Multicast Overlay for High-Bandwidth Applications Over Optical WDM Networks

Arush Gadkar, Jeremy Plante, and Vinod M. Vokkarane

***Abstract*—Multicast communication in wavelength division multiplexed (WDM) networks is traditionally supported by the assumption that the optical crossconnects are multicast capable, i.e., they are capable of switching an incoming signal to more than one output interface. A naïve method of supporting this functionality in a multicast-incapable (MI) environment is by creating a virtual topology consisting of lightpaths from the multicast source to each destination of the multicast session. For large sets of multicast requests, however, the network bandwidth consumed by such a scheme may become unacceptable due to the unicasting nature of the lightpaths. We refer to this method as achieving multicast via WDM unicast (MVWU). To support users' multicast requests (from higher electronic layers) in MI networks, we propose two overlay solutions: drop at member node (DMN) and drop at any node (DAN). In these solutions, we achieve multicasting by creating a set of lightpath routes (possibly multiple hops) in the overlay layer from the source node of a request to each destination member. In the DMN case, we allow a lightpath route to originate/terminate only at source and destination members of a request, whereas in the DAN model we impose no such restrictions. We first consider a static traffic model, wherein the set of multicast requests is known ahead of time, and present integer linear programs (ILPs) to solve these problems (MVWU, DMN, and DAN) with the goal of minimizing the total number of wavelengths required to service the set. We also present an efficient heuristic and compare its performance to the ILP for a small network, and run simulations over real-world, large-scale networks. Moreover, we present lower bounds to calculate the minimum number of wavelengths required by the DMN and DAN models. Finally, we evaluate the performance of the heuristic (minimization of the number of wavelengths) under a dynamic traffic scenario and also evaluate the blocking performance for a fixed number of wavelengths.**

***Index Terms*—IP-over-WDM; Lightpath; Multicasting; Overlay networks; WDM.**

I. Introduction

Recently, a growing number of scientific applications require large amounts of data (usually on the scale of petabytes) generated by experiments to be accessible and analyzed by a large number of geographically dispersed users. To support such applications, not only is large bandwidth required, but the underlying network must be much more intelligent and efficient than the traditional best-effort Internet. Because of the enormous bandwidth offerings available in optical wavelength division multiplexed (WDM) networks, they prove to be a potential candidate to support such applications.

As pointed out in [1], one of the requirements of these bandwidth-hungry applications is to use the network as a data cache, i.e., the ability for a particular site to store the data generated by different experiments at different geographical locations across the network. This calls for the use of a point-to-multipoint communication paradigm, which is becoming increasingly important to service next-generation applications. Multicasting is such a communication paradigm that supports transfer of data between a single source and a set of pre-selected destinations.

There is abundant work in the literature on multicasting in optical networks. The benefits of supporting multicast at the optical layer have been discussed in [2–4]. To efficiently support multicast at the optical layer, the network creates *light-trees* [3]. A *light-tree* is a generalization of a lightpath, which starts at the source node of a multicast request and reaches all the destinations by possibly branching (splitting) the signal at intermediate nodes. The problem of finding an optimal route for the *light-tree* is equivalent to finding a Steiner tree, which is known to be NP-complete [5], although efficient approximation algorithms exist [6]. In order to support light-trees, the nodes in an optical network must be able to split an incoming photonic signal to multiple output ports. This can be accomplished by using splitter-and-delivery-(SaD-) based switches [7,8]. These switches are known as multicast capable-optical crossconnects (MC-OXC). In [9], the authors have investigated the routing and wavelength assignment (RWA) problem in networks employing *tap-and-continue* switches which have limited multicast capabilities. The authors of [2,10] have investigated the RWA problem under various *fan-out* splitting policies, and in [3] the authors formulate the RWA problem of multicast routing in packet-switched networks as a mixed integer linear programing (MILP) problem.

Multicast applications are becoming more popular and will make up an important part of future Internet traffic [11]. Examples of multicast applications include video-conferencing, interactive distance learning, streaming media, distributed data processing, storage area networks, and e-Science applications. Large-scale scientific experiments, such as those conducted at CERN [12], require the reservation of large amounts of bandwidth to transfer data sets to possibly multiple

Manuscript received December 6, 2011; revised March 24, 2012; accepted May 1, 2012; published July 10, 2012 (Doc. ID 159584).

Arush Gadkar and Jeremy Plante are with the Department of Computer and Information Science (CIS), University of Massachusetts, Dartmouth, Massachusetts, USA.

Vinod M. Vokkarane (e-mail: vvokkarane@ieee.org) is with the Department of Computer and Information Science (CIS), University of Massachusetts, Dartmouth, Massachusetts, USA, and is currently a Visiting Scientist at the Claude E. Shannon Communication and Network Group (Research Laboratory of Electronics), Massachusetts Institute of Technology, Cambridge, Massachusetts, USA.

Digital Object Identifier 10.1364/JOCN.4.000571

locations worldwide. A fundamental obstacle in supporting such communications over the DOE's Energy Sciences network (ESnet) [1] is that the underlying optical layer is not multicast capable [13], i.e., the optical crossconnects (OXCs) deployed in the ESnet are not capable of all-optically splitting an incoming signal to multiple output ports. Such networks may be referred to as multicast-incapable (MI) networks. Hence MI networks do not inherently support the multicast communication paradigm. This problem can be overcome by a naïve scheme by providing a simple point-to-point (unicast) communication channel (lightpath) between the source and each multicast destination. More specifically, for a given multicast request we establish individual end-to-end lightpaths from the source node of the request to each and every destination member of the request. As scientific applications evolve and more collaborations occur (between scientists and laboratories), the multicast destination set size increases and the network bandwidth consumption of such an approach will prove to be inefficient. Hence, to service the higher layer (user's) request on MI networks, we propose a efficient overlay approach to achieve multicast communication.

Two traffic models are usually considered for wavelength-routed networks: static and dynamic [14]. A static traffic model provides all the traffic demands between source and destination(s) before provisioning begins. Given a traffic matrix, the goal is typically to find an RWA scheme that can meet all the demands while minimizing overall cost (e.g., using the least number of transmitters/receivers). In the dynamic traffic model, requests arrive one-by-one according to some stochastic process and use network resources only temporarily as they only require access to the network for a finite amount of time. When dynamic traffic is considered, the number of transmitters and receivers is fixed and the goal is to minimize request blocking. A request is said to be blocked if there are not enough resources available to route it. There is a significant amount of work for the multicast problem with these types of traffic demands [9,10,15–19]. The traffic models can be further classified as *immediate reservation* (IR) or *advance reservation* (AR) [20] requests. The data transmission of an IR demand starts immediately upon arrival of the request and the holding time is typically assumed to be infinite for static traffic. AR demands, in contrast, typically specify a data transmission start time that is some time in the future and also specify a finite holding time. AR may also be referred to as scheduled demands [21], especially when considering static traffic. In this work we consider IR traffic only. AR is the focus of our on-going work.

In this paper we consider both static and dynamic IR traffic and propose an overlay approach to service the higher layer multicast requests on MI networks. Specifically, we create a set of multi-hop lightpath routes in the overlay network from the source node of the requests to the destination nodes and aim at minimizing the total number of wavelengths used to satisfy the requests. We present two solutions: drop at member node (DMN), wherein we restrict the termination of a lightpath only to members of the multicast session, whereas in the drop at any node (DAN) model, we allow a lightpath to be terminated at any node. For the case of static traffic, we present integer linear programs (ILPs) to solve the DMN and DAN overlay problems that minimize the number of wavelengths used to satisfy the multicast requests. We also present an efficient heuristic to approximate the above ILP solutions over realistic amounts of time and compare its performance to that of the ILP for small networks. A theoretical lower bound to calculate the minimum number of wavelengths is also introduced, by which we can evaluate the performance of the aforementioned heuristic. Moreover, we present an ILP for the case of achieving multicast via WDM unicast (MVWU), i.e., the naïve case of establishing single-hop unicast connections, and compare its performance to both DMN and DAN. For dynamic traffic, we evaluate the performance of the heuristic (minimize the number of wavelengths) on real-world large-scale networks. Finally, we evaluate the blocking performance on these networks for a fixed number of wavelengths. The rest of the paper is organized as follows. Section II defines each of the multicast overlay problems formally. In Section III, we present our ILP formulations for MVWU, DMN, and DAN, while Section IV presents sub-optimal heuristics that efficiently tackle these problems. The lower bounds are presented in Section V, after which we evaluate the ILPs and heuristics on various network topologies and under several traffic scenarios in Section VI. Section VII concludes the paper.

II. Problem Definition

Given a network topology graph $G = (V, E)$, where V is the set of nodes in the network and E is a set of fiber links (edges) that connect nodes, a multicast request may be defined as $R = (s_r, D_r)$, where s_r is the source node of the request ($s_r \in V$) and D_r is the set of destination members ($D_r \subseteq V - \{s_r\}$). For a multicast request with K destination members, we represent the destination set as $D_r = \{d_{r_1}, d_{r_2}, \ldots, d_{r_K}\}$. According to the MVWU approach to solving the overlay RWA problem, we establish lightpaths from the source node of a multicast request to each destination member of the request, i.e., we establish lightpaths from s_r to each d_i ($d_i \in D_r$) $\forall i = 1, 2, \ldots, K$. Note that in doing so, we create unicast routes (single logical hop) from the source node to each destination node. In the DMN model, we find a set of lightpath routes that start at the source node and reach each destination member, possibly via multiple logical hops. While creating these lightpath routes we take into account that we can terminate (i.e., drop) a lightpath only at nodes that belong to the destination set.

