

Algorithms for Multicast Traffic Grooming in WDM Mesh Networks*

Ahmed E. Kamal
Electrical and Computer Engineering Department
Iowa State University
Ames, IA 50011, U.S.A.
E-mail: kamal@iastate.edu

Abstract

Several of the new applications in high-performance networks are of the multicast traffic type. Since such networks employ an optical network infrastructure, and since most of these applications require subwavelength bandwidth, several streams are usually groomed on the same wavelength. This paper presents an account of recent advances in the design of optical networks for multicast traffic grooming in WDM mesh networks. The paper addresses network design and session provisioning under both static and dynamic multicast traffic. Under static traffic conditions, the objective is to accommodate a given set of multicast traffic demands, while minimizing the implementation cost. Optimal and heuristic solution techniques for mesh network topologies will be presented. Under dynamic traffic conditions, techniques for dynamic routing and session provisioning of multicast sessions, whose objective is to minimize session blocking probabilities will be explained. The paper will also present a number of open research issues.

I Introduction

During the last decade, Wavelength Division Multiplexing (WDM) Networks have emerged as an attractive architecture for backbone networks. WDM networks provide high bandwidth, on the order of tens of Gigabits per second per channel. With dense WDM (DWDM), an aggregate of several Terabits per second rates are achievable [1]. However, most applications have bandwidth requirements which are far less than that provided by a lightpath. It is therefore economical to use a lightpath to concurrently support multiple connections. The process of allocating subwavelength traffic demands to lightpaths such that the resources are shared is known as *traffic grooming*.

Client layers of the optical layer, e.g., SONET, IP, or MPLS, must terminate lightpaths using their own equipment, e.g., SONET ADMs, IP router ports, or MPLS LSR ports (to make the discussion general, we refer to such ports as a Line Terminating Equipment (LTE)). Such equipment is costly, and in designing networks for traffic grooming every effort should be made to intelligently support the traffic demands with the minimal number of LTEs. Optimal network design for traffic grooming is an emerging field, and progress has already been made in this venue [2].

*This research was supported in part by grants ANI-0087746 and CNS-0626741 from the National Science Foundation.

Several of the new and emerging applications using high-performance networks are of the multicast traffic type, where sessions include a single source transmitting to a number of destinations. Such networks are based on optical network infrastructures. Providing multicast service on optical networks was traditionally considered in two domains, namely, optical multicasting on wavelength routing networks, where a session uses a full wavelength capacity [3, 4], and multicasting on passive star optical couplers in broadcast and select networks [5]. In this paper, we consider multicasting on second generation optical networks, i.e., wavelength routed networks, with mesh topologies. Most applications require transmission rates which are much less than those provided by lightpaths. Therefore, traffic grooming techniques can be applied to those applications. However, most traffic grooming models in the literature, and network design for traffic grooming have only considered unicast traffic. This treatment of unicast traffic grooming may not be directly applicable to the treatment of multicast traffic, since there is a difference in traffic handling between multicast traffic, and unicast traffic. Under multicast traffic, traffic may have to be duplicated. Electronic equipment functionalities are therefore different, and network design, as well as session provisioning must take this into account. Therefore, this paper addresses the design and operation of optical networks to support traffic grooming of applications of the multicast type on WDM mesh networks. The paper will provide a survey of the techniques recently introduced in the literature for this purpose. These techniques will be illustrated using examples. It will also present some open research problems.

It should be pointed out that multicast traffic grooming can be regarded as a low overhead, reduced processing and cost effective approach for facilitating a Layer 1 virtual private network (VPN), with multicast communication support. It is well known that any to any (which includes one to many) communication on a virtual private network is difficult to support at higher layers, especially with a Layer 3 VPN [6]. Multicast traffic grooming provides this functionality easily. Providing multicast capable VPN services at layer 1 rather than on layers 2 or 3 helps the providers in aspects like traffic engineering, and multiprotocol, and multi-customer support.

This paper is organized as follows. In Section II we provide some background on traffic grooming, including multicast traffic grooming. In Section III, we consider node design for the support of multicast traffic grooming. Then, in Section IV, optimal and heuristic design and session provisioning techniques under static traffic will be presented. Also, techniques for session provisioning under dynamic multicast traffic grooming will be presented in Section V. Examples of the surveyed protocols will be presented. In Section VI we outline several open research problems in multicast traffic grooming, and finally, in Section VII we conclude this paper with a few remarks.

II Background

II.1 Traffic Grooming

Grooming is the term used to describe how different traffic streams are switched and packed into higher-speed streams. Nodes in an optical network are equipped with optical add/drop multiplexers (OADMs), or optical crossconnects (OXC), lightpaths carrying traffic streams that are sourced or terminated at the node are terminated at those nodes using LTEs. A lightpath termination at a node can be for the purpose of delivery to a destination, or for further processing and multiplexing. The terminated lightpaths are converted to electronic form and are each processed by an LTE, which adds and/or drops low-rate traffic tributaries from the aggregate channel stream and sends the channel back to the OADM/OXC in the optical form. The OADM then multiplexes the wavelength with other wavelengths in the outgoing fiber. Current research studies have concentrated on reducing the higher-layer electronic processing equipment (LTEs) whose cost dominates over the cost of the optical equipment (transceivers) in WDM networks, e.g., see [7, 2].

