

Dynamic Provisioning of Optical Networks with Many-to-Many Traffic Grooming

(Invited Paper)

Mohammad A. Saleh Ahmed E. Kamal

Department of Electrical and Computer Engineering, Iowa State University, Ames, IA 50011

Abstract—A large number of network applications today allow several users to interact together using the *many-to-many* service mode. In many-to-many communication, a session consists of a group of users (we refer to them as *members*), where each member transmits its traffic to all other members in the same group. In this paper, we address the problem of dynamic provisioning of optical WDM networks to support many-to-many traffic grooming. The objective is to minimize the overall blocking probability of arriving many-to-many sessions. We address the problem in both *non-splitting networks* where the nodes do not have optical splitting capabilities and in *splitting networks* where the nodes have optical splitting capabilities. In each of the two networks, we propose a number of dynamic provisioning heuristics and we provide extensive experiments to evaluate and compare their performance.

I. INTRODUCTION

Early internet applications such as TELNET and FTP are characterized as unicast or “one-to-one”. Several network applications today, however, are of the multipoint type. For example, video distribution and file distribution are examples of multicast or “one-to-many” applications, while resource discovery and data collection are examples of many-to-one or “inverse multicasting” applications. Recently, another set of network applications has emerged such as multimedia conferencing, e-science applications, distance learning, distributed simulations, and collaborative processing. In these applications, each of the participating entities both contributes and receives information to and from the other entities in the same communication session, and therefore are characterized as “many-to-many”. In many-to-many communication [1], a session consists of a group of users (we refer to them as *members*), where each member transmits its traffic to all other members in the same group (see Fig. 1).

In optical wavelength routing networks, using wavelength division multiplexing (WDM), it is feasible to have hundreds of wavelengths, each operating at 10 to 40 Gbps, per fiber. Bandwidth requirements of user sessions, however, are usually of sub-wavelength granularities. For example, an MPEG compressed HDTV channel requires less than 20 Mbps of bandwidth. In order to reduce this huge bandwidth gap, *traffic grooming* was introduced to allow a number of sessions with sub-wavelength granularities to share the bandwidth of a wavelength channel.

The main two resources in an optical WDM network are the wavelengths available on each fiber link and the higher




Fig. 1. A many-to-many session with members $\{A, B, C, D\}$ each with traffic denoted as a, b, c and d , respectively.

layer electronic ports (we refer to them as *transceivers*) available at each node in the network. A transceiver is needed for each initiation or termination of an optical channel. For example, a lightpath requires two transceivers while a light-tree with N endpoints requires N transceivers. Also, an optical channel occupies a whole wavelength on each fiber link it traverses, and due to the wavelength continuity constraint this wavelength must be the same on all the fiber links traversed (assuming no wavelength conversion capabilities).

In this paper, we address the problem of provisioning and grooming of dynamic many-to-many traffic in optical WDM mesh networks. Traffic grooming has mainly been considered for unicast traffic [2]-[4]. It has also been considered recently for multicast [5]-[7] and many-to-one [8]-[9] traffic types. All of the above references have dealt with the static traffic grooming problem where traffic demands are known in advance. In [10]-[11], the authors considered the dynamic unicast traffic grooming problem, while in [12]-[13], the authors considered the dynamic multicast traffic grooming problem in optical WDM networks. For a survey of advances in unicast and multicast traffic grooming, the reader is referred to [4], [7], respectively.

Many-to-many traffic grooming in optical WDM mesh networks is a new field of research that has been only considered in [14]-[16]. In [14]-[16], MILP formulations, heuristic solutions, and approximation algorithms were introduced for the static many-to-many traffic grooming problem in WDM mesh networks. This paper is the first to address the dynamic many-to-many traffic grooming problem in WDM mesh networks. We also address the problem in both *non-splitting networks* where the nodes do not have optical splitting capabilities, and in *splitting networks* where the nodes have optical splitting

capabilities. Note that in non-splitting networks only lightpaths are supported, while in splitting networks both lightpaths and light-trees are supported.

The rest of the paper is organized as follows. In Section I we formally define the dynamic many-to-many traffic grooming problem and introduce the assumptions and notations used in the paper. In Sections III and IV, we introduce a number of heuristic algorithms for the dynamic many-to-many traffic grooming problem in non-splitting and splitting networks respectively. In Section V, we conduct extensive experiments to evaluate and compare the performance of the proposed heuristics. In Section VI, we conclude the paper.

