W wavelengths (rather than just for the case when each link has only a number of wavelengths left over by the first $k - 1$ sub-topologies)¹. Then, we have that:

$$L^{SEP} = L^{INT}. \quad (1)$$

Proof. Let L_k^{INT} denote the number of lightpaths in the integrated formulation solution dedicated to class- k traffic, $k = 1, \dots, K$. Since the sub-topologies dedicated to each traffic class are disjoint by construction, we have that:

$$L^{INT} = \sum_{k=1}^K L_k^{INT}. \quad (2)$$

By assumption, the value of $L_k^{SEP}, k = 1, \dots, K$, is optimal, hence:

$$L_k^{SEP} \leq L_k^{INT} \quad \forall k. \quad (3)$$

From expressions (2)-(3), we obtain:

$$L^{SEP} = \sum_{k=1}^K L_k^{SEP} \leq \sum_{k=1}^K L_k^{INT} = L^{INT}. \quad (4)$$

On the other hand, since L^{INT} represents the jointly optimal solution for all classes of traffic, we have that:

$$L^{INT} \leq L^{SEP}. \quad (5)$$

Finally, from expressions (4) and (5) we have the result in (1), establishing that the two formulations are equivalent under the condition of the lemma. ■

C. Baseline Formulation

In this approach, we simply solve the formulation in [10] with the traffic matrix $T = \sum_k T_k$. In other words, we assume that all traffic offered to the network belongs to the same class, hence there are no class-specific grooming constraints. This baseline formulation provides a lower bound on the number of lightpaths required to accommodate the traffic demands, as well as a reference value in terms of the computational effort (running time) necessary to obtain a solution. Hence, it can be used to characterize the additional resources involved in providing flow isolation guarantees.

IV. NUMERICAL RESULTS

In this section, we investigate the relative performance of the three formulations above (i.e., integrated, separate, and baseline) on the two network topologies shown in Figure 2.

¹Note that this condition is always true for $k = 1$.




Fig. 2. Sample networks used to compare grooming methods: (a) a 6-node topology, (b) the 11-node Cost239 topology

The first is a small 6-node, 8-link topology, and the second is the larger 11-node, 26-link Cost 239 topology. Each link in the two networks is assumed to be bidirectional, consisting of a single fiber in each direction. A fiber has $W = 3$ wavelength channels, each of capacity $C = 48$. We assume that there are $K = 3$ traffic classes. The amount t_k^{sd} of class- k traffic, $k = 1, 2, 3$, originating at node s and terminating at node d is a random integer uniformly distributed in $[0, t_{max}]$. We vary the value of parameter t_{max} to obtain results for different traffic loads.

It is well known that the grooming problem is NP-hard [4], [11], hence ILP formulations are only solvable for small networks. Therefore, we imposed a limit of 24 hours for CPLEX to find a solution; if it failed to do so within the 24-hour limit, we recorded the best (not necessarily optimal) solution that CPLEX had found at that time, as well as the optimality gap. All the results were obtained using the CPLEX 11 software on a cluster of compute nodes with dual Woodcrest Xeon processors running at 2.33GHz with 1333MHz bus, 4GB of memory and 4MB L2 cache.

A. Results for the 6-node topology

Figure 3 plots the number of lightpaths used by each grooming method as t_{max} increases from 2 to 16, for the 6-node topology shown in Figure 2(a). As Table I shows, for this small network CPLEX was able to obtain the optimal solution for all instances when using the baseline and separate methods;




Fig. 3. Number of lightpaths used by each grooming method for the 6-node topology of Figure 2(a)

TABLE I
AVERAGE CPLEX OPTIMALITY GAP OF EACH GROOMING METHOD FOR THE 6-NODE TOPOLOGY OF FIGURE 2(A)

t_{max}	Baseline	Integrated	Separate
2	0	0	0
4	0	0	0
6	0	0.87%	0
8	0	9.01%	0
10	0	12.36%	0
12	0	13.50%	0
14	0	13.27%	0
16	0	13.15%	0

however, under the integrated method, problem instances with $t_{max} \geq 6$ could not be solved to optimality within the 24-hour time limit, resulting in small to moderate optimality gaps as the table indicates. We emphasize that each data point in Figure 3 and Table I is the average of 30 problem instances randomly generated for the stated value of t_{max} .

As expected, as the total amount of traffic (i.e., the value of t_{max}) increases, the number of lightpaths required to carry the traffic increases accordingly for all three methods. Also, the baseline formulation uses fewer lightpaths than the other two formulations across the entire range of traffic loads; as we explained earlier, the baseline formulation views all traffic as belonging to the same class, hence it has more opportunities for traffic grooming and leads to better utilization and fewer lightpaths.

