




ELSEVIER

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Survey Paper

Hierarchical traffic grooming: A tutorial

Hui Wang^{a,1}, George N. Rouskas^{b,c,d,*}^a Quantcast, San Francisco, CA, USA^b Operations Research, North Carolina State University, Raleigh, NC, USA^c The Department of Computer Science, North Carolina State University, Raleigh, NC 27695-8206, USA^d King Abdulaziz University, Saudi Arabia

ARTICLE INFO

Article history:

Received 23 September 2013

Received in revised form 21 April 2014

Accepted 22 April 2014

Available online 14 May 2014

Keywords:

Traffic grooming

Hierarchical grooming

Network design

ABSTRACT

Traffic grooming is concerned with the design, operation, and control of networks with multigranular bandwidth demands. As the number of resources in a multigranular network increases rapidly with the network size, wavelength capacity, and load, a scalable framework for managing these entities becomes essential. Hierarchical traffic grooming facilitates the control and management of multigranular WDM networks. In this paper, we present a survey of traffic grooming schemes for optical networks that make use of architectures, algorithms and design techniques that impose a hierarchical structure on the network topology.

© 2014 Elsevier B.V. All rights reserved.

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 [1]. 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 essential 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 [2–10].

Typically, an integer linear programming (ILP) formulation serves as the basis for reasoning about and tackling the problem. One crucial concern about such formulations is that they are solvable only for small network topologies [11]. For larger topologies representative of commercial networks, the ILP formulation cannot be solved to optimality within a reasonable amount of time (for instance, within a few hours). Therefore, the offline traffic grooming problem has mostly been addressed using heuristic algorithms [12] whose quality cannot be easily characterized. Other approaches tackle the problem by manipulating the ILP formulation using decomposition or column generation techniques [13].

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

* Corresponding author at: The Department of Computer Science, North Carolina State University, Raleigh, NC 27695-8206, USA. Tel.: +1 919 515 3860; fax: +1 919 515 7925.

E-mail addresses: hwang@quantcast.com (H. Wang), rouskas@ncsu.edu (G.N. Rouskas).

¹ Formerly with Operations Research and the Department of Computer Science, North Carolina State University, Raleigh, NC, USA.

of decomposing the network, scale well and are more compatible with the manner in which networks operate in practice.

In this paper, we present a survey of traffic grooming schemes for optical networks that make use of architectures, algorithms and design techniques that impose a hierarchical structure on the network topology. In Section 2, we examine early research in hierarchical traffic grooming for networks with symmetric or acyclic topologies, including rings, torus, trees, and stars. In Section 3, we review a general framework for hierarchical grooming in networks of general topology, that decomposes the problem into three logical subproblems: clustering, hierarchical virtual topology and traffic routing (H-VTTR), and routing and wavelength assignment (RWA). In Section 4, we discuss and compare several variants of the H-VTTR subproblem. We conclude the survey in Section 5.

2. Hierarchical grooming in elemental topologies

2.1. Ring networks

Early research in traffic grooming focused on ring topologies [2–4], 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 in [3,4]. 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 [2] 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 [2]. In this architecture, shown in Fig. 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




Fig. 1. Hierarchical ring architecture with 12 access and 4 backbone nodes.

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, 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 [14], 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 [15]. Specifically, the N ring nodes are grouped into K super-nodes, where each super-node consists of several consecutive ring nodes, as shown in Fig. 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, [2] also proposes the decomposition of a ring into contiguous segments; these are similar to the super-nodes of [15] 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




Fig. 2. Ring architecture with 4 super-nodes, each of size 4.

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 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.

2.2. Torus, tree, and star networks

A hierarchical approach for networks with a torus or tree topology was presented in [14], 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 [14]. 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, Chen and Modiano [14] 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 Fig. 1.

A traffic grooming algorithm for networks with a star topology was developed in [16]. 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 [16] that this solution is close to optimal for a wide range of problem instances.

3. Hierarchical grooming in mesh 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). A framework for hierarchical traffic grooming that is applicable to networks with a general topology was developed in [17,18], and builds upon the star grooming algorithm of [16]. 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) [19]. 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 [1]:

- (1) *virtual 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. We will refer to these two subproblems as the *virtual topology and traffic routing (VTTR) SP*.

The hierarchical grooming approach, first described in [18], 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 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 (i.e., at nodes within the cluster). 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 Fig. 3. The figure shows a partition of the network into eight clusters, B_1, \dots, 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 in [17] is to solve the first and second subproblems of the traffic grooming problem (i.e., construct the virtual 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 [17], each discussed in the following three subsections.

