Dynamic Traffic Grooming in WDM Mesh Networks Using a Novel Graph Model *

Hongyue Zhu+, Hui Zang++, Keyao Zhu+, and Biswanath Mukherjee+

+ Department of Computer Science, Univ. of California, Davis, CA 95616, USA
Tel: +1-530-752-5129, Fax: +1-530-752-4767, Email: {zhuh, zhuk, mukherje}@cs.ucdavis.edu
++ Sprint Advanced Technology Laboratories, Burlingame, CA 94010, USA
Tel: +1-650-375-4423, Fax: +1-650-375-4330, Email: hzang@sprintlabs.com

Abstract

As our fiber-optic backbone networks migrate from interconnected SONET rings to arbitrary mesh topology, traffic grooming on WDM mesh networks becomes an extremely important research problem. (Traffic grooming refers to the problem of efficiently packing a set of low-speed connection requests onto high-capacity channels.) To address this problem, we employ a new, generic graph model for dynamic traffic grooming in WDM mesh networks, where connections arrive one at a time and hold for random time durations. The novelty of this model is that, by only manipulating the edges of an auxiliary graph created by the model and the weights of these edges, the model can achieve various objectives using different grooming policies, while taking into account various constraints such as transceivers, wavelengths, wavelength-conversion capabilities, and grooming capabilities. Based on the auxiliary graph, we develop a dynamic traffic-grooming algorithm which employs the simple shortest-path computation method to solve the traffic-grooming problem. Different grooming policies can be implemented by different weight functions assigned to the edges in the auxiliary graph. We propose four fixed grooming policies, whose performance is compared in various network scenarios. We also develop an Adaptive Grooming Policy (AGP), in which the weights of the edges are dynamically adjusted according to the current network state, and our results show that it outperforms the fixed grooming policies.

Keywords

Optical network, mesh network, WDM, dynamic traffic grooming, graph model

I. INTRODUCTION

Wavelength-division multiplexing (WDM) is a key approach to increase the bandwidth of an optical network [1]. As WDM technology continues to mature, there exists a large gap between the capacity of a WDM channel (e.g., OC-48, or OC-192, or OC-768) and the bandwidth requirement of a typical connection request (e.g., STS-1, OC-3, OC-12, etc.). In order to use the network resources efficiently, low-speed traffic streams need to be multiplexed and switched onto high-speed lightpaths, which is also known as traffic-grooming problem [2].

The traffic-grooming problem can be formulated as follows [3]. Given a network configuration (including physical topology, where each edge is a physical link; number of transceivers at each node; number of wavelengths on each fiber; and the capacity of each wavelength) and a set of connection requests with different bandwidth granularities, such as OC-12, OC-48, etc., we need to determine how to set up lightpaths to satisfy the connection requests. Because of the sub-wavelength granularity of the connection requests, one or more connections can be multiplexed on the same lightpath.

* This work has been supported by the National Science Foundation (NSF) under Grant Nos. NCR-9508239 and ANI-9805285, and by Sprint Advanced Technology Laboratories.

In dynamic traffic grooming, the connection requests arrive one at a time with different starting time and holding period. The objective is to minimize the network resources used for each request, which implicitly attempts to minimize the overall blocking probability.

When a connection request arrives, the network operator should determine the following: (1) Should this connection be routed on the current set of lightpaths, i.e., virtual topology, if it is possible to do so? Sometimes, it may be better to set up a new lightpath even though the connection can be carried on the current set of lightpaths. (2) How to change the virtual topology to accommodate the connection? i.e., between which two nodes should we set up a new lightpath, if any? In some cases, we can set up a lightpath directly from the source of the traffic to the destination. In other cases, it may not be necessary or possible to set up this direct lightpath; instead, we may need to set up some other lightpath(s) and route the connection onto these lightpath(s) and/or some existing lightpaths. Different decisions on these questions can result in different network performance. These decisions reflect the intentions of the network operator, and they are referred to as *grooming policies* [4], [5]. By using different grooming policies, a network operator can achieve various objectives. As the network state changes, the optimization objective may also need to change. Dynamically evolving the grooming policy according to the network state is a challenge for traffic grooming.

A. Previous Work

Traffic grooming is an important and practical problem for designing WDM networks and it is receiving increasing research attention both in academe and in industry. The work in [6] reviews most of the recent research work on traffic grooming in WDM ring and mesh networks.

Past research efforts on traffic grooming have mainly focused on SONET/WDM ring networks. The major cost of such a network is considered to be dominated by SONET add-drop multiplexers (ADMs). Therefore, minimizing the number of SONET ADMs has been the objective of static traffic grooming in recent research. The general traffic-grooming problem in a SONET/WDM ring network is proven to be NP-complete [7], [8]. An optimal algorithm for a single-hub ring is proposed in [7] and several optimal or near-optimal algorithms for traffic grooming and wavelength assignment to reduce the number of wavelengths and SONET ADMs are proposed in [9]. As a network design problem, the authors in [10] attempt to minimize the network cost, which is dominated by SONET ADMs, in an optical add-drop wavelength-division-multiplexed (OADM) ring network. Six optical WDM ring architectures are provided in [10] and the cost of different architectures, as well as the switching capabilities of different architectures under various traffic assumptions are compared. The maximum terminal-equipment savings using wavelength ADMs are quantified in [11] for WDM rings carrying uniform and distance-dependent traffic. Grooming with arbitrary traffic in Bidirectional-Line-Switched-Rings (BLSRs) is addressed in [8]. In [12], based on a general formulation of the virtual topology problem, a framework used to evaluate the performance of heuristics and requiring less computation than evaluating the optimal solution is presented. The authors in [13] formulate the grooming optimization problem as an integer linear program (ILP) and compare single-hop grooming and multi-hop grooming. Instead of single-ring architectures, interconnected WDM rings are studied in [14] and several strategies for traffic grooming in such networks are compared. All the above references except [10] focus on static traffic only. The authors of [15] study the dynamic traffic-grooming problem in SONET/WDM rings and formulate it as a bipartite graph-matching problem.

As our fiber-optic backbone networks migrate from rings to mesh, traffic grooming on WDM mesh networks becomes an extremely important area of research. The authors in [3] propose several node architectures for supporting traffic grooming in WDM mesh networks and formulate the static traffic-grooming problem as an ILP. They present two heuristics and compare the performance with that of the ILP. The works in [4], [16], [17], [18], [19], [20], [21] consider a dynamic traffic pattern in WDM mesh networks. In [4], two route computation algorithms are proposed and compared, and the results indicate that, in order to achieve good performance in a dynamic environment, different grooming policies and route-computation algorithms need to be used under different network states. In [16], the authors propose a connection admission control scheme to ensure fairness in terms of connection blocking. A theoretical capacity correlation model is presented in [17] to compute the blocking probability for WDM networks with constrained grooming capability. The work in [18] addresses how to dynamically establish reliable low-rate traffic in WDM mesh networks with traffic-grooming capability and two grooming schemes are compared. The work in [19] studies how to plan and design a WDM mesh network with certain forecast traffic demands to satisfy all the connections as well as minimize the network cost. In [20], the authors investigate how to design multi-layer mesh networks to satisfy the connections' bandwidth and protection requirement while minimizing the overall network cost. In [21], the authors develop a heuristic to groom dynamic traffic based on the concept of collaboration group.

