Dynamic Provisioning of Low-Speed Unicast/Multicast Traffic Demands in Mesh-Based WDM Optical Networks

Ahmad Khalil, Antonis Hadjiantonis, Chadi M. Assi, Abdallah Shami, *Member, IEEE*, George Ellinas, and Mohamed A. Ali

***Abstract*—This paper addresses the problem of dynamically provisioning both low-speed[1] unicast and multicast connection requests in mesh-based wavelength division multiplexing (WDM) optical networks. Several routing/provisioning schemes to dynamically provision both unicast and multicast connection requests are presented. In addition, a constraint-based grooming strategy is devised to utilize the overall network resources as efficiently as possible. Based on this strategy, several different sequential multicast grooming heuristics are first presented. Then, we devise a hybrid grooming approach and combine it with sequential approaches to achieve a grooming scheme that is biased toward serving multicast traffic demands in comparison with all other sequential grooming approaches. To achieve our objective, we decompose the problem into four subproblems: 1) routing problem; 2) light-tree-based logical-topology-design problem; 3) provisioning problem; and 4) traffic-grooming problem. The simulation results of the proposed schemes are compared with each other and with those of conventional nongrooming approaches. To the best of our knowledge, this is the first detailed paper to address and examine the problem of grooming dynamic multicast traffic demands.**

***Index Terms*—Hybrid provisioning, multicast, traffic grooming, wavelength division multiplexing (WDM) networks.**

I. Introduction

THE ABUNDANCE of bandwidth propelled by the explosion of wavelength division multiplexing (WDM) along with recent advances in optical networking technology has successfully provided the required capacity to meet the phenomenal growth in Internet traffic. Commercially available WDM transmission systems can support greater than 1 Tb/s over a single fiber, by means of multiplexing more than a hundred channels at 10 Gb/s each [1]. In a wavelength-routed network, data are transported in all-optical WDM channels (lightpaths) where the bandwidth granularity is at the full wavelength level. A lightpath is an end-to-end connection that may span a number of physical links. If no wavelength converters are used, a lightpath is associated with the same wavelength on all physical links spanning the entire path from source to destination (wavelength continuity constraint). The set of established lightpaths forms the logical topology of a WDM network. Data can be processed electronically (added/terminated) only at the endpoints of a lightpath, and switched optically (cut through) at intermediate nodes of the underlying physical topology.

A major shortcoming in current WDM networks is the large disparity between the coarse/fixed granularity bandwidth offered by the optical layer to clients [full wavelength level, e.g., OC-48 (2.5 Gb/s), OC-192 (10 Gb/s), and OC-768 (40 Gb/s)] and the bandwidth requirement of a typical connection request, which is only a fraction of a wavelength [e.g., STS-1 (51 Mb/s), OC-3 (155 Mb/s), OC-12 (622 Mb/s), etc.]. Clearly, traffic demands with finer bandwidth granularity are the rule and those requiring full wavelength capacity are the exceptions.

Therefore, in order to efficiently utilize the capacity of each wavelength channel (lightpath), several independent lower speed traffic streams (i.e., subwavelength connections) must be multiplexed onto a single lightpath. The process of combining low-rate traffic streams onto high-capacity optical channels (lightpaths) is known in the literature as "traffic grooming" [2], [3]. To support traffic grooming, the cross-connect fabric of each optical node should have the capability of switching traffic at the wavelength granularity as well as at finer granularities [2], [3].

Most early work on traffic grooming has focused on synchronous optical network (SONET) rings, where traffic is often static and known in advance [3]–[6]. More recently, traffic grooming in mesh-based WDM networks has attracted an increased amount of research effort [7]–[11]. Most of these studies have assumed only unicast traffic. However, as networks evolve to support more bandwidth-intensive applications, and as rich multimedia and real-time services become more popular, next-generation networks are expected to support both unicast and multicast applications (e.g., multiparty conferencing, software and video distribution, and distributed computing, etc.). To support multicasting at the physical layer of WDM networks, the concept of a light tree has been introduced [12]–[20]. A light tree is a point-to-multipoint extension of a lightpath, where the branching nodes of a light tree are equipped with optical power splitters.

Manuscript received November 29, 2004; revised October 2, 2005.

A. Khalil, A. Hadjiantonis, and M. A. Ali are with the Electrical Engineering Department, The City College of New York, New York, NY 10031 USA (e-mail: akhalil@ccny.cuny.edu).

C. M. Assi is with the Concordia Institute for Information Systems Engineering (CIISE), Concordia University, Montreal, QC H3G 1M8, Canada.

A. Shami is with the Department of Electrical and Computer Engineering, University of Western Ontario, London, ON N6A 5B9, Canada.

G. Ellinas is with the Department of Electrical and Computer Engineering, University of Cyprus, Nicosia, Cyprus.

Digital Object Identifier 10.1109/JLT.2005.861922

[1]A low-speed connection is a connection request with only a fraction of a wavelength capacity.

Similar to the case of unicast traffic demands, some of these multicast applications require only a fraction of the channel capacity [for example, high-definition TV (HDTV) needs only 20 Mb/s]. Thus, dedicating an entire light tree to a single user or even a few users may lead to a huge waste of network resources. To improve the network throughput, one would need to bundle several low-rate unicast and multicast traffic streams efficiently onto a single high-capacity light tree so that the number of wavelengths that have to be processed at each node is minimized. Hence, the problem of multicast traffic grooming is expected to become an important area for future research work. Furthermore, as network architectures change from ring based to mesh based, both unicast and multicast traffic grooming in mesh-based networks will become an important extension to current ring-based grooming algorithms.

Although the problem of all-optical multicasting has received considerable attention in the literature [12]–[17], however, the problem of grooming multicast traffic has received little attention and has only considered static multicast traffic [21], [22]. This paper addresses the problem of dynamically provisioning low-speed unicast and multicast connection requests in mesh-based WDM optical networks. Specifically, this work focuses on building a dynamic logical topology where lightpaths/light trees are set up and torn down in response to dynamic multicast traffic demands. We develop several routing/provisioning schemes to dynamically provision low-speed unicast/multicast connection requests. A constraint-based grooming strategy is devised to utilize the overall network resources as efficiently as possible by selecting the most appropriate combination(s) of the existing multiple routing/provisioning schemes. Based on this strategy, several different sequential multicast grooming heuristics are first presented. Then, we devise a hybrid grooming approach and combine it with sequential approaches to achieve a grooming scheme that is biased toward serving multicast traffic demands in comparison with all other sequential grooming approaches.

To achieve our objective, we decompose the problem into four subproblems: 1) routing problem; 2) design of a light-tree-based logical topology; 3) provisioning problem; and 4) traffic-grooming problem. The simulation results of the proposed schemes are compared with each other and with those of conventional nongrooming approaches. The rest of this paper is organized as follows. Section II presents a background and an overview of the multicasting problem in WDM networks along with a review of related work. Section III presents the network model and design. Different multicast grooming methodologies and algorithms as well as the used heuristics are presented in Section IV. A performance evaluation is presented in Section V and the paper is concluded in Section VI.

II. Background and Overview

A. Multicasting in WDM Networks

Multicasting is the ability to transmit a message from a single source node to multiple destination nodes. Multicasting has emerged as one of the essential features in current and future networks with the development of computer and communication applications such as distributed computing, audio and video conferencing, software and video distribution, and database replication.

The simplest approach to serve multicast traffic is to treat every multicast request as a set of separate unicast requests and route each request independently. This approach, however, results in a huge waste of network resources. Hence, efficient multicast approaches should be applied to satisfy near-optimal usage of the various network resources (i.e., link bandwidth, number of receivers, etc.). Using these multicast schemes, information is transmitted along a set of physical links, which constitute a multicast tree.

IP multicasting has been proposed to provide efficient one-to-many (many-to-many) data-delivery services [23]. More recently, IP-over-WDM multicasting becomes very attractive due to the transition of next-generation networks to WDM networks. In fact, IP multicasting can be used to support data multicasting in IP-over-WDM networks. In this scenario, each IP router on a multicast tree (constructed by the IP layer) makes copies of a data packet and transmits a copy to its immediate downstream routers. However, this requires optical/electronic/optical (O/E/O) conversions of every data packet at every router on the multicast tree, which may be inefficient due to the latency problem, and undesirable because of data transparency. Another way to support multicasting in IP-over-WDM networks is WDM multicasting, where multicasting is supported at the WDM layer by making copies of data packets in the optical domain via light splitting, avoiding the O/E/O conversions and allowing significant bandwidth savings over the shared links of the multicast tree (i.e., light tree) [13]–[15].