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, \dots, 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 achieves 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.² 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 [17,20] that the clustering and hub selection subproblem bears similarities to the classical k -center problem [21,22]. 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 . The k -center problem is NP-Complete, and the 2-approximation algorithm of [21] was adapted in [17,20] for the traffic grooming context.

A clustering algorithm designed specifically for hierarchical traffic grooming was presented in [23]. This work identified several grooming-specific factors 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,

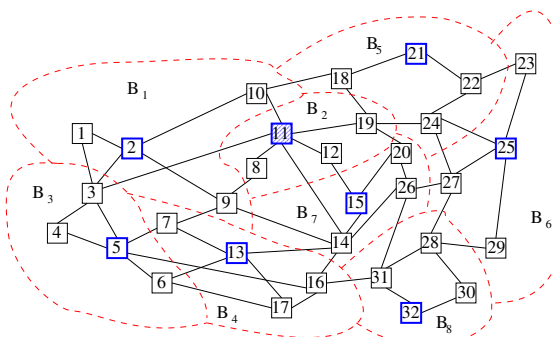


Fig. 3. A 32-node network, partitioned into eight first-level clusters B_1, \dots, B_8 , with the corresponding hubs at the second level of the hierarchy.

² 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.

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 [23] that this algorithm outperforms the k -center algorithm in terms of both the port and wavelength costs.

3.2. Hierarchical virtual topology and traffic routing

During this phase, the VTTR subproblem of the traffic grooming problem is solved. 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. The formation of the hierarchical virtual topology for traffic grooming follows three steps: formation of direct lightpaths, intra-cluster lightpaths, and inter-cluster lightpaths.

Direct lightpaths for large traffic demands. During this step, “direct-to-destination” lightpaths are created between two nodes that exchange a large amount 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 virtual 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. Consider some cluster B with hub h . B is viewed as a *virtual star* such that the intra-cluster lightpaths within the cluster 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 2.2 to cluster B in isolation. Having all inter-cluster traffic originate or terminate at the hub imposes a hierarchical structure to the virtual 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. 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 virtual 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 2.2 to this cluster in isolation. As with intra-cluster lightpaths, the routing of the individual traffic components is implicit in the topology.

3.3. Routing and wavelength assignment

The outcome of the virtual topology phase is a set of lightpaths and an implicit routing of the original traffic components over these lightpaths. 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 [8,9,24–26], and any existing algorithm may be used in this case. 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 vast body of research on RWA algorithms.

3.4. Discussion

The hierarchical grooming framework for general topology networks that was presented in [17] and summarized above 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 virtual 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.

Note, however, that the algorithm proposed in [17] to solve the VTTR subproblem imposes the following constraints on the virtual topology:

- It defines a cluster of nodes around each hub, and requires non-hub nodes to groom inter-cluster traffic towards their local hub only.
- In constructing the virtual topology, it considers each cluster at the first level of the hierarchy, as well as the cluster of hubs at the second level, in isolation.
- It constructs the virtual topology for each cluster by viewing it as a *virtual star* and applying the algorithm in Section 2.2, despite the fact that the physical topology of the cluster may be very different than that of a *physical star*.

With these constraints, the algorithm only takes a few seconds to construct the virtual topology for networks with fifty or more nodes [17]. On the other hand, the constraints exclude a large number of potential solutions hence the algorithm explores only a small fraction of the hierarchical virtual topology solution space. For instance, solutions that groom inter-cluster traffic from a node through a hub other than the node's local hub are disallowed. Similarly, at the second level of the hierarchy, grooming of traffic between hubs is only allowed via a *virtual star* topology.

In the next section, we examine several variants of the hierarchical VTTR problem so as to explore the spectrum of solutions between (1) the flat grooming approach that is the subject of most studies, and (2) the hierarchical grooming algorithm of [17].

4. The hierarchical VTTR problem and variants

Consider a connected graph $G = (\mathcal{V}, \mathcal{L})$, where \mathcal{V} denotes the set of nodes and \mathcal{L} denotes the set of directed links (arcs) in the network. Let $L = |\mathcal{L}|$ denote the number of links. Each directed link $l \in \mathcal{L}$ consists of an optical fiber that may support W distinct wavelengths indexed as $1, 2, \dots, W$. Let $T = [t^{sd}]$ denote the traffic demand matrix, where t^{sd} is a non-negative integer representing the traffic demand units to be established from source node s to destination node d . In general, traffic demands may be asymmetric, i.e., $t^{sd} \neq t^{ds}$. Finally, let C denote the capacity of a single wavelength channel in terms of traffic demand units.