In Fig. 1 we show a simple six-node network with bi-directional links and illustrate how the MVWU and DMN models service a set of two multicast requests: $R_1 : \{1, (2, 5, 6)\}$ and $R_2 : \{4, (2, 3, 5)\}$. Figure 1(a) shows how the MVWU model establishes a single lightpath from the source node to every destination node (member) of the multicast request. For example, for request R_1, MVWU establishes a single lightpath from the source (node 1) to each destination member (nodes 2, 5, 6). It can be seen that the minimum number of wavelengths required by MVWU to service requests R_1 and R_2 is 3. Let us name these three wavelengths: λ_1, λ_2, and λ_3. The lightpaths from node 4 to nodes 2 and 5 both use λ_1. Lightpaths from 1 to 2 and 4 to 3 reserve λ_2, and the remaining paths must use λ_3. Figure 1(b) shows how the set of lightpath routes is created for the DMN model. For example, for request R_1, the set of lightpath routes created is as follows: a lightpath from node 1 to node 2, from node 2 to node 5, and from node 5 to

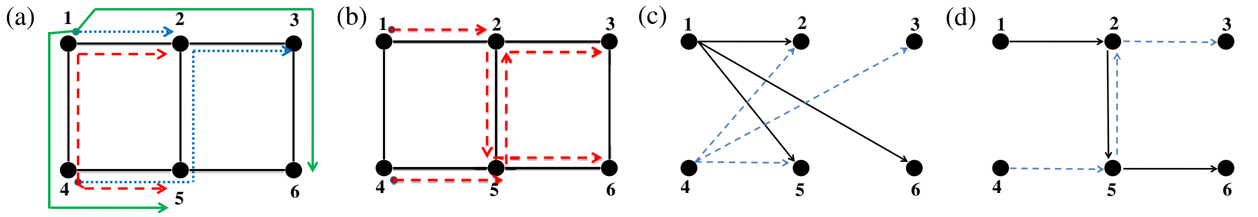

Fig. 1. (Color online) Illustration: number of wavelengths required to satisfy the multicast requests $R_1 : \{1,(2,5,6)\}$ and $R_2 : \{4,(2,3,5)\}$ using (a) MVWU and (b) DMN; the corresponding logical overlays for (c) MVWU and (d) DMN.

node 6. Notice that since node 2 is a destination member of the multicast request, we can drop a lightpath at this node. To service these multicast requests, DMN utilizes only a single wavelength. This decrease in wavelength usage comes at the expense of an increase in the average number of logical hops (2 for DMN and 1 for MVWU), which results in a greater end-to-end delay. We calculate these average logical hops as the average number of lightpaths needed to reach each destination in the multicast set. Furthermore, in Figs. 1(c) and 1(d), we represent the corresponding logical overlays for the MVWU and DMN models, respectively.

The overlay RWA multicast problems may be defined as follows.

Definition (MVWU): Given a network $G = (V,E)$ and a set of multicast requests $\overline{R} = \{R_1,R_2,\ldots,R_R\}$, the solution must assign, for each request, a lightpath and a wavelength from the source node of the request to each destination member of the request, in such a manner that the number of wavelengths required is minimized while satisfying the wavelength continuity constraint (WCC).[1]

Definition (DMN): Given a network $G = (V,E)$ and a static set of multicast requests $\overline{R} = \{R_1,R_2,\ldots,R_R\}$, the solution must assign a set of lightpath routes (starting at the source node) reaching each member of the multicast request in such a way that the number of wavelengths required is minimized while satisfying the WCC. The solution must take into consideration that no lightpath terminates/originates at a non-member node, i.e., all lightpaths must end at a node belonging to the multicast destination set.

Definition (DAN): The DAN problem is similar to DMN, except that it allows the flexibility to terminate/originate a lightpath even at/from a non-member of the request set.

Considering the example in Fig. 1, it can be easily verified that DAN also requires a single wavelength to service both the requests. There are, however, some cases where DAN tends to outperform DMN. Note that we are minimizing the network-wide wavelength count, not the maximum number of wavelengths on any single link [22].

III. ILP Formulations

In this section we formulate the ILPs to solve the multicast overlay problem, with the objective of minimizing the total number of wavelengths used throughout the network. Note that this is different from minimizing the maximum number of wavelengths on any particular link. We present two ILPs to solve the DMN and DAN problems. The ILPs are not practical for large networks, so we introduce a sub-optimal heuristic to approximate the ILP performance in the following section. We can apply the ILPs to smaller networks to compare the results of our heuristic to the optimal solution. We first formulate the MVWU ILP and then present the DMN and DAN ILPs, each with the objective of minimizing wavelength consumption while servicing the multicast requests. In all situations, we allow a multicast source to transmit to each destination on possibly different wavelengths. However, for any single lightpath established in the overlay network, we adhere to the WCC. In DMN and DAN, we service a given multicast request by setting up lightpath routes (multi-hop) in the overlay network, wherein we potentially terminate the lightpath at some intermediate node. This incurs an O–E–O conversion at that node, and it is thus fair to allow a free wavelength conversion there. For all the following ILPs, we use the notations u and v to denote the source node and a destination node of a multicast session in the overlay network, and use i and j to denote the end points of a physical link (WDM layer).

A. Multicast Via WDM Unicast (MVWU)

In [22], the authors propose an ILP for the static lightpath establishment problem wherein they restrict the number of lightpaths for a given pair of nodes to one, and aim at minimizing the maximum number of wavelengths on any link in the network. In their other work [3], the authors remove the above-mentioned restriction and formulate an ILP with no WCC, i.e., they assume wavelength converters in the optical network. In what follows, we build on the ILP from [3] taking into account the WCC and formulate an ILP to satisfy a set of multicast requests by setting up unicast (single-hop) lightpaths from the source node of the request to each destination member of the request individually, with the goal of minimizing the number of wavelengths required.

The ILP can be formulated as follows. Given the following input parameters:

$\overline{N}$ is the set of nodes in the network, which are numbered 1 through N, i.e., $\overline{N} = \{1,2,\ldots,N\}$.

A_{ij} is 1 if a physical link exists between node i and node j.

$\overline{R}$ is the set of multicast requests, which are numbered 1 through R, i.e., $\overline{R} = \{R_1,R_2,\ldots,R_R\}$.

[1] A lightpath must occupy the same wavelength across every physical link of the network it traverses.

$\overline{W}$ is the set of wavelengths per link. The wavelengths are numbered 1 through W, i.e., $\overline{W} = \{1, 2, \ldots, W\}$.

K is the number of destination nodes (members) of a multicast request.

We denote the source node of a multicast request (r) as s_r and the set of destination nodes as D_r, i.e., $D_r = \{d_{r_1}, d_{r_2}, \ldots, d_{r_K}\}$. The input parameter for lightpath setup may be defined as

$T^r_{u,v}$ is 1 if a lightpath is to be established for request r from node u to destination node v.

The ILP will solve for the following variables:

$L^{r,w}_{u,v}$ is 1 if a lightpath is established for request r from node u to node v on wavelength w.

$F^{r,w}_{u,v,i,j}$ is 1 if there is a flow on the physical link from node i to node j on wavelength w for lightpath request r, from node u to node v.

MaxIndex represents the largest wavelength index used.

Objective Function:

$$\text{minimize : MaxIndex}$$

Subject to:

$$\text{MaxIndex} \geq L^{r,w}_{u,v} \times w \quad \forall r \in \overline{R},\ \forall u, v \in \overline{N}, \forall w \in \overline{W}. \tag{1}$$

$$\sum_w L^{r,w}_{u,v} = T^r_{u,v} \quad \forall r \in \overline{R},\ \forall u, v \in \overline{N}. \tag{2}$$

$$\sum_r \sum_u \sum_v F^{r,w}_{u,v,i,j} \leq 1 \quad \forall i, j \in \overline{N},\ \forall w \in \overline{W}. \tag{3}$$

$$F^{r,w}_{u,v,i,j} \leq A_{ij} \times L^{r,w}_{u,v} \quad \forall u, v, i, j \in \overline{N},\ \forall w \in \overline{W}. \tag{4}$$

$$\sum_i F^{r,w}_{u,v,i,j} - \sum_k F^{r,w}_{u,v,j,k} = \begin{cases} 0 & \text{if } j \neq u, v \\ L^{r,w}_{u,v} & \text{if } j = v \\ -L^{r,w}_{u,v} & \text{if } j = u. \end{cases} \tag{5}$$

Constraint (1) is used to keep track of the maximum wavelength index used. Constraint (2) specifies that for a given multicast request, a lightpath has to be established from the source node of the request to each destination node of the request. Constraints (3)–(5) are the physical-layer constraints. Constraint (3) states that a wavelength on every physical link can be used by at most one multicast request. Constraint (4) specifies that for a given lightpath set up in the overlay layer, routing in the WDM layer is possible only if a physical link exists. Constraint (5) is the standard flow conservation constraint, which states that the sum of inflow is equal to the sum of outflow for any intermediate node on the route of a lightpath, and it is set to ± 1 if the node in question is either the source node or the destination node of the lightpath.

B. Drop at Member Node (DMN)

In this section we present the ILP formulation for the DMN overlay problem. In addition to the input parameters stated in Subsection III.A, we define

Q is the number of non-destination nodes (non-members) of a multicast request ($Q = (|\overline{N}| - K) - 1$).

For a given multicast request, we denote the source node of the request as s_r and the set of destination member nodes is represented as D_r, i.e., $D_r = \{d_{r_1}, d_{r_2}, \ldots, d_{r_K}\}$. The set of non-destination members is denoted as $D'_r = \{d'_{r_1}, d'_{r_2}, \ldots, d'_{r_Q}\}$. To solve the ILP we use the same notations for $L^{r,w}_{u,v}$ and $F^{r,w}_{u,v,i,j}$ and also define the following additional variables:

C^r_w is 1 if wavelength w is used to service multicast request r, else it is 0.

X^r_u: the order node u has been added to the multicast request r. This prevents loops from being formed by the set of lightpath routes in the overlay layer.

Objective Function:

$$\text{minimize : MaxIndex}$$

Subject to:

$$\text{MaxIndex} \geq C^r_w \times w \quad \forall r \in \overline{R},\ \forall w \in \overline{W}. \tag{6}$$

$$\sum_w C^r_w \geq 1 \quad \forall r \in \overline{R}. \tag{7}$$

$$L^{r,w}_{u,v} \leq C^r_w \quad \forall r \in \overline{R},\ \forall w \in \overline{W},\ \forall u, v \in \overline{N}. \tag{8}$$

$$\sum_w \sum_v L^{r,w}_{s_r,v} \geq 1 \quad \forall r \in \overline{R}. \tag{9}$$

$$\sum_w \sum_v L^{r,w}_{v,u} = 1 \quad \forall r \in \overline{R},\ \forall u \in D_r. \tag{10}$$

$$\sum_w \sum_v L^{r,w}_{v,s_r} = 0 \quad \forall r \in \overline{R}. \tag{11}$$

$$\sum_w \sum_v L^{r,w}_{u,v} = 0 \quad \forall r \in \overline{R},\ \forall u \in D'_r. \tag{12}$$

$$\sum_w \sum_v L^{r,w}_{v,u} = 0 \quad \forall r \in \overline{R},\ \forall u \in D'_r, u = s_r. \tag{13}$$

$$X^r_u - X^r_v + |N| \times L^{r,w}_{u,v} \leq |N| - 1 \quad \forall r \in \overline{R}, w \in \overline{W}, u, v \in \overline{N}. \tag{14}$$

We use the same physical-layer constraints as Subsection III.A, i.e., constraints (3)–(5). The constraints for the DMN model can be explained as follows. Constraint (6) is used to keep track of the maximum wavelength index used. Constraints (7)–(14) are used to build the set of lightpath routes to satisfy a multicast request. Constraints (7) and (8) ensure that we build the set of lightpath routes for a given multicast request by using at least one wavelength. Constraint (9) specifies that there must be at least one lightpath emerging from the source node of a request. Constraint (10) ensures that in order to service each destination member of the multicast request, there must be a lightpath on some wavelength that terminates at each member node. Constraint (11) specifies that there are no lightpaths that terminate at the source node of a request. Constraints (12) and (13) ensure that no lightpaths terminate at/originate from a non-member of a multicast request set. Finally, constraint (14) is used to prevent the formation of routing loops.

C. Drop at Any Node (DAN)

To formulate the ILP for the DAN problem, we use the same variable definitions used in DMN. To provide for the flexibility of dropping a lightpath at any node, we present below the necessary changes to be made to the DMN ILP (Subsection III.B). We remove constraints (12) and (13) and replace them with the following constraints:

$$\sum_{u}\sum_{w} L_{u,v}^{r,w} \le 1 \quad \forall r \in \overline{R},\; u \neq s_r. \tag{15}$$

$$\sum_{v}\sum_{w} L_{u,v}^{r,w} - |N| \times \sum_{v}\sum_{w} L_{v,u}^{r,w} \le 0 \quad \forall r \in \overline{R}, u \neq s_r. \tag{16}$$

$$\sum_{u}\sum_{w} L_{u,v}^{r,w} - \sum_{u}\sum_{w} L_{v,u}^{r,w} \le 0 \quad \forall r \in \overline{R}, \forall v \in D'_r. \tag{17}$$

Constraint (15) specifies that all nodes (except the source) can have at most one incoming lightpath, while constraint (16) specifies that a node can have an outgoing lightpath only if it has at least one lightpath that terminates at it. Finally, constraint (17) specifies that nodes not in the candidate destination set that have an incoming lightpath must have at least one outgoing lightpath.

IV. Multicast Overlay Heuristics

Algorithm 1: Shortest Path Overlay Heuristic (SPOH)

input : Multicast Request: $R = (s_r, D_r) = s : \{d_1, d_2, \ldots, d_K\}$
List of appropriate SP routes $routeList$
output: Multicast overlay tree yielding the fewest number of additional wavelengths to the network

1 $AltTrees[K] = NULL$
2 $routedDestinations = 0$
3 $sort(routeList)$
4 **for** *each* $d_i \in D_r$ **do**
5   $Tree_k = NULL$
6   $Tree_k.add(SP(s, d_i))$
7   $update\ routedDestinations$
8   **for** *each* $route_i \in routeList$ **do**
9   **if** $routedDestinations = K$ **then**
10       $break$
11   **if** $(route_i.source \in Tree_k)$ & $(route_i.dst \notin Tree_k)$ & $(route_i.dst \in D_r)$ **then**
12       $Tree_k.add(route_i)$
13       $update\ routedDestinations$
14   $AltTrees.add(Tree_k)$
15 **return** $min(AltTrees)$

In the previous section, we formulated ILPs to solve the MVWU, DMN, and DAN problems. The primary disadvantage of modeling the multicasting functionality in the overlay network using an ILP is its inherent complexity, which limits the scope of request set sizes to a few tens of requests. In this section we propose a heuristic for providing multicast capability over MI networks, which implements in polynomial time both the DMN and DAN solutions. Our shortest path overlay heuristic (SPOH), described in Algorithm 1, can be applied to a much greater number of requests and still provide acceptable run time complexity for realistic request sets. In Algorithm 1, we depict the steps taken by SPOH to provision a single multicast request. For a static set of multicast requests, we provision the set by considering a sequential order of the request i.d., i.e., we first provision request R_1, and then proceed to request R_2, and so on until all requests have been provisioned. For the case of dynamic traffic, we provision the requests based on their arrival pattern which follows some stochastic distribution.

The input to SPOH is a set of predetermined routes called the *routeList*. Note that in the case of DMN, we allow a signal to be dropped only at a member (node) of the multicast request. Hence, in this case, the *routeList* contains the set of all shortest path (SP) routes from a multicast request member to every other member of the request. In DAN, we have no such restrictions and potentially allow the signal to be dropped at any node in the network. In this case the *routeList* contains the set of all unicast SP routes from every node in the network to every other node. The routes in this list are sorted (as shown in line 3) with the shortest routes (in terms of physical hop count) first. The basic idea behind SPOH is to generate K-alternate logical multicast trees, where $K = |D_r|$, i.e., the number of nodes in the destination set of a multicast request. These alternate trees are stored in an array of size K named *AltTrees* (line 1). Each of these alternate trees is determined within the loop starting at line 4. For each multicast tree corresponding to a particular multicast destination d_i, the first logical hop is mandated to be the SP from the multicast source s_r to d_i, and this hop is added to the tree at line 6. Note that it is possible that this first hop to d_i may actually route through one or more other $d_j \in D_r$, and this appropriate number of routed multicast destinations is updated in line 7. SPOH then inspects each route in the now sorted *routeList* to determine its suitability for inclusion in the multicast tree. In order for a given route to be included in the tree, its source node must already be in the tree, and its destination must not. However, the route's destination must be a member of D_r. The first route in the *routeList* which satisfies these criteria will be added to the tree as in line 12. This continues until all K destinations have been routed, at which point the tree will be added to the collection of alternate trees as in line 14.

Once K-alternate trees have been identified, the least-cost tree (defined as the tree that results in the smallest number of additional wavelengths required to provision a request) is selected as the multicast overlay tree. Note that if more than one tree shares this minimum wavelength count, the tie is broken by selecting the tree which incurs the fewest logical hops from s_r to all K destinations.[2] In what follows, we first illustrate, with the help of an example, how SPOH generates K-alternate trees for a given multicast request and then conclude this section by outlining the complexity of our heuristic.

In Fig. 2, we consider a simple six-node network and a single multicast request, $R = 2 : \{4, 5, 6\}$, with $K = 3$ destination nodes and show how SPOH generates $K = 3$ multicast overlay trees to provision this request. Note that the SPs in the *routeList* are ordered according to a simple bakery algorithm. For example, if node 1 is to route to node 5, it will route along node 2 rather than node 4 since node 2 has a lower index. Similarly assume that if a route from s_r to a destination d_i has the same weight

[2] In Section VI, we also consider request blocking in our evaluation of SPOH, wherein we consider each link to be equipped with a fixed number of wavelengths. In this case, SPOH blocks a multicast request if no tree can be added to *AltTrees* due to wavelength unavailability.

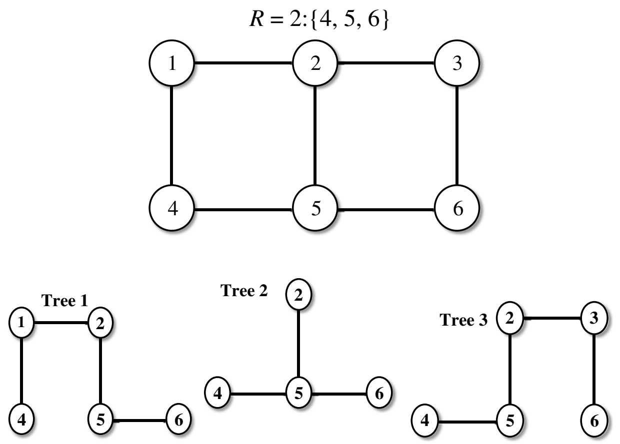

Fig. 2. Illustrative example of SPOH for a specific multicast request.

as a route from d_j to d_i, the route from s_r is chosen over its competitor. The first tree requires that its first hop be from node 2 (source node) to node 4. The corresponding route is via node 1. The next shortest route is directly from node 2 to node 5, and the last hop will be from node 5 to node 6. The second tree will first route from node 2 to node 5. This now means that the SPs to both nodes 4 and 6 are sourced at the closest member, in this case node 5. Finally, the third tree will first route from node 2 to node 6 (via node 3). The next two routes included are from node 2 to node 5, and from node 5 to node 4, respectively. These distinct alternate trees are shown in Fig. 2. The current state of the network largely affects which of these trees is chosen for provisioning the request. We select the tree that requires the smallest number of wavelengths to be added to the network in order to provision the request. That is, the tree added is the one which adds fewest wavelengths to the existing topology. Let us assume that trees 1 and 2 require no new wavelengths, and tree 3 requires 1 new wavelength. Tree 3 is therefore no longer in contention, and the tie (between trees 1 and 2) is broken by calculating the number of logical hops required to provision the multicast request. Note that tree 2 contains 1 logical hop to destination node 5, and two logical hops each to reach nodes 4 and 6. Thus it requires a total of 5 logical hops to provision the request. Tree 1, on the other hand, requires a total of 4 logical hops (1 logical hop each to reach nodes 4 and 5, and 2 logical hops to reach node 6) to provision the request. Hence for the given request, tree 1 is selected as the multicast overlay tree.

Complexity: The network topology can be represented as a graph $G = (V,E)$ with V nodes and E links. Without considering the complexity of the underlying unicast lightpath reservations for an overlay multicast tree, the complexity of identifying a single tree can be represented by the product of the number of routes traversed in the *routeList*, $O(V^2)$, and the maximum number of nodes to traverse in a single tree, $O(V)$. Considering that SPOH establishes K-alternate trees, the worst case complexity of Algorithm 1 can be expressed as $O(KV^3)$.

In the following section, we compare the performance of the DMN, DAN, and MVWU ILPs on a simple 6-node network. Similarly, we also compare the performance of the SPOH heuristic for DMN and DAN with the performance of a simple shortest path unicast heuristic (SPUH), which provides a sub-optimal solution to the MVWU problem by establishing unicast lightpaths from the source to every destination member such that there is always exactly one logical hop to each member. Due to space restrictions and its inherent simplicity, we do not include a full description of SPUH here.

V. Lower Bounds

In this section we define a lower bound on the number of wavelengths required for the multicast overlay models (DMN/DAN). Our derivations are based on the work in [23]. We derive two lower bounds, one based on comparing the logical (overlay) nodal degree and the physical nodal degree of the nodes, LB_1, and one based on congestion of the links, LB_2. Let $d_{(out)_j}$ be the physical nodal degree (out-degree) of node j and $d_{(in)_j}$ be the physical in-degree of node j. Since we consider an undirected graph, $d_{(out)_j} = d_{(in)_j} = d_j$, i.e., the in-degree and out-degree are the same for each node. For a given set of requests $\overline{R} = \{R_1, R_2, \ldots, R_R\}$, we calculate the following:

$O(j)$: the number of lightpaths originating from node j in the overlay layer.

$T(j)$: the number of lightpaths terminating at node j in the overlay layer.

For example, given a request set $\overline{R} = \{R_1, R_2, R_3\}$, with $R_1 : \{1,(4,5,6)\}$, $R_2 : \{3,(1,4,5)\}$, and $R_3 : \{2,(1,3,5)\}$, $O(1) = 1$, $T(1) = 2$, $O(5) = 0$, $T(5) = 3$, and so forth. With this in mind the lower bound LB_1 is calculated as follows:

$$LB_1 = \max_j \left\{ \max \left\{ \left\lceil \frac{O(j)}{d_j} \right\rceil, \left\lceil \frac{T(j)}{d_j} \right\rceil \right\} \right\}.$$

Note that we take the ceiling since the number of lightpaths must be an integer. The lower bound simply finds the node that requires the most wavelengths. This lower bound does not take the routing of requests into account; therefore we propose another lower bound that does.

The next lower bound, LB_2, is derived by taking the routing of trees into account, providing a lower bound on congestion over any link, which is in turn a lower bound for the number of wavelengths required. Let L_r be the minimum number of links required to route request r. In [6], it is shown that $L_r = (\min_d \{SP(s_r, d)\}, d \in D_r) + |D_r| - 1$, where SP is the number of links along the shortest path. Note that in the case of the DMN model, $d \in D_r$ necessarily consists of members of the multicast request. However, in the case of the DAN model, node d may be any node in the network.

$$LB_2 = \left\lceil \frac{\sum_{r \in \overline{R}} L_r}{2|E|} \right\rceil.$$

This lower bound finds the congestion and uses it as a lower bound for the number of wavelengths required. This provides the total minimum number of links required by all of the

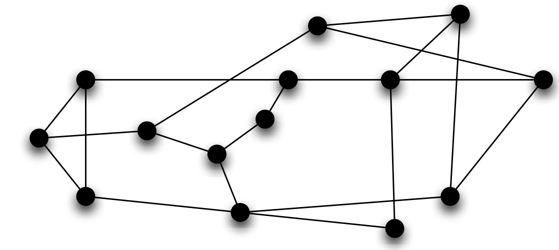

(a) 14-node NSFnet.

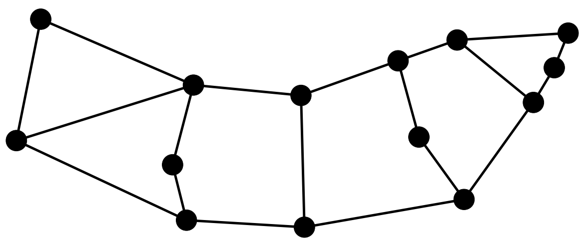

(b) Augmented Energy Sciences Network (ESnet).

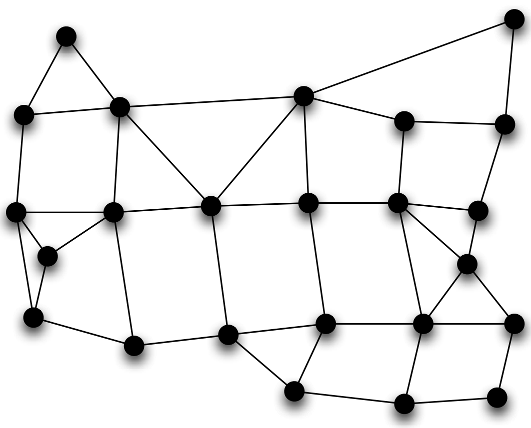

(c) 24-node network.

Fig. 3. WDM mesh networks used for heuristic evaluation.

requests. The average number of requests using each physical edge is this total divided by twice the number of edges (since we assume bi-directional links). The final lower bound evaluated in the next section is $LB = \max\{LB_1, LB_2\}$.

VI. Performance Evaluation

In this section we evaluate the performance of DMN and DAN versus MVWU. We first evaluate the performance of the ILPs by comparing them to the MVWU ILP and then to both SPOH and SPUH on the small 6-node network shown in Fig. 1. The heuristics are then subjected to large sets of static requests and then compared to each other on the three realistic networks shown in Fig. 3. Both the optimal ILP solutions and the heuristic approximations are compared to the lower bounds calculated in the previous section. Furthermore, the heuristics are subjected to dynamic multicast traffic and compared in terms of both wavelength consumption and blocking probability when the number of network-wide wavelength resources is fixed. In each subsection that follows, we present the average values calculated for 30 independent request sets.[3]

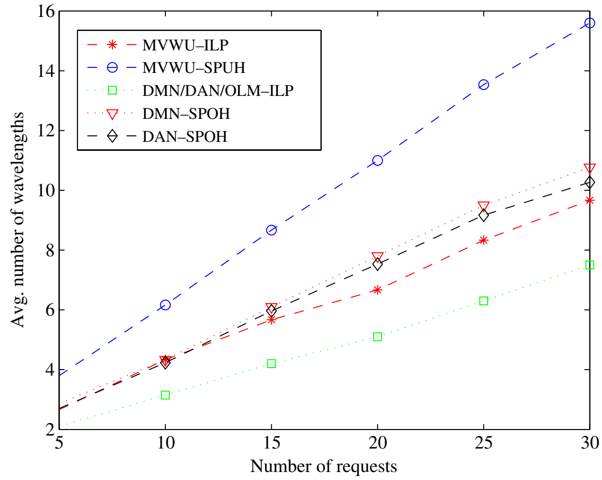

Fig. 4. (Color online) Performance comparison of MVWU and DMN for $K = 2$.

A. ILP Results

Because of the inherent complexity of the MVWU/DMN/DAN ILPs, results can only be obtained on small networks. All ILPs are solved with CPLEX version 12.1. Both the ILP and heuristics are run on a machine with a 2.33 GHz Quad Core Xeon processor and 8 GB of RAM. The processor also has hyper-threading so CPLEX is able to use eight threads while solving the ILPs. We run the ILP and the heuristics (SPOH and SPUH) for request set sizes of $R = 5, 10, \ldots, 30$. The source node of each multicast request is uniformly distributed and the destination set size, K, is fixed at 2, 3, and 4 destinations. We imposed a 24 h limit to ILP executions. Consequently, there were a few instances (only 4 of 30 for $K = 3$ and 4 and number of requests = 30) where the 24 h limit was exceeded without finding the optimal solution. In these rare cases, we used the calculated feasible solution. It is likely, however, that given the size of the problem, this feasible solution is the same as the elusive optimal solution (this feasible solution was never higher than the corresponding solution the heuristic found).

In Figs. 4 and 5 we show the minimum number of wavelengths required by the ILPs and the heuristics to satisfy all requests with no blocking for values of K = 2 and 3, respectively. It can be observed that for the 6-node network, both the DMN and DAN ILPs produce identical results. We also compare our overlay approach for achieving multicasting to the case of optical-layer multicasting (OLM) (i.e., nodes capable of splitting at the optical layer) [24]. We run the ILP proposed in [24] on the same network and compare its performance to those of DMN and DAN. For the range of

[3] We also compute confidence intervals, but do not plot them here on any of the curves to avoid cluttering of the figures. However, the 95% confidence intervals are very narrow in all cases.

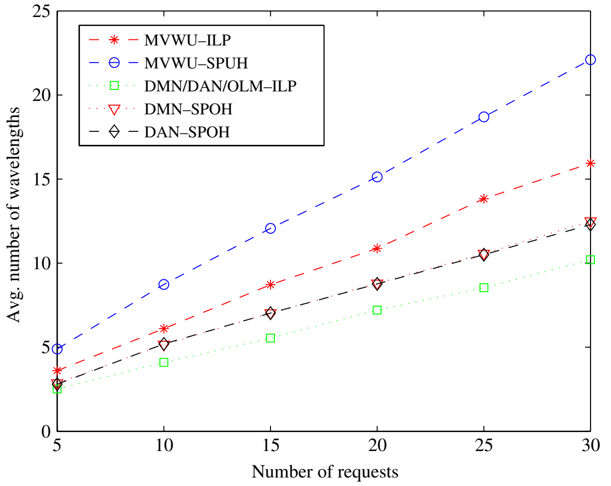

Fig. 5. (Color online) Performance comparison of MVWU and DMN for $K = 3$.

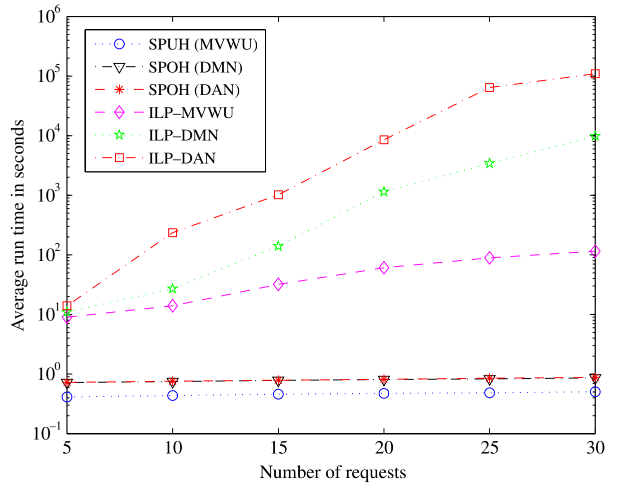

Fig. 6. (Color online) Comparison of run times for ILP and heuristic (SPOH) for $K = 2$.

request sizes considered, both overlay models produce identical results to each other and to the OLM approach.

It can be observed that for values of $K = 2$ and 3, DMN/DAN use 2 to 5 wavelengths fewer than MVWU at higher loads. Also note that our proposed SPOH heuristic performs well as compared to the ILPs. In fact, both the DMN and DAN approximations of SPOH outperform the MVWU ILP for $K = 3$. We have also simulated the above ILPs with a variable destination set for each request, i.e., the destination set size is uniformly distributed between 2 and 4 for any given request. We do not present those results here, as a similar trend was observed.

In Fig. 6, we plot the run times of the ILPs and SPOH for $K = 2$. We can observe the large increase in run times for the ILPs as the request set gets larger. Note the logarithmic scale of the y-axis. SPOH takes time of the order of a few milliseconds to complete, whereas the run time of the ILP grows rapidly. This shows that the ILP is not practical, given the run times for the small set sizes and small network, and given the relatively small difference in wavelength consumption from heuristic to ILP, SPOH provides a reasonable trade-off given the execution time savings.

In Fig. 7 we compare the average number of logical hops taken by DMN and DAN (ILPs) for $K = 2$ and 3. It is seen that DAN takes a slightly higher number of logical hops and produces the same results for the average number of wavelengths required to satisfy a given multicast request set. Note that MVWU has an average number of logical hops of 1 as we are establishing end-to-end lightpaths from the source node of a request to each destination member of the multicast request.

B. Lower Bound Results

In this section we compare the DMN/DAN ILPs and SPOH heuristics to the theoretical lower bounds introduced in Section V. As mentioned earlier, we have defined two lower bounds and take the maximum of the two as the actual bound, LB. This lower bound is not the actual minimum number of wavelengths required, but a theoretical lower bound. We first compare our ILP with the lower bound to show that the bound is reasonable, then compare our heuristics with the bound.

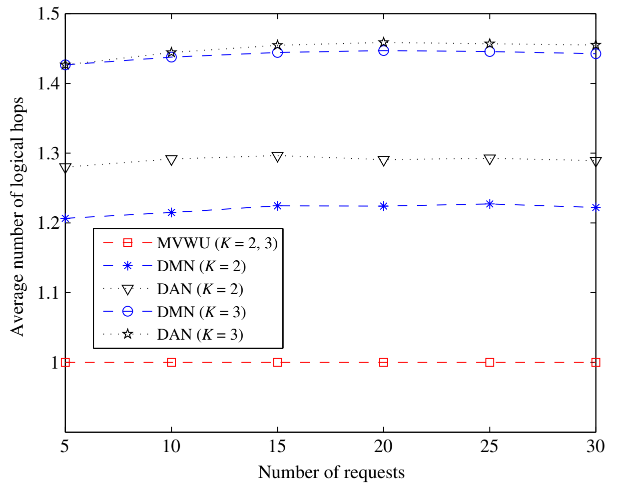

Fig. 7. (Color online) Performance comparison: average number of logical hops.

TABLE I
COMPARISON OF OPTIMAL ILP SOLUTION TO LOWER BOUNDS FOR THE NETWORK SHOWN IN FIG. 1 ($K = 2$)

Requests	LB_{DMN}	LB_{DAN}	ILP	Δ_{DMN}	Δ_{DAN}
5	1.97	1.93	2.1	0.13	0.17
10	3.07	3.00	3.15	0.08	0.15
15	3.93	3.83	4.2	0.27	0.37
20	4.87	4.80	5.1	0.23	0.30
25	5.93	5.80	6.3	0.37	0.50
30	6.83	6.80	7.5	0.67	0.70

In Table I we compare the ILP optimal solutions to the lower bounds for the six-node network shown in Fig. 1. We represent the lower bounds for the DMN and DAN models as LB_{DMN} and LB_{DAN}, respectively, and compute the difference in wavelengths for them as compared to the ILP (represented

TABLE II
COMPARISON OF HEURISTIC APPROXIMATION TO LOWER BOUNDS FOR THE NETWORK SHOWN IN FIG. 1 ($K = 2$)

Requests	LB_{DMN}	LB_{DAN}	$\mathrm{SPOH}_{\mathrm{DMN}}$	$\mathrm{SPOH}_{\mathrm{DAN}}$	Δ_{DMN}	Δ_{DAN}
5	1.97	1.93	2.83	2.67	0.86	0.74
10	3.07	3.00	4.37	4.27	1.30	1.27
15	3.93	3.83	6.17	6.00	2.24	2.17
20	4.87	4.80	7.80	7.50	2.93	2.7
25	5.93	5.8	9.47	9.10	3.54	3.30
30	6.83	6.8	10.80	10.27	3.97	3.47

TABLE III
AVERAGE NUMBER OF WAVELENGTHS REQUIRED ($K = 2$)

		$K = 2$				
Network	R	MVWU	DMN	%DMN	DAN	%DAN
NSFnet	1000	176.00	142.53	19.02	131.43	25.32
	2500	425.43	342.37	19.53	318.73	25.08
	5000	843.30	675.63	19.88	631.07	25.17
	7500	1263.60	1005.87	20.40	941.03	25.53
	10,000	1677.20	1337.87	20.23	1250.17	25.46
24-node	1000	161.23	119.70	25.76	114.50	28.98
	2500	395.33	291.33	26.31	278.93	29.44
	5000	771.77	570.77	26.04	546.67	29.17
	7500	1158.77	850.30	26.62	815.93	29.59
	10,000	1547.73	1131.37	26.90	1084.63	29.92
ESnet	1000	329.07	231.03	29.79	224.83	31.68
	2500	808.60	567.13	29.86	550.87	31.87
	5000	1612.60	1126.37	30.15	1088.90	32.48
	7500	2419.60	1687.97	30.24	1631.40	32.58
	10,000	3216.80	2244.40	30.23	2166.37	32.65

by Δ_{DMN} and Δ_{DAN}, respectively). Recall that for this six-node network, the ILP results for the DMN and DAN models are identical as shown in Fig. 4. From Table I, we can conclude that the lower bound is very accurate, with a maximum difference of 0.6–0.7 wavelengths as compared to the ILP. A similar trend in the results (not shown here) is achieved for $K = 3$ and 4.

We compare the performance of SPOH with the lower bounds in Table II. The table is structured similarly to Table I. The wavelength counts required by SPOH for the DMN and DAN models are shown in columns 4 and 5 ($\mathrm{SPOH}_{\mathrm{DMN}}$ and $\mathrm{SPOH}_{\mathrm{DAN}}$). It can be observed that SPOH performs well as compared to the lower bounds with the maximum difference in the number of wavelengths being approximately 3.4 to 4.0 for higher numbers of requests. It should be noted that as the number of requests grows (or as the network size grows larger), the bounds become less accurate because the routing possibilities become increasingly complicated.

C. Performance Evaluation of Heuristics: Static Traffic

Given the infeasibility of the ILPs, we simulate our heuristics on realistic WDM networks: the 14-node National Science Foundation network (NSFnet), an augmented ESnet topology, and a 24-node mesh network with larger request sets. These topologies are given in Fig. 3. We consider static multicast request set sizes, R, of up to 10,000 requests. The source node of each multicast request is uniformly distributed and the destination set size, K, is set to 2, 3, or 4. Unless explicitly mentioned, we will hereafter refer to SPOH as simply DMN or DAN, and will similarly refer to SPUH as MVWU. Tables III, IV, and V show the average number of wavelengths required to satisfy all requests in a multicast set for $K = 2$, 3, and 4, respectively. As in the previous subsection, we assume that the network is equipped with sufficient wavelength resources to eliminate blocking completely. We also show the percentage improvement (in terms of wavelength usage) of DMN (%DMN) and DAN (%DAN) as compared to MVWU. It is observed that DMN clearly outperforms MVWU and achieves a 20%–50% improvement in average wavelength usage. Note that DAN achieves an additional 2%–5% improvement as compared to DMN, but at the expense of a slightly higher average number of logical hops as shown in Table VI.

D. Performance Evaluation of Heuristics: Dynamic Traffic

In this subsection we present the results of the performance evaluations for dynamically arriving multicast connection requests. Requests arrive according to a Poisson process with average arrival rate λ and exponentially distributed holding times with average service rate μ. The source node of each request is uniformly distributed over all nodes in the network. For each multicast request, the size of the destination set is uniformly distributed across the interval $[2, D_{\max}]$ (where $D_{\max}$ is a parameter which represents the maximum number of multicast destinations). For each request, the destination nodes are also uniformly distributed across the network. We simulated each heuristic with different values for $D_{\max}$, and recorded the average number of wavelengths. Note that each

TABLE IV
AVERAGE NUMBER OF WAVELENGTHS REQUIRED ($K = 3$)

Network	R	$K=3$				
		MVWU	DMN	%DMN	DAN	%DAN
NSFnet	1000	265.23	176.63	33.40	166.53	37.21
	2500	642.20	432.90	32.59	410.33	36.11
	5000	1266.57	856.50	32.38	812.60	35.84
	7500	1886.63	1280.77	32.11	1216.93	35.50
	10,000	2506.83	1704.73	32.00	1619.80	35.38
24-node	1000	236.20	148.53	37.12	140.53	40.50
	2500	584.97	362.27	38.07	343.77	41.23
	5000	1158.83	716.00	38.21	680.17	41.31
	7500	1736.33	1070.43	38.35	1017.40	41.41
	10,000	2310.83	1423.03	38.42	1353.93	41.41
ESnet	1000	496.90	278.63	43.93	271.40	45.38
	2500	1235.40	690.90	44.07	674.50	45.40
	5000	2438.67	1369.70	43.83	1340.90	45.02
	7500	3639.13	2049.17	43.69	2008.80	44.80
	10,000	4841.77	2730.80	43.60	2679.00	44.67

TABLE V
AVERAGE NUMBER OF WAVELENGTHS REQUIRED ($K = 4$)

Network	R	$K=4$				
		MVWU	DMN	%DMN	DAN	%DAN
NSFnet	1000	347.97	205.93	40.82	197.50	43.24
	2500	844.53	504.67	40.24	485.27	42.54
	5000	1681.07	1004.60	40.24	964.60	42.62
	7500	2505.37	1502.80	40.02	1445.40	42.31
	10,000	3333.47	2000.37	39.99	1924.73	42.26
24-node	1000	319.13	169.63	46.85	161.47	49.40
	2500	780.20	415.70	46.72	396.90	49.13
	5000	1549.43	827.00	46.63	789.33	49.06
	7500	2320.30	1234.77	46.78	1177.37	49.26
	10,000	3084.37	1642.77	46.74	1567.13	49.19
ESnet	1000	663.17	320.83	51.62	311.73	52.99
	2500	1628.53	792.00	51.37	772.77	52.55
	5000	3225.60	1583.90	50.90	1546.17	52.07
	7500	4838.53	2370.40	51.01	2313.90	52.18
	10,000	6447.93	3157.30	51.03	3081.90	52.20

TABLE VI
COMPARISON OF AVERAGE NUMBER OF LOGICAL HOPS

Network	R	$K=2$			$K=3$			$K=4$		
		DAN	DMN	Δ	DAN	DMN	Δ	DAN	DMN	Δ
	1000	1.37007	1.17220	0.19787	1.62524	1.35008	0.27517	1.78038	1.53151	0.24887
NSFnet	2500	1.37404	1.17495	0.19909	1.62325	1.34783	0.27541	1.78350	1.53123	0.25227
	5000	1.37523	1.17504	0.20020	1.62287	1.34722	0.27565	1.78509	1.53141	0.25368
	7500	1.37529	1.17532	0.19997	1.62215	1.34695	0.27520	1.78485	1.53123	0.25362
	10,000	1.37545	1.17549	0.19996	1.62258	1.34678	0.27580	1.78484	1.53107	0.25377
	1000	1.39920	1.19125	0.20795	1.67567	1.36301	0.31266	1.89514	1.54682	0.34833
24-node	2500	1.39616	1.19245	0.20371	1.67545	1.36250	0.31295	1.89979	1.54705	0.35274
	5000	1.39647	1.19274	0.20373	1.67530	1.36214	0.31316	1.89907	1.54674	0.35233
	7500	1.39654	1.19286	0.20368	1.67579	1.36167	0.31413	1.90051	1.54779	0.35272
	10,000	1.39717	1.19302	0.20415	1.67514	1.36146	0.31368	1.90026	1.54758	0.35268
	1000	1.38762	1.22253	0.16508	1.67300	1.41984	0.25316	1.89726	1.62033	0.27693
ESnet	2500	1.38939	1.22498	0.16441	1.67279	1.42102	0.25176	1.89466	1.61846	0.27620
	5000	1.38996	1.22617	0.16378	1.67380	1.42188	0.25192	1.89363	1.61797	0.27566
	7500	1.39037	1.22662	0.16375	1.67299	1.42088	0.25210	1.89335	1.61797	0.27537
	10,000	1.38999	1.22668	0.16331	1.67340	1.42024	0.25316	1.89294	1.61778	0.27516

fiber of the network is still assumed to have a sufficient number of wavelengths to eliminate blocking. We present the results of blocking (with a fixed number of wavelengths) in the next subsection. The network load in Erlang is calculated as the

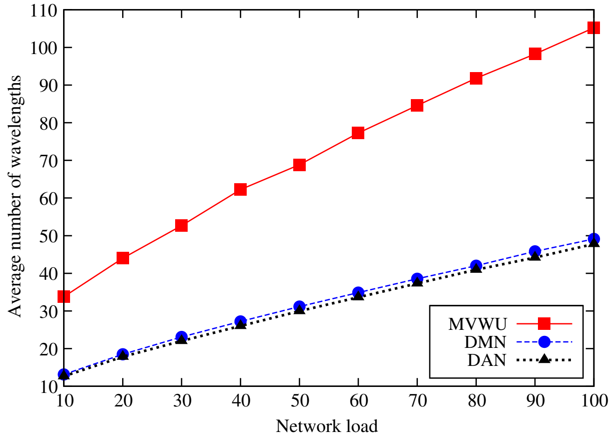

Fig. 8. (Color online) NSFnet: average number of wavelengths ($D_{\max} = 10$).

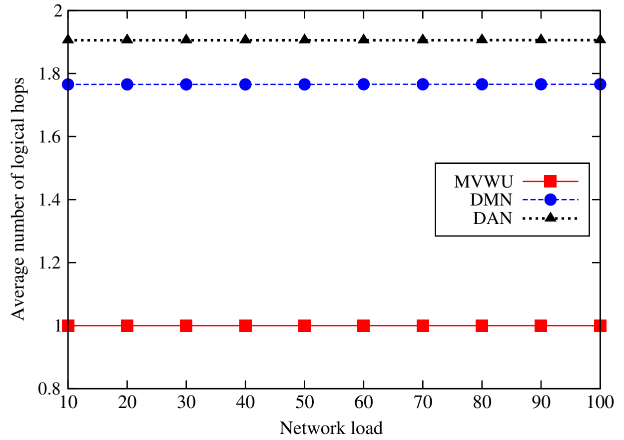

Fig. 10. (Color online) NSFnet: comparison of logical hops ($D_{\max} = 10$).

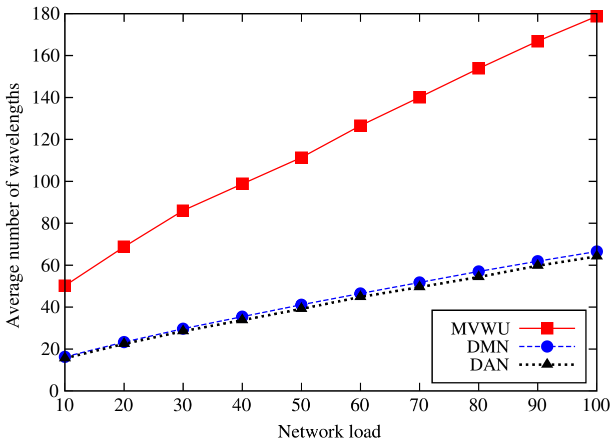

Fig. 9. (Color online) ESnet: average number of wavelengths ($D_{\max} = 10$).

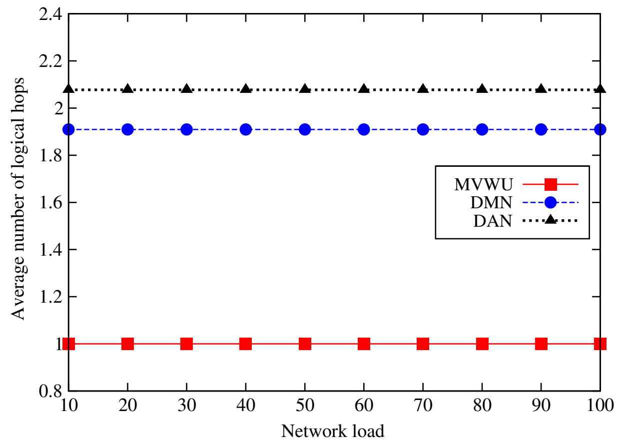

Fig. 11. (Color online) ESnet: comparison of logical hops ($D_{\max} = 10$).

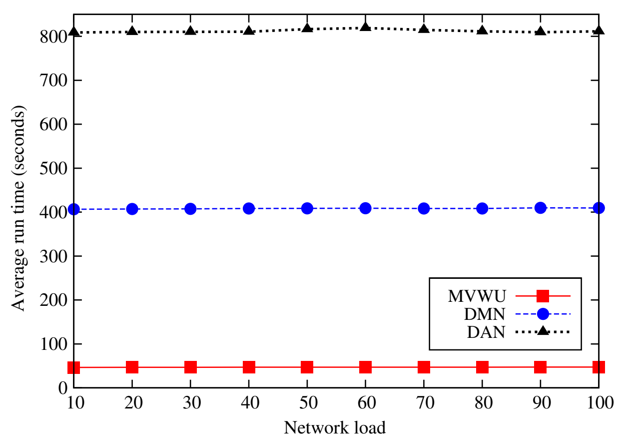

Fig. 12. (Color online) ESnet: comparison of run times ($D_{\max} = 10$).

ratio of the average arrival rate to the average service rate (λ/μ). Each request set consists of 10^5 multicast requests. The results presented in this section represent the average of 30 unique request sets.[4]

In Figs. 8 and 9, we show the average number of wavelengths required by the three heuristics for the NSFnet and the ESnet for $D_{\max} = 10$. For the NSFnet, it is observed that DMN achieves approximately a 53%–61% improvement (in terms of number of wavelengths) over the MVWU heuristic. Similarly, in Fig. 9, it is observed that for the ESnet, DMN achieves 62%–68% improvement over MVWU. For both networks, it is seen that DAN achieves marginal performance improvement (approximately 1%–2%) over DMN. This improvement comes at the expense of a slightly higher average number of logical hops (as shown in Figs. 10 and 11). We have also evaluated the maximum number of logical hops required to provision the multicast requests. Due to space restrictions, we have not included these evaluations, although a similar trend is observed.

[4] We also compute confidence intervals, but do not plot them here on any of the curves to avoid cluttering of the figures. However, we noticed that the 95% confidence intervals were quite narrow in all cases.

TABLE VII
NSFNET: AVERAGE NUMBER OF WAVELENGTHS REQUIRED

	$D_{max}=6$					$D_{max}=8$				
Load	MVWU	DMN	DAN	%DMN	%DAN	MVWU	DMN	DAN	%DMN	%DAN
10	22.90	12.53	11.33	45.27	50.51	28.67	12.97	11.90	54.77	58.49
20	29.90	17.13	15.43	42.70	48.38	37.50	17.53	16.63	53.24	55.64
30	36.43	20.33	19.13	44.19	47.48	45.57	21.70	20.83	52.38	54.28
40	42.53	24.13	22.73	43.26	46.55	52.53	25.97	24.50	50.57	53.36
50	48.47	27.97	26.10	42.30	46.15	59.13	29.83	27.83	49.55	52.93
60	53.33	30.97	29.27	41.94	45.13	65.90	33.17	31.30	49.67	52.50
70	57.83	33.93	32.30	41.33	44.15	72.70	36.30	34.53	50.07	52.50
80	63.27	37.03	35.20	41.46	44.36	78.17	39.77	37.57	49.13	51.94
90	68.27	40.20	37.77	41.11	44.68	83.10	42.93	40.70	48.34	51.02
100	72.83	43.43	41.00	40.37	43.71	88.93	46.40	44.10	47.83	50.41

TABLE VIII
ESNET: AVERAGE NUMBER OF WAVELENGTHS REQUIRED

	$D_{max}=6$					$D_{max}=8$				
Load	MVWU	DMN	DAN	%DMN	%DAN	MVWU	DMN	DAN	%DMN	%DAN
10	33.67	15.10	14.27	55.15	57.62	41.50	15.73	14.63	62.09	64.74
20	45.37	21.70	20.33	52.17	55.18	56.87	22.93	21.17	59.67	62.78
30	56.53	27.53	25.87	51.30	54.25	70.07	28.70	26.93	59.04	61.56
40	67.27	32.43	30.93	51.78	54.01	82.47	34.67	32.07	57.96	61.12
50	77.97	38.07	35.57	51.18	54.38	94.27	40.03	37.40	57.53	60.33
60	87.50	42.73	40.07	51.16	54.21	106.43	45.33	42.27	57.41	60.29
70	97.60	47.30	44.53	51.54	54.37	119.13	50.00	47.30	58.03	60.30
80	105.73	51.93	48.90	50.88	53.75	131.30	55.33	51.93	57.86	60.45
90	115.03	56.57	53.20	50.83	53.75	142.20	60.10	56.13	57.74	60.53
100	122.93	61.27	57.17	50.16	53.50	151.77	64.47	60.77	57.52	59.96

In Fig. 12 we plot the average run times for the three heuristics on the ESnet. It can be verified that the run times for MVWU are of the order of a few tens of seconds. The DMN/DAN heuristics not only create logical lightpath trees in the overlay network, but generate K-alternate trees for every request, and then select the best tree for use in establishing lightpaths on the underlying layer. This increases the computational complexity and it can be observed that the two heuristics take time of the order of a few hundreds of seconds to complete. A similar trend is observed for the NSFnet (not shown).

Tables VII and VIII show the results obtained for the NSFnet and ESnet for $D_{max}=6$ and 8. We show the percentage improvement (in terms of wavelength usage) of DMN (%DMN) and DAN (%DAN) as compared to MVWU. For both networks, it can be observed that DMN clearly outperforms MVWU and DAN achieves marginal improvement over DMN. As the number of multicast destinations increases, the performance of both SPOH implementations improves relative to MVWU, while the improvement of DAN over DMN remains relatively consistent across network loads and request set sizes.

E. Performance Evaluation: Blocking

All aforementioned simulation results assume a sufficient quantity of wavelengths on any given link in the network topology so as to eliminate blocking entirely. In this section we consider the performance of the SPOH and SPUH heuristics in situations with realistic resource availability. Figure 13 shows the average blocking probability for request sets on the NSFnet for a D_{max} value of 10, with 16 wavelengths per link throughout the network. It can be observed that both DMN and DAN greatly outperform MVWU. This is particularly obvious at loads below 50 Erlang. For example, at 50 Erlang, DMN improves blocking performance by about an order of magnitude over MVWU. This improvement becomes increasingly drastic as load decreases (nearly 2 orders of magnitude at 30 Erlang, and 3 orders at 20 Erlang). In all configurations, DAN achieves slightly better performance than the more restrictive DMN heuristic.

Figure 14 shows the blocking probability for the ESnet topology with the same D_{max} and wavelength configurations. A similar trend to Fig. 13 is observable here. It should be noted, however, that for lower loads, all heuristics yield higher blocking values. For instance, MVWU blocks, on average, more than 10% of the incoming multicast requests even at an offered network load value of just 10 Erlang. This in turn leads to more dramatic relative improvement in the performance of DMN and DAN than is observed for the NSFnet. At such low loads, the ESnet allows DMN to outperform MVWU by more than 4 orders of magnitude. As load increases, however, the blocking probability of DMN and DAN ramps up very quickly, so that for higher loads, although still providing noticeable improvement over MVWU, DMN and DAN yield greater blocking than they do for the NSFnet. It can also be observed that on the ESnet, the relative improvement in performance of DAN over DMN is very small, and, in fact, non-existent at very high loads.

We also show how the logical hop counts of the various overlay solutions are affected in realistic resource scenarios. For the same configurations as shown in the previous two figures, we depict the average logical hop counts of all three

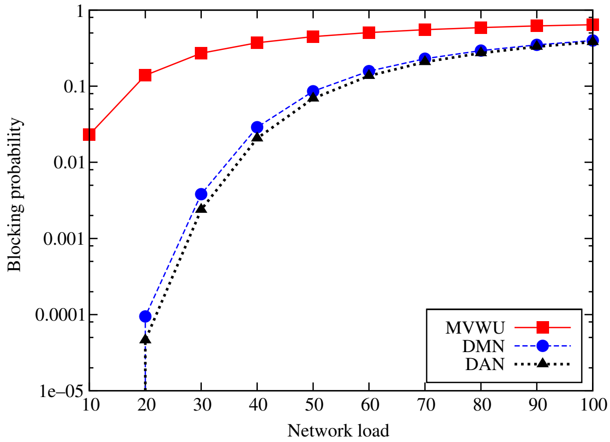

Fig. 13. (Color online) NSFnet: blocking probability ($D_{\max} = 10$, $W = 16$).

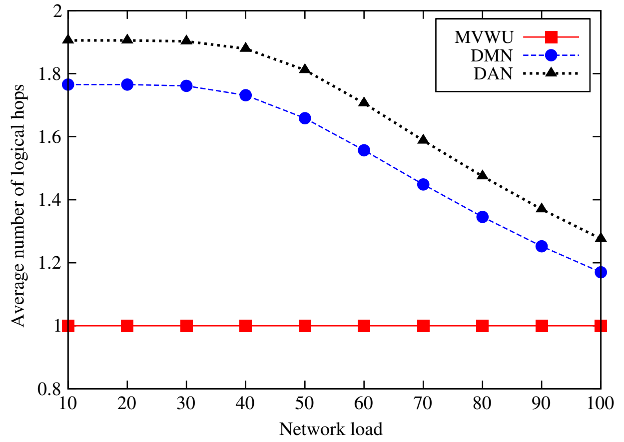

Fig. 15. (Color online) NSFnet: average number of logical hops ($D_{\max} = 10$, $W = 16$).

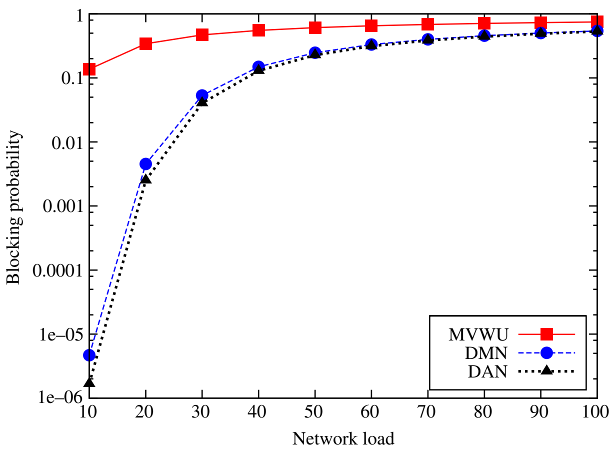

Fig. 14. (Color online) ESnet: blocking probability ($D_{\max} = 10$, $W = 16$).

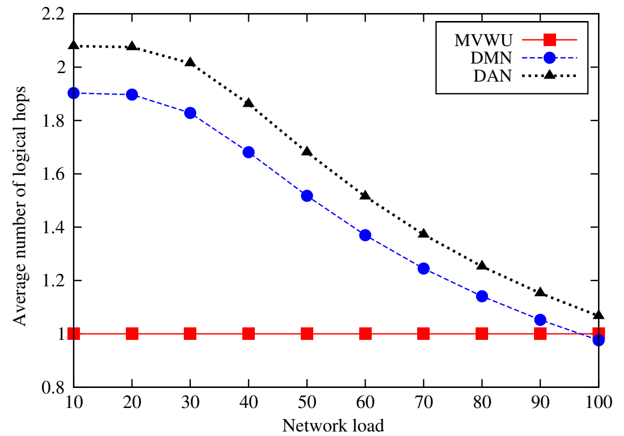

Fig. 16. (Color online) ESnet: average number of logical hops ($D_{\max} = 10$, $W = 16$).

heuristics for the NSFnet and ESnet in Figs. 15 and 16, respectively. In general, both figures follow similar trends. The nature of the MVWU approach mandates that the logical distance from the source to any destination is always 1 hop. As expected, DMN and DAN require increased hop counts. As load increases, so too does blocking (refer once again to Figs. 13 and 14), and consequently the hop count decreases for DMN and DAN. As more requests are fed into the network, more resources (wavelengths) are in use at any given time, thus blocking longer paths which require more wavelengths and provisioning the shorter paths. This leads to an overall decrease in logical hop count as blocking probability increases.

VII. Conclusion

Recent growth in worldwide bandwidth-intensive, large-scale scientific applications has mandated the need for multicast communication abilities in optical WDM networks. An immediate obstacle to provisioning multicast requests is the incompatibility between the communication paradigm and the physical limitations of MI networks. Under such limitations, one must implement multicasting logically as a virtual overlay to the optical layer. In this paper, we have presented two multi-hop solutions to the multicast overlay problem, DMN and DAN, and through extensive simulations have demonstrated their potential to provide substantial savings in terms of network resources over a more traditional naïve, single-hop MVWU overlay mechanism. We have formulated ILPs for these solutions, and by subjecting them to static traffic have shown the significant reduction in wavelength consumption across the network. We have further developed sub-optimal heuristics to compare our multicast overlay solutions for larger request sets on real-world, large-scale networks and shown that DMN and DAN provide superior performance in terms of the number of wavelengths needed to provision requests while eliminating blocking. Further, we have shown that our conclusions for static traffic performance are able to carry over to a more realistic dynamic traffic pattern. By limiting the number of available wavelength resources, we identified tremendous improvements in blocking performance of DMN as compared to MVWU. We also showed that having the flexibility of dropping a signal at any node in the network (DAN) resulted in a marginal blocking performance improvement (at the expense of a higher average number of logical hops) as compared to DMN. Our conclusions provide a strong foundation for the adoption of the DMN overlay solution on real MI networks, such as the DOE's ESnet.

Acknowledgments

This work has been supported by the Department of Energy (DOE) COMMON project under grant DE-SC0004909. Parts of this paper are presented in [25,26].

References

[1] *Networking for the Future of DOE Science* [Online]. Available: https://es.net/assets/Uploads/ESnet4-Networking-for-the-Future-of-Science-2008-05-05.NP.v1.pdf.

[2] R. Malli, X. Zhang, and C. Qiao, "Benefits of multicasting in all-optical networks," *Proc. SPIE*, vol. 3531, pp. 209–220, Nov. 1998.

[3] L. 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.

[4] G. Rouskas, "Optical layer multicast: rationale, building blocks, and challenges," *IEEE Network*, vol. 17, no. 1, pp. 60–65, Feb. 2003.

[5] R. M. Karp, "Reducibility among combinatorial problems," in *Proc. Complexity of Computer Computations*, 1972, pp. 85–103.

[6] H. Takahashi and A. Matsuyama, "An approximate solution for the Steiner problem in graphs," *Math. Japonica*, vol. 24, no. 6, pp. 573–577, 1980.

[7] W. S. Hu and Q. J. Zeng, "Multicasting optical cross connects employing splitter-and-delivery switch," *IEEE Photon. Technol. Lett.*, vol. 10, pp. 970–972, 1998.

[8] J. Leuthold and C. H. Joyner, "Multimode interference couplers with tunable power splitting ratios," *J. Lightwave Technol.*, vol. 19, no. 5, pp. 700–707, May 2001.

[9] M. Ali and S. Deogun, "Power-efficient design of multicast wavelength routed networks," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 10, pp. 1852–1862, Oct. 2000.

[10] G. Sahin and M. Azizoglu, "Multicast routing and wavelength assignment in wide-area networks," *Proc. SPIE*, vol. 3531, pp. 196–208, 1998.

[11] T. Szymanski, "Future internet video multicasting with essentially perfect resource utilization and QoS guarantees," in *Proc. IEEE 19th Int. Workshop on Quality of Service*, 2011, pp. 1–3.

[12] *Worldwide LHC Computing Grid* [Online]. Available: http://lcg.web.cern.ch/lcg/.

[13] *Workshop Report on Advanced Networking for Distributed Petascale Science: R&D Challenges and Opportunities*, 2008 [Online]. Available: http://science.energy.gov/~/media/ascr/pdf/program-documents/docs/Network_research_workshop_report_08.pdf.

[14] H. Zang, J. P. Jue, and B. Mukherjee, "A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks," *Opt. Networks Mag.*, vol. 1, no. 1, pp. 47–60, Jan. 2000.

[15] J. Bermond, L. Gargano, S. Perennes, A. A. Rescigno, and U. Vaccaro, "Efficient collective communication in optical networks," *Theor. Comput. Sci.*, vol. 233, no. 1–2, pp. 165–189, 2000.

[16] X. Zhang, J. Y. Wei, and C. Qiao, "Constrained multicast routing in WDM networks with sparse light splitting," *J. Lightwave Technol.*, vol. 18, no. 12, pp. 1917–1927, Dec. 2000.

[17] R. Libeskind-Hadas and R. Melhem, "Multicast routing and wavelength assignment in multihop optical networks," *IEEE/ACM Trans. Netw.*, vol. 10, no. 5, pp. 621–629, 2002.

[18] B. Chen and J. Wang, "Efficient routing and wavelength assignment for multicast in WDM networks," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 1, pp. 97–109, Jan. 2002.

[19] S. Sankaranarayanan and S. Subramaniam, "Comprehensive performance modeling and analysis of multicasting in optical networks," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 9, pp. 1399–1413, Nov. 2003.

[20] N. Charbonneau, V. M. Vokkarane, C. Guok, and I. Monga, "Advance reservation frameworks in hybrid IP-WDM networks," *IEEE Commun. Mag.*, vol. 49, no. 5, pp. 132–139, May 2011.

[21] J. Kuri, N. Puech, M. Gagnaire, E. Dotaro, and R. Douville, "Routing and wavelength assignment of scheduled lightpath demands," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 8, pp. 1231–1240, 2003.

[22] D. Banerjee and B. Mukherjee, "A practical approach for routing and wavelength assignment in large wavelength-routed optical networks," *IEEE J. Sel. Areas Commun.*, vol. 14, no. 5, pp. 903–908, June 1996.

[23] R. Ramaswami and K. N. Sivarajan, "Design of logical topologies for wavelength-routed optical networks," *IEEE J. Sel. Areas Commun.*, vol. 14, no. 5, pp. 840–851, June 1996.

[24] N. Charbonneau and V. M. Vokkarane, "Routing and wavelength assignment of static manycast demands over all-optical wavelength-routed WDM networks," *J. Opt. Commun. Netw.*, vol. 2, no. 7, pp. 427–440, July 2010.

[25] A. Gadkar, J. Plante, and V. M. Vokkarane, "Static multicast overlay in WDM unicast networks for large-scale scientific applications," in *Proc. ICCN*, July 2011.

[26] A. Gadkar and J. Plante, "Dynamic multicasting in WDM optical unicast networks for bandwidth-intensive applications," in *Proc. Globecom*, Dec. 2011.

Arush Gadkar (S'05, M'10) received his B.E. degree in Electronics and Telecommunication Engineering from the University of Pune, India, in 2002; his M.S. degree in Electrical Engineering from The George Washington University (GWU) in 2004; and his Ph.D. degree in Electrical Engineering from The George Washington University in 2009. He served as an Instructor in the Electrical and Computer Engineering Department at GWU from 2009 to 2010. Currently, he is a Post-Doctoral Research Associate in the Computer and Information Science Department at the University of Massachusetts, Dartmouth. He has been on the technical program committees of the IEEE International Conference on Advanced Networks and Telecommunication Systems (ANTS) 2011 and the International Conference of Communications (ICC) 2012. His research interests include, but are not limited to, future optical network architectures, design and analysis of communication networks, and network optimization.

Jeremy Plante received his B.S. degree with high honors in Computer and Information Science from the University of Massachusetts at Dartmouth in 2011. He is currently pursuing an M.S. degree in Computer Science from the same institute. His research interests include the investigation of non-traditional communication paradigms over large-scale optical core networks.

Vinod M. Vokkarane (S'02-M'04-SM'09) received his B.E. degree with honors in Computer Science and Engineering from the University of Mysore, India, in 1999; his M.S. degree in Computer Science from the University of Texas at Dallas in 2001; and his Ph.D. degree in Computer Science from the University of Texas at Dallas in 2004. Dr. Vokkarane is an Associate Professor of Computer and Information Science at the University of Massachusetts, Dartmouth. Currently, he is a Visiting Scientist at the Research Laboratory of Electronics (RLE) at Massachusetts Institute of Technology (MIT). His primary areas of research include design and analysis of architectures and protocols for optical and wireless networks. Dr. Vokkarane is the co-author of a book, *Optical Burst Switched Networks*, Springer, 2005. He is a recipient of the Texas Telecommunication Engineering Consortium Fellowship 2002–03 and the University of Texas at Dallas Computer Science Dissertation of the Year Award 2003–04. He is a recipient of the Best Paper Award at the IEEE GLOBECOM 2005 Symposium on Optical Systems and Networks. He currently serves on the Editorial Board of *IEEE Communication Letters* and has served as a Guest Editor for the Elsevier *Optical Switching and Networking* journal. He has been on the technical program committees of several IEEE conferences including INFOCOM, ICC, GLOBECOM, ICCCN, HSN, and ANTS, and served as TPC Co-Chair for the Optical Networks and Systems (ONS) symposia at ICCCN 2007 and 2010, GLOBECOM 2011, ICC 2012, and INFOCOM High-Speed Networks (HSN) workshop 2011.