II. PROBLEM DESCRIPTION

We define the dynamic many-to-many traffic grooming problem as follows. Given:

- 1) An arbitrary optical WDM network topology, where the optical nodes either do not support optical splitting (non-splitting networks) or they fully support optical splitting (splitting networks).
- 2) The current network state represented by the set of optical channels (lightpaths and light-trees) that are currently established, the amount of bandwidth available on each of them, the set of free wavelengths on each fiber link, and the number of free transceivers at each node in the network.
- 3) An arriving many-to-many session request with an arbitrary subwavelength traffic demand (we assume all members in the session have the same traffic demand).

Provision the many-to-many session on the optical WDM network with the objective of minimizing blocking probability of future many-to-many sessions. In order to minimize the blocking probability of future sessions, the provisioning algorithm must provision the session with the minimum number of new resources used (bandwidth on existing optical channels, transceivers, and wavelengths). Note that the provisioning of the session may not include the use of any new wavelength or transceiver if we can route and groom the session's traffic on the existing virtual topology without adding new lightpaths or light-trees.

We now introduce the notations and assumptions used in the paper. The optical WDM network has an arbitrary topology represented by an undirected graph $G(V, E)$, with a set of nodes V ($N = |V|$) and a set of physical links E . Each physical link $e \in E$ is composed of two unidirectional fibers in opposite directions. The number of wavelengths per fiber is the same among all fibers and is denoted by W , the grooming factor is denoted by g , and the number of transceivers available at each node is the same among all nodes and is denoted by R . An arriving many-to-many session request is denoted by s with a set of members $m_s \subseteq V$ with cardinality $N_s = |m_s|$. Each member in m_s has the same traffic demand t_s , where $1 \leq t_s \leq g$. We also define $H_s = \lceil (N_s - 1)t_s / g \rceil$ to be a lower bound on the number of incoming channels to a member in m_s in order to receive the traffic from the other $N_s - 1$ members in the same session.




Fig. 2. (a): LC for a session s_1 where $m_{s_1} = \{A, B, C\}$ each with one traffic unit denoted as a_1, b_1 and c_1 , respectively ($g = 4; H_{s_1} = 1$). (b): Optimal provisioning of session s_2 (while s_1 in service) where $m_{s_2} = \{B, C, D\}$ each with one traffic unit denoted as b_2, c_2 and d_2 . (c) Alternative non-optimal provisioning of s_2 (while s_1 in service).

III. HEURISTICS FOR NON-SPLITTING NETWORKS

In this section, we introduce heuristic solutions for the dynamic many-to-many traffic grooming problem in non-splitting networks.

A. Lightpath Cycles Heuristic (LCH)

In non-splitting networks, only lightpaths are supported. A direct lightpath (that may span multiple physical links) can be established between any two nodes in the network and it may groom traffic from different sessions and traffic from different members within the same session.

Definition 1. A lightpath cycle (LC) for a many-to-many session s is a simple cycle of N_s lightpaths that visits each member in m_s exactly once.

An example of a LC for a many-to-many session s_1 with a set of members $m_{s_1} = \{A, B, C\}$ is shown in Fig. 2.(a). Note that the LC for a session s only describes a virtual topology and it always contains N_s lightpaths regardless of the order of the members and regardless of the underlying physical topology. In our previous work on the static many-to-many traffic grooming problem [15], we have shown that LCs serve as an optimal virtual topology (in terms of minimizing the total number of transceivers used) to provision a single many-to-many session, as indicated by the following theorem.

Theorem 1. An optimal virtual topology that minimizes the total number of transceivers required to provision a single many-to-many session s consists of H_s identically ordered LCs for s .

Proof: See [15]. ■

Although this theorem is derived for the static many-to-many traffic grooming problem, it is quite useful in the dynamic version of the problem. For example, consider a network state where there are no sessions in the network and consider an arrival of a many-to-many session s . Based on Theorem 1, the optimal way to provision s with the minimum number of new resources used is through H_s LCs for s . Here, we focus on the number of new transceivers used since the number of new wavelengths used depends on the routing and wavelength assignment approach used. The way the traffic is routed on the LCs is as follows. Each member in m_s transmits its traffic through the H_s identically ordered LCs for s until it