B. Our Contribution

The WDM backbone network is expected to emerge as a multi-vendor, heterogeneous mesh network. Some of the nodes may have full wavelength-conversion capability (so that any incoming wavelength can be converted into any outgoing wavelength), some may have no wavelength-conversion capability (so that traffic must stay on the same wavelength when bypassing these nodes), and some may have partial wavelength-conversion capability (where some wavelengths can be converted into some other wavelengths). At the same time, a WDM mesh network may consist of systems from multiple vendors, and different vendors may employ different node architectures and technologies, which may have different grooming capabilities. Some architectures may have full grooming capabilities, while some may impose constraints such as the number of grooming ports (represented by the number of transceivers used for originating and terminating groomable wavelength channels) on the grooming capability. In addition, some nodes may have no grooming capability. These partial and sparse wavelength-conversion capability and grooming-capability scenarios are very practical and should be considered when solving the traffic-grooming problem.

In this work, we employ a novel graph model for dynamic traffic grooming in a WDM mesh network. The model consists of an auxiliary graph and a dynamic traffic-grooming algorithm. Various network constraints, such as the number of transceivers at each node, the number of wavelengths on each fiber-link, wavelength-conversion capabilities and grooming capabilities of each node, are represented by different edges in the auxiliary graph. An on-line grooming algorithm is developed based on the auxiliary graph, which applies a shortest-path routing algorithm to the auxiliary graph for each traffic demand and updates the auxiliary graph accordingly. This model can achieve various objectives under different grooming policies. Instead of designing a route-computation algorithm for each grooming policy, simple shortest-path route-computation algorithms can be used to achieve various objectives by carefully choosing the weight functions for the edges in the auxiliary graph. Four different fixed grooming policies are proposed and their performances are compared. We also propose an adaptive grooming policy (AGP) that dynamically adjusts the weight functions of the edges in the auxiliary graph according to the current network state. We compare the performance of AGP to those of the fixed grooming policies.

In [5], we first proposed this novel graph model and studied static traffic grooming using this model. In our present work, we use this graph model to solve the dynamic traffic-grooming problem and evaluate the performance of different fixed grooming policies under dynamic environment. Another contribution of our present work is the new proposal for an adaptive grooming policy (AGP).

The paper is organized as follows. In Section II, we demonstrate how to construct an auxiliary graph according to the network state. Based on the auxiliary graph, a dynamic traffic-grooming algorithm is proposed and an illustrative example is given in Section III. Four fixed grooming policies are proposed in Section IV and their performance is compared through numerical examples in Section V. We propose an adaptive grooming policy (AGP) and compare it with the fixed grooming policies in Section V. Section VI concludes the paper.

II. Construction of an Auxiliary Graph

In order to solve the dynamic-grooming problem, we first construct an auxiliary graph according to the given network configuration.

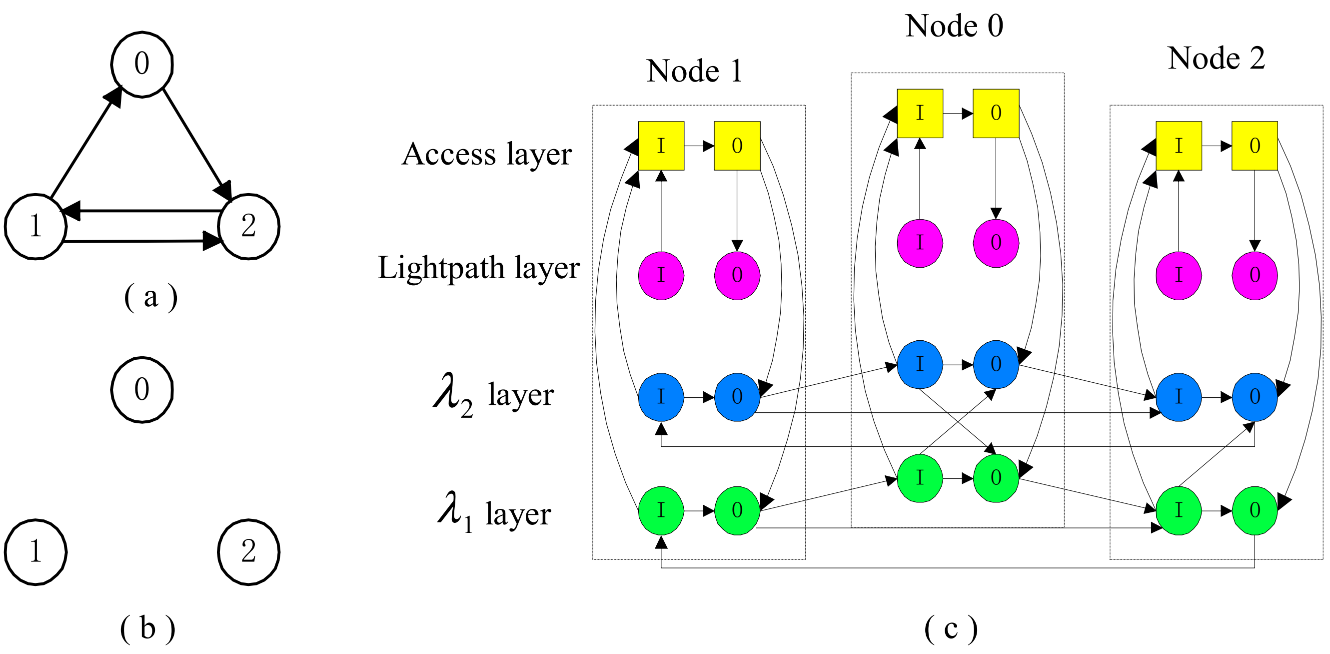

Fig. 1 (a) Physical topology of Network 1. (b) Virtual topology of Network 1. (c) Auxiliary graph of Network 1.

An illustrative example is shown in Fig. 1. In order to make the constructed auxiliary graph clear enough, we choose a very simple network topology. Network 1 (Fig. 1(a)) is a three-node network with four unidirectional fiber-links, each of which has two wavelengths. All nodes are assumed to have grooming functionality. Node 0 has wavelength converters with full wavelength-conversion capability, node 1 has no wavelength converter, and node 2 has wavelength converters with limited wavelength-conversion capability in the sense that only wavelength λ_1 can be converted to λ_2. At the beginning, there is no lightpath in the network, so there is no edge in the virtual topology of Network 1, as shown in Fig. 1(b). An auxiliary graph is constructed as in Fig. 1(c).

In general, a network can be represented by a graph $G_0(V_0, E_0)$, where V_0 and E_0 are its node set and link set, respectively. We construct the corresponding auxiliary graph $G(V, E)$ with vertex set V and edge set E. Assuming that each link has W wavelengths, λ_1 through λ_W, $G(V, E)$ is constructed as follows.

Auxiliary graph G is a layered graph with (W+2) layers. Layers 1 through W denote the W *wavelength layers*, layer (W+1) is called the *lightpath layer*, and layer (W+2) is called the *access layer*, where a traffic flow starts and terminates. Each node has two ports on each layer, denoted by two vertices, an input port (a vertex marked with "I") and an output port (a vertex marked with "O"). The edges are inserted in auxiliary graph G as follows.