The concept of traffic grooming on optical channels can be implemented in a number of technologies. Traditionally, it has been applied to SONET/SDH channels. The channel granularity in this case is well defined. However, more recent implementations of traffic grooming apply to IP over WDM, and MPLS label switched paths (LSPs), where the reserved bandwidth per channel can be flexibly allocated. Moreover, with Next Generation SONET (NGS), Generic Framing Procedure (GFP), Virtual Concatenation (VCAT) and Link Capacity Adjustment Scheme (LCAS) provide mechanisms for flexibly grooming tributaries with different rates [8]. This provides more control on the efficiency of channel provisioning, and more flexibility in providing value-added services, such as the support of differentiated services. In this paper, however, we do not make any distinction between the different implementations, since we address the problem of traffic grooming at an abstract level, which can be applied to any of these technologies.

Excellent surveys of traffic grooming techniques have appeared in [7, 2].

II.2 Multicast traffic grooming

Support of multicast traffic has recently received significant attention, since several of the emerging high-performance network applications are of the multicast traffic type. Multicast applications belong to the class of multipoint applications which include [9]:

One-to-many or multicast. This type of service is very well known, and several applications

belong to this class. These include document distribution, on-demand video distribution, network news distribution, file distribution and caching.

Many-to-one which corresponds to data delivered from a group of users to a single destination.

Applications of this type include resource discovery, data collection at a central location, auctions, group polling, and accounting.

Many-to-many where several users interact together. Applications requiring this type of traffic

include the combination of one-to-many and many-to-one interactions cited above between a speaker and a group of users. They also include multimedia conferencing, synchronized resources, distance learning, distributed simulations, and collaborative processing.

Since most multipoint traffic carried by high-performance networks is of the multicast type, we concentrate in this paper on multicast traffic grooming.

Most of the above applications require transmission rates which are on the order of subwavelength rates. This makes the multiplexing of several application streams on the same wavelength a logical, and a cost effective approach. This multiplexing, however, must be done in a way that optimizes the use of resources, and minimizes the cost of the network. The allocation of resources should therefore take into account the session patterns, and their traffic requirements.

Based on the above, the importance of grooming of multicast traffic in WDM networks should be evident. Very few studies in this area exist and will be reviewed in this paper.

III Node Design for Multicast Traffic Grooming

In order to support multicast traffic in general, data must be copied and duplicated using special hardware, which may 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 in Figure 1 performs this operation. In this figure an example is shown in which the traffic to be duplicated is received on wavelength λ_2 , and is then duplicated in the electronic domain, before being routed back through the digital cross-connects to the two LTEs which transmit this traffic on λ_1 and λ_2 , and on two different outgoing OXC ports. A copy of the traffic may also be dropped at the considered node for delivery to attached end users, if needed. It is to be noted that

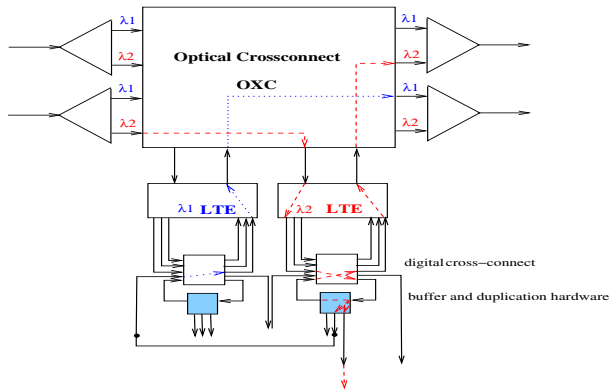


Figure 1: Multicast traffic duplication in the electronic domain

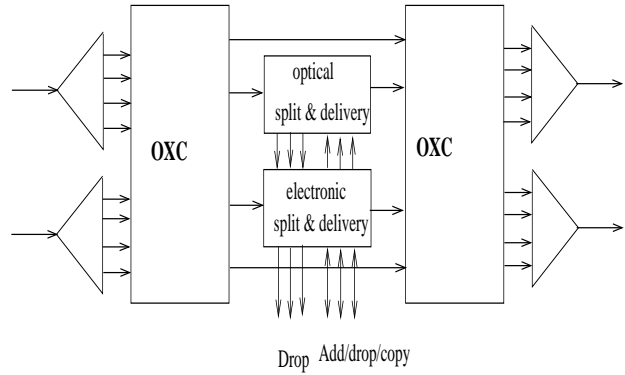


Figure 2: Node architecture employing optical and electronic duplication proposed in [10]

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, but this will introduce an added delay, which may result in bandwidth wastage, or more complex synchronization.

In references [10, 11], it was argued that implementing some of the traffic duplication in the optical domain might be less expensive, since the cost of (passive) optical splitters is considerably less than the cost of electronic LTEs. This is especially true if the duplicated traffic continues to use the same wavelength on different OXC output ports, and if the LTEs are only used for traffic duplication, and not for traffic dropping at a destination. However, electronic devices are still needed if traffic needs to be added to wavelength channels. Moreover, optical splitting can result in over usage of wavelength channels. Therefore, the authors introduced the node architecture shown in Figure 2 which implements both electronic, and optical duplication, based on need and cost. Such nodes are known as translucent nodes.

In this paper, we will present algorithms for multicast traffic grooming which assume both types of electronic (opaque) and translucent node architectures.

IV Static Multicast Traffic Grooming

Under the static traffic grooming problem, a physical network is given in terms of fiber links and wavelength channels, and a set of sub-wavelength traffic demands between the node pairs in the network is to be provisioned. The traffic is assigned to wavelengths such that the cost of the network, in terms of LTEs, is minimized. In some instances of the problem, the number of

wavelength channels is also a design parameter that should be minimized. In another version of the problem, the network resources are given, and the objective is to maximize the amount of traffic that can be accommodated. Lightpaths are therefore viewed as the transport channels in the backbone network, onto which traffic from multiple user applications is multiplexed.

