Multicast Routing and Wavelength Assignment in Multihop Optical Networks

Ran Libeskind-Hadas, *Member, IEEE*, and Rami Melhem, *Fellow, IEEE*

***Abstract*—This paper addresses multicast routing in circuit-switched multihop optical networks employing wavelength-division multiplexing. We consider a model in which multicast communication requests are made and released dynamically over time. A multicast connection is realized by constructing a multicast tree which distributes the message from the source node to all destination nodes such that the wavelengths used on each link and the receivers and transmitters used at each node are not used by existing circuits. We show that the problem of *routing and wavelength assignment* in this model is, in general, NP-complete. However, we also show that for any given multicast tree, the *wavelength assignment* problem can be solved in linear time.**

***Index Terms*—Multicast communication, optical networks, routing and wavelength assignment, wavelength-division multiplexing (WDM).**

I. Introduction

WAVELENGTH-DIVISION multiplexing (WDM) is emerging as a key technology in communication networks. In WDM networks, the fiber bandwidth is partitioned into multiple data channels which may be transmitted simultaneously on different wavelengths. Thus, WDM permits the use of enormous fiber bandwidth by providing data channels whose individual bandwidths more closely match those of the electronic devices at their endpoints [18].

In this paper, we consider circuit-switched WDM networks. Such networks can be classified as either *single-hop* or *multihop* networks [10], [11]. In single-hop (or *all-optical*) networks, each message is transmitted from the source to the destination without any optical-to-electronic conversion within the network. Single-hop communication can be realized by using a single wavelength to establish a connection, but such connections may, in general, be difficult or impossible to find in the presence of other network traffic [15]. Alternatively, all-optical wavelength converters may be used to convert from one wavelength to another within the network but such converters are likely to be prohibitively expensive for most applications in the foreseeable future [18].

Manuscript received March 27, 2001; revised January 26, 2002; approved by IEEE/ACM Transactions on Networking Editor G. Rouskas. This work was supported by the National Science Foundation under Grant CCR-9900491 to Harvey Mudd College and Grant MIP 96-33729 to the University of Pittsburgh. A preliminary version of this work appeared in the Proceedings of the 2001 IEEE International Conference on Networks, Colmar, France.

R. Libeskind-Hadas is with the Department of Computer Science, Harvey Mudd College, Claremont, CA 91711 USA (e-mail: hadas@cs.hmc.edu).

R. Melhem is with the Department of Computer Science, University of Pittsburgh, Pittsburgh, PA 15260-3803 USA (e-mail: melhem@cs.pitt.edu).

Digital Object Identifier 10.1109/TNET.2002.803907.

In *multihop* communication networks, a message entering an intermediate node on a particular wavelength can be converted into the electronic medium by a receiver and retransmitted on a new wavelength by a transmitter. We define the *hops* incurred by a message to be the number of transmissions used to send the message from the source to the destination. For example, an all-optical connection from the source to the destination incurs one hop, whereas a connection requiring one retransmission incurs two hops, and so on. Multihop networks have been shown to enjoy higher utilization of bandwidth and lower probability of blocking than single-hop networks [2]. However, a multihop connection may use more transmitters and receivers than a single-hop connection and, depending on the network architecture, each hop can contribute significantly to the communication latency. Therefore, it is generally desirable to find multihop connections that minimize the number of transmitters and receivers and/or the number of hops used.

Finally, a network may support *unicast* (or *one-to-one*) communication as well as *multicast* (or *one-to-many*) communication. Multicast communication arises in a wide variety of applications such as video distribution and teleconferencing [12]. Given a source node and a set of destination nodes, the multicast *routing and wavelength assignment* (RWA) problem is that of finding a set of links and wavelengths on these links on which to establish the connection from the source to the destination nodes. Similarly, the *wavelength assignment* (WA) problem is that of finding a set of wavelengths to use on a predetermined multicast tree. In general, it is desirable to find an RWA or a WA that is optimal with respect to some cost metric.

Various multicast RWA and WA problems have been investigated recently for both packet- and circuit-switched WDM networks [1], [3], [7], [8], [12], [13], [15]–[17], [19]–[21]. For example, Bermond *et al.* [3] have investigated RWA for all-optical multicast in networks with only unicast-capable switches. Ali and Deogun [1] have investigated RWA in networks employing *tap-and-continue* switches which have limited multicast capabilities. Liang and Shen [8] have investigated the problem of finding minimum-cost RWAs in a model that associates a specific cost for each wavelength on each link as well as a cost for converting between wavelengths at intermediate nodes. Sahin and Azizoglu [14] have investigated multicast RWA under various fanout splitting policies, and Malli *et al.* [9] have investigated the problem under a sparse splitting model. Sahasrabuddhe and Mukherjee [13] have formulated the RWA problem for multihop multicast routing in packet-switched networks as a mixed-integer linear programming problem.

In this paper, we investigate multicast communication in circuit-switched multihop networks. We consider networks with an arbitrary number of nodes, a fixed number of transmitters and receivers at each node, and a fixed number of wavelengths on each link. Multicast communication requests are made and released over time. A multicast connection may be realized by constructing a multicast tree which distributes the message from the source node to all destination nodes such that the wavelengths used on each link and the receivers and transmitters used at each node are not used by existing circuits. We assume a *multicast-capable* switch model [6], [9], [13] in which a wavelength on an input link may be routed to the same wavelength on any number of output links and, optionally, to a receiver at the local node. A message received at a local node can be retransmitted from that node on another wavelength (thereby incurring a hop).

The RWA problem in this model is that of selecting a multicast tree, the wavelengths on the links in the tree, and, thus, the intermediate nodes that will perform wavelength conversion. In the WA problem, a multicast tree is given and the problem is that of selecting the wavelengths on the links in the tree and the intermediate nodes for wavelength conversion. In general, it may not always be possible to find an RWA or a WA for a given multicast request because the needed wavelengths on particular links, as well as transmitters and receivers at intermediate nodes, are being used by other multicast connections.