- *Wavelength Bypass Edges* (*WBE*). There is an edge from the input port to the output port on each wavelength layer l at node i, denoted as $WBE(i, l)$.
- *Grooming Edges* (*GrmE*). There is an edge from the input port to the output port on access layer at node i if node i has grooming capability, denoted as $GrmE(i)$.
- *Mux Edges* (*MuxE*). There is an edge from the output port on the access layer to the output port on the lightpath layer at each node, denoted as $MuxE(i)$.
- *Demux Edges* (*DmxE*). There is an edge from the input port on the lightpath layer to the input port on the access layer at each node, denoted as $DmxE(i)$.
- *Transmitter Edges* (*TxE*). There is an edge from the output port on the access layer to the output port on wavelength layer l, denoted as $TxE(i, l)$, if there are transmitters available on wavelength λ_l at node i.
- *Receiver Edges* (*RxE*). There is an edge from the input port on wavelength layer l to the input port on the access layer, denoted as $RxE(i, l)$, if there are receivers available on wavelength λ_l at node i.
- *Converter Edges* (*CvtE*). There is an edge from the input port on wavelength layer l_1 to the output port on wavelength layer l_2 at node i, denoted as $CvtE(i, l_1, l_2)$, if wavelength λ_{l_1} can be converted to wavelength λ_{l_2} at node i. For example (Fig. 1), there are two converter edges at node 0 because it has full wavelength-

conversion capability; however, there is only one such edge at node 2 because only wavelength λ_1 can be converted to λ_2, and no such edge at node 1 because it does not have wavelength-conversion capability.

- *Wavelength-Link Edges* (*WLE*). There is an edge from the output port on wavelength layer l at node i to the input port on wavelength layer l at node j, denoted as $WLE(i, j, l)$, if there is a physical link from node i to node j and wavelength λ_l on this link is not used.
- *Lightpath Edges* (*LPE*). There is an edge from the output port on the lightpath layer at node i to the input port on the lightpath layer at node j, denoted as $LPE(i, j)$, if there is a lightpath from node i to node j. There is no such edge in Fig. 1(c) because there is no lightpath set up yet.

Each edge in the auxiliary graph G has a property tuple $P(c,w)$ associated with it, where c denotes the capacity of this edge and w denotes its weight. For a wavelength-link edge, its capacity is the capacity of the corresponding wavelength on the corresponding link. For a lightpath edge, its capacity is the residual capacity of the corresponding lightpath. For all the other types of edges, we set the capacity to ∞. The weights of edges can reflect the cost of each network element (transceiver, wavelength-link, wavelength converter, etc.), and/or a certain grooming policy. The weights can either be fixed, or be adjusted according to the current network state. A fixed weight assignment reflects a fixed grooming policy, while an adjustable weight assignment reflects an adaptive grooming policy.

From the above procedure, it should be clear that the auxiliary graph reflects the current state of the network, and the network can be heterogeneous, with different nodes having different resources and capabilities.

III. DYNAMIC TRAFFIC-GROOMING ALGORITHM

Based on the auxiliary graph, we develop a dynamic-grooming algorithm. The input includes the initial network state and a set of traffic demands, each of which has different arrival and departure time, and can be represented by $T(s, d, g, m)$, where s and d are the source and destination nodes, respectively; g is the granularity of the traffic demand, for instance, OC-48; and m is the amount of the traffic in units of g. The algorithm works as follows.

- Initialization:

  Construct the corresponding auxiliary graph G according to the initial network state.
- When a connection request T arrives:

Step 1: Compute the shortest path p from the output port on the access layer of the source to the input port on the access layer of the destination of T on graph G, ignoring the edges whose capacities are less than the requirement of the request. If such a path does not exist, block the traffic demand; otherwise, continue with the following steps.

Step 2: If p contains wavelength-link edges, set up one or more lightpaths going through the corresponding wavelength-links.

Step 3: Route T along the pre-existing lightpaths in p and/or lightpaths newly set up according to p.

Step 4: Update graph G as follows:

- For each newly setup lightpath, a lightpath edge from the output port of the starting node of the lightpath to the input port of the ending node of the lightpath is added on the lightpath layer.
- The wavelength-link edges used by the lightpath are removed from the corresponding wavelength layers.
- If there is no more transmitter/receiver available at node i on wavelength λ_l, the corresponding transmitter/receiver edge will be removed from G, i.e., this node cannot source/sink a lightpath on wavelength λ_l any more and can only be bypassed by a lightpath.
- If there is no more wavelength converter which can convert wavelength λ_{l_1} to wavelength λ_{l_2} available at node i, the converter edge will be removed from G.
- Update the property tuple $P(c, w)$ of the edges. For the lightpaths carrying the traffic T, the capacities of the corresponding lightpath edges are decreased by the amount of the carried traffic. How to update the weights of the edges in the graph is determined by the grooming policies, which will be addressed in Section IV.

- When a connection terminates:

Step 1: Remove the traffic from the network.

Step 2: Tear down all the lightpaths that do not carry any traffic.

Step 3: Update graph G by applying the reverse of the update method used in Step 4 above. We omit the details for terseness.

It can be observed that the algorithm routes a given traffic request under the current network state and updates the network state after routing and terminating the traffic, making the auxiliary graph always reflect the current network state.

A. An Illustrative Example

To illustrate how the graph model and the algorithm work, we give an example based on the network in Fig. 1. Suppose the capacity of each wavelength is OC-48 and each node has grooming capability and two tunable transceivers.

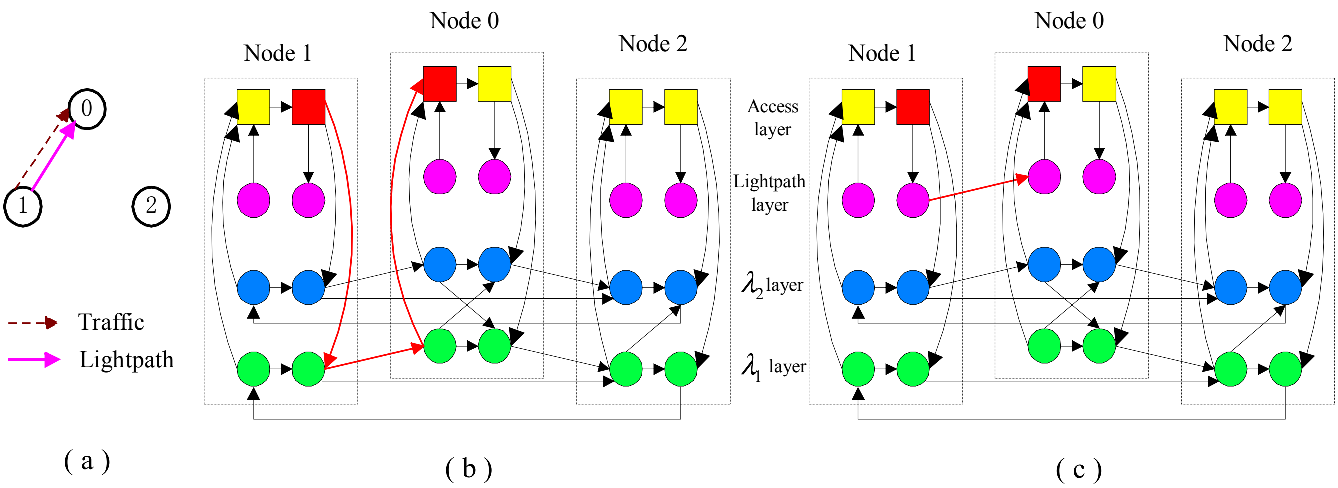

Fig. 2 (a) Virtual topology of Network 1. (b) Corresponding auxiliary graph before routing the first traffic request T_1. (c) Corresponding auxiliary graph after routing the first traffic request T_1.