Supporting multicasting at the WDM layer has several advantages. The first advantage is the ability to construct more efficient multicast trees (due to the knowledge of the optical layer, which may not be the same as that seen at the upper electronic layer). The second advantage is the inherent light-splitting capability of some optical switches (optical switches that do not have an inherent light-splitting capability may be augmented with such a capability at a reasonable cost). In general, it is more effective to perform light splitting at the optical layer than copying IP datagrams in electronics (latency problem due to the O/E/O conversion and IP header processing). Furthermore, performing multicast in optics provides consistent support of format and bit-rate transparencies across both unicast and multicast transmissions [13]–[15].

Multicast routing in all-optical wavelength-routed networks implies that, given the network topology along with limited network resources, the objective is to construct a multicast tree (light tree) from the source to the destination nodes in order to lower the blocking probability (BP) and reduce the number of electronic components [12]–[20]. A light tree is an all-optical point-to-multipoint virtual connection in which the source of a light tree transmits the data on a particular wavelength, and the data reaches all the destination nodes through the use of optical splitters that can split the optical signal from a single incoming port to multiple outgoing ports [17]. Therefore, using these splitters, an optical signal can be delivered to multiple destinations along a set of physical links, which constitute a multicast tree with the source node as the root.

B. Multicast Routing and Wavelength Assignment (MC-RWA)

In wavelength-routed networks and without the presence of wavelength converters, a multicast tree requires a dedicated wavelength for each of its branches (links). This problem is referred in the literature as MC-RWA [12], [16], [20]. As in the case of RWA for unicast connections (UCs), this problem is partitioned into two independent subproblems, namely multicast-routing problem and wavelength-assignment problem.

The multicast-routing subproblem (multicast-tree problem) is often modeled as the Steiner minimum tree (SMT) problem, an NP-complete problem [24]. Many heuristics were proposed to find an approximate (suboptimal) solution to this problem [25], [26]. Once the multicast-routing subproblem is solved and the multicast tree is found, the wavelength-assignment subproblem will therefore be the problem of assigning a wavelength along all the links of the corresponding multicast tree [19], [27].

Let $G = (V, E, c)$ denote an undirected graph where V is the set of n vertices in the network, E is the set of m edges (links) in the network, and c is a positive cost function on E.

The SMT problem is defined as the problem to find a connected subgraph T of G that includes a subset of the vertices $M \subseteq G$ that minimizes $z = \sum_{e \in T} c(e)$, where $c(e)$ is the cost of the edge e. In other words, we need to find T, such that there is a path between a pair of T vertices, and the total cost of T (i.e., the sum of all its edge costs) is minimized.

In this paper, the traffic is assumed to be a combination of unicast and multicast connections according to a given ratio, and each multicast session has only one single source, but a node can be a source of multiple sessions. For each multicast session with source s, $G_{\rm p}$ represents the maximum percentage of the network nodes that could be destinations, and the number of destinations is uniformly distributed in the interval $[2, d]$, where: $d = G_{\rm p} \times n/100$, and n is the network size.

The shortest path tree (SPT) heuristic is used for multicast routing [16]. In SPT, all shortest paths from the source of a multicast session to all its destinations are calculated (using Dijkstra algorithm), and then the paths are combined to form the multicast tree. Since the time needed to calculate all the shortest paths is $O(n^2)$, the time complexity of the SPT algorithm is $O(dn^2)$. However, if we calculate all the shortest paths (from any source to any destination) *a priori*, the time complexity to calculate a multicast tree using the SPT algorithm reduces to $O(dn)$. Once the multicast tree is found, and assuming there are no wavelength converters, the wavelength-assignment subproblem becomes finding a single available wavelength on all the multicast-tree branches. In this paper, the first-fit scheme is used for wavelength assignment.

C. Review of Related Work

The problem of grooming multicast traffic in optical networks is an important problem that has received little attention given its immense practical importance [21], [22]. In [21], the authors proposed an integer linear programming (ILP) formulation in order to minimize the number of wavelength channels used and the cost of the network in terms of the number of SONET add/drop multiplexers (ADMs). In that work, the network was represented as three different levels, namely the physical, the lightpath, and the connection levels. The authors considered nonuniform static traffic; they also introduced heuristics to solve the problem by obtaining first an initial solution using SPT and first-fit wavelength assignment, and then iteratively improving it by exploring other routes. The authors in [22] formulated an optimization problem for the design of a light-tree-based logical topology. That problem consisted of two subproblems, namely the MC-RWA, and the design of a light-tree-based logical topology for multicast streams. In that work, ILP formulation was used for the design of optimum light trees and then the light-tree-based logical topology was modeled as a hypergraph over which static multicast streams were routed. To the best of our knowledge, the problem of grooming dynamic multicast traffic and designing a light-tree-based hypergraph logical topology (HGLT) for dynamic multicast traffic in WDM networks has not been previously considered.

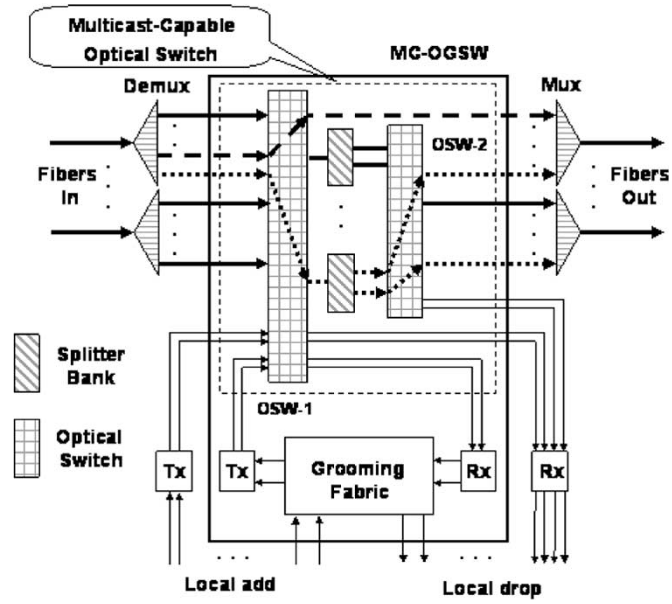

Fig. 1. MC-OGSW architecture.

III. Network Model and Design

A. Node Architecture

We consider a WDM network with n optical nodes interconnected by bidirectional fiber links, where each link carries w wavelengths and each node is equipped with a multicast-capable all-optical switch (MC-OSW) with p input ports (fibers) and p output ports. The multicast capability is supported via optical splitter banks, where a $p \times p$ optical splitter bank splits the input signal into p identical output signals [17]. Note that splitter banks may be enhanced to compensate for the power loss, wavelength conversion, and signal regeneration. A $p \times p$ MC-OSW that supports w wavelengths consists of p $(1 \times w)$ demultiplexers, p $(w \times 1)$ multiplexers, w splitter banks (one for each wavelength), and a two-stage OSW. Note that the design of the splitter banks along with the second-stage optical switch (OSW-2) ensures that the MC-OSW is strictly nonblocking (Fig. 1).

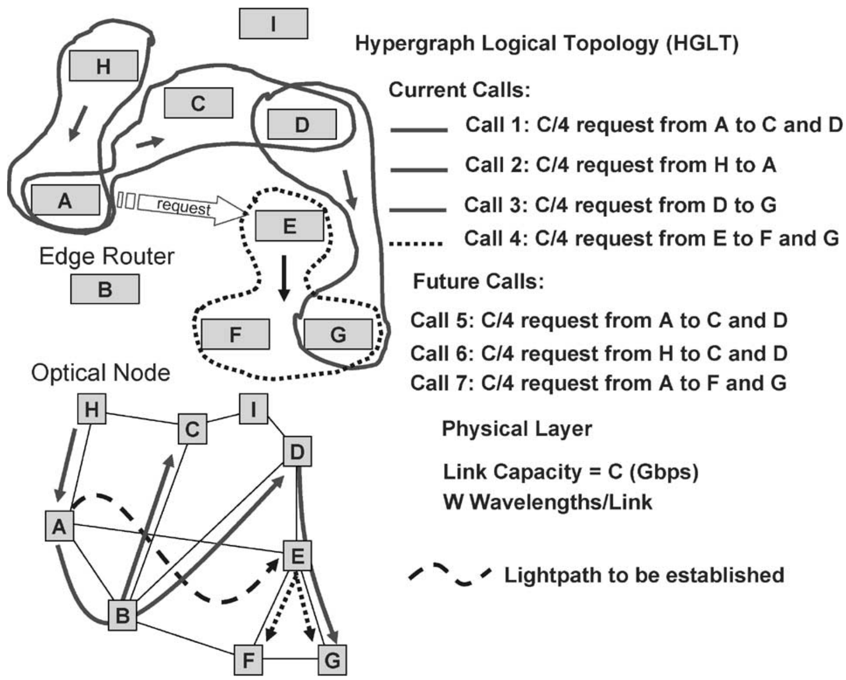

Fig. 2. HGLT design for a multicast-capable WDM network.