In this paper, we show that the RWA problem in this model is, in general, NP-complete, but that the WA problem can be solved in linear time.[1] Moreover, we show that the linear-time WA algorithm can be extended to find "optimal" solutions under various definitions of optimality such as minimizing the maximum number of hops or minimizing any linear combination of the number of transmitters and receivers used.

The remainder of this paper is organized as follows. In Section II, we formally describe the model under consideration and define notation. In Section III, we investigate the complexity of the RWA problem. In Section IV, we give a linear-time algorithm for the WA problem and generalize the algorithm to find optimal multicasts with respect to minimizing the maximum number of hops and minimizing any linear combination of the number of transmitters and receivers used. Section V describes experimental results using these algorithms. Conclusions are given in Section VI.

II. Model and Notation

We represent an interconnection network by a connected directed graph $G = (V, E)$ where the vertices represent switches and the directed edges represent links between pairs of switches. Each switch may be connected to a node or network access station. Except where the distinction is necessary, we henceforth use the terms "switch," "node," and "vertex" interchangeably and let n denote $|V|$. Similarly, we use "link" and "edge" interchangeably. Each link can carry some number w of different wavelengths denoted by $\Lambda = \{\lambda_1, \ldots, \lambda_w\}$. Each node v has $T(v)$ tunable transmitters and $R(v)$ tunable receivers, each of which can tune to any of the w wavelengths. Let $d_{\text{in}}(v)$ and $d_{\text{out}}(v)$ denote the number of incoming and outgoing links, respectively, at node v. We assume that the number of nodes n in the network is variable but that parameters w, $T(v)$, $R(v)$, $d_{\text{in}}(v)$, and $d_{\text{out}}(v)$ are bounded by constants dictated by the technology.

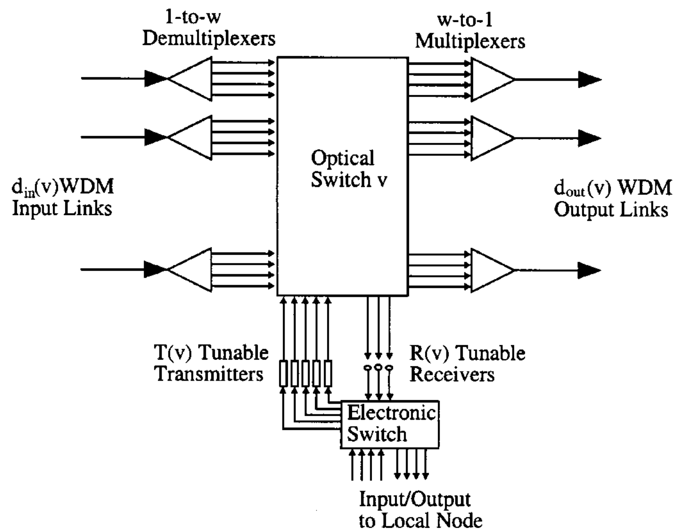

Fig. 1. Schematic of switch model.

A wavelength on an input link may be routed to the same wavelength on any number of output links and, optionally, to a receiver at the local node. Similarly, a message transmitted on a particular wavelength by a transmitter at a node may be routed on this wavelength to any number of output links. Routing must satisfy the constraint that two messages using the same wavelength cannot share the same link. A switch model with these properties is shown in Fig. 1. Switches with some similar characteristics were described by Kovačević and Acampora [6] and by Sahasrabuddhe and Mukherjee [13]. We note that the results described in this paper can be adapted to a number of other switch models.[2]

A *multicast communication request* is an ordered pair (s, D) where $s \in V$ is the source of the multicast and $D \subseteq V - s$ is the set of destination nodes. We assume that multicast communication requests are made and released dynamically. At the time that a particular multicast communication request is made, there may be some limits imposed on the routing resources available in the network. Specifically, each node v has some available number $t(v)$ of transmitters and $r(v)$ of receivers that can be used to implement the multicast where $t(v) \leq T(v)$ and $r(v) \leq R(v)$. In addition, each link (v, x) has some set $w(v, x) \subseteq \Lambda$ of available wavelengths. Let $W(v)$ denote the total number of distinct wavelengths available on all outgoing links from node v. These resource limits may reflect the actual available resources, due to utilization of resources by

[1]Note that the NP-completeness result described here is unrelated to the NP-completeness of finding a minimum-cost RWA in a weighted graph [8], [15]. In the case described here, the computational intractability arises from the limited number of transmitters and receivers at intermediate nodes and the constraints on which wavelengths are available on each link. In the case of weighted graph, the intractability is due to the fact that the RWA problem contains the NP-complete Steiner tree problem as a special case.

[2]For example, one possible variant of this model has a dedicated receiver associated with each transmitter. In this model, a message arriving on an input link in Fig. 1 can be routed by the optical switch directly to a receiver/transmitter pair, bypassing the electronic switch. This model is similar to the share-per-node switch architecture described in [7].

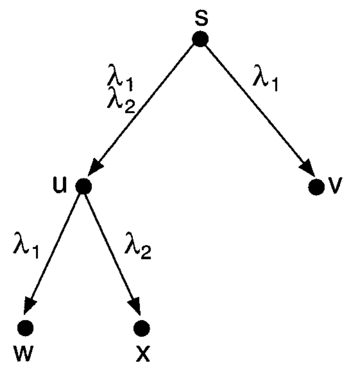

Fig. 2. Node s has two transmitters and all other nodes have no transmitters.

existing connections, or these limits may be imposed in order to reduce cost or to leave resources available for subsequent connection requests.