We consider a network with N nodes, where the network graph can assume any arbitrary mesh topology. Each edge in the graph corresponds to a pair of fibers between a pair of nodes, which are used for transmission in two different directions. Each fiber has W wavelength channels. The capacity of each channel is expressed in terms of g , the grooming factor, which is the number of basic units of traffic which can be carried on a channel. For example, if the transmission rate on a wavelength channel is OC-48, or 2.4 Gbps, and the basic transmission unit is OC-3, or 150 Mbps, then $g = 16$. There are K multicast sessions in the network. Each session, a , where $1 \leq a \leq K$, is identified by the tuple $\langle s_a, D_a, m_a \rangle$. s_a is the source of the session, D_a is the set of destinations of the session, and m_a is the number of traffic units to be delivered from s_a to all destinations $d \in D_a$.

IV.1 Exact Approaches

We first consider that all traffic duplication is carried out in the electronic domain, using the node architecture shown in Figure 1. An exact approach for the static multicast traffic grooming problem was presented in [12] assuming the above node architecture. We present a simplified version of this approach, in the form of an Integer Linear Program (ILP), which minimizes the number of electronic equipment, LTEs. We present the ILP formulation qualitatively, while the formal ILP is available in [12]:

Objective function:

$$\text{Minimize : total number of LTEs} \tag{1}$$

Subject to:

Number of LTEs:

- Number of LTEs at node i is sufficient to terminate lightpaths starting or ending at $i \forall i$

Lightpath level constraints:

- A lightpath between two LTEs must start and end the LTEs

- A lightpath between two LTEs must not terminate at an intermediate OXC
- Number of lightpaths between two nodes i and j = Total number of lightpaths between i and j on all wavelengths

Multicast session topology constraints:

- Multicast sessions must start at a lightpath starting at the source
- Multicast sessions must terminate on a lightpath terminating at a destination
- Multicast sessions flow within the network at intermediate nodes is conserved

The above constraints are similar to the lightpath constraints, except that they apply to the multicast tree traffic, and the continuity is over the lightpaths, and not on the physical links.

Bandwidth constraints:

- For each lightpath:

The sum of bandwidths of all connections sharing the lightpath does not exceed the physical capacity of the wavelength

The above ILP is comprehensive, and produces the locations of the LTEs, as well as session provisioning, in terms of routes and wavelength assignment.

Other exact approaches were introduced in [13, 10, 11]. The model in [13] considered multicast traffic grooming on a mesh network, where the objective was to reduce the number of wavelength links. However, the model in [13] assumed that multicast sessions were routed in advance, which removes the routing problem from the formulation. In addition, wavelength conversion was used, which removes the constraints on wavelength continuity. Moreover, optical splitting can be used. The problem therefore reduced to a bin packing problem on one link at a time.

Using the translucent node architecture shown in Figure 2, the authors in [10] introduced an exact approach, in the form of an Integer Nonlinear Programming formulation, for the network design and session provisioning, using this node architecture. The objective function was to also minimize the cost of the electronic equipment. The formulation bears resemblance to the above one, except in one aspect, which is used to determine whether or not an add/drop port is required at a node. This involves terms, for each wavelength, and for each pair of links having a common node, which are the products over all sessions of other terms which are 1 if an add/drop is not

needed for each of these sessions. Using a similar architecture, the authors in [11] also introduced an ILP formulation for multicast traffic grooming on ring networks with one traffic unit per session.

The above multicast models assume that all destinations within a destination set have the same traffic requirements. However, different modes of operation of multicasting service include:

- Multicasting with *partial receiver reachability*, in which a subset of the receivers must receive the traffic originating from the source, while every attempt should be made to deliver the traffic to the remaining receivers, if possible, and without increasing the network cost. This corresponds to applications in which users are prioritized, and lower priority users can be accommodated only if this imposes no additional cost.
- Multicasting in which *traffic may be pruned, or thinned* while propagating downstream, depending on the receivers' needs. Applications of this type include multi-layer coded video streams, in which different receivers may require different stream qualities.

The above two cases were also considered in [12], in which the above ILP was modified in order to accommodate these two cases. The destination set for session a , D_a , is partitioned into two classes, D'_a , and D''_a , such that $D_a = D'_a \cup D''_a$, and $D'_a \cap D''_a = \phi$. All destinations $d' \in D'_a$ have a traffic requirement of m'_a , while destinations $d'' \in D''_a$ have a traffic requirement of m''_a .

For the partial destination reachability case, $m'_a = m''_a$, but D''_a can be reached only if this is not going to increase the network cost. The objective function in this case is the same as that in equation (1), except that it includes another negative term which is the sum of reachability indicators for all nodes in the D''_a , for all a . This maximizes the reachability of nodes in D''_a .

Also, for the traffic thinning case, $m'_a \neq m''_a$, and usually $m'_a > m''_a$. The objective function remains the same as in equation (1). However, the constraints must be revised by including a binary variable for each session, a on each lightpath, (i, j) , in order to select the correct traffic level, depending on the requirements of the downstream destinations. That is, the traffic delivered on a lightpath used by a session, is by default equal to the m''_a . If the lightpath is used to deliver a traffic level of m'_a , which will cause a binary variable to be equal to 1, a constant quantity equal to $m'_a - m''_a$ is added through multiplying it by the above binary variable.