Algorithm 1. Lightpath Cycles Heuristic (LCH)**input** : An arriving many-to-many session request s and the current network state**output**: Provisioning of s

```

1 Separate members in  $m_s$  into two disjoint sets  $\mathcal{O}$  and  $\mathcal{N}$ .
2  $order(\mathcal{O})$ .
3  $order(\mathcal{N})$ .
4 for ( $i = 0, 1, \dots, |\mathcal{O}| - 2$ ) do
5   Provision  $\min\{(N_s - 1)t_s, C_{\mathcal{O}[i], \mathcal{O}[i+1]}\}$  on existing lightpaths
   between members  $\mathcal{O}[i]$  and  $\mathcal{O}[i+1]$  in  $VT$ .
6   Establish  $\lceil ((N_s - 1)t_s - C_{\mathcal{O}[i], \mathcal{O}[i+1]})/g \rceil$  new lightpaths
   between members  $\mathcal{O}[i]$  and  $\mathcal{O}[i+1]$  to provision the remaining
   unprovisioned traffic (if any).
7 end
8 for ( $i = 0, 1, \dots, |\mathcal{N}| - 2$ ) do
9   Establish  $H_s$  lightpaths between members  $\mathcal{N}[i]$  and  $\mathcal{N}[i+1]$ .
10  if ( $|\mathcal{O}| = 0$ ) then
11    Establish  $H_s$  lightpaths between members  $\mathcal{N}[|\mathcal{N}| - 1]$  and
     $\mathcal{N}[0]$ .
12  end
13  else
14    if ( $|\mathcal{N}| = 0$ ) then
15      Establish  $H_s$  lightpaths between members  $\mathcal{O}[|\mathcal{O}| - 1]$ 
      and  $\mathcal{O}[0]$ .
16    end
17    else
18      Establish  $H_s$  lightpaths between members  $\mathcal{O}[|\mathcal{O}| - 1]$ 
      and  $\mathcal{N}[0]$  and  $H_s$  lightpaths between members
       $\mathcal{N}[|\mathcal{N}| - 1]$  and  $\mathcal{O}[0]$ .
19    end
20  end
21 end

```

Procedure 1. $order(\mathcal{X})$ select a member in \mathcal{X} randomly as the *current member*;**while** there is at least one unselected member in \mathcal{X} **do**Case 1: $\mathcal{X} = \mathcal{O}$ select the *next member* (from the remaining unselected members) as the member who has the shortest logical distance in VT from the *current member*.Case 2: $\mathcal{X} = \mathcal{N}$ select the *next member* (from the remaining unselected members) as the member who has the shortest physical distance in G from the *current member*.*current member* = *next member*.**end**

reaches the member just before it in the LCs (see Figure 2.(a)). Using this routing strategy, we guarantee two things. First, exactly $(N_s - 1)t_s$ traffic units are groomed between each pair of consecutive members in the LCs and therefore H_s lightpaths are sufficient to groom this traffic. Second, each member in m_s receives the traffic from the other $N_s - 1$ members in the same session. Another useful property of the H_s LCs is that it equally distributes the use of new transceivers among all the members in the session. This is very important in a dynamic environment where resources (i.e., transceivers) are usually distributed equally among all the nodes in the network.

Although Theorem 1 proves the optimality of lightpath cycles in a special case where there are no sessions in the current network state, many-to-many sessions tend to be provisioned through lightpath cycles even in a network state where there are other sessions that are already provisioned in the network.

For example, consider a network state where session s_1 in Fig. 2.(a) is still in service and a new many-to-many session request s_2 with a set of members $m_{s_2} = \{B, C, D\}$ and $t_{s_2} = 1$ arrives. The optimal provisioning of s_2 in terms of the number of new transceivers used is shown in Fig. 2.(b). Note that s_2 is also provisioned through a LC for s_2 ($B - C - D - B$) and that the LCs for s_1 and s_2 share the lightpath $B \rightarrow C$. More precisely, the lightpath $B \rightarrow C$ grooms the two traffic units $a1, b1$ belonging to session s_1 and the two traffic units $b2, d2$ belonging to session s_2 . Note that the order of the members in the LC for s_2 is significant. For example, if the order of the members in the LC for s_2 was $B - D - C - B$ instead of $B - C - D - B$, then the two LCs for s_1 and s_2 will not share a lightpath and we would require six lightpaths instead of five (see Fig. 2.(c)).

In this subsection, we design a heuristic algorithm that assumes that sessions are provisioned through LCs. More precisely, we assume that each arriving many-to-many session s is provisioned through LCs for s where $(N_s - 1)t_s$ traffic units are groomed between each pair of consecutive members in the LCs. Based on this assumption, the heuristic needs to determine two things for each arriving session s : 1) How to order members in the TLCs for s , and 2) How to provision the $(N_s - 1)t_s$ traffic units between each pair of consecutive members in the LCs.

