Comprehensive Performance Modeling and Analysis of Multicasting in Optical Networks

Srivatsan Sankaranarayanan and Suresh Subramaniam, *Member, IEEE*

***Abstract*—Multicasting is becoming increasingly important in today's networks. In optical networks, optical splitters facilitate the multicasting of optical signals. By eliminating the transmission of redundant traffic over certain links, multicasting can improve network performance. However, in a wavelength-division multiplexed optical network, the lack of wavelength conversion necessitates the establishment of a single multicast circuit (lighttree) on a single wavelength. On the other hand, establishing several unicast connections (lightpaths) to satisfy a multicast request, while requiring more capacity, is less constrained in terms of wavelength assignment. The objective of this paper is to evaluate the tradeoff between capacity and wavelength continuity in the context of optical multicasting. To this end, we develop accurate analytical models with moderate complexity for computing the blocking probability of multicast requests realized using lighttrees, lightpaths, and combinations of lighttrees and lightpaths. Numerical results indicate that a suitable combination of lighttrees and lightpaths performs the best when no wavelength conversion is present.**

***Index Terms*—Blocking probability, hybrid approach, lightpath, lighttree, multicasting, optical splitters, wavelength conversion, wavelength routing.**

I. Introduction

WAVELENGTH-DIVISION multiplexed (WDM) optical networks that use wavelength routing may be used to set up on-demand circuit-switched all-optical connections called as *lightpaths*. Lightpaths can be used as high-speed point-to-point links in a higher layer service network. The proliferation of services such as video distribution and conferencing requires future networks to be multicast capable. In a wavelength-routing optical network, multicasting is easily enabled by optical power splitters which split an optical signal into many "copies." Using these splitters, an optical signal can be delivered to multiple destinations over a tree (consisting of the set of links on which the signal is going to propagate) with the source node's transmitter as the root. However, splitting results in power loss and, in practice, there may be a limit to the number of times a signal may be split.

In this paper, we consider a WDM wavelength-routing network on which we wish to set up point-to-point (unicast) requests and point-to-multipoint (multicast) requests. Unicast requests can be set up using lightpaths. It is an interesting question to ask how multicast requests should be realized.

Manuscript received January 2, 2003; revised August 11, 2003. This work was supported in part by the National Science Foundation (NSF) under Grant ANI-9973111.

The authors are with the Department of Electrical and Computer Engineering, The George Washington University, Washington, DC 20052 USA (e-mail: vatsan@gwu.edu; suresh@gwu.edu).

Digital Object Identifier 10.1109/JSAC.2003.818229

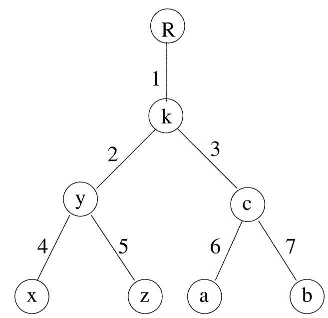

Fig. 1. A multicast tree.

They may be realized using an optical tree (*lighttree* [1][1]), or a set of lightpaths (one from the source to each destination), or a combination of lightpaths and lighttrees. In this paper, we define the three ways of realizing a multicast request as the lighttree approach, the lightpath approach, and the hybrid approach, respectively.

Let us illustrate the various tradeoffs involved in the three approaches with the help of an example multicast tree shown in Fig. 1. Suppose no wavelength conversion is available and a multicast request is to be set up over the given tree. For the lighttree approach, we need to find a single wavelength that is free on all the seven links of the multicast tree. For the lightpath approach, we need to find four different wavelengths, one free on links $\{1,2,4\}$, another free on links $\{1,2,5\}$, a third free on links $\{1,3,6\}$, and a fourth free on links $\{1,3,7\}$. The constraint due to wavelength continuity is the most severe in the lighttree approach while the capacity requirement is the least (seven wavelengths—one on each of the seven links of the tree). On the other hand, the lightpath approach is least constrained in terms of wavelength continuity because each of the four wavelengths needs to be free on only three links, but the capacity requirement is the highest (12 wavelengths—four on link 1, two each on links 2 and 3, and one each on links $\{4,5,6,7\}$. The hybrid approach is in between the lighttree and lightpath approaches in terms of the capacity and wavelength continuity constraints. One hybrid approach could be to split the multicast tree into two subtrees, one consisting of links $\{1,2,4,5\}$, and the other consisting of links $\{1,3,6,7\}$. Now, the capacity requirement is only eight but the wavelength continuity constraint is more pronounced than in the lightpath approach. Another hybrid approach could be to further split one of the subtrees (say the one consisting of links $\{1,3,6,7\}$) into two paths $\{1,3,6\}$ and $\{1,3,7\}$. In this case, we have one subtree and two paths which span the multicast tree, and

[1]A lighttree is an all-optical tree in which an optical signal from the root (or source) of the tree is split and delivered to every destination.

the capacity requirement is ten with somewhat more relaxed constraints on wavelength continuity. Several other hybrid realizations are also possible.

It is important to note that the hybrid approach does not include the lighttree and lightpath approaches. By "a hybrid approach," we mean that every multicast call (for a particular source and destination set) is set up using the *same* combination of lighttrees and lightpaths. Two hybrid approaches are different if the lighttree/lightpath combinations are different. This is analogous to using a fixed routing and wavelength assignment algorithm in unicast. For example, one hybrid approach for the multicast tree shown in Fig. 1 consists of two subtrees spanning the links $\{1, 2, 4, 5\}$ and $\{1, 3, 6, 7\}$. Another hybrid approach consists of one subtree $\{1, 2, 3, 4, 5\}$ and two subpaths $\{1, 6\}$ and $\{1, 7\}$.

A. Previous Work

Optical multicasting has attracted a fair amount of research attention recently. We now give an overview of the previous work done in optical multicasting. We classify the work into three categories, viz., multicast node architectures, multicast routing and wavelength assignment, and performance modeling. The work in each of these three areas is briefly reviewed below.

1) Node Architectures: Multicasting can be achieved without using any optical splitting by just establishing separate lightpaths from the source to every destination. But from the point of view of resource optimization, it is more economical to use optical splitters thereby eliminating redundant traffic on certain links. Such a multicast-capable (MC) node can have different degrees of splitting capabilities. The most basic MC node architecture is the splitter and delivery (SaD) [also called the tap and continue (TaC)] architecture [2]. Here, most of the energy in the incoming signal is passed to the output but a very small amount of it is tapped for local use. This can then be amplified and delivered to the local node. The advantage of doing this is that it is simple to implement and also does not result in significant power losses for the downstream nodes.

Three architectures for MC nodes are presented in [3]. These are called the multicast with same wavelength (MSW), multicast with the same destination wavelength (MSDW), and multicast with any wavelength (MAW). MSW corresponds to no wavelength conversion (WC), i.e., the incoming signal and outgoing signals must be on the same wavelength. In MSDW, all output signals must use the same wavelength but this may be different from the input wavelength. MAW corresponds to a splitting node with full WC, i.e., the incoming signal and the outgoing signals may all use different wavelengths. Power efficient designs of multicast cross-connects that provide unicast service without imposing power splitting loss on it are studied in [4].

In this paper, we assume the MSW node model for MC nodes without WC and the MAW node model for MC nodes with WC. For reader convenience, we present the MSW and MAW node architectures below by essentially reproducing them from [3]. The multicast node architecture for no WC is shown in Fig. 2(a) for one wavelength (wavelength 1). The input wavelengths are assumed to have been demultiplexed prior to the inputs shown in the figure, which carry only one wavelength. The splitter splits

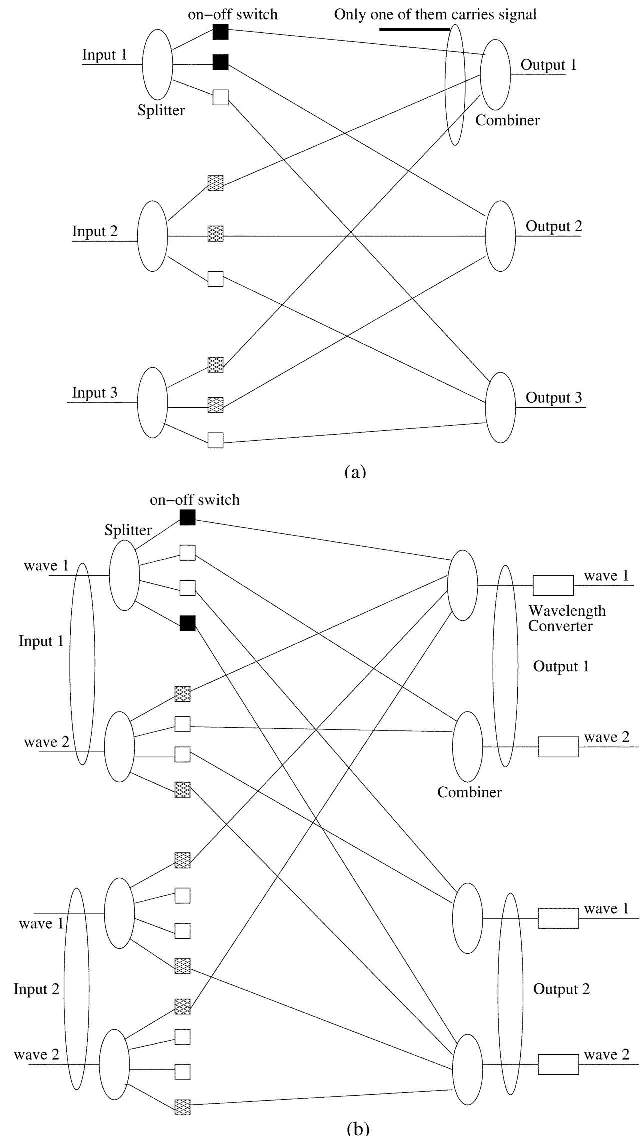

Fig. 2. Multicast node architecture. (a) No wavelength conversion. (b) Full wavelength conversion.

the input optical signal on the given wavelength into three copies with equal power and each of these is connected to one output combiner. The combiner essentially selects the active input and passes it on to its output. It is necessary that only one of the several combiner inputs are active at any given time. This is ensured by programming the on–off switches (implemented, for example, using semiconductor optical amplifiers) appropriately. To multicast the optical signal on input 1 and wavelength 1 to outputs 1 and 2 and the same wavelength, we need to activate those switches that are shown in dark in the figure and disable those that are shown by shaded crossed lines. This ensures that there is no contention at the combiners.

The multicast node architecture for full WC is shown in Fig. 2(a) when there are two wavelengths per fiber. For clarity, complete connections to the combiners are shown only for those two combiners that correspond to output1-wave1 and output2-wave2. The only additional component needed now is the full wavelength converter which can convert any input wavelength to a fixed output wavelength. Every output-wavelength combination needs such a converter. Suppose that the signal on input 1 and wavelength 1 needs to be multicast to outputs 1 and 2 on the wavelengths 1 and 2, respectively. This can be achieved by the activation pattern (dark shaded switches are on) for the on–off switches shown in Fig. 2(b).

2) Multicast Routing and Wavelength Assignment: Multicast routing and wavelength assignment (RWA) is the problem of choosing a tree and assigning wavelength(s) to it to satisfy a multicast request. Typically, constraints on the splitting degree of an MC node are assumed. Constrained multicast routing with sparse light splitting, where only some nodes in the network are capable of splitting a signal while others are not, is studied in [5]–[7]. Specifically, four WDM multicast routing algorithms are proposed and their relative performances are compared in [5]. Multicast RWA to minimize the number of wavelengths in a WDM network with splitter constraints is studied in [8]. The problem of multicast routing in circuit-switched multihop optical networks is studied in [9]. Dynamic traffic is considered in that paper. The problem is to find suitable wavelengths on the links of the tree on which to establish the call in such a way that the transmitters and receivers used by the call are not already used by existing circuits. It is shown that this problem is NP-complete and efficient heuristics are proposed and evaluated. In [2], an alternative to implementing multicasting without using power splitters is considered. The node architecture assumed here is the TaC architecture which avoids the power losses due to splitting. Heuristics to construct a trail to cover all the destinations in a multicast group are given and evaluated.