The first connection request T_1 is *T*(1, 0, OC-12, 2). To satisfy this request, we need to find in the auxiliary graph a path from the output port on the access layer at node 1 to the input port on the access layer at node 0, shown as red ports in Figs. 2(b) and 2(c). It is easy to see that there exists a path along the edges *TxE*(1,1), *WLE*(1,0,1), and *RxE*(0,1), shown as red lines in Fig. 2(b). Since this path contains a wavelength-link edge *WLE*(1,0,1), which denotes a wavelength-link, we need to set up a lightpath L_1 using λ_1 on the fiber-link from node 1 to node 0. After setting up L_1, we need to add a lightpath edge *LPE*(1,0) into the graph, shown as the red line in Fig. 2(c), which means that there is a lightpath from node 1 to node 0. Meanwhile, the wavelength-link edge *WLE*(1,0,1) must be removed from the graph since this wavelength-link cannot be used to set up another lightpath later on. This connection T_1 then can be routed onto lightpath L_1 and the residual capacity of L_1 is 2 × OC-12. So the capacity of edge *LPE*(1,0) is 24, which means that the capacity is equivalent to 24 OC-1s. The current virtual topology and the updated auxiliary graph are shown in Figs. 2(a) and 2(c), respectively.

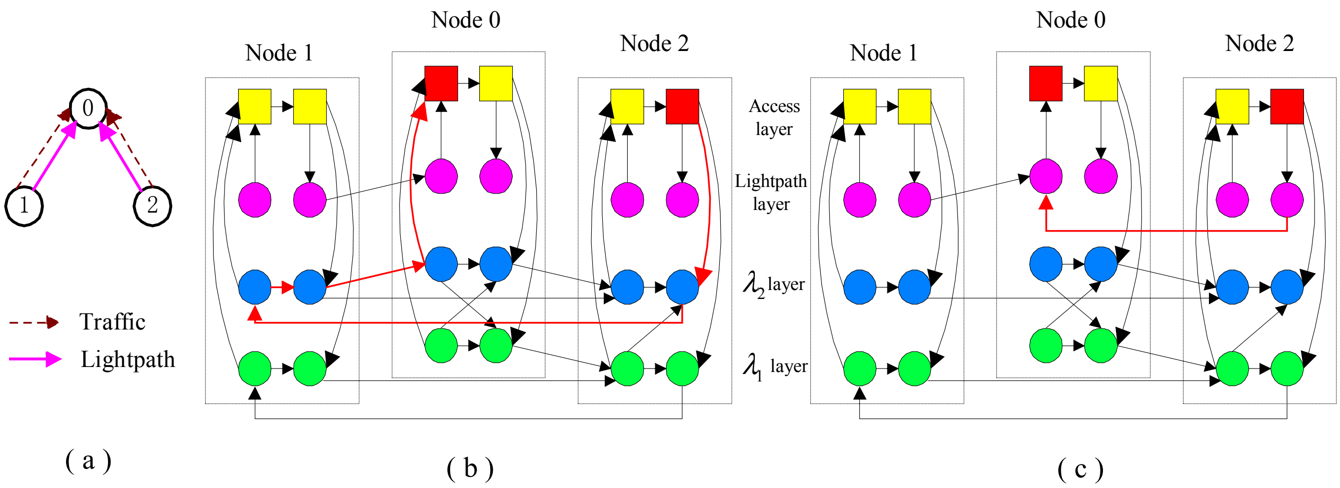

Fig. 3 (a) Virtual topology of Network 1. (b) Corresponding auxiliary graph before routing the second traffic request T_2. (c) Corresponding auxiliary graph after routing the second traffic request T_2 using single-hop grooming.

Suppose the second connection request T_2 is *T*(2, 0, OC-12, 1). Following the same procedure as above, we need to determine a path from the output port on the access layer at node 2 to the input port on the access layer at node 0, shown as the red ports in Figs. 3(b) and 3(c). There exist several paths in the auxiliary graph.

- Case 1 (Single-hop grooming). One path is along the edges *TxE*(2,2), *WLE*(2,1,2), *WBE*(1,2), *WLE*(1,0,2), and *RxE*(0,2), shown as red lines in Fig. 3(b). This path contains edges *WLE*(2,1,2) and *WLE*(1,0,2), which denote wavelength λ_2 on the fiber-links from node 2 to node 1 and from node 1 to node 0, respectively. If this path is chosen, a lightpath L_2 consisting of these two wavelength-links needs to be set up. As a result, a lightpath edge *LPE*(2,0) is added into the graph, shown as the red line in Fig. 3(c), and the two wavelength-link edges *WLE*(2,1,2) and *WLE*(1,0,2) are removed from the graph. Since both receivers at node 0 are used, we remove all the receiver edges, i.e., edges *RxE*(0,1) and *RxE*(0,2), which means that node 0 cannot sink lightpaths any more. After the traffic is routed onto lightpath L_2, the capacity of lightpath edge *LPE*(2,0) is 36 units. In this case, we set up one lightpath using two wavelength-links. Since the connection traverses a single lightpath, we call this approach *single-hop grooming*. Figures 3(a) and 3(c) show the current virtual topology and the updated auxiliary graph, respectively.

- Case 2 (Multi-hop grooming). Another path is along the edges *TxE*(2,1), *WLE*(2,1,1), *RxE*(1,1), *GrmE*(1), *MuxE*(1), *LPE*(1,0), and *DmxE*(0), shown as red lines in Fig. 4(b). This path contains edges *WLE*(2,1,1) and *LPE*(1,0), which denote wavelength λ_1 on the fiber-link from node 2 to node 1 and the lightpath from node 1 to node 0, respectively. If choosing this path, we need to set up a lightpath L_3 from node 2 to node 1 using wavelength λ_1 on the fiber-link from node 2 to node 1, and a lightpath edge *LPE*(2,1) is added, shown as the red line in Fig. 4(c), and wavelength-link edge *WLE*(2,1,1) removed. Then we route T_2 onto the newly setup lightpath L_3 and the pre-existing lightpath L_3. The capacities of lightpath edge *LPE*(2,1) and *LPE*(1,0) are 36 and 12, respectively. In this case, we have to route the connection onto two lightpaths, but only one more wavelength-link is required for satisfying this traffic. Since the connection traverses multiple lightpaths, we call this approach *multi-hop grooming*. However, this kind of multi-hop grooming will add burden on the electrical devices, which are the bottleneck and major cost in a WDM network, at the intermediate node(s) (node 1 in this case). Figures 4(a) and 4(c) show the current virtual topology and the updated auxiliary graph, respectively.