The current virtual topology is represented in the heuristic as a directed graph VT with a set of nodes that includes every node in G that at least has one lightpath incoming or outgoing. A directed edge from node i to node j exists in VT only if there exists at least one lightpath from node i to node j in G . Each edge (i, j) in VT has a capacity $C_{i,j}$ representing the remaining capacity on lightpaths from node i to node j in G .

The heuristic (shown in Algorithm 1) orders members in session s by first separating the members in m_s into two disjoint sets \mathcal{O} and \mathcal{N} , and then orders each set independently (lines 1-3). The set \mathcal{O} includes members in m_s that already exist in the current virtual topology VT , while the set \mathcal{N} includes the remaining members in m_s that do not exist in VT . The heuristic orders members in the \mathcal{O} set by minimizing the logical hop distance between each pair of consecutive members, while it orders members in the \mathcal{N} set by minimizing the physical hop distance between each pair of consecutive members (see Procedure 1). Afterwards, between each pair of consecutive members in the \mathcal{O} set, the heuristic attempts to provision as much traffic as possible out of the $(N_s - 1)t_s$ traffic units using existing lightpaths in VT (lines 4-5). For the remaining unprovisioned traffic t' (if any), the heuristic establishes $\lceil ((N_s - 1)t_s - C_{\mathcal{O}[i], \mathcal{O}[i+1]})/g \rceil$ lightpaths between the two members (line 6). Between each pair of consecutive members in the \mathcal{N} set, the heuristic establishes H_s lightpaths to provision the $(N_s - 1)t_s$ traffic units (line 8-9). Finally, the heuristic completes the cycle for each session s by connecting the \mathcal{O} set and the \mathcal{N} set by H_s lightpaths at both ends (lines 10-21). It is to be noted that all the new established lightpaths in Algorithm 1 (lines 6-20) are routed using shortest path routing and assigned a wavelength according to first fit

Algorithm 2. Multicast Heuristic (MH)**input** : An arriving many-to-many session request s and the current network state**output**: Provisioning of s

```

1 for ( $i = 0, 1, \dots, |m_s| - 1$ ) do
2   Let  $s_i$  be a multicast session with source  $m_s[i]$  and destinations  $m_s \setminus m_s[i]$ .
3   Construct the shortest path tree for  $s_i$  ( $SPT_i$ ).
4   Provision as much traffic as possible out of the  $t_s$  traffic units from the source  $m_s[i]$  to each of destinations  $m_s \setminus m_s[i]$  on  $SPT_i$  using existing lightpaths.
5   For the remaining unprovisioned traffic (if any), establish new lightpaths on  $SPT_i$  using first fit wavelength assignment.
6 end

```

Algorithm 3. Unicast Heuristic (UH)**input** : An arriving many-to-many session request s and the current network state**output**: Provisioning of s

```

1 for ( $i = 0, 1, \dots, |m_s| - 1$ ) do
2   for ( $j = 0, 1, \dots, |m_s| - 1; j \neq i$ ) do
3     Let  $s_{ij}$  be a unicast session with source  $m_s[i]$  and destination  $m_s[j]$ .
4     Find the shortest path from  $m_s[i]$  to  $m_s[j]$ .
5     Provision as much traffic as possible out of the  $t_s$  traffic units on the shortest path using existing lightpaths.
6     For the remaining unprovisioned traffic (if any), establish a new lightpath on the shortest path using first fit wavelength assignment.
7   end
8 end

```

wavelength assignment.

B. Multicast Heuristic (MH)

Note that a many-to-many session s with N_s members can be viewed as a set of N_s multicast sessions each sourced at one of the members and destined to the remaining $N_s - 1$ members in the same session. Therefore, one approach to provision a many-to-many session s is to first break it into N_s multicast sessions, and then provision each multicast session independently. Multicast traffic grooming has been extensively studied in the literature and many heuristic algorithms have been proposed. A well known heuristic for the dynamic multicast traffic grooming problem is to provision an arriving multicast session on its shortest path tree (SPT). The description of the heuristic is shown in Algorithm 2. The heuristic first breaks the many-to-many session s into N_s multicast sessions and then finds the corresponding SPT for each multicast session (lines 1-3). Then, for each multicast session, the heuristic tries to provision as much traffic as possible from the source to each of the destinations on the SPT using existing lightpaths (line 4). Finally, for the remaining unprovisioned traffic (if any), new lightpaths are added on the SPT using first fit wavelength assignment (line 5).