3) Performance Modeling: The concept of a lighttree for improved performance in optical networks is introduced in [1]. There has been some work in modeling the blocking performance of multicast requests when they are implemented using lighttrees. An analytical multicast blocking model for so-called homogeneous trunk switched networks (TSNs) is proposed in [10]. A path decomposition approach for evaluating the call blocking probability of multicast calls is proposed in [11]. Multicast blocking models for a fully connected network with constraints on the number of hops used for routing are presented in [12]. Blocking probability analysis in WDM switching networks with limited wavelength conversion is presented in [13]. Multicast communication in a class of multicast-capable WDM networks with regular topologies under some commonly used routing algorithms is addressed in [14]. The problem of multicast-capable node placement in wavelength-routed optical networks is addressed in [15]. Point-to-multipoint connection establishment with an upper bound on the number of wavelength converters is addressed in [16]. A point-to-multipoint generalization of a lightpath is given in [17]. The problem of minimizing the cost of establishing multicast calls under dynamic traffic is considered in [18] and [19]. The blocking performance of multicast traffic in wavelength routing networks under various fanout splitting policies is addressed in [20]. Bounds on wavelength requirements for multicasting in general topology WDM networks are presented in [21].

B. Contributions of This Paper

Our example earlier illustrated the tradeoff between capacity requirement and wavelength continuity for the three approaches for realizing a multicast request, viz., the lighttree, lightpath, and hybrid approaches. It is far from clear which of these approaches is the most suitable from a network performance point of view even if there are no limitations imposed by the physical layer in the form of splitting losses. Our objective in this paper is to develop a framework to compare these three approaches. To this end, we develop analytical models using a common framework for computing the probability that a multicast request is blocked in each of the three approaches, and studying the impact of wavelength conversion on their performance. The models will be shown to be quite accurate and their complexities are low enough to employ them to quickly obtain numerical results. This is the main contribution of this paper. There has been no effort until now to analyze the lightpath and hybrid approaches. To the best of our knowledge, the tradeoffs between wavelength continuity and capacity have not been recognized by previous researchers. As the discussion accompanying the example shows, a careful evaluation of the performance of all three approaches is required before any conclusions regarding optical multicasting can be made. We will see later that, in fact, the lighttree approach is not desirable in certain situations. The rest of the paper is organized as follows. We present our common framework for the analysis of the three approaches in the next section. The details of the analytical models are presented in the following three sections. Numerical results validating the models and comparing the three approaches are presented in Section VI, and finally, the paper is concluded in Section VII.

II. A Common Framework for Multicasting Analysis

In this section, we develop a common framework that is used by the three models presented in the next three sections. We start by describing the network model and the assumptions used in our analytical models.

- Each link of the network is assumed to have one fiber and each fiber has F wavelengths.
- Unicast call requests arrive according to a Poisson process and are set up using lightpaths on fixed routes.
- We consider a single multicast tree $\mathcal{T}$, and assume that multicast requests arrive according to a Poisson process for this tree. The arrival rate of multicast requests is assumed to be small relative to that of unicast requests. This is practical and reasonable because multicast requests are expected to arrive much more infrequently than unicast requests. A similar assumption was made in [11], wherein no more than a single multicast call was allowed at any time for a given multicast group.
- Random wavelength assignment is used for the unicast calls. This assumption is necessary for analytical tractability.
- Multicast calls may be set up using lighttrees, lightpaths, or a combination of them.
- For multicast calls set up as a lighttree, a wavelength is selected randomly from the set of free wavelengths on $\mathcal{T}$ (if there is no WC).
- Different wavelengths may be assigned for different branches at the root of the tree.

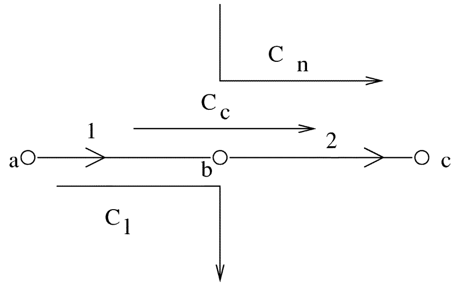

Fig. 3. A two-hop path.

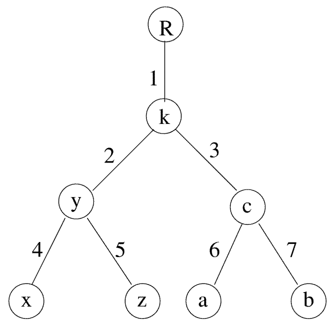

Fig. 4. A multicast tree.

- We do not use a reduced-load model [22], and assume that the effect of request blocking on the load offered to a link is negligible, as in [23].

For multicasting using lighttrees, only the MSW (corresponding to no WC) and MAW (corresponding to full WC) node models are considered. The network with the locations of the MSW and MAW nodes is assumed to be given. To illustrate our wavelength conversion model, consider the node k in Fig. 4 to be a converter. Then, the tree is partitioned into three *segments* as follows: one segment consisting of links 2, 4, and 5, another consisting of 3, 6, and 7, and the other consisting of only the link 1. There are no converter nodes inside a segment. Wavelength assignment can be done independently for these segments. Note that a segment can itself be a tree rather than a path.

The analytical models are based on the unicast blocking model in [23]. The following notation is from [23]. Consider the two-hop path shown in Fig. 3. A call that traverses both the links is defined as a *continuing call*, one that traverses only the first link is defined as a *leaving call*, and one that traverses only the second link is called an *entering call*. All the probabilities defined in this paper are steady-state probabilities. Let

- $\pi(c_l, c_c, c_n) = \Pr\{c_l$ leaving calls, c_c continuing calls and, c_n entering calls in the two-hop path$\}$;
- $Q(w) = \Pr\{w$ wavelengths free on a link$\}$;
- $S(y\,|\,x) = \Pr\{y$ wavelengths are free on the second link of the path $|\,x$ wavelengths are free on the first link of the path$\}$;
- $U(z\,|\,x, y) = \Pr\{z$ same wavelengths busy on both the links $|\,x$ and y wavelengths free on the first and the second link, respectively$\}$;
- $R(n\,|\,x, y, z) = \Pr\{n$ wavelengths free on both the links $|\,x$ and y wavelengths, respectively, are free on the first and the second link and z wavelengths are used by the same call on both the links$\}$.

The expressions for the above quantities from [23] are reproduced here for convenience. In [23], the joint occupancy distribution of the continuing, entering, and leaving calls on a two-hop path was obtained by modeling the (c_l, c_c, c_n)-tuple as a truncated Markov chain. By solving the Markov chain, it was shown in [23] that

$$\pi(a, b, c) = \frac{\frac{\rho_l^a}{a!}\frac{\rho_c^b}{b!}\frac{\rho_n^c}{c!}}{\sum_{z=0}^{F}\sum_{x=0}^{F-z}\sum_{y=0}^{F-z}\frac{\rho_l^x}{x!}\frac{\rho_c^z}{z!}\frac{\rho_n^y}{y!}}$$

and

$$Q(w) = \sum_{z=0}^{F-w}\sum_{y=0}^{F-z}\pi(F-w-z, z, y) \quad (1)$$

where ρ_l, ρ_c, and ρ_n are the Erlang loads of the leaving, continuing, and entering calls on the two-hop path. Also, from [23]

$$S(y\,|\,x) = \frac{\sum_{z=0}^{\min(F-x,F-y)}\pi(F-x-z, z, F-y-z)}{\sum_{z=0}^{F-x}\sum_{u=0}^{F-z}\pi(F-x-z, z, u)}$$

$$R(n\,|\,x, y, z) = \frac{\binom{x}{n}\binom{F-x-z}{y-n}}{\binom{F-z}{y}}$$

and

$$U(z\,|\,x, y) = \frac{\pi(F-x-z, z, F-y-z)}{\sum_{u=0}^{\min(F-x,F-y)}\pi(F-x-u, u, F-y-u)}.$$

Let R be the root of the multicast tree $\mathcal{T}$. R, as well as all the leaf nodes of $\mathcal{T}$ can be assumed to be WC nodes for the purpose of the model (even though they may not actually be). For simplicity, we assume that the leaf nodes of $\mathcal{T}$ are the only destination nodes. We assume $\mathcal{T}$ to not branch at R since, otherwise, it can be split into multiple different trees rooted at R for wavelength assignment purposes. We describe some new notation now. A node refers to a multicast tree node in the notation below.

- P_b is the probability that the multicast call over $\mathcal{T}$ is blocked. We assume that all the given destinations must be served for the multicast call to be successful.
- The pivot of node x, $\mathrm{pivot}(x)$ is that WC node closest to x on the reverse path from x to R on $\mathcal{T}$.
- $\mathcal{N}_x$ is the set of same wavelengths that are free on all the links of $\mathcal{T}$ which connect $\mathrm{pivot}(x)$ to x. Let $N_x = |\mathcal{N}_x|$.
- $L_x(w_x, n_x) = \Pr\{w_x$ wavelengths free on the link connecting x to its parent and $N_x = n_x\}$.This is defined for all nodes.
- $V_x(n_x\,|\,n_y) = \Pr\{N_x = n_x\,|\,N_y = n_y\}$, where y is the parent node of x. $V(\cdot\,|\,\cdot)$ is defined for every node x of $\mathcal{T}$ which is at least two hops away from its pivot node.
- f_x is a fringe node of x if it is the first WC node on a downstream path (on $\mathcal{T}$) from x. There may be several fringe nodes for a given node x. A WC node is a fringe node of itself. F_x is the set of all fringe nodes of x. When x is a WC node or a leaf node, $F_x = \{x\}$.
- A segment with root x is said to be a downstream segment of node y if either y is the same as x, or y lies on the reverse path from x to R. S_x is the collection of all the downstream segments of x.

Let us explain the above notation with the help of the tree shown in Fig. 4. Suppose that there are no converters. Then, the pivot of all the nodes is R. w_x is the number of wavelengths free on link 4 and n_x is the number of same wavelengths free on links 1, 2, and 4. $V_x(n_x \mid n_y)$ is the $\Pr\{n_x$ same wavelengths free on links 4, 2, and $1 \mid n_y$ same wavelengths are free on links 2 and 1$\}$.

But if y were a converter, $V_x(\cdot \mid \cdot)$ is not defined, and nodes x and z would be y's fringe nodes. Then, the segment consisting of link 4 is a downstream segment of the nodes R, k, and y. Another downstream segment of the above nodes would be link 5.

If k were the only WC node, then k is the pivot of the nodes y, x, z, a, b, and c. The pivot of k is R. Then, the two downstream segments of k consist of links 2, 4, and 5, and 3, 6, and 7, respectively. These two are also downstream segments of R. The downstream segments of the leaf nodes of $\mathcal{T}$ are nonexistent.

III. Lighttree Model

We now present the model for computing the blocking probability of a multicast request on $\mathcal{T}$ assuming that a lighttree is used to set up the request. We first give a high-level view of how our model works for no WC. Consider the multicast tree of Fig. 4. Initially, we start with the wavelength sets $\mathcal{N}_x, \mathcal{N}_z, \mathcal{N}_a$, and $\mathcal{N}_b$. We find the probability distribution of the number of same wavelengths belonging to both $\mathcal{N}_a$ and $\mathcal{N}_b$ (i.e., the intersection of $\mathcal{N}_a$ and $\mathcal{N}_b$). Let us call this $\mathcal{N}_{ab}$. Similarly, we find $\mathcal{N}_{xz}$. At this point, we find the probability distribution of the number of same wavelengths belonging to both $\mathcal{N}_{ab}$ and $\mathcal{N}_{xz}$ (called $\mathcal{N}_{xzab}$.) The success probability of establishing the multicast call is nothing but the probability that $|\mathcal{N}_{xzab}|$ is at least one.