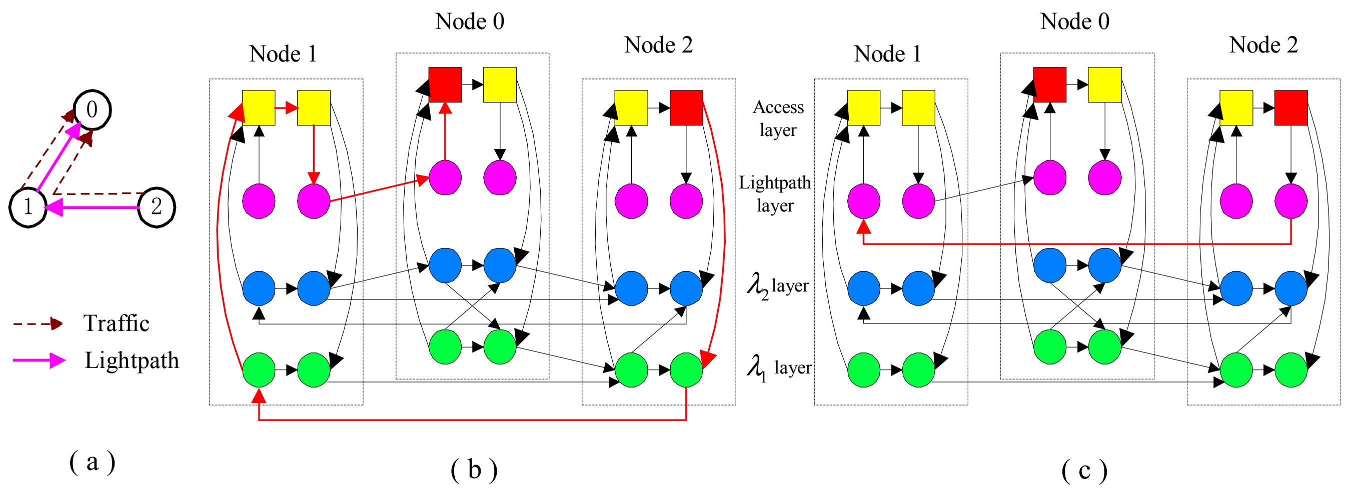

Fig. 4 (a) Virtual topology of Network 1. (b) Corresponding auxiliary graph before routing the second traffic request T_2. (c) Corresponding auxiliary graph after routing the second traffic request T_2 using multi-hop grooming.

Which path should be chosen depends on the grooming policy. Since our algorithm chooses the shortest path, the weights of the edges in the auxiliary graph will determine how to carry a connection in the network. Therefore, the grooming policy should be reflected in the weight-assignment function. We will discuss the grooming policies in Section IV.

Suppose the third traffic demand T_3 is T(1,0,OC-48,1). If we use single-hop grooming for the second connection, we cannot find a path from the output port on the access layer at node 1 to the input port on the access layer at node 0 after removing all the lightpath edges since they cannot accommodate this traffic request and it will be blocked. However, if we use multi-hop grooming for the second connection, we can still find a path in the graph since there is a wavelength λ_2 available which can be used to set up a lightpath from node 1 to node 0 to carry the T_3 traffic.

IV. GROOMING POLICIES

A grooming policy determines how to carry the traffic in a certain situation. In general, for a traffic demand $T(s, d, g, m)$ in a network, there are four possible operations that can be used to carry the traffic without altering the existing lightpaths. Note that we do not consider reconfiguring existing lightpaths because, then, the traffic on the network would be interrupted.

- *Operation 1*: Route the traffic onto an existing lightpath directly connecting the source *s* and the destination *d*.
- *Operation 2*: Route the traffic through multiple existing lightpaths.
- *Operation 3*: Set up a new lightpath directly between the source *s* and the destination *d* and route the traffic onto this lightpath. Using this operation, we set up only one lightpath if the amount of the traffic is less than or equal to the capacity of the lightpath.
- *Operation 4*: Set up one or more lightpaths that do not directly connect source *s* and destination *d*, and route the traffic onto these lightpaths and/or some existing lightpaths. Using this operation, we need to set up at least one new lightpath. However, since some existing lightpaths may be utilized, the number of wavelength-links used to set up the new lightpaths could be less than the number of wavelength-links needed to set up a lightpath directly connecting source *s* and destination *d*.

Each operation has certain prerequisites for it to be applied. For instance, if there is no lightpath between the source and the destination that can accommodate the traffic, then Operation 1 cannot be used. In some situations, all the operations are applicable, while in other situations, only some of them are. If none of them can be applied, the traffic must be blocked without reconfiguring the existing lightpaths.

In a situation where multiple operations can be applied, how to choose the appropriate operation is a matter of the grooming policy. By combining the various operations in different priority order, we can achieve different grooming policies.

Below, we present four different grooming policies.

- *Minimize the Number of Traffic Hops on the Virtual Topology (MinTHV)*

  We first use Operation 1. If Operation 1 fails, we always try to set up a lightpath from *s* to *d* and route the traffic onto this lightpath (Operation 3). Only when such a direct lightpath cannot be set up, we use multi-hop grooming by either Operation 2 or Operation 4, and choose the one with fewer hops on the virtual topology (number of lightpaths). This policy chooses the route with the fewest lightpaths for a connection.

- *Minimize the Number of Traffic Hops on the Physical Topology (MinTHP)*

  We compare the number of wavelength-links used by all the four operations and choose the one with the fewest wavelength-links.

- *Minimize the Number of Lightpaths (MinLP)*

  This policy is similar to MinTHV but it tries to set up the minimal number of *new* lightpaths to carry the traffic. Operation 1 is attempted first. If it fails, we try to route the traffic using multiple existing lightpaths (Operation 2). If Operation 2 also fails, we try to set up one lightpath with the minimal number of wavelength-links either by Operation 3 or Operation 4. If such a lightpath is not feasible, we go with Operation 4 and set up two or more lightpaths.

- *Minimize the Number of Wavelength-Links (MinWL)*

  This policy is similar to MinTHP but it tries to consume the minimum number of *extra* wavelength-links, i.e., wavelength-links not being used by any lightpaths for now, to carry the traffic. The difference between MinLP and MinWL is that, if both Operations 1 and 2 fail, MinWL compares the number of wavelength-links used by Operations 3 and 4, and chooses the one requiring fewer wavelength-links; MinLP, on the other hand, compares the number of lightpaths used by Operations 3 and 4, chooses the one requiring fewer lightpaths, and uses the number of wavelength-links for tie-breaking.

Our graph model can easily implement these grooming policies by applying different weight-assignment functions. We call an edge in an auxiliary graph the *dominant edge* if this edge satisfies the following condition: if a path p_1 in the graph contains more of this kind of edges than another path p_2, then the weight of p_1 is always larger than that of p_2. Here, the weight of a path is the summation of the weights of the edges it traverses. To achieve MinTHV, we just need to make GrmEs the dominant edges. To achieve MinLP, we should make TxEs and RxEs the dominant edges. To achieve MinWL, WLEs should be the dominant edges. In all the above three policies, the weight of the LPEs is fixed and less than that of WLEs. To achieve MinTHP, we also make WLEs the dominant edges. In addition, a LPE is considered as a concatenation of WLEs whose corresponding wavelength-links are used by the lightpath represented by the LPE in the auxiliary graph, and the weight of a LPE is the summation of the weight of those WLEs.

In dynamic grooming, the network state changes as connection requests come and go. To achieve good performance, the grooming policy should be adjusted according to the current network state. For instance, if transceivers are becoming the more scarce resource, we should make full use of existing lightpaths to accommodate the new traffic and avoid setting up new lightpaths. This requirement can be easily satisfied by modifying the weights of edges according to the current network state. This capability of easily adjusting grooming policies makes the graph model very suitable for dynamic traffic grooming.

V. NUMERICAL EXAMPLES

We compare the performance of different grooming policies on the network topology shown in Fig. 5, which has 19 nodes and 31 links. All the nodes have grooming capability but no wavelength-conversion capability. Each link is bidirectional with W=16 wavelengths in each direction, and each wavelength has a capacity of OC-192. The traffic arrival is a Poisson process and the connection holding time is exponentially distributed (whose average value is normalized to unity in our studies reported here). The traffic is uniformly distributed among all node pairs. There are four types of connection requests: OC-3, OC-12, OC-48, and OC-192, and the proportion of the number of these connections is 6:6:6:1. For a connection request $T(s, d, g, m)$, m, the amount of the traffic in unit of g, is uniformly distributed between 1 and 32, 1 and 16, 1 and 8, and 1 and 2 for OC-3, OC-12, OC-48, and OC-192 types of connections, respectively. We simulate 100,000 connection requests to obtain the network performance under a certain scenario and a grooming policy. We ran our simulation experiments on a Linux PC with a 1.5 GHz Pentium IV processor and 512 MB memory. Each data point reported in the illustrations in this section took between 6-9 minutes of running time on this computer.