C. Unicast Heuristic (UH)

A many-to-many session s with N_s members can also be viewed as a set of $N_s(N_s - 1)$ unicast sessions each sourced at one of the N_s members and destined to one of the remaining $N_s - 1$ members. Therefore, one approach to provision a many-to-many session s is to first break it into $N_s(N_s - 1)$ unicast sessions, and then provision each unicast session independently. Unicast traffic grooming has been extensively studied in the literature and many heuristic algorithms have been proposed. A well known heuristic for dynamic unicast traffic grooming is to provision an arriving unicast session on its shortest path (SP) from the source to the destination. The description of the heuristic is shown in Algorithm 3. The heuristic first breaks the many-to-many session s into $N_s(N_s - 1)$ unicast sessions and finds the corresponding shortest path for each unicast session (lines 1-4). Then, for each unicast session, the heuristic tries to provision as much traffic as possible on the shortest path using existing lightpaths (line 5). Finally, for the remaining

unprovisioned traffic (if any), a new lightpath is added on the shortest path using first fit wavelength assignment (line 6).

IV. HEURISTICS FOR SPLITTING NETWORKS

In this section, we introduce heuristic solutions for the dynamic many-to-many traffic grooming problem in splitting networks.

A. Hub-Based Heuristic (HBH)

In our previous work on the static many-to-many traffic grooming problem [15], we have introduced a heuristic algorithm for splitting networks that is based on a hub node that collects traffic from members using lightpaths and then distributes the traffic back to the members using light-trees. In this subsection, we extend the heuristic to the dynamic many-to-many traffic grooming problem. More precisely, for each arriving many-to-many session s , the heuristic selects a hub node from the session's set of members. Each member besides the hub transmits as much traffic as possible out of its t_s traffic units to the hub on the shortest path using existing lightpaths. For the remaining unprovisioned traffic (if any), a new lightpath is added on the shortest path using first fit wavelength assignment (upstream traffic). Using the new technique of *network coding*, the hub then linearly combines the traffic units received together with its own t_s traffic units to generate $N_s - 1$ linearly independent combinations. These combinations must also be linearly independent from each of the original t_s traffic units received from the members. Afterwards, the $N_s - 1$ combinations are groomed and delivered back to the members on the shortest path tree (SPT) from the hub to the members using existing light-trees. For the remaining unprovisioned traffic (if any), new light-trees are added on the SPT using first fit wavelength assignment (downstream traffic), see Fig. 3.(a).

According to this heuristic, each member is guaranteed to recover the original traffic units transmitted by all other members in the same session by linearly combining its own t_s traffic units with the received combinations (i.e., solving N_s linearly independent equations). To perform network coding at the hub, we may need to buffer traffic units that arrive early until all traffic units from the $N_s - 1$ members arrive. Using




Fig. 3. Provisioning of a session s where $m_s = \{A, B, C\}$ each with traffic denoted as a, b and c , respectively ($H_s = 1$) using (a): HBH where $Hub(s) = A$. (b): AOH.

Next Generation SONET, multiservice provisioning platform (MSPP) equipments allow up to 128ms differential delay between different traffic streams.

We propose two simple schemes for selecting the hub for an arriving many-to-many session. The first one, most transceivers used (MTU), selects the member with the largest number of used transceivers. The intuition behind this scheme is to select a member with a large number of lightpaths and light-trees originating and terminating. This increases the likelihood of finding existing lightpaths and light-trees to provision the new session's traffic. The second scheme, least transceivers used (LTU), selects the member with the fewest number of used transceivers. The intuition behind this scheme is to distribute the use of transceivers among all the nodes in the network, and not to make certain nodes a bottleneck.

The description of the heuristic (which we refer to as the hub-based heuristic (HBH)) is shown in Algorithm 4. The heuristic first selects a hub node for the arriving many-to-many session s according to MTU or LTU (line 1). Then, for each member in s besides the hub, the heuristic finds the shortest path to the hub and provisions as much traffic as possible out of the t_s traffic units using existing lightpaths (lines 2-4). Afterwards, for the remaining unprovisioned traffic (if any), the heuristic establishes a new lightpath on the shortest path using first fit wavelength assignment (line 5). The heuristic then provisions as much traffic as possible out of the $(N_s - 1)t_s$ traffic units (linear combinations) from the hub h to the remaining members on the SPT using existing light-trees (line 7). Finally, for the remaining unprovisioned traffic (if any), the heuristic establishes new light-trees on the SPT using first fit wavelength assignment (line 8).

The advantage of network coding in the HBH is the reduction of downstream traffic for each arriving session s from $N_s t_s$ to $(N_s - 1)t_s$ traffic units. Therefore, the total number of transceivers saved for session s equals the total number of light-trees saved ($\lceil N_s t_s / g \rceil - \lceil (N_s - 1)t_s / g \rceil$) times the number of transceivers per light-tree (N_s), which is given by the following equation:

$$R_{saved}(s) = N_s(\lceil N_s t_s / g \rceil - \lceil (N_s - 1)t_s / g \rceil) \quad (1)$$

B. All-Optical Heuristic (AOH)

In this heuristic, the t_s traffic units from each member in an arriving many-to-many session s are delivered directly to

Algorithm 4. Hub-Based Heuristic (HBH)

input : An arriving many-to-many session request s and the current network state
output: Provisioning of s

- 1 let h be the hub node for session s selected according to MTU or LTU.
- 2 **for** $(i = 0, 1, \dots, |m_s| - 1; m_s[i] \neq h)$ **do**
- 3 find the shortest path from $m_s[i]$ to h .
- 4 Provision as much traffic as possible out of the t_s traffic units from member $m_s[i]$ to hub h on the shortest path using existing lightpaths.
- 5 For the remaining unprovisioned traffic (if any), establish a new lightpath on the shortest path using first fit wavelength assignment.
- 6 **end**
- 7 Provision as much traffic as possible out of the $(N_s - 1)t_s$ traffic units (linear combinations) from hub h to the remaining members on the SPT using existing light-trees.
- 8 For the remaining unprovisioned traffic (if any), establish new light-trees on the SPT using first fit wavelength assignment.

Algorithm 5. All-Optical Heuristic (AOH)

input : An arriving many-to-many session request s and the current network state
output: Provisioning of s


- 1 **for** $(i = 0, 1, \dots, |m_s| - 1)$ **do**
- 2 Construct the shortest path tree from $m_s[i]$ to $m_s \setminus m_s[i]$ (SPT_i).
- 3 Provision as much traffic as possible out of the t_s traffic units from the source $m_s[i]$ to each of destinations $m_s \setminus m_s[i]$ on SPT_i using existing light-trees.
- 4 For the remaining unprovisioned traffic (if any), establish a new light-tree on SPT_i using first fit wavelength assignment.
- 5 **end**

the other $N_s - 1$ members in the same session using a light-tree, see Fig. 3.(b). For each member, the heuristic attempts to provision as much traffic as possible out of the t_s traffic units to the other $N_s - 1$ members using existing light-trees on the shortest path tree (SPT). For the remaining unprovisioned traffic (if any), a new light-tree is added on the SPT. The description of the heuristic (which we refer to as the all-optical heuristic (AOH)) is shown in Algorithm 5. For each member in session s , the heuristic first finds the SPT to all other members in the same session (lines 1-2). Then, it provisions as much traffic as possible out of the t_s traffic units on the SPT using existing light-trees (line 3). Finally, for the remaining unprovisioned traffic (if any), the heuristic establishes a new light-tree on the SPT using first fit wavelength assignment (line 4).

According to this heuristic, traffic grooming is only performed when two or more many-to-many sessions with the same member set exist in the network at the same time. Otherwise, no traffic grooming is performed. Therefore, we expect this heuristic to be suitable when traffic demands of user sessions almost fill the capacity of a wavelength channel.


C. Complexity Analysis

The time complexity of the LCH, MH, UH, HBH, and the AOH is dominated by the step of finding the SP/SPT (that



(a): EON Topology

(b): USNET Topology

Fig. 5. Blocking probability comparison between LCH, MH and UH on the EON network ($W = 48$ and $R = 30$).

has a time complexity of $O(N^2)$) and the step of performing first fit wavelength assignment (that has a time complexity of $O(W|E|)$). These two steps are repeated for each member in the session which drives the time complexity of the LCH, MH, UH, HBH, and the AOH to $O(N^3 + NW|E|)$.