We now generalize this procedure for arbitrary WC and present a formal description of the model. We start by defining some more notation that is used exclusively by this model.

- Θ_x is the random variable representing the number of same wavelengths that are free on the linear segment from $\mathrm{pivot}(x)$ to x, as well as on all the links of the subtree rooted at x and extending up to the fringe nodes of x. This quantity is defined for all nodes. For WC nodes, the subtree is empty.
- We define S^x as the event that there exists, for each segment in S_x, at least one wavelength free on all the links of the segment. We also call this as a *success event*. Recall that S_x is the collection of all the downstream segments of x.
- $T_x(\theta_x, S^x \mid n_x) = \Pr\{\Theta_x = \theta_x$, success event for all the downstream segments of $x \mid N_x = n_x\}$. This quantity is defined for all nodes.
- Let $1, 2, \ldots, k$ be the k children of x. Then, $M_{\text{int}}(n \mid n_x, n_1, n_2, \ldots, n_k) = \Pr\{n$ same wavelengths free over the linear segment from $\mathrm{pivot}(x)$ to x, as well as on the links between x and the nodes $1, 2, \ldots, k \mid N_x = n_x, N_1 = n_1, N_2 = n_2, \ldots, N_k = n_k$, and x not a converter$\}$. We can easily see that n is just the cardinality of $\mathcal{N}_1 \cap \cdots \cap \mathcal{N}_k$.

We once again use Fig. 4 to explain the notation. Assuming no WC nodes, θ_k is the number of same wavelengths free on all the links of $\mathcal{T}$. If k is a converter, θ_k is the number of wavelengths free on link 1 alone, as the subtree rooted at k is considered empty. Assuming no WC nodes, $T_k(\theta_k, S^k \mid n_k)$ is the $\Pr\{\theta_k$ wavelengths free on all the links of the tree (link 1 + the subtree rooted at $k) \mid N_k = n_k\}$. The event S^k is not of significance here because the set S_k is empty. If y is the only converter, then $T_k(\theta_k, S^k \mid n_k)$ is the $\Pr\{\theta_k$ wavelengths free on the links 1, 2, 3, 6, and 7 (link 1 + the subtree rooted at k limited by its fringe nodes) and the multicast call succeeds on the two segments consisting of the links 4 and 5, respectively $\mid N_k = n_k\}$.

We proceed by first finding $L(\cdot \mid \cdot)$ and $V(\cdot \mid \cdot)$ for the nodes in $\mathcal{T}$. The computation of $L(\cdot \mid \cdot)$ is recursive as shown below.

If the parent of x, say y, is a converter

$$L_x(w_x, n_x) = Q(n_x), \qquad \text{if } w_x = n_x \tag{2}$$

and is 0, otherwise.

If y is not a converter

$$L_x(w_x, n_x) = \sum_{n_y=n_x}^{F} \sum_{w_y=n_y}^{F} \sum_{w_z=0}^{\min(F-w_x, F-w_y)} L_y(w_y, n_y) \cdot S(w_x \mid w_y) \cdot U(w_z \mid w_x, w_y) \cdot R(n_x \mid w_z, w_x, n_y). \tag{3}$$

$V_x(\cdot \mid \cdot)$ is computed from $L_y(\cdot \mid \cdot)$ of its parent y as

$$V_x(n_x \mid n_y) = \sum_{w_x=n_x}^{F} \sum_{w_y=n_y}^{F} \sum_{w_z=0}^{\min(F-w_x, F-w_y)} L_y(w_y, n_y) \cdot S(w_x \mid w_y) \cdot U(w_z \mid w_x, w_y) \cdot R(n_x \mid w_z, w_x, n_y). \tag{4}$$

We now compute $M_{\text{int}}(\cdot \mid \cdot \cdots \cdot)$ by recursion. Essentially, we combine two $\mathcal{N}_i$'s into a single effective $\mathcal{N}_c$, which is then combined with other $\mathcal{N}_i$'s, shown in the equation at the bottom of the next page, where

$$M_{\text{int}}(n_c \mid n_x, n_1, n_2) = \frac{\binom{n_2}{n_c}\binom{n_x - n_2}{n_1 - n_c}}{\binom{n_x}{n_1}}$$

for $\max(0, n_1 + n_2 - F) \leq n_c \leq \min(n_1, n_2)$, and is 0, otherwise.

It is easy to verify the above expression for $M_{\text{int}}(n_c \mid n_x, n_1, n_2)$ with the help of the wavelength distribution diagram shown in Fig. 5. The windows of wavelengths $\mathcal{N}_1$ and $\mathcal{N}_2$ (of sizes n_1 and n_2, respectively) are randomly distributed within the window $\mathcal{N}_x$ (of size n_x). (Recall that random wavelength assignment is used.) The intersection window, of size n_c, is also randomly distributed

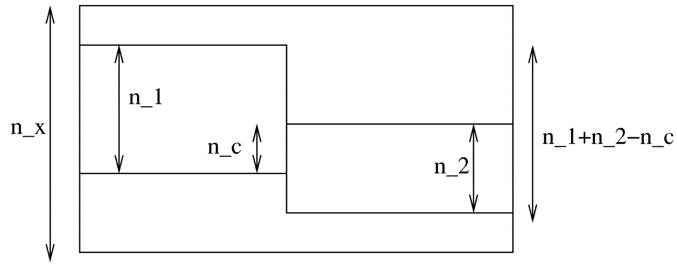

Fig. 5. Wavelength distribution diagram for computing $M_{\text{int}}(n \mid n_x, n_1, n_2)$.

within $\mathcal{N}_x$. The union window of size $n_1 + n_2 - n_c$ which is randomly distributed within $\mathcal{N}_x$ is also shown. This will be used in the lightpath model analysis.

Then

$$P_b = 1 - \sum_{n_x=1}^{F} \sum_{\theta_x=1}^{n_x} T_x(\theta_x, S^x \,|\, n_x) \cdot Q(n_x) \tag{5}$$

where x is the child of R. Recall that R has only one child.

$T_x(\cdot,\cdot\,|\,\cdot)$ is computed from the $T(\cdot,\cdot\,|\,\cdot)$ of its children in three different cases as follows.

Case 1) x is a converter but not a leaf node.

Then, $T_x(\theta_x, S^x \,|\, n_x) = 0$ if $\theta_x \neq n_x$. Otherwise,

$$T_x(\theta_x, S^x \,|\, n_x) = \prod_{i=1}^{k} \sum_{n_i=1}^{F} \sum_{\theta_i=1}^{n_i} \left\{ T_i(\theta_i, S^i \,|\, n_i) \cdot \sum_{w_x=1}^{F} \{ S(n_i \,|\, w_x) \cdot \Pr\{w_x \,|\, n_x\}\} \right\} \tag{6}$$

where $1, 2, \ldots, k$ are the children of x. $\Pr\{w_x \,|\, n_x\}$ can be easily obtained from $L(w_x, n_x)$.

Case 2) x is a leaf node.

Then, $T_x(\theta_x, S^x \,|\, n_x) = 1$ if $\theta_x = n_x \geq 1$, and is 0, otherwise.

Case 3) x is not a converter.

$$T_x(\cdot,\cdot\,|\,\cdot) = \sum_{\theta_1=1}^{n_x} \sum_{\theta_2=1}^{n_x} \cdots \sum_{\theta_k=1}^{n_x} \left\{ \prod_{i=1}^{k} \sum_{n_i=1}^{n_x} \{ T_i(\theta_i, S^i \,|\, n_i) \cdot V_i(n_i \,|\, n_x)\} \right\} \cdot M_{\text{int}}(\theta_x \,|\, n_x, \theta_1, \ldots, \theta_k) \tag{7}$$

where $1, 2, \ldots, k$ are the children of x. This completes the lighttree model description.

A. *Pseudocode for Computing Lighttree Model P_b*

We now briefly outline the steps involved in computing P_b for our example tree of Fig. 1 for the lighttree model. We assume no conversion and, therefore, ignore S^x in $T_x(\cdot,\cdot\,|\,\cdot)$.

Step 1) Compute $Q(\cdot)$, $L_x(\cdot,\cdot)$, and $V_x(\cdot\,|\,\cdot)$ for all nodes except R and k, using (1), (2), (3), and (4).

Step 2) Compute $T_c(\theta_c \,|\, n_c)$ from $V_a(n_a \,|\, n_c)$ and $V_b(n_b \,|\, n_c)$ using (7). Note that $T_a(\cdot\,|\,\cdot)$ and $T_b(\cdot\,|\,\cdot)$ are trivial.

Step 3) Similar to the above step, compute $T_y(\theta_y \,|\, n_y)$ from $V_x(n_x \,|\, n_y)$ and $V_z(n_z \,|\, n_y)$ using (7).

Step 4) Compute $T_k(\theta_k \,|\, n_k)$ from $T_y(\theta_y \,|\, n_y)$, $T_c(\theta_c \,|\, n_c)$, $V_y(n_y \,|\, n_k)$, and $V_c(n_c \,|\, n_k)$ using (7).

Step 5) Finally, compute P_b from $T_k(\theta_k \,|\, n_k)$ and $Q(n_k)$ using (5).

IV. Lightpath Model

We now present the model for the lightpath approach. The multicast call succeeds if all the lightpaths on the multicast tree can be set up. Let us illustrate the condition under which the call can succeed. First, consider the case of no wavelength conversion. Let the destination nodes be numbered $1, 2, \ldots, k$. Then, we have to find one wavelength from each of the set of wavelengths $\mathcal{N}_i, 1 \leq i \leq k$, and these must be distinct (since they all share the link from the root of the tree).[2] It is not difficult to see that wavelength assignment to lightpaths corresponds to finding a matching in a bipartite graph. Consider the bipartite graph $G = (U, V, E)$, where $u \in U$ corresponds to a wavelength set $\mathcal{N}_i$, and V is simply the set of wavelengths, $\{1, 2, \ldots, F\}$. There is an edge between $u \in U$ and $v \in V$ iff the wavelength $v \in u$. The call can be established iff a matching of size k can be found in this bipartite graph.

A necessary and sufficient condition for a k-cardinality matching to exist in a bipartite graph is given by Hall's theorem [24]: if $U_0 \subseteq U$ such that $|U_0| \leq k$, and the range set $R(U_0) \subseteq V$ is the set of nodes adjacent to some node in U_0, then $|U_0| \leq |R(U_0)|$. A condition on each subset is referred to as a Hall's condition. In our graph, this implies that the union of any $m \leq k$ sets of wavelengths $\mathcal{N}_i$ must have cardinality at least m.

To compute the blocking probability of the call, we need to find the joint probability that all the 2^k Hall's conditions are satisfied. This observation indicates the difficulty of analytically computing the blocking probability of the call if it is set up using lightpaths. In our model, we approximate the success probability of establishing the call by the joint probability that

[2]This is an indication of why a simple model which considers the multicast call to be successful if each of the lightpaths is *independently* successful will not work.

$$\begin{aligned}
M_{\text{int}}(u_x \,|\, n_x, n_1, \ldots, n_k) = \sum_{n_{1(k-1)}=1}^{F} \Bigg\{ \sum_{n_{1(k-2)}=1}^{F} \cdots \Bigg\{ \sum_{n_{13}=1}^{F} \Bigg\{ \sum_{n_{12}=1}^{F} M_{\text{int}}(n_{12} \,|\, n_x, n_1, n_2) \\
\cdot M_{\text{int}}(n_{13} \,|\, n_x, n_{12}, n_3) \Bigg\} \cdot M_{\text{int}}(n_{14} \,|\, n_x, n_{13}, n_4) \Bigg\} \cdots \\
\cdot M_{\text{int}}(u_x \,|\, n_x, n_{1(k-1)}, n_k) \Bigg\}
\end{aligned}$$

a certain subset of Hall's conditions are satisfied. We will see later that this approximation is quite accurate.