IV.2 Heuristic Approaches

Unicast traffic grooming is known to be an NP-complete problem [2, 7]. Multicast traffic grooming is therefore an even harder problem. Therefore, the above exact approaches are only useful for use with small to, at best, medium sized networks, and with limited traffic levels. For large networks,

and for higher levels of traffic, other approximate approaches are therefore needed. In addition, the exact approach is still needed in order to assess the accuracy of the approximate approaches. Several heuristic approaches were proposed in the literature, and will be reviewed here.

Grooming with Lightpath Replacement (GLR)

Reference [12] introduced a heuristic approach based on some observations made from the exact solution of small sized examples. It was found that many sessions are routed along non-shortest paths, and several destinations were reached through lightpaths carrying traffic for such sessions. Those lightpaths were provisioned in a way that makes use of LTEs which were already in place. Therefore, the heuristic employed a strategy that is given below in Algorithm 1.

Algorithm 1: Grooming with Lightpath Replacement (GLR)

```

for each session do
  | Construct shortest path tree for the session;
while at least one session is not routed do
  | if SPT of session can be routed then
  | | route SPT of session and evaluate number of LTEs;
  | else
  | | introduce a new wavelength and route session;
  | for each link in the tree do
  | | remove link, reroute tree and reevaluate number of LTE;
  | | if number of LTEs is reduced then
  | | | use the new routing;
  | | else
  | | | keep the previous routing;
  | mark session as routed;

```

To illustrate the concept of exploring other trees which can reduce the number of LTEs, consider the example in Figure 3.(a). Two lightpaths were established to deliver traffic (which belong to different sessions) from nodes A to B and C. The grooming factor, g , is 2 units, and the bandwidth requirement per destination is 1 unit. This results in using 4 LTEs. However, when removing the link (A,C), C can be serviced using two branches from A to B, and then from B to C, hence requiring only 3 LTEs, as shown in Figure 3.(b). Since the search space can be very large, the step of exploring other routes and wavelengths is performed by inspecting one lightpath at a time, removing one link on the lightpath, and inspecting whether using existing LTEs can lead to reaching the affected destination while using fewer LTEs. The performance of GLR is usually within 30% of the optimal solution.

k -Shortest Path Trees (k -SPT)

Reference [10] introduced three heuristics for networks using a combination of optical and electronic duplication, but they still work with electronic duplication only.




Figure 3: An example to show the two steps of the GLR heuristic

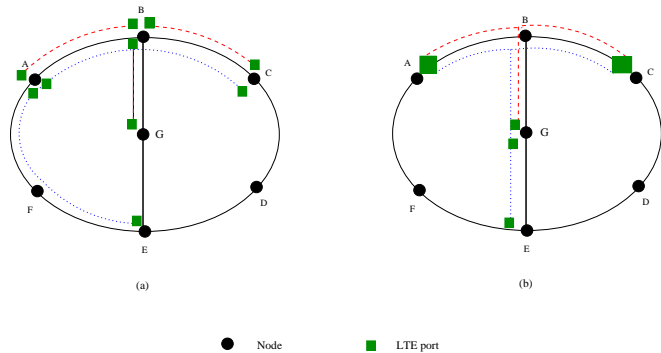


Figure 4: An example of the k-SPT heuristic

The first algorithm is the k-Shortest Path Tree (k-SPT) which works as follows:

Algorithm 2: k-Shortest Path Trees (k-SPT)

```

for each session do
  └ Construct k SPTs;
while at least one session is not rerouted do
  └ if no SPT can be routed on available wavelengths then
    └ introduce a new wavelength;
    route the SPT that requires the fewest number of LTEs;
    mark session using SPT as routed;
    remove remaining k-1 SPTs for routed session;

```

Note that the k-SPTs are constructed by constructing the SPT, and then removing the links from this SPT, one at a time, in order to obtain the remaining $k - 1$ SPTs. An example is shown in Figure 4, where there are two sessions: $\langle A, \{C, G\}, 1 \rangle$ and $\langle A, \{E, G\}, 1 \rangle$. The grooming factor, g , is 2. Figure 4.(a) shows routing under SPT for each session, which requires 10 LTE ports, which are shown as squares. However, in Figure 4.(b), optical splitting is used at node B, which reduces the number of LTEs to 5. Because of optical splitting at B, in Figure 4.(b), traffic from the second session is delivered to node C, which is not part of the destination set of this session.

Grooming with Rerouting Sessions (GRS)

In the second algorithm in [10], sessions are rerouted in order to find free resources. The algorithm consists of the following steps:

Algorithm 3: Grooming with Rerouting Sessions (GRS)

```
for each session do
  └ Construct shortest path tree;
select a session randomly, and route it;
for each unrouted session do
  └ if session can be routed with available wavelengths then
    └ route the session;
  └ else
    └ find bottleneck links, i.e., links which do not have available capacity;
    └ if number of bottleneck links  $\leq 2$ , and rerouting bottleneck traffic will accommodate session then
      └ reroute bottleneck traffic;
    └ else
      └ └ add a new wavelength and route session;
    └ mark session as routed;
```

The number of bottleneck links in the above algorithm is chosen to be at most two in order to limit the number of sessions that need to be rerouted.