V. PERFORMANCE EVALUATION

In this section, we evaluate and compare the performance of the proposed heuristics for both non-splitting and splitting networks. We consider two sample networks in our experiments. One is the European Optical Network (EON) shown in Fig. 4.(a) and the other is the USNET network shown in Fig. 4.(b). Many-to-many sessions arrive according to a Poisson distribution with rate λ and they stay in the network for a time that is exponentially distributed with rate μ . The capacity of a wavelength channel is OC-48 while the basic unit of traffic is OC-1, and hence the grooming factor is $g = 48$. The traffic demand of members in a session is uniformly chosen from the set $\{OC-1, OC-3, OC-9, OC-12, OC-24, OC-36, OC-48\}$ which represent the recommended rates for OC streams. The number of members in a session is uniformly distributed between $[2, N]$, while a member in a session is randomly selected between $[0, N-1]$. The number of wavelengths per fiber W and the number of transceivers at each node R are set to ($W = 48$, $R = 30$) in the EON network experiments, while they are set to ($W = 64$, $R = 40$) in the USNET network experiments. Finally, the number of sessions in each simulation run is set to 1000.

Figs. 5 and 6 show the blocking probability for the three heuristics for non-splitting networks (LCH, MH and UH) for different values of network traffic load in the EON and the USNET networks, respectively. We can see from the figures that the LCH outperforms both the MH and the UH. This demonstrates the effectiveness of lightpath cycles in




Fig. 6. Blocking probability comparison between LCH, MH and UH on the USNET.





Fig. 7. Blocking probability comparison between HBH-MTU, HBH-LTU, and AOH on the EON network.

Fig. 8. Blocking probability comparison between HBH-MTU, HBH-LTU, and AOH on the USNET network ($W = 64$ and $R = 40$).

provisioning many-to-many sessions. It also demonstrates that a many-to-many session better be viewed as a single session rather than a set of multicast or unicast sessions.

Figs. 7 and 8 show the blocking probability of the three heuristics for splitting networks (HBH-MTU, HBH-LTU, and AOH) for different values of network traffic load in the EON and the USNET networks, respectively. We can see from the figures that the HBH heuristics, through the novel approach of combining optical splitting and network coding, outperform the AOH. We can also see from the figures that the HBH-LTU outperforms the HBH-MTU. The intuition behind this is that the HBH-LTU distributes the use of transceivers among all the nodes in the network which avoids making certain nodes a bottleneck. Although the HBH-MTU better utilizes existing lightpaths and light-trees, it makes certain nodes in the network a bottleneck which increases the blocking probability.

Next, we compare the performance of the heuristics for non-splitting networks with the heuristics for splitting networks.

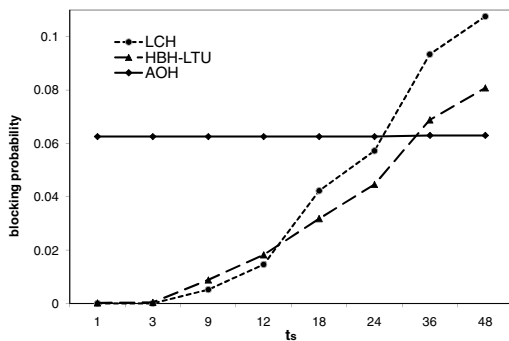


Fig. 9. Blocking probability comparison between LCH, HBH-MTU and AOH on the EON network ($W = 48$ and $R = 30$).

We will show when each of the heuristics is the most suitable choice for dynamic many-to-many traffic grooming. Since the grooming capabilities of the heuristics are varied, their performance will be dependent on traffic granularities of sessions in the network. Therefore, we should compare them for different values of traffic granularities. To make this comparison, we perform eight simulation runs where we fix the traffic demand t_s of arriving many-to-many sessions in each run to one of the following eight values $\{OC-1, OC-3, OC-9, OC-12, OC-24, OC-36, OC-48\}$, respectively. All other settings of the eight runs are exactly the same as the settings described earlier at the beginning of this section and the network traffic load of all runs is fixed to 0.5. Figs. 9 and 10 compare the blocking probability of the LCH, HBH-LTU and AOH for different values of t_s on the EON and the USNET networks, respectively.

We can see from Figs. 9 and 10 that the heuristic for non-splitting networks, LCH, is the most suitable choice when traffic granularities of sessions are relatively low (e.g., $t_s \leq g/4$). This is intuitive since lightpaths are more efficient than light-trees in grooming and packing low granularity traffic. This is a result of the point-to-point nature of a lightpath where it is possible to route many sessions or members with sub-wavelength granularities through it. We can also see from Figs. 9 and 10 that the heuristic for splitting networks, AOH, is the most suitable choice when traffic granularities of sessions are relatively high (e.g., $t_s \geq 3g/4$). This is also intuitive since when traffic granularities of sessions are relatively high, then light-trees are more efficient than lightpaths (a light-tree from a source to a set of destinations requires fewer transceivers than a set of lightpaths each from the source to one of the destinations). Finally, the heuristic for splitting networks, HBH-LTU, which uses both lightpaths and light-trees is the most suitable choice when traffic granularities of sessions are in the middle (e.g., $g/4 < t_s < 3g/4$).

