Design and Provisioning of WDM Networks with Multicast Traffic Grooming *

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

In this paper we consider the optimal design and provisioning of WDM networks for the grooming of multicast subwavelength traffic. We develop a unified framework for the optimal provisioning of different practical scenarios of multicast traffic grooming. We also introduce heuristic solutions. Optimal solutions are designed by exploiting the specifics of the problems to formulate Mixed Integer Linear Programs (MILPs). Specifically, we solve the generic multicast problem in which, given a set of multicast sessions and all destination nodes of a multicast session requiring the same amount of traffic, all demands need to be accommodated. The objective is to minimize the network cost by minimizing the number of higher layer electronic equipment and, simultaneously, minimizing the total number of wavelengths used. We also solve two interesting and practical variants of the traditional multicast problem, namely, multicasting with partial destination set reachability and multicasting with traffic thinning. For both variants, we also provide optimal as well as heuristic solutions. Also, the paper presents a number of examples based on the exact and heuristic approaches.

Keywords: Wavelength Division Multiplexing (WDM) networks; traffic grooming; Add/Drop Multiplexers (ADM); IP routers; multicast traffic; optimization; Mixed Integer Linear Programming (MILP); optimal design and provisioning; heuristics.

---

*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.

‡Corresponding author: Tel.: (515) 294-3580; Fax: (515) 294-1152; E-mail: kamal@iastate.edu.

I Introduction

Wavelength Division Multiplexing (WDM) has emerged as a means for exploiting the enormous amount of bandwidth available on optical fibers. WDM helps to utilize the bandwidth of a fiber by dividing it into several channels, which are used to support multiple concurrent transmissions. However, each channel still has a transmission rate on the order of several tens of Gigabits per second. Currently, most of the applications have bandwidth requirements that are far less than that provided by a single wavelength channel. Therefore, in order to efficiently utilize resources, multiple subwavelength traffic requirements are multiplexed onto a single wavelength channel. This technique is known as traffic grooming on WDM networks. Since traffic grooming can be implemented in several ways, it must be done in a way that reduces the cost of the network. Earlier, research on traffic grooming emphasized the reduction of the number of required wavelength channels. However, with the realization that the dominant cost factor is not the number of wavelength channels, but is rather the number of higher layer electronic components, such as SONET Add/Drop Multiplexers (ADMs), Multiservice Provisioning Platforms (MSPPs), IP router ports, or MPLS Label Switching Router (LSR) ports, most of the research now is focused on reducing the total number of the electronic components in the network while accommodating a given set of the traffic.

The traffic grooming problem with arbitrary traffic, even for simple topologies, has been proved to be NP-Complete [1]. The traffic grooming problem is regarded to be even harder than the combined virtual topology design and Routing and Wavelength Assignment (RWA) problem [2]. When considering topologies in which the traditional RWA problem can be solved within polynomial time, the traffic grooming problem turns out to be NP-Complete [3]. To reduce the complexity of the problem, many relaxations have been considered in the literature. For example, most studies allow the traffic between a source-destination pair to be (vertically) split, and each component is routed separately, a condition known as bifurcation. Allowing bifurcation provides flexibility in traffic allocation and thus simplifies the problem, leading to a possible reduction in the number of wavelength channels as well as the number of higher layer light-path-terminating electronic equipment. However, bifurcation increases the complexity and the cost of traffic reassembly, and may also introduce delay jitter at the application layer. Many applications, especially real-time applications, require that their traffic be kept intact, i.e., without demultiplexing at the source, independent switching at intermediate nodes, and multiplexing at the destination. In this paper we have considered the case in which traffic bifurcation is not allowed, either among different wavelengths or even on the same wavelength but different physical routes[1]. Note that a specific traffic demand, as a whole, still can traverse different wavelengths while making its way to the destination (which we refer to as horizontal splitting). The only restriction here is that the whole traffic demand needs to be intact on each wavelength it traverses.

Most studies on traffic grooming in WDM networks have exclusively dealt with unicast traffic. However, it is expected that a sizable portion of the traffic in future high performance networks will be multicast in nature, for example, multi-party conferencing, video distribution, network news distribution, collaborative processing, and web content distribution to proxies [5, 6]. Interestingly, most multicast service applications require only sub-wavelength capacity. For example, HDTV can work well with just 20 Mbps per channel, while a normal TV channel typically requires less than 2 Mbps per channel, when compressed using MPEG-2, as in digital television. Hence, many such connections can be groomed together onto a single wavelength. It is therefore important to design and dimension networks in order to be able to support traffic of the multicast type, while grooming subwavelength traffic demands. This will provide flexibility and facilitates the implementation of Layer 1 Virtual Private Network (VPN) that can support one-to-many communication.

In this paper, we consider the optimal design and provisioning of WDM networks for multicast traffic grooming. We develop a unified framework to solve different practical scenarios of multicast traffic grooming. Each scenario is solved optimally, and approximate solutions are obtained. Optimal solutions are designed by exploiting the specifics of the problems to formulate Mixed Integer Linear Programs (MILPs). Approximate solutions are also obtained by using heuristic approaches. Specifically, we solve the generic multicast problem in which, given a set of multicast demands and all destination nodes of a multicast session requiring the same amount of traffic, all demands need to be accommodated. To lower the cost of the network, dominant cost factors are minimized. In this paper, we minimize the higher layer electronic line terminating equipment (LTs), be it ADMs, or IP or MPLS router ports, while simultaneously minimizing the total number of required wavelengths, such that the total network cost is minimized. Furthermore, we also solve two interesting and practical variants of the traditional multicast problem, namely, *multicasting with partial destination set reachability* and *multicasting with traffic thinning.* For both problems the destination set of each

[1]With Next Generation SONET (NGS) equipment and protocols [4], the Virtual Concatenation (VCAT) scheme supports traffic bifurcation. However, the traffic can be decomposed into a subset of a given set of virtual tributaries. Therefore, an exact formulation that allows bifurcation under VCAT and the given set of virtual tributaries will increase the complexity of the problem.

multicast session consists of two subsets. In multicasting with partial destination set reachability only one subset of each multicast session must be accommodated while the other subset can be accommodated only if this results in no additional cost. On the other hand, in traffic thinning both subsets of each multicast session must be accommodated. However, each subset has different bandwidth requirements. All three types of multicast traffic grooming will be further explained with the help of examples in Section V of this paper.

Multicasting with partial destination set reachability is useful when design and provisioning need to be done under a tight budget, and destinations of the multicast sessions can be classified as critical or non-critical. Destinations that are part of the critical set need to be accommodated in all cases, while destinations that are part of the non-critical set are accommodated only if this accommodation can be done without incurring any additional resources[2]. Traffic thinning, however, is helpful when a subset of the destinations of a session can be satisfied by a bandwidth less than that of other destinations of the session. In practice, this can happen when destinations of a multicast session have different Quality-of-Service (QoS) requirements, e.g., cases in which destinations are served with levels of multi-layer video coding [7]. Provisioning all destinations of a session with the maximum amount of bandwidth required by any of the destinations of the session could result in potential wastage of resources. The latter issue is also reconfirmed by our experiments in Section VIII of the paper. Hence, entertaining the destinations with different bandwidth levels, i.e., traffic thinning, can be helpful to reach all the destinations at an overall reduced network cost.

The contributions of this paper can be summarized as:

- It provides a unified frame work for optimal network dimensioning and channel provisioning for multicast traffic grooming, for arbitrary topologies and traffic.
- Non-bifurcation of traffic is assured.
- Network cost is minimized by minimizing multiple resources (i.e., the number of LTs and wavelengths) instead of a single resource.

The rest of the paper is organized as follows. In Section II, we provide an overview of the related work. In Section III we discuss possible implementation of multicast traffic grooming under first, and next generation SONET. In Section IV, we explain the network model used in this paper. In Section V, we explain the multicast traffic grooming problem and its variants with the help of an

[2]Without loss of generality, we here assume that resources are represented by the LTs and wavelength channels.

example. In Section VI we develop the MILPs for the multicast traffic grooming problems. In Section VII we develop heuristic techniques for multicast traffic grooming. In Section VIII, we present a few experimental results. Finally, Section IX concludes the paper.

II Related Work

As indicated earlier, the traffic grooming problem with arbitrary traffic, even for simple topologies, has been proved to be NP-Complete [1]. Recent survey papers [8, 9, 10] cover most of the related work in the field of traffic grooming on WDM networks.

At a broader level, the related work for the design and provisioning of WDM networks with traffic grooming can be categorized based on: (1) traffic patterns, e.g., uniform traffic [1, 12, 13] and arbitrary traffic [2, 14, 16, 17, 18], (2) network topologies, e.g., unidirectional rings [1, 2], bidirectional rings [13, 14, 15, 16, 17, 18], and random topologies [11, 24, 19, 20, 21], (3) solution approaches, e.g., heuristics [1, 16, 17, 18], ILPs [2, 11, 17, 24, 20], and bounds [1, 12, 13, 14, 15]

As we are interested in arbitrary mesh topologies, we will overview only relevant work. For work on regular topologies, the interested readers are referred to the surveys in references [9] and [10]. The authors in [11, 19, 20, 21] considered unicast traffic grooming problem on arbitrary mesh topologies. In [11], an ILP was presented that maximizes the network throughput on a mesh topology. Two heuristics were also presented that try to maximize the one hop traffic and the utilization of the lightpaths, respectively. In [19], a few simplifying assumptions are made, which include categorizing the nodes into two types, translucent nodes, i.e., nodes where traffic can originate or terminate, and transparent, i.e., nodes where traffic cannot originate or terminate. It was further assumed that the set of the transparent nodes in the network forms an independent set. The problem was then reduced to the finding of a maximum weighted independent set such that the cost, which is the number of the LTs, be minimized. A heuristic was then presented to determine such a maximum weighted independent set. In [20] also, a few simplifying assumptions were made. For example, it was assumed that all virtual topologies are implementable on the given physical topology. An ILP was then presented to minimize the number of lightpaths, thus minimizing the number of transceivers. Also, a heuristic was presented for large sized problems. Finally, in [21] an iterative greedy algorithm was presented that aimed at reducing the total number of light paths while accommodating the given traffic set.

All the above mentioned studies assumed that the set of traffic demands is known a priori.

However, some studies assumed a dynamic traffic model in which traffic demands arrive one at a time, and decisions are taken without waiting for future traffic demands. These models are more suited to the operational mode of WDM networks, and hence factors like network utilization or blocking probability are optimized, e.g., [22, 23].

In reference [22], the authors developed an analytical model, using discrete-time Markov chains, to study the performance of different types of grooming nodes in a mesh network. They considered two types of nodes: nodes which can and cannot groom at the subwavelength level. From the analytical and simulation results, they concluded that the networks with grooming capable nodes experience lower levels of blocking probabilities than networks without grooming capable nodes. Reference [23] explored different routing strategies for dynamic traffic grooming. The authors considered Widest-Shortest path routing, Shortest-Widest path routing, Shortest-MaxSum path routing and MaxSum-Shortest path routing. The objective was to evaluate the performance of each routing strategy in terms of network utilization on a mesh topology. It was found that Shortest-Widest path routing offered better network utilization than the other routing strategies.

Multicast traffic grooming, however, is a new field that emerged in the authors' work in [24]. Besides this, reference [25], presented a heuristic approach for routing and wavelength assignment of multicast sessions with subwavelength traffic demands on a WDM ring network. The authors assumed that the traffic demand of each multicast session is one unit. They also allowed duplication of traffic streams at the optical as well as electronic level. However, their objective was to minimize the number of electronic signal duplication instances. For that purpose, they tried different routing strategies, and constructed circles by grouping non-overlapping arcs and combining those circles into different wavelengths. They showed through simulation that their approach leads to a reduced equipment cost than that obtained by routing each multicast session along its minimum spanning tree and then using a well known heuristic [16] for circle construction. Reference [26] solved the multicast traffic grooming in mesh network with sparse nodal light splitting capabilities, with an objective of minimizing the total number of wavelengths. The authors provided a heuristic solution by providing an ILP to compute the minimum number of wavelengths on a link, and then estimating the total number of wavelengths for the entire network. Their model assumed that the multicast routing trees are given. They also provided a heuristic approach for constructing multicast routing trees and a first-fit wavelength assignment algorithm to perform traffic grooming.

The work presented in this paper is different from other work in the literature in that it presents an optimal solution to the multicast traffic grooming problem. The problem is solved using the more realistic condition of non-bifurcating traffic. Also, this work explores the variants of generic multicast traffic grooming.

III Implementation

In order to support grooming of multicast traffic, and to provision multicast sessions with subwavelength requirements, a number of functionalities must be provided, namely:

1. Dropping the traffic at the end of a lightpath, if the lightpath terminates at a leaf node in the multicast tree.

2. Dropping a copy of the traffic at an LT which terminates a lightpath at a node to which a destination is connected, but this destination is not a leaf node in the multicast tree. This means that the traffic will be provisioned on one or more lightpaths starting at this node.

3. Duplication of traffic at an LT installed at a branching node.

The above functionalities can be applied selectively to different parts of the traffic stream, e.g., under the traffic thinning protocol.

In this section, we briefly discuss the implementation of these functionalities in First Generation SONET, using SONET ADMs, and in Next Generation SONET, using Data over SONET (DoS) switches.

III.1 Implementation in First Generation SONET

In order to support multicast traffic in general, data must be copied and duplicated using special hardware. This hardware can be electronic, optical or a combination of both. When multicast traffic grooming is involved, it may happen that at a node in the network, some of the tributaries aggregated on a certain wavelength need to be duplicated, while others need not be duplicated. In this case, it is natural to use an approach in which the optical signal is terminated at an LTE, and the tributaries are accessed. Tributaries that need to be copied, are then duplicated in the electronic domain. The LTE shown in Figure 1, together with the digital crossconnect (DXC) and buffer and duplication hardware, perform this operation. Notice that this figure shows an example in which the traffic to be duplicated is received on wavelength λ_2, and is then converted to the electronic domain. It is then routed to the duplication hardware by the DXC before being routed back through the DXCs to the two LTEs that transmit this traffic on λ_1 and and λ_2, and on two different outgoing

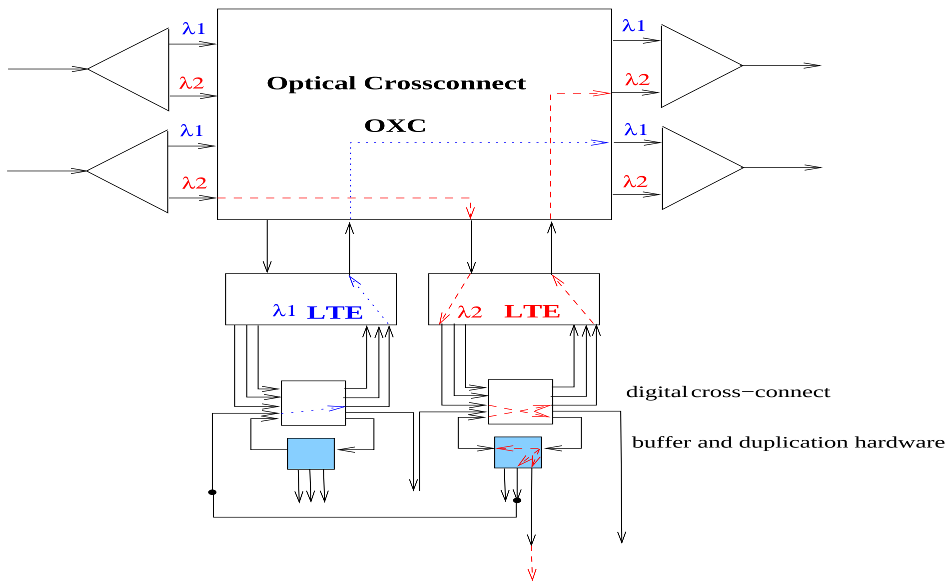

Figure 1: Multicast traffic duplication in the electronic domain

OXC ports. A copy of the traffic is also dropped at the node for delivery to attached end users, if needed, and as shown in the figure. It is to be noted that the traffic duplication hardware may also include buffering and regeneration circuitry. The traffic duplication hardware may not be required if the traffic is transmitted multiple times from the same digital cross-connect input port to multiple outputs. However, this will introduce an added delay, which may result in bandwidth wastage, or more complex synchronization.

III.2 Implementation in Next Generation SONET

In this section we discuss the implementation of the above functionalities in DoS switches used to implement Next Generation SONET (NGS) protocols. In particular, we briefly discuss how these functionalities can be implemented using the Generic Framing Procedure (GFP) and the Virtual Concatenation (VCAT) protocols [27, 28].

Our approach consists of two steps:

1. Using VCAT, the traffic substreams which will be forwarded selectively are assigned to different (virtual) tributaries, and the entire set of tributaries corresponding to the multicast stream are virtually concatenated. This approach lends itself to an easy way of dealing with different parts of the stream by allowing a simple filtering process.

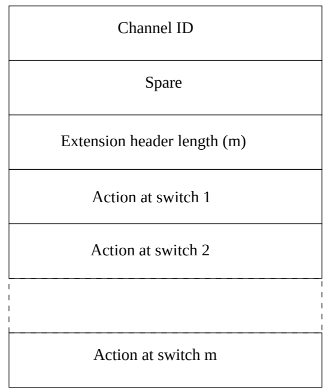

Figure 2: The multicast extension header.

2. Using a new extension header in GFP frame, that we refer to as a *multicast* extension header[3]. This header identifies how this stream should be processed at the DoS switch terminating the lightpath in order to support the multicasting operation. The extension header would consist of m fields, where $m \geq 1$, which are preceded by a header length field, i.e., the value of m (Figure 2). Each of the fields identifies how this traffic should be directed at the switch at the end of the lightpath:

   - A value of 0 means that a copy of traffic should be dropped at this node, i.e., a destination is connected at this node.
   - A non-zero value identifies the lightpath on which a copy of the traffic should be sent next.

   This means that at a multicast leaf node, the extension header should have a single field that has a value of 0. Since the payload of the GFP frame is decapsulated at the end of each lightpath, a new extension header is formatted at the source of each lightpath.

This extension header approach can be considered a generalization of the ring extension header [28], which was extended in reference [29] to the Multidrop GFP (GFP-MD) header. At each DoS switch, the GFP-MD frame is inspected, and an extension header in the GFP-MD frame indicates whether a copy of the GFP frame should be dropped at the current node, or the frame should be forwarded to the next node. The only way this extension header can support multicasting is by establishing a trail from the source to all destinations. However, this can be very inefficient. The

[3]The new extension header type can be defined, and included in the Extension Header Identifier in the GFP Payload Type Field.

approach proposed here is a more versatile and more efficient approach for supporting multicasting.

The actual duplication of traffic can be implemented in the DoS switch, and several commercially available framers like the Intel IXF19302 framer, implement a crossbar that can support multicasting.

IV The Model

In this section we will describe the network model that we consider for the multicast traffic grooming problems. The network model consists of three levels of abstraction: the physical, the lightpath and the connection, or session, levels, as shown in Figure 3:

1. The physical level corresponds to the network topology consisting of physical links between nodes, and is an input parameter. We assume that each physical link is composed of two fibers that are used to communicate in opposite directions, and that each fiber can support W wavelengths in one direction only.

2. The lightpath level represents the virtual topology, made of all-optical lightpaths, and is an output from the MILP. Each lightpath can span several physical links. Also, more than one lightpath may exist between a pair of nodes. If a pair of such lightpaths uses the same wavelength, then they must follow link disjoint physical routes. We assume that no optical wavelength conversion is available, and hence we maintain the wavelength continuity constraint at the lightpath level.

3. Connection level links between nodes represent the traffic demands. Each link at the connection level may span several lightpaths.

In Figure 3 nodes d_1, d_2 and d_3 are the members of the destination set of the multicast group originating at source node s. We consider the case in which branching for multicast traffic is implemented in the electronic domain; therefore, a multicast tree consists of several connection level links. Note that due to the implementation of the branching at electronic level, we do not require any special features at a node for duplication of the traffic, except that required for an electronic-multicast-capable node, as explained in the previous section.

Regarding notations, we will use s and d to represent source and destination of a connection, 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. Let K be the total number of sessions from all

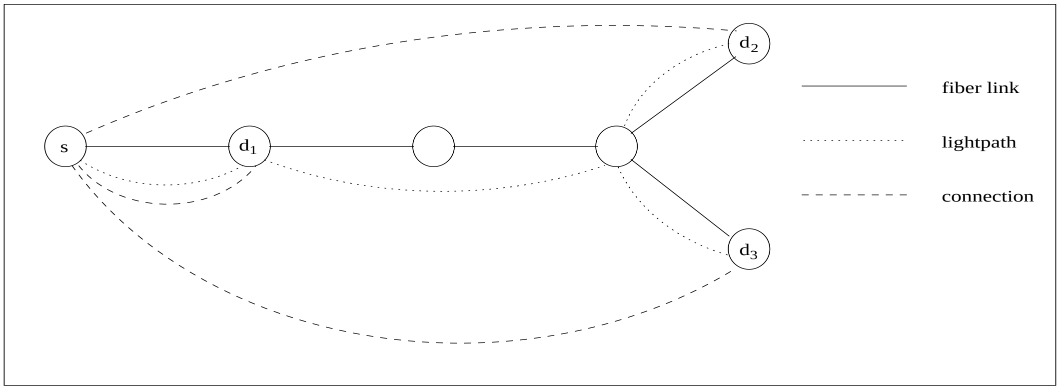

Figure 3: The three level network representation, source s, and destination set d_1, d_2 and d_3.

sources. Then, each connection c_a, where $1 \leq a \leq K$, corresponds to an ordered pair (s, k), where k represents the k^{th} (unicast or multicast) session originating from source s. Let D_{c_a} represent the destination set of session c_a, and d represent a destination in the destination set. For the two variants of generic multicast traffic grooming problem we assume that each destination set, D_{c_a}, is partitioned into two disjoint subsets, represented by $D^{'}_{c_a}$ and $D^{''}_{c_a}$, such that $D_{c_a} = D^{'}_{c_a} \bigcup D^{''}_{c_a}$ and $D^{'}_{c_a} \bigcap D^{''}_{c_a} = \emptyset$. Let $m^{'}_{c_a}$ represent the number of basic units of traffic required by each member of the destination set $D^{'}_{c_a}$, and $m^{''}_{c_a}$ represent the number of basic units of traffic required by each member of the destination set $D^{''}_{c_a}$. If $m^{'}_{c_a} = m^{''}_{c_a}$, then let m_{c_a} represent the number of basic units of traffic. The importance and use of these destination sets and subsets, as well as the rates, will be discussed in the next section.

V Problem Description

In this section we will describe the multicast traffic grooming problem and will show how it differs significantly from the multicast problem in an all-optical network.

In an all-optical network, multicasting is supported by developing a Steiner Minimum Tree (SMT), where the optimality criterion is usually the number of hops and the number of splitters. In the problem at hand, an SMT based on the above metrics may not be optimal as a certain link can be traversed multiple times, in order to save on the number of LTs. With the help of an example, we show that in case of multicast traffic grooming, using an SMT in terms of the number of hops, will not necessarily give an optimal solution in terms of the number of required LTs. Consider the six-node bidirectional ring shown in Figure 4. Let us assume that the capacity of each wavelength is 2 units and there exists 3 traffic sessions as follows.

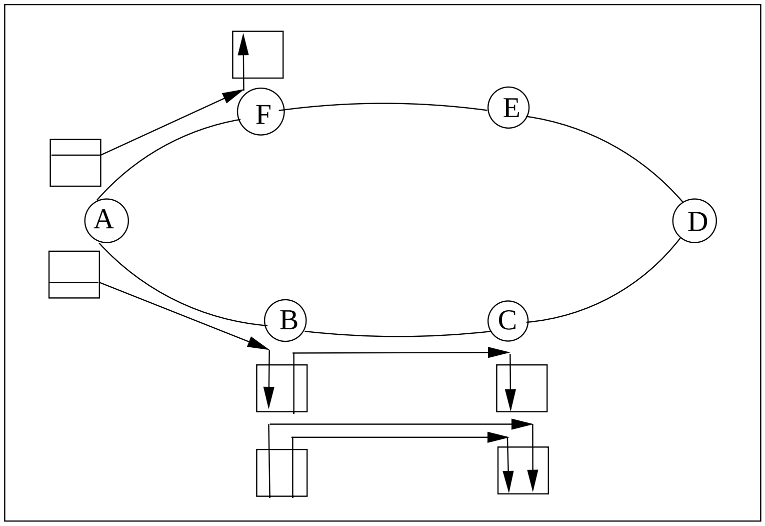

Figure 4: Routing using Steiner Minimum Tree.

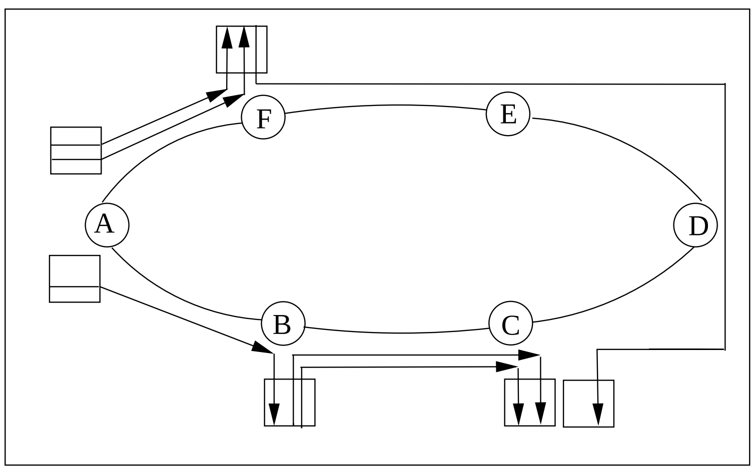

Figure 5: Routing in order to minimize the number of LTs.

Session 1: Source = A; Destination = {B, C}; Traffic demand = 1 unit;

Session 2: Source = B; Destination = {C}; Traffic demand = 2 unit;

Session 3: Source = A; Destination = {F}; Traffic demand = 1 unit;

Routing the demands using an SMT requires 7 LTs and two wavelengths, as shown in Figure 4. However, using the routing shown in Figure 5 costs just 6 LTs and one wavelength, which proves our claim. Hence, we need to take a totally different approach for designing WDM network to support multicast traffic grooming. One simple, but expensive technique to handle multicast traffic, especially in the absence of multicast enabled routers, is to treat every multicast demand as a set of unicast demands from the source to each of the destinations. However, it is obvious that such a policy will not lead to a minimum number of LTs in most cases.

A given traffic matrix typically consists of a number of unicast and multicast sessions[4], from a source node to a set of destination nodes. Using the notations defined earlier, we state that in a *generic* multicast traffic grooming problem all destinations in the destination set of some multicast session c_a, namely, $d \in D_{c_a}$, have the same bandwidth requirements. Also, all destinations, $d \in D_{c_a}, \forall c_a$ need to be served. Hence, no differential treatment is meted out to the destinations in a destination set in terms of delivery of traffic and the amount of the delivered traffic. However, in two variants of the generic multicast traffic grooming problem, for each multicast session, c_a, the destination set, D_{c_a}, consists of two subsets, D'_{c_a} and D''_{c_a}.

1. In the **partial destination set reachability problem**, we are required to accommodate all the destinations in subset D'_{c_a} only, while minimizing the network cost. The destinations in D''_{c_a} would be accommodated only if this action does not require any additional LTs or wavelength channels. The objective is also to maximize such an accommodation of $d \in D''_{c_a}, \forall c_a$, which does not increase the network cost.

2. In the **traffic thinning case**, a differential treatment is meted out to the destinations of the two subsets in terms of bandwidth requirements, with one of the two subsets receiving lower quality signal, i.e., lower bandwidth. More specifically, for each c_a, $d \in D''_{c_a}$ will be entertained with a bandwidth m''_{c_a}, while $d \in D'_{c_a}$ will be entertained with a bandwidth m'_{c_a}, such that $m''_{c_a} < m'_{c_a}$.

We will explain both variants with the help of an example. A 4-node network is shown in Figure 6. Each edge corresponds to tow fibers which are used for communication in the two opposite directions. Suppose that a single wavelength channel can accommodate two basic units of traffic, and there is a total of 4 sessions, one from each node. The details of the sessions are given below:

$c_1 : s = 0; D'_{c_1} = \{1\}; D''_{c_1} = \{2\}; m'_{c_1} = m''_{c_1} = \{1\};$

$c_2 : s = 1; D'_{c_2} = \{2\}; D''_{c_2} = \{\}; m'_{c_2} = \{2\};$

$c_3 : s = 2; D'_{c_3} = \{3\}; D''_{c_3} = \{0\}; m'_{c_3} = m''_{c_3} = \{1\};$

$c_4 : s = 3; D'_{c_4} = \{0\}; D''_{c_4} = \{\}; m'_{c_4} = \{2\};$

Figure 6 also shows the solution for the partial destination reachability set problem. Each square box represents an LT. Thus, the number of LTs required are 5, while a single wavelength channel

[4] In the model we make no distinction between unicast and multicast sessions. A unicast session is treated as a multicast session with one destination.

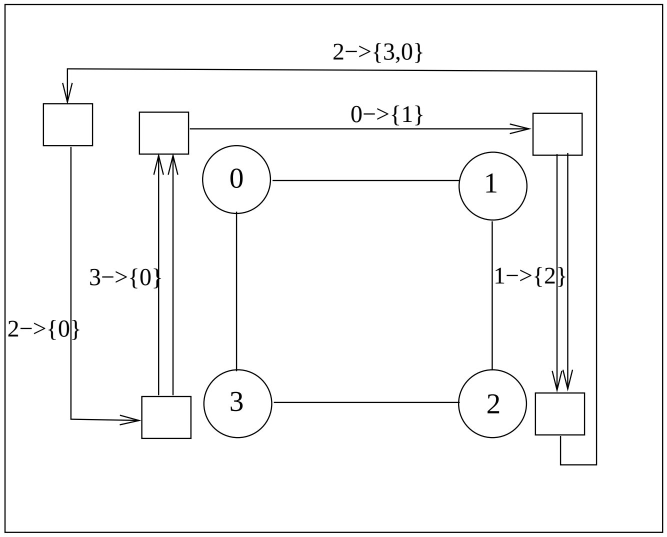

Figure 6: Partial destination set reachability problem on a 4-node network.

is enough to accommodate the requests for the destinations $d \in D^{'}_{c_a}$. Note that the solution accommodates the request to $D^{''}_{c_3} = \{0\}$ because this does not increase the network cost. However, accommodating the request to $D^{''}_{c_1} = \{2\}$ would have required an additional wavelength channel, and is therefore not served.

For the traffic thinning problem, both $d \in D^{'}_{c_a}$ and $D^{''}_{c_a}$ need to be accommodated. The above example is used again, however, to illustrate the point, the traffic demands are changed as follows.
$m^{'}_{c_1} = \{2\}; m^{''}_{c_1} = \{1\};$
$m^{'}_{c_2} = \{1\};$
$m^{'}_{c_3} = \{2\}; m^{''}_{c_3} = \{1\};$
$m^{'}_{c_4} = \{2\};$
The solution is shown in Figure 7, and needs 6 LTs and 2 wavelength channels. Note that if $D^{''}_{c_1} = \{2\}$ would have also been served with 2 units of traffic, as was the case for $D^{'}_{c_1} = \{1\}$, the number of required LTs would be 8. This shows that traffic thinning potentially can reduce the cost of the network, while still serving all requests.

VI Problem Formulation

In this section we will present the MILP for the generic multicast traffic grooming problem and its two variants: partial destination set reachability and traffic thinning problems. First, we will define all the variables used in the MILP.

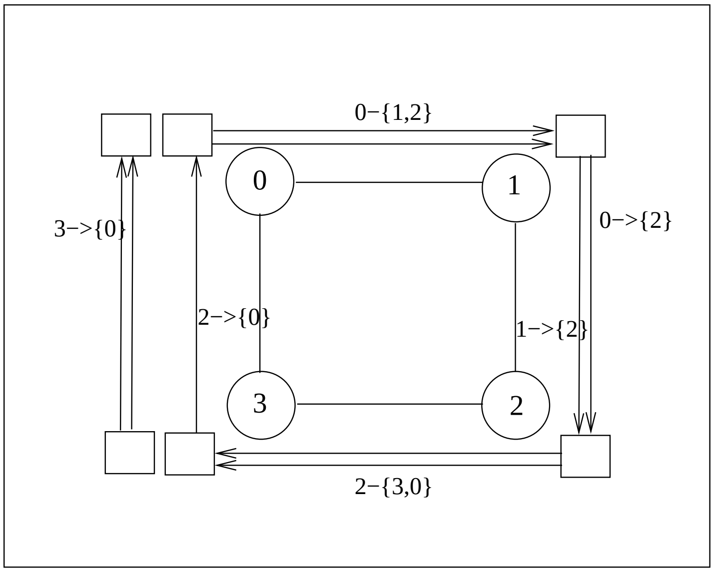

Figure 7: Traffic Thinning problem on a 4-node network.

VI.1 Definitions

The parameters involved can be divided into two classes: input parameters to the MILP, and parameters that are determined by the MILP, and hence are an output from the MILP. The definition of such parameters is given below.

- *Input parameters:*

  N: total number of nodes in the network

  W: maximum number of wavelengths per fiber [5]

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

  α: cost of an LT

  β cost of a wavelength

  γ: a scalar, smaller than both α and β, and is used for partial destination reachability

  Q: a very large integer number, (it suffices to set Q such that $Q \geq N^2 - N$)

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

  m_{c_a}: number of basic units of traffic required by each member of the destination set D_{c_a}

  m'_{c_a}: number of basic units of traffic required by each member of the destination set D'_{c_a}

  m''_{c_a}: number of basic units of traffic required by each member of the destination set D''_{c_a}

- *Variables of the MILP:*

  LT_n: number of LTs at node n

  ψ: highest index of wavelengths used over all fiber links

  y_w: a binary indicator; equals 1 if and only if wavelength w is used on at least one lightpath

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

  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, $L_{ij} = \sum_w L_{ij}^w$

  $Z_{ij}^{c_a,d}$: a real number between 0 and 1, which takes non-zero values if and only if connection c_a, destined to d, is employing a lightpath from i to j as an intermediate virtual link

  $M_{ij}^{c_a}$: a binary indicator; is 1 if and only if connection c_a is using a lightpath between nodes i and j to reach at least one destination d, where $d \in D_{c_a}$. This means that if for any $d \in D_{c_a}$, $Z_{ij}^{c_a,d}$ is greater than zero, then $M_{ij}^{c_a} = 1$.

  $G_{ij}^{c_a}$: a binary indicator; is 1 if and only if connection c_a is using lightpath between nodes i and j to reach at least one d, where $d \in D_{c_a}^{'}$.

  $X_{ij}^{c_a}$: a real number; capacity used by connection c_a on lightpath(s) between nodes i and j.

  $J_{ij}^{c_a,c_b}$: a binary indicator; is 1 if and only if connections 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 $X_{ij}^{c_b}$.

VI.2 Common Constraints

The common set of constraints for all the above mentioned problems are as follows:

- *Number of LTs:*

  The following two constraints ensure that for each originating or terminating lightpath at a node an LT is present:

$$LT_i \geq \sum_{w} \sum_{j,j \neq i} L_{ij}^w \quad \forall i \tag{1}$$

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

  The minimization of the cost of the LTs in the objective function (equation (23)) means that LT_i is equal to the maximum number of LTs required at node i.

- *Number of wavelengths:*

  The following constraints ensure that ψ will be the index of the highest numbered wavelength used on any fiber link in the network. Notice that the right hand side of equation (3) equals w if and only if w is used in the network. This is satisfied since equations (4) and (5) (and the minimization of the objective function) mean that y_w is 1 if and only if w is used in the network.

$$\psi \geq w * y_w \quad \forall w \tag{3}$$

$$y_w \geq \sum_i \sum_{j,j \neq i} L_{ij}^w / Q \quad \forall w \tag{4}$$

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

  We use the highest index of the number of wavelengths since usually all fibers are provisioned with the same set of channels.

- *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:

$$\sum_{m, for P_{mi}=1} F_{mi}^{ij,w} = \sum_{n, for P_{jn}=1} F_{jn}^{ij,w} = 0 \quad \forall i, j, w \tag{6}$$

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

$$\sum_{m, for P_{mj}=1} F_{mj}^{ij,w} = \sum_{n, for P_{in}=1} F_{in}^{ij,w} = L_{ij}^w \quad \forall i, j, w \tag{7}$$

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

$$\sum_{m, for P_{mx}=1} F_{mx}^{ij,w} = \sum_{n, for P_{xn}=1} F_{xn}^{ij,w} \quad \forall w, i, j, x; x \neq i, j \tag{8}$$

  Equations (7) and (8) together ensure that for each lightpath there exists a corresponding physical path, while maintaining wavelength continuity over the physical path. Equation (9) then ensures that wavelength w, on a fiber from node m to node n, accommodates at most one lightpath.

$$\sum_i \sum_{j, for P_{mn}=1} F_{mn}^{ij,w} \leq 1 \quad \forall m, n, w \tag{9}$$

It should be mentioned that our constraints involving the $F_{in}^{ij,w}$ variables are similar to those in [11]. However, due to the different nature of the problem, the rest of the constraints are different.

- *Multicast connection topology constraints:*

  The following two constraints ensure that for connection c_a destined for d, no traffic is coming in (going out) the source (destination), respectively

$$\sum_i Z_{is}^{c_a,d} = 0 \quad \forall c_a, d \in D_{c_a} \tag{10}$$

$$\sum_j Z_{dj}^{c_a,d} = 0 \quad \forall c_a, d \in D_{c_a} \tag{11}$$

  The following constraint preserves the continuity of the traffic of connection c_a on multiple lightpaths

$$\sum_{i,i\neq x} Z_{ix}^{c_a,d} = \sum_{j,j\neq x} Z_{xj}^{c_a,d} \quad \forall c_a, d \in D_{c_a}, x, (x \neq s, d) \tag{12}$$

  For a multicast session, delivery to two or more members of the session's destination set can be made by sending the traffic only once over a lightpath. Then, traffic duplication can take place after terminating the lightpath. The following constraints set the variable $M_{ij}^{c_a}$ to 1, if and only if at least one destination, $d \in D_{c_a}$, is reached through a lightpath between nodes i and j.

$$M_{ij}^{c_a} \geq \sum_{d\in D_{c_a}} Z_{ij}^{c_a,d}/Q \quad \forall d \in D_{c_a}, c_a, i, j \tag{13}$$

$$M_{ij}^{c_a} \leq \sum_{d\in D_{c_a}} Z_{ij}^{c_a,d} \quad \forall d \in D_{c_a}, c_a, i, j \tag{14}$$

  Notice that the above two sets of equations mean that $M_{ij}^{c_a}$ is just the disjunction of the binary variables $Z_{ij}^{c_a,d}$ for all values $d \in D_{c_a}$. Later, $M_{ij}^{c_a}$ will also be used to guarantee non-bifurcation.

  The following constraint ensures that the capacity, represented by $X_{ij}^{c_a}$, used by connection c_a on lightpath(s) between nodes i and j, does not exceed the physical capacity of the lightpaths on which the traffic is accommodated. The value of $X_{ij}^{c_a}$ itself will be computed for each problem separately.

$$\sum_{c_a=1}^{K} X_{ij}^{c_a} \leq L_{ij} * g \quad \forall i, j \tag{15}$$

- *Non-Bifurcation:*

  Bifurcation (vertical split) of a traffic demand can happen at three levels:

  1. Among lightpaths between different nodes,
  2. Among lightpaths between the same pair of nodes but on different wavelengths, and
  3. Among lightpaths between the same pair of nodes and on the same wavelength, but with each lightpath taking a different physical route.

  Routing constraints to be provided in equations (19), (20) together with constraint provided in (12) prevent only the first type of bifurcation. However, these constraints do not preclude bifurcation of the second and third type. Note that one of the objectives of the formulation in [11] was to ensure non-bifurcation of traffic. Although the formulation in [11] works perfectly for the case in which g is an integer multiple of any traffic demand, it does not guarantee non-bifurcation for arbitrary integer values of g. However, the formulation in [11] also guarantees non-bifurcation of the first type only for any value of g. Hence to obtain a complete non-bifurcated solution, one needs the following three steps, which apply to lightpaths between a pair of nodes i and j, $\forall i,\ j$:

  1. Assume that sessions are not sharing any lightpaths,
  2. Identify possible combinations of sessions sharing lightpaths so that the total number of lightpaths is equal to that obtained from constraint (15), and
  3. Choose the lightpath sharing combination that results in each lightpath capacity accommodating the entire bandwidth of all sessions sharing this lightpath, hence preventing bifurcation.

  The first step is partly satisfied by the constraint

$$J_{ij}^{c_a,c_b} \leq (M_{ij}^{c_a} + M_{ij}^{c_b})/2 \quad \forall c_a, c_b, i, j \tag{16}$$

  which establishes the necessary condition for the sessions c_a and c_b to share a lightpath between node pair (i,j) ($J_{ij}^{c_a}$ is 1). This condition is that both $M_{ij}^{c_a}$ and $M_{ij}^{c_b}$ are 1. For sessions not satisfying this condition, $J_{ij}^{c_a,c_b} = 0$. However, for other pairs of sessions, $J_{ij}^{c_a,c_b}$ may or may not be 1.

In the second step, we use the following equation to find feasible combinations of the remaining $J_{ij}^{c_a,c_b}$ variables.

$$L_{ij} = J_{ij}^{c_1,c_1} + \sum_{a=2}^{K} (J_{ij}^{c_a,c_a} - \vee_{b=1}^{a-1} J_{ij}^{c_b,c_a}) \quad \forall i,j \tag{17}$$

This equation counts lightpaths, such that each lightpath between nodes i and j is assigned to the least numbered session not using another lower numbered lightpath. To understand this equation, let session c_a use a lightpath between i and j. Therefore, $J_{ij}^{c_a,c_a} = 1$. If c_a is not sharing a lightpath with a lower numbered session (session c_1 is a special case), then the disjunction on the right hand side is 0, and and the lightpath corresponds to c_a, which is counted. However, if any other session, c_b, for $b < a$ is sharing this lightpath with c_a ($J_{ij}^{c_a,c_b} = 1$), then session c_a should not be counted since this lightpath has already been counted; hence, the subtraction of the disjunction on the right hand side.

Finally, the third step is achieved using the constraint

$$X_{ij}^{c_a} + \sum_{c_b, c_b \neq c_a} X_{ij}^{c_b} J_{ij}^{c_a,c_b} \leq g \quad \forall c_a, i, j \quad . \tag{18}$$

This constraint says that for all sessions sharing a lightpath, the lightpath capacity must accommodate the entire bandwidth of all such sessions, hence preventing bifurcation. The bandwidth used on a lightpath between i and j by session c_a, $X_{ij}^{c_a}$, will be determined according to the type of the multicast problem, as will be explained in the next three sections.

It is to be finally noted that the disjunction $\vee_{b=1}^{a-1} J_{ij}^{c_b,c_a}$ in equation (17) is implemented using linear constraints as follows. If $V = \vee_{b=1}^{a-1} J_{ij}^{c_b,c_a}$, then it is computed as:

$$V \leq \sum_{b=1}^{a-1} J_{ij}^{c_b,c_a} \quad \text{and} \quad V \geq \sum_{b=1}^{a-1} J_{ij}^{c_b,c_a} / (a-1)$$

VI.3 The Generic Multicast Problem

In addition to the above constraints, the set of the constraints required for the generic multicast problem is provided in this subsection. Under the generic multicast problem, all destinations must receive the same amount of traffic. Hence, we need the following two constraints to ensure delivery:

$$\sum_{j, j \neq s} Z_{sj}^{c_a,d} = 1 \quad \forall c_a, d \in D_{c_a} \tag{19}$$

$$\sum_{i, i \neq d} Z_{id}^{c_a,d} = 1 \quad \forall c_a, d \in D_{c_a} \tag{20}$$

which mean that for each session, c_a, the source must transmit to each destination, and each destination must receive from the source of the session. Together with constraint (12), flow conservation is achieved.

The bandwidth used by connection c_a on lightpath(s) between nodes i and j is computed as follows:

$$X_{ij}^{c_a} = m_{c_a} M_{ij}^{c_a} \tag{21}$$

This bandwidth is used in equation (18), which becomes:

$$m_{c_a} + \sum_{c_b, c_b \neq c_a} m_{c_b} J_{ij}^{c_a, c_b} \leq g \quad \forall c_a, i, j \tag{22}$$

Objective function for the generic multicast problem:

$$Minimize : \alpha * \sum_n LT_n + \beta * \psi \tag{23}$$

In the objective function, α and β represent the relative cost of an LT and a wavelength channel, respectively. In Section VIII, the choice of the values of α and β in the numerical examples will reflect actual cost factors of LTs and wavelength provisioning, respectively.

VI.4 Multicasting with Partial Destination Set Reachability

In addition to the common set of constraints, the set of the constraints required for the partial destination set reachability problem is provided in this subsection. In case of partial destination set reachability, the amount of the traffic delivered to the two classes of destination set, $D_{c_a}^{'}$ and $D_{c_a}^{''}$, is the same, i.e., $m_{c_a}^{'} = m_{c_a}^{''} = m_{c_a}$. Since source traffic delivery to the destinations, $d \in D_{c_a}^{''}$, need only be done if it entails no additional cost, the following set of constraints ensures the delivery to the destinations $d \in D_{c_a}^{'}$, while only providing the possibility of delivery to the destinations $d \in D_{c_a}^{''}$:

$$\sum_{j, j \neq s} Z_{sj}^{c_a, d} = 1 \quad \forall c_a, d \in D_{c_a}^{'} \tag{24}$$

$$\sum_{j, j \neq s} Z_{sj}^{c_a, d} \leq 1 \quad \forall c_a, d \in D_{c_a}^{''} \tag{25}$$

$$\sum_{i, i \neq d} Z_{id}^{c_a, d} = 1 \quad \forall c_a, d \in D_{c_a}^{'} \tag{26}$$

$$\sum_{i, i \neq d} Z_{id}^{c_a, d} \leq 1 \quad \forall c_a, d \in D_{c_a}^{''} \tag{27}$$

The value of $X_{ij}^{c_a}$ is the same as that of generic multicast, and equation (18) is also the same as equation (22), which is reproduced here for completeness:

$$m_{c_a} + \sum_{c_b, c_b \neq c_a} m_{c_b} J_{ij}^{c_a, c_b} \leq g \quad \forall c_a, i, j \tag{28}$$

Objective function for the partial destination set reachability problem:

$$Minimize : \alpha * \sum_{n} LT_n + \beta * \psi - \gamma \sum_{c_a} \sum_{d \in D''_{c_a}} \sum_{j, j \neq s} Z_{sj}^{c_a, d} \tag{29}$$

In addition to choosing α and β similar to the generic multicast problem, γ is chosen such that the accommodation of the destination set D''_{c_a} will not be done at the expense of any additional LT or wavelength. The following assignments capture these objectives:

$$\gamma = 1$$

$$\beta = \gamma + \sum_{c_a} |D''_{c_a}| * (N - 1)$$

$$\alpha = b * \beta,$$

where b is a positive integer number and is the ratio of the cost of an LT to the cost of a wavelength.

In equation (29), our objective is to maximize the number of served destinations which belong to the subset D''_{c_1}. This is why their sum is negated and added to the objective function which is minimized. However, serving any such destinations should not result in increasing the cost of the network in terms of wavelength channels or LTs. The relation between α, β and γ above, will make sure that an LT or a wavelength channel is more expensive than accommodating all destinations in all D''_{c_a}, for all c_a.

VI.5 Multicasting with Traffic Thinning

In this subsection, we will provide the constraints specifically needed for the multicasting under traffic thinning. In this case, $m''_{c_a} < m'_{c_a}$. Also, as all destinations in D'_{c_a} and D''_{c_a} must be accommodated, we need equation(19) and equation(20) which are reproduced in the following for the sake of completeness:

$$\sum_{j, j \neq s} Z_{sj}^{c_a, d} = 1 \quad \forall c_a, d \in D_{c_a} \tag{30}$$

$$\sum_{i, i \neq d} Z_{id}^{c_a, d} = 1 \quad \forall c_a, d \in D_{c_a} \tag{31}$$

Since $G_{ij}^{c_a}$ is 1 if and only if a lightpath between nodes i and j is used by connection c_a to reach at least one destination d, where $d \in D'_{c_a}$, then $G_{ij}^{c_a}$ is the disjunction of $Z_{ij}^{c_a,d}$ for all $d \in D'_{c_a}$ This is given by the following constraints:

$$G_{ij}^{c_a} \geq \sum_{d} Z_{ij}^{c_a,d}/Q \quad \forall d \in D'_{c_a}, c_a, i, j \tag{32}$$

$$G_{ij}^{c_a} \leq \sum_{d} Z_{ij}^{c_a,d} \quad \forall d \in D'_{c_a}, c_a, i, j \tag{33}$$

$X_{ij}^{c_a}$, which is the bandwidth used by connection c_a on lightpaths between nodes i and j, can be expressed as a function of $M_{ij}^{c_a}$, $G_{ij}^{c_a}$, m'_{c_a} and m''_{c_a} as follows:

$$X_{ij}^{c_a} = m''_{c_a} M_{ij}^{c_a} + (m'_{c_a} - m''_{c_a}) G_{ij}^{c_a} \tag{34}$$

Hence the value of $X_{ij}^{c_a}$ depends on a total of four combinations of the binary variables $M_{ij}^{c_a}$ and $G_{ij}^{c_a}$. However, the combination $M_{ij}^{c_a} = 0$ and $G_{ij}^{c_a} = 1$ cannot take place, and hence this reduce the possible number of combinations to three as follows:

- When $M_{ij}^{c_a} = 1$ and $G_{ij}^{c_a} = 0$, $X_{ij}^{c_a}$ will be m''_{c_a}.
- When $M_{ij}^{c_a} = 1$ and $G_{ij}^{c_a} = 1$, $X_{ij}^{c_a}$ will be m'_{c_a}.
- When $M_{ij}^{c_a} = 0$ and $G_{ij}^{c_a} = 0$, $X_{ij}^{c_a}$ will be zero.

Equation (18), which is given by

$$X_{ij}^{c_a} + \sum_{c_b, c_b \neq c_a} X_{ij}^{c_a} * J_{ij}^{c_a,c_b} \leq g \quad \forall c_a, i, j \tag{35}$$

is now nonlinear since $X_{ij}^{c_a}$ given by equation (34) is a function of $M_{ij}^{c_a}$ and $G_{ij}^{c_a}$. We therefore map the non-linear term, $X_{ij}^{c_a} * J_{ij}^{c_a,c_b}$, to a linear representation by defining $Y_{ij}^{c_a,c_b} = X_{ij}^{c_a} * J_{ij}^{c_a,c_b}$, and rewriting equation (35) as follows

$$X_{ij}^{c_a} + \sum_{c_b, c_b \neq c_a} Y_{ij}^{c_a,c_b} \leq g \quad \forall c_a, i, j \tag{36}$$

The product term, $Y_{ij}^{c_a,c_b}$, can now be computed using the following set of linear equations.

$$Y_{ij}^{c_a,c_b} \geq Q * J_{ij}^{c_a,c_b} - Q + X_{ij}^{c_b} \quad \forall c_a, c_b, i, j \tag{37}$$

$$Y_{ij}^{c_a,c_b} \leq X_{ij}^{c_b} \quad \forall c_a, c_b, i, j \tag{38}$$

$$Y_{ij}^{c_a,c_b} \leq Q * J_{ij}^{c_a,c_b} \quad \forall c_a, c_b, i, j \tag{39}$$

Notice that if $J_{ij}^{c_a,c_b} = 1$, then $\max(X_{ij}^{c_a}, 0) \leq Y_{ij}^{c_a,c_b} \leq \min(Q, X_{ij}^{c_a})$, which means that $Y_{ij}^{c_a,c_b} = X_{ij}^{c_a}$. However, when if $J_{ij}^{c_a,c_b} = 0$, $\max(0, -Q) \leq Y_{ij}^{c_a,c_b} \leq \min(0, X_{ij}^{c_a})$, hence resulting in $Y_{ij}^{c_a,c_b} = 0$.

Objective function for Traffic Thinning: The objective function is this case is the same as that of the generic multicast problem, and simply corresponds to a weighted combination of the total number of LTs, and the total number of wavelength channels used in the network, as given below.

$$Minimize : \alpha * \sum_{n} LT_n + \beta * \psi \tag{40}$$

VI.6 MILP Complexity

The complexity of the MILP, for all formulations, in terms of the number of variables is $O(N^4W + N^2K^2)$ and in terms of the number of constraints is $O(N^3W + N^2K^2)$, where K is the total number of multicast connections. If $K > NW$, then both the number of variables and the constraints are $O(N^2K^2)$, while if $K^2 < NW$, then the number of variables and the constraints are $O(N^4W)$ and $O(N^3W)$, respectively.

Note that the variable $Z_{ij}^{c_a,d}$ can be defined as an integer. However, experiments show that removing the integer constraint results in significant reduction in computation time. In general, by reducing the number of integer variables, the number of branching variables is reduced, which helps in reducing the complexity of the problem. We declared the $Z_{ij}^{c_a,s}$ variables as real numbers in the range [0,1]. However, the constraints imposed by equations (13) and (14) together with equations (19) and (20) for generic multicast traffic grooming, equations (24-27) for partial destination reachability, and equations (30) and (31) for traffic thinning, will essentially force the $Z_{ij}^{c_a,s}$ variables to take either 0 or 1 value[6].

Finally, we introduce two sets of modifications that help speed up the MILP, while not affecting the order of the computational complexity. Note that for the J variables, the following two equations hold, which help reduce the number of variables

$$J_{ij}^{c_a,c_a} = M_{ij}^{c_a} \quad \forall c_a, i, j \tag{41}$$

$$J_{ij}^{c_a,c_b} = J_{ij}^{c_b,c_a} \quad \forall c_a, c_b, i, j \tag{42}$$

Also, all those variables that sum to zero, e.g., in equation (10) and equation (11), can be simply removed while generating the constraints.

[6]Notice that if $Z_{ij}^{c_a,d}$ assumes any non-integer value in the range (0,1), then by equations (13) and (14), $M_{ij}^{c_a}$ cannot be either 0 or 1, which contradicts its definition as a binary variable.

VII A Heuristic Approach

In this section we will present a heuristic approach to solve the multicast traffic grooming problem. The basic idea is to incorporate the observations made from the results of the MILP. An examination of the MILP solutions for small sized examples reveals that many lightpaths are routed along non-shortest paths at their corresponding physical levels, e.g., see the example in Figure 14. Therefore, it makes sense to explore non-shortest path physical routes for lightpaths that will result in cost reduction. However, since the search space to explore all the possible physical routes is huge, we devise a two step heuristic solution approach to explore a subset of the search space. In the first step, an initial solution is constructed. This solution is constructed using two different approaches, and will be elaborated on below. In the second step, a systematic method is employed to explore many non-shortest routes, and accepting those routes that result in the maximum saving in terms of the *cost* of the network. Below, we present a description of the heuristic. Instead of presenting a different algorithm for the generic problem and each of its variants, we devise a generic algorithm that handles all of the above mentioned multicast traffic grooming problems.

First we will define a few terms:

- Let the term *CurrentNetwork* represent the state (or snapshot) of network, i.e., it captures the traffic demands which have been accommodated at that instant, and the set of LTs and wavelengths used to accommodate those traffic demands.

- Let the *cost* of the network be defined by:

$$cost(network) = \alpha \sum_{n} LT_n + \beta\psi \tag{43}$$

  We also use a variable, *CurrentCost*, to temporarily memorize the cost of the network.

- Each multicast session c_a is routed over a multicast tree. Let any such tree consist of a set of paths $p_{c_a,d}$, one for each destination $d \in D_{c_a}$, and $l_{c_a,d,i}$ represent a set of links of path $p_{c_a,d}$ from node i to destination d.

- Let $H_{c_a} = D'_{c_a}$ if the problem to be solved is multicasting with partial destination set reachability; otherwise, let $H_{c_a} = D_{c_a}$.

The pseudo-code of the algorithm is shown in Figure 11. The algorithm starts by finding an initial solution and then employs an iterative improvement algorithm. Figures 8 and 9 shows two different algorithms used for finding an initial solution[7]. In Figure 8, an initial solution is found by Algorithm I which constructs a Shortest Path Tree (SPT) for each request c_a. Note that the SPT is based on the physical topology. The SPT may consist of multiple hops, where each hop corresponds to a lightpath, which is routed over the shortest physical path. Each lightpath is then routed on the SPT using the first-fit wavelength, that is, the wavelength on the shortest route, that has available capacity to accommodate this lightpath. If a lightpath cannot be accommodated on existing wavelengths, then a new wavelength is added. In Figure 9, Algorithm II routes each traffic demand first on its shortest path and then employs Algorithm III (Figure 10) to search a set of alternate paths for the path with the least cost. For a specific traffic demand (c_a, d), the alternate paths are determined by excluding the physical links $l_{c_a,d,i}$ from the network, and incrementing i from the source node to the destination node on the shortest path $p_{c_a,d}$. The cost of accommodating the demand on each of the alternate paths using first-fit wavelength assignment is determined. If a cost reduction is possible, the demand is re-routed on the alternate path that provides the maximum reduction.

The iterative improvement starts with the current state of the network and then for each destination of each session (i.e., $\forall d \in c_a, \forall c_a$) finds a set of alternate paths and the cost of accommodating these paths using Algorithm III, hence incorporating an alternate path that provides the maximum cost reduction. Once all traffic demands have been considered, the same procedure is repeated until no further cost reduction is achieved. The rationale for repeating the procedure is that if in the first iteration a traffic demand i gets a chance to benefit from the re-routing of traffic demands j, for $j < i$, then in the next iteration traffic demands j, for $j < i$ can benefit from re-routing i, which has been re-routed in the first iteration. Obviously, traffic demand i will again get a chance to benefit in the second iteration, as it will be considered after the routing of j. This indeed increased the chances to take advantage of each other's routing positions, and this was verified by the experimental results.

Finally, if the problem at hand is the partial destination set reachability, then for each c_a and for each $d \in D''_{c_a}$ we call a slightly modified version of Algorithm III, that we refer to as Algorithm III$'$. The only difference is that Algorithm III$'$ uses the condition $BestSaving \geq 0$, while Algorithm III uses $BestSaving > 0$. Therefore, the accommodation of the optional traffic streams destined to $d \in D''_{c_a}$, will be done only if this does not increase the cost of the network.

[7]We tried the two approaches, and found that neither of the two algorithms performs consistently better than the other. Each of the two algorithms may produce better results in different scenarios.

```
ALGORITHM I - Shortest Path Trees
BEGIN
1. For each traffic session c_a
2.          Construct a Shortest Path Tree, SPT_{c_a};
3.          For each d ∈ H_{c_a}
4.                    Route traffic demand (c_a, d) on SPT_{c_a}, using first-fit wavelength;
END
```

Figure 8: Algorithm I; initail accommodation using shortest path trees.

```
ALGORITHM II - Accommodation by exploring alternate paths
BEGIN
1. For each traffic session c_a
2.          For each d ∈ H_{c_a}
3.                    Route traffic demand (c_a, d) on the shortest path, p_{c_a,d}, using first-fit wavelength;
4.                    ALGORITHM III(CurrentNetwork,p_{c_a,d});
END
```

Figure 9: Algorithm II; initail accommodation by exploring alternate paths.

VII.1 Complexity of the heuristic

The time complexity of Algorithm I is $O(N^2K \log N + N^2KW)$. The time complexity of Algorithm III is $O(N^2 \log N)$, which guides the time complexity of the Algorithm II to be $O(N^3K \log N)$. The WHILE loop in Multicast traffic grooming heuristic iterates a constant number of times (simulation results show 2-4 times). Therefore, if $W < N \log N$, the time complexity of the Multicast traffic grooming heuristic is $O(N^3K \log N)$.

VIII Experimental Results

In this section we will present the results of the MILP model and the heuristic for the multicast traffic grooming problems.

VIII.1 Optimal Network Design

Using the MILP, we conduct experiments on two different network topologies, namely, the six-node network shown in Figure 12, and the the 14-node NSF network shown in Figure 13. The traffic demands in all cases consist of integer multiples of OC-1 connections. The capacity of a wavelength is OC-48, hence $g = 48$.

```
Algorithm III - Alternate path exploration
INPUT: CurrentNetwork, p_{c_a,d}
BEGIN
1. CurrentCost ← cost(CurrentNetwork);
2. For each of the node i ∈ p_{c_a,d}
3.         TempNetwork_i ← CurrentNetwork;
4.         Remove the traffic demand (c_a, d) from l_{c_a,d,i};
5.         Find an alternate path from node i to d, such that new path does not include any of the physical
link ∈ l_{c_a,d,i}

6.         Accommodate the traffic demand (c_a, d) between node i and d over the alternate path using
first-fit wavelength;
7.         saving_i ← CurrentCost − cost(TempNetwork_i);
8. i* ← arg(min_i(saving_i));
9. BestSaving ← saving_{i*};
10. IF (BestSaving > 0) THEN
11.         CurrentNetwork ← TempNetwork_i
END
```

Figure 10: Algorithm III; alternate path exploration.

Since our objective is to reduce the network cost in terms of the LTs and wavelength channels, we used the following:

- Unless otherwise stated explicitly, it is assumed that the LT is an OC-48 SONET ADM, which costs about \$25,000.

- If a wavelength is used in the network, all fibers are provisioned with transceivers for such a wavelength. Therefore, using a 2.5 GHZ laser which costs about \$175, and photodectors which cost about \$75 each, a wavelength would cost \$250/fiber. This leads to a total cost per wavelength of

  $$\$250 \times \text{number of edges in the graph} \times 2,$$

  since each edge in the network graph corresponds to two fibers. This means that the cost per wavelength in the 6 node, 8 edge network in Figure 12 is \$4,000, while in the 14 node, 21 edge NSF network in Figure 13 it is \$10,500.

We first experiment with the network in Figure 12. A sample traffic that consists of a mix of multicast and unicast sessions is generated. For comparison purposes, we use a traffic matrix for generic multicasting, and then modify it for the partial destination set reachability and the traffic thinning problems. We divide each destination set, D_{c_a}, which is a part of the original traffic matrix, into two subsets, $D^{'}_{c_a}$ and $D^{''}_{c_a}$, such that $D^{'}_{c_a} \bigcup D^{''}_{c_a} = D_{c_a}$ and $D^{'}_{c_a} \bigcap D^{''}_{c_a} = \emptyset$. The generated traffic demands are shown in Table 1. For the generic and partial destination set reachability problems,

Algorithm Multicast Traffic Grooming
BEGIN //initial solution
1. Determine initial solution using Algorithm I and Algorithm II, and select the best of the two
//iterative improvement
2. $CurrentCost \leftarrow cost(CurrentNetwork)$;
3. $NewCost \leftarrow 0$;
4. WHILE ($NewCost < CurrentCost$) DO BEGIN
5. For each traffic session c_a
6. For each of the destination $d \in H_{c_a}$
7. ALGORITHM III(CurrentNetwork,$p_{c_a,d}$);
8. $NewCost \leftarrow$ least cost found so far; END WHILE
9. IF(partial destination set reachability problem) THEN
10. For each traffic session c_a
11. For each $d \in D^{''}_{c_a}$
12. ALGORITHM III$'$(CurrentNetwork,$p_{c_a,d}$) [with BestSaving ≥ 0];
END

Figure 11: Multicast traffic grooming heuristic.

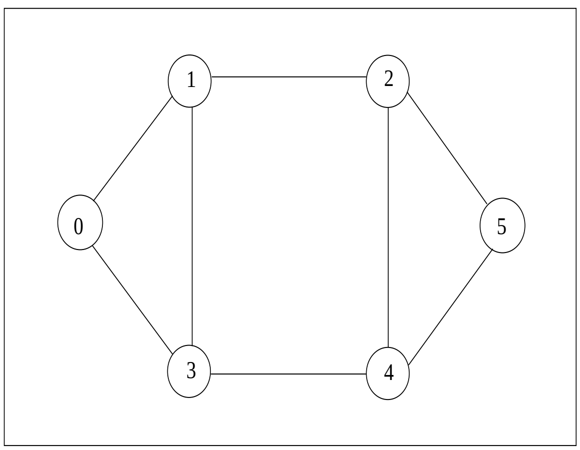

Figure 12: A six node network.

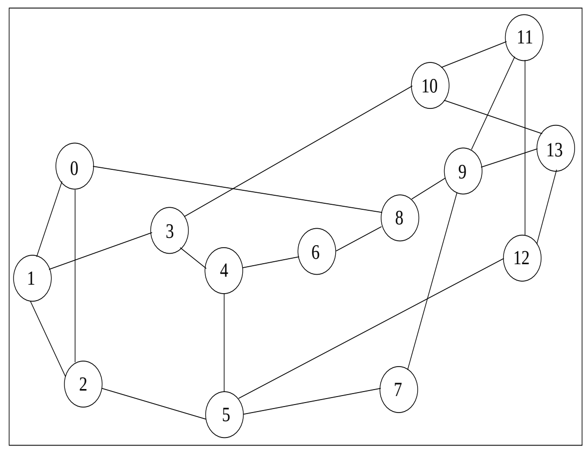

Figure 13: NSF Network topology.

the amount of traffic of each multicast connection is listed under the column $m^{'}_{c_a}$ of Table 1, while for the traffic thinning problem the traffic demands of the two destination subsets $D^{'}_{c_a}$ and $D^{''}_{c_a}$ are given by $m^{'}_{c_a}$ and $m^{''}_{c_a}$ of Table 1, respectively. The MILP is solved using the Cplex linear programming package [30]. The problem is run using four wavelengths for the generic, partial destination reachability, and traffic thinning problems. Table 2 summarizes the total number of LTs and the total number of wavelengths obtained for each problem. The last column of the Table 2 shows the number of LTs at each node.

As expected, the total number of LTs and wavelengths are lesser under both the partial destination set reachability and traffic thinning strategies than for the generic multicast traffic grooming

Table 1: Multicast traffic demands on the 6-node network in Figure 12

c_a	D_{c_a}	D'_{c_a}	D''_{c_a}	m'_{c_a}	m''_{c_a}
(0,1)	{5}	{5}	{}	12	-
(0,2)	{1,2,4}	{1,2}	{4}	24	18
(0,3)	{1,2,5}	{5}	{1,2}	9	3
(0,4)	{2,3,4,5}	{2,4}	{3,5}	36	24
(1,1)	{3}	{3}	{}	24	-
(1,2)	{3,5}	{3}	{5}	18	-
(1,3)	{1,2,5}	{2,5}	{1}	36	24
(2,1)	{1}	{1}	{}	9	-
(2,2)	{5}	{5}	{}	18	-
(2,3)	{0,3,5}	{0}	{3,5}	24	18
(2,4)	{0,1,3,4,5}	{1,3,4}	{0,5}	12	9
(3,1)	{0,1}	{0,1}	{}	3	-
(3,2)	{4,5}	{4}	{5}	48	36
(4,1)	{0}	{0}	{}	12	-
(4,2)	{0}	{0}	{}	24	-
(4,3)	{0,1,2}	{0,2}	{1}	36	24
(5,1)	{2}	{2}	{}	48	-

problem. The small difference between the number of required LTs for generic multicast traffic grooming and traffic thinning is primarily due to the fact that in the traffic thinning problem all the destinations need to be reached, and the difference in the amount of traffic between generic and traffic thinning problem is small. The difference in the number of LTs in case of partial destination set reachability problem, however, is profound. One of the possible explanations is that the MILP, after accommodating the must-accommodate sets, i.e., D'_{c_a}, has the flexibility to accommodate only those destinations which will require no additional LTs to be reached. There is a total of 37 destinations in the original traffic matrix, and a total of 13 destinations in optionally-accommodate destination sets, i.e., D''_{c_a}. Out of these 13 destinations, the MILP manages to accommodate 7 destinations. In other words, out of a total of 37 destinations, partial destination set reachability problem accommodated 31 destinations while using only 13 LTs as compared to 21 LTs used by the

Table 2: The number of wavelengths, total number of LTs, and location of LTs in the network of Figure 12 as generated by MILP for the generic, partial destination set reachability, and the traffic thinning problems, and using the traffic in Table 1.

	wavelengths	Total LTs	Location of LTs
Generic	3	21	LT_0=4, LT_1=2, LT_2=4, LT_3=3, LT_4=3, LT_5=5
Partial	2	13	LT_0=3, LT_1=1, LT_2=4, LT_3=2, LT_4=2, LT_5=1
Thinning	2	20	LT_0=3, LT_1=2, LT_2=5, LT_3=3, LT_4=3, LT_5=4

generic multicast traffic grooming problem. This translates into a saving of 38% of the LTs while accommodating 83% of the destinations. Moreover, the total number of wavelengths also reduces to 2, as compared to 3 wavelengths required for the generic multicast traffic grooming problem[8]. Thus,when working under a tight budget, dividing the destinations into critical and non-critical sets and following an approach similar to partial destination set reachability, can result in a design that meets the financial constraints.

In Figure 14, we show the routing of the multicast tree from source 0 to destination set {1,2,5}, under generic multicast traffic grooming. The figure shows the routing at the lightpath level, and also the corresponding physical links traversed by the lightpaths. Moreover, all the LTs required by the optimal solution of the generic multicast problem for the traffic matrix given in Table 1, are also shown. Note that the lightpaths between nodes 0 and 1, and between nodes 2 and 5, are not using shortest path at the physical level. In Table 3, we list the lightpaths used by each unicast or multicast connection, as determined by the optimal solution of the generic multicast problem for the traffic matrix given in Table 1. Similarly, in Table 4, we list the physical paths of all the lightpaths required by the optimal solution of the generic multicast problem for the traffic matrix given in Table 1. Also, the corresponding wavelengths used by the lightpaths are mentioned. A detailed inspection of the solutions produced by the MILPs reveals the following information.

- At the lightpath level, the multicast traffic is delivered either directly, from the source to the destination, or through another destination in the multicast destination set.

[8]Note that it may seem an anomaly that adding 6 more LTs to the partial destination reachability problem can result in reaching all destinations, and using a total of 19 LTs, as opposed to 21 LTs in solving the generic multicast problem. However, it is not anomaly, since it cannot happen. The reason is that with two wavelength channels there is no sufficient bandwidth to reach all remaining destinations. Hence, an additional wavelength is needed, on which additional LTs are required.

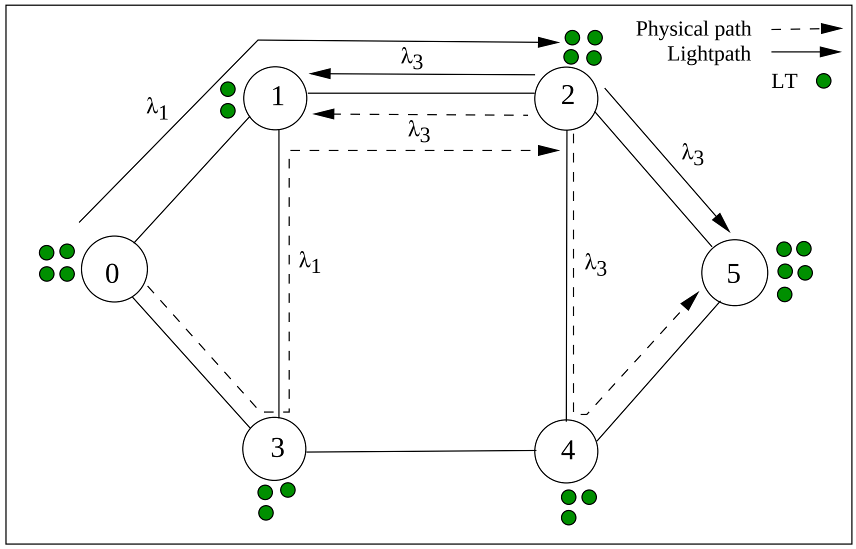

Figure 14: A part of the solution by the MILP for the generic multicast traffic grooming problem when the traffic matrix given in Table 1 is used.

- Instead of establishing multiple individual lightpaths between a source and each of its destinations, many lightpaths carry the (multicast) traffic to more than one destination in the same session, while simultaneously grooming the traffic to other sessions.
- An inspection of the physical paths corresponding to lightpaths revealed that not all the lightpaths are routed over shortest physical path. This shows that to obtain the minimum number of LTs, one needs to explore routes other than the shortest-path routes.

To run the optimal approach on a real network, we chose the 14-node, 21-edge NSF network shown in Figure 13. We selected a traffic matrix, shown in Table 5, such that every node is selected as a destination at least once. Also, out of 14 nodes, 10 nodes are acting as source nodes. However, each source node is establishing only a single unicast or multicast session. Moreover, out of a total of 10 sessions, 5 sessions are unicast sessions, while the rest of the sessions are multicast. We ran the MILP for the generic multicast traffic grooming problem with 4 wavelength channels. The MILP reduced the number of wavelengths to 1, and the corresponding number of required LTs to accommodate the whole matrix was found to be 16. An inspection of the solution, in terms of routing of lightpaths and their corresponding physical links, leads to the same conclusions drawn for the 6-node topology network.

Table 3: Lightpaths generated by the MILP for the generic multicast traffic grooming problem on the network topology shown in Figure 12, and given the traffic demands in Table 1 .

c_a	D_{c_a}	Lightpaths
(0,1)	{5}	0-5
(0,2)	{1,2,4}	0-2, 2-1, 2-4
(0,3)	{1,2,5}	0-2, 2-1, 2-5
(0,4)	{2,3,4,5}	0-2, 0-5, 5-3, 5-4
(1,1)	{3}	1-3
(1,2)	{3,5}	1-3, 3-5
(1,3)	{1,2,5}	1-0, 0-2, 0-5
(2,1)	{1}	2-1
(2,2)	{5}	2-5
(2,3)	{0,3,5}	2-4, 4-3, 3-5, 5-0
(2,4)	{0,1,3,4,5}	2-5, 5-0, 5-3, 5-4, 4-1
(3,1)	{0,1}	3-0, 0-2, 2-1
(3,2)	{4,5}	3-5, 5-4
(4,1)	{0}	4-0
(4,2)	{0}	4-3, 3-0
(4,3)	{0,1,2}	4-1, 4-0, 0-2
(5,1)	{2}	5-2

VIII.2 Verification of the Heuristic

To accommodate the traffic demands given in Table 1 on the network topology shown in Figure 12, the heuristic approaches for the generic multicast traffic grooming, partial destination set reachability, and traffic thinning problems requires 29, 19, and 27, LTs and 4, 3, and 3, wavelengths, respectively. The results are summarized in Table 6. In the same table we show the corresponding results obtained from the MILP. It is apparent that the number of LTs obtained by the heuristic solution are within 38%, 46%, and 35% of the optimal number of LTs for the generic multicast traffic grooming, partial destination set reachability, and traffic thinning problems, respectively. Moreover, for the partial destination set reachability problem, out of 13 destinations, which are members of the optionally-accommodate destination sets, i.e., D''_{c_a}, the heuristic solution accommodated 5 destinations. Hence, the heuristic managed to accommodate 29 destinations out of a total of 37 destinations, while using 19 LTs and 3 wavelengths. When compared to the solution generated by the heuristic for generic multicast traffic grooming problem, this translates into a saving of 34% of the LTs while accommodating 78% of the destinations.

In the same table, we also show the results obtained by applying the shortest path tree (SPT). The results include the number of wavelengths, number of LTs, and the number of reached destinations as determined by the initial SPT solution. As evident from the table, the iterative improvement helped reduce the number of required LTs by almost 6%, 9% and 10% for the generic multicast traffic grooming, partial destination set reachability, and traffic thinning problems, respectively. Notice that the number of wavelengths, and number of reached destinations, however, remained the same.

When we ran the heuristic on the traffic matrix shown in Table 5 for the NSF network shown in Figure 13, and for the generic multicast traffic grooming problem, it produced a solution that requires 22 LTs and 2 wavelengths. The heuristic solution, in terms of the number of LTs, is within 37% of the optimal value. It should be mentioned that the experimental results from other medium sized examples, show that the heuristic solutions are within 30-40% of the optimal values.

VIII.3 Large Network Design using the Heuristic

To study the performance of the heuristic with large traffic demands, we chose the NSF network topology and randomly generated the traffic for the generic multicast problem, with the following parameters for 2 scenarios:

- A number of sessions from each node is generated uniformly between 0 and 14.
- 50% of the sessions carry multicast traffic, while the remaining sessions carry unicast traffic.
- For each multicast session, the destination set size is uniformly distributed between 2 and 8 for the first scenario, and between 7 and 13 for the second scenario.
- The destinations are chosen randomly among all nodes, excluding the source, for both unicast and multicast sessions.
- The generated traffic, for both unicast and multicast sessions, is an OC-i stream, which is an integer multiple of OC-1, and i is uniformly chosen from the set {1,3,9,12,18,24,36,48}. These values represent the recommended rates for OC streams.

Table 7 shows the results of the experiments. For comparison purposes, experiments are conducted for the original traffic load, generated as mentioned above, and also by doubling the load. To double the load, we duplicate all the requests from each source. Moreover, the results from the heuristic are compared to the case when all multicast traffic is accommodated using multiple unicast connections[9]. In Table 7 the two scenarios correspond to different destination set sizes. As we mentioned earlier, for the first scenario the destination set size is uniformly distributed between 2 and 8, and for the second scenario it is uniformly distributed between 7 and 13. Hence the destination size in the second scenario is almost twice that of the first scenario. The total amount of traffic generated for the first scenario is equivalent to 504 OC-1 streams, while the total amount of traffic generated for the second scenario is equivalent to 924 OC-1 streams. From Table 7, it is evident that both the number of LTs and the number of wavelengths increases by almost 100% when the load is doubled (the doubled load for the first and second scenario is equivalent to 1008 OC-1 streams and 1848 OC-1 streams, respectively). This shows that our heuristic is grooming the traffic effectively. Moreover, Table 7 shows that when the multicast traffic is accommodated by employing multiple unicast connections, the number of required LTs is about 30% than that of our heuristic. Similarly, using multiple unicast connections one ends up using more than twice the number of wavelengths computed by our heuristic to accommodate the same traffic demands. Finally, the comparison between the first and the second scenarios shows that if the traffic is accommodated by constructing multicast trees and the size of the destination set is doubled, the number of required LTs increases by almost 60%, while the increment in the number of wavelengths is less than twice. However, if multiple unicast connections are employed to accommodate the traffic, then doubling the size of the destination set increases the number of LTs by almost 100% while the number of wavelengths increases by almost 75%. Thus, multicast traffic grooming through the introduced heuristic has a clear advantage over considering the multicast traffic as multiple unicast connections and then grooming them together. To show how this heuristic improves over the simple SPT, we have included results from the SPT in parentheses in the same table. The improvement introduced by the heuristic over SPT can reach 10% in terms of the number of LTs.

We also conducted experiments on the NSF network topology for partial destination set reachability and traffic thinning. We randomly generated traffic with the above mentioned parameters and with the destination set size uniformly distributed between 2 and 8. From this traffic matrix, we then generated the traffic matrices for the partial destination set reachability problem and for

[9]This is the technique used to carry multicast traffic in the Internet in the absence of multicast capable routers.

the traffic thinning problem using the following guidelines:

- The destination set is equally divided into two subsets, $D^{'}_{c_a}$ and $D^{''}_{c_a}$.
- For traffic thinning an immediate lower rate is selected from the set {1,3,9,12,18,24,36,48}, e.g., if $m^{'}_{c_a}$ is 12 units of traffic, then $m^{''}_{c_a}$ is selected to be 9 units of traffic.

The results obtained by running the MILP on the modified traffic matrices will essentially capture the reduction in the cost of the network due to either partial destination set reachability or traffic thinning. Traffic generation as outlined above generated a total of 350 connections. Out of these connections, 218 connections belong to sets $D^{'}_{c_a}$ (all unicast connections belong to these sets too) and 132 connections belong to sets $D^{''}_{c_a}$. As shown in Table 8, the proposed heuristic, for the partial destination set reachability problem, is able to accommodate 55 connections that belong to sets $D^{''}_{c_a}$ without an increase in the cost of the network. Hence, a total of 78% of the connections are accommodated while achieving close to a 31% lower network cost. In case of traffic thinning each connection is served while achieving a 10% lower network cost. Results from the SPT are also shown in the table, and in parentheses. It is evident that the use of the heuristic results in a noticeable reduction in the number of LTs and wavelength channels when compared to results from the SPT.

IX Conclusions

In this paper, we developed a unified framework to design and provision WDM networks to groom sub-wavelength multicast traffic. We introduced possible implementation for multicast traffic grooming in First and Next Generation SONET networks. We then solved the design and provisioning problem for different multicast grooming problems, which include generic multicast traffic grooming, partial destination reachability, and traffic thinning. The paper developed both optimal and approximate heuristic solutions. The optimal solutions exploited the specifics of the problem and introduced a mixed integer linear program formulation. The formulation is generic and also ensures the non-bifurcation of traffic. The heuristic technique was motivated by the observation that feeding a destination through another destination following a non-shortest path at the physical level can reduce the total number of required LTs. The heuristic first developed an initial feasible solution based on an SPT, and an alternate path selection procedure, and then chose the best of the two solution. The heuristic then improved it iteratively by exploring alternate paths in a systematic fashion.

A number of experiments were conducted using the exact and heuristic approaches. Experimental results verified that taking into account the multicast nature of the subwavelength connections can reduce the number of required LTs by substantial numbers as compared to the number of LTs required when each multicast connection is treated as consisting of multiple unicast connections. Results also showed that the iterative improvement of the proposed heuristic reduce the cost in terms of the number of LTs by up to 10% over routing using the shortest-path-tree. Moreover, it was shown that both partial destination set reachability and traffic thinning require lower network cost and hence can be used while dimensioning the network under tight budget constraints. In most cases, the heuristic approach produced results which were within 30% of the optimal solution.

This work can be extended in a number of directions. One direction involves implementing multicasting at both the optical layer and the electronic layer can provide some cost saving, as well as flexibility. This is especially true if the traffic granularities are diverse in the sense that some can be supported by electronic layer multicasting, and others can be supported in a more cost effective manner using optical layer multicasting. The problem in this case, especially the exact formulation, is much more difficult and involved, and different approaches to reduce its complexity need to be explored. For example, relaxing the integerality constraints can significantly reduce the complexity of the problem. Approaches similar to those in [31] and [32], as well as other approaches, can be employed. Moreover, since the control plane needs to handle multicasting at two different layers, more intelligence needs to be incorporated in the control plane, without unduly increasing its complexity. In another direction, exploring cost effective, and flexible implementation techniques of multicast traffic grooming is a task that is worth undertaking. Supporting the reverse multicasting, or many-to-one traffic grooming is another interesting problem that need to be addressed.

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] 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.

[3] R. Dutta, Shu Huang, and G. N. Rouskas, "On Optimal Traffic Grooming in Elemental Network Topologies," In *Proceedings of Opticomm'03*, Dallas, Texas, Vol. 5285, pp. 13-24, October 2003.

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

[5] B. Quinn and K. Almeroth, "IP Multicast Applications: Challenges and Solutions", Internet Engineering Task Force (IETF) Request for Comments (RFC) 3170, Sep. 2001.

[6] 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.

[7] H. M. Radha, M. van der Schaar, and Y. Chen, "The mpeg-4 fine-grained scalable video coding method for multimedia streaming over ip," *IEEE TRANSACTIONS ON MULTIMEDIA*, vol. 3, pp. 53–68, Mar. 2001.

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

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

[10] 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.

[11] 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.

[12] 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.

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

[14] X. -Y. Li, L. -W. Liu, P. -J. Wan, and O. Frieder, "Practical Traffic Grooming Scheme for Single-Hub SONET/WDM Rings," In *Proceedings of IEEE 25th Annual Conference on Local Computer Networks (LCN)*, Tampa, Florida, pp. 556-564, November 2000.

[15] 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.

[16] 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.

[17] J. Wang, W. Cho, V. Vermuri, and B. Mukherjee, "Improved Approaches for Cost-Effective Traffic Grooming in WDM Ring Networks: ILP Formulations and Single-hop and Multihop Connections," *IEEE Journal of Lightwave Technology*, Vol. 19, No. 11, pp. 1645-1653, November 2001.

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

[19] 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.

[20] 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.

[21] M. Brunato and R. Battiti, "A Multistart Randomized Greedy Algorithm for Traffic Grooming on Mesh Logical Topologies," In *Proceedings of Optical Network Design and Modelling*, February 2002.

[22] S. Thiagarajan and A. K. Somani, "Performance Analysis of WDM Optical Networks with Grooming Capabilities," In *Proceedings of SPIE Technical Conference on Terabit Optical Networking: Architecture, Control, and Management Issues*, Boston, November 2000.

[23] 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.

[24] A. E. Kamal and R. Ul-Mustafa, "Multicast Traffic Grooming in WDM Networks," In *Proceedings of Opticomm 03*, Dallas, Tx, October 2003.

[25] H. Madhyastha et al., "Grooming of Multicast Sessions in WDM Ring Networks", In *Proceedings of Opticomm 03*, Dallas, Tx, October 2003.

[26] A. Billah, B. Wang, and A. Awwal "Multicast Traffic Grooming in WDM Optical Mesh Networks", In *Proceedings of Globecom 03*, San Francisco, CA, Vol. 5, pp. 2755-2760, December 2003.

[27] P. Bonenfant and A. Rodriguez-Moral, "Generic Framing Procedure (GFP): The Catalyst for Efficient Data over Transport", IEEE Communications, Vol. 40, No. 5, May 2002, pp. 72–79.

[28] 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.

[29] K. Seppanen, "Multidrop Generic Framing Procedure (GFP-MD)", in the proceedings of *SSGRR 2002s - International Conference on Advances in Infrastructure for e-Business, e-Education, e-Science, And e-Medicine on the Internet*, 2002.

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

[31] A. E. Ozdaglar and D. P. Bertsekas, "Routing and wavelengthassignment in optical networks," *IEEE/ACM Transactions on Networking*, vol. 11, pp. 259–272, Apr. 2003.

[32] J.-Q. Hu, "Traffic grooming in wavelength-division-multiplexing ring networks: A linear programming solution," *Journal of Optical Networks*, vol. 1, pp. 397–408, Nov. 2002.

Table 4: Lightpaths and their corresponding physical paths generated by the MILP for the generic multicast traffic grooming problem for the traffic demands given in Table 1 on network topology shown in Figure 12.

Lightpaths	Physical links
0-2	0-3, 3-1, 1-2 (λ_1)
	0-3, 3-1, 1-2 (λ_2)
	0-1, 1-2 (λ_3)
0-5	0-3, 3-4, 4-2, 2-5 (λ_3)
1-0	1-3, 3-0 (λ_1)
1-3	1-3 (λ_3)
2-1	2-1 (λ_3)
2-4	2-4 (λ_1)
2-5	4-5, 2-4 (λ_3)
3-0	3-0 (λ_3)
3-5	3-4, 4-2, 2-5 (λ_1)
	3-4, 4-5 (λ_2)
4-0	4-3, 3-1, 1-0 (λ_3)
4-1	4-2, 2-1 (λ_2)
4-3	4-3 (λ_1)
5-0	5-2, 2-1, 1-0 (λ_1)
5-2	5-2 (λ_2)
5-3	5-4, 4-3 (λ_2)
5-4	5-4 (λ_1)
	5-4 (λ_3)

Table 5: Multicast traffic demands on the NSF network.

c_a	D_{c_a}	m_{c_a}
(0,1)	{2}	12
(1,1)	{5,9}	24
(3,1)	{8}	36
(4,1)	{0,6}	24
(6,1)	{1}	36
(7,1)	{2,11,12,13}	18
(9,1)	{10}	24
(10,1)	{8,11}	24
(12,1)	{4}	3
(13,1)	{3,7}	12

Table 6: Comparison of the heuristic under generic, partial destination set reachability, and the traffic thinning problems, to SPT and optimal results, using the traffic matrix shown in Table 1; H = heuristic, S = SPT, and O = Optimal.

	Number of wavelengths			Number of LTs			Reached destinations		
	H	S	O	H	S	O	H	S	O
Generic	4	4	3	29	31	21	37	37	37
Partial	3	3	2	19	21	13	29	29	31
Thinning	3	3	2	27	30	20	37	37	37

Table 7: The number of LTs and wavelengths, obtained by the generic traffic grooming heuristic and by accommodating the traffic demands using multiple unicast connections; results from the SPT are in paranetheses.

		Original load	Doubled load	Multiple unicasts with original load	Multiple unicasts with doubled load
First Scenario	LTs	187 (208)	362 (392)	239 (257)	475 (498)
	wavelengths	9 (9)	18 (17)	14 (14)	27 (27)
Second Scenario	LTs	277 (313)	555 (605)	412 (443)	823 (871)
	wavelengths	15 (13)	26 (25)	23 (26)	46 (52)

Table 8: The results of the heuristic for generic traffic grooming, partial destination set reachability, and traffic thinning problems on NSF network; results from the SPT are in parentheses.

	wavelengths	Total LTs	Percentage accommodated	Percentage saving
Generic	12 (15)	281 (300)	100	0
Partial	10 (12)	195 (210)	78	31
Thinning	12 (14)	265 (273)	100	6