1

Hierarchical Traffic Grooming

George N. Rouskas and Rudra Dutta

Department of Computer Science, North Carolina State University, Raleigh, NC, USA {`rouskas,rdutta`}`@ncsu.edu`

1.1 Introduction

Traffic grooming is the field of study that is concerned with the development of algorithms and protocols for the design, operation, and control of networks with multigranular bandwidth demands [12]. As the number of logical entities (including sub-wavelength channels, wavelengths, wavebands, and fibers) that need to be controlled in a multigranular network increases rapidly with the network size, wavelength capacity, and load, a scalable framework for managing these entities becomes a *sine qua non* for future wide area WDM networks.

Several variants of the traffic grooming problem have been studied in the literature under a range of assumptions regarding the network topology, the nature of traffic, and the optical and electronic switching model [11,13,16,18–21,25,26]. Typically, an integer linear programming (ILP) formulation serves as the basis for reasoning about and tackling the problem. Unfortunately, solving the ILP directly does not scale to instances with more than a handful of nodes, and cannot be applied to networks of practical size covering a national or international geographical area. Consequently, either the ILP is tackled using standard relaxation techniques, or the problem is decomposed into subproblems which are solved using heuristics.

Most of the above studies regard the network as a flat entity for the purposes of lightpath routing, wavelength assignment, and traffic grooming. It is well-known, however, that in existing networks resources are typically managed and controlled in a hierarchical manner. The levels of the hierarchy either reflect the underlying organizational structure of the network or are designed in order to ensure scalability of the control and management functions. Accordingly, several studies have adopted a variety of hierarchical approaches to traffic grooming that, by virtue of decomposing the network, scale well and are more compatible with the manner in which networks operate in practice.

The rest of this chapter is organized as follows. In Section 1.2 we survey hierarchical traffic grooming techniques for networks with a special topology.

In Section 1.3 we present a hierarchical framework for traffic grooming in networks of general topology with either static or dynamic traffic. We present a performance study of hierarchical grooming in Section 1.4, and we conclude the chapter in Section 1.5.

1.2 Hierarchical Grooming in Special Topology Networks

1.2.1 Ring Networks

Early research in traffic grooming focused on ring topologies [11,13,25], mainly due to the practical importance of upgrading the existing SONET/SDH infrastructure to support multiple wavelengths. A point-to-point WDM ring is a straightforward extension of a SONET/SDH network, but requires that each node be equipped with one add-drop multiplexer (ADM) per wavelength. Clearly, such a solution has a high ADM cost and cannot scale to more than a few wavelengths. Therefore, much of the research in this context has been on reducing the number of ADMs by grooming sub-wavelength traffic onto lightpaths that optically bypass intermediate nodes, and several near-optimal algorithms have been proposed [11, 25]. However, approaches that do not impose a hierarchical structure on the ring network may produce traffic grooming solutions, in terms of the number of ADMs and their placement, that can be sensitive to the input traffic demands.

The study in [13] was the first to present several hierarchical ring architectures and to characterize their cost in terms of the number of ADMs (equivalently, electronic transceivers or ports) and wavelengths for non-blocking operation under a model of dynamic traffic. In a single-hub ring architecture, each node is directly connected to the hub by a number of lightpaths, and all traffic between non-hub nodes goes through the hub. In a double-hub architecture, there are two hub nodes diametrically opposite to each other in the ring. Each node is connected to both hubs by direct lightpaths, and non-hub nodes send their traffic to the hubs for grooming and forwarding to the actual destination.

A more general hierarchical ring architecture was also proposed in [13]. In this architecture, shown in Figure 1.1, ring nodes are partitioned into two types: *access* and *backbone*. The set of wavelengths is also partitioned into access and backbone wavelengths. The access wavelengths are used to connect all nodes, including access and backbone nodes, in a point-to-point WDM ring that forms the first level of the hierarchy. At the second level of the hierarchy, the backbone wavelengths are used to form a point-to-point WDM ring among the backbone nodes only. This hierarchy determines the routing of traffic between two access nodes as follows. If the two access nodes are such that there is no backbone node along the shortest path between them, their traffic is routed using single-hop lightpaths over the access ring along the shortest path. Otherwise, suppose that b_1 and b_2 are the first and last backbone nodes,

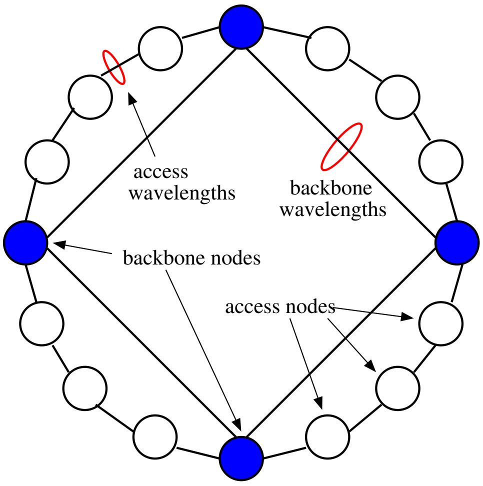

Fig. 1.1. Hierarchical ring architecture with 12 access and 4 backbone nodes

respectively, along the shortest path between two access nodes a_1 and a_2 (note that b_1 and b_2 may coincide). Then, traffic from a_1 to a_2 is routed to b_1 over the access ring, from there to b_2 over the backbone ring, and finally over the access ring to a_2.