Let us now compare the integrated and separate formulations. For $t_{max} = 2, 4, 6$, and 8 , the two formulations produce the same results, i.e., the same number of lightpaths. Since the traffic load is relatively low for these values of t_{max} , it is likely that the conditions of Lemma 3.1 hold true in this case, hence the two formulations are equivalent. Interestingly, however, for larger values of t_{max} , the separate formulation produces *better* solutions than the integrated one. This is counter to intuition, as the integrated formulation is the one that is jointly optimal for all traffic classes. Nevertheless, this result can be explained

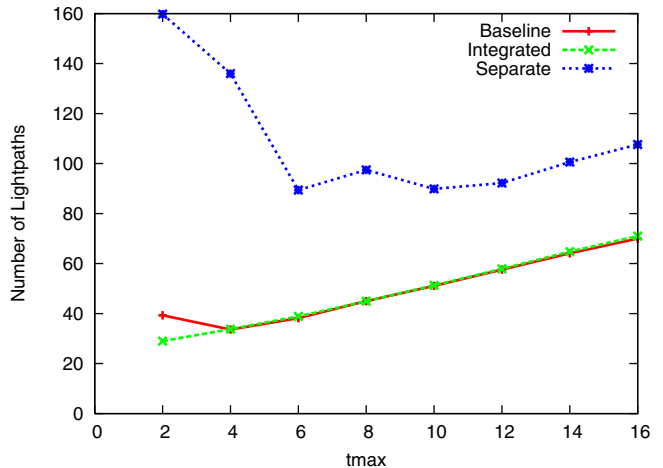


Fig. 4. Number of lightpaths used by each grooming method for the 11-node topology of Figure 2(b)

TABLE II
AVERAGE CPLEX OPTIMALITY GAP OF EACH GROOMING METHOD FOR THE 11-NODE TOPOLOGY OF FIGURE 2(B)

t_{max}	Baseline	Integrated	Separate
2	56.89%	53.19%	88.25%
4	26.73%	27.05%	72.26%
6	13.03%	14.69%	48.16%
8	7.61%	7.60%	42.73%
10	4.72%	5.16%	31.43%
12	3.30%	4.11%	24.02%
14	1.82%	3.66%	20.85%
16	1.40%	3.56%	18.24%

by referring to Table I which indicates that instances of the integrated formulation have not been solved to optimality for $t_{max} \geq 6$. The size of the integrated formulation is larger by roughly a factor of K than that of the baseline one. As a result CPLEX does not reach the optimal solution within the 24-hour limit, hence the lightpath values reported in Figure 3 are *not* optimal, explaining why the separate formulation shows lower lightpath values.

B. Results for the 11-node Cost 239 topology

Let us now turn our attention to Figure 4 which is similar to Figure 3, but presents results for the 11-node Cost 239 topology shown in Figure 2(b), and to Table II which lists the corresponding average optimality gaps. Each data point in Figure 4 and Table II is the average of 10 problem instances randomly generated for the stated value of t_{max} .

We immediately observe that there are similarities but also significant differences between the relative behavior of the curves in Figures 4 (the 11-node network) and 3 (the smaller 6-node network). To explain these differences, let us first discuss the average optimality gaps in Table II. As we can see, with the larger network, CPLEX cannot find the optimal solution within the 24-hour limit we imposed on the running time. Also, the optimality gap depends on both the solution method and the amount of traffic. Specifically, the optimality gaps for the

baseline and integrated methods are similar, while the separate method has a much larger optimality gap for the same t_{max} value. This result can be explained by the fact that, with the separate method, $K = 3$ different formulations are solved separately (refer to Section III-B). For a fair comparison, we let each of these three different formulations run for at most eight hours (for a total of 24 hours, as with the other two solution methods). Consequently, CPLEX stops when the optimality gap is higher than if it had run for 24 hours for each formulation.

The optimality gap also decreases as the amount of traffic (which is directly tied to the value of t_{max}) increases, and this behavior is consistent across the three solution methods. Note that, in a large network, when traffic demands are low, the solver may take a long time to evaluate all the traffic grooming possibilities so as to minimize the number of lightpaths. As traffic increases, however, the number of ways that the various demands can be groomed together to reduce the number of lightpaths decreases, hence the solver needs less time to evaluate them, resulting in smaller optimality gaps.