Grooming by Computing Overlapped Trees (GCOT)

The third and last algorithm tries to pack sessions onto wavelengths in a greedy manner, and consists of the following steps:

Algorithm 4: Grooming by Computing Overlapped Trees (GCOT)

```
for each session do
  └ Construct shortest path tree;
Order SPTs in descending order of number of links, and put them in list  $L$ ;
route the first SPT in  $L$ ;
while  $L$  is not empty do
  └ find SPT  $T_i$  with the largest overlap degree in terms of the number of links with already
    └ routed SPTs;
  └ if  $T_i$  can be routed without adding a new wavelength then
    └ route it;
    └ remove  $T_i$  from  $L$ ;
    └ break;
  └ else
    └ └ choose next SPT from  $L$ , call it  $T_i$ , and repeat if statement;
  └ if no SPT could be routed then
    └ add a new wavelength;
    └ route SPT with largest overlap;
    └ mark SPT as routed;
```

It was shown in [10] that the k-SPT algorithm outperforms both the GRS and GCOT algorithm, both in terms of the number of LTEs and the number of wavelength channels. Also, as the value of k increases, the performance of k-SPT improves further, but of course at the expense of an added

complexity. By comparing the performance of the algorithms to the optimal solution obtained by solving the nonlinear program mentioned earlier, it was shown that the results of the algorithm are within 66% of the optimal. However, in most cases they are within 10 to 20% of the optimal.

The GLR algorithm described above is very much related to the k-SPT algorithm, since it also investigates several shortest paths. However, there are two differences between the two algorithms. First, under GLR, the number of alternate shortest paths is not fixed, while it is fixed under k-SPT. Second, under k-SPT, shortest path trees are constructed after removing a link, while under GLR, the original tree is augmented in order to reach the destinations which have been disconnected.

Multicast Traffic Grooming with Sparse Optical Splitters

Reference [13] considered a network where nodes may be equipped with optical splitters, which are used in the construction of multicast trees. Such nodes are known as Multicast Capable (MC) nodes. Otherwise, a node is called Multicast Incapable (MI). It is assumed that each MI node has at least one adjacent MC node, which it can use as virtual root, where the signal is split in the optical domain. It is also assumed that each node is equipped with a bank of wavelength converters, hence, wavelength continuity is not a requirement. The paper presented a two phase algorithm for provisioning multicast traffic sessions, namely, *Multicast Tree Construction (MTC)* and *First Fit Grooming Algorithm (FFGA)*. MTC constructs shortest path trees from each MC node, and allows an MI node to use an adjacent MC node as a virtual source, or root. The MI node, therefore, routes its multicast traffic to that adjacent MC node, and the MC node will split the traffic and route it to the destinations. The FFGA is a first fit greedy algorithm that allocates traffic streams on a link to the first available wavelength channel on that link. If no wavelength channel is available, a new wavelength channel is created. The algorithms are shown as algorithms 5 and 6, respectively.

Algorithm 5: Multicast Tree Construction (MTC)

```

for each MC node do
  Use Dijkstra's algorithm to find SPT in the network;
Sort all MCs in ascending order in terms of number of hops of their SPTs;
for each MI node do
  Choose the neighboring MC node with highest node connectivity, and smallest number
  of hops as the virtual root;

```

Algorithm 6: First Fit Grooming Algorithm (FFGA)

```
for each multicast traffic session do
  | Map traffic from the source onto appropriate links;
for each link do
  | for each traffic stream on the link do
  |   | if traffic stream can fit on an available channel then
  |   |   | Groom stream onto the first available channel;
  |   | else
  |   |   | Create a new wavelength channel;
  |   |   | Groom stream onto the new channel;
```

Table 1: Summary of static multicast traffic grooming heuristic algorithms

Algorithm	Duplication domain	Complexity
GLR [12]	electronic	$O(K \cdot M \cdot N^2 \log_2 N)$
k-SPT [10]	electronic-optical	$O(K \cdot k \cdot N^2 \log_2 N)$
GRS [10]	electronic-optical	$O(K^2 \cdot M \cdot N^2 \log_2 N)$
GCOT [10]	electronic-optical	$O(K^2 \cdot M \cdot W)$
MTC+FFGA [13]	optical	$O(K \cdot N^2 \log_2 N + K \cdot M \cdot W)$

V Dynamic Multicast Traffic Grooming

With the use of IP over WDM, MPLS over WDM, or with NGS, traffic sessions now tend to exhibit a dynamic nature, as opposed to the static nature assumed in the previous section. Sessions arrive according to a certain arrival process, and they are characterized by holding times, which are taken from a certain distribution. Since it is practically impossible to design optical networks such that they accommodate all such dynamic sessions, most efforts have concentrated on devising call acceptance, and session provisioning strategies that will try to reduce the session blocking probabilities.

A number of session provisioning strategies for multicast traffic were introduced in the literature, and will be presented below.

The Maximizing Minimum Freeload (MMFL) Algorithm

Reference [14] introduced a session provisioning strategy for dynamic multicast traffic, assuming translucent node architecture in Figure 2. The objective of this algorithm is to increase resource utilization, and to minimize the blocking probability for future arriving requests. This is done by using paths which will maximize the bandwidth capacity left after routing a multicast tree. In other words, the problem can be considered as a max-min problem, where the minimum left over bandwidth on all links is maximized.

The bandwidth capacity left after routing the tree is called the *freeload*, and the freeload on link (i, j) which has an available capacity $c_{i,j}$ on wavelength w after routing a session with bandwidth requirement B on this link, and this wavelength, is given by

$$\frac{c_{i,j} - B}{g} \quad (2)$$

where g is the grooming factor. The minimum freeload on all wavelengths is calculated assuming that the session is provisioned, and the routing that yields the maximum over these minima is used.