VI. CONCLUSIONS

We have addressed the dynamic many-to-many traffic grooming problem in optical WDM mesh networks. We have introduced different heuristic solutions for the problem in both non-splitting and splitting networks. It was shown that the LCH, that is based on lightpath cycles, outperforms the multi-




Fig. 10. Blocking probability comparison between LCH, HBH-MTU and AOH on the USNET network ($W = 64$ and $R = 40$).

cast and the unicast heuristics (MH and UH, respectively), and that it is the most suitable choice when traffic granularities of sessions are relatively low (e.g., $t_s \leq g/4$). It was also shown that the HBH-LTU, through the novel use of network coding, is the most suitable choice when traffic granularities of sessions lie in the middle (e.g., $g/4 < t_s < 3g/4$), and that the AOH is the most suitable choice when traffic granularities of sessions are relatively high (e.g., $t_s \geq 3g/4$).

REFERENCES

- [1] B. Quinn and K. Almeroth, "IP multicast applications: challenges and solutions," IETF Request for Comments (RFC) 3170, Sept. 2001.
- [2] S. Antonakopoulos and L. Zhang, "Approximation Algorithms for Grooming in Optical Network Design," in *Proc. IEEE Infocom'09*.
- [3] K. Zhu and B. Mukherjee, "Traffic Grooming in an Optical WDM Mesh Network," *IEEE J. Select. Areas Commun.*, vol. 20, pp. 122-133, 2002.
- [4] R. Dutta and G. N. Rouskas, "Traffic grooming in WDM networks: past and future," *IEEE Network*, vol. 16, no. 6, pp. 46-56, Nov./Dec. 2002.
- [5] H. Madhyastha et al., "Grooming of multicast sessions in metropolitan WDM ring networks," *Computer Networks*, vol. 49, pp. 561-579, 2005.
- [6] R. Ul-Mustafa and A. E. Kamal, "Design and provisioning of WDM networks with multicast traffic grooming," *IEEE J. Select. Areas Commun.*, vol. 24, no. 4, pp. 37-53, Apr. 2006.
- [7] A. E. Kamal, "Algorithms for Multicast Traffic Grooming in WDM Mesh Networks," *IEEE Communications* pp. 96-105, Nov. 2006.
- [8] R. Dutta, A. Kamal, and G. Rouskas, "Traffic Grooming for Optical Networks: Foundations, Techniques and Frontiers" *Springer*, August 2008.
- [9] R. Ul-Mustafa and A. E. Kamal, "Many-to-one Traffic Grooming with Aggregation in WDM Networks" *IEEE J. Select. Areas Commun.*, vol. 24, pp. 68-81, 2006.
- [10] S. Ramasubramanian and A. K. Somani, "MICRON: A framework for connection establishment in optical networks," *IEEE/ACM Transactions on Networking*, vol. 16, no. 2, April 2008, pp. 473-485.
- [11] C. Xin, "Blocking analysis of dynamic traffic grooming in mesh WDM optical networks," *IEEE/ACM Transactions on Networking*, Vol. 15, no. 3, June 2007 pp. 721-733.
- [12] G. Chowdhary and C. S. R. Murthy, "Dynamic Multicast Traffic Engineering in WDM Groomed Mesh Networks," in *Workshop on Traffic Grooming*, 2004.
- [13] A. Khalil et al., "Sequential and Hybrid Grooming Approaches for Multicast Traffic in WDM Networks," in *Proc. GLOBECOM'04*
- [14] M. Saleh and A. Kamal, "Many-to-Many Traffic Grooming in WDM Networks," *IEEE/OSA Journal of Optical Communications and Networking*, pp. 376-391, vol. 1, no. 5, Oct. 2009.
- [15] M. Saleh and A. Kamal, "Design and Provisioning of WDM Networks for Many-to-Many Traffic Grooming," in *proc. IEEE Globecom'09*.
- [16] M. Saleh and A. Kamal, "Approximation Algorithms for Many-to-Many Traffic Grooming in WDM Mesh Networks," in *proc. IEEE Infocom 2010*.