Due to these resource constraints, it may not be possible to realize a multicast communication request, in which case the request is said to be *blocked*. Moreover, in some cases it may not be possible to realize a request when only one wavelength may carry the message on each link, while the connection may be realizable when multiple wavelengths are permitted to carry the message on the same link. Let ℓ denote the maximum number of wavelengths that may be used to transmit the same message over any single link. Fig. 2 illustrates an example of a network in which node s is the source of the multicast and all the remaining nodes are destinations. In this example, node s has two transmitters while all remaining nodes have zero transmitters. When $\ell = 1$, node s may use only a single wavelength on each link. Since node u has no transmitters, it is not possible for the message to be delivered to both destinations w and x. On the other hand, when $\ell = 2$ node s may transmit on both wavelengths λ_1 and λ_2 and both of these wavelengths can be used on link (s, u). The message is then delivered to w using wavelength λ_1 on link (u, w) and on wavelength λ_2 on link (u, x). Node v receives the message on wavelength λ_1 on link (s, v). In this case, all destination nodes can be reached. Thus, by limiting the value of ℓ, we may limit the wavelength resources allocated to the multicast but consequently decrease the probability of satisfying a given multicast communication request. A summary of notation defined to this point is given in Table I.

We now formalize the definitions of the RWA and WA problems.

Definition 1: Let $G = (V, E)$ be a directed graph and (s, D) a multicast communication request in this graph. An RWA is a collection of links, wavelengths on these links, and wavelength settings for transmitters and receivers at each node such that each $v \in D$ receives the message from s, at most ℓ wavelengths from $w(v, x)$ are used on each link $(v, x) \in E$, and no more than $t(v)$ transmitters and $r(v)$ receivers are used at each node $v \in V$.

Definition 2: Let $G = (V, E)$ be a directed graph, (s, D) a multicast communication request in this graph, and τ a unidirectional subtree of G with root s and containing all vertices in D. A WA with respect to τ is a set of wavelengths on the links in τ and wavelength settings for transmitters and receivers at each node in τ such that each $v \in D$ receives the message from s, at most ℓ available wavelengths from $w(v, x)$ are used on each link (v, x) in τ, and no more that $t(v)$ transmitters and $r(v)$ receivers are used at each node v in τ.

TABLE I
SUMMARY OF NOTATION

Notation	Definition
$G = (V, E)$	Directed graph representing an optical network
n	Number of nodes ($\lvert V\rvert$)
w	Number of wavelengths
$\Lambda = \{\lambda_1, \ldots, \lambda_w\}$	Set of wavelengths
$T(v)$	Total number of transmitters at v
$R(v)$	Total number of receivers at v
$d_{\text{in}}(v)$	In-degree of v
$d_{\text{out}}(v)$	Out-degree of v
s	Source node of multicast
D	Set of destination nodes of multicast
$t(v)$	Number of available transmitters at v
$r(v)$	Number of available receivers at v
$w(v, x)$	Set of available wavelengths on link (v, x)
$W(v)$	Total number of wavelengths available on all outgoing links from node v
ℓ	Maximum number of wavelengths permitted to carry the same message on a single link

Note that in both the RWA and WA problems it is assumed that a single multicast communication request is made at one time. The RWA used to satisfy this request will reduce the set of available resources which may be used for the next multicast communication request. Also note that in the WA problem it is assumed that the tree is unidirectional with edges oriented from the source to the destinations.

In some cases, we may wish to find an RWA or WA that is optimal with respect to a given measure. For example, for some switch models we may wish to find a routing that minimizes the maximum number of hops from the source to any destination. Alternatively, to minimize use of resources, we may wish to find a routing that minimizes some linear combination of the number of receivers and transmitters used.

In Section III, we show that the problem of finding an RWA can be solved in polynomial time for a special case but, in general, is NP-complete. Thus, the problem of finding an optimal RWA according to a given measure of optimality is also, in general, NP-complete. We then show that the problem of finding a WA can be solved in linear time. Moreover, this algorithm can be adapted to find optimal WAs for different measures of optimality, also in linear time.

III. Routing and Wavelength Assignment Problem

In this section, we investigate the complexity of the RWA problem.

Theorem 1: For any value of $\ell \geq 1$, if $t(v) \geq \min\{W(v), d_{\text{out}}(v)\}$ and $r(v) \geq 1$ for all $v \in V$, then an RWA can be found, or it can be determined that none exists, in time $O(n)$.

Proof: Remove from $G = (V, E)$ any edge that contains no available wavelengths. Since $t(v) \geq \min\{W(v), d_{\text{out}}(v)\}$, each node v may transmit the message to all of its neighbors in G such that exactly one wavelength is used on each outgoing edge from v. Since each node has at least one receiver available, each neighbor of v can receive the message from v. Thus, v

may transmit the message to all of its neighbors for any value of $\ell \geq 1$.

Apply a breadth-first search in the graph beginning at source node s. From the observations above, if all destination nodes are reached by the search, then an RWA exists. Otherwise, no RWA exists. The running time of a breadth-first search is $O(|V|+|E|)$ and since the degree of each node is upper bounded by some constant, $O(|E|) = O(|V|)$. Thus, an RWA can be found or it can be determined that none exists in time $O(|V|) = O(n)$. □

In general, we cannot assume that the number of transmitters available at each node is at least as large as the number of available wavelengths on the outgoing links. In the next theorem we show that when the number of transmitters available at each node is not necessarily as large as the number of available wavelengths, the problem of finding an RWA is NP-complete.

Theorem 2: For the general case, the problem of determining if there exists an RWA is NP-complete.