Informally, our model works as follows. Consider the multicast tree of Fig. 4. We start with the wavelength sets $\mathcal{N}_x, \mathcal{N}_z, \mathcal{N}_a$, and $\mathcal{N}_b$ as with the lighttree approach. We now find the probability distribution of the number of wavelengths that belong to either $\mathcal{N}_a$ or $\mathcal{N}_b$ (i.e., the union of $\mathcal{N}_a$ and $\mathcal{N}_b$). Let us call this $\mathcal{N}_{ab}$. We also impose the constraint that $|\mathcal{N}_a|$ and $|\mathcal{N}_b|$ must be at least one. Similarly, we find $\mathcal{N}_{xz}$ and impose the constraint that $|\mathcal{N}_x|$ and $|\mathcal{N}_z|$ must be at least one. At this point, we find the probability distribution of the number of wavelengths that belong to either $\mathcal{N}_{ab}$ or $\mathcal{N}_{xz}$ (called $\mathcal{N}_{xzab}$) and impose the constraint that $|\mathcal{N}_{ab}|$ and $|\mathcal{N}_{xz}|$ should be at least two each. The success probability of establishing the multicast call now is nothing but the probability that $|\mathcal{N}_{xzab}|$ obtained thus is at least four. What this procedure does in effect is to find the joint probability that all the constraints mentioned above are satisfied. There are a total of seven constraints and these are the seven Hall's conditions that our model accounts for in this particular multicast tree.

We now formally describe the model by giving some more notation first.

- A downstream node of x is one that has x on its reverse path to R.
- $\min_x =$ the number of leaf nodes of the subtree of $\mathcal{T}$ rooted at x. This is the minimum value of N_x that is required for the call to possibly succeed, as every leaf node of $\mathcal{T}$ that has x on its reverse path would require a wavelength on all the links that connects it to R. If x is a leaf node, then $\min_x = 1$.
- Let $\mathcal{U}_x = \cup_{i=1}^{c} \mathcal{N}_i$, where $1, 2, \ldots, c$ are the nodes in F_x, and let $U_x = |\mathcal{U}_x|$.
- Let $1, 2, \ldots, k$ be the children of x. Then, for $1 \leq j \leq k, h_{x,j}$ denotes the Hall's condition that $|\cup_{i=1}^{j} \mathcal{U}_i| \geq \sum_{i=1}^{j} \min_i$. Let h_x be the set of all Hall's conditions $h_{x,j}, 1 \leq j \leq k$, (denoted by $\cup_{j=1}^{k} h_{x,j}$). If x is a converter, $h_{x,j}$'s are undefined and h_x is the condition that $U_x \geq \min_x$.
- Let H_x be the set of Hall's conditions $h_x \cup h_1 \cup h_2 \ldots \cup h_m$, where $1, 2, \ldots, m$ are the downstream nodes of x that belong to the same segment as x. All the nodes other than the root of a segment are said to belong to that segment. If x is a converter or a leaf node, then H_x consists only of h_x.
- H^x is the event that all the Hall's conditions belonging to H_x are satisfied.
- We define S^x as the event that, for each segment in S_x, all the Hall's conditions that are taken care of by our lightpath model for that segment, are satisfied. We also call this as a *success event*. Recall that S_x is the collection of all the downstream segments of x.
- $B_x(u_x, H^x, S^x \,|\, n_x)$ is the $\Pr\{U_x = u_x$, all conditions in H_x are true, and the call succeeds on all the segments downstream of $x \,|\, N_x = n_x\}$. This quantity is defined for all nodes.
- $C_x(S^x \,|\, n_x)$, where x is a nonleaf converter node, is the $\Pr\{$the call succeeds in all the downstream segments of $x \,|\, N_x = n_x\}$. This quantity is defined for all converting nodes.
- $C_{x,i}(S_i^x \,|\, n_x)$ is the $\Pr\{$the call succeeds in the segment containing the ith child of x (as a nonroot node), as well as on all the downstream segments of $i \,|\, N_x = n_x\}$.
- $M_{\text{uni}}(u \,|\, n_x, n_1, n_2, \ldots, n_k) = \Pr\{\,|\mathcal{N}_1 \cup \mathcal{N}_2 \cup \ldots \cup \mathcal{N}_k| = u$, all conditions $h_{x,i}$ are satisfied for $2 \leq i \leq k-1 \,|\, N_x = n_x, N_1 = n_1, N_2 = n_2, \ldots, N_k = n_k$, and x not a converter$\}$, where $1, 2, \ldots, k$ are the children of x.

We explain the notation with the help of Fig. 4. Assume no WC nodes (except the root and leaf nodes) unless otherwise specified. The downstream nodes of k are y, c, x, z, a, and b. The downstream nodes of y are x and c. There are no downstream nodes for x, z, a, and b. $\min_k = 4, \min_y = \min_c = 2$, and $\min_x = \min_z = \min_a = \min_b = 1$. These numbers do not depend on whether the nodes of the tree are WC or not. $h_{k,c}$ is the Hall's condition that $|\mathcal{N}_x \cup \mathcal{N}_z \cup \mathcal{N}_a \cup \mathcal{N}_b| \geq \min_y + \min_c = 4$. $h_{k,y}$ is the Hall's condition that $|\mathcal{N}_x \cup \mathcal{N}_z| \geq \min_y = 2$. Note that $h_{k,y}$ is the same condition as $h_{y,z}$.

If k is a converter, h_k is the Hall's condition that $n_k \geq 1$. This is because u_k is the same as n_k. If y and c were the only converters, then $H_k = h_k \cup h_y \cup h_c, H_y = h_y$, and $H_c = h_c$. Since x, z, a, and b are leaf nodes which are assumed to be converters, $H_p = h_p$, if p is one of x, z, a, and b. S_k is empty. Thus, $B_k(u_k, H^k, S^k \,|\, n_k)$ is the probability that $U_k = u_k$, all the Hall's conditions belonging to H_k are true, and the call succeeds in all the downstream segments of k given n_k wavelengths are free on link 1.

If k is the only converter, then $C_k(S^k \,|\, n_x)$ is the probability that the call succeeds on the two downstream segments of k consisting of the links $\{2, 4, 5\}$ and $\{3, 6, 7\}$, respectively, given that n_k wavelengths are free on link 1. $C_{k,1}(S_1^k \,|\, n_k)$ is the probability that the call succeeds on the segment consisting of the links $\{2, 4, 5\}$ given $N_k = n_k$ and $C_{k,2}(S_2^k \,|\, n_k)$ is the probability that the call succeeds on the segment consisting of the links $\{3, 6, 7\}$ given $N_k = n_k$.

$M_{\text{uni}}(u_x \,|\, n_x, n_1, \ldots, n_k)$ is computed as shown in the equation at the bottom of the next page, where $M_{\text{uni}}(u \,|\, n_x, n_1, n_2) = M_{\text{uni}}(n_1 + n_2 - u \,|\, n_x, n_1, n_2)$ for $\max(n_1, n_2) \leq u \leq \min(n_x, n_1 + n_2)$, and is 0, otherwise, and $\min_{1i} \stackrel{\text{def}}{=} \sum_{j=1}^{i} \min_j$.

The Hall's conditions that are accounted for by our model in a given segment of $\mathcal{T}$ are the ones belonging to H_x, where x is the child of the root of the segment. Note that the root of every segment has only one child because of our wavelength conversion model. The number of Hall's conditions accounted for by our model is linear in the size of the tree. For a binary tree with k leaf nodes, the model takes care of only $2k-1$ conditions while there are 2^k conditions in total.

$L(\cdot, \cdot)$ and $V(\cdot \,|\, \cdot)$ for the nodes are computed as in the lighttree model. Let x be the child of R. Then, $P_b = 1 - \Pr\{H^x, S^x\}$. Note that H_x is the same as H_R since x is the only child of R. $\Pr\{H^x, S^x\}$ is given by

$$\Pr\{H^x, S^x\} = \sum_{n_x=\min_x}^{F} \sum_{u_x=\min_x}^{n_x} B_x(u_x, H^x, S^x \,|\, n_x) \cdot Q(n_x). \tag{8}$$

$B_x(\cdot, \cdot, \cdot \,|\, \cdot)$ is computed in three different cases as follows.

Case 1) x is not a converter. Let its children be $1, 2, \ldots, k$. Then, $B_x(u_x, H^x, S^x \mid N_x)$ can be expressed in terms of the B_j's, $j = 1, 2, \ldots, k$, as

$$B_x(u_x, H^x, S^x \mid n_x) = \sum_{u_1=\min_1}^{n_x} \sum_{u_2=\min_2}^{n_x} \cdots \sum_{u_k=\min_k}^{n_x} \left\{ \prod_{j=1}^{k} \left\{ \sum_{n_j=u_j}^{n_x} B_j(u_j, H^j, S^j \mid n_j) \;\; V(n_j \mid n_x) \right\} \right\} \cdot M_{\text{uni}}(u_x \mid n_x, u_1, \ldots, u_k) \tag{9}$$

if $u_x \geq \min_x$, and is equal to 0, otherwise.

Case 2) x is a leaf node. Then, $B_x(u_x, H_x, S^x \mid n_x) = 1$ if $u_x = n_x \geq \min_x$, else is 0.

Case 3) x is a converter but not a leaf node. We observe that, $B_x(u_x, H_x, S^x \mid n_x) = 0$ if $u_x \neq n_x$ or $u_x = n_x < \min_x$, and $B_x(u_x, H_x, S^x \mid n_x) = C_x(S^x \mid n_x)$, otherwise.

$C_x(S^x \mid n_x) = \prod_{i=1}^{k} C_{x,i}(S_i^x \mid n_x)$, where x has i children, where we have assumed statistical independence between the segments rooted at x. Let y be the ith child of x. Then, $C_{x,i}(S_i^x \mid n_x) = \sum_{w_x=n_x}^{F} \sum_{n_y=\min_y}^{F} \Pr\{w_x \mid n_x\} \cdot S(n_y \mid w_x) \cdot B_y(H_y, S^i \mid n_y)$.

This completes the description of the lightpath model.

A. Pseudocode for Computing Lightpath Model P_b

We now briefly outline the steps involved in computing P_b for our example tree of Fig. 1 for the lightpath model. Again, we assume no conversion, and therefore ignore S^x in $B_x(\cdot, \cdot, \cdot \mid \cdot)$.

Step 1) Compute $Q(\cdot), L_x(\cdot, \cdot)$ and $V_x(\cdot \mid \cdot)$ for all nodes except R and k using (1), (2), (3), and (4).

Step 2) Compute $B_c(u_c, H^c \mid n_c)$ from $V_a(n_a \mid n_c)$ and $V_b(n_b \mid n_c)$ using (9). Note that $B_a(\cdot, \cdot \mid \cdot)$ and $B_b(\cdot, \cdot \mid \cdot)$ are trivial. $B_c(u_c, H^c \mid n_c)$ is the $\Pr\{U_c = u_c, N_a \geq 1$, and $N_b \geq 1 \mid N_c = n_c\}$.

Step 3) Similar to the above step, compute $B_y(u_y, H^y \mid n_y)$ from $V_x(n_x \mid n_y)$ and $V_z(n_z \mid n_y)$ using (9). $B_y(u_y, H^y \mid n_y)$ is the $\Pr\{U_y = u_y, N_x \geq 1$, and $N_z \geq 1 \mid N_y = n_y\}$.