The algorithm in [14] is called *Maximizing Minimum Freeload* (MMFL). It is assumed that a freeload graph for each wavelength, w , is formed for each session. The current such graph is referred to as G_w , and the new graph, G'_w , is formed for each session by applying equation (2) on all links of the graph using the session requirement. After deciding on the routing and wavelength assignment of the session, only G_w on which the session is routed is updated by discounting the allocated link capacities. MMFL is given below, and is executed for every arriving multicast session:

Algorithm 7: Maximizing Minimum Freeload (MMFL)

```

for each arriving multicast session do
    for each wavelength  $w$  do
        generate freeload graph  $G'_w$  from  $G_w$ ;
        find paths from source to each destination on  $G'_w$ , such that average freeload per link
        is maximum;
        calculate the minimum freeload in the network (MFN) when routing the tree;
        select the tree on the wavelength that maximizes the MFN and route it;

```

Figure 5 shows an example of the MMFL in which there are two sessions, the first is $\langle A, \{C, G\}, 1 \rangle$, while the second is $\langle A, \{C, E\}, 1 \rangle$. Both sessions have bandwidth requirements of one unit, while $g = 2$. It is assumed that there is a single wavelength in the network. In Figure 5.(a), session 1 arrives, and the freeload factor for all fibers is calculated as $\frac{2-1}{2} = 0.5$. Session 1 is routed on the SPT shown by the dashed lines in Figure 5.(b). When session 2 arrives, in Figure 5.(b) the freeload factors for session 2 are calculated as shown in the figure. Fibers used by session 1 have a freeload factor of $\frac{1-1}{2} = 0$, while all other fibers have a freeload factor of 1. Session 2 is provisioned on the SPT shown by the dotted lines in Figure 5.(b). If there was another wavelength in the network, the freeload diagram for the second wavelength when session 2 arrives would be identical to that in Figure 5.(a), and session 2 would be provisioned similar to that in Figure 4.(b), but on the second wavelength.

It was shown in [14] that MMFL outperforms both fixed and adaptive SPT provisioning techniques in terms of call acceptance ratio and resource utilization, since MMFL attempts to uniformly

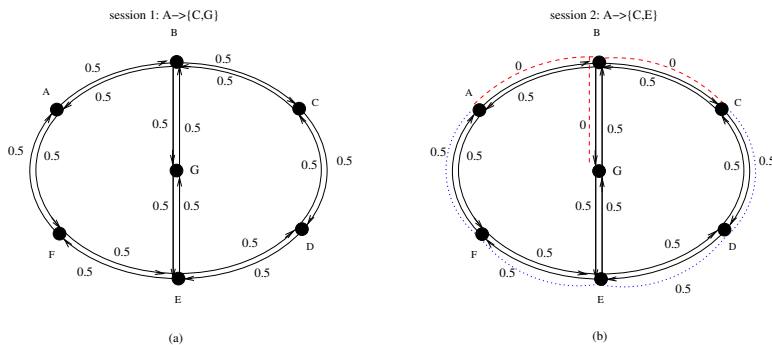


Figure 5: An example of the operation of the MMFL algorithm




Figure 6: Examples of the approaches used in sequential multicast traffic grooming

distribute the session loads over wavelengths, hence resulting in a better chance of call acceptance.

It is to be noted that this algorithm does not consider the effect of using optical splitting on accommodating sessions. For example, the route provisioning shown in Figure 4.(b) will not be selected by this algorithm.

Sequential and Hybrid Multicast Traffic Grooming

In references [15, 16] a node architecture similar to that in Figure 2 was assumed, and multicasting is always supported on a light-tree. All traffic on the same light-tree will be delivered to the same set of destinations on the light-tree, whether they are part of the destination set or not.

In order to support dynamic traffic, once a multicast session is to be accommodated, one of four traffic grooming approaches may be used. These approaches are described below, with the aid of the example given in Figure 6, assuming that there is one wavelength channel, and all sessions have a traffic requirement of 1 unit, while $g = 2$. It is also assumed that sessions 1, 2 and 3 are already established on the network.

1. *Sequential Single-Hop Provisioning*: in which an available light-tree, with the same source, and reaching the same set of destinations, with sufficient bandwidth is used to provision the

arriving session. For example, in Figure 6, when session 4 arrives, it can be accommodated on the light-tree used by session 1.

2. *Sequential Multi-Hop Provisioning*: which is used when there is a light-tree serving the set of destinations (called the *to-destinations light-tree* (TDLT)), but with traffic from a different source. In this case, an available single-hop lightpath from the source of the session to the source of the light-tree (called *from-source light-path* (FSLP)) is used to send the target session traffic to the TDLT. The TDLT and FSLP must have enough capacity to accommodate the arriving session. For example, when session 5 arrives, it can be provisioned using an FSLP from F to A (which is used to accommodate session 3), and then TDLT from A to C and G (which is used to accommodate session 1).
3. *Hybrid Provisioning*: which is similar to the above approaches, but new light-paths, or light-trees are provisioned when needed. For example, when session 6 arrives, the light-tree used by session 2 from G to {D,E,F} can be used to reach B and D. In addition, a light-path from G to B is needed, and is therefore provisioned. Alternatively, a light-path from either D or F can be also used to reach B.
4. *Non-Restricted Sequential Multi-Hop Provisioning*: which allows light-trees to carry traffic destined to only a subset of the destinations of the light-tree (this includes an empty subset, which corresponds to the case in which the light-tree is used as a first hop). For example, when session 7 arrives, the signal can be transmitted on the light-tree used by session 1. Since none of the destinations of session 7 are on this tree, at node G the signal is carried on the light-tree used by session 2 to reach D and F. Notice here that the signals also reach nodes C, and E, which are not in the destination set of session 7.