A similar hierarchical ring structure was considered in [8], and it was shown that using local (access) and bypass (backbone) wavebands, P-port dynamic traffic (in which each node is allowed to send and receive at most P wavelengths worth of traffic) can be supported with a minimum number of wavelengths.

A different hierarchical approach for grooming sub-wavelength traffic in ring networks was introduced in [23]. Specifically, the N ring nodes are grouped into K *super-nodes*, where each super-node consists of several consecutive ring nodes, as shown in Figure 1.2. The idea behind this partitioning is to pack (groom) all traffic from some super-node x to another super-node y onto lightpaths that are routed directly between the two super-nodes, bypassing intermediate nodes and hence, reducing the number of ADMs required. The study considered both uniform and distance-dependent traffic patterns, and, for each pattern, derived the number K of super-nodes, as a function of the number N of ring nodes and the granularity $C \geq 1$ of each wavelength, so as to minimize the number of ADMs; the granularity C is the number of unit traffic components that can be carried on a single wavelength.

Finally, [13] also proposes the decomposition of a ring into contiguous segments; these are similar to the super-nodes of [23] but are referred to as *subnets*. The ring network is organized in a hierarchical manner as a tree of subnets, where the root of the tree corresponds to a segment that consists of the entire ring. A tree node corresponding to a non-empty subnet s may be subdivided recursively into contiguous subsegments (subnets), and these

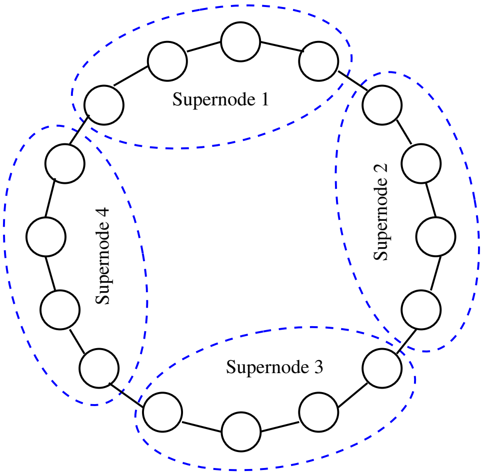

Fig. 1.2. Ring architecture with 4 super-nodes, each of size 4

become the children of subnet s in the tree. The set of wavelengths is also recursively partitioned into *transit* and *internal* sets at each node. Internal wavelengths are used within each subnet child of a node to carry traffic local to this subnet, whereas transit wavelengths are used to carry traffic between the subnet children of a node.

1.2.2 Torus, Tree, and Star Networks

A hierarchical approach for networks with a torus or tree topology was presented in [8], and is based on embedding rings on the underlying topology and then selecting hub nodes along each ring and using bypass wavelengths to interconnect the hubs. Consider first a $N \times M$ torus network, whose nodes are logically arranged on a grid of N rows and M columns. The network is viewed as a collection of N row-rings and M column-rings, and several nodes on each ring are designated as hubs; the hub selection is performed using an algorithm described in [8]. Traffic demands from some source s to a destination d are routed in three steps: from s to a hub h_1 in the same row as s along the appropriate row-ring; from h_1 to a hub h_2 in the same column as h_1 along the column-ring; and finally from h_2 to the destination d in the same row along a row-ring. This approach imposes a two-level hierarchy with non-hub nodes at the first level and hub nodes at the second level.

For tree networks, [8] proposes to embed a *virtual ring* in two steps: (1) using depth-first search to visit every node in the tree, and (2) locally arranging the tree nodes in a ring such that the nodes are connected in the ring in the order in which the corresponding tree nodes were first visited by the depth-first search. By defining hubs along the virtual ring, traffic components can be routed using the same algorithm we described for the hierarchical ring in Figure 1.1.

A traffic grooming algorithm for networks with a star topology was developed in [3]. The algorithm starts by creating lightpaths between the hub and each non-hub node s to carry all traffic originating and terminating at s. Such a solution provides maximum flexibility in terms of grooming, since traffic can be packed efficiently for transmission to the hub, and it can be groomed effectively there for forwarding to the destination. However, it usually requires a large number of lightpaths (equivalently, electronic ports). The algorithm then considers all traffic components in decreasing order of magnitude. Let t be a traffic component from some node s to another node d. The algorithm creates a direct lightpath from s to d to carry t, if there is an available wavelength for doing so; otherwise, no such lightpath is created. A direct lightpath is optically switched at the hub, bypassing electronic switching and grooming, and creating one has the potential to decrease the number of lightpaths by eliminating two lightpaths to/from the hub. The algorithm proceeds until all traffic components have been considered, and returns the solution with the minimum number of lightpaths. It was shown in [3] that this solution is close to optimal for a wide range of problem instances.

1.3 Hierarchical Grooming in General Topology Networks

All the approaches we have discussed so far were developed for networks with topologies that are either symmetric (i.e., ring or torus) or contain no cycles (i.e., tree or star). In this section we describe a framework for hierarchical traffic grooming that is applicable to networks with a general topology. The framework can be used for static or dynamic traffic, and for either sub-wavelength demands (to be groomed into lightpaths) or full-wavelength demands (to be groomed into wavebands). Although our discussion will consider only two levels of hierarchy, this approach can be extended in a straightforward manner to three or more levels of hierarchy to deal with networks of large size.

The traffic grooming problem involves the following conceptual subproblems (SPs) for sub-wavelength demands [12]:

1. *logical topology SP:* find a set of lightpaths to carry the offered traffic;
2. *traffic routing SP:* route the traffic components over the lightpaths; and
3. *lightpath routing and wavelength assignment (RWA) SP:* assign a wavelength and path over the physical topology to each lightpath.

This is only a conceptual decomposition that helps in understanding and reasoning about the problem; in an optimal approach, the subproblems would be considered jointly in the solution. The first and second subproblems together constitute the grooming aspect of the problem.

The hierarchical grooming approach, first described in [5], emulates the hub-and-spoke model used by the airline industry to “groom” passenger traffic onto connecting flights. Specifically, the network is first partitioned into

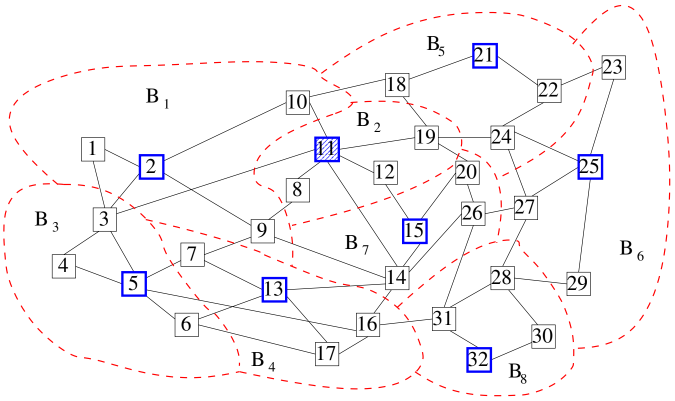

Fig. 1.3. A 32-node network, partitioned into eight first-level clusters $B_1, \cdots, B_8$, with the corresponding hubs at the second level of the hierarchy

clusters (or islands) of nodes, where each cluster consists of nodes in a contiguous region of the network. The clusters form the first level of the hierarchy, and may either correspond to independent administrative entities (e.g., autonomous systems), or may be created solely for the purpose of simplifying resource management and control functions (e.g., as in partitioning a single OSPF administrative domain into multiple areas). Within each cluster, one node is designated as the *hub*, and is responsible for grooming intra-cluster traffic as well as inter-cluster traffic originating or terminating locally. Hub nodes collectively form the second level of the hierarchy, and are expected to be provisioned with more resources (e.g., larger number of switching ports and higher capacity for grooming traffic) than non-hub nodes. Returning to the airline analogy, a hub node is similar in function to airports that serve as major hubs; these airports are typically larger than non-hub airports, in terms of both the number of gates ("ports") and physical space (for "switching" passengers between gates).

To illustrate this approach, let us consider the 32-node network in Figure 1.3. The figure shows a partition of the network into eight clusters, $B_1, \cdots, B_8$, each cluster consisting of four nodes. These clusters represent the first level of the hierarchy. Within each cluster, one node is the hub; for instance, node 2 is the hub for cluster B_1. The hub nodes of the eight first-level clusters form the second level of the hierarchy, and are responsible for grooming and routing inter-cluster traffic.

The main idea behind the hierarchical grooming strategy is to solve the first and second subproblems of the traffic grooming problem (i.e., construct the logical topology and determine the routing of traffic components on it) separately for each level of the hierarchy. In the first step, each cluster is considered independently of the others, and a set of lightpaths is created to route local (intra-cluster) traffic, as well as inter-cluster traffic to and from the local hub. In the second step, lightpaths are created between the hub nodes to carry the inter-cluster traffic. Consequently, the problem of routing inter-cluster demands is divided into three simpler subproblems: routing the component to the local hub, from there to the remote hub, and then to the ultimate destination. Finally, given the set of inter- and intra-cluster lightpaths, the third subproblem can be solved on the underlying physical topology of the network using a standard RWA algorithm.

The hierarchical grooming algorithm for sub-wavelength demands consists of three phases:

1. **Clustering and hub selection.** In this phase, the network is partitioned into clusters and one node in each cluster is designated as the hub.
2. **Hierarchical logical topology formation and traffic routing.** During this phase, the first and second subproblems of the traffic grooming problem are solved in an integrated manner. The outcome of this phase is a set of lightpaths for carrying the traffic demands, and a routing of individual traffic components over these lightpaths.
3. **Routing and wavelength assignment.** Each of the lightpaths in the logical topology are assigned a wavelength and path on the underlying physical topology of the network.

This approach has the following desirable characteristics:

- it is hierarchical, facilitating control, management, and security functions;
- it decouples the grooming of traffic components into lightpaths from the routing and wavelength assignment for these lightpaths: grooming is performed on a logical hierarchy of clusters while RWA is performed directly on the underlying physical topology;
- it provisions only a few nodes (the hubs) for grooming traffic they do not originate or terminate;
- it handles efficiently small traffic demands: at the first level of hierarchy, nodes pack their traffic on lightpaths to the local hub; at the second level, demands among remote clusters are packed onto lightpaths between the corresponding hubs; and
- it allows for large traffic components to be routed on direct lightpaths, eliminating the cost of terminating and switching them at intermediate nodes.

The following subsections discuss each of the three phases of the algorithm in more detail.

1.3.1 Clustering and Hub Selection

The objective of this phase is twofold: (1) to partition the network nodes into some number k of clusters, denoted $B_1, \cdots, B_k$, and (2) to select one node in each cluster B_i as the hub, denoted h_i. Clearly, the number of clusters, their composition, and the corresponding hubs must be selected in a way that helps achieve the goal of minimizing the number of lightpaths and wavelengths required to carry the traffic demands. Therefore, the selection of clusters and hubs is a complex and difficult task, as it depends on both the physical topology of the network and the traffic matrix T. To illustrate this point, consider the tradeoffs involved in determining the number k of clusters[1]. If k is very small (but greater than one), the amount of inter-cluster traffic generated by each cluster will likely be large. Hence, the k hubs may become bottlenecks, resulting in a large number of ports at each hub and possibly a large number of wavelengths (since many lightpaths may have to be carried over the fixed number of links to/from each hub). On the other hand, a large value for k implies a small number of nodes within each cluster. In this case, the amount of intra-cluster traffic will be small, resulting in inefficient grooming (i.e., a large number of lightpaths); similarly, at the second-level cluster, $O(k^2)$ lightpaths will have to be set up to carry small amounts of inter-cluster traffic.

It was observed in [4, 7] that the clustering and hub selection subproblem bears similarities to the classical k-center problem [14, 15]. The objective of the k-center problem is to find a set S of k nodes (centers) in the network, so as to minimize the maximum distance from any network node to the nearest center. Thus, the set S implicitly defines k clusters with corresponding hub nodes in S. A solution to the k-center problem may be useful for hierarchical traffic grooming since it is likely to lead to short lightpaths within a cluster, thus requiring fewer wavelengths. Also, this type of clustering tends to avoid physical topologies with a large diameter for each cluster; such topologies are not a good match for hierarchical grooming that requires nodes to send their traffic to the hub.

The k-center problem is NP-Complete, and the 2-approximation algorithm of [14] was used in [4,7], with one modification. The modification was based on the observation that the k-center problem takes only the physical topology as input, and its only goal is to minimize the maximum node-to-hub distance; in the traffic grooming context, on the other hand, hub capacity should also be considered. Since hubs are responsible for originating and terminating a larger number of lightpaths than non-hub nodes, it is generally desirable to select as hubs the nodes with the largest bandwidth capacity, i.e., those with the largest physical degree, so as to keep the wavelength requirements for the network low. Therefore, whenever the algorithm of [14] selects a new hub arbitrarily among a set of candidate nodes, it was modified to select the candidate node with the maximum physical degree.

A clustering algorithm designed specifically for hierarchical traffic grooming was presented in [6]. This work identified several grooming-specific factors

---

[1] Note that in the special case of $k = 1$, there is a single cluster with one hub and $N - 1$ non-hub nodes, whereas in the special case $k = N$, there are N clusters, each with a single hub and no non-hub nodes.

affecting the selection of clusters and hubs, and developed a parameterized algorithm that can achieve a desired tradeoff among various goals. The algorithm partitions the network into clusters by considering: (1) the intra- and inter-cluster traffic, attempting to cluster together nodes with "dense" traffic in order to reduce the number of long inter-cluster lightpaths; (2) the capacity of the cut links connecting each cluster to the rest of the network, selecting clusters with a relatively large cut size so as to keep the number of wavelengths low; and (3) the physical shape of each cluster, attempting to avoid clusters with a large diameter. The algorithm also selects hubs on the basis of their physical degree, to prevent hub links from becoming bottlenecks. It was shown in [6] that this algorithm outperforms the k-center algorithm in terms of both the port and wavelength costs.

1.3.2 Hierarchical Logical Topology Formation and Traffic Routing

The formation of the hierarchical logical topology for traffic grooming follows three steps: formation of direct lightpaths, intra-cluster lightpaths, and inter-cluster lightpaths. The following discussion assumes that traffic demands are static; at the end of the section we explain how a hierarchical topology may be formed for dynamic traffic.

Direct lightpaths for large traffic demands. During this step, "direct-to-destination" lightpaths are created between two nodes that exchange large amounts of traffic, even if these nodes belong to different clusters. Similarly, "direct to/from remote hub" lightpaths are created between some node s and a remote hub h if there is a sufficiently large amount of traffic between s and the nodes in h's cluster. Setting up such lightpaths to bypass the local and/or remote hub node has several benefits: the number of lightpaths in the logical topology is reduced, the number of ports and switching capacity required at hub nodes is reduced (leading to higher scalability), and the RWA algorithm may require fewer wavelengths (since hubs will be less of a bottleneck).

Intra-cluster lightpaths. At this step, each cluster is considered independently of the others, and intra-cluster lightpaths are formed by viewing each cluster as a *virtual star* (despite the fact that, in general, the actual topology of a cluster is very different than that of a *physical* star, as Figure 1.3 illustrates). Consider some cluster B with hub h. The intra-cluster lightpaths within cluster B are formed by (1) having all traffic to (respectively, from) any node s of B from (to) nodes outside the cluster originate (terminate) at the hub h, and (2) applying the algorithm for star networks discussed in Section 1.2.2 to cluster B in isolation. Having all inter-cluster traffic originate or terminate at the hub imposes a hierarchical structure to the logical topology of lightpaths: inter-cluster traffic, other than that carried by direct lightpaths set up earlier, is first carried to the local hub, groomed there with inter-cluster traffic from other local nodes, carried on lightpaths to the destination hub (as we discuss shortly), groomed there with other local and non-local traffic, and finally carried to the destination node.

At this stage, the lightpaths to be created are simply identified; the routing of these lightpaths over the physical topology is performed during the RWA phase discussed in the next section. Depending on the actual topology of the cluster B, which may be quite different than that of a physical star, once routed, these lightpaths may follow paths that do not resemble at all the paths of a physical star. Also, recall that the lightpaths created by the star algorithm are either "single-hop" (i.e., from a non-hub node to the hub, or vice versa), or "two-hop" (i.e., from one non-hub node to another). Hence, the routing of the individual traffic components is implicit in the hierarchical logical topology of each cluster.

Inter-cluster lightpaths. At the end of the intra-cluster grooming step, all traffic (other than that carried by the initial direct lightpaths) from the nodes of a cluster B with destination outside the cluster, is carried to its hub h for grooming and transport to the destination hub. In order to carry this traffic, a second-level cluster is considered, consisting of the k hub nodes of the first-level clusters. This cluster is also viewed as a *virtual star* with an associated traffic matrix representing the inter-cluster demands only. The inter-cluster lightpaths to carry these demands are then obtained by applying the star algorithm of Section 1.2.2 to this cluster in isolation. As with intra-cluster lightpaths, the routing of the inter-cluster lightpaths is performed on the underlying physical topology during the RWA phase.

Dynamic Traffic. While the above discussion assumed a static traffic scenario, the three-step approach to forming the logical topology can be adapted to accommodate dynamic traffic. If a new connection request is for a sufficiently large traffic demand, then a direct lightpath is setup, otherwise grooming must be considered. For an intra-cluster connection request that can be accommodated (groomed) on existing intra-cluster lightpaths (either directly to the destination or through the local hub), no changes in the logical topology are required; otherwise one or two lightpaths will need to be created (from the source to the local hub, and/or from the hub to the destination). Similarly, for an inter-cluster request that cannot be accommodated on the current logical topology, up to three new lightpaths may have to be setup (from the source to the local hub, from there to the remote hub, and finally to the destination). We emphasize again that any lightpaths that need to be created are simply identified in this phase; the routing of these new lightpaths is discussed next.

1.3.3 Routing and Wavelength Assignment

For static traffic demands, the outcome of the logical topology phase is a set of lightpaths and an implicit routing of the original traffic components over these lightpaths. In this case, the objective of this phase is to route the lightpaths over the underlying physical topology, and color them using the minimum number of wavelengths. The static RWA problem on arbitrary network topologies has been studied extensively in the literature [1, 9, 16, 18, 24], and any existing algorithm may be used in this case. With dynamic traffic demands, new lightpaths need to be added to the logical topology each time an arriving connection request cannot be accommodated. An existing dynamic RWA algorithm [10, 22] may be used in this case to find a path and wavelength for the new lightpaths. Hence, by decoupling the grooming and routing of sub-wavelength traffic components onto lightpaths from the routing and wavelength assignment for these lightpaths, hierarchical grooming may capitalize on the the vast body of research on RWA algorithms.

1.3.4 Extension to Lightpath Grooming

The above hierarchical approach may also be applied to networks with multi-granular optical switching capabilities, in which multiple wavelengths may be groomed into wavebands and all wavelengths in a waveband be switched as a group [2]. Let us define a *bandpath*, a generalization of the lightpath concept, as a (waveband, path) pair that uniquely identifies the path over which the set of wavelengths included in the waveband will travel. It was observed in [17] that the traffic grooming problem on a set of full-wavelength demands involves the following three subproblems that are similar to the ones we described earlier:

1. *logical topology SP:* find a set of bandpaths to carry the offered full-wavelength (lightpath) traffic;
2. *lightpath routing SP:* route the lightpaths over the bandpaths; and
3. *bandpath routing and wavelength assignment (RWA) SP:* assign a waveband and path over the physical topology to each bandpath.

Therefore, the hierarchical grooming algorithm can be applied to this problem with only small modifications. Specifically, after partitioning the network and assigning a hub to each cluster, a hierarchical logical topology can be formed by creating bandpaths (instead of single lightpaths) (1) from each node to its local hub; (2) between hubs to carry inter-cluster traffic; and (3) from each hub to the nodes in its cluster. Finally, an existing (static or dynamic) RWA algorithm may be used to assign a waveband and path to each bandpath; since each waveband carries a unique set of wavelengths, assigning a waveband to each bandpath implicitly assigns a wavelength to each lightpath in the bandpath. It was shown in [17] that this hierarchical approach is effective, scalable, and outperforms an existing algorithm for forming and routing bandpaths.

1.4 Performance of Hierarchical Grooming

We now present a small set of experimental results to illustrate the performance of the hierarchical grooming algorithm for static traffic demands. The following methodology was employed in this study. First, the modified k-center algorithm (refer to Section 1.3.1) was used on the 32-node, 53-link network shown in Figure 1.3 to obtain three different clusterings with two, four, and eight clusters, respectively; the special case of a single cluster comprising all

#Clusters	Avg Lightpath Length	Avg Max Hub Degree	Avg #Wavelengths
1	3.17	266	60
2	2.93	231	57
4	2.87	182	56
8	2.75	145	53

Table 1.1. Aggregate statistics over all 30 instances

network nodes was also considered. A random traffic pattern was assumed, and thirty problem instances (i.e., random traffic matrices) were generated. The hierarchical logical topology for each of the 120 instance-clustering pairs was determined by applying the methodology described in Section 1.3.2. The LFAP RWA algorithm [24] was used to route and assign a wavelength to each lightpath; LFAP is fast, conceptually simple, and has been shown to use a number of wavelengths that is close to the lower bound. In order to characterize the performance of the hierarchical grooming solutions, lower bounds on the number of lightpaths and wavelengths necessary to carry a given traffic demand matrix were also obtained; these bounds were obtained independent of the manner (hierarchical or otherwise) in which grooming is performed, as explained in [7].

In order to compare results among different problem instances, two performance metrics were defined: the *normalized lightpath count* and the *normalized wavelength count.* For a given problem instance, the normalized lightpath count is computed as the ratio $lp_h/lp_l \geq 1$, where lp_l is the lower bound on the number of lightpaths, and lp_h is the actual number of lightpaths required in the hierarchical grooming solution. Clearly, the closer this value is to one, the closer the hierarchical solution is to the optimal. The normalized wavelength count is computed in a similar manner.

Figures 1.4 and 1.5 and Table 1.1 present experimental results for the random traffic pattern. Figures 1.4 and 1.5 plot the normalized lightpath and wavelength count, respectively, for each problem instance and corresponding clustering, while Table 1.1 presents aggregate statistics over all thirty problem instances regarding the average lightpath length, the average maximum hub degree (i.e., the maximum of the number of incoming or outgoing wavelengths at the hub), and the average number of wavelengths.

We observe that as the number of clusters into which the network is partitioned increases, the total number of lightpaths in the resulting topology increases gradually (Figure 1.4). On the other hand, the number of required wavelengths generally decreases as the number of clusters increases (Figure 1.5), and so do the average lightpath length and the maximum hub degree. These results can be explained by noting that, as the number of clusters increases, the size of each cluster decreases. With a smaller cluster size, more lightpaths are necessary for both intra-cluster traffic (since the amount of traffic within a cluster is relatively small and lightpaths are not utilized

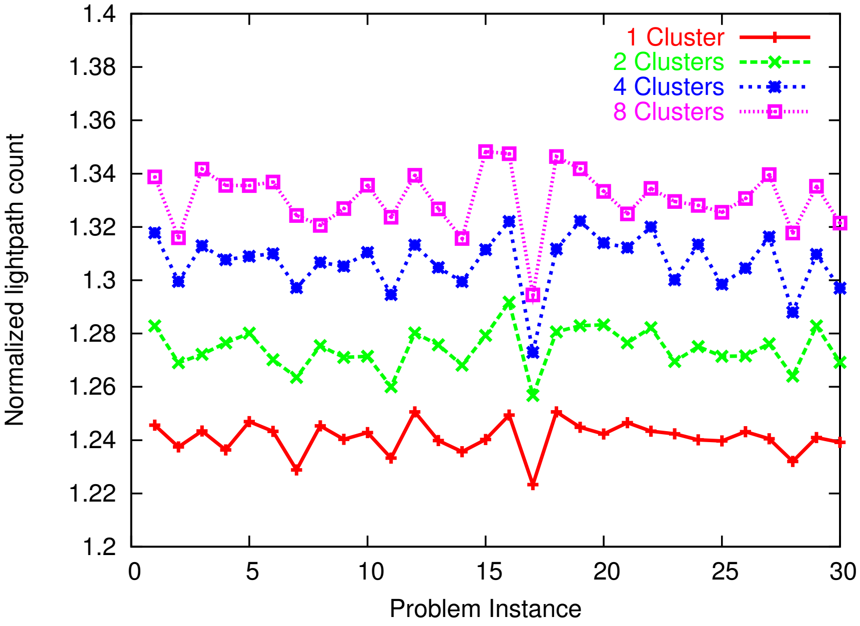

Fig. 1.4. Lightpath comparison

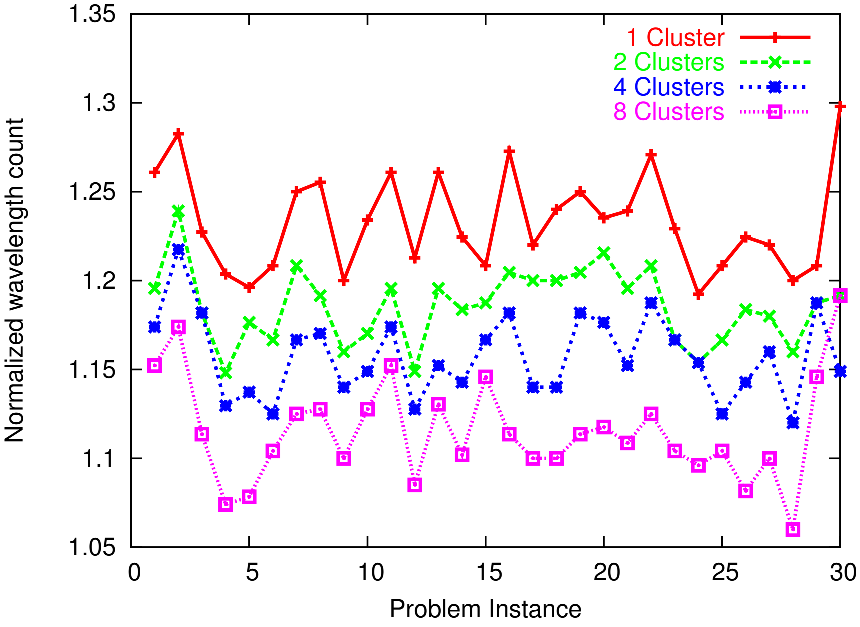

Fig. 1.5. Wavelength comparison

efficiently) and inter-cluster traffic (since each hub has to establish lightpaths to a larger number of hubs in other clusters). Also, intra-cluster lightpaths are shorter when clusters are small, and these short lightpaths are less likely to share links, resulting in fewer wavelengths. At the same time, there is relatively less traffic to be groomed at each hub, hence hub degrees (and hub cost) decrease; the fact that hubs are less of a bottleneck also reduces the wavelength requirements.

From Figure 1.4, we note that the number of lightpaths created by the hierarchical grooming approach are only about 25-35% above the lower bound, and this behavior is consistent across all problem instances. From Figure 1.5, we observe that, with appropriate clustering, the wavelength requirements of this approach are close to the lower bound. Recall that both lower bounds have been computed in a manner that is independent of the grooming methodology employed. Consequently, these results demonstrate that despite its hierarchical nature, this approach produces grooming solutions that are close to optimality.

Similar results regarding the performance of the hierarchical grooming algorithm can be found in [6, 7, 17]. Overall, these results demonstrate that hierarchical traffic grooming can be applied efficiently to large size networks, and produces logical topologies whose lightpath and wavelength requirements are close to the corresponding lower bounds.

1.5 Concluding Remarks

Hierarchical traffic grooming is an efficient and scalable approach to grooming multigranular traffic in large-scale WDM networks with a general topology. We presented a hierarchical framework that is applicable to both static and dynamic contexts, and for either sub-wavelength or full-wavelength traffic. The framework generalizes earlier techniques developed for networks with a special topology (e.g., ring, torus, or tree), and consists of three phases: clustering and hub selection, hierarchical logical topology design, and routing and wavelength assignment. Hierarchical grooming has been shown to perform well over a range of network topologies and traffic patterns, and scales to networks of realistic size.

References

1. S. Baroni and P. Bayvel. Wavelength requirements in arbitrary connected wavelength-routed optical networks. *IEEE/OSA J. Lightwave Technol*, 15(2):242–251, Feb. 1997.
2. X. Cao, V. Anand, and C. Qiao. Waveband switching in optical networks. *IEEE Communications Magazine*, 41(4):105–112, 2003.
3. B. Chen, R. Dutta, and G. N. Rouskas. Traffic grooming in star networks. In *Proceedings of the First Traffic Grooming Workshop*, October 2004.

4. B. Chen, R. Dutta, and G. N. Rouskas. On the application of k-center algorithms to hierarchical traffic grooming. In *Proceedings of the Second Traffic Grooming Workshop*, pages 295–301, October 2005.
5. B. Chen, G. N. Rouskas, and R. Dutta. A framework for hierarchical grooming in WDM networks of general topology. In *Proceeding of BROADNETS 2005*, pages 167–176, October 2005.
6. B. Chen, G. N. Rouskas, and R. Dutta. Clustering and hierarchical traffic grooming in large scale WDM networks. In *Proceedings of ONDM 2007*, May 2007.
7. B. Chen, G. N. Rouskas, and R. Dutta. On hierarchical traffic grooming in WDM networks. *IEEE/ACM Transactions on Networking*, 2007. (Submitted).
8. L.-W. Chen and E. Modiano. Dynamic routing and wavelength assignment with optical bypass using ring embeddings. *Optical Switching and Networking*, 1(1):35–49, 2005.
9. I. Chlamtac, A. Ganz, and G. Karmi. Lightpath communications: An approach to high bandwidth optical WANS. *IEEE Transactions on Communications*, 40(7):1171–1182, July 1992.
10. X-W. Chu, B. Li, and Z. Zhang. A dynamic RWA algorithm in a wavelength-routed all-optical network with wavelength converters. In *Proceedings of IEEE INFOCOM*, pages 1795–1802, 2003.
11. R. Dutta and G. N. Rouskas. On optimal traffic grooming in WDM rings. *IEEE Journal on Selected Areas in Communications*, 20(1):110–121, January 2002.
12. R. Dutta and G. N. Rouskas. Traffic grooming in WDM networks: Past and future. *IEEE Network*, 16(6):46–56, November/December 2002.
13. O. Gerstel, R. Ramaswami, and G. Sasaki. Cost-effective traffic grooming in WDM rings. *IEEE/ACM Transactions on Networking*, 8(5):618–630, October 2000.
14. T. Gonzalez. Clustering to minimize the maximum inter-cluster distance. *Theoret. Comput. Sci.*, 38:293–306, 1985.
15. D. Hochbaum and D.B. Shmoys. A best possible heuristic for the k-center problem. *Math. Oper. Res.*, 10:180–184, 1985.
16. J.Q. Hu and B. Leida. Traffic grooming, routing, and wavelength assignment in optical WDM mesh networks. *Proceedings of the IEEE INFOCOM 2004*, 23(1):495–501, March 2004.
17. M. Iyer, G. N. Rouskas, and R. Dutta. A hierarchical model for multigranular optical networks. 2007. (Submitted).
18. Z.K.G. Patrocínio Jr. and G. R. Mateus. A lagrangian-based heuristic for traffic grooming in WDM optical networks. In *IEEE GLOBECOM 2003*, volume 5,1-5, pages 2767–2771, Dec 2003.
19. V. R. Konda and T. Y. Chow. Algorithm for traffic grooming in optical networks to minimize the number of transceivers. In *IEEE Workshop on High Performance Switching and Routing*, pages 218–221, 2001.
20. C. Lee and E. K. Park. A genetic algorithm for traffic grooming in all-optical mesh networks. In *IEEE SMC 2002*, volume 7, 6-9, Oct 2002.
21. D. Li, Z. Sun, X. Jia, and K. Makki. Traffic grooming on general topology WDM networks. *IEE Proc. Commun.*, 150(3):197–201, June 2003.
22. A. Mokhtar and M. Azizoglu. Adaptive wavelength routing in all-optical netowrks. *IEEE/ACM Transactions on Networking*, 6(2):197–206, April 1998.

23. J. Simmons, E. Goldstein, and A. Saleh. Quantifying the benefit of wavelength add-drop in WDM rings with distance-independent and dependent traffic. *IEEE/OSA Journal of Lightwave Technology*, 17:48–57, Jan 1999.
24. H. Siregar, H. Takagi, and Y. Zhang. Efficient routing and wavelength assignment in wavelength-routed optical networks. *Proc. 7th Asia-Pacific Network Oper. and Mgmt Symposium*, pages 116–127, Oct. 2003.
25. P-J. Wan, G. Calinescu, L. Liu, and O. Frieder. Grooming of arbitrary traffic in SONET/WDM BLSRs. *IEEE Journal on Selected Areas in Communications*, 18(10):1995–2003, 2000.
26. K. Zhu and B. Mukherjee. Traffic grooming in an optical WDM mesh network. *IEEE Journal on Selected Areas in Communications*, 20(1):122 –133, Jan 2002.