Step 4) Compute $B_k(u_k, H^k \mid n_k)$ from $B_y(u_y \mid n_y)$, $B_c(u_c \mid n_c), V_y(n_y \mid n_k)$ and $V_c(n_c \mid n_k)$ using (9). $B_k(u_k, H^k \mid n_k)$ is the $\Pr\{U_k = u_k, N_x \geq 1, N_z \geq 1, N_a \geq 1, N_b \geq 1, U_y \geq 2$, and $U_c \geq 2 \mid N_k = n_k\}$.

Step 5) Finally, compute P_b from $B_k(u_k, H^k \mid n_k)$, and $Q(n_k)$ using (8). Note that during this computation, an additional condition that $U_k \geq 4$ is imposed.

V. Hybrid Model

We now present the hybrid model, which is actually a combination of the lightpath and lighttree models. We note that the extension to the hybrid approach provided herein applies to only some of the several possible combinations of lighttrees and lightpaths for satisfying a given multicast call. The sets of destination nodes served by these lighttrees and lightpaths are assumed to not overlap. First, we specify the conditions under which the extension is possible. We are given the following information. The hybrid implementation consists of a set of lighttrees and lightpaths, each having a given demand, i.e., the number of wavelengths, that are free on all the links of the given tree/path in order to successfully establish the multicast call. A demand of more than one is needed if wavelength conversion is present. Thus, we see that wavelength assignment for a hybrid implementation boils down again to the bipartite matching problem (after we instantiate as many copies of each of these tree/paths as its demand). These trees and paths can have a very complicated link intersection (or overlap) pattern and, as such, the general problem (corresponding to an arbitrary lighttree/lightpath combination) is very difficult to analyze. Nevertheless, we have found that accurate modeling of the situation is possible for a subset of all the possible hybrid implementations for the given multicast call. We call each of the trees and paths as a subtree and a subpath, respectively, and also refer to both of them by the common terminology "subcall." Let the subcalls be denoted by $\mathcal{C}_i$ for $i = 1, \ldots, n$. Each of these is either a tree or a path. Now, accurate modeling is possible if the following two conditions are satisfied.

Condition 1: Consider any two child links of a nonconverting node a of $\mathcal{T}$ that are part of the same subtree $\mathcal{T}_*$ (within a given segment). Let the corresponding child nodes be x and y. If x and y are nonconverters, then the subtrees of $\mathcal{T}$ rooted at the nodes x and y are also part of $\mathcal{T}_*$. If either of them (say x) is a converter, then the condition is still applicable to the subtree at y, which is a nonconverter. The subtrees rooted at leaf nodes are taken to be *NULL*.

Condition 2: Consider any nonconverting node a on the multicast tree. Let us form n sets of links, $G_1, G_2, \ldots, G_n$ out of

$$M_{\text{uni}}(u_x \mid n_x, n_1, \ldots, n_k) = \sum_{n_{1(k-1)}=\min_{1(k-1)}}^{F} \left\{ \sum_{n_{1(k-2)}=\min_{1(k-2)}}^{F} \cdots \left\{ \sum_{n_{13}=\min_{13}}^{F} \left\{ \sum_{n_{12}=\min_{12}}^{F} M_{\text{uni}}(n_{12} \mid n_x, n_1, n_2) \cdot M_{\text{uni}}(n_{13} \mid n_x, n_{12}, n_3) \right\} \cdot M_{\text{uni}}(n_{14} \mid n_x, n_{13}, n_4) \right\} \ldots M_{\text{uni}}(u_x \mid n_x, n_{1(k-1)}, n_k) \right\}$$

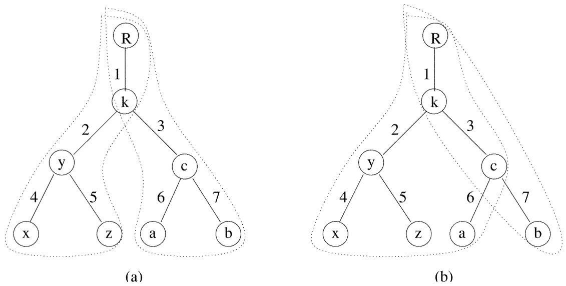

Fig. 6. (a) An analyzable hybrid scenario irrespective of wavelength conversion. (b) A nonanalyzable hybrid scenario under no WC (analyzable if k is a converter).

the child links of a, with each group containing those child links that belong to a given subcall, where n is the number of subcalls in the given hybrid realization. Some of the above sets can be *NULL*. Any two of these n sets of links should either overlap completely or not overlap at all.

We speak of the above conditions being satisfied or not in the context of the node a.

Note that these conditions do not come into effect when there is full WC. Thus any hybrid implementation can be analyzed by our model in the case of full WC.

We illustrate the above conditions with three hybrid realization examples. Let us assume no wavelength conversion. Consider the following two hybrid realizations of a multicast call. One consists of the set of subcalls $\{(1,2,4,5) \text{ and } (1,3,6,7)\}$ shown in Fig. 6(a) and the other consists of the set $\{(1,2,4,5,3,6) \text{ and } (1,3,7)\}$ shown in Fig. 6(b). The former combination (of two lighttrees) satisfies both the conditions but the latter does not satisfy either condition. The latter does not satisfy condition 1 at k because k has two child links 2 and 3 that belong to the subcall (1, 2, 4, 5, 3, 6) but the subtree at c (corresponding to child link 3) has a link (link 7) that does not belong to the subcall. It does not satisfy condition 2 at k because we can form two groups G_1 and G_2 out of the child links of k, where $G_1 = \{2,3\}$ and $G_2 = \{3\}$. G_1 corresponds to the first subcall and G_2 corresponds to the second subcall. They overlap (i.e., have a common link) but not completely (i.e., they are not identical). If k is a converter, then the latter combination is still analyzable by our hybrid model as neither condition is applicable at k.

Now, consider another hybrid realization consisting of two subtrees, one spanning the links (1, 2, 4, 3, 7) and the other spanning the links (1, 2, 3, 5, 6). Under no WC, condition 1 is clearly not satisfied at k. When y and c are converters, condition 1 at k is not concerned about these nodes. Condition 2 is satisfied at k, since both the subcalls span the links 2 and 3, the two groups of child links of k are identical ($=\{2,3\}$).

We now describe the model briefly. The notation is almost the same as that for the lightpath model, the main difference being in how the variable $\mathcal{U}_x$ is defined. We define $\mathcal{U}_x$ recursively here. Consider an arbitrary nonroot node of $\mathcal{T}$, say x. Let the child links of x be organized into several groups each containing only those links that are part of a subtree/subpath corresponding to a subcall. We note that the same child links of x belonging to a given group can be part of two different subcalls. In this case, we represent both the groups by a single effective group. Let the groups be $G_1,\dots,G_g$ and let D_i for $i=1,\dots,g$ be the number of subcalls that span the links of group G_i.

- Θ_x^i for $i=1,\dots,g$ is the random variable representing the intersection (or number of same wavelengths that belong to all) of the sets $\mathcal{U}_j$ for all nodes j that connect the links of group G_i to x. This definition is for a nonconverting x. If x is a converter or a leaf node, Θ_x^i is the same as $\mathcal{N}_x$. The variables Θ_x^i play the same role as $\mathcal{N}_i$ in the lightpath model for defining $\mathcal{U}_x$.
- Let $\mathcal{U}_x = \cup_{i=1}^{g} {\Theta_x}^i$ and let $U_x = |\mathcal{U}_x|$. For leaf nodes and converting nodes, $\mathcal{U}_x$ is the same as $\mathcal{N}_x$.
- $\min_x =$ the number of subcalls that span the linear segment from $\mathrm{pivot}(x)$ to x.
- For $1 \le j \le g, h_{x,j}$ denotes the Hall's condition that $|\cup_{i=1}^{j} \Theta_i| \ge \sum_{i=1}^{j} D_i$. Let h_x be the set of all Hall's conditions $h_{x,j}, 1 \le j \le g$, (denoted by $\cup_{j=1}^{g} h_{x,j}$). If x is a converter, $h_{x,j}$'s are undefined and h_x is the condition that $U_x \ge \min_x$.
- $H_x, H^x, S^x, B_x(u_x, H^x, S^x \,|\, n_x), \quad C_x(S^x \,|\, n_x), \quad C_{x,i}$ $(S_i^x \,|\, n_x)$ are the same as for the lightpath model.
- Let $1,2,\dots,k$ be the children of x. Let $G_1, G_2, \dots, G_g$ be the different groups of the child links of x as explained before. Let $M_{\text{hyb}}(u \,|\, n_x, n_1, n_2, \dots, n_k) = \Pr\{\, |\theta_x^1 \cup \theta_x^2 \cup \dots \cup \theta_x^k| = u$, all conditions $h_{x,i}$ are satisfied for $G_2 \le i \le G_{(g-1)} \,|\, N_x = n_x, N_1 = n_1, N_2 = n_2, \dots, N_k = n_k$, and x not a converter$\}$.

P_b and $\Pr\{H^x, S^x\}$ are the same as for the lightpath model. The computation of $B_x(\cdot,\cdot,\cdot\,|\,\cdot)$ is also the same except for Case 1 which is given below.

Case 1: x is not a converter. Let its children be $1,2,\dots,k$. Then, $B_x(u_x, H^x, S^x \,|\, N_x)$ can be expressed in terms of the B_j's, $j = 1,2,\dots,k$, as

$$B_x(u_x, H^x, S^x \,|\, n_x) = \sum_{u_1=\min_1}^{n_x} \sum_{u_2=\min_2}^{n_x} \cdots \sum_{u_k=\min_k}^{n_x} \left\{ \prod_{j=1}^{k} \left\{ \sum_{n_j=u_j}^{n_x} B_j(u_j, H^j, S^j \,|\, n_j) \cdot V(n_j \,|\, n_x) \right\} \right\} \cdot M_{\text{hyb}}(u_x \,|\, n_x, u_1, \dots, u_k) \tag{10}$$

if $u_x \ge \min_x$, otherwise $=0$.

We now compute $M_{\text{hyb}}(u_x \,|\, n_x, u_1, \dots, u_k)$. We divide the child links of x into g groups as explained before. Let the links that belong to group 1 be $1,\dots,l_1$, those that belong to group 2 be $l_1+1,\dots,l_2$ and so on and those that belong to group g be $l_{g-1}+1,\dots,k$.

To find $M_{\text{hyb}}(u_x \,|\, \cdot, \dots)$, we first find the intersection of the u_i's for various i's that belong to a single group (which gives the corresponding θ variable) and then take the union of each of these. The intersection of a given set of u_i's is obtained along the same lines as $M_{\text{int}}(u_x \,|\, n_x, \dots)$ in the lighttree model. Then, $M_{\text{hyb}}(u_x \,|\, \cdot, \dots)$ is obtained along the same lines as $M_{\text{uni}}(u_x \,|\, \cdot, \dots)$ was obtained in the lightpath model. For the latter computation, we use $\min_{1j} = \sum_{i=1}^{j} D_i$, for $j = 1,\dots,g$.

A. Pseudocode for Computing Hybrid Model P_b

We now briefly outline the steps involved in computing P_b for our example tree of Fig. 1 for the hybrid model. We assume that the hybrid approach consists of two subtrees, one consisting of the links 1, 2, 4, and 5 and the other consisting of the links 1, 3, 6, and 7. Again, we assume no conversion, and therefore ignore S^x in $B_x(\cdot,\cdot,\cdot\,|\,\cdot)$.

Step 1) Compute $Q(\cdot), L_x(\cdot,\cdot)$ and $V_x(\cdot\,|\,\cdot)$ for all nodes except R and k using (1), (2), (3), and (4).