Using the above four approaches, the authors introduced four heuristics for provisioning multicast calls with subwavelength traffic rates. These are given below:

1. *Logical-First Sequential Routing (LFSR)*: in which an already provisioned single or multihop logical route is searched for, and if sufficient bandwidth is available, it is used to accommodate the arriving session. Otherwise, a new light-tree is established.
2. *Physical-First Sequential Routing (PFSR)*: in which the creation of a new light-tree from the source to the destination will be attempted first. Only if this fails, existing (single or multihop) logical routes will be inspected.

Algorithm 8: Logical-First Sequential Routing (LFSR)

```
while uninspected logical routes do
| Search for logical route from source to destination;
| if route has sufficient bandwidth then
| | provision the new session on the route;
| | exit;
Establish a new light-tree from source to destination;
```

Algorithm 9: Physical-First Sequential Routing (PFSR)

```
Establish a new light-tree from source to destination;
if light-tree is unsuccessful then
| while uninspected logical routes do
| | Search for logical route from source to destination;
| | if route has sufficient bandwidth then
| | | provision the new session on the route;
| | | exit;
```

3. *Logical-First Hybrid Routing (LFHR)*: in which the *LFSR* is invoked to find an existing lightpath to provision an FSLP to deliver the traffic from the source of the session to the root of a TDLT that reaches the destinations. If this fails, a new light-path is created and used to reach the source of the TDLT.

Algorithm 10: Logical-First Hybrid Routing (LFHR)

```
Use LFSR to find a lightpath with sufficient bandwidth to be used;
if LFSR is unsuccessful then
| Establish a new FSLP from source to to source of TDLT;
```

4. *Non-Restricted Logical-first Sequential Routing (NRLFSR)*: which is similar to LFSR, but a non-restricted multi-hop provisioning approach is used when LFSR fails.

Using a simulation study, and assuming that only one third of the calls are multicast, it was shown in [16] that among the restricted algorithms (LFSR, PFSR and LFHR), PFSR has the smallest blocking probability. The use of multihop session provisioning reduces the blocking probability even further, and provides a reduction in the blocking probability that can be very close to 20%. On the other hand, it was also shown that the non-restricted routing always outperforms constrained routing at light load. The last observation is expected since the routing problem has fewer constraints. However, as the load is increased, the restricted LFSR outperforms NRLFSR, since it is more conservative in using bandwidth to accommodate new sessions. When most of the sessions are multicast sessions, the non-restricted routing approach prevails.

Table 2: Summary of dynamic multicast traffic grooming heuristic algorithms

Algorithm	Duplication domain	Complexity
MMFL [14]	electronic-optical	$O(K \cdot N^2 \log_2 N)$
LFSR [16]	optical	$O(W \cdot N^2)$
PFSR [16]	optical	$O(W \cdot N^2)$
LFHR [16]	optical	$O(W \cdot N^2)$
NRLFSR [16]	optical	$O(W \cdot N^2)$

In Table 2 we provide a summary of the above protocols. Similar to Table 1, we also provide the duplication domain for which the algorithm was designed as well as the complexity of the algorithm.

VI Research Issues in Multicast Traffic Grooming

Multicast traffic grooming is a new field of research, and there are several open research problems in this field. Below, we give a list of some open research problems in this field.

- It has been indicated earlier that the traffic duplication process can be implemented using three different types of nodes, namely, electronic, optical, and translucent. A careful analysis of the cost factors involved in the node design and implementation needs to be carried out. With the use of Multiservice Provisioning Platform (MSPP) elements, a simple count of the number of ADMs is not an accurate indicator of the network cost. Since transport blades operating at different transmission rates have different cost factors, and as the cost of a transport blade increases with the transmission rate, using a translucent node may not necessarily result in a lower network cost, even if it results in the use of fewer ADMs, or transport blades. For example, although optical splitting may reduce the number of electronic ports, the ports may be operating at transmission rates which are higher than those needed with the use of electronic signal duplication. This is due to the fact that the electronic ports will be supporting traffic which could have been otherwise dropped, or not delivered to this port at all. For example, referring back to Figure 4.(b), because of the use of optical splitting at node B, traffic from the second session, viz., $\langle A, \{E, G\}, 1 \rangle$, is delivered to node C, which is not part of the destination set. Although this reduces the number of required LTE ports, the ports are of a higher transmission rate in this case.
- Implementation strategies of multicast traffic grooming in different technologies, such as IP, MPLS, traditional SONET and NGS need to be developed. Buffering, synchronization and multiplexing are issues that need to be taken into account.