Returning to Figure 4, we note that as the traffic demand increases, both the integrated and baseline methods require more lightpaths, as expected. The integrated method also uses only slightly more lightpaths than the baseline method, except for $t_{max} = 2$, when it requires significantly fewer lightpaths (29 versus 39.3, on average across the 10 problem instances). We believe that this result can be explained as follows. On one hand, the high optimality gaps for lower values of t_{max} indicate that the values plotted in Figure 4 are not optimal. Hence, the results for low traffic demands can be skewed in favor of the integrated method. On the other hand, for large t_{max} values when the values plotted are closer to optimal, there is sufficient traffic within each traffic class so that grooming solutions obtained with the integrated formulation are comparable to ones obtained with the baseline formulation.

Finally, the counter-intuitive behavior of the curve in Figure 4 corresponding to the separate formulation can also be explained by referring to the corresponding optimality gaps. Specifically, the high lightpath values for $t_{max} \leq 8$ are due to the fact that they are far from optimal; we expect that, given sufficient time to run, CPLEX would have found much better solutions. This conjecture is supported by the fact that, for higher t_{max} values CPLEX is able to find better solutions and the curve starts trending upwards, as one would expect given the higher traffic that needs to be carried. Even in this case, however, the difference between the solutions of the separate formulation and the solutions of the baseline and integrated formulations cannot be explained by the optimality gap alone; in other words, the latter two methods produce intrinsically better results. Also note that, we have assumed a small number $W = 3$ of wavelengths, therefore it is likely that the conditions of Lemma 3.1 do not hold true, and the optimal solution of the integrated method is strictly better than that of the separate method.

Overall, the results presented in this section indicate that this variant of traffic grooming is effective in capturing flow

isolation considerations. Furthermore, the additional penalty for providing isolation guarantees is relatively small, and may become negligible for larger topologies (Figure 4). The separate formulation achieves a good trade-off between optimality and scalability in small networks: its running time is comparable to that of the baseline formulation that does not take flow isolation into account, and the number of lightpaths it requires are close to that of the optimal integrated formulation. However, for medium size networks, the integrated formulation is computationally more efficient, and is able to find solutions close to those obtained by the baseline formulation. We expect that these observations remain true for larger networks, making the integrated formulation the clear choice whenever the problem cannot be solved to optimality.

V. CONCLUSION

Ensuring the integrity and privacy of communications in optical networks is likely to become an important requirement as critical and sensitive transactions move to the cloud. In this paper, we explored a new variant of the well-known traffic grooming problem as a means to providing flow isolation guarantees. Using a multi-class traffic model, we developed two grooming solutions to guarantee that only traffic components within a given class are permitted to share a particular lightpath. The results obtained on two sample network topologies suggest that it is possible to achieve a good balance between scalability and optimality, and that the amount of additional resources required to ensure flow isolation is sufficiently low to make this a practical approach.

REFERENCES

- [1] A. L. Chiu and E. Modiano. Traffic grooming algorithms for reducing electronic multiplexing costs in WDM ring networks. *IEEE/OSA Journal of Lightwave Technology*, 18(1):2–12, January 2000.
- [2] R. Dutta, A. E. Kamal, and G. N. Rouskas (Eds.). *Traffic Grooming for Optical Networks: Foundations, Techniques and Frontiers*. Springer, New York, 2008.
- [3] R. Dutta and G. N. Rouskas. Traffic grooming in WDM networks: Past and future. *IEEE Network*, 16(6):46–56, November/December 2002.
- [4] S. Huang, R. Dutta, and G. N. Rouskas. Traffic grooming in path, star, and tree networks: Complexity, bounds, and algorithms. *IEEE Journal on Selected Areas in Communications*, 24(4):66–82, April 2006.
- [5] B. Jaumard, C. Meyer, B. Thiongane, and X. Yu. ILP formulations and optimal solutions for the RWA problem. In *Proceedings of IEEE GLOBECOM'04*, volume 3, pages 1918–1924, November 29–December 3 2004.
- [6] 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.
- [7] B. Mukherjee. *Optical WDM Networks*. Springer, 2006.
- [8] R. Ramaswami and K. N. Sivarajan. *Optical Networks: A Practical Perspective, 2nd Edition*. Morgan Kaufmann Publishers, San Francisco, California, 2002.
- [9] T. E. Stern, G. Ellinas, and K. Bala. *Multiwavelength Optical Networks: Architectures, Design, and Control*. Cambridge University Press, 2009.
- [10] E. Yetginer and G. N. Rouskas. Power efficient traffic grooming in optical WDM networks. In *Proceedings of IEEE GLOBECOM 2009*, December 2009.
- [11] 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.