To support grooming of multicast traffic, the node architecture shown in Fig. 1 comprises a grooming fabric (GF). In this architecture OSW-1 is connected to the GF through a set of transponders. An optical switch with both multicast and grooming capabilities is referred to as multicast-capable optical-grooming switch (MC-OGSW). The provisioning of a lightpath/light-tree connection is enabled by the MC-OGSW, where OSW-1 provides all-optical bypass for a lightpath/light tree passing through the node without any electronic processing. Alternatively, a lightpath/light tree can be dropped if the node is the final destination of all the traffic carried by the lightpath/light tree. Otherwise, the lightpath/light tree is dropped to the GF where other traffic can be multiplexed (groomed) onto this lightpath/light tree (i.e., multihop grooming) to increase the bandwidth efficiency of wavelength channels. Upon grooming, the traffic is switched either to an output link if this is unicast traffic, or to a corresponding splitter to split the optical signal and forward it to OSW-2, which in turn routes the signals to their different output links. Note that this will allow traffic to be forwarded from one lightpath/light tree to another lightpath/light tree to reach its ultimate destination(s). This is referred to as multicast multihop grooming.

B. Design of HGLT

Unlike traditional mesh-based WDM networks where only unicast (point–point) traffic is considered, the logical topology of WDM networks with multicast capability must support point-to-multipoint traffic. While the point-to-point-based logical topologies are usually modeled by directed graphs whose edges represent the lightpaths, multicast-capable networks are best modeled by directed hypergraphs [18], [22], in which every two vertices are connected by a hyperarc that represents a light tree.

Fig. 2 shows a hypergraph with four established calls, where there are two wavelength channels available on every link (solid and dotted), and every call is a sublambda request of 25% of the channel capacity. In this example, Call#1 is from source A to destinations $\{C, D\}$, Call#2 is from source H to destination $\{A\}$, and Call#3 is from source D to destination $\{G\}$, and are all served on the solid wavelength. Call#4, from source E to destinations $\{F, G\}$, is served on the dotted wavelength. Note that electronic processing and multiplexing of traffic is performed only at the end nodes of the HGLT edges (i.e., at the source and destinations of the light trees), and the traffic is switched/split optically at intermediate nodes of the underlying physical network.

IV. Multicast Grooming Methodologies, Algorithms, and Heuristics

The multicast grooming strategy adopted in this paper has two main objectives: 1) to utilize the overall network resources as efficiently as possible; and 2) to devise grooming algorithms that favor provisioning of multicast traffic demands versus that of unicast demands (i.e., more suitable for serving multicast connection requests). To achieve our objectives, we divide the overall network topology and resources into three different sets: 1) physical topology and its associated available resources; 2) logical topology and its associated available resources; and 3) hybrid topology (combined physical and logical topologies) and its associated available resources. Having identified the various available network topologies and resources, we first devise three different sets of routing and provisioning algorithms, one for each of the three different topologies defined above. The routing/provisioning algorithms developed for each set have only limited access to their own topology and resources. We define the process of selecting the most appropriate blend from among these multiple available routing/provisioning schemes to efficiently serve a given unicast/multicast connection request, as our grooming policy. Based on this definition, we then devise a constraint-based grooming strategy that enables us to select the most appropriate blend of these multiple routing/provisioning schemes. We impose the following two constraints on the proposed grooming strategy.

1) The number of logical hops (i.e., lightpaths/light trees) a call can be routed over is limited to a maximum of two hops.
2) Since in multihop grooming a node may groom unicast and multicast traffic demands on the same output channel (i.e., same lightpath/light tree), some destinations on the multicast tree may end up receiving unintended unicast data, leading to a waste of network resources. To alleviate this problem, multihop grooming is constrained such that unicast traffic may not be groomed onto multicast traffic; however, multicast traffic can always be groomed onto unicast traffic.

A connection request can be served using the physical-layer resources (MC-RWA), the logical-layer resources (logical routing/provisioning), or a combination of both (hybrid routing/provisioning).

A. Routing/Provisioning Problem

1) Physical Routing/Provisioning: This paper assumes a static multicast routing and dynamic wavelength-assignment algorithm that attempts to solve the routing problem and the wavelength-assignment problem independently by dividing the problem into two subproblems: the multicast-routing problem and the wavelength-assignment problem. This is the MC-RWA problem that was described earlier in Section II, where the cost of a physical link (i, j) is defined as follows

$$P_{ij} = \begin{cases} 1, & \text{link } (i,j) \text{ exists} \\ \infty, & \text{link } (i,j) \text{ does not exist} \end{cases} . \tag{1}$$

As explained previously, the SPT multicast-routing algorithm and the first-fit wavelength-assignment scheme are used to solve the MC-RWA problem.

2) Logical Routing/Provisioning: The logical routing/provisioning of a unicast/multicast request is achieved by considering the established lightpaths and light trees as directional logical links that comprise the HGLT. The HGLT construction is performed whenever a call is attempted to be served logically, therefore, the HGLT is dynamically changing every time a lightpath/light tree is set up or torn down in response to the dynamic unicast/multicast traffic demands. Typically, a multicast request may either be routed over a single light tree (i.e., single hyperarc or single hop) or it may span multiple light trees (i.e., multihyperarcs, or multihops).

In the single-hop case, a connection is allowed to traverse a single lightpath/light tree [the same wavelength is used throughout the entire route from source to destination(s)], which means that only end-to-end traffic grooming (multiplexing) is allowed. In the multihop case, a connection is allowed to traverse multiple lightpaths/light trees, i.e., a connection can be dropped/terminated at an intermediate node and multiplexed with other low-speed unicast/multicast connections on different lightpaths/light trees (wavelengths) before it reaches its destination(s). In the case of a single-hop route, only calls with the same source–destination(s) can be multiplexed onto one lightpath/light tree. On the other hand, in the case of a multihop route, calls with different source–destination(s) can be multiplexed into the same lightpath/light tree.

The cost function used for routing calls over the HGLT (logical routing) between a node i and its destination group d_i is defined as the following

$$L_{i,d_i} = \begin{cases} 1, & R_{\mathrm{BW}} >= \mathrm{RQ}_{\mathrm{BW}} \\ \infty, & R_{\mathrm{BW}} < \mathrm{RQ}_{\mathrm{BW}} \text{ or } R_{\mathrm{BW}} = C \end{cases} \tag{2}$$

where R_{BW} is the residual bandwidth of the light tree connecting the source i and its destination group d_i, $\mathrm{RQ}_{\mathrm{BW}}$ is the required bandwidth of a multicast connection request, and C is the channel capacity.

A unicast/multicast connection request can be logically provisioned using either a single hop (existing lightpath/light tree) or a multihop (more than one existing lightpath/light tree). Note that when the residual bandwidth of a light tree reaches the channel capacity, the light tree will be torn down, the corresponding hyperarc will be deleted from the HGLT, and the wavelengths assigned along the links of the corresponding light tree will be released. Also, pruning is performed every time the mechanism is invoked. That is, logical hyperarcs that do not have enough bandwidth to accommodate the call are deleted from the topology. In other words, the routing algorithm used for the logical provisioning approach is the minimum-cost routing algorithm (i.e., Dijkstra) over HGLT, where the cost of a hyperarc is given as defined in (2) (i.e., logical shortest paths).

a) Single-hop approach: In this approach, the algorithm checks for an existing light tree on the hypergraph with available bandwidth to support the new request. If it finds such a light tree, grooming will be performed at the logical layer and the new multicast session will be served using a single-hop logical hyperarc. Note that, for this approach, the new multicast connection must have the same source and multicast destinations as those associated with the logical hyperarc. An illustrative example is shown in Fig. 2. When Call#5 arrives, the HGLT is checked for available single-hop (direct) connection from A to $\{C, D\}$ (with enough residual capacity on the light tree). A direct connection is found (Call #1) with enough residual bandwidth, and Call#5 is then groomed with Call#1 on the solid wavelength.

b) Multihop approach: A multicast session can be provisioned on the logical topology by routing data on more than one light tree (only two hops are allowed here). In this approach, the algorithm searches for an existing light tree whose destinations are the same as those of the new multicast session; we call this light tree "to-destinations light tree" (TDLT). Once a TDLT is found, the algorithm searches for a single-hop lightpath whose source is the same as the source of the new request and its destination is the source of the TDLT; we call this lightpath "from-source lightpath" (FSLP). If the algorithm succeeds, the new multicast connection request is served on the combination of FSLP and TDLT. If it does not succeed, (i.e., no lightpath was found) the call is blocked. An illustrative example is shown in Fig. 2. When Call#6 arrives, the single-hop scheme fails to find a single-hop (direct) connection to serve the multicast call. Then, the multihop scheme is invoked, and a TDLT is searched to the multicast destinations $\{C, D\}$ with enough bandwidth. Such a light tree (hyperarc) is found (from A to $\{C, D\}$ on solid wavelength), and then FSLP is searched from the source of Call#6 (i.e., H) to the source of the TDLT (i.e., A). A single lightpath is again found (on solid wavelength) and Call#6 is then groomed on the multihop route H to $\{A\}$ (on solid wavelength) and A to $\{C, D\}$ (on solid wavelength).

3) Hybrid Routing/Provisioning: A unicast/multicast request can be routed/provisioned over a combination of existing lightpaths/light trees (logical routing) and a newly created lightpath/light tree (physical routing, e.g., RWA/MC-RWA). In this approach, the optical layer must keep updated databases about the connectivity of both the logical and physical topologies as well as the resource utilization across both layers [28], [29]. In this paper, the idea of a hybrid provisioning approach is defined to find a combined route of an existing logical segment (light tree) and an unprovisioned physical segment (lightpath) that needs to be set up.

To illustrate the concept of hybrid provisioning, assume the case depicted in Fig. 2. Let Call#1 to Call#6 arrive at the network and be served as shown in the figure. Suppose now that Call#7 arrives at the network requesting service from A to $\{F, G\}$. Since the link between E and $\{G\}$ is completely utilized at the wavelength level (both wavelengths are used), and no logical connectivity exists from A to $\{F, G\}$, traditionally, this call would have been blocked. Under the hybrid approach, however, node A knows that there exists a TDLT to $\{F, G\}$ that has originated at E. Therefore, A requests from the physical topology an FSLP to $\{E\}$, and if enough resources are found and a successful connection is established, Call#7 is served by utilizing the new established FSLP from A to $\{E\}$ and the already established TDLT from E to $\{F, G\}$.

B. Multicast Traffic Grooming Strategy and Heuristics

The critical remaining question that needs to be addressed in this section is how to combine the above developed multiple routing/provisioning approaches to serve unicast/multicast traffic demands as efficiently as possible using the overall global network resources; in other words, how these global network resources should be allocated to a given request if multiple routing/provisioning schemes are available. Should a new request be accommodated using one or more existing lightpaths/light trees first? Is it more appropriate to set up a new lightpath/light tree first? Or should it be accommodated using a combination of both in a sequential order similar to conventional sequential routing/provisioning approaches used for unicast traffic grooming [7], [11], [30]? It is important to emphasize that existing sequential unicast grooming strategies and algorithms along with associated simulation results [7], [11], [30] cannot be taken as guidelines when considering multicast traffic. Thus, conventional sequential routing/provisioning approaches used for unicast traffic grooming should be revisited (after being modified to tailor provisioning of multicast traffic) and thoroughly examined, to test its effectiveness for grooming multicast traffic as well.

To answer these questions, we first develop several different sequential multicast grooming heuristics by means of interchanging the search space between the physical and logical layers. Second, guided by the simulation results of the proposed sequential multicast grooming heuristics, we augment the sequential grooming heuristics (the one that gives best performance results) by a hybrid approach to implement a grooming scheme that is biased toward serving multicast traffic demands in comparison with all other sequential grooming approaches.

C. Sequential Multicast Traffic Grooming Heuristics

1) Logical-First Sequential Routing: With this algorithm, the network first tries to accommodate the call on the logical topology making use of the already-existing connections. Depending on the grooming approach to be used (i.e., single hop or multihop), if no available single/multihop route were found on the logical topology, the source node sends a request to the ingress optical cross connect (OXC) to set up a light tree on the physical topology to all the multicast destinations. Then, the ingress OXC invokes its MC-RWA algorithm to set up a new light tree to the destinations. If the MC-RWA is unsuccessful, the call is blocked. Depending on the number of logical hops, this algorithm can further be classified into the following two schemes.

1) Logical-first sequential routing with single-hop grooming (LFSEQSH): With this scheme, the network first tries to service the call using a single-hop logical hyperarc only (i.e., single light tree). If the search is successful, the call is serviced. Otherwise (i.e., if the search fails), the ingress OXC invokes its MC-RWA algorithm to set up a light tree on the physical topology to all the multicast destinations. If the MC-RWA is unsuccessful, the call is blocked.
2) Logical-first sequential routing with multihop grooming (LFSEQMH): This algorithm is similar to the "LFSEQSH;" however, the search on the logical topology is allowed to include two logical hops instead of a single hop. Therefore, multihop grooming is performed in a sequential way by attempting the logical layer first. Note that throughout the remaining part of this paper, "multihop" means only two hops. The algorithm is summarized below in Table I.

2) Physical-First Sequential Routing: With this approach, the network attempts to accommodate a call on the physical layer first. If the new light tree is established successfully, a new logical/virtual hyperarc (light tree) is created in the logical layer. If the physical routing fails, then, routing on logical topology is attempted. Similar to logical-first approach,

TABLE I
LFSEQMH ALGORITHM

Step 1: For each multicast connection request, check whether a logical single light-tree with available bandwidth and with the same source and destinations as those of the multicast connection request exists. If yes, go to **Step 5**, else go to **continue**.
Step 2: Search logical topology for an existing single light-tree with available bandwidth whose destinations are the same destinations as the multicast connection request (i.e., **TDLT**). If yes, go to **Step 4**, else go to **continue**.
Step 3: Invoke physical provisioning (i.e., **MC-RWA**), if a multicast tree was found and a wavelength was available, go to **Step 5**, else **stop** and block the connection request.
Step 4: Identify the source of **TDLT** and search the logical topology for a lightpath (i.e., **FSLP**) between the original source of the request and the source of TDLT. If yes, go to **continue,** else go back to **Step 3**
Step 5: Serve the multicast request by reserving the bandwidth along the corresponding lightpath and/or light-tree and update the logical topology; the connection request is satisfied.

depending on the number of logical hops, this algorithm can also be classified into the following two schemes.

1) Physical-first sequential routing with single-hop grooming (PFSEQSH): This scheme is similar to the LFSEQSH; however, the search is now attempted on the physical layer first. Therefore, single-hop grooming is performed in a sequential way by attempting the physical layer first.
2) Physical-first sequential routing with multihop grooming (PFSEQMH): This scheme is similar to the LFSEQMH; however, the search is now attempted on the physical layer first. Therefore, multihop grooming is performed in a sequential way by attempting the physical layer first.

3) Combined Sequential and Hybrid Routing: In this scheme, we combine the hybrid provisioning approach with sequential provisioning approaches to achieve a grooming scheme that is biased toward serving multicast traffic demands in comparison with all other sequential grooming approaches. Specifically, we combine the hybrid approach with that of the LFSEQMH sequential approach and denote the resultant grooming scheme as logical-first hybrid routing (LFHYB). The reason for choosing the LFSEQMH algorithm is that, as will be shown below, this scheme gives the best performance results from among all the sequential schemes described above.

The steps for implementing this scheme are exactly the same as LFSEQMH except that upon the failure of **Step 4** in Table I [no existing lightpath (FSLP) that directly connects the source of the requested multicast session and the source of a TDLT is found on the logical topology], rather than going back to **Step 3**, the algorithm tries to set up a new lightpath (by invoking MC-RWA on the physical topology) between the source of the requested multicast session and the source of a TDLT. If the MC-RWA succeeds, the combination of the newly provisioned lightpath and the existent TDLT will be used to serve the new multicast request. If it fails, then the algorithm goes back to **Step 3** in Table I. Note that more than one TDLT can be found; however, we only consider the "first available." The algorithm is summarized below in Table II.

V. PERFORMANCE EVALUATION

The performances of the proposed grooming algorithms are evaluated through the simulation of several network topologies that demonstrated similar conclusions. We present results here for the 14-nodes 21-link National Science Foundation (NSF) network shown in Fig. 3. Unless otherwise specified, the following assumptions and parameters are used throughout the simulation.

1) Each node in the network is an MC-OGSW with full splitting and full grooming capability, but with no wavelength-conversion capability.
2) Each adjacent node pair is interconnected by one (bidirectional) fiber and each fiber carries 64 wavelengths.
3) Multicast traffic is assumed to constitute half of the total traffic and the other half constitutes unicast traffic.
4) Signal power loss due to light splitting is neglected because optical amplifiers are used.
5) A dynamic traffic model is used. Call requests arrive at each node according to a Poisson process and an arrival session is equally likely to be destined to any node(s) in the network. For both unicast and multicast sessions, the low-speed connection requests are 25% of the wavelength capacity.

TABLE II
LFHYB ALGORITHM

Step 1: For each multicast connection request, check whether a logical single light-tree with available bandwidth and with the same source and destinations as those of the multicast connection request exists. If yes, go to **Step 6**, else go to **continue**.
Step 2: Search logical topology for an existing single light-tree with available bandwidth whose destinations are the same destinations as the multicast connection request (i.e., **TDLT**). If yes, go to **Step 4**, else go to **continue**.
Step 3: Invoke physical provisioning (**MC-RWA**), if a multicast tree was found and a wavelength was available, go to **Step 6**, else **stop** and block the connection request.
Step 4: Identify the source of **TDLT** and search the logical topology for an existing lightpath (i.e., **FSLP**) between the original source of the request and the source of **TDLT**. If yes, go to **Step 6,** else **continue.**
Step 5: Invoke physical provisioning (**MC-RWA**), between the original source of the request and the source of **TDLT**, if succeeded**, continue;** else go back to **Step 3.**
Step 6: Serve the multicast request by reserving the bandwidth along the related lightpath and/or light-tree and update the logical topology, thus the connection request is satisfied.

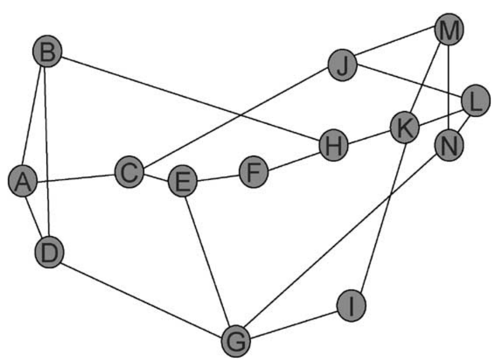

Fig. 3. NSFNET.

6) For each multicast session, the destination group size G_p is equal to 30%, where G_p is the maximum percentage of the network nodes that could be destinations.

The simulation results of the proposed grooming algorithms will be compared with those of two baseline conventional multicast schemes. In the first one, a multicast connection request is served as multiple independent UCs. In the second scheme, each multicast request is served on a separate light tree (MC-RWA) without taking into account the grooming capability of the optical node.

Fig. 4 shows the BP of the proposed sequential grooming schemes. As it can be seen from the figure, the performance of the "LFSEQMH" scheme outperforms all other schemes. Note also that the performance of both logical-first schemes (LFSEQMH and LFSEQSH) outperform those of physical-first schemes (PFSEQMH and PFSEQSH). This is in sharp contrast to the sequential unicast grooming algorithms where it has been shown that the performance of the physical-first scheme is the one that outperforms all other sequential unicast schemes [7], [11], [30].

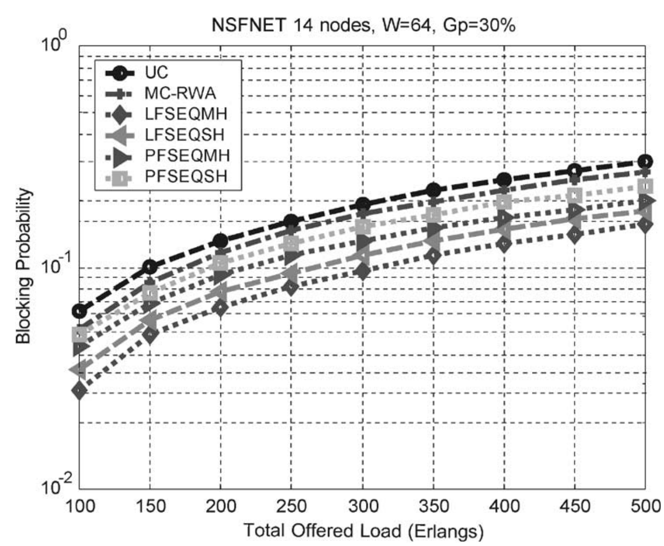

Fig. 4. BP of proposed sequential approaches and conventional approaches (i.e., UC and MC-RWA).

We define the grooming gain as the percentage in network performance gain (in terms of BP) achieved when a grooming scheme is used versus a nongrooming scheme. For instance, if "LFSEQMH" is used as a grooming algorithm, and

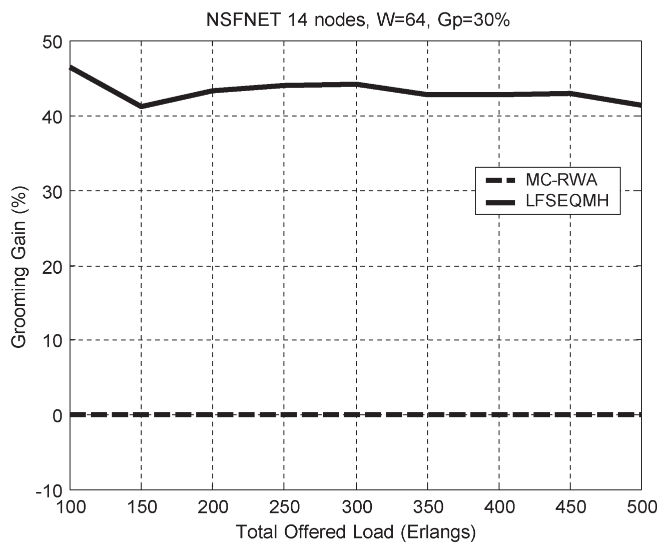

Fig. 5. Grooming gain versus network load.

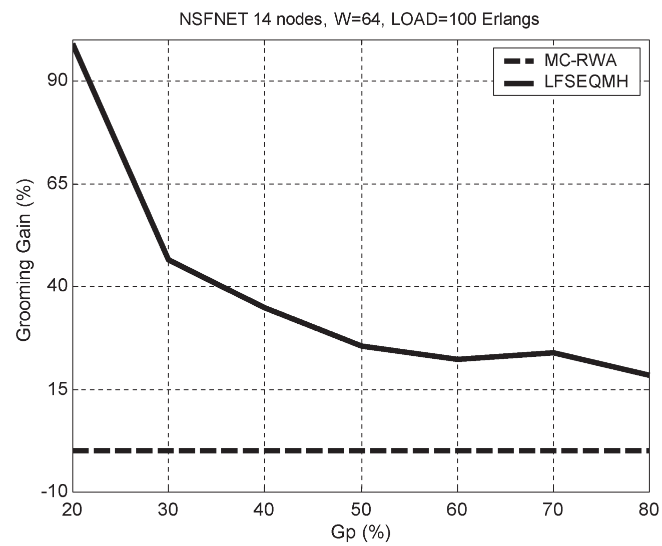

Fig. 6. Grooming gain versus multicast group size.

"MC-RWA" is used as a conventional nongrooming algorithm, then, the grooming gain of "LFSEQMH" is defined as follows

$$G_{\text{LFSEQMH}} = \frac{\text{BP}_{\text{MC-RWA}} - \text{BP}_{\text{LFSEQMH}}}{\text{BP}_{\text{MC-RWA}}} \times 100. \quad (3)$$

Using a grooming scheme such as "LFSEQMH" results in at least a 40% gain in network performance (grooming gain) when compared to a nongrooming scheme such as "MC-RWA" (see Fig. 5). Note that the grooming gain is almost independent of the total network offered load since the gain is relative to "MC-RWA," and as the load increases, both schemes will have a higher BP.

Fig. 6 shows the grooming gain versus the multicast group size (G_{p}). As can be seen from Fig. 6, the grooming gain decreases with the increase of multicast group size. This is because, as the number of multicast destinations increases, the probability of finding a single-hop light tree or a combination of single-hop lightpath and a single-hop light tree with the same exact multicast destinations decreases. Hence, as expected, the grooming gain will decrease.

Fig. 7 shows the percentage of multicast traffic served at the logical layer (both single hop and multihop) as well as at the physical layer when "LFSEQMH" is used. As can be seen from the figure, the total number of groomed multicast sessions that are served on the HGLT (logical layer) is between 30% and 50% of the total multicast connection requests. Note that the number of single-hop multicast sessions being served on the HGLT is higher than those of the multihop. Note also that the percentage of single-hop multicast sessions increase with offered load while the percentage of the multihop multicast sessions do not change much. This is due to the fact that, as the offered load increases, more multicast requests will arrive, which leads to highly connected HGLT. Hence, it is more probable to find existing light trees for grooming single-hop sessions, leading to the observed increase of multicast sessions that are groomed with single-hop routes.

To assess the performance of a multihop grooming scheme versus that of a single-hop one, we define the multihop gain (MHG) to be the percentage in performance gain (in terms of BP) achieved when using a multihop routing approach over a single-hop approach. Thus, MHG is defined as follows

$$\text{MHG} = \frac{\text{BP}_{\text{SH}} - \text{BP}_{\text{MH}}}{\text{BP}_{\text{SH}}} \times 100. \quad (4)$$

Using (4), we calculate the MHG of LFSEQMH with respect to LFSEQSH. Fig. 8 shows the MHG versus percentage of multicast traffic demands. As can be seen from Fig. 8, with 20% multicast traffic, there is no significant advantage as most of the traffic is unicast. As the percentage of multicast traffic increases, the probability of finding a combination of existing lightpath/light tree for multihop routing increases, resulting in a corresponding increase of the MHG, but up to a certain limit (percentage of multicast sessions are 50% or less). As the percentage of multicast sessions exceeds 50%, the gain decreases again. This is because with more multicast traffic, the probability of finding a single connection (single-hop lightpath) between the source of the multicast session and an intermediate source (the source of a TDLT) will decrease (less unicast traffic).

We now examine the impact of combining both the hybrid and sequential approaches (LFHYB) on the grooming capability of multicast sessions and compare its performance with that of the sequential approach (LFSEQMH). Fig. 9 shows the percentage of additional multicast sessions that are groomed on the logical topology when the hybrid approach (LFHYB) is used with respect to those that would have been groomed if the sequential logical-first approach (LFSEQMH) is used. As expected, the percentage of multicast sessions that are served using the hybrid scheme (LFHYB) is always greater than that of the sequential scheme (LFSEQMH). For instance, at low loads (100 erlangs), the figure shows that the hybrid approach is able to groom almost 85% more multicast sessions than the logical

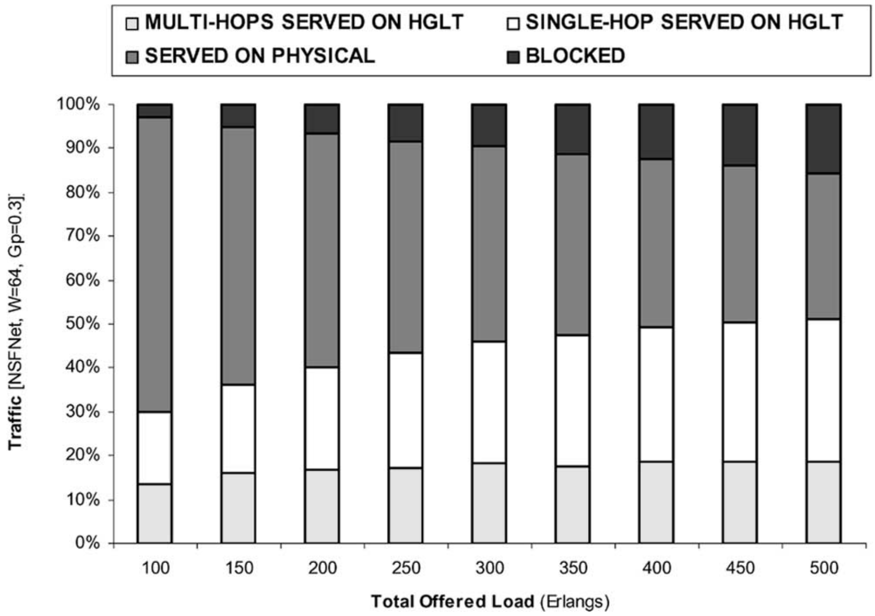

Fig. 7. Percentage of single-hop and multihop groomed sessions for "LFSEQMH."

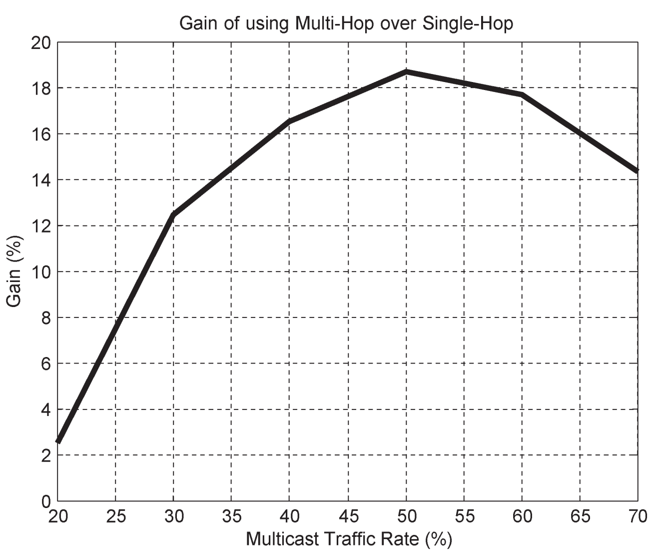

Fig. 8. Multihop gain (Load $=$ 100 erlangs; $G_{\mathrm{p}} = 30\%$).

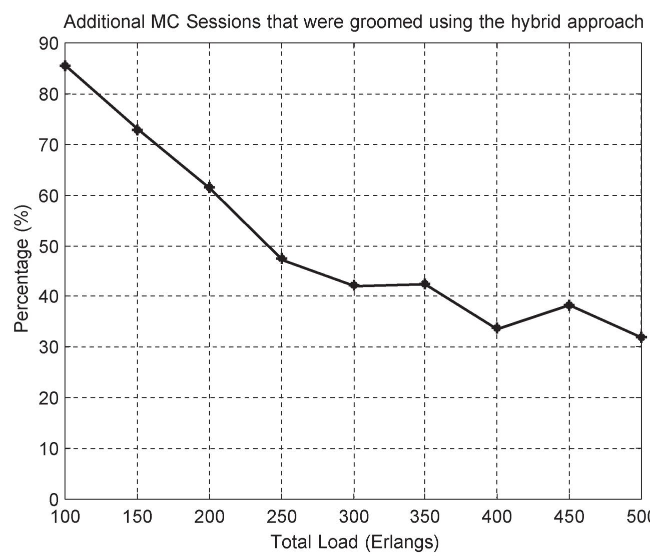

Fig. 9. Additional groomed multicast sessions when "LFHYB" is used instead of "LFSEQMH."

sequential approach. This is very critical because blocking a unicast session is usually less expensive than blocking a multicast session. Hence, by allowing more multicast sessions to be served, we are improving the network performance. Note, however, as the load increases, the additional percentage of groomed multicast sessions decreases (down to almost 30%). This is because the hybrid approach uses both logical and physical resources to groom these additional multicast sessions. At high loads, physical resources are almost depleted and, therefore, the hybrid approach performance approaches that of the logical sequential approach (high loads means high logical connectivity and less available physical resources, which are required for hybrid provisioning).

Fig. 10 compares the percentage of the total traffic (both unicast and multicast) that is groomed on the logical topology only, for both LFHYB and LFSEQMH schemes. The results shown in Fig. 10 further confirm the results of Fig. 9. With more calls served purely logically, the hybrid approach minimizes the total cost of the network by maximizing the utilization of the existent logical resources.

Finally, we assess the performance of the above multihop grooming approaches when the second constraint imposed on our grooming strategy (unicast traffic may not be groomed onto multicast traffic) is relaxed. When the second constraint is relaxed, we redefine the original restricted "LFSEQMH" scheme as a nonrestricted logical-first sequential routing (NRLFSEQMH) scheme. Fig. 11 compares the performance of the restricted "LFSEQMH" scheme with that of the "NRLFSEQMH" scheme. As can be seen from the figure, at

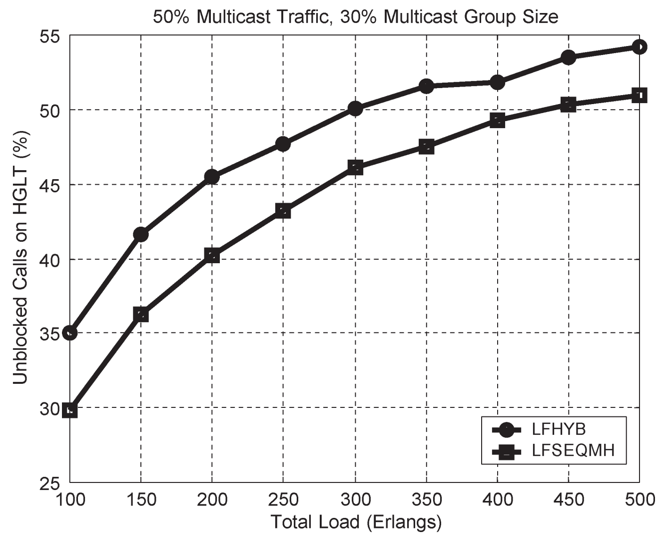

Fig. 10. Percentage of calls that are purely served on HGLT.

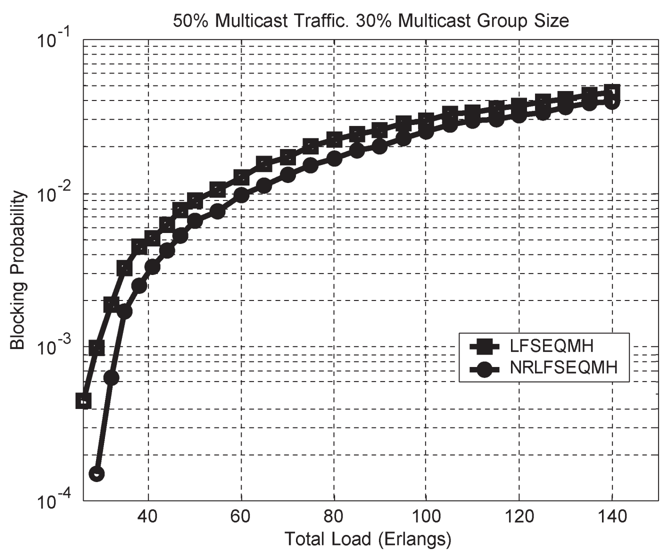

Fig. 11. Performance of nonrestricted approach at low loads.

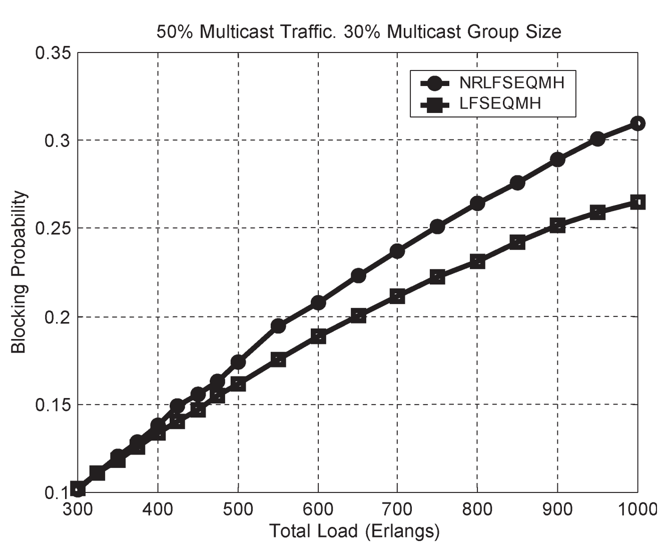

Fig. 12. Performance of nonrestricted approach at high loads.

low loads, the performance of the "NRLFSEQMH" scheme outperforms that of the original "LFSEQMH" scheme. This is because, at low loads, more resources are still available and the impact of wasting bandwidth (some destinations on the multicast tree end up receiving unintended unicast data) is negligible. However, at higher loads, as shown in Fig. 12, the performance of the original restricted grooming algorithm "LFSEQMH" outperforms that of the "NRLFSEQMH" scheme. This is because, at higher loads, available resources are now scarce and the impact of wasting bandwidth has a deleterious effect (increases BP) on the overall network performance.

VI. Conclusion and Future Work

This paper has addressed the problem of dynamically provisioning both low-speed unicast and multicast connection requests in mesh-based WDM optical networks. Several routing/provisioning schemes to dynamically provision multicast connection requests have been presented. We have devised a constraint-based grooming strategy to utilize the overall network resources as efficiently as possible. Based on this strategy, several different sequential multicast grooming heuristics have been developed. Guided by the simulation results of the sequential multicast grooming heuristics, we have augmented the sequential grooming heuristics by a hybrid approach to implement a grooming scheme that is biased toward serving multicast traffic demands in comparison with all other sequential grooming approaches. Note that our proposed algorithms can be modified to allow a user to be added or removed from a multicast session; however, this falls outside the scope of the paper, where the focus is more on the proposed grooming strategies and the hybrid approach.

Numerical results have indicated that the proposed grooming approaches use the network resources more efficiently compared to the nongrooming and unicast approaches. The results have also shown that the percentage of multicast sessions that are served on the logical topology using the hybrid scheme is always greater than that of the sequential approaches. This is very important because blocking a unicast session is usually less expensive than blocking a multicast session. Finally, the hybrid approach increased the percentage of both unicast and multicast calls that were served purely on the logical topology; hence, it minimized the total cost of the network by maximizing the utilization of the existent logical resources. We are currently investigating the effect of wavelength converters in the network, as well as the effect of sparse wavelength splitters.

References

[1] R. Ramaswami and K. N. Sivarajan, *Optical Networks: A Practical Perspective*, 2nd ed. San Mateo, CA: Morgan Kaufman, Nov. 2001.
[2] K. Zhu and B. Mukherjee, "A review of traffic grooming in WDM optical networks: Architectures and challenges," *SPIE Opt. Netw. Mag.*, vol. 4, no. 2, pp. 55–64, Mar./Apr. 2003.
[3] E. Modiano and P. J. Lin, "Traffic grooming in WDM networks," *IEEE Commun. Mag.*, vol. 39, no. 7, pp. 124–129, Jul. 2001.
[4] R. Dutta and G. N. Rouskas, "On optimal traffic grooming in WDM rings," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 1, pp. 110–121, Jan. 2002.

[5] O. Grestel, R. Ramaswami, and G. Sasaki, "Cost effective traffic grooming in WDM rings," *IEEE/ACM Trans. Netw.*, vol. 8, no. 5, pp. 618–630, Oct. 2000.
[6] X. Zhang and C. Qiao, "An effective and comprehensive approach for traffic grooming and wavelength assignment in SONET/WDM rings," *IEEE/ACM Trans. Netw.*, vol. 8, no. 5, pp. 608–617, Oct. 2000.
[7] S. Thiagarajan and A. Somani, "Capacity fairness of WDM networks with grooming capabilities," *SPIE Opt. Netw. Mag.*, vol. 2, no. 3, pp. 24–31, May/Jun. 2001.
[8] H. Zhu, H. Zang, K. Zhu, and B. Mukherjee, "A novel, generic graph model for traffic grooming in heterogeneous WDM mesh networks," *IEEE/ACM Trans. Netw.*, vol. 11, no. 2, pp. 285–299, Apr. 2003.
[9] K. Zhu and B. Mukherjee, "Traffic grooming in an optical WDM mesh network," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 1, pp. 122–133, Jan. 2002.
[10] B. Doshi *et al.*, "Optical network design and restoration," *Bell Labs Tech. J.*, vol. 4, no. 1, pp. 58–84, Jan.–Mar. 1999.
[11] C. Assi, Y. Ye, A. Shami, S. Dixit, and M. Ali, "Integrated routing algorithms for provisioning "Sub-Wavelength" connections in IP-over-WDM networks," *Photonic Netw. Commun.*, vol. 4, no. 3/4, pp. 377–390, Jul. 2002.
[12] G. Sahin and M. Azizoglu, "Multicast routing and wavelength assignment in wide-area networks," in *Proc. SPIE—All Optical Networking*, Boston, MA, Nov. 1998, pp. 197–208.
[13] R. Malli, X. Zhang, and C. Qiao, "Benefits of multicasting in all-optical networks," in *Proc. SPIE—All Optical Networking*, Boston, MA, Nov. 1998, pp. 209–220.
[14] X. Zhang, J. Wei, and C. Qiao, "On fundamental issues in IP over WDM multicast," in *Proc. IEEE Int. Conf. Computer Communications and Networks (ICCCN)*, Boston, MA, Oct. 1999, pp. 84–90.
[15] ——, "Constrained multicast routing in WDM networks with sparse light splitting," in *Proc. Int. Conf. Computer Communications (INFOCOM)*, Tel-Aviv, Israel, Mar. 2000, pp. 100–110.
[16] Y. Sun, J. Gu, and D. H. K. Tsang, "Multicast routing in all-optical wavelength-routed networks," *SPIE Opt. Netw. Mag.*, vol. 2, no. 4, pp. 101–109, Jul./Aug. 2001.
[17] L. H. Sahasrabuddhe and B. Mukherjee, "Light-trees: Optical multicasting for improved performance in wavelength-routed networks," *IEEE Commun. Mag.*, vol. 37, no. 2, pp. 67–73, Feb. 1999.
[18] S. Jiang, "Multicast multihop lightwave network: Design and implementation," Center Telecommun. Res., Columbia Univ., New York, CTR Tech. Rep. 418 95 24, 1995.
[19] K. Bala, "Routing in linear networks," Ph.D. dissertation, Dept. Elect. Eng., Columbia Univ., New York, 1992.
[20] K. Bala, K. Petropoulos, and T. E. Stern, "Multicasting in linear lightwave networks," in *Proc. IEEE Int. Conf. Computer Communications (INFOCOM)*, San Francisco, CA, Mar. 1993, pp. 1350–1358.
[21] A. Kamal and R. Ul-Mustafa, "Multicast traffic grooming in WDM networks," in *Proc. SPIE/IEEE Optical Networking and Communications (OptiComm)*, Dallas, TX, Oct. 2003, pp. 25–36.
[22] D.-N. Yang and W. Liao, "Design of light-tree based logical topologies for multicast streams in wavelength routed optical networks," in *Proc. IEEE Int. Conf. Computer Communications (INFOCOM)*, San Francisco, CA, Mar. 2003, vol. 22, no. 1, pp. 32–41.
[23] S. Deering, "Multicast routing in Internetworks and extended LANs," Ph.D. dissertation, Dept. Comput. Sci., Stanford Univ., Stanford, CA, Aug. 1988.
[24] M. R. Garey, R. L. Graham, and D. S. Johnson, "The complexity of computing Steiner minimal trees," *SIAM J. Appl. Math.*, vol. 32, no. 4, pp. 835–859, 1977.
[25] F. K. Hwang and D. Richards, "The Steiner tree problem," *Networks*, vol. 22, no. 1, pp. 55–89, Jan. 1992.
[26] H. Takahashi and A. Matsuyama, "An approximate solution for the Steiner problem in graphs," *Math. Jpn.*, vol. 24, no. 6, pp. 573–577, 1980.
[27] A. Mokhtar and M. Azizoglu, "Adaptive wavelength routing in all-optical networks," *IEEE/ACM Trans. Netw.*, vol. 6, no. 2, pp. 197–206, Apr. 1998.
[28] A. Khalil, A. Hadjiantonis, G. Ellinas, and M. A. Ali, "Optical layer-based unified control plane for emerging IP/MPLS over WDM networking architecture," in *Proc. 29th Eur. Conf. Optical Communication*, Rimini, Italy, Sep. 2003, pp. 836–837.
[29] ——, "A novel IP-over-optical network interconnection model for the next-generation optical internet," in *Proc. IEEE Global Telecommunications (GLOBECOM)*, San Francisco, CA, Dec. 2003, pp. 3984–3989.
[30] Y. Ye, C. Assi, S. Dixit, and M. A. Ali, "A simple dynamic integrated provisioning/protection scheme in IP over WDM networks," *IEEE Commun. Mag.*, vol. 39, no. 11, pp. 174–182, Nov. 2001.

Ahmad Khalil received the B.E. and M.S degrees in mechanical engineering from the Lebanese University, Beirut, Lebanon, and Ohio University, Athens, respectively. In May 2005, he received the Ph.D. degree in electrical engineering from the Graduate School and University Center, City University of New York (CUNY).

He joined the mathematics and engineering science department at LaGuardia College, CUNY, in September 2004. He has been an Assistant Professor since September 2005. He is also a research scientist at the Next-Generation Networking Group, City College, CUNY. He received the Robert E. Gilleece Fellowship, from 2000 to 2004, for graduate studies at The Graduate Center/CUNY, and has authored or co-authored more than 15 journal and conference papers. He has conducted research in the areas of routing and wavelength assignment algorithms, provisioning and protection, traffic grooming, and multicasting in optical networks for over than four years. His current research interests are in the areas of optical as well as wireless *ad hoc* networks.

Antonis Hadjiantonis was born in Morphou, Cyprus, in 1973. He received the B.E.E.E. and M.E.E.E. degrees from the City College, City University of New York (CUNY), in 1998 and 2000, respectively. In 2005, received the Ph.D. degree from the Graduate School and University Center, CUNY.

He is currently involved in teaching and research projects in Cyprus. His research interests include next generation all-optical networks, integration of optical and electronic layers, and signaling and routing algorithms.

Chadi M. Assi received the B.S. degree in engineering from the Lebanese University, Beirut, Lebanon, in 1997 and the Ph.D. degree from the Graduate Center, City University of New York (CUNY), in April 2003.

He was a Visiting Researcher at Nokia Research Center, Boston, MA, from September 2002 to August 2003, working on quality-of-service in optical access networks. He joined the Concordia Institute for Information Systems Engineering (CIISE), Concordia University, Montreal, QC, Canada, in August 2003 as an Assistant Professor.

Dr. Assi received the Mina Rees Dissertation Award from CUNY in August 2002 for his research on wavelength-division-multiplexing optical networks. His current research interests are in the areas of optical networks, provisioning and restoration, wireless and *ad hoc* networks, and quality of service.

Abdallah Shami (S'01–M'03) received the B.E. degree in electrical and computer engineering from the Lebanese University, Beirut, Lebanon, in 1997 and the Ph.D. degree in electrical engineering from the Graduate School and University Center, City University of New York (CUNY), in 2003.

In September 2002, he joined the Department of Electrical Engineering, Lakehead University, Thunder Bay, ON, Canada, as an Assistant Professor. Since July 2004, he has been with the University of Western Ontario, London, ON, where he is currently an Assistant Professor with the Department of Electrical and Computer Engineering. His current research interests are in the area of data/optical networking, EPONs, QoS in wireless local area networks, and software tools.

Dr. Shami received the Irving Hochberg Dissertation Fellowship Award, CUNY, for his research in optical networking.

George Ellinas received the B.S., M.S., M.Phil, and Ph.D. degrees in electrical engineering from Columbia University, New York, NY.

He is currently an Assistant Professor with the Department of Electrical and Computer Engineering, University of Cyprus, Cyprus, Greece. Prior to joining the University of Cyprus, he was an Associate Professor of electrical engineering at City College, City University of New York. Before joining the academia, he was a senior network architect at Tellium Inc. In this role, he worked on lightpath provisioning and fault restoration algorithms in optical mesh networks and the architecture design of the MEMS-based all-optical switch. He also served as a senior research scientist with Telcordia Technologies' (formerly Bellcore) Optical Networking Research Group. He performed research for the Optical Networks Technology Consortium (ONTC), Multiwavelength Optical Networking (MONET), and Next-Generation Internet (NGI) projects from 1993 to 2000. He also served as an Adjunct Assistant Professor at Columbia University and the University of Maryland, College Park, teaching courses on multiwavelength optical networking in 1999 and 2000, respectively.

Dr. Ellinas received a Fulbright fellowship, from 1987 to 1991, for undergraduate studies at Columbia University and has authored or co-authored 75 journal and conference papers. He is also the holder of 24 U.S. and international patents on optical networking and has four U.S. patent applications currently pending.

Mohamed A. Ali received the M.S. and Ph.D. degree in electrical engineering from the City College, City University of New York (CUNY), in 1985 and 1989, respectively.

He joined the faculty of electrical engineering at the City College, CUNY, in 1989, where he is currently a Professor. His research interest is in the general area of broadband information, networking architecture, and technology. His current research activities involve next-generation data-centric networking architecture, traffic engineering, GigE networking technology and architecture, IP/MPLS-based VPNs, DWDM and TDM multiple-access broadband networks; high-performance IP/MPLS routers, CATV distribution over fiber-based local access networks, and ATM/Ethernet-based passive optical networks. He has consulted for several major carriers. He has published over 100 papers in professional journals and international conferences.

Dr. Ali received the NSF Faculty Career Development Award.