- The blocking probability under unicast traffic has a simple definition: the probability that the entire session cannot be accommodated. However, with multicast traffic, the definition of the blocking probability may need to be revised, and may take different forms, which take both the user and the operator perspectives. In addition, the user perspective may itself depend on whether that user is a source or a receiver. The blocking probability may refer to either the ability to deliver data to all destinations in a session, or it may be measured in terms of the fraction of destinations which cannot be served.
- The static multicast traffic grooming problem is known to be NP-complete. Obtaining exact solutions to the network design and session provisioning problem is nevertheless important under this type of traffic for a number of reasons. First, it can be used with limited size networks, and second, it can be used to assess the accuracy of heuristic solutions. The exact formulation presented in Subsection IV.1 was a generic one. However, it needs to be extended in order to take a number of practical issues into account. For example, although constraints to prevent signal bifurcation were introduced in [12], there are cases in which a combination of bifurcation and non-bifurcation is used. This is especially important since with VCAT bifurcation is allowed at the sources only, while non-bifurcation within the network must be guaranteed.
- Another direction in obtaining exact solutions would be the derivation of exact solutions for special cases, which can reduce the complexity of the problem, hence facilitating the solution of larger networks. Such cases can include special cases of traffic, where for example, the traffic of all sessions can be the same, which can be regarded as the unit of traffic; or, special cases in which g can be a small number, such as 2. Other cases may include special cases of network topology, such as linear, or a tree topology. It is to be noted that in [11] algorithms for multicast traffic grooming on a ring topology were introduced.
- For static multicast traffic grooming, heuristics are certainly needed. Although a number of heuristics for multicast traffic grooming have been presented, usually these heuristics provide errors between 10 and 30% from the optimal solution. Therefore, more accurate heuristics are still needed. Such heuristics can be based on observations from the exact solutions. Or, they can be relaxations of exact solutions. Also, the use of combinatorial search heuristics, such as genetic algorithms, or simulated annealing should be investigated.
- In the case of dynamic multicast traffic grooming, open research problems can be classified

into two categories. The first one is the development of call admission algorithms which will reduce the blocking probabilities. With the new definitions of blocking probabilities mentioned above, different strategies may need to be developed for each definition of blocking probability. The second class of open research issues pertains to developing mathematical models to evaluate the session, or user blocking probability. This problem may not be easy due to the multicast nature of the traffic, even if blocking is defined in terms of the number of users who can be served within a session.

VII Conclusions

This paper has considered the grooming of multicast traffic in optical mesh networks. The paper has discussed the optimal design and session provisioning approaches presented for grooming of static multicast traffic sessions. In addition to the optimal approaches, approximate heuristic approaches were also considered. Grooming of dynamic multicast traffic was discussed, and the approaches presented in the literature were reviewed. A comparison between approaches that appeared in the literature was presented, and a set of open research problems was introduced.

References

- [1] P. Bayvel, "Future high-capacity optical telecommunication networks," *Philosophical Trans. of the Royal Soc. of London: Math., Phys. and Eng. Sciences*, vol. 358, pp. 303–329, Jan. 2000.
- [2] R. Dutta and G. N. Rouskas, "Traffic grooming in wdm networks: Past and future," *IEEE Network*, vol. 16, pp. 45–56, Nov./Dec. 2002.
- [3] G. Rouskas, "Optical layer multicast: Rationale, building blocks, and challenges," *IEEE Network*, vol. 17, no. 1, pp. 60–65, 2003.
- [4] A. Hamad, T. Wu, A. E. Kamal, and A. K. Somani, "Multicasting protocols for wavelength-routing networks," *Computer Networks*, to appear.
- [5] A. Hamad and A. E. Kamal, "A survey of multicasting protocols for broadcast-and-select single-hop networks," *IEEE Network*, vol. 16, pp. 36–48, July 2002.
- [6] T. Takeda *et al.*, "Layer 1 virtual private networks: Service concepts, architecture requirements, and related advances in standardization," *IEEE Communications*, vol. 42, pp. 132–138, June 2004.

- [7] E. Modiano, "Traffic grooming in wdm networks," *IEEE Communications*, pp. 124–129, July 2001.
- [8] D. Cavendish *et al.*, "New transport services for next-generation sonet/sdh systems," *IEEE Communications*, vol. 40, pp. 80–87, May 2002.
- [9] B. Quinn and K. Almeroth, "Ip multicast applications: Challenges and solutions." Internet Engineering Task Force (IETF) Request for Comments (RFC) 3170, Sept. 2001.
- [10] G. Chowdhary and C. S. R. Murthy, "Grooming of multicast sessions in wdm mesh networks," in *Workshop on Traffic Grooming*, 2004.
- [11] H. V. Madhyastha, G. V. Chowdhary, N. Srinivas, and C. S. R. Murthy, "Grooming of multicast sessions in metropolitan wdm ring networks," *Computer Networks*, vol. 49, no. 4, pp. 561–579, 2005.
- [12] R. Ul-Mustafa and A. E. Kamal, "Design and provisioning of wdm networks with multicast traffic grooming," *IEEE Journal on Selected Areas in Communications, Part II: Optical Communications and Networking*, vol. 24, pp. 37–53, Apr. 2006.
- [13] A. R. Billah, B. Wang, and A. A. Awwal, "Multicast traffic grooming in wdm optical mesh networks," in *Proceedings of the Conference on Global Communications (GLOBECOM)*, 2003.
- [14] G. Chowdhary and C. S. R. Murthy, "Dynamic multicast traffic engineering in wdm groomed mesh networks," in *Workshop on Traffic Grooming*, 2004.
- [15] A. Khalil, A. Hadjiantonis, G. Ellinas, and M. Ali, "On multicast traffic grooming in wdm networks," in *IEEE Int. Symp. on Computers and Communications*, pp. 282–287, 2004.
- [16] A. Khalil, A. Hadjiantonis, G. Ellinas, and M. Ali, "Sequential and hybrid grooming approaches for multicast traffic in wdm networks," in *Proceedings of the Conference on Global Communications (GLOBECOM)*, 2004.