Step 2) Compute $B_c(u_c, H^c\,|\,n_c)$ from $V_a(n_a\,|\,n_c)$ and $V_b(n_b\,|\,n_c)$ using (10). Note that $B_a(\cdot,\cdot\,|\,\cdot)$ and $B_b(\cdot,\cdot\,|\,\cdot)$ are trivial. $B_c(u_c, H^c\,|\,n_c)$ is the $\Pr\{\mathcal{N}_a$ and $\mathcal{N}_b$ have u_c common elements, $N_a \geq 1$, and $N_b \geq 1\,|\,N_c = n_c\}$.

Step 3) Similar to the above step, compute $B_y(u_y, H^y\,|\,n_y)$ from $V_x(n_x\,|\,n_y)$ and $V_z(n_z\,|\,n_y)$ using (10). $B_y(u_y, H^y\,|\,n_y)$ is the $\Pr\{\mathcal{N}_x$ and $\mathcal{N}_z$ have u_y common elements, $N_x \geq 1$ and $N_z \geq 1\,|\,N_y = n_y\}$.

Step 4) Compute $B_k(u_k, H^k\,|\,n_k)$ from $B_y(u_y, H^y\,|\,n_y)$, $B_c(u_c, H^c\,|\,n_c), V_y(n_y\,|\,n_k)$, and $V_c(n_c\,|\,n_k)$ using (10). $B_k(u_k, H^k\,|\,n_k)$ is the $\Pr\{U_y \cup U_c = u_k, N_x \geq 1, N_z \geq 1, N_a \geq 1, N_b \geq 1, U_y \geq 1$, and $U_c \geq 1\,|\,N_k = n_k\}$.

Step 5) Finally, compute P_b from $B_k(u_k, H^k\,|\,n_k)$ and $Q(n_k)$ using (8). Note that during this computation, an additional condition that $U_k \geq 2$ is imposed.

B. Model Complexity

For the lighttree model, the computation of the quantity $L_x(\cdot,\cdot)$ takes $O(F^5)$ time and so does the computation of the quantity $V_x(\cdot\,|\,\cdot)$. The computation of $M_{\rm int}(\cdot\,|\,\cdot,\cdot,\cdot)$ takes $O(F^4)$ time. The computation of $R(\cdot\,|\,\cdot,\cdot,\cdot)$ takes $O(F^4)$ time. All other quantities, such as $U(\cdot\,|\,\cdot,\cdot)$, take only $O(F^3)$ time or less. Hence, the computational complexity of the model is $O(F^5)$. The other two models can also be shown to have the same complexity.

VI. Performance Evaluation

Performance evaluation results are presented in this section. We first validate the three models using simulation results and then proceed to the evaluation of the three approaches.

A. Simulation Details

We start with a brief description of the simulation set up. We consider the NSFNET topology shown in Fig. 7. Unicast call requests arrive for each of the 182 ordered pairs of nodes according to a Poisson process, and the load for every node pair is the same. Fixed shortest-path routing is assumed for the unicast calls. There is a single multicast tree consisting of seven links whose links are shown in bold in Fig. 7. The source node of the multicast tree is 2. The multicast call load is set to $1/30$ of the unicast call load for a single node pair. In assigning wavelengths for the multicast calls in the lightpath approach, we do not use truly random assignment since this would require the generation of all the matchings as described in Section IV. Instead, we use an arbitrary assignment that is obtained using a piece of public-domain code for generating a matching in a bipartite graph. Generation of matchings is also needed for the hybrid approach. For the lightpath approach, we have several sets of wavelengths each corresponding to the same wavelengths that are free on all the links constituting a single subpath (i.e., a unicast path from the root to a destination node) and we need to find a full cardinality matching on these sets of wavelengths. For the hybrid approaches, these individual sets need not correspond to unicast paths but can be trees themselves. But for wavelength assignment, a full cardinality matching has to be found on these sets as in the case of the lightpath approach.

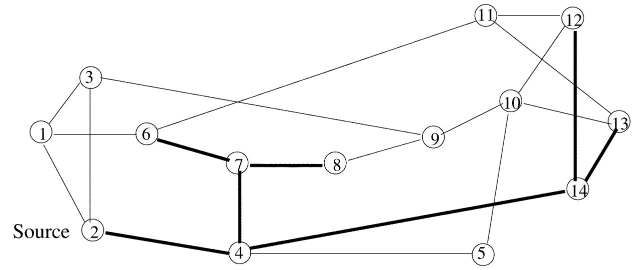

Fig. 7. NSFNET topology. The multicast tree links are in bold. Node 2 is the root.

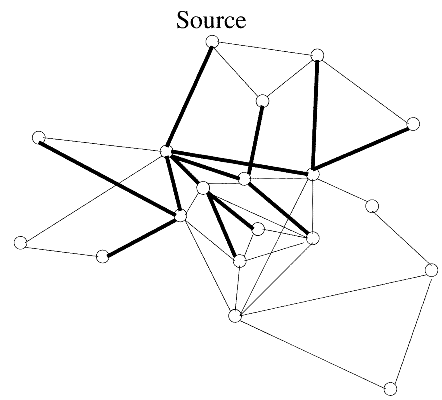

Fig. 8. EON topology. The multicast tree links are in bold. The source node is indicated in the figure.

We have also generated results for a 13-link tree in the European optical network topology (EON) shown in Fig. 8. These results are available in [25] and are not produced here due to space constraints, and because they do not lead to significantly different observations. We found that the models predict the blocking of the various approaches accurately. This confirms the validity (and the computational tractability) of our model for even large trees. In all our experiments, we have chosen a range of load values that give a multicast P_b between 10^{-3} and 10^{-1}, which is generally accepted in the literature to be the range of interest.

B. Model Validation

In all the plots presented in this section, solid curves denote analytical results obtained using the models (denoted by M-) and dotted curves denote simulation results (denoted by S-). Moreover, we present only multicast performance results since unicast performance is not our concern in this paper. Fig. 9(a) and (b) show the simulation and model results for the lighttree and lightpath approaches when $F = 16$, respectively. Results

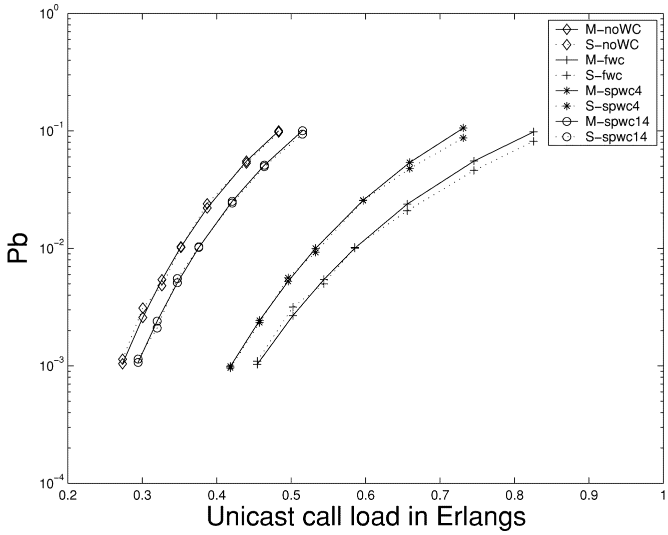

(a)

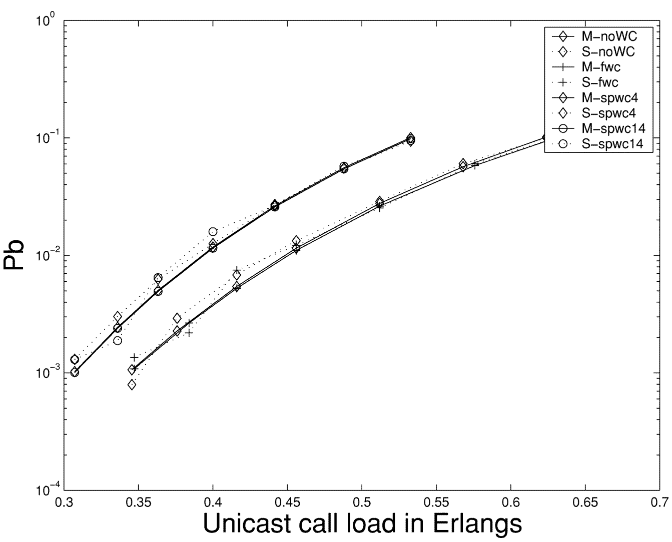

(b)

Fig. 9. Multicast blocking probability for $F = 16$ with no WC, full WC, and sparse WC (case 1: one converter placed at node 4, case 2: one converter placed at node 14) for the 7-link binary tree in NSFNET (a) using the lighttree approach and (b) using the lightpath approach.

are shown for no WC, full WC (converters placed at all the internal nodes of the tree), and two sparse WC scenarios (one converter placed at the child of the root node of the multicast tree, i.e., node 4, and one converter placed at node 14). Note that we plot unicast call load (per node pair) values on the x axis, and the multicast call load is actually $1/30$ of this load. As can be seen, the model is very accurate for both the tree and the path approaches. We make the following observations from these results. For the lighttree approach, we see that a large reduction in blocking probability is obtained when just a single converter is placed at node 4. However, there is not much more reduction in the blocking probabilities when full conversion is used. Thus, much of the performance improvement is obtained by using just one converter at the child of the root node (i.e., node 4), even though there is a nonnegligible improvement achievable by going to full conversion. Also note that when a single converter is placed at node 14, there is not much improvement compared with the case of no conversion. Similar observations can be made for the lightpath approach. Here, there is almost no performance improvement compared with the no conversion case, when a single converter is placed at node 14. When a single converter is placed at node 4, the performance is almost the same as in the case of full conversion. These observations point out the importance of proper wavelength converter placement.

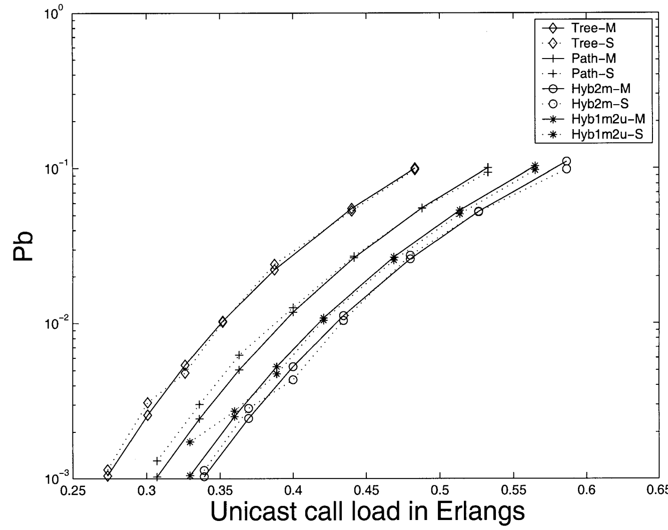

Fig. 10. Model validation for two hybrid approaches and P_b comparison for the tree, path, and the two hybrid approaches for the NSF tree such as the one in Fig. 4 at $F = 16$ and no WC. One hybrid approach has two multicast subtrees $\{$(links 1, 2, 4, 5) and (links 1, 3, 6, 7)$\}$ (denoted by Hyb2m) and the other has one multicast subtree (links 1, 2, 4, 5) and two unicast paths $\{$(links 1, 3, 6) and (links 1, 3, 7)$\}$ (denoted by Hyb1m2u).

Also, the benefits of conversion are not as much as in the lighttree approach where a drastic reduction in blocking is observed when sparse conversion is introduced. This is because the lightpath approach is mainly constrained by capacity (number of free wavelengths) and conversion does not do much to help it, whereas the performance of the lighttree approach is severely constrained by wavelength continuity and even a very small amount of conversion can greatly help.

Fig. 10 validates the results for two different hybrid approaches for satisfying multicast requests—one uses two subtrees to satisfy a request and the other uses one subtree and two lightpaths to satisfy a multicast request. In the figure, $F = 16$ and no WC is assumed. The corresponding blocking probabilities for the lightpath and lighttree approaches are also plotted for comparison. Once again, we note that the models are quite accurate in predicting the blocking performance. We have done extensive experimentation with the developed models and have found similarly accurate results that cannot be presented here due to space considerations. An interesting observation that can be made from Fig. 10 is that the best performance for the considered parameters is obtained for the two hybrid approaches. While this may seem somewhat surprising at first, it may be explained as follows. The lighttree approach has the worst performance because of the strong wavelength continuity constraint, and the lightpath approach improves performance by relaxing the wavelength continuity constraint as much as possible (at the expense of requiring more wavelengths). It

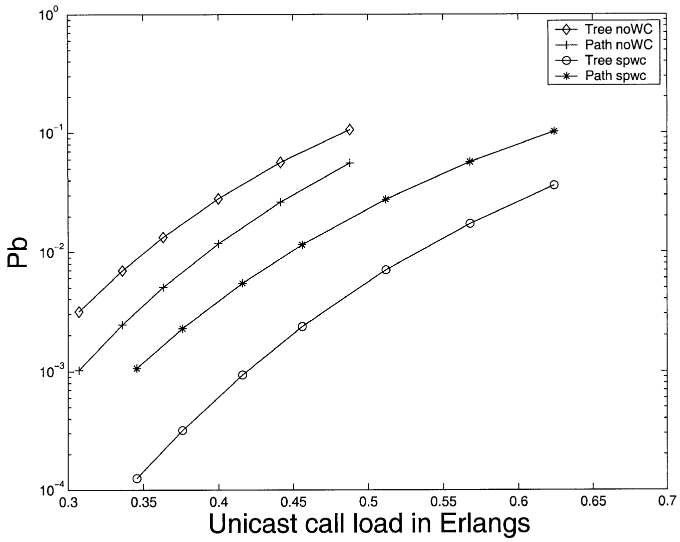

Fig. 11. P_b comparison between the lighttree and the lightpath approaches for $F = 16$, no WC, and sparse WC (one converter placed at node 4).

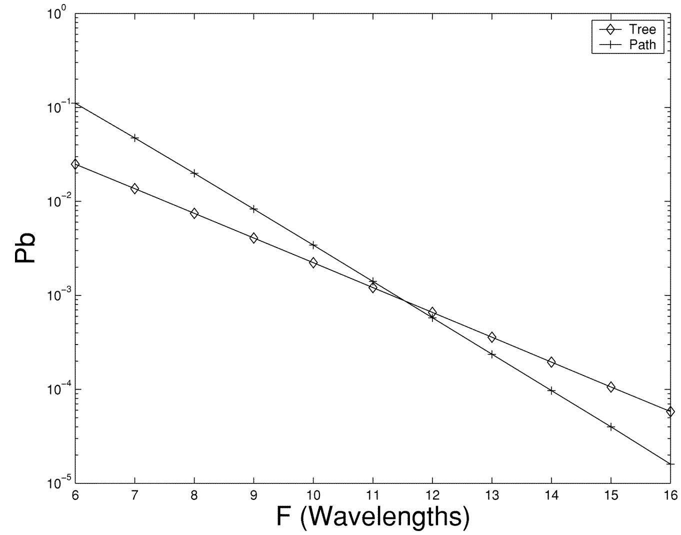

Fig. 12. P_b comparison between the lighttree and the lightpath approaches for no WC and a unicast call load per node pair per wavelength of 0.013 Erlangs.

is quite natural to then consider the possibility of improving performance even further by trying to balance the tradeoff between wavelength continuity and wavelength requirement. This is exactly what the hybrid approaches do and, indeed, an improved performance is seen.

C. Performance Evaluation

Having presented sample results to validate the performance models, we now present some more results to compare the three approaches and the benefits that wavelength conversion provides to each of the three approaches. Throughout this section, we present only analytical results, except where indicated.

It was observed from Fig. 9(a) and (b) that wavelength conversion improved the performance of the lighttree approach much more than that of the lightpath approach. We present a more direct comparison of the two approaches in Fig. 11, where we plot the blocking probability against the call load for no WC and sparse WC (with one converter at the root's child, as before) for $F = 16$. As can be seen, the substantial performance improvement with WC for the lighttree approach is applicable for a large range of loads.

We noticed in Fig. 10 that the lighttree approach performed worse than the lightpath approach for $F = 16$ without WC. As we see below in Fig. 12, when the number of wavelengths is sufficiently small, the opposite happens and the lighttree approach outperforms the lightpath approach. This is because, as the number of wavelengths decreases, the capacity constraint increasingly dominates the wavelength continuity constraint.

In order to evaluate the effects of the various parameters on the performance of the three approaches, we consider the performance of the lighttree approach as a baseline for comparison and define *relative utilization* as the ratio of the load supported using a lightpath/hybrid realization to the load supported using the lighttree approach at a given blocking level.

We plot the relative utilization for the lightpath approach and the two hybrid realizations as a function of the number of wavelengths in Fig. 13(a), where we have fixed the P_b at 5×10^{-3} and there is no WC. For $F = 6$, as we have already seen, the lighttree approach outperforms the lightpath approach as indicated by a relative utilization value of less than unity for the lightpath approach. We also see that the realization using a subtree and two lightpaths performs better, but is still worse than the lighttree approach. However, even for $F = 6$, we see that using two subtrees to realize a multicast request is better than using a single lighttree. The relative utilizations of the lightpath and the hybrid approaches increase with increasing number of wavelengths, as expected. Fig. 13(b) shows the relative utilization for sparse conversion [one converter placed at the child (i.e., node 4) of the root]. Now, the tree approach remains the best even when the number of wavelengths is large, as the effect of wavelength continuity is now mitigated to a large extent. It is interesting to note that the utilization curves flatten out beyond a point. This means that the supportable loads increase in the same order for all the approaches regardless of the presence or absence of wavelength conversion.

We next plot the relative utilization against the number of converters within the multicast tree for $F = 16$ and a P_b of 5×10^{-3} in Fig. 14. The lightpath and both hybrid realizations outperform the lighttree approach when there is no WC, but when even a single converter is introduced into the multicast tree (at the child of the root node) the lighttree approach outperforms the other realizations as the relative utilizations drop below one. There is no further significant decrease in relative utilization as more converters are added. This may be explained as follows. The presence of a converter mainly helps the lighttree approach as it is the one which is drastically constrained by the wavelength continuity requirement. For a simple and intuitive explanation, let us assume (reasonably) that the performances of the other approaches are essentially unaffected by the presence of conversion. maintained. Then, the characteristics shown in Fig. 14, depend only on how the performance of the lighttree approach varies with the number of converters. The first converter is placed at node 4, where the first splitting takes place. Intuitively, one can expect a very good improvement in performance which is confirmed by Fig. 9(a). Moreover, it is also

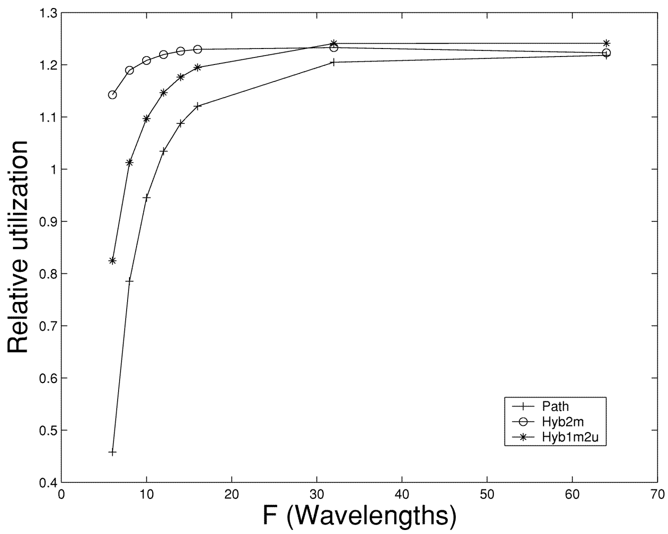

(a)

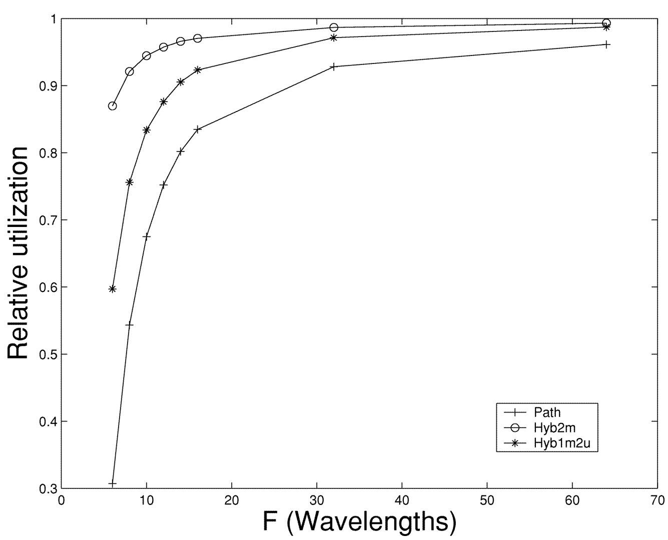

(b)

Fig. 13. Relative utilization versus F (Wavelengths) at $P_b = 5 \times 10^{-3}$ for (a) no WC, and (b) for sparse conversion (one converter at the child of the root).

reasonable to expect that most benefits of conversion are obtained by placing a single converter at the appropriate node, i.e., the child of the root. This again is confirmed by the same figure. Furthermore, we see that the relative ordering of the lightpath and the two hybrid realizations is

There has been much work recently on forming degree-constrained multicast trees in optical networks because of the increased power losses due to multiway optical splitting. What our results here have shown is the surprising fact that realizing multicast requests using multiple subtrees (essentially assuming that the degree of each subtree is constrained) may actually be *better* even from a network (blocking) performance point of view (i.e., when physical layer implications are not considered) in certain cases. However, if there is some wavelength conversion present, then the capacity advantage provided by trees with relaxed nodal degree constraints tilts the balance toward using such trees, and splitting loss would then be the limiting factor. We would like to mention here that we do not recommend any specific approach

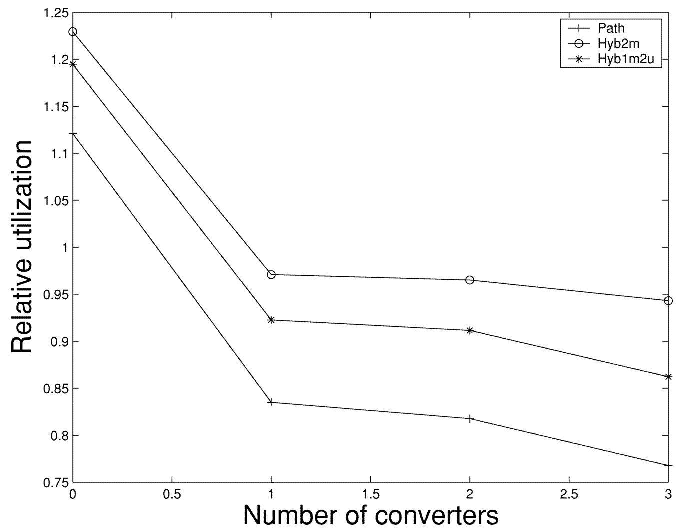

Fig. 14. Relative utilization versus the number of converters within the multicast tree for $F = 16$ and $P_b = 5 \times 10^{-3}$. One converter case corresponds to a single converter at node 4, two converters case corresponds to converters placed at nodes 4 and 14, and three converters case corresponds to converters placed at nodes 4, 14, and 7.

over the others. Which approach would perform better depends on the actual parameters and our models can be used to predict this.

All of the results presented thus far assumed that wavelengths were assigned randomly from the set of available wavelengths for the unicast calls. When a more systematic wavelength assignment algorithm is used for the unicast calls, one would expect the performance improvement due to wavelength conversion to diminish. In order to study this, we have also obtained through simulations the blocking probabilities for the multicast calls using the lighttree and lightpath realizations when the well-known first-fit wavelength assignment algorithm is used for the unicast calls. Fig. 15(a) and (b) shows these results for the cases of no WC and a single converter at the child of the root node, respectively. As expected, under first-fit assignment, introducing conversion does not improve performance as much for the tree approach as it does under random assignment. For the path approach, when there is no WC, the first-fit assignment improves multicast blocking probability by a fair amount, but there is negligible improvement when there is a converter present as now the effect of wavelength continuity is mitigated and the improved wavelength usage correlation does not do much to help reduce the blocking probability. Indeed, under low loads (and sparse conversion), random assignment performs *better* than first fit. This is because, under random assignment, the wavelength usages by the different unicast calls (including those using the multicast tree links) are less correlated, and so the overlap among the range sets is reduced in the context of the matching problem for multicast assignment. This helps enhance the blocking performance.

VII. CONCLUSION

Multicasting support is a fundamental requirement in future networks. Optical power splitters facilitate multicasting by creating many copies of the same signal all optically. By

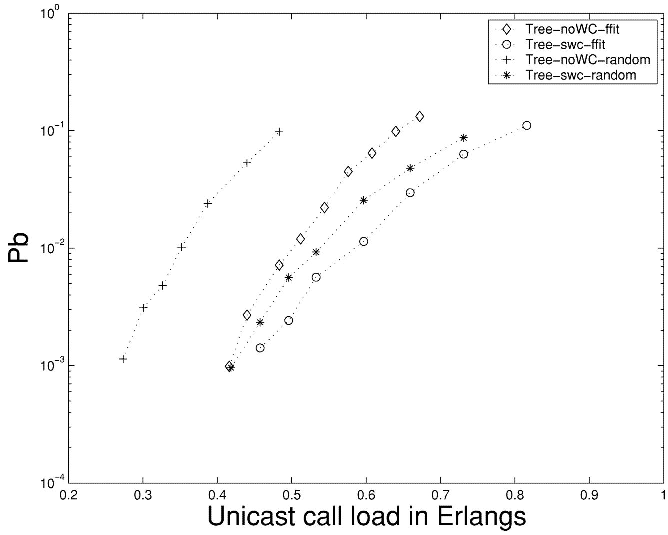

(a)

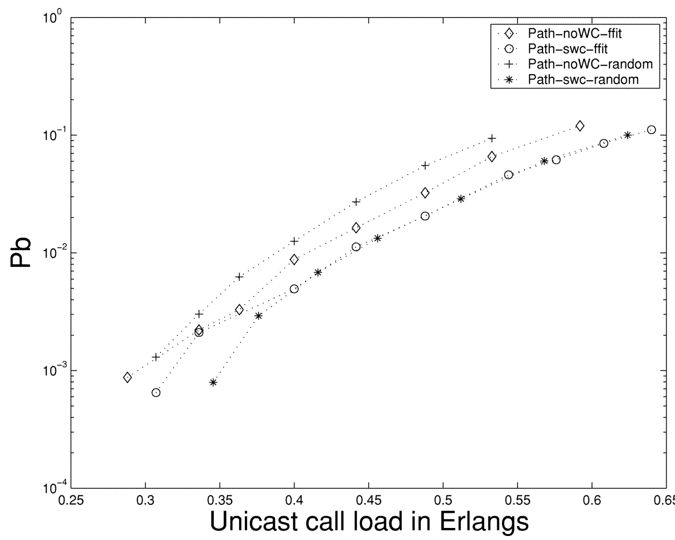

(b)

Fig. 15. Comparison of the performance under random and first-fit wavelength assignment for no WC and sparse WC for (a) the lighttree approach and (b) the lightpath approach.

eliminating the transmission of redundant traffic over certain links, multicasting can improve network performance. However, the lack of wavelength conversion necessitates the establishment of a lighttree on a single wavelength. On the other hand, establishing lightpaths on different wavelengths to satisfy a multicast request, while requiring more capacity, is less constrained in terms of wavelength assignment. Hybrid approaches that combine lightpaths and lighttrees to set up a multicast call may give the best performance.

In this paper, we developed a common framework for analyzing the blocking performance of multicast call requests when they are implemented using either the lighttree, lightpath, or a hybrid approach. The latter two may be needed in the case when only some nodes in the network are multicast-capable (with either full or partial splitting) as, then, there may not exist a single multicast tree which satisfies the splitting constraints and at the same time serves all the destinations in the destination set. The models were validated using simulation results. The results from the models and the simulations suggest that, under no wavelength conversion, it is preferable to realize the multicast call with a set of trees which have small nodal degrees. Such a realization finds the right balance between wavelength continuity and required capacity. Between the lightpath and the lighttree approaches, the former is better when there is no wavelength conversion, but even a small amount of wavelength conversion helps shift the advantage to the lighttree approach.

In our work, we did not consider alternate realizations, i.e., using one realization for one multicast request and another approach for another multicast request. All of the multicast requests within an experiment were assumed to be implemented using either the lighttree, lightpath, or the (same) hybrid approach. It would be interesting to study the performance when such alternate realizations are allowed. Multicast wavelength assignment for dynamic traffic is another topic for future study. Finally, routing and wavelength assignment for multicast requests by jointly considering the blocking performance and physical layer constraints seems to be yet another important problem to consider in the future.

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for their comments that helped improve the paper.

REFERENCES

[1] L. H. Sahasrabuddhe and B. Mukherjee, "Light-trees: Optical multicasting for improved performance in wavelength-routed networks," *IEEE Commun. Mag.*, vol. 37, pp. 67–73, Feb. 1999.

[2] M. Ali and J. S. Deogun, "Cost-effective implementation of multicasting in wavelength-routed networks," *J. Lightwave Technol.*, vol. 18, pp. 1628–1638, Dec. 2000.

[3] Y. Yang, J. Wang, and C. Qiao, "Nonblocking WDM multicast switching networks," *IEEE Trans. Parallel Distrib. Syst.*, vol. 11, pp. 1274–1287, Dec. 2000.

[4] M. Ali and J. Deogun, "Power-efficient design of multicast wavelength-routed networks," *IEEE J. Select. Areas Commun.*, vol. 18, pp. 1852–62, Oct. 2000.

[5] X. Zhang and C. Qiao, "Constrained multicast routing in WDM networks with sparse light splitting," *J. Lightwave Technol.*, vol. 18, pp. 1917–1927, Dec. 2000.

[6] W.-Y. Tseng and S.-Y. Kuo, "All-optical multicasting on wavelength routed WDM networks with partial replication," in *Proc. 15th Int. Conf. Information Networking*, 2001, pp. 813–818.

[7] R. Malli, X. Zhang, and C. Qiao, "Benefits of multicasting in all-optical networks," in *Proc. SPIE Conf. All-Optical Networking*, vol. 2531, Nov. 1998, pp. 209–220.

[8] N. Sreenath and C. Murthy, "Virtual source based multicast routing in wdm optical networks," *Photonic Network Commun.*, vol. 3, no. 3, pp. 213–226, 2001.

[9] R. Libeskind-Hadas and R. Melhem, "Multicast routing and wavelength assignment in multi-hop optical networks," *IEEE Trans. Networking*, vol. 10, pp. 621–629, Oct. 2002.

[10] R. Srinivasan and A. K. Somani, "A generalized framework for analyzing time-space switched optical networks," *IEEE J. Select. Areas Commun.*, vol. 20, pp. 179–88, Apr. 2001.

[11] S. Ramesh and G. N. Rouskas, "Computing blocking probabilities in multiclass wavelength-routing networks with multicast calls," *IEEE J. Select. Areas Commun.*, vol. 20, pp. 89–96, Jan. 2002.

[12] A. E. Kamal and A. Al-Yatama, "Blocking probabilities in circuit-switched WDM networks under multicast service," *Perform. Eval.*, vol. 47, no. 2, pp. 43–71, Jan. 2002.

[13] X. Qin and Y. Yang, "Blocking probability of WDM switching networks with limited wavelength conversion," in *Proc. ICC*, 2002, pp. 454–459.

[14] Y. Wang and Y. Yang, "Multicasting in a class of multicast-capable WDM networks," *J. Lightwave Technol.*, vol. 20, pp. 350–359, Mar. 2002.

[15] M. Ali, "Optimization of splitting node placement in wavelength-routed optical networks," *IEEE J. Select. Areas Commun.*, vol. 20, pp. 1571–1579, Oct. 2002.

[16] T. Znati, T. Alrabiah, and R. Melhem, "Point-to-multipoint path establishment schemes to support multicasting in WDM networks," presented at the Proc. IFIP Working Conf. Optical Network Design and Modeling, 1999.

[17] N. Singhal and B. Mukherjee, "Architectures and algorithm for multicasting in WDM optical mesh networks using opaque and transparent optical cross-connects," presented at the Proc. OFC, Mar. 2001.

[18] Y. Sun, J. Gu, and D. H. K. Tsang, "Multicast routing in all-optical wavelength routed networks," *Opt. Networks Mag.*, pp. 101–109, July/Aug. 2001.

[19] X. Jia, D. Du, X. Hu, M. Lee, and J. Gu, "Optimization of wavelength assignment for QoS multicast in WDM networks," *IEEE Trans. Commun.*, vol. 49, pp. 341–350, Feb. 2001.

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

[21] R. K. Pankaj and R. G. Gallager, "Wavelength requirements of all-optical networks," *IEEE/ACM Trans. Networking*, vol. 3, pp. 269–280, June 1995.

[22] A. Girard, *Routing and Dimensioning in Circuit-Switched Networks*. Reading, MA: Addison-Wesley, 1990.

[23] S. Subramaniam, M. Azizoĝlu, and A. K. Somani, "All optical networks with sparse wavelength conversion," *IEEE/ACM Trans. Networking*, vol. 4, pp. 544–557, Aug. 1996.

[24] J. A. Bondy and U. S. R. Murty, *Graph Theory with Applications*. New York: American Elsevier, 1976.

[25] S. Sankaranarayanan, "Performance evaluation of multicasting in optical networks," Master's Thesis, George Washington Univ., Washington, DC, 2003.

Srivatsan Sankaranarayanan received the B.Tech. degree in electrical engineering from Indian Institute of Technology, Madras, in 1999, and the M.S. degree in electrical engineering from The George Washington University, Washington, DC, in 2003. Currently, he is a doctoral student at Virginia Polytechnic Institute and State University, Blacksburg.

His research interests are in wireless networks, optical networks, and performance analysis of communication networks and protocols.

Suresh Subramaniam (S'95–M'97) received the Ph.D. degree in electrical engineering from the University of Washington, Seattle, in 1997.

He is an Associate Professor in the Department of Electrical and Computer Engineering, George Washington University, Washington, DC. His research interests include performance analysis of and algorithm design for communication networks, with particular emphasis on optical and wireless ad hoc networks. He is a coeditor of the book *Optical WDM Networks—Principles and Practice* (Norwell, MA: Kluwer, 2000).

Dr. Subramaniam has been on the program committees of several conferences including INFOCOM, ICC, GLOBECOM, and Opticomm, and currently serves on the Editorial Boards of the *Journal of Communications and Networks*, the *IEEE Communications Surveys and Tutorials*, and *Optical Networks Magazine*. He is a corecipient of the Best Paper Award at the 1997 SPIE Conference on All-Optical Communication Systems.