Proof: This decision problem is clearly in the class NP since a solution can be easily verified in polynomial time. The reduction is from a restricted version of 3-Satisfiability (3SAT), in which each variable occurs in at most five clauses. This restricted version of 3SAT is known to be NP-complete. For a given instance of the restricted 3SAT problem, let $x_1, \ldots, x_v$ denote the variables and let $C_1, \ldots, C_k$ denote the clauses, each of which contains the disjunction of exactly three literals over the set of variables. Corresponding to an instance of restricted 3SAT, we construct a network as follows. Vertex s is the source of the multicast. For each variable x_i, $1 \leq i \leq v$, there is a corresponding vertex labeled with the name of the variable. For each clause C_j, $1 \leq j \leq k$, there is a corresponding vertex labeled with the name of the clause. Construct a directed path originating at vertex s and passing through vertices $x_1, \ldots, x_v$. On each of the v links on this path, wavelength λ_1 is the only available wavelength. For each occurrence of literal x_i ($\overline{x}_i$) in clause C_j, there is an edge with wavelength λ_{true} (λ_{false}) from vertex x_i to vertex C_j. All vertices have one receiver and one transmitter. All vertices other than s are destination vertices. This reduction can clearly be performed in time polynomial in the size of the restricted 3SAT instance. In addition, the number of transmitters and receivers at each node, the number of wavelengths, and the in and out degrees of each vertex are upper bounded by constants, as required by our model. Specifically, no vertex has more than one receiver or one transmitter, the total number of wavelengths is three, and no vertex has in or out degree larger than five. Finally, since each link has exactly one available wavelength, the reduction holds for any $\ell \geq 1$. An illustration of this construction is shown in Fig. 3.

We claim that an RWA exists for the multicast problem instance if and only if the given restricted 3SAT instance is satisfiable. Assume that the given restricted 3SAT instance is satisfiable. Construct a corresponding solution to the multicast problem by transmitting a message from s to $x_1, \ldots, x_v$ using wavelength λ_1. For each variable x_i, if x_i is true in the satisfying assignment, then vertex x_i uses its single transmitter to transmit the message on wavelength λ_{true} and otherwise transmits it on wavelength λ_{false}. Vertex x_i transmits on the selected wavelength to each vertex C_j that has not yet received the message. Consider an arbitrary destination vertex C_j, $1 \leq j \leq k$. The corresponding clause in the restricted 3SAT instance contains at least one literal which evaluates to true. If literal $x_i \in C_j$ evaluates to true in the restricted 3SAT instance, then vertex C_j receives the message on wavelength λ_{true} in the constructed multicast problem instance. If literal $\overline{x}_i \in C_j$ evaluates to true, then x_i is false and vertex C_j receives the message on wavelength λ_{false}. Thus, every destination vertex receives a copy of the message.

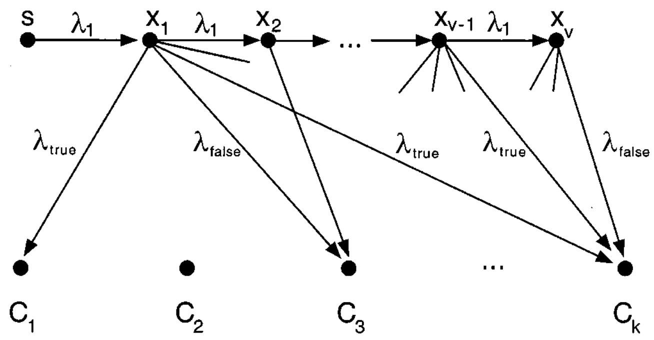

Fig. 3. Construction used in the proof of Theorem 2.

Conversely, assume that every destination vertex receives the message. Then vertex s delivers the message to all of $x_1, \ldots, x_v$ on wavelength λ_1. Each vertex $x_1, \ldots, x_v$ must then transmit the message on at most one of wavelengths λ_{true} or λ_{false}. For each vertex x_i, assign the corresponding variable in the restricted 3SAT instance to be true if the vertex transmits on wavelength λ_{true} and assign the variable to be false otherwise. Each clause C_j is satisfied by this assignment since at least one of its literals evaluates to true. □

IV. Wavelength Assignment Problem

In this section, we show that the WA problem can be solved in linear time. Throughout this section, the following assumptions are made.

1) A fixed multicast tree is given with source node s at the root. All destination nodes are in the tree, although the tree may also contain nondestination nodes.

2) All leaves in the multicast tree are destination nodes. (Otherwise, leaf nodes can be repeatedly removed until this property is true.)

3) For each destination node v in the multicast tree, the number of receivers is greater than 0. That is, $r(v) > 0$. (Otherwise, no WA exists.)

We begin in Section IV-A by examining the case that $\ell = 1$. In Section IV-B, we show how the algorithm can be adapted to find optimal WAs for different criteria of optimality when $\ell = 1$. In Section IV-C, we examine the case $\ell > 1$.

A. Wavelength Assignment for $\ell = 1$

The algorithm is based on dynamic programming. For each nonroot node v, let $p(v)$ denote the parent of v in the given multicast tree. Then $(p(v), v)$ denotes the link from the parent of v to v. Define the predicate $m_v(\Lambda) \rightarrow \{\text{true}, \text{false}\}$ by $m_v(\lambda) = \text{true}$ if and only if wavelength λ is available on the link $(p(v), v)$ and node v can deliver the message to all destinations in its subtree if it receives the message on wavelength λ. Recall that every leaf is a destination node. Thus, from the above definition it follows that for each leaf v in the tree

$$m_v(\lambda) = \begin{cases} \text{true}, & \text{if } \lambda \in w(p(v), v) \\ \text{false}, & \text{otherwise.} \end{cases} \tag{1}$$

In other words, if v is a leaf, then $m_v(\lambda)$ is true if and only if wavelength λ is available on the link from v's parent to v.

Next, consider an internal nonroot node v which has no receivers available. Since $r(v) = 0$, node v may forward the message on the incoming wavelength to its children using the same wavelength but it may not receive the message and then retransmit it on other wavelengths. Let $C(v)$ denote the set of children of v. Let $\bigwedge$ and $\bigvee$ denote the Boolean AND and OR operators, respectively. If $r(v) = 0$, then

$$m_v(\lambda) = \begin{cases} \bigwedge_{x \in C(v)} m_x(\lambda), & \text{if } \lambda \in w(p(v), v) \\ \text{false}, & \text{otherwise.} \end{cases} \tag{2}$$

This rule asserts that v can deliver a message received on wavelength λ to all destinations in its subtree if and only if λ is available on the link entering v from its parent and each child x of v can deliver the message to all destinations in its subtree if x receives the message on wavelength λ.

Next, consider the case that $r(v) > 0$. In this case, node v can use wavelength λ to deliver the message to its children and, in addition, node v can receive the message and retransmit the message to its children using up to $t(v)$ wavelengths other than λ. Define a *wavelength selection set* with respect to λ to be a subset of Λ which contains λ. Let $\mathcal{A}_{\lambda, c}$ denote the set of all wavelength selection sets with respect to λ of size at most $c + 1$. Thus, every set in $\mathcal{A}_{\lambda, c}$ comprises λ and up to c additional wavelengths. Then

$$m_v(\lambda) = \begin{cases} \bigvee_{A \in \mathcal{A}_{\lambda, t(v)}} \bigwedge_{x \in C(v)} \bigvee_{\lambda' \in A} m_x(\lambda'), & \text{if } \lambda \in w(p(v), v) \\ \text{false}, & \text{otherwise.} \end{cases} \tag{3}$$

This rule asserts that $m_v(\lambda)$ is true if and only if wavelength λ is available on the link entering v from its parent and there exists some wavelength selection set A comprising λ and up to $t(v)$ additional wavelengths (to be transmitted at v) with the following property: Every child x of v can deliver the message to all of its descendant destinations if it receives the message on one of the wavelengths λ' in set A.

Finally, consider the case of the root node s. Unlike the other nodes in the tree, node s does not receive the message from a parent node. Instead, node s transmits the message using up to $t(s)$ different wavelengths. Let $\mathcal{B}_c$ denote the set of all subsets of Λ of size at most c. Define $M = \text{true}$ if and only if a WA exists originating at the source node. Then

$$M = \bigvee_{B \in \mathcal{B}_{t(s)}} \bigwedge_{x \in C(s)} \bigvee_{\lambda' \in B} m_x(\lambda'). \tag{4}$$

This rule is analogous to the one in (3) except that node s now transmits all wavelengths itself rather than receiving one on an incoming link.

The dynamic programming algorithm is shown in Algorithm 1. Recall that given an acyclic directed graph with n vertices $v_1, \ldots, v_n$, a topological ordering of the vertices is a permutation $v_{i_1}, \ldots, v_{i_n}$ of the vertices such that if there is a directed edge from v_{i_j} to v_{i_k} then $j < k$. Since the multicast tree is acyclic, there exists a topological ordering of the vertices [4]. Note that by visiting the vertices in the order $v_{i_n}, \ldots, v_{i_1}$, a node is only visited if all of its descendants have been visited. Note that the dynamic programming algorithm computes the value of $m_v(\lambda)$ from the bottom of the tree upwards. Therefore, when $m_v(\lambda)$ is computed for a given vertex v, the required values for all of its children in the tree have already been computed.

```
Algorithm 1
Compute a topological ordering v_{i_1}, ..., v_{i_n}
  of the n nodes in the multicast tree
for j = n down to 1
    Let v = v_{i_j}
    for each wavelength λ
      if v is a leaf node then compute m_v(λ)
  using (1)
      if j > 1 and r(v) = 0 then compute m_v(λ)
  using (2)
      if j > 1 and r(v) > 0 then compute m_v(λ)
  using (3)
      if j = 1 then compute M using (4)
    end for (Comment: End inner for loop)
end for (Comment: End outer for loop)
return (M)
```

Note that the actual WA can be found, if one exists, by recording the WAs in addition to the values of $m_v(\lambda)$ and M.

We now derive an upper bound on the running time of the algorithm. In general, computing a topological ordering takes time $O(n + m)$ where $n = |V|$ and $m = |E|$ [4]. Since $m = n - 1$ in a tree, the topological ordering can be computed in time $O(n)$. There are a total of wn iterations through the **for** loops. Among the computations performed inside the **for** loops, the computation in (3) requires the largest number of steps. An upper bound on the number of steps required to compute $m_v(\lambda)$ in (3) can be derived as follows: For each wavelength λ at most $\sum_{i=0}^{t(v)} \binom{w-1}{i}$ distinct wavelength selection sets are considered because there are $\binom{w-1}{i}$ ways of choosing i wavelengths other than λ from Λ. For each wavelength selection set A, consider the set of children $C(v)$ of node v. Set $C(v)$ has size at most $d_{\text{out}}(v)$, and for each $x \in C(v)$, at most $t(v) + 1$ steps are required to determine if there exists a wavelength $\lambda' \in A$ such that $m_x(\lambda')$ is true. Therefore, in the worst case the number of steps required to compute $m_v(\lambda)$ is bounded by $[\sum_{i=0}^{t(v)} \binom{w-1}{i}] d_{\text{out}}(v)(t(v)+1)$. Letting $t = \max_{v \in V} t(v) + 1$, $C = \sum_{i=0}^{t-1} \binom{w}{i}$, and $d = \max_{v \in V} d_{\text{out}}(v)$, the running time of the computations performed inside the **for** loops is upper bounded by $[wCdt]n$. Thus, the algorithm has $O(n)$ running time, with the constant term depending on constants w, t, and d. The impact of these constants on the running time, in practice, is discussed in Section V.

B. *Optimal Multicast for* $\ell = 1$

Here, we show that the dynamic programming solution described in the previous section can be adapted to solve the problem of finding optimal WAs, or determining that no WA exists, under different measures of optimality. For concreteness, we demonstrate this for two specific measures: WAs which minimize the maximum number of hops required to reach all destination nodes and WAs which minimize the total number of transmitters used. The latter result can be easily generalized to minimize any linear combination of the number of transmitters and receivers used.

1) Minimizing the Maximum Number of Hops: Recall that the number of hops incurred from the source to a destination is defined to be the number of transmissions. Thus, each message incurs at least one hop. For each nonroot node v, $h_v(\lambda)$ is defined to be the minimum value k such that there exists a path from v to every destination node in the subtree rooted at v which uses at most k hops, assuming the message enters v on wavelength λ. If wavelength λ is not available on link $(p(v), v)$ or it is not possible for v to reach all of the destination nodes in its subtree when the message enters v on wavelength λ, then define $h_v(\lambda) = \infty$. (This is equivalent to the case that $m_v(\lambda) = \text{false}$.) For example, if wavelength λ is available on link $(p(v), v)$ and wavelength λ is available on every link in the subtree rooted at v, then $h_v(\lambda) = 0$ since no hops need to be performed inside the subtree rooted at v in order to reach all destination nodes in that subtree.

From the definition, it follows that for each leaf v in the tree

$$h_v(\lambda) = \begin{cases} 0, & \text{if } \lambda \in w(p(v), v) \\ \infty, & \text{otherwise.} \end{cases} \tag{5}$$

Next, consider an internal nonroot node v. If $r(v) = 0$, node v cannot receive and retransmit the message but may only distribute the message to its children using wavelength λ. Thus, if $r(v) = 0$, then

$$h_v(\lambda) = \begin{cases} \max\limits_{x \in C(v)} h_x(\lambda), & \text{if } \lambda \in w(p(v), v) \\ \infty, & \text{otherwise.} \end{cases} \tag{6}$$

If $r(v) > 0$, node v may distribute the message to its children on wavelength λ without incurring an additional hop. In addition, node v may receive the message and retransmit it to v's remaining children using up to $t(v)$ wavelengths other than λ. Each child which receives the message on a wavelength λ' other than λ incurs an additional hop. Then, for $r(v) > 0$

$$h_v(\lambda) = \begin{cases} \min\limits_{A \in \mathcal{A}_{\lambda,\, t(v)}} \max\limits_{x \in C(v)} \min\limits_{\lambda' \in A} \begin{cases} h_x(\lambda'), & \text{if } \lambda' = \lambda \\ 1 + h_x(\lambda'), & \text{if } \lambda' \neq \lambda \end{cases}, & \text{if } \lambda \in w(p(v), v) \\ \infty, & \text{otherwise.} \end{cases} \tag{7}$$

Let H denote the minimum number of hops required. One hop is incurred by the initial transmission of the message at node s. Therefore

$$H = 1 + \min_{B \in \mathcal{B}_{t(s)}} \max_{x \in C(s)} \min_{\lambda' \in B} h_x(\lambda'). \tag{8}$$

Finally, in Algorithm 1, $m_v(\lambda)$ and M are replaced by $h_v(\lambda)$ and H, respectively, and (1)–(4) are replaced by (5)–(8), respectively. The asymptotic running time and constants are easily verified to be the same as that of the original algorithm.

2) Minimizing the Total Number of Transmitters: To simplify our discussion, we restrict our attention here to that of minimizing the total number of transmitters. A straightforward extension of this formulation can be used to minimize any linear combination of the number of transmitters and receivers.

For each nonroot node v, $tr_v(\lambda)$ is defined to be the minimum number of transmitters that must be used in the subtree rooted at v in order to reach all of the destination nodes in the subtree if the message enters v on wavelength λ. Define $tr_v(\lambda) = \infty$ if wavelength λ is not available on link $(p(v), v)$ or it is not possible to reach all destinations in the subtree when v receives the message on wavelength λ.

For each leaf v in the tree

$$tr_v(\lambda) = \begin{cases} 0, & \text{if } \lambda \in w(p(v), v) \\ \infty, & \text{otherwise} \end{cases} \tag{9}$$

because a leaf node has no descendants and, thus, incurs no transmissions. Next, consider an internal nonroot node v. If $r(v) = 0$, node v cannot receive the message and, therefore, cannot employ transmitters; it may only distribute the message to its children on wavelength λ. Thus, the total number of transmitters used in this subtree is the sum of the number of transmitters used by the subtrees rooted at each of the children of v. Thus, for $r(v) = 0$

$$tr_v(\lambda) = \begin{cases} \sum\limits_{x \in C(v)} tr_x(\lambda), & \text{if } \lambda \in w(p(v), v) \\ \infty, & \text{otherwise.} \end{cases} \tag{10}$$

If $r(v) > 0$, then node v may receive the message and retransmit the message on up to $t(v)$ wavelengths in addition to λ. Then, if $r(v) > 0$

$$tr_v(\lambda) = \begin{cases} \min\limits_{A \in \mathcal{A}_{\lambda,\, t(v)}} \left\{ |A| - 1 + \sum\limits_{x \in C(v)} \min\limits_{\lambda' \in A} tr_x(\lambda') \right\}, & \text{if } \lambda \in w(p(v), v) \\ \infty, & \text{otherwise.} \end{cases} \tag{11}$$

Similarly, let T denote the minimum total number of transmitters and receivers over all WAs. Then

$$T = \min_{B \in \mathcal{B}_{t(s)}} \left\{ |B| + \sum_{x \in C(s)} \min_{\lambda' \in B} tr_x(\lambda') \right\}. \tag{12}$$

Finally, in Algorithm 1, $m_v(\lambda)$ and M are replaced by $tr_v(\lambda)$ and T, respectively, and (1)–(4) are replaced by (9)–(12), respectively. The running time is easily verified to be the same as that of the original algorithm.

C. *Wavelength Assignment for* $\ell > 1$

As illustrated in the example in Fig. 2, a WA may not exist when each link is permitted to send the message on only one wavelength but may exist when more than one wavelength may be used per link. In this section, we show how the algorithm described in Section IV-B can be generalized for the case that $\ell > 1$. Algorithms for optimal multicast problems, such as those described in Section IV-B, can be generalized in a similar manner.

For a given value of ℓ, let $\mathcal{S}_\ell$ denote the set of all nonempty subsets of Λ of size ℓ or less. We generalize the definition of function m_v as follows: $m_v(\mathcal{S}_\ell) \rightarrow \{\text{true, false}\}$ such that for each $S \in \mathcal{S}_\ell$, $m_v(S) = \text{true}$ if and only if all wavelengths in S are available on link $(p(v), v)$ and node v can deliver the message to all destinations in its subtree if it receives the message on all wavelengths in set S. From this definition it follows that for each leaf v in the tree

$$m_v(S) = \begin{cases} \text{true}, & \text{if } S \subseteq w(p(v), v) \\ \text{false}, & \text{otherwise.} \end{cases} \tag{13}$$

Next, consider an internal nonroot node v such that $r(v) = 0$.

$$m_v(S) = \begin{cases} \bigwedge_{x \in C(v)} \bigvee_{S' \subseteq S} m_x(S'), & \text{if } S \subseteq w(p(v), v) \\ \text{false}, & \text{otherwise.} \end{cases} \tag{14}$$

This rule asserts that v can deliver the message to all destinations in its subtree when it receives the message on all of the wavelengths in S if and only if all of the wavelengths in S are available on the link entering v from its parent and each child x of v can deliver the message to the descendants in its subtree when it receives the message on some subset S' of the wavelengths in S.

Next, consider the case that $r(v) > 0$. In this case, the message arriving at v on the wavelengths in S can be forwarded to any of its children on any of the wavelengths in S. In addition, node v can receive the message and retransmit the message to its children using up to $t(v)$ wavelengths other than those in S. Define a wavelength selection set with respect to S to be a subset of Λ that contains S. Let $\mathcal{A}_{S,c}$ denote the set of all wavelength selection sets with respect to S of size up to $c + |S|$. Thus, every set in $\mathcal{A}_{S,c}$ comprises the elements of S and up to c additional wavelengths. Then

$$m_v(S) = \begin{cases} \bigvee_{A \in \mathcal{A}_{S,\, t(v)}} \bigwedge_{x \in C(v)} \bigvee_{S' \subseteq A,\, |S'| \le \ell} m_x(S'), & \text{if } S \subseteq w(p(v), v) \\ \text{false}, & \text{otherwise.} \end{cases} \tag{15}$$

Finally, consider the case of the root node s. Unlike the other nodes in the tree, node s does not receive the message from a parent node. Instead, node s transmits the message using up to $t(s)$ different wavelengths. Recall that $\mathcal{B}_c$ denotes the set of all subsets of Λ of size less than or equal to c. Define $M = \text{true}$ if and only if a WA exists originating at the source node. Then

$$M = \bigvee_{B \in \mathcal{B}_{t(s)}} \bigwedge_{x \in C(s)} \bigvee_{S' \in B,\, |S'| \le \ell} m_x(S'). \tag{16}$$

The dynamic program shown in Algorithm 1 can now be applied to this case by replacing (1)–(4) with (13)–(16), respectively, and by replacing the inner **for** loop with a loop which iterates over all nonempty subsets $S \subseteq \Lambda$ such that $|S| \le \ell$. The running time of the algorithm remains $O(n)$ and the constant term can be derived in a fashion similar to that of the original algorithm.

V. Experimental Results

In this section, we describe experimental results using the algorithms presented in the previous section. We begin by describing results under the restriction that only one wavelength per channel can be used to transmit the message ($\ell = 1$). The first set of experiments used the WA algorithm described in Section IV-A to measure the number of multicast requests that were successfully realized as a function of the number of available wavelengths per link and number of available transmitters per node. Specifically, a single random multicast tree with 100 nodes ($n = 100$) was generated in which each node had a random number of children selected uniformly between 0 and 3 ($0 \le d_{\text{out}}(v) \le 3$). The generated tree had height 8 and the destination nodes comprised the 53 leaves of the tree. Each link was assumed to carry ten distinct wavelengths ($w = 10$). Very similar results to those reported below were obtained for other randomly generated multicast trees with other values of these parameters.

In one group of experiments the number of available transmitters per node was chosen at random from the uniform [0, 2] distribution and in the second group the uniform [1, 3] distribution was used. In all experiments, the number of available receivers per node was set to 1. In each group of experiments, the set of available wavelengths on each link was also selected at random where the size of the set was taken from the uniform $[x-1, x+1]$ distribution for a given value of x. For each value of x ranging from 2 to 9, 100 runs were performed. The data labeled "Exact Solution" in Fig. 4 shows the results of these experiments for the two groups of experiments.

In some situations, it may be desirable to use heuristics that are faster or simpler than the dynamic programs described here. The exact solutions found by the dynamic programming algorithms can then be used off-line to evaluate the quality of such heuristics. As an example, we have investigated a simple greedy heuristic for finding WAs. The heuristic operates as follows. The source node s determines the available wavelength that can be used to reach the largest number of its children, breaking ties arbitrarily. Then the available wavelength is found that reaches the largest number of remaining children. This process is repeated until a set S of wavelengths is found that can be used to reach all of the children of s. If the number of wavelengths in S exceeds the number of transmitters available at s, the heuristic fails to satisfy the multicast request and terminates. Otherwise, each child x of the source node may "choose" to receive the message on any one of the wavelengths in $S \cap w(s, x)$. For each child x of s, the heuristic determines which $\lambda \in S \cap w(s, x)$ reaches the largest number of children of x. This wavelength is then used to deliver the message from s to x and then from x to as many of its children as possible. Next, the heuristic repeatedly selects the wavelength that can be used to reach the largest number of remaining children of x until all children of x are reachable with the selected

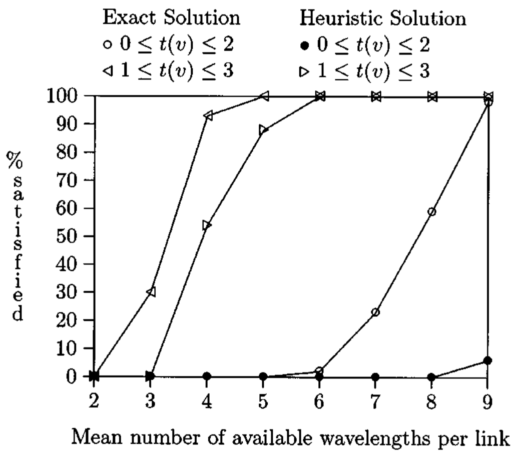

Fig. 4. Percentage of multicast requests satisfied as a function of the mean number of available wavelengths.

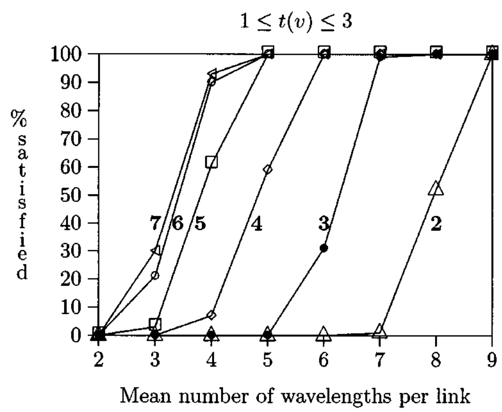

Fig. 5. Percentage of multicast requests satisfied using at most h hops as a function of the mean number of available wavelengths. Curves are labeled with h.

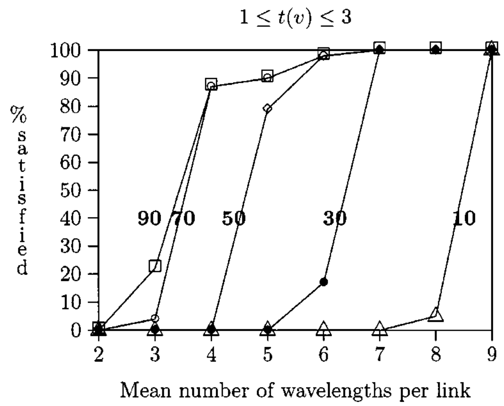

Fig. 6. Percentage of multicast requests satisfied using at most t transmitters and receivers, other than receivers at destination nodes, as a function of the mean number of available wavelengths. Curves are labeled with t.

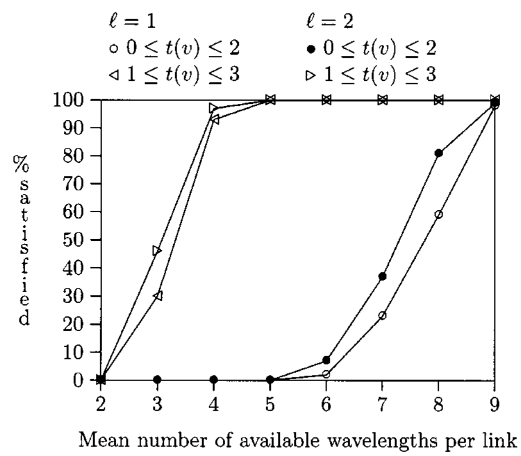

Fig. 7. Percentage of multicast requests satisfied using the optimal dynamic programming solutions for $\ell = 1$ and $\ell = 2$.

wavelengths. If the number of wavelengths selected is larger than the number of transmitters at x, then the heuristic fails to satisfy the request and terminates. Otherwise, this process is repeated until all destination nodes are reached.

This heuristic has $O(n)$ running time but has a significantly smaller constant term than that of the dynamic program. The results of running this greedy heuristic for the data used above are shown in Fig. 4 for comparison with the exact solutions obtained using the dynamic programming algorithm. Although the exact solutions are generally better than those found by the heuristic, the data also indicates that for some cases the heuristic performs very well.

Next, the dynamic programming formulation from Section IV-B1 was used to measure the number of hops required for the same parameters used in the above experiments. The results are shown in Fig. 5 for the case that the number of transmitters was selected from the uniform [1, 3] distribution. Each curve labeled with a value h indicates the percentage of multicast requests satisfied using at most h hops from the source to any destination. We note that for this data set no multicast request could be satisfied using fewer than two hops and no multicast request required more than seven hops.

Next, a dynamic programming formulation similar to the one described in Section IV-B2 was used to measure the total number of transmitters and receivers required for the same parameters used in the above experiments.[3] The results are shown in Fig. 6 for the case that the number of transmitters was selected from the uniform [1, 3] distribution. Each curve labeled with a value t indicates the percentage of multicast requests satisfied using at most t transmitters and receivers, except that the 53 receivers required at the 53 destination nodes are not included in this count.

Finally, as indicated by the example in Fig. 2, the probability of blocking may be reduced by allowing each link to use more than one wavelength to transmit a multicast message. In the last set of experiments, shown in Fig. 7, the dynamic programming formulation described in Section IV-C was used to investigate the benefits of increasing the value of ℓ from 1 to 2 for the same data set as used in Fig. 4. The improvement in the number of satisfied requests for $\ell = 3$ over $\ell = 2$ was negligible and is, therefore, not shown here.

VI. Conclusion and Future Research

In this paper, we have investigated the problems of multicast RWA. We have shown that the RWA problem is, in general,

[3]The dynamic program in Section IV-B2 minimizes the total number of transmitters used. Here, we used a generalization of that dynamic program to minimize the sum of the number of transmitters and receivers used.

NP-complete. We have also shown that the WA problem for any fixed multicast tree can be solved in linear time in the number of nodes when the number of wavelengths per link, transmitters and receivers per node, and switch degree are constants. Moreover, we have demonstrated that the dynamic programming algorithm for the WA problem can be adapted to find optimal WAs for different measures of optimality such as minimizing the number of hops or minimizing the number of transmitters used. The algorithms described in this paper can be used either to find exact solutions to the WA problem or to evaluate solutions found by faster and simpler heuristics.

ACKNOWLEDGMENT

The authors gratefully acknowledge the assistance of A. Fineman who implemented the algorithms described here and performed the computational experiments. The authors also thank the anonymous referees who made a number of suggestions which significantly improved this paper.

REFERENCES

[1] M. Ali and J. Deogun, