Many-to-one Traffic Grooming with Aggregation in WDM Networks *

Raza Ul-Mustafa[†] Ahmed E. Kamal[‡]

Department of Electrical and Computer Engineering
Iowa State University
Ames, IA 50011
U.S.A.
{raza,kamal}@iastate.edu

Abstract

Most of the network applications bandwidth requirements are far less than the bandwidth offered by a full wavelength in WDM networks. Hence, traffic grooming is needed to make efficient use of the available resources. In this paper we address the grooming of many-to-one traffic demands in WDM networks on arbitrary topologies. Traffic streams from different sources, but part of the same session and thus terminating at the same destination, can be aggregated using arbitrary, but application dependent, aggregation ratios. We provide optimal as well as heuristic solutions to the problem. The objective is to minimize the cost of the network, by minimizing the total number of the higher layer components and the total number of the wavelengths used in the network. One of the main contributions of this work is to provide a mixed integer linear solution, to an otherwise non-linear problem, by exploiting the specifics of routing and aggregation sub-problems, while still maintaining the optimality of the solution. The formulation is generic and can handle varying amounts of traffic from each source to a common destination, as well as arbitrary aggregation fractions of the data coming from the different sources. This fraction is made to be a function of the number of the streams participating in the aggregation. For the heuristic solution we developed a Dynamic Programming style approach that builds the solution progressively, going through a number of stages, while choosing the best partial solutions among a number of possible partial solutions at each stage.

Keywords: Wavelength Division Multiplexing (WDM) networks; traffic grooming; Line Terminating (LT) equipment; IP routers, many-to-one traffic; data aggregation; Mixed Integer Linear Programming (MILP);

*This work is supported in part by grant ANI-0087746 from the National Science Foundation.

[†]R. Ul-Mustafa was with Iowa State University. He is now with Microsfot Corporation, One Microsoft Way, Redmond, WA 98052-6399, e-mail: razaulm@microsoft.com.

[‡]Corresponding author: e-mail: kamal@iastate.edu; telephone: (515) 294-3580; fax: (515) 294-1152.

I Introduction

Currently, Wavelength Division Multiplexing (WDM) networks are the most prevalent mode of utilizing the enormous amounts of bandwidth offered by the wide spread deployment of optical fibers. However, the bandwidth offered by a single wavelength in WDM networks is of the order of tens of Gigabits per second, while most of the applications' bandwidth requirements are still subwavelength. For example, HDTV can work well with just 20 Mbps, while a normal TV channel typically requires less than 2 Mbps, when compressed using MPEG-2. This disparity between the capacity of a single wavelength and the applications' bandwidth requirement has recently brought the field of traffic grooming in WDM networks to the spotlight. In traffic grooming, subwavelength traffic streams are groomed onto available wavelengths. The grooming can be implemented to achieve different objectives, including network cost minimization. The dominant cost factor in WDM networks is considered to be the higher layer network components [6], such as SONET Add/Drop multiplexers (ADMs), IP or MPLS router ports. In this paper, we will refer to such equipment as line terminating equipment (LT).

Because of the hard nature of the traffic grooming problem [6], many simplifying assumptions have been used in the literature. Some of the studies have allowed the traffic between each source-destination pair to be (vertically) split over multiple wavelengths - a condition known as bifurcation[1]. For example, in bifurcation, an OC-48 traffic demand between a source-destination pair may be split into four OC-12 (or even lower granularity) demands. Now each component of the original traffic demand may be independently routed over different physical links and/or lightpaths. This provision provides flexibility in traffic allocation, which may lead to a reduction in the number of the wavelengths as well as the number of the LTs. However, bifurcation increases the complexity and the cost of traffic reassembly, and may also introduce delay jitter at the application layer. Our intention in this paper is to work with more realistic models, and hence we do not allow traffic bifurcation.

Most of the current studies in traffic grooming have dealt exclusively with unicast traffic. However, as WDM networks are extending their presence from merely in backbone networks to metro and local area networks, there is a need to accommodate other types of traffic efficiently. In the foreseeable future, a sizable portion of the traffic is expected to be multipoint. Multipoint traffic can take different forms, e.g., one-to-many and many-to-one. One-to-many, or multicast applica-

[1]We use the term *vertical* splitting of traffic to mean bifurcation of traffic, or splitting the traffic over multiple lightpaths starting from the same node. We also use the term *horizontal* splitting to mean carrying the traffic over multiple virtual hops, or lightpaths.

tions, include multi-party conferencing, video distribution, network news distribution, collaborative processing, and web content distribution to proxies [8]. The many-to-one service model is used by applications such as resource discovery, data collection, auctions, polling and audience to speaker communication. Although both traffic types have the multipoint factor in common, each type has different characteristics and hence demands different design and operational approaches.

Many-to-one traffic propagates from a set of sources to a particular destination. Depending on the application, each source may generate a different amount of data. There could be some overlap or redundancy between the data streams heading towards the destination. Examples of scenarios in which aggregation may be performed include:

- Aggregation of video streams from monitoring cameras feeding a control center. When the number of streams changes, the resolution of each of the streams needs to be adjusted, especially if all streams are to to displayed on one monitor. This can be implemented using multilayer coding of the video streams, where the data corresponding to the chosen layers are forwarded, and the rest are blocked [9]. In this case, the blocked traffic is a fraction of the original stream.

- In a multiparty video conference, a combination of a many-to-one and one-to-many service modes can be used to provide the required many-to-many service. Due to the ubiquitous single speaker mode, video and audio streams may be aggregated on the many-to-one tree such that only the active stream is delivered, while other streams are blocked. Other aggregation strategies are also possible, in which the active stream is delivered in full resolution, while the incative streams are delivered with a reduced resolution. In the first aggregation strategy, if the number of users is n, then $1/n$th of the traffic is delivered, while the rest is blocked. In the second strategy, the delivered traffic is also a fraction of the total traffic.

- Backup of a replicated database, where different copies of the same data are transmitted to the storage center. Only one copy needs to be delivered, and the rest can be blocked.

To efficiently use the resources, such data streams might be aggregated on their way to the destination. Hence a many-to-one routing tree, sourced at different nodes and terminating at a common destination can be created such that data streams are aggregated at each merging point of the many-to-one tree[2]. Compared to a typical multicast tree where a single stream of data coming

[2] As mentioned above, many-to-many service can be provided using a many-to-one tree and a one-to-many tree. The destination of the many-to-one tree acts as the source of the one-to-many tree.

from a source is duplicated onto multiple streams at branching points, the many-to-one tree merges and may aggregate the data coming from different sources at an aggregation point. This changes the bandwidth requirements and equipment functionalities for a many-to-one tree. Hence, the algorithms developed for multicast trees are not very useful for designing the many-to-one trees.

In this paper, we address the many-to-one traffic grooming problem such that given a traffic matrix all the traffic should be accommodated with the least number of required higher layer light terminating equipment, while also trying to minimize the total number of used wavelengths. We assume that aggregation, when performed, will be implemented at layer 1, e.g., by delineating and flagging data blocks which may be dropped during aggregation. This can be implemented using protocols in Next Generation SONET (NGS) [10], e.g., using extension headers [11]. To the best of our knowledge this problem has not been addressed before in the literature. We illustrate the concept of many-to-one traffic grooming with the help of a simple example. Figure 1 shows a network consisting of 7 nodes. Suppose that there is only one session which consists of 4 source nodes, S_1, S_2, S_3 and S_4, located at node 1, 3, 4 and 6, respectively. The destination of this session, d, is located at node 7. Let the capacity of a single wavelength be 48 units (e.g., OC-48) and the traffic originating from each source be 24 units. We refer the traffic originating from each source as a stream. Without traffic grooming, the session $\{S_1, S_2, S_3, S_4 \rightarrow d\}$ must be accommodated using 3 wavelengths and 8 LTs (2 LT for each source-destination pair). However, with traffic grooming the number of LTs and wavelengths can be reduced. The streams from different sources destined to same destination, d, may be aggregated on their way. If the aggregation ratio is 1[3], i.e., no data is removed from any participating stream at an aggregation point then the session $\{S_1, S_2, S_3, S_4 \rightarrow d\}$ can be accommodated as shown in Figure 1.a. Streams from S_1 and S_2 are aggregated at node 3. This accommodation requires only 2 wavelengths and 7 LTs. However, if at an aggregation point the data of each stream is reduced by a fixed ratio of 0.5, then session $\{S_1, S_2, S_3, S_4 \rightarrow d\}$ can be accommodated as shown in Figure 1.b. Streams from S_1 and S_2 are first aggregated at node 3 and then streams from S_1, S_2 and S_3 are aggregated at node 4. This accommodation requires only 1 wavelength and 6 LTs. Therefore, many-to-one traffic grooming with aggregation can result in a substantial reduction of the cost of the network in terms of the number of LTs and the number of wavelengths. Although in Figure 1.b we assumed a fixed aggregation ratio of 0.5 irrespective of the number of streams aggregated at a point, in reality the aggregation ratios at each aggregation point depend on the number of sources (of the same many-to-one session) participating at that

[3]The aggregation ratio will be defined formally below. However, for the purpose of this example, the aggregation ratio refers to the proportion of the aggregated traffic that will be delivered.

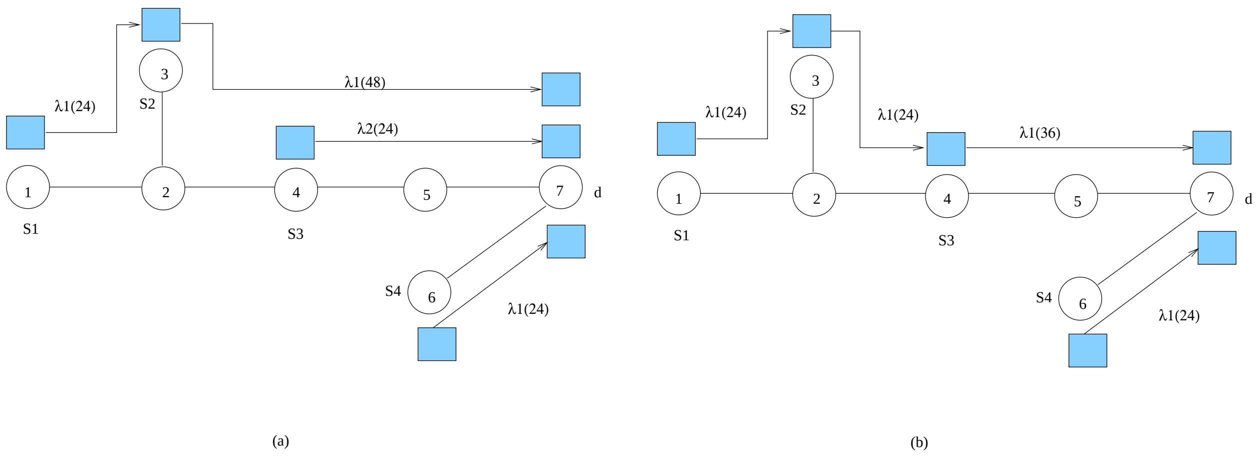

Figure 1: An example of many-to-one traffic grooming: (a) without aggregation and traffic reduction; (b) with aggregation and traffic reduction.

specific aggregation point. Our model takes such factor into consideration, and we develop both optimal as well as heuristic solutions to the routing and provisioning problems of many-to-one traffic grooming with traffic aggregation. One of the main contributions of the work is to provide a Mixed Integer Linear Programming (MILP) solution, to an otherwise non-linear problem. We achieve this task by exploiting the specifics of routing and aggregation sub-problems, while still maintaining the optimality of the solution. The formulation of the MILP is generic and can handle varying amounts of traffic from each source to a common destination, as well as a varying number of independent sessions. Also, at each merging point in a many-to-one tree, an arbitrary fraction of the data coming from each source can be selected. Thus, if the data streams of a source set are highly correlated, then, at a merging point, more data from each stream can be dropped by assigning a lower value to an aggregation factor. Moreover, to capture a more realistic model we successfully make this fraction a function of the number of the streams participating in the aggregation. Thus, the fractional factor for each stream decreases when the number of the streams participating in the aggregation increases, reflecting more opportunity to remove the redundant data. For the heuristic solution we develop a Dynamic Programming style approach that builds the solution going through a number of stages, while choosing the best partial solutions among a number of possible partial solutions at each stage. In fact, instead of selecting just one partial solution, the algorithm maintains a number of best partial solutions at each stage. Any stage, however, always depends on the previous stage only. Finally, in the final stage the solution with the least cost is selected.

We summarize the contributions of this work as follows.

- We provide optimal as well as heuristic solutions to the design of optical networks supporting many-to-one traffic grooming.

- The models work for random topologies, and maintain the non-bifurcation of traffic demands.
- The models provide the flexibility in aggregating the traffic streams with arbitrary ratios, such that the ratios are a function of the number of the traffic streams participating in the aggregation.

It should pointed out that many-to-one, together with multicast traffic grooming serve as low overhead, reduced processing and cost effective approaches for facilitating a Layer 1 virtual private network (VPN), which provide any-to-any communication support. It is well known that any-to-any communication on a virtual private network is difficult to support at higher layers, especially with a Layer 3 VPN [12]. The proposed approaches for traffic grooming provide this functionality easily. Separation between different VPNs is natural under multipoint traffic grooming, which simplifies routing, facilitates secure communication, and provides bandwidth guarantees seamlessly. Moreover, layer 1 VPNs have advantages over Layer 2 and 3 VPNs, especially from the providers point of view in aspects like traffic engineering, and multiprotocol, and multi-customer support [12].

The rest of the paper is organized as follows. In Section II, we briefly review the work done in the area of traffic grooming in WDM networks. In Section III, we present our model and the MILP formulation for the many-to-one traffic grooming network design problem. In Section IV, we present our heuristic approach. In Section V, we present some experimental results, while Section VI concludes the paper.

II Related Work

The traffic grooming problem with arbitrary traffic, even for simple topologies, has been proven to be NP-Complete [1]. To the best of our knowledge the earliest papers that explicitly dealt with the topic appeared in 1998 [2, 4, 3, 5]. By now, few nice survey papers have appeared in the literature [13, 14, 15], which cover most of the related work in the field of traffic grooming in WDM networks.

At a broader level, related work can be categorized based on the traffic patterns, for example, uniform traffic [1, 19, 16] and arbitrary traffic [17, 20, 6, 18], network topologies, for example, unidirectional rings [1, 6], bidirectional rings [16, 20, 17, 18], and random topology [24, 7, 21, 22, 23], and solution approaches, for example, heuristics [1, 17, 18], ILPs [24, 7, 6, 22], and bounds [1, 19, 16, 20]. Most of the research in the area focuses on reducing the number of required higher layer components [14]. However, there are some notable exceptions in which factors like blocking probability and network utilization are considered, e.g., [25, 26].

Multipoint traffic grooming is a recent field. To the best of our knowledge this paper is the first to address the many-to-one traffic grooming problem. However, the problem of routing many-to-one trees and the identification of the merging points of flows has been treated in other contexts, most notably within the domain of Multiprotocol Label Switching (MPLS). For example, in [27] the optimal design of many-to-one label switched paths (LSP) which have inverse tree forms in MPLS neteworks in order to accommodate a set of given routes from ingress to egress points was formulated as an integer linear program with the objective of minimizing the number of many-to-one trees. This has the effect of minimizing the number of MPLS labels. A followup problem which minimizes link capacities given the trees obtained from the above optimizaiton problem, and sets of working and backup routes was formulated as a mixed integer linear program. Although this approach formulates two optimal design approaches, the final solution is not optimal because of two reasons. First, the formulation of the trees is not included in the ILP. A set of candidate trees for each egress node is first constructed, and the ILP chooses from among those trees. Second, the capacity planning phase is disjoint from the tree design phase. In reference [28] the two phases were combined in a mixed integer program. However, a set of feasible many-to-one trees were also precomputed, and the optimizatin program also chooses from among those trees. This paper also introduced a heuristic solution approach. Reference [29] introduced algorithms for the creation of many-to-one LSPs, given a set of point-to-point LSPs. A distinguishing operational feature that we consider in this study is the provision for aggregating traffic from sources within the same session, and reducing the amount of traffic due to this aggregation, while guaranteeing that aggregated streams do not split after combining. This mode of operation has not been considered within the MPLS domain. Moreover, the provisioning phase and the tree design phase are treated as one joint problem, while guaranteeing non-bifurcation.

We should point out that some related work in the area of multicast traffic grooming has also appeared recently, for example, [7, 31, 32]. Those studies provided optimal and heuristic solutions for traffic grooming on mesh and ring topologies.

III Problem Formulation

In this section we will explain the model used in the paper and formulate an MILP for the many-to-one traffic grooming problem on arbitrary network topologies.

III.1 The Model

We allow each many-to-one session to traverse multiple lightpath hops, while each lightpath hop itself may span multiple physical links. Hence, we visualize the network at three different levels: the physical level, the lightpath level and the session level. The physical level represents the topology made of physical links between nodes, and is an input parameter. In our model each physical link represents two fibers that are used to communicate in opposite directions. Also each fiber can support W wavelengths in each direction. The lightpath level represents the virtual topology made of all-optical lightpaths. Each lightpath can span several physical links, and is an output from the MILP. Also, between any node pair multiple lightpaths are allowed to exist, but if they use the same wavelength they must be routed on different physical routes. Session level virtual links between nodes represent the traffic demand between such nodes. Each such virtual link is allowed to span multiple lightpaths.

The bandwidth requirements of a many-to-one routing tree is very different from that of a typical multicast tree. A set of sources transmit data to a common destination over a many-to-one routing tree, while aggregating the data streams at each merging point on the tree. Redundant data among aggregated data streams may get removed at the merging points. In our model, we consider the case in which aggregation is implemented in the electronic domain, and at an LT. We will then minimize the number of the LTs, hence leading to a minimal network cost. We consider a generic model that allows the selection of a fraction of the data of each stream at each aggregation point[4]. Moreover, this fraction can be made a function of the number of the streams participating in the aggregation. We provide the flexibility of choosing any arbitrary set of values for the aggregation ratios, as compared to fixed values. This, however, makes the problem much more difficult. For each many-to-one session, at each aggregation point, MILP has to determine the total capacity after aggregation. This involves selection of an appropriate aggregation ratio, which itself depends on the number of sources of the many-to-one session participating in the aggregation at that specific aggregation point. We will formulate the problem in a linear fashion while still maintaining the optimality of the solution.

Regarding notation, we will use s and d to represent the source and the destination of a session, i and j to represent the source and destination nodes of a lightpath, m and n to represent the source and destination nodes of a physical link, respectively. Each session consists of a number of sources and a destination. There could exist multiple sessions with the same source set and the same

[4]Since aggregation is done at merging points, we will use aggregation and merging points interchangeably.

destination. Therefore, to differentiate between such multiple sessions, the kth session is referred to by c_k, where $1 \leq k \leq K$, where K be the total number of sessions in the system. Moreover, let the source set of session c_k be S_{c_k}, where s represents *a* source in the source set, S_{c_k}. Finally, let $\phi_{s,d,k}$ represents a *traffic request* or *stream* originating at s, destined to d, and belonging to session k. A session c_k will then consists of a set of streams $\phi_{s,d,k}, \forall s \in S_{c_k}$. The rest of the notations used in the paper are defined below.

- *Input parameters:*

  N : total number of nodes in the network

  W : number of wavelengths per fiber [5]

  g : Capacity of a wavelength in terms of the number of basic units of traffic (also called grooming factor)

  $r_f^{c_a}$: aggregation ratio; the fraction of the capacity of each stream from $s \in S_{c_a}$ selected for aggregation, when the number of the sources participating in aggregation is f; $r_1^{c_a} = 1, \forall c_a$ [6]

  $m_{c_a,s}$: number of basic units of traffic originating from source s of session c_a

  P_{mn} : number of physical fiber links (1 or 0) connecting nodes m and n.

  α : cost of an LT

  β : cost of a wavelength, $\beta < \alpha$

  Q: A very large integer, $Q \geq \max(\max_{C_a}(\sum_{s \in S_{c_a}} m_{c_a,s}), \max_i degree_i)$, where $degree_i$ is the degree of node i

- *Variables of the MILP:*

  LT_n : number of LTs at node n

  L_{ij}^w : number of lightpaths from node i to node j on wavelength w

  L_{ij} : number of lightpaths from node i to node j on all the wavelengths, $L_{ij} = \sum_w L_{ij}^w$

  ψ : highest index of wavelength used over any fiber link in the network

  y_w : a binary indicator; must be 1 if there is at least one lightpath in the network on wavelength w

  $F_{mn}^{ij,w}$: a binary indicator; is 1 if and only if there is a lightpath between node pair (i, j) routed through fiber (m, n) on wavelength w

[5] Although we start with W wavelengths, the MILP minimizes the total number of required wavelengths.

[6] The aggregation ratio applies to the total number of aggregated streams, even if they were aggregated earlier at an upstream aggregation point. That is, in the example of Figure 1.(b), at node 3 the aggregation ratio $r_2^{c_a}$ is applied, while at node 4, the ratio $r_3^{c_a}$ is applied.

$Z_{ij}^{c_a,s}$: a real number between 0 and 1, which takes non zero values if and only if the traffic stream from source $s \in S_{c_a}$, is using a lightpath from i to j

$M_{ij}^{c_a}$: a binary indicator; is 1 if and only if at least one of the sources of session c_a is using a lightpath from node i to j to reach the destination, i.e, $\exists s \in S_{c_a}$, such that $Z_{ij}^{c_a,s} = 1$

$I_{ij}^{c_a,f}$: a binary indicator; is 1 if and only if the number of sources $s \in S_{c_a}$, on a lightpath from node i to j, is $\geq f$

$X_{ij}^{c_a}$: a real number which represents the amount of traffic on a lightpath from node i to j due to all sources in S_{c_a}

$J_{ij}^{c_a,c_b}$: a binary indicator; is 1 if and only if session c_a and c_b are groomed on the same lightpath from i to j

$Y_{ij}^{c_a,c_b}$: a real number and is a product of $J_{ij}^{c_a,c_b}$ and $X_{ij}^{c_b}$

III.2 The MILP Formulation

In this subsection we present the MILP for the many-to-one grooming problem. Our objective is to minimize the total cost of the higher layer equipment by minimizing the total number of LTs, as well the maximum number of different wavelengths on a fiber.

Objective function:

$$Minimize : \alpha * \sum_{n} LT_n + \beta * \psi \tag{1}$$

The above function is a multi-objective function, in which LTs are given higher priority for minimization than the wavelengths.

Subject to:

- *Number of LTs:*

  The following two constraints ensure that for each originating (terminating) lightpath from (at) a node, an LT is present

$$LT_i \geq \sum_{w} \sum_{j,j \neq i} L_{ij}^{w} \quad \forall i \tag{2}$$

$$LT_i \geq \sum_{w} \sum_{j,j \neq i} L_{ji}^{w} \quad \forall i \tag{3}$$

- *Number of wavelengths:*

  The following constraints compute the variable ψ such that it is the index of the highest numbered wavelength used on any fiber link in the network. Minimizing ψ will then minimize the total number of different wavelengths in the network and not just the total number of the wavelengths between each node pair

$$\psi \geq w * y_w \quad \forall w \tag{4}$$

$$y_w \geq \sum_i \sum_{j,j\neq i} L_{ij}^w / Q \quad \forall w \tag{5}$$

$$y_w \leq \sum_i \sum_{j,j\neq i} L_{ij}^w \quad \forall w \tag{6}$$

- *Lightpath level constraints:*

  The following constraint ensures that the origin node, i, and the terminating node, j, of lightpath(s) between node i and j have no incoming and outgoing traffic carried on such lightpaths, respectively

$$\sum_{m, forP_{mi}=1} F_{mi}^{ij,w} = \sum_{n, forP_{jn}=1} F_{jn}^{ij,w} = 0 \quad \forall i, j, w \tag{7}$$

  The following constraint determines the total number of the lightpaths on wavelength w between node i and j, supported by the underlying physical topology.

$$\sum_{m, forP_{mj}=1} F_{mj}^{ij,w} = \sum_{n, forP_{in}=1} F_{in}^{ij,w} = L_{ij}^w \quad \forall i, j, w \tag{8}$$

  The following constraint preserves wavelength continuity of lightpaths over multiple physical links

$$\sum_{m, forP_{mx}=1} F_{mx}^{ij,w} = \sum_{n, forP_{xn}=1} F_{xn}^{ij,w} \quad \forall w, i, j, x; x \neq i, j \tag{9}$$

  Equations (8) and (9) together ensure that for each lightpath there exists a corresponding physical path, while maintaining wavelength continuity over the physical path. Equation (10) then ensures that on a wavelength, w, on fiber from node m to node n, only one lightpath is present.

$$\sum_i \sum_j F_{mn}^{ij,w} \leq 1 \quad \forall w, m, n, whereP_{mn} = 1 \tag{10}$$

- *Session topology constraints:*

  The following constraint ensures that for the traffic request sourced at s and part of the many-to-one session c_a, no traffic is coming in (going out) the source (destination), respectively

$$\sum_i Z_{is}^{c_a,s} = \sum_j Z_{dj}^{c_a,s} = 0 \quad \forall c_a, s \in S_{c_a} \tag{11}$$

The following constraint ensures that the traffic request between s and d, which is part of the many-to-one session c_a, is originating (terminating) at s (d), respectively

$$\sum_{j,j\neq s} Z_{sj}^{c_a,s} = \sum_{i,i\neq d} Z_{id}^{c_a,s} = 1 \quad \forall c_a, s \in S_{c_a} \tag{12}$$

The following constraint preserves the continuity of session traffic on multiple lightpaths

$$\sum_{i,i\neq x,i\neq d} Z_{ix}^{c_a,s} = \sum_{j,j\neq x,j\neq s} Z_{xj}^{c_a,s} \quad \forall c_a, s \in S_{c_a}, x, (x \neq s, d) \tag{13}$$

Once data streams, from more than one source, are aggregated at some node, these should remain intact for the rest of their travel to the destination, i.e., these should follow the very same physical path for the rest of the journey. The following set of constraints ensures that the aggregated streams will remain intact at the lightpath level. As we are ensuring non-bifurcation too (using a separate set of constraints to be described later), these data streams will remain intact at the physical level too. The following set of constraints essentially also determine the routing of the many-to-one sessions.

$$M_{ij}^{c_a} \geq \sum_{s\in S_{c_a}} Z_{ij}^{c_a,s}/Q \quad \forall c_a, i, j \tag{14}$$

$$M_{ij}^{c_a} \leq \sum_{s\in S_{c_a}} Z_{ij}^{c_a,s} \quad \forall c_a, i, j \tag{15}$$

$$\sum_{j} M_{ij}^{c_a} \leq 1 \quad \forall c_a, i \tag{16}$$

Although the variable $Z_{ij}^{c_a,s}$ is assumed to be a real number between 0 and 1, equations (12), (14) and (15) will guarantee that it must either 0 or 1. If the variable $Z_{ij}^{c_a,s}$ was to assume a real number other than 0 or 1, then by equations (14) and (15), the variable $M_{ij}^{c_a}$ cannot be either 0 or 1, which contradicts its definition as a binary variable. Note that by reducing the number of integer variables, the number of branching variables is reduced, which helps in reducing the complexity of the problem.

The next task is to determine the amount of traffic before and after aggregation on each of the lightpath. For this purpose, we first use the following constraints to set the binary indicator $I_{ij}^{c_a,f}$ on the lighpaths from i to j to one if and only if the number of sources $s \in S_{c_a}$ is greater than or equal to f, i.e., $I_{ij}^{c_a,f}$ decreases from 1 to 0 at $\sum_{s\in S_{c_a}} Z_{ij}^{c_a,s} + 1$:

$$\sum_{s\in S_{c_a}} Z_{ij}^{c_a,s} = \sum_{f=1}^{|S_{c_a}|} I_{ij}^{c_a,f} \quad \forall c_a, i, j \tag{17}$$

$$I_{ij}^{c_a,f} \leq I_{ij}^{c_a,f-1} \quad 2 \leq f \leq |S_{c_a}|, \forall c_a, i, j \tag{18}$$

Once the $I_{ij}^{c_a,f}$ variables have been determined by the above constraints, they are used to determine the exact amount of traffic after aggregation on lightpath(s) from node i to node j due to the sources in S_{c_a}. This is done using the following set of constraints.

$$X_{ij}^{c_a} \geq r_f^{c_a} * \sum_{s \in S_{c_a}} m_{c_a,s} Z_{ij}^{c_a,s} - Q \sum_{s \in S_{c_a}} Z_{ij}^{c_a,s} + Q * f * I_{ij}^{c_a,f} \quad 1 \leq f \leq |S_{c_a}|, \forall c_a, i, j \tag{19}$$

$$X_{ij}^{c_a} \leq r_f^{c_a} * \sum_{s \in S_{c_a}} m_{c_a,s} Z_{ij}^{c_a,s} + Q \sum_{s \in S_{c_a}} Z_{ij}^{c_a,s} - Q * f * I_{ij}^{c_a,f} \quad 1 \leq f \leq |S_{c_a}|, \forall c_a, i, j \tag{20}$$

To understand how the above constraints will compute the exact amount of traffic after aggregation on lightpath(s) from node i to node j due to the sources in S_{c_a}, we divide the range of f into three sub-ranges: (1) $f < \sum_{s \in S_{c_a}} Z_{ij}^{c_a,s}$, (2) $f > \sum_{s \in S_{c_a}} Z_{ij}^{c_a,s}$, and (3) $f = \sum_{s \in S_{c_a}} Z_{ij}^{c_a,s}$. Note that the value $\sum_{s \in S_{c_a}} Z_{ij}^{c_a,s}$ is the number of the sources ($s \in S_{c_a}$) arriving at the node i and hence taking part in the aggregation. For sub-range (1), $X_{ij}^{c_a}$ will be greater than a large negative number and smaller than a large positive number. For sub-range (2), the values of $I_{ij}^{c_a,f}$ will be zero and again $X_{ij}^{c_a}$ will be greater than a large negative number and smaller than a large positive number. However, for sub-range (3) the last two terms will cancel each other and we will obtain an exact value for the amount of traffic after aggregation on lightpath(s) from node i to node j due to sources in S_{c_a}. The whole range $1 \leq f \leq |S_{c_a}|$ needs to be considered because the number of the data streams on lightpath(s) from node i to node j, due to the sources in S_{c_a}, itself is a variable determined by the MILP.

Finally, to bound the total capacity on lightpath(s) between node i and j, we add the following constraint

$$\sum_{c_a=1}^{K} X_{ij}^{c_a} \leq L_{ij} * g \quad \forall i, j \tag{21}$$

- *Non-Bifurcation:*

  By now we have ensured that each traffic request of a many-to-one session does not bifurcate (vertically split) its traffic between lighpaths of different node pairs. However, traffic requests can still get bifurcated over multiple lightpaths between a node pair. Such lighpaths can either be on different wavelengths or can even be on the same wavelength using different physical routes. We will make one assumption here that the total capacity of all aggregated traffic streams of a session is less than or equal to that of a single wavelength, i.e., $\sum_{s \in S_{c_a}} m_{c_a,s} \leq g$. We will observe this restriction while generating the traffic for the MILP and the heuristic in Section V. The following set of constraints, now, together with the above constraints, will ensure a complete non-bifurcated solution.

$$J_{ij}^{c_a,c_b} \leq (M_{ij}^{c_a} + M_{ij}^{c_b})/2 \quad \forall c_a \neq c_b, i, j \tag{22}$$

$$J_{ij}^{c_a,c_a} = M_{ij}^{c_a} \tag{23}$$

$$L_{ij} = J_{ij}^{c_1,c_1} + \sum_{a=2}^{K}(J_{ij}^{c_a,c_a} - \vee_{b=1}^{a-1} J_{ij}^{c_b,c_a}) \quad \forall i, j \tag{24}$$

$$X_{ij}^{c_a} + \sum_{c_b, c_b \neq c_a} Y_{ij}^{c_a,c_b} \leq g \quad \forall c_a, i, j \tag{25}$$

Equations (22) and (23) ensure that if $J_{ij}^{c_a,c_b}$ is 1 (i.e., sessions c_a and c_b share the same lightpath between nodes i and j), then at least one stream of session c_a and one stream of session c_b must be utilizing the lightpath from node i to node j. Equation (24) ensures that the total number of shared and unshared lighpaths from node i to node j must be equal to the total number of lightpaths from node i to node j. This equation counts all lightpaths from i to j by counting all the sessions sharing each lightpath from i to j only once. Equation (25) ensures that the capacity of each lightpath is not exceeded. The first term on the left hand side of the inequality (25) computes the effective bandwidth used by all the streams of the session c_a that share the lightpath between nodes i and j, while the second term computes the effective bandwidth used by all such streams, of all other sessions ($c_b, c_b \neq c_a$), which share the very same lightpath with streams of session c_a. Note that we have a non-linear term $Y_{ij}^{c_a,c_b} = J_{ij}^{c_a,c_b} * X_{ij}^{c_b}$ in the equation. We, however, compute the non-linear term $Y_{ij}^{c_a,c_b}$ using the following set of linear constraints

$$Y_{ij}^{c_a,c_b} \geq Q J_{ij}^{c_a,c_b} - Q + X_{ij}^{c_b} \quad \forall c_a, c_b, i, j \tag{26}$$

$$Y_{ij}^{c_a,c_b} \leq X_{ij}^{c_b} \quad \forall c_a, c_b, i, j \tag{27}$$

$$Y_{ij}^{c_a,c_b} \leq Q J_{ij}^{c_a,c_b} \quad \forall c_a, c_b, i, j \tag{28}$$

Finally, the disjunction operation, say, $V = \vee_{i=1}^{n} H_i$, $V \in \{0, 1\}$, is implemented using the following two constraints

$$V \leq \sum_{i=1}^{n} H_i \quad \text{and} \quad V \geq \sum_{i=1}^{n} H_i/n$$

The complexity of the MILP in terms of the number of the variables is $O(N^4W + N^2K^2)$ and in terms of the number of the constraints is $O(N^3W + N^2K^2)$, where K is the total number of many-to-one sessions. If $K > NW$, then both the number of the variables and the constraints are $O(N^2K^2)$. While if $K^2 < NW$, then the number of the variables and the constraints are $O(N^4W)$ and $O(N^3W)$, respectively.

IV Heuristic Approach

In this section we will present a heuristic approach for many-to-one traffic grooming on random topologies that also performs data aggregation and does not bifurcate the traffic originating at each source.

The search space for the many-to-one traffic grooming problem is huge. To design a heuristic using non-exhaustive search, one needs to make decisions, which may be greedy in nature, at a number of points during the search process. For example, the sequence in which each session is selected and accommodated will affect the cost associated with the accommodation of other sessions. Also, the decisions taken regarding the routing and wavelength assignment of each session will determine the many-to-one tree and thus the aggregation points. Moreover, data aggregation will take place on those physical links where two or more streams of the same many-to-one session share the very same wavelength. Once data aggregation takes place, we need to ensure that all the aggregated streams follow the same physical route afterwards until they reach the destination.

To explore the search space we designed a Dynamic Programming (DP) style algorithm, which we refer to as the Many-to-one Traffic Grooming heuristic based on Dynamic Programming (MTG-DP). Before describing the algorithm, it should be pointed out that this algorithm has the style of Dynamic Programming algorithms because it proceeds in stages, where in stage i different subsets of i streams each are routed and provisioned. The number of subsets, as will be explained below, is equal to the total number of streams. In stage i, stream j is provisioned together with $i-1$ streams, where such streams are chosen as one of the subsets of streams provisioned in stage $i-1$, such that the provisioning cost is minimized. This makes the computations at stage i dependent on stage $i-1$ only, hence the dynamic programming style.

An example of the DP-style algorithm is shown in Figure 2. Let each many-to-one session, c_a, consist of $|S_{c_a}|$ streams, one from each source $s \in S_{c_a}$, and let Φ be the total number of streams from all the many-to-one sessions, i.e., $\Phi = \sum_{c_a} |S_{c_a}|$. As shown in the figure, the algorithm consists of Φ stages, while each stage further consists of Φ steps. The sequence in which the streams are stage ordered, or are ordered within each stage, is taken arbitrarily. However, the algorithm checks several combinations during its execution. At each step of a stage, we accommodate a new stream (i.e., assign it a specific route and a specific wavelength) after conducting a systematic search, such that the decisions at each step of a stage depends only on the previous stage. The outcome of each step is a selection of a set of accommodated streams; let us call such a set, an assignment. At the end of each stage we have a collection of several best assignments. After the final stage, we simply select the assignment corresponding to the least cost.

For the systematic search at each stage, we examine all the assignments of previous stage and determine which assignment can accommodate the current stream with the least cost. In Figure 2, steps 2 and 4 of stages 2 and 3 are shown in detail.

Stage 1 is an initialization stage. At each step only one stream is accommodated. The I.D. of the stream is same as that of the step number. In stage 2 at each step only one new stream is to be accommodated. The I.D. of the to-be-accommodated stream is same as that of the step number. However, this to-be-accommodated stream is accommodated, one by one, with each of the assignments corresponding to each of the steps of stage 1. In the end the accommodation that results in the least cost is selected as the outcome of the current step of stage 2. For example, at step 2 of stage 2, the stream with I.D. 2 is the to-be-accommodated stream. However, in order to do that, all the assignments from stage 1, i.e., $\{1\}, \{3\}, \{4\}, ..., \{\Phi\}$, are considered one by one and are combined with stream 2. Hence, the new assignments explored in step 2 of stage 2 are $\{1,2\}, \{3,2\}, \{4,2\}, ..., \{\Phi,2\}$, Where $\{1,2\}$ means that the streams with I.D. 1 and 2 are accommodated. After evaluating the cost of each of the new assignments the outcome of step 2 in stage 2 is the assignment $\{b^*,2\}$, where b^* could be any of the previous assignments $\{1\}, \{3\}, \{4\}, ..., \{\Phi\}$, which when accommodated together with stream 2, incurs the least cost. Similarly, the outcome of the step 4 of the stage 2 is the assignment $\{d^*,4\}$.

To present the heuristic formally we define the additional terms given in Table 1. Also, we define a procedure Evaluate($\phi_{s,d,k}, P, \Lambda_{s,d,k}(P)$) which computes the total number of the LTs in the network, if $\phi_{s,d,k}$ is routed along path P using wavelength(s) specified by $\Lambda_{s,d,k}(P)$.

The pseudo-code of the heuristic is shown in Figure 3. First, we generate $\mathbf{K}_{s,d,k}$ alternate shortest paths from each source s in a session k to its destination d (the value of $\mathbf{K}_{s,d,k}$ and the manner in which the alternate shortest paths are obtained will be discussed below). The algorithm then considers combinations of shortest paths from all sources to the destination within a session, such that one shortest path from each source to the destination is considered in a combination. For each such combination, we repeat the following procedure. In the very first stage, at each step, only one stream is accommodated using shortest path (line 4). This creates Φ independent assignments, each consisting of a single stream. The dynamic programming approach is implemented in lines

Symbol	Meaning
$\phi_{s,d,k}$	a stream originating at s, destined to d, and belonging to session k
$\gamma_{d,k}$	a set of streams, $\phi_{s,d,k}$, corresponding to a many-to-one session, destined to d, and belonging to session k,i.e., $\{\phi_{s,d,k} \mid s \in S_{C_k}\}$
$\Pi_{i,j}$	an assignment of accommodated streams at step j of stage i
P	a set of physical links corresponding a path
$\Lambda_{s,d,k}(P)$	a set of wavelengths used by $\phi_{s,d,k}$ on the physical links of path P
$\Delta_{i,j}$	cost of accommodating a stream at stage i, given the set $\Pi_{i-1,j}$

Table 1: Variables used in the heuristic

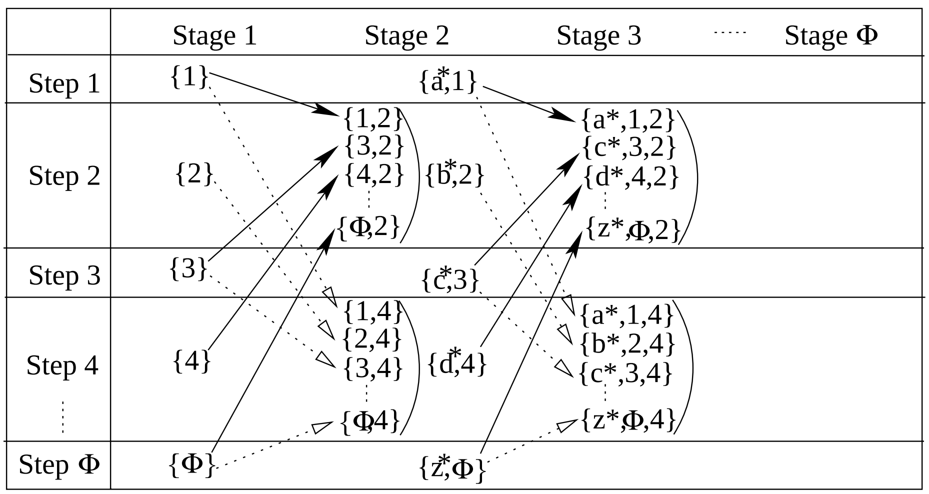

Figure 2: Graphical depiction of the many-to-one traffic grooming heuristic. It is assumed that $a^* \neq 1, b^* \neq 2, c^* \neq 3, d^* \neq 4$, and $z^* \neq \Phi$

5-12, where each stage, i for $i = 2to\Phi$, is evaluated in terms of the pervious stage, $i-1$ only. The algorithm then, in each following stage, iteratively adds other streams by selecting the best assignments (lines 5-12). The stream to-be accommodated at a step is $\phi_{s,d,k}$. To select the best assignment for $\phi_{s,d,k}$, a search procedure is carried out (lines 6-12). All those assignments, from the previous stage only, which do not include the $\phi_{s,d,k}$ are considered (lines 6-7). The assignment that contributes the least cost if $\phi_{s,d,k}$ is added to it is selected, by using the procedure compute_cost(), and $\phi_{s,d,k}$ is made part of that assignment (lines 10-12). In the final stage, the assignment that incurs the least number of LTs is selected.

To compute the cost (line 9) in Figure 3, we employ the procedure shown in Figure 4. In Figure 4 the shortest path for the stream $\phi_{s,d,k}$ is used (line 1), and different wavelength-assignments are explored (lines 2-9). The cost corresponding to first-fit wavelength-assignment, on the shortest-path, is stored in the variable δ_0 (line 3). Note that we select such a first-fit wavelength that does not share the wavelength with any of the other streams of the same many-to-one session. Consequently, no data aggregation can take place on the selected first-fit wavelength. We then explore wavelength-assignment of each stream $\phi_{s',d,k} \in \gamma_{d,k}$, provided $\phi_{s',d,k}$ is also part of the wavelength-assignment under consideration (lines 4-8). We divide the wavelength-assignment of the physical links corresponding to the path of stream $\phi_{s,d,k}$ into two subsets, P_1 and P_2. An illustration is given in Figure 5. The subset P_1 consists of physical links that are common between stream $\phi_{s,d,k}$ and $\phi_{s',d,k}$, while the subset P_2 consists of physical links that are part of the path of stream $\phi_{s,d,k}$ but are not part of the path of stream $\phi_{s',d,k}$. We use first-fit wavelength-assignment for the members of the subset P_2, while we follow the wavelength-assignment of the other stream, $\phi_{s',d,k}$, for the members of the subset P_1, using aggregated capacity. In case the wavelength used by the other stream $\phi_{s',d,k}$ is not accommodating, an infinite cost is recorded. Note that as the wavelength-assignment of the other stream, $\phi_{s',d,k}$, could also be using the same wavelength as determined by the first-fit wavelength-assignment, we eventually end up aggregating the streams even on the first-fit wavelength. The only restriction is that aggregation is allowed when streams follow the same wavelength from the point of their first merging till the destination. This ensures non-bifurcation of streams once they are aggregated. Finally, we select the wavelength-assignment that results in the minimum cost (line 9). This procedure returns: 1) the least cost for accommodation of stream $\phi_{s,d,k}$, 2) the corresponding physical path, and 3) the corresponding wavelength-assignment. The heuristic thus explores, in a systematic way, the different ways all the streams can be accommodated and explores the ways of constructing the different many-to-one trees that employ aggregation.

The reason alternate shortest path were explored is to allow different streams from the same session to share common links, and may therefore be aggregated, hence resulting in fewer LTs. We explored different approaches for choosing alternate shortest paths. We found that using disjoint shortest paths for the same stream results in no improvement in most cases, and may in fact result in an increased cost. The approach that consistently resulted in a better cost reduction is to choose alternate paths which consist of the concatenation of: 1) the shortest path to reach a path used by another stream in the same session, and 2) the remainder of the second path to the destination. The number of alternate paths per stream is therefore on the order of session size, $|S_{c_k}|$, for sessoin c_k. This means that the procedure of the algorithm in Figure 3 will have to be repeated for a number of steps that is $O(\Pi_{c_k}|S_{c_k}|^{S_{c_k}})$, which makes this approach computationally infeasible. However, we found that the approach shown in Figure 6 produced good results, while being computationally tractable. In this procedure, we start with all streams routed on the shortest paths from sources to destination. Then, for each stream, an alternate path is explored, and if it results in a lower cost, it is used for routing this stream. Of course, this procedure is not optimal since there may be possible combinations of stream routing strategies which are not explored. The cost reduction realized by this approach reached 10%, and the number of repetitions of the algorithm in Figure 3 is therefore $O(\sum_{c_k} |S_{c_k}|^2)$. Since the *compute_cost*() is called Φ^3 times, and the complexity of the procedure itself is $O(N^2)$, then the complexity of the heuristic is $O(\Phi^3 N^2 \sum_{c_k} |S_{c_k}|^2)$.

Algorithm: Many-to-one Traffic Grooming

Input: Many-to-one traffic matrix

Output: Number of LTs and wavelengths required to accommodate the input traffic matrix, and corresponding dimensioning of the network

BEGIN

1. FOR each $\phi_{s,d,k}$
2.   Generate $\mathbf{K}_{s,d,k}$ shortest paths from source s to destination d in session k
3. FOR each combination of shortest paths
4. $\Pi_{1,j} = j;\qquad 1 \leq j \leq \Phi$
5.   FOR $i = 2$ to Φ //for each stage
6.   FOR each $\phi_{s,d,k}$ //for each stream
7.   FOR $j = 1$ to Φ //from each of the steps in previous stage
8.   IF($\phi_{s,d,k} \notin \Pi_{i-1,j}$)   //if current stream is not a part of the assignment
9. $(\Delta_{i,j}, P^{(j)}, \Lambda^{(j)}_{s,d,k}(P)) \leftarrow compute\_cost(\Pi_{i-1,j}, \phi_{s,d,k});$
10. $j^* \leftarrow arg(min_j \Delta_{i,j})$ //select least cost assignment;
11. $\Pi_{i,<s,d,k>} = \Pi_{i-1,j^*} \cup \phi_{s,d,k};$
12.   Route $\phi_{s,d,k}$ using $P^{(j^*)}$ and $\Lambda^{(j^*)}_{s,d,k}(P)$;
13. Choose the assignment with least number of LTs, breaking the tie with the least number of wavelengths.

END

Figure 3: Many-to-one traffic grooming heuristic.

V Experimental Results

In this section we will present the results of the MILP as well as the heuristic approach for the many-to-one traffic grooming problem. To experiment with the MILP, we used the network shown in Figure 7, while for the heuristic we also considered the NSF network topology shown in Figure 8. The traffic demands consist of integer multiples of OC-3 sessions. The capacity of a wavelength compute_cost($\Pi_{i-1,j}, \phi_{s,d,k}$)

1. $P_{s,d,k}$ = shortest path for $\phi_{s,d,k}$ in the combination of paths;
2. $t = 0$;
3. $\delta_t = Evaluate(\phi_{s,d,k}, P_{s,d,k}, FF)$; //determine the cost using First-Fit wavelength on path $P_{s,d,k}$
4. FOR each $\phi_{s',d,k} \in (\gamma_{d,k} \cap \Pi_{i-1,j})$
5. $t \leftarrow t + 1$;
6. $P_1 = P_{s,d,k} \cap P_{s',d,k}$;
7. $P_2 = P_{s,d,k} - P_{s',d,k}$;
   //use aggregated capacity on $\Lambda_{s',d,k}(P_1)$
8. $\delta_t = Evaluate(\phi_{s,d,k}, P_2, FF) + Evaluate(\phi_{s,d,k}, P_1, \Lambda_{s',d,k}(P_1))$;
9. $\Lambda_{s,d,k}(P_{s,d,k}) \leftarrow min_t \delta_t$; // wavelength-assignment that results in minimum cost
10. return($min_t \delta_t, P_{s,d,k}, \Lambda_{s,d,k}(P_{s,d,k})$)

Figure 4: Procedure for computing the cost of accommodating a stream given an assignment of already accommodated streams.

is OC-48, and therefore, the grooming factor, g, is 16.

V.1 Exact Solutions

The MILP problem is solved using the CPLEX linear programming package [33]. The values of α and β are selected to be 100 and 1, respectively. We selected these values to give higher priority to the minimization of the number of LTs over minimization of the number of wavelengths in the network. A sample traffic that consists of a mix of many-to-one and one-to-one sessions is generated and is shown in Table 2. The last column shows the fractions, from each data stream, to be accepted

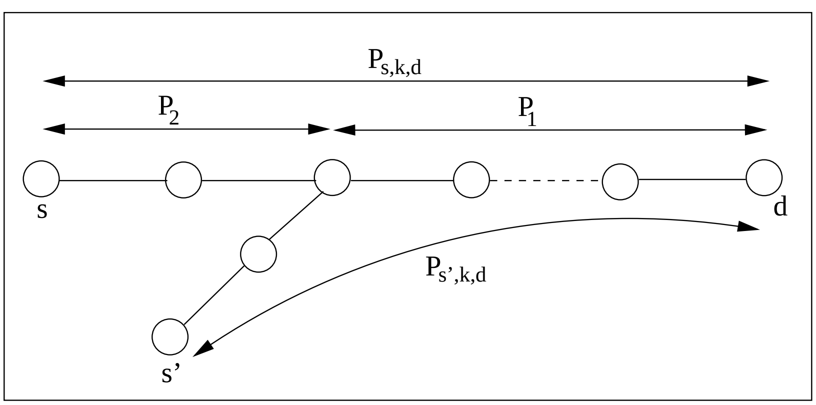

Figure 5: Illustration of different paths for the wavelength-assignment in the many-to-one traffic grooming heuristic

1. Route all streams on the shortest paths to the destination
2. Compute cost Δ
3. For all streams $\phi_{s,d,k}$
4.   For i=1 to $\mathbf{K}_{s,d,k}$
5.   Use alternate shortest path i
6.   Compute Cost $\Delta^{(i)}_{s,d,k}$
7.   if $\Delta^{(i)}_{s,d,k} < \Delta$
8.   Route session $\phi_{s,d,k}$ on alternate shortest path i
9. $\Delta = \Delta^{(i)}_{s,d,k}$

Figure 6: Procedure for exploring alternate paths

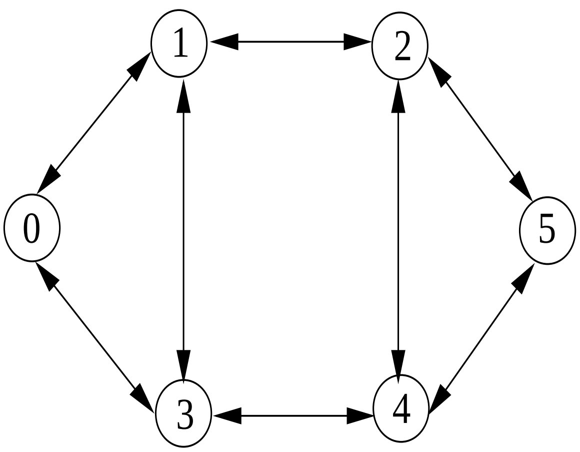

Figure 7: A six node network.

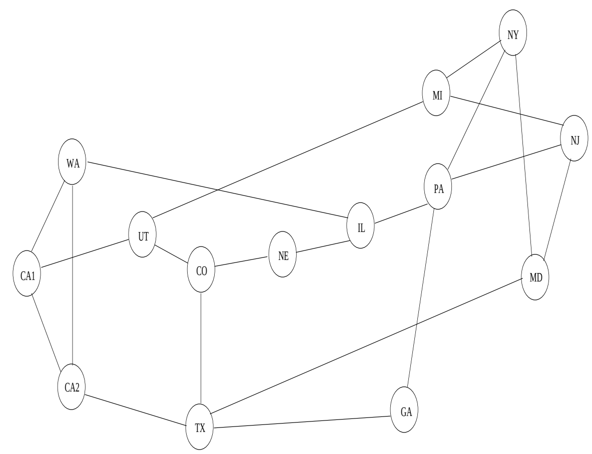

Figure 8: NSF Network topology.

for aggregation corresponding to the number of the streams being aggregated. These values are generated randomly between 0 and 1 such that $r_1^{c_a} = 1$, and $r_f^{c_a} \leq r_{f-1}^{c_a}, 2 \leq f \leq |S_{c_a}|$. Multiple fractions correspond to the number of the streams being aggregated. Hence, for example, for a many-to-one session terminating at node 0 and belonging to session 1, the fractional value is 1 if a single stream is aggregated, while it is 0.69 if two streams are aggregated and 0.484 if three streams are aggregated. Say if the capacity of 2 streams are m_1 and m_2 and these are being aggregated, then the total capacity of aggregated stream will be $0.69 * (m_1 + m_2)$. Note that each many-to-one session has an independent set of values. It is possible that different many-to-one sessions correspond to different applications, and hence aggregate the traffic using different fractions. A higher value of the fraction represents less correlation between the data streams of a many-to-one session, and vice versa. For example, note that the streams destined to node 3 have a smaller degree of correlation. We run the problem using 4 wavelengths. However, the optimal solution reduced the total number

Table 2: Many-to-one traffic demands supported on the network topology in Figure 7

Destination	Session	Source set	Traffic (multiples of OC-3)	Aggregation ratios
0	1	{1,2,4}	{6,4,6}	{1, 0.690, 0.484}
1	2	{0}	{12}	{1}
	3	{3,5}	{8,6}	{1, 0.698}
2	4	{1}	{8}	{1}
	5	{0,3,5}	{3,8,4}	{1, 0.889, 0.863}
3	6	{0}	{16}	{1}
	7	{0,1}	{12,3}	{1, 0.997}
5	8	{3}	{8}	{1}
	9	{1,4}	{3,8}	{1, 0.666}

of wavelengths to 2. The total number of LTs required for this example is 13. Figure 9.a shows the solution in terms of the LTs and how the two wavelengths are provisioned on the different physical links. In Figure 9.b we also show how session 1 is provisioned. In this case, the traffic from sources 1 and 4 is aggregated at node 1. However, the traffic from node 2 is not aggregated. The source traffic is shown below the node's number, while the total traffic from the session served by a particular lightpth is shown on the lightpath, and it may or may not be aggregated.

We also conducted an experiment on the NSF Network topology shown in Figure 8. The traffic demands are shown in Table 3. The number of sessions per destination was uniformly distributed between 0 and 3, and the number of sources per session was also uniformly distributed between 1 and 4. The demand per source was chosen from the following set of demands: {1, 3, 9, 12}. The grooming factor was chosen as 12. The aggregation ratio of aggregating k streams is given by $1 - 0.1 \times (k - 1)$, as shown in the table. The optimal formulation produced a solution of 29 LTs and 5 wavelength channels.

V.2 Results From the Heuristic Solution

In this section we show results obtained from the heuristic solution. First, we verify the accuracy of the heurisitc solution by comparing results obtained from the heuristic to the two sets of optimal results obtained in Subsection V.1. Then, we show results obtained when the heuristic is run for large sets of demands on the NSF network.

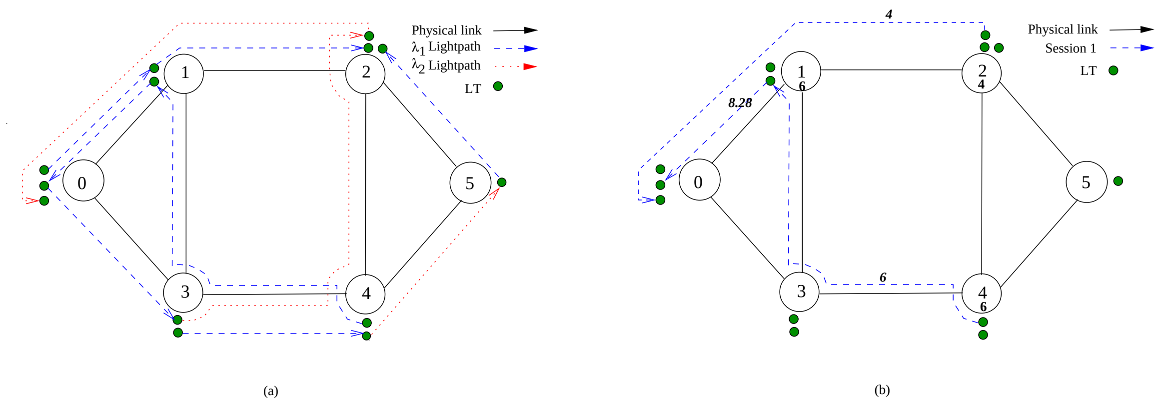

Figure 9: The exact solution for the traffic in Table 2 on the network in Figure 7: (a) Location of LTs and wavelength provisioning on physical links; (b) The provisioning of session 1.

Table 3: Many-to-one traffic demands supported on the network topology in Figure 8

Destination	Session	Source set	Traffic (multiples of OC-3)	Aggregation ratios
1	1	{9}	{1}	{1}
	2	{0,8,11}	{1,3,9}	{1, 0.9, 0.8}
3	3	{8,10}	{1,12}	{1, 0.9}
	4	{12}	{9}	{1}
	5	{5,7,10,12}	{3,3,9,9}	{1, 0.9, 0.8, 0.7}
4	6	{1,2,10}	{1,3,9}	{1, 0.9, 0.8}
	7	{7,8}	{12,1}	{1, 0.9}
5	8	{2,8,11}	{9,3,1}	{1, 0.9, 0.8}
8	9	{6}	{1}	{1}
9	10	{12}	{9}	{1}
	11	{0, 2, 13}	{9, 3, 1}	{1, 0.9, 0.8}
10	12	{13}	{3}	{1}
	13	{1, 4}	{12, 1}	{1, 0.9}
	14	{0}	{9}	{1}
12	15	{10}	{9}	{1}

V.2.i Verification of the Heuristic Accuracy

Due to the high complexity of the problem, one cannot obtain optimal solutions within reasonable time for large topologies. Therefore, one needs to resort to the heuristic approach. However, the above MILP can act as a baseline to determine the performance of any heuristic designed for many-to-one traffic grooming problem. Our heuristic approach, when run on the topology shown in Figure 7 and the inputs shown in Table 2, produced a solution with 16 LTs and 2 wavelengths with one shortest path. While the number of wavelength channels is the same as that produced by the optimal solution, the number of LTs is 23% more than the optimal number of LTs. However, when the heuristic was run with $\mathbf{K} = 3$, i.e., a maximum of three alternate shortest paths, the heuristic produced a solution of 15 LTs and 2 wavelengths. The solution in this case is within 15% of the optimal solution. Thereefore, and as expected, the use of alternate shortest paths for accommodating the different streams can result in a better solution.

We also applied the heuristic to the traffic demands in Table 3 which are carried on the NSF network in Figure 8. The heuristic produced 40 LTs, and 6 wavelength channels. This is an increase of about 38% more LTs over the optimal solution.

For other small size, and limited traffic examples, we observed that the heuristic approach produces results, in terms of the number of LTs, which are within 10%-38% of the optimal values.

V.2.ii Large Network Examples

To collect the results on some real network topologies, we chose the NSF network topology shown in Figure 8. The grooming factor, g, used is 48. To study the effect of aggregation, traffic granularity, grooming factor, and source set size, we conducted many experiments and divided them into four scenarios.

We generated a traffic matrix randomly with the following parameters, and used it for all the experiments:

- The number of sessions destined to each node is generated uniformly between 0 and 3.
- 50% of the sessions carry many-to-one traffic, while the remaining sessions carry unicast traffic.
- For each many-to-one session, the source set size is uniformly distributed between 2 and 8.
- For each many-to-one session terminating at a destination, the sources are chosen randomly among all the nodes, excluding the destination, for both unicast and many-to-one sessions.
- The generated traffic, for both unicast and many-to-one sessions, is an integer multiple of OC-1, and is uniformly chosen from the set $\{1,3,9,12,18,24,36,48\}$. These values represent the recommended rates for OC-1 streams.

Since the type and amount of aggregation can be application dependent, we use three different ways to compute the aggregation ratios, as follows.

- Aggregation Type 1: All aggregation ratios are unity, i.e, $r_f^{c_a} = 1, \forall f$. Hence when two or more streams are aggregated, no capacity reduction will take place
- Aggregation Type 2: Aggregation ratios decrease linearly with a step σ, i.e., $r_f^{c_a} = r_{f-1}^{c_a} - \sigma, 2 \leq f \leq |S_{c_a}|$. We assume that $|S_{c_a}| < 10$, in this case we have chosen σ to be 0.1.
- Aggregation Type 3: Aggregation ratios are computed using the formula, $r_f^{c_a} = 1 - \ln(f)/\rho$, which provides logarithmically decreasing values. We have chosen ρ to be 5. Hence, quantitatively, when the number of aggregated streams increases, the difference in the reduction of the capacity of each stream decreases. In other words, the benefit we achieve, in terms of reduction in capacity, by aggregating two streams over one stream, is more than that of aggregating three streams over two aggregated streams.

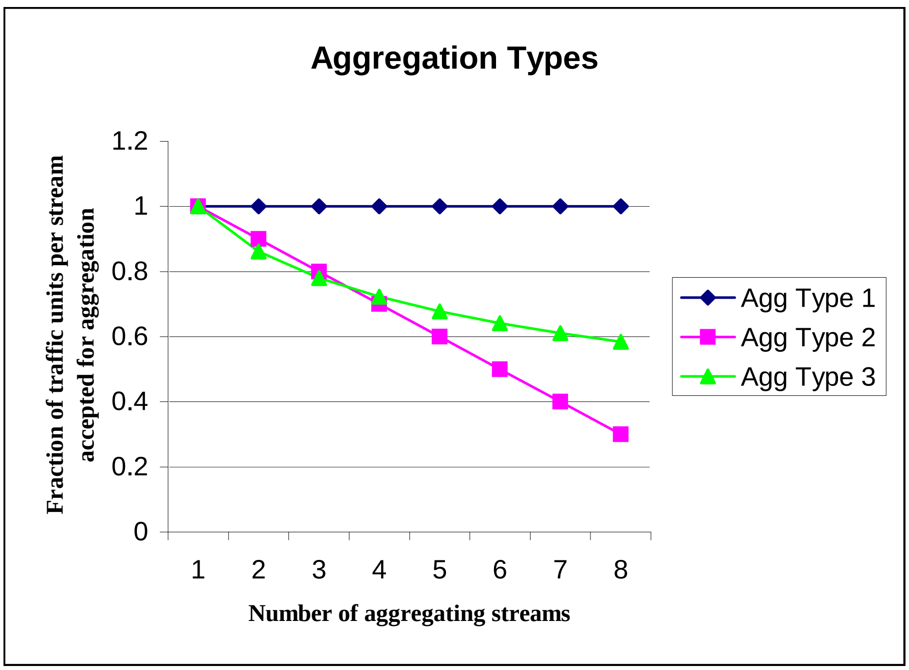

Figure 10: Different aggregation types used for collecting experimental results

The three aggregation types are graphically shown in Figure 10. Note that aggregation type 3 provides the maximum per stream traffic reduction if the number of the aggregated streams are 2 or 3, but if the number of the aggregated streams are more than 3 then aggregation type 2 provides more reduction.

We have conducted experiments using four scenarios, where for each scenario, we assumed that all many-to-one sessions have the same aggregation type and ratios. Now we will discuss experiments in each scenario. The experiments in scenario 1 are conducted to study the effect of the aggregation factor on the network cost. We conducted three experiments using the three aggregation types. The results are shown in Table 4. MTG-DP with alternate shortest paths outperformed the MTG-DP with one shortest path in all cases. The comparison between both heuristics shows the benefit of the extra computations we are performing in MTG-DP with alternate shortest paths. For aggregation types 1, i.e., no aggregation, the MTG-DP with alternate shortest paths resulted in a 1.5% saving in the number of LTs, while for aggregation types 2 and 3, the saving in the number of LTs when MTG-DP with alternate shortest paths was used is 8%, and 5%, respectively. Notice that under the fixed shortest path approach, under aggregation type 2 the number of LTs is greater than the number of LTs under aggregation type 1, but with one wavelength less. Using alternate shortest paths, the number of LTs was reduced under aggregation type 2, and is even less than the number of LTs under aggregation type 1.

The experiments in scenario 2 are conducted to study the effect of the traffic granularity on the cost of the network. We modified the original traffic matrix by revisiting each of the sources and replacing its current traffic amount with the next higher traffic amount in the vector {1,3,9,12,18,24,36,48}, unless the traffic was already 48 units, in which case it was not changed. For example, if a source was previously sending out 18 units of traffic, it will now be sending out 24 units of traffic. We used this modified traffic matrix to run three different experiments with aggregation types 1, 2 and 3, respectively. The results are shown in Table 5. Again, the alternate shortest path heuristic outperformed the MTG-DP approach with a single shortest path, and without increasing the number of wavelength channels. One should also observe that more aggressive aggregations, i.e., smaller aggregation ratios result in offsetting the increase in the traffic granularity. For example, with aggregation types 1 and 2, the number of additional required LTs is 10, while under aggregation type 3, it is only limited to 7.

The experiments in scenario 3 are conducted to study the effect of the grooming factor on the cost of the network. The results are shown in Table 6. We used the traffic matrix with the original load, but this time we set the grooming factor to 96. Hence the capacity of each participating wavelength is now doubled. As a result, we see a substantial decrease in the number of wavelength channels, which exceeds 50% in the case of fixed shortest paths. With alternate shortest paths, the reduction in the number of wavelength channels is still significant, and is about 25%. The reduction in the number of LTs, is also significant, and is on the order of 25%. It is not be expected that a 50% reduction in the number of LTs can be achieved in this case due to a number of factors such as the non-bifurcation of traffic, and the limited number of streams to be groomed on the same

Table 4: Results of experiments in scenario 1; grooming factor is 48.

		Agg. type 1	Agg. type 2	Agg. type 3
MTG-DP with shortest paths	LTs	67	69	65
	wavelengths	5	4	5
MTG-DP with alternate shortest paths	LTs	66	63	62
	wavelengths	4	4	4

Table 5: Results of experiments in scenario 2; grooming factor is 48 and traffic granularity is increased

		Agg. type 1	Agg. type 2	Agg. type 3
MTG-DP with shortest paths	LTs	79	77	72
	wavelengths	5	5	5
MTG-DP with alternate shortest paths	LTs	76	73	69
	wavelengths	5	5	5

wavelength. As the number of streams directed from a source to the same desination increases, we should expect a greater reduction in the number of LTs. Another thing that is to be noticed here is that increasing the grooming factor results in an incremental improvement in the cost of the network under the case of no aggregation, i.e., type 1 aggregation, that is greater than the incremental improvement under types 2 and 3. This can be justified by noting that aggregation is more useful when a small reduction in the total traffic granularity can result in fitting more streams on a wavelength. With increasing the grooming factor this mode of improvement becomes less pronounced since the wavelength channel bandwidth is no longer a bottleneck.

Finally, the experiments in scenario 4 are conducted to study the effect of varying the size of the source set on the network cost. We modified the original traffic matrix by randomly adding 3 additional sources per session. We employ aggregation types 1, 2 and 3 for three different experiments. The results are shown in Table 7. Similar to previous scenarios, the proposed heuristic MTG-DP with alternate shortest paths resulted in further savings in terms of the number of LTs and the number of wavelength channels when compared to the MTG-DP algorithm with one shortest path. One thing to be noticed here is that alternate shortest paths did not result in a significant

Table 6: Results of experiments in scenario 3; grooming factor is 96.

		Agg. type 1	Agg. type 2	Agg. type 3
MTG-DP with shortest paths	LTs	50	50	50
	wavelengths	2	2	2
MTG-DP with alternate shortest paths	LTs	49	49	47
	wavelengths	3	3	3

Table 7: Results of experiments in scenario 4; 3 additional sources per session are added to the original traffic matrix.

		Agg. type 1	Agg. type 2	Agg. type 3
MTG-DP with shortest paths	LTs	89	81	77
	wavelengths	6	5	6
MTG-DP with alternate shortest paths	LTs	82	80	75
	wavelengths	5	5	5

saving under aggregation types 2 and and 3. The reason for this is that with an increased number of sources per session, the shortest paths for streams within the same session will most probably overlap, hence allowing aggregation naturally. Investigating alternate shortest paths will not result in any further significant overlap of streams. From this we may conclude that with a large number of source per stream the shortest path approach may result in good performance.

VI Conclusions

In this paper we addressed the design of optical network which supports many-to-one traffic grooming, using optimal as well as heuristic approaches. We developed a mixed integer linear program of an otherwise non-linear problem, by exploiting the specifics of routing and aggregation subproblems. The model used is quite generic and supports arbitrary many-to-one sessions. Moreover, at merging points in many-to-one trees, a fraction of the traffic from each source can be selected such that the fraction itself depends on the number of the streams being aggregated. Furthermore, we ensured the non-bifurcation of the traffic. We also designed a Dynamic Programming style heuristic that explored the solution through a number of stages. At each stage the heuristic selected a number of best partial solutions from a large number of candidate partial solutions. Alternate shortest paths from the sources to the destination within a session were explored, and it was found that such exploration can result in reducing the cost by a factor that ranges from 5% to 10%. Experimental results showed that the MILP can be used for small to medium sized examples, while the heuristic approach can be used for large sized networks. Our proposed heuristic approach produced results that are within 10% to 38% of the optimal solution within reasonable times. The degree of sub-optimality was found to increase as the size of the network increases.

Our future work includes exploring approaches to improve the performance of the heuristic, while also exploring other heuristic approaches. Moreover, we also plan to investigate parallelizing the execution of the MILP solver in order to use our MILP formulation to solve large networks. In addition, we plan to investigate traffic grooming of any-to-any communication. This problem is more involved since it requires supporting both many-to-one and one-to-many types of traffic.

Acknowledgements

The authors would like to thank the editor and the reviewers for their constructive comments which lead to improving the presentation of the paper.

References

[1] A. L. Chiu and E. H. Modiano, "Traffic Grooming Algorithms for Reducing Electronic Multiplexing Costs in WDM Ring Networks," *IEEE Journal of Lightwave technology*, Vol. 18, No. 1, pp. 2-12, January 2000.

[2] J. M. Simmons, E. L. Goldstein, and A. Saleh, On the Value of Wavelength-add/drop in WDM Rings With Uniform Traffic, In *Proceedings of Optical Fiber Communication '98*, pp. 361-362, 1998.

[3] O. Gerstel, R. Ramaswami, and G. Sasaki, Cost-Effective Traffic Grooming in WDM Rings, In *Proceedings of IEEE INFOCOM'98*, Vol. 1, pp. 69-77, March 1998.

[4] Angela Chiu and Eytan Modiano, Reducing Electronic Multiplexing Costs in Unidirectional SONET/WDM Ring Networks Via Efficient Traffic Grooming, In *Proceedings of Globecom'98*, Sydney, Australia, Vol. 1, pp. 322-327, November 1998.

[5] X. Zhang and C. Qiao, An Effective and Comprehensive Approach to Traffic Grooming and Wavelength Assignment in SONET/WDM Rings, In *Proceedings of SPIE Conference On All-optical Networking*, Boston, MA, Vol. 3531, pp. 221-232, September 1998.

[6] R. Dutta and G. N. Rouskas, "On Optimal Traffic Grooming in WDM Rings," *IEEE Journal on Selected Areas in Communication*, Vol. 20, No. 1, pp. 110-121, January 2002.

[7] R. Ul-Mustafa and A. E. Kamal, "Design and Provisioning of WDM Networks with Multicast Traffic Grooming", to appear in the IEEE Journal on Selected Areas in Communications.

[8] G. Feng, C. Siew, and T-S. Yum, "Architectural Design and Bandwidth Demand Analysis for Multiparty Videoconferenceing on SONET/ATM Rings," *IEEE Journal on Selected Areas in Communications*, Vol. 20, No. 8, pp. 1580-1588, October 2002.

[9] S. McCanne and V. Jacobson and M. Vetterli, "Receiver-Driven Layered Multicast", ACM SIGCOMM 1996, pp. 117–130.

[10] D. Cavendish et al., "New Transport Services for Next-Generation SONET/SDH Systems", IEEE Communications, Vol. 40, No. 5, May 2002, pp. 80–87.

[11] E. Hernandez-Valencia, M. Scholten and Z. Zhu, "The Generic Framing Procedure (GFP): An Overview", IEEE Communications, Vol. 40, No. 5, May 2002, pp.63–71.

[12] T. Takeda et al., "Layer 1 Virtual Private Networks: Service Concepts, Architecture Requirements, and Related Advances in Standardization", IEEE Communications, Vol. 42, No. 6, June 2004, pp. 132–138.

[13] E. Modiano, "Traffic Grooming in WDM Networks," *IEEE Communications Magazine*, pp. 124-129, July 2001.

[14] R. Dutta and G. N. Rouskas, "Traffic Grooming in WDM Networks: Past and Future," *IEEE Network*, Vol. 16, No. 6, pp. 46-56, November/December 2002.

[15] K. Zhu and B. Mukherjee, "A Review of Traffic Grooming in WDM Optical Networks: Architectures and Challenges," *Optical Networks Magazine*, Vol. 4, No. 2, March/April 2003.

[16] O. Gerstel, P. Lin, and G. Sasaki, "Combined WDM and SONET Design, In *Proceedings of IEEE INFOCOM'99*, pp. 734-743, 1999.

[17] X. Zhang and C. Qiao, "An Effective and Comprehensive Approach for Traffic Grooming and Wavelength Assignment in SONET/WDM Rings," *IEEE/ACM Transactions on Networking*, Vol. 8, No. 5, pp. 608-617, October 2000.

[18] R. Ul-Mustafa and A. E. Kamal, "WDM Network Design with Non-Uniform Traffic Grooming," In *Proceedings of Optical Network Design and Modeling*, Vol. 2, pp. 965-984, Budapest, Hungary, February 2003.

[19] O. Gerstel, R. Ramaswami, and G. Sasaki, "Cost-Effective Traffic Grooming in WDM Rings," *IEEE/ACM Transactions on Networking*, Vol. 8, No. 5, pp. 618-630, October 2000.

[20] P. -J. Wan, G. Calinescu, L. Liu, and O. Frieder, "Grooming of Arbitrary Traffic in SONET/WDM BLSRs," *IEEE Journal of Selected Areas in Communications*, pp. 1995-2003, 2000.

[21] G. Sasaki and T. Lin, "A Minimal Cost WDM Network for Incremental Traffic," In *Proceedings of SPIE Conference on All-Optical Networking*, Vol. 3843, September 1999.

[22] V. Konda and T. Chow, "Algorithm for Traffic Grooming in Optical Networks to Minimize the Number of Transceivers," In *Proceedings of 2001 IEEE Workshop on High Performance Switching and Routing*, pp. 218-221, May 2001.

[23] M. Brunato and R. Battiti, "A Multistart Randomized Greedy Algorithm for Traffic Grooming on Mesh Logical Topologies," In *Proceedings of Optical Network Design and Modeling*, Vol. 242, Torino, Italy, February 2002.

[24] K. Zhu and B. Mukherjee, "Traffic Grooming in a WDM Mesh Network," *IEEE Journal on Selected Areas in Communications*, Vol. 20, No. 1, pp. 122-133, January 2002.

[25] S. Thiagarajan and A. K. Somani, "Traffic Grooming for Survivable WDM Mesh Networks," In *Proceedings of OPTICOMM 2001*, August 2001.

[26] R. Srinivasan and A. K. Somani, "Dynamic Routing in WDM Grooming Networks," In *Photonic Network Communications*, Vol. 5, No. 2, pp. 123-135, March 2003.

[27] H. Saito, Y. Miyao and M. Yoshida, "Traffic Engineering Using Multiple Multipoint-to-Point LSPs", in the proceedings of IEEE Infocom 2000, pp. 894–901.

[28] A. Srikitja and D. Tipper, "Topological Design of Multiple VPNs over MPLS Networks", in the proceedings of IEEE Globecom 2002, pp. 243–247.

[29] S. Bhatnagar, S. Ganguly and B. Nath, "Creating Multipoint-to-Point LSPs for Traffic Engineering", IEEE Communications, Vol. 43, No. 1, January 2005, pp. 95–100.

[30] R. Ramaswami and K. Sivarajan, *Optical Networks: A Practical Perspective*, Morgan Kaufmann, 2nd edition, San Francisco, CA, 2001.

[31] H. V. Madhyastha, G. V. Chowdhary, N. Srinivas and C. Siva Ram Murthy, "Grooming of Multicast Sessions in Metropolitan WDM Ring Networks", Computer Networks, Vol. 49, No. 4, Nov. 2005, pp. 561-579.

[32] A. Billah, B. Wang, and A. Awwal "Multicast Traffic Grooming in WDM Optical Mesh Networks," In *Proceedings of Globecom 03*, San Francisco, CA, December 2003.

[33] http://www.ilog.com/products/cplex/