Consider a hierarchical approach to traffic grooming that consists of the three phases we described in Section 3. More specifically, assume that the first phase (hub selection) and the third phase (routing and wavelength assignment) are solved using the algorithms discussed in Sections 3.1 and 3.3, respectively. Therefore, for the remainder of this section we focus on hierarchical solutions to the virtual topology and traffic routing (VTTR) problem of the second phase. The objective of interest is to minimize the total number of lightpaths used in the network; such an objective decreases the use of critical resources and provides ample flexibility for future expansion of the network.

To define the *hierarchical VTTR (H-VTTR)* problem, we assume that a set $\mathcal{H} \subset \mathcal{V}$ of hub nodes in the network is given. Hub nodes are nodes with traffic grooming capabilities. However, in contrast with the work in [17] (and the problem variant we discuss in the following subsection), no clusters are defined in the network; in other words, non-hub nodes are *not* assumed to be assigned to clusters and associated with a “local” hub. We also let $\mathcal{N} = \mathcal{V} \setminus \mathcal{H}$ be the set of non-hub nodes, and $K = |\mathcal{H}|$ be the number of hubs.

Definition 4.1 (H-VTTR). Given the set \mathcal{V} of nodes in the graph G , the set of hubs \mathcal{H} , the wavelength capacity C , and the traffic demand matrix T , establish the minimum number of lightpaths to carry all traffic demands, under two constraints: (1) only hub nodes may groom traffic that

they do not themselves originate or terminate, and (2) no direct lightpaths between two non-hub nodes (i.e., nodes in \mathcal{N}) are allowed.

H-VTTR is a generalization of the VTTR problem defined and studied in [27]. Specifically, VTTR allows grooming of traffic to take place at any node in the network, as well as lightpaths to exist between any pair of nodes in the network. Therefore, if we let $\mathcal{H} = \mathcal{V}$ and $\mathcal{N} = \emptyset$, i.e., each node to be a hub node, H-VTTR reduces to VTTR. Note also that, because of the constraint on direct lightpaths, traffic between two non-hub nodes has to be carried on at least two lightpaths via at least one hub node. For the ILP formulation of the H-VTTR problem, the reader is referred to [28].

In the following subsections we discuss two variants of the H-VTTR problem, both inspired by the airline industry's hub-and-spoke model.

4.1. H-VTTR with clustering (HC-VTTR)

The hierarchical VTTR with clustering (HC-VTTR) problem is a variant of H-VTTR that adopts the concept of clustering considered in [17]. Specifically, we assume that the set \mathcal{V} of network nodes is partitioned into $K = |\mathcal{H}|$ clusters, v_1, \dots, v_K , and that node $h_i \in \mathcal{H}$ is the hub node of cluster v_i . In HC-VTTR, traffic originating from, or terminating at, a non-hub node in cluster v_i may only be groomed with other traffic at the local hub h_i . More formally, we have the following definition.

Definition 4.2 (HC-VTTR). Given the set \mathcal{V} of nodes in the graph G , the set of hubs \mathcal{H} , a set of $K = |\mathcal{H}|$ clusters $\{v_1, \dots, v_K\}$ such that each node $h_i \in \mathcal{H}$ is the hub of cluster v_i , the wavelength capacity C , and the traffic demand matrix T , establish the minimum number of lightpaths to carry all traffic demands, under three constraints: (1) only hub nodes may groom traffic that they do not originate or terminate, (2) traffic originating from, or terminating at, a non-hub node in cluster v_i may only be groomed with other traffic at the local hub h_i , and (3) no direct lightpaths between two non-hub nodes (i.e., nodes in \mathcal{N}) are allowed.

The key idea in HC-VTTR is to ensure that grooming of traffic takes place “near” non-hub nodes (i.e., at their local hub). Local grooming handles small traffic demands efficiently, and it prevents solutions with long underutilized lightpaths. On the other hand, traffic between two non-hub nodes in different clusters must be carried on at least three lightpaths: from the source node to its local hub, then to the remote hub, and finally to the destination node. The ILP formulation of the HC-VTTR problem is similar to that of the H-VTTR problem with additional constraints to prevent the establishment of lightpaths between a non-hub node and hubs other than the one in its cluster.

4.2. Hierarchical grooming with direct lightpaths

The H-VTTR and HC-VTTR problems explicitly prevent direct lightpaths between non-hub nodes. Note, however, that if there is sufficient traffic between two non-hub nodes to fill a lightpath, forcing this traffic to travel via a

hub node results in more lightpaths: sending the traffic directly to its destination requires only one lightpath, whereas sending it through one or more hubs requires at least two lightpaths without improving the grooming of other traffic (since this traffic takes up the whole capacity of these lightpaths). Our experience [17] also indicates that it is often cost-effective to establish partially filled direct lightpaths as long as these lightpaths have high utilization (i.e., the traffic between the two non-hub nodes is close to the capacity of a lightpath). Such high direct traffic demands may not present effective opportunities to groom other traffic on the same lightpaths; furthermore, including partially filled lightpaths in the solution makes it possible to accommodate future increases in traffic demands without the need to establish new lightpaths, an important consideration for long-term network planning.

We now formally define the H-VTTR problem with direct lightpaths (H-VTTR/DL); its ILP formulation is available in [28].

Definition 4.3 (H-VTTR/DL). Given the set \mathcal{V} of nodes in the graph G , the set of hubs \mathcal{H} , the wavelength capacity C , the traffic demand matrix T , and a threshold θ , $0 < \theta \leq 1$, establish the minimum number of lightpaths to carry all traffic demands, under two constraints: (1) only hub nodes may groom traffic that they do not originate or terminate, and (2) direct lightpaths between two non-hub nodes (i.e., nodes in \mathcal{N}) are allowed only if the traffic between these nodes is at least equal to θC .

A similar HC-VTTR/DL problem with clustering can be defined, in which direct lightpaths between non-hub nodes, or a non-hub node and a remote hub, are allowed as long as the traffic between these nodes is at least equal to θC . The ILP formulation of the problem is similar to the formulation of H-VTTR/DL with additional constraints.

Note that, by disallowing direct lightpaths between non-hub nodes, the ILP formulations of the H-VTTR and HC-VTTR problems restrict the solution space substantially. Accordingly, the scalability of the traffic grooming problem improves to the point that it can be solved to optimality within minutes for network sizes representative of national or international backbones (refer also to the next subsection). The H-VTTR/DL and HC-VTTR/DL variants, on the other hand, allow for certain direct lightpaths as determined by the threshold value θ . Specifically, the smaller the value of θ , the larger the size of the solution space that the formulations include. At the limit, i.e., when $\theta = 0$, the H-VTTR/DL and HC-VTTR/DL problems reduce to taking a flat (i.e., non-hierarchical) approach to traffic grooming; as we mentioned earlier, such flat approaches cannot be solved to optimality for anything but toy networks. Therefore, the value of θ represents a tradeoff between the size of the solution space (and, hence, solution quality) and scalability. It has been observed [28] that the number of lightpaths in the optimal solution is not affected when the value of the threshold decreases below $\theta = 0.6$. This observation is consistent with intuition and represents a balance between solution quality and running time.

Finally, we note that the HC-VTTR/DL problem is identical to the one studied in [17]. But whereas the virtual

topology algorithm developed in [17] treated each cluster in isolation as a virtual star and used a heuristic to determine the lightpaths, the ILP formulation considers the clusters in an integrated manner and solves the HC-VTTR problem optimally.

4.3. Performance comparison

We now present results to evaluate the performance of hierarchical solutions to the VTTR problem in terms of two metrics: quality of solution (i.e., the number of lightpaths produced by the solution) and running time. We compare the following five ILP formulations: H-VTTR, H-VTTR/DL, HC-VTTR, HC-VTTR/DL, and VTTR. Note that the VTTR ILP formulation is similar to the one for H-VTTR in [28], but takes a flat view of the network such that grooming may take place at any node, not just hubs, and lightpaths are allowed between any pairs of nodes without any threshold constraints on the traffic demands. Since the four hierarchical formulations are derived from the VTTR formulation by adding appropriate constraints, the solution to the VTTR formulation provides a lower bound for the solutions to the hierarchical formulations. (In fact, as shown in [29], the solution to the VTTR formulation is a lower bound to the solution of the original traffic grooming problem, and it is optimal whenever the network is not wavelength limited.) Hence, we are interested in characterizing the performance of the hierarchical solutions relative to the baseline VTTR formulation.

Four network topologies were used in this study (link counts refer to directed links): the 14-node, 42-link NSFNet [30]; the 17-node, 52-link German network [31]; the 32-node, 106-link network shown in Fig. 3; and the 47-node, 192-link network from [32]. For each problem instance, the traffic matrix $T = [t^{sd}]$ was generated by drawing each traffic demand T^{sd} uniformly and randomly in the interval $[0, t_{max}]$. Each data point in the following figures is the average of ten problem instances. For the experiments, the wavelength capacity was set to $C = 16$, the threshold value was fixed to $\theta = 0.6$, and the value of parameter t_{max} was varied as $t_{max} = 10, 20, 30, 40, 50, 60$, to investigate various traffic loads. The results were obtained by running the IBM CPLEX 12 optimization tool on a cluster of identical compute nodes with dual Woodcrest Xeon CPU at 2.33 GHz with 1333 MHz memory bus, 4 GB of memory and 4 MB L2 cache. A 3% relative optimality gap was imposed in solving the optimization problems with CPLEX.

Figs. 4 and 5 compare the five formulations above across the four network topologies, in terms of the objective value and the CPU time it takes CPLEX to solve them, respectively. For these experiments, we set $t_{max} = 40$, and we used the k -center algorithm [21] to determine the hubs for each topology. Specifically, we set the number of hubs to four for the 14- and 17-node topologies, and eight for the 32- and 47-node networks. We also set a time limit of two hours. As we can see, CPLEX was able to solve all the formulations within the time limit, except for the VTTR formulation on the 47-node network for which no integer solution was found within two hours; hence, the two figures do not present results for this formulation and topology.

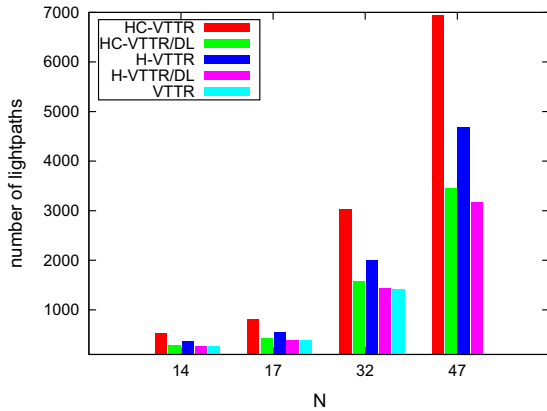


Fig. 4. Objective value comparison, $t_{max} = 40$.




Fig. 5. CPU time comparison, $t_{max} = 40$.

Let us first refer to Fig. 4 that compares the five formulations in terms of solution quality. We first note that the objective value increases with the size of the network topology, as expected: for a given value of t_{max} , a larger network has more traffic to carry than a smaller one, requiring a larger number of lightpaths. We also note that the objective value obtained by solving the HC-VTTR formulation is always higher than that obtained by H-VTTR. Recalling the problem definitions, HC-VTTR includes more constraints than H-VTTR: in the former, traffic from a non-hub node must be groomed at the local hub, whereas in the latter it may be groomed at any hub node. Therefore, the solution to HC-VTTR cannot be better than that to H-VTTR. Also, the variants that allow for direct lightpaths (H-VTTR/DL and HC-VTTR/DL) lead to solutions that are better than variants that do not allow direct lightpaths (H-VTTR and HC-VTTR, respectively). Again, this result can be explained by the fact that allowing direct lightpaths increases the space of candidate solutions. Finally, the original VTTR formulation produces the best solution, as expected, for the three topologies for which a solution to this formulation was obtained within the time limit. However, in all three cases, the solution to H-VTTR/DL is very close to that of VTTR. Overall, the relative performance of the five formulations is consistent

across the four topologies: VTTR leads to the best solution, followed by H-VTTR/DL, HC-VTTR/DL, H-VTTR, and HC-VTTR, in this order.

Let us now turn our attention to Fig. 5 that compares the running time for the five formulations. We observe that solving the HC-VTTR formulation takes the least amount of time, less than a second, on average, even for the 47-node network. Among the hierarchical formulations, the next fastest solution time is achieved by HC-VTTR/DL, followed by H-VTTR and H-VTTR/DL. We also note that, for a given formulation, the running time is similar for the 14- and 17-node networks, and is also similar (but higher) for the 32- and 47-node networks. On the other hand, for the two small networks, the VTTR formulation that does not impose any hierarchical structure on the topology, takes about the same time as H-VTTR/DL, the hierarchical formulation with the worst running time. But whereas the running time of H-VTTR/DL increases by a small factor as we move from the 17- to the 32-node network, the running time of VTTR increases by almost three orders of magnitude; similarly, the running time of H-VTTR/DL increases slightly from the 32- to the 47-node network, but the running time of VTTR increases significantly and exceeds the two-hour limit we imposed. From these results, we conclude that imposing a hierarchical structure on the virtual topology is not beneficial in terms of running time when the size of the network is relatively small (in our study, up to 17 nodes). However, as the network size grows, flat solutions (i.e., VTTR) do not scale whereas hierarchical solutions scale quite well; indeed, it is at larger network sizes that one would expect the benefits of hierarchical structures to materialize. Overall, these results indicate that H-VTTR/DL represents the best tradeoff between running time and solution quality, as it takes, on average, about 100 s or less to obtain solutions close to the optimal (i.e., that obtained by VTTR).

In Fig. 6 we compare the solution quality of the five formulations as a function of t_{max} , i.e., the traffic load. The results shown are for the 32-node network with eight hubs; the relative behavior of the various curves is representative of that for the other topologies. As the traffic load increases, the number of lightpaths increases almost

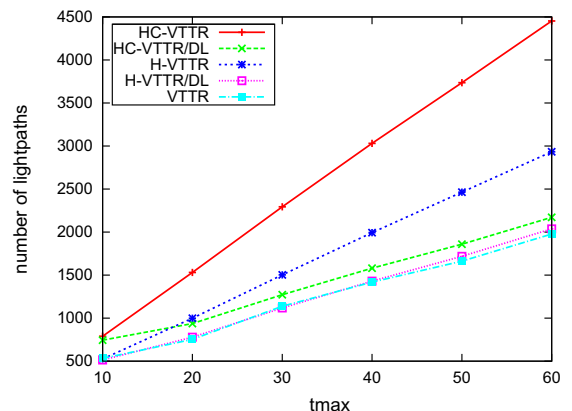


Fig. 6. Objective value comparison, 32-node network, $K = 8$ hubs.

linearly, but the rate of increase depends on the particular formulation. We also observe that the relative performance across the various values of t_{max} is similar to that in Fig. 4, i.e., HC-VTTR requires the largest number of lightpaths, followed by H-VTTR and HC-VTTR/DL, while H-VTTR/DL and VTTR have very similar objective values. These results further support our earlier conclusion that H-VTTR/DL provides the best tradeoff between running time and solution quality.

Finally, Figs. 7 and 8 show the effect of the number K of hubs on the objective value and running time, respectively, for the four hierarchical grooming formulations. As we can see, the number of hubs has little effect on the number of lightpaths for formulations that allow direct lightpaths; this result is due to the fact that a good amount of traffic is sent over such direct lightpaths and hence the number of hubs is not very important. For H-VTTR, as the number of hubs increases, the solution improves; since H-VTTR allows a node to send traffic to any hub, increasing the number of hubs also increases the number of candidate solutions. HC-VTTR, on the other hand, requires each node to send its traffic to the local hub, hence increasing the

number of hubs may also initially increase the overall number of lightpaths. As Fig. 8 indicates, the running time increases with the number of hubs across all formulations, between one and two orders of magnitude $K = 2$ to $K = 8$. Therefore, if one of the formulations that allow direct lightpaths is adopted, these results indicate that a smaller number of hubs should be used.

5. Concluding remarks

We have presented a comprehensive review of research in hierarchical traffic grooming techniques. Such techniques present an efficient and scalable approach to grooming multigranular traffic in large-scale WDM networks with a general topology. With the emergence of the elastic optical network paradigm, exciting new directions open up for the traffic grooming problem. We believe that new hierarchical approaches to attack the problem more ambitiously and in more general contexts remain, and we are certain that many significant results are forthcoming.

Acknowledgements

This work was supported in part by the National Science Foundation under Grant CNS-1113191, and in part by the Deanship of Scientific Research (DSR), King Abdulaziz University, under Grant No. 2-611-1434-HiCi.

References

- [1] R. Dutta, G.N. Rouskas, Traffic grooming in WDM networks: past and future, *IEEE Netw.* 16 (6) (2002) 46–56.
- [2] O. Gerstel, R. Ramaswami, G. Sasaki, Cost-effective traffic grooming in WDM rings, *IEEE/ACM Trans. Network.* 8 (5) (2000) 618–630.
- [3] P.-J. Wan, G. Calinescu, L. Liu, O. Frieder, Grooming of arbitrary traffic in SONET/WDM BLSRs, *IEEE J. Sel. Area. Commun.* 18 (10) (2000) 1995–2003.
- [4] R. Dutta, G.N. Rouskas, On optimal traffic grooming in WDM rings, *IEEE J. Sel. Area. Commun.* 20 (1) (2002) 110–121.
- [5] V.R. Konda, 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*, 2001, pp. 218–221.
- [6] D. Li, Z. Sun, X. Jia, K. Makki, Traffic grooming on general topology WDM networks, *IEE Proc. Commun.* 150 (3) (2003) 197–201.
- [7] K. Zhu, B. Mukherjee, Traffic grooming in an optical WDM mesh network, *IEEE J. Sel. Area. Commun.* 20 (1) (2002) 122–133.
- [8] J.Q. Hu, B. Leida, Traffic grooming, routing, and wavelength assignment in optical WDM mesh networks, in: *Proceedings of IEEE INFOCOM 2004*, March 2004, pp. 495–501.
- [9] Z.K.G. Patrocínio Jr., G.R. Mateus, A lagrangian-based heuristic for traffic grooming in WDM optical networks, in: *IEEE GLOBECOM 2003*, vol. 5, 1–5, December 2003, pp. 2767–2771.
- [10] C. Lee, E.K. Park, A genetic algorithm for traffic grooming in all-optical mesh networks, in: *IEEE SMC 2002*, vol. 7, 6–9, October 2002.
- [11] E. Yetginer, G.N. Rouskas, Power efficient traffic grooming in optical WDM networks, in: *Proceedings of IEEE GLOBECOM 2009*, December 2009.
- [12] K. Zhu, B. Mukherjee, A review of traffic grooming in WDM optical networks: architectures and challenges, *Opt. Netw. Mag.* 4 (2) (2003) 55–64.
- [13] H. Wang, G.N. Rouskas, Traffic grooming in optical networks: decomposition and partial linear programming (LP) relaxation, *J. Opt. Commun. Network.* 5 (8) (2013) 825–835.
- [14] L.-W. Chen, E. Modiano, Dynamic routing and wavelength assignment with optical bypass using ring embeddings, *Opt. Switch. Network.* 1 (1) (2005) 35–49.

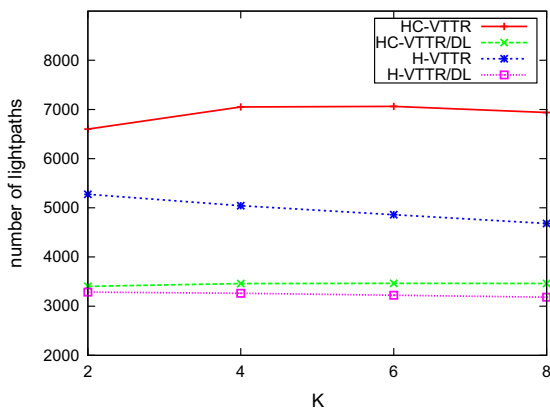


Fig. 7. Objective value comparison, 47-node network, $t_{max} = 40$.




Fig. 8. CPU time comparison, 47-node network, $t_{max} = 40$.

- [15] J. Simmons, E. Goldstein, A. Saleh, Quantifying the benefit of wavelength add-drop in WDM rings with distance-independent and dependent traffic, *IEEE/OSA J. Lightwave Technol.* 17 (1999) 48–57.
- [16] B. Chen, R. Dutta, G.N. Rouskas, Traffic grooming in star networks, in: *Proceedings of the First Traffic Grooming Workshop*, October 2004.
- [17] B. Chen, G.N. Rouskas, R. Dutta, On hierarchical traffic grooming in WDM networks, *IEEE/ACM Trans. Network.* 16 (5) (2008) 1226–1238.
- [18] B. Chen, G.N. Rouskas, R. Dutta, A framework for hierarchical grooming in WDM networks of general topology, in: *Proceeding of BROADNETS 2005*, October 2005, pp. 167–176.
- [19] M. Iyer, G.N. Rouskas, R. Dutta, A hierarchical model for multigranular optical networks, in: *Proceedings of BROADNETS 2008*, September 2008.
- [20] B. Chen, R. Dutta, G.N. Rouskas, On the application of k-center algorithms to hierarchical traffic grooming, in: *Proceedings of the Second Traffic Grooming Workshop*, October 2005, pp. 295–301.
- [21] T. Gonzalez, Clustering to minimize the maximum inter-cluster distance, *Theor. Comput. Sci.* 38 (1985) 293–306.
- [22] D. Hochbaum, D.B. Shmoys, A best possible heuristic for the k-center problem, *Math. Oper. Res.* 10 (1985) 180–184.
- [23] B. Chen, R. Dutta, G.N. Rouskas, Clustering and hierarchical traffic grooming in large scale WDM networks, in: *Proceedings of ONDM 2007*, LNCS, vol. 4534, May 2007, pp. 249–258.
- [24] I. Chlamtac, A. Ganz, G. Karmi, Lightpath communications: an approach to high bandwidth optical WANs, *IEEE Trans. Commun.* 40 (7) (1992) 1171–1182.
- [25] H. Siregar, H. Takagi, Y. Zhang, Efficient routing and wavelength assignment in wavelength-routed optical networks, in: *Proc. 7th Asia-Pacific Network Oper. and Mgmt Symposium*, October 2003, pp. 116–127.
- [26] S. Baroni, P. Bayvel, Wavelength requirements in arbitrary connected wavelength-routed optical networks, *IEEE/OSA J. Lightwave Technol.* 15 (2) (1997) 242–251.
- [27] H. Wang, G.N. Rouskas, An efficient algorithm for solving traffic grooming problems in optical networks, in: *Proceedings of IEEE ICC 2013*, June 2013.
- [28] H. Wang, G.N. Rouskas, Hierarchical traffic grooming formulations, in: *Proceedings of IEEE GLOBECOM 2013*, December 2013.
- [29] H. Wang, Z. Liu, G.N. Rouskas, Scalable traffic grooming in optical networks, in: *Proceedings of ACP 2012*, November 2013.
- [30] B. Mukherjee et al., Some principles for designing a wide-area WDM optical network, *IEEE/ACM Trans. Network.* 4 (5) (1996) 684–696.
- [31] W.D. Grover, D.P. Onguetou, A new approach to node-failure protection with span-protecting p-cycles, in: *Proceedings of ICTON 2009*, June 2009.
- [32] P. Baran, On distributed communications networks, *IEEE Trans. Commun.* 12 (1) (1964) 1–9.



Hui Wang received her Ph.D. degree in Operations Research, with co-major in Computer Science, at North Carolina State University. She received the Master in Statistics and Master in Operations Research degrees from NCSU in 2011 and 2009, respectively, and the B.S. degree in Applied Mathematics from Peking University, Beijing, China in 2007. Her research interests include optical networks, network design and optimization, and performance evaluation. She joined Quantcast in August 2013.



George N. Rouskas is a Professor of Computer Science at North Carolina State University and an IEEE Fellow. He received the Diploma in Computer Engineering from the National Technical University of Athens (NTUA), Athens, Greece, and the M.S. and Ph.D. degrees in Computer Science from the College of Computing, Georgia Institute of Technology, Atlanta, GA. His research interests include network architectures and protocols, optical networks, network design and optimization, and performance evaluation. He is co-editor of

the book “Next-Generation Internet Architectures and Protocols” (Cambridge University Press, 2011), author of the book “Internet Tiered Services” (Springer, 2009), and co-editor of the book “Traffic Grooming for Optical Networks” (Springer, 2008). He is founding co-editor-in-chief of the *Optical Switching and Networking Journal* and he has served on the editorial boards of the *IEEE/ACM Transactions on Networking*, *IEEE/OSA Journal of Optical Networking*, *Computer Networks*, and *Optical Networks*. He was the General Chair for ICCCN 2013, and he has served as TPC or general chair for several conferences, including ICCCN 2011, the IEEE GLOBECOM 2010 ONS Symposium, BROADNETS 2007, IEEE LANMAN 2004 and 2005, and IFIP NETWORKING 2004. He is the recipient of a 1997 NSF CAREER Award, the 2004 ALCOA Foundation Engineering Research Achievement Award, and the 2003 NCSU Alumni Outstanding Research Award, and he was inducted in the NCSU Academy of Outstanding Teachers in 2004. He serves as the Secretary of the IEEE Optical Networking Technical Committee (ONTC) in 2012–2013, and he served as a Distinguished Lecturer for the IEEE Communications Society in 2010–2011.