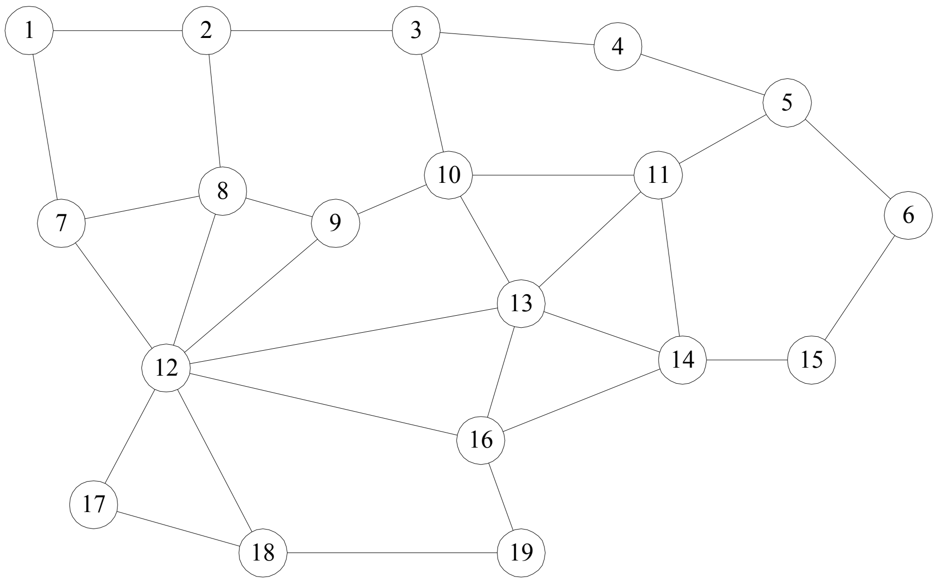

Fig. 5 A 19-node telecom network.

Table 1 Average utilization of wavelength-links and transceivers when *W*=16 and *L*=300 Erlangs.

	Tx = 16		Tx = 32		Tx = 40	
	U_W	U_{Tx}	U_W	U_{Tx}	U_W	U_{Tx}
MinTHV	0.7819	0.9858	0.8878	0.7264	0.8905	0.5884
MinTHP	0.5674	0.9901	0.7354	0.8165	0.7361	0.6910
MinLP	0.7403	0.9807	0.8890	0.8007	0.8918	0.6651
MinWL	0.6201	0.9859	0.8133	0.8683	0.8120	0.7825

Table 1 shows the average utilization of wavelength-links (U_W) and the average utilization of transceivers (U_{Tx}) when the network load L is 300 Erlangs. When each node has only 16 transceivers, the utilization of transceivers is very high since they are the more constrained resources. When there are 32 transceivers at each node, the utilizations of both transceivers and wavelengths are quite balanced and high as well. If there are 40 transceivers at each node, we have relatively more transceivers; hence, wavelength-links become the more constrained resources, so the utilization of the transceivers is relatively lower. Fig. 6 shows the network performance under different grooming policies with *Tx*=32 transceivers at each node. We observe that, as the network load increases, the percentage of blocked traffic also increases, but different grooming policies have different blocking probabilities. Grooming policy MinTHV performs best, followed by MinTHP, MinLP, and MinWL in sequence.

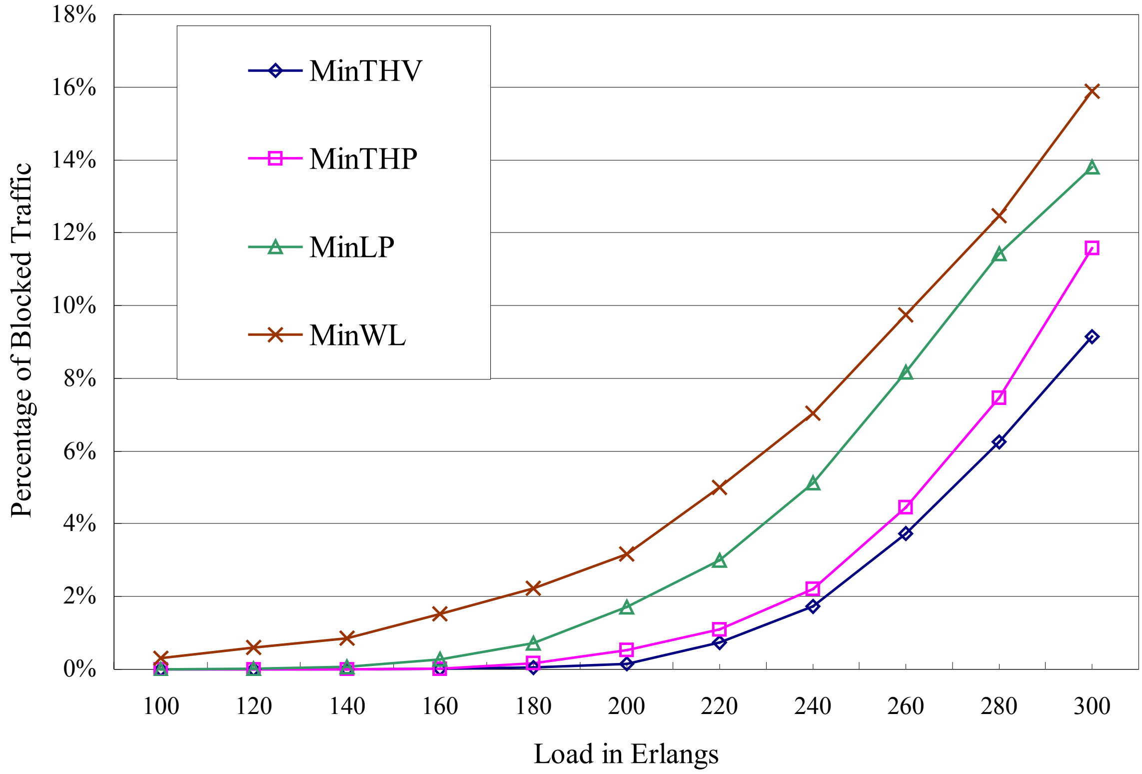

Fig. 6 Percentage of blocked traffic when *Tx*=32.

When we alter the network configuration by adding more transceivers at each node, the performance of each of the policies also changes, as shown in Fig. 7. In this scenario, each node has 40 transceivers instead of 32. We observe that, now, MinTHP outperforms MinTHV and achieves the best results, and MinLP becomes the poorest-performing policy. This is because, in this network configuration, there are relatively more transceivers in the network so that wavelength-links become the more constrained resources. Recall that MinTHP utilizes wavelength-links more efficiently than other policies; hence, it performs the best in this case.

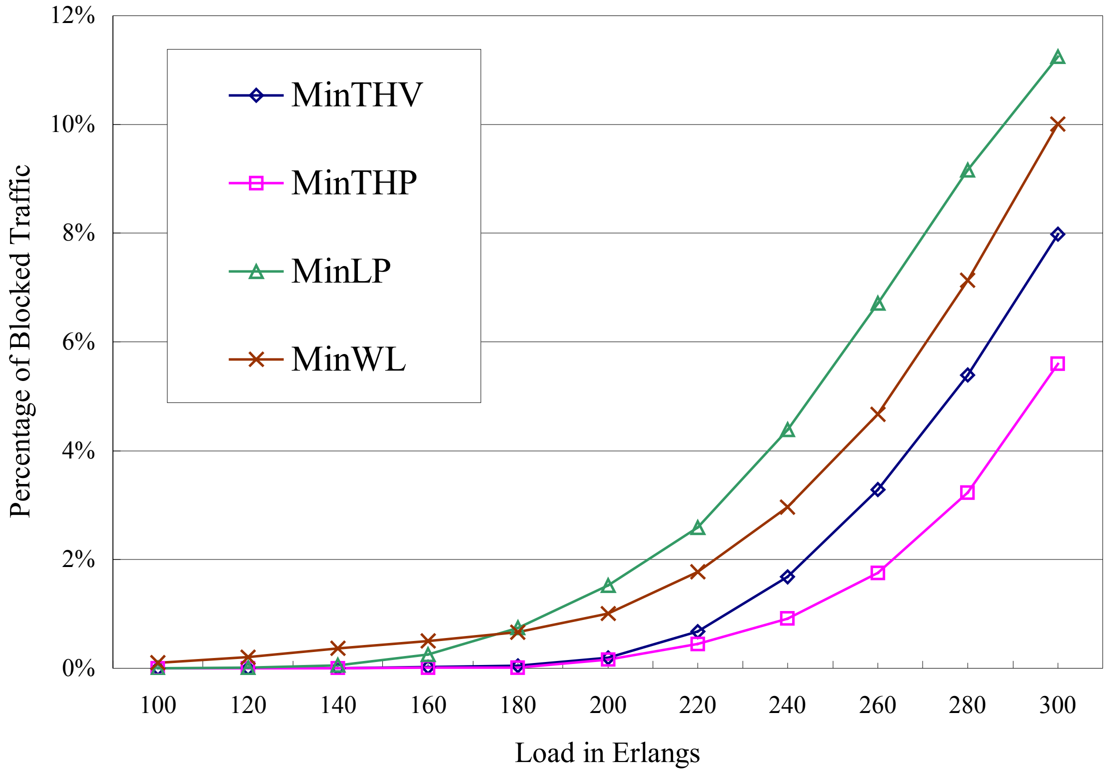

Fig. 7 Percentage of blocked traffic when *Tx*=40.

Another observation from both Figs. 6 and 7 is that MinTHV and MinTHP always perform better than MinLP and MinWL in terms of percentage of blocked traffic. This is because MinTHV and MinTHP examine the *overall* resource requirement of a given connection, while MinLP and MinWL only consider the new lightpaths to be set up or the extra wavelength-links used by these new lightpaths while setting up the connection. Therefore, MinTHV and MinTHP are more resource-efficient.

From the above results, we can observe that different grooming policies have different performance under various network configurations, which suggests that a grooming policy should be adjusted according to the current network state.

A. Adaptive Grooming Policy (AGP)

Since MinTHV performs best when transceivers are the more constrained resources and MinTHP gives the best results when wavelength-links become more scarce resources, we try to utilize the advantages of these two grooming policies by combining them together. Here we present an *Adaptive Grooming Policy* (*AGP*) which, for each connection request, switches between MinTHV and MinTHP according to the current network state.

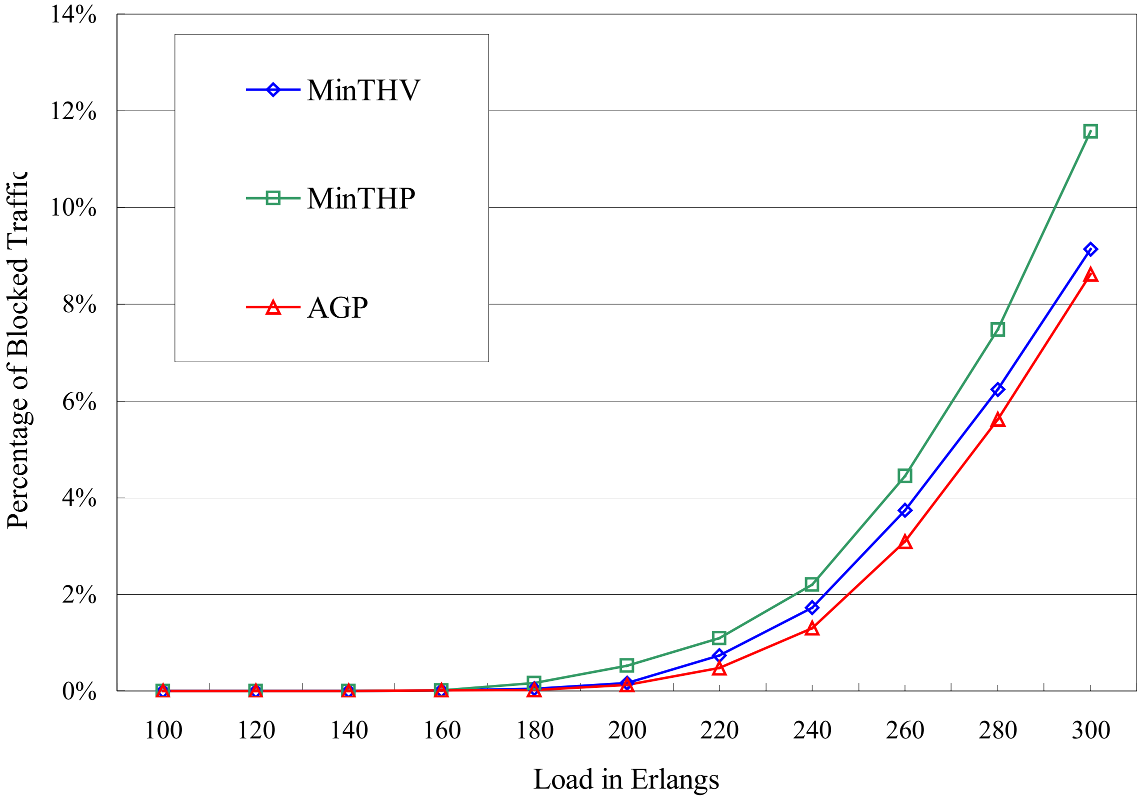

Fig. 8 Performance of AGP when *Tx*=32.

We use the ratio of the number of unused wavelength-links in the network to the total number of available transceivers at all nodes as an indicator of the network state. If the ratio is larger than δ_1, MinTHV will be used to avoid setting up lightpaths since transceivers are more scarce resources at this time; if the ratio is less than δ_2, MinTHP will be employed to try to save wavelength-links as much as possible; if the ratio is in between, the policy will not be changed.

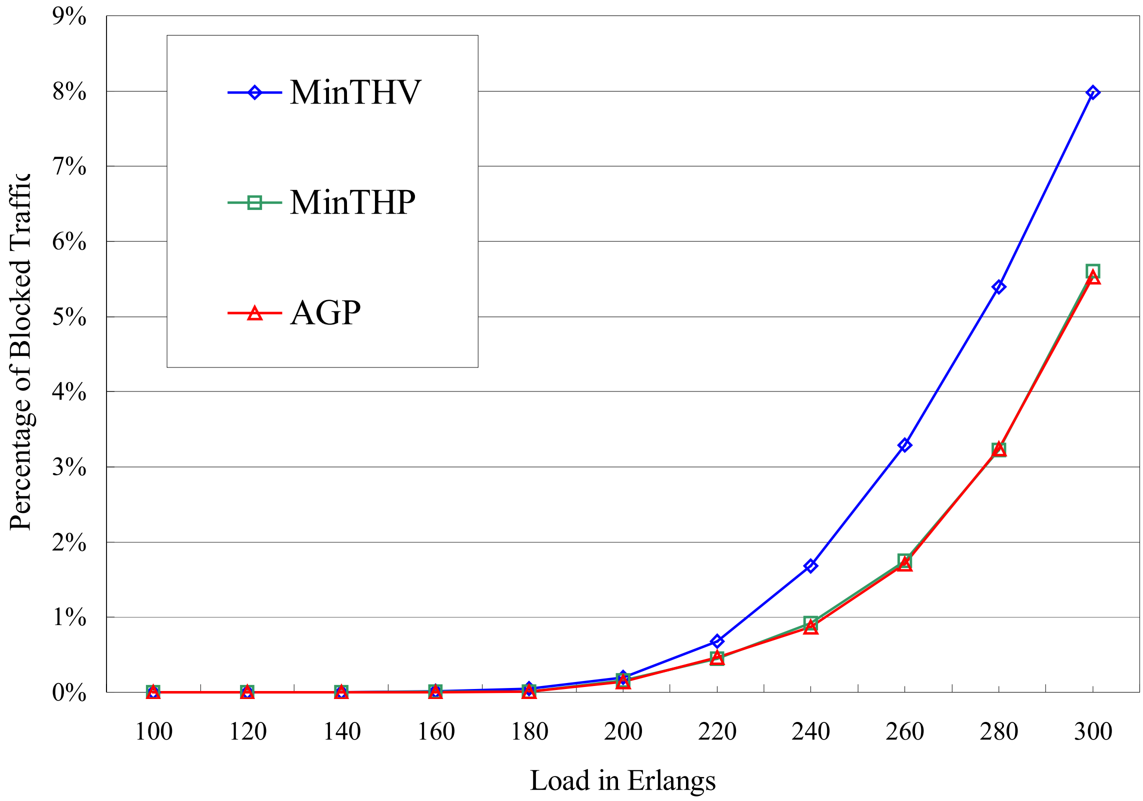

Fig. 9 Performance of AGP when *Tx*=40.

For our numerical examples, we report results for $\delta_1 = 1.2$ and $\delta_2 = 1.0$. (We experimented with other combinations of these two parameters, and found these choices of values to perform the best for the network topology in Fig. 5.) Our results are shown in Figs. 8 and 9. We observe that the Adaptive Grooming Policy (AGP) achieves the best results under different network configurations.

VI. Conclusion

In this study, we used a novel graph model for dynamic traffic grooming in a WDM optical mesh network. This model takes various constraints into account and can achieve various objectives by using different grooming policies. Moreover, the ability to easily adjust the grooming policy by manipulating the weights of the edges in the auxiliary graph according to the current network state make this model very suitable for dynamic traffic grooming. We proposed four different fixed grooming policies and compared their performance in various network scenarios for dynamic traffic grooming. We also developed an Adaptive Grooming Policy (AGP), which switches between different grooming policies according to the current network state. Our results show that different grooming policies have various performance under different network configurations and AGP can achieve the best performance.

REFERENCES

[1]. B. Mukherjee, *Optical Communication Networks*. New York: McGraw-Hill, 1997.

[2]. E. Modiano and P. Lin, "Traffic grooming in WDM networks," *IEEE Communications Magazine*, vol. 39, pp. 124–129, July 2001.

[3]. K. Zhu and B. Mukherjee, "Traffic grooming in WDM optical mesh networks," *IEEE Journal on Selected Areas in Communications*, vol. 20, pp. 122–133, January 2002.

[4]. K. Zhu and B. Mukherjee, "On-line approaches for provisioning connections of different bandwidth granularities in WDM mesh networks," in *Proc., OFC '02*, Anaheim, CA, pp. 549–551, March 2002.

[5]. H. Zhu, H. Zang, K. Zhu, and B. Mukherjee, "A novel, generic graph model for traffic grooming in heterogeneous WDM mesh networks," submitted for publication.

[6]. K. Zhu and B. Mukherjee, "A review of traffic grooming in WDM optical networks: Architectures and challenges," *SPIE Optical Networks Magazine*, to appear, 2002.

[7]. A. L. Chiu and E. H. Modiano, "Traffic grooming in algorithms for reducing electronic multiplexing costs in WDM ring networks," *IEEE/OSA Journal of Lightwave Technology*, vol. 18, pp. 2–12, January 2000.

[8]. 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*, vol. 18, pp. 1995--2003, October 2000.

[9]. X. Zhang and C. Qiao, "An effective and comprehensive approach for traffic grooming and wavelength assignment in SONET/WDM rings," *IEEE/ACM Transactions on Networking*, vol. 8, pp. 608–617, October 2000.

[10]. O. Gerstel, R. Ramaswami, and G. H. Sasaki, "Cost-effective traffic grooming in WDM rings," *IEEE/ACM Transactions on Networking*, vol. 8, pp. 618–630, October 2000.

[11]. 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*, vol. 17, pp. 48--57, January 1999.

[12]. R. Dutta and G. N. Rouskas, "On optimal traffic grooming in WDM rings," *IEEE Journal on Selected Areas in Communications*, vol. 20, pp. 110--121, January 2002.

[13]. J. Wang, W. Cho, V. R. Vemuri, and B. Mukherjee, "Improved approaches for cost-effective traffic grooming in WDM ring networks: ILP formulations and single-hop and multihop connections," *IEEE/OSA Journal of Lightwave Technology*, vol. 19, pp. 1645–1653, November 2001.

[14]. J. Wang and B. Mukherjee, "Interconnected WDM ring networks: Strategies for interconnection and traffic grooming," *SPIE Optical Networks Magazine*, to appear, 2002.

[15]. R. Berry and E. Modiano, "Reducing electronic multiplexing costs in SONET/WDM rings with dynamically changing traffic," *IEEE Journal on Selected Areas in Communications*, vol. 18, pp. 1961--1971, October 2000.

[16]. S. Thaigaran and A. Somani, "Capacity fairness of WDM networks with grooming capabilities," *SPIE Optical Networks Magazine*, vol. 2, pp. 24--31, May/June 2001.

[17]. R. Srinivasan and A. K. Somani, "A generalized framework for analyzing time-space switched optical networks," *IEEE Journal on Selected Areas in Communications*, vol. 20, pp. 202--215, January 2002.

[18]. S. Thiagarajan and A. Somani, "Traffic grooming for survivable mesh networks," in *Proc., OPTICOMM '01*, Denver, CO, August 2001.

[19]. L. A. Cox and J. Sanchez, "Cost savings from optimized packing and grooming of optical circuits: Mesh versus ring comparisons," *SPIE Optical Networks Magazine*, vol. 2, pp. 72--90, May/June 2001.

[20]. A. Lardies, R. Gupta, and R. A. Patterson, "Traffic grooming in a multi-layer network," *SPIE Optical Networks Magazine*, vol. 2, pp. 91--99, May/June 2001.

[21]. C. Xin, Y. Ye, S. Dixit, and C. Qiao, "An integrated lightpath provisioning approach in mesh optical networks," in *Proc., OFC '02*, Anaheim, CA, pp. 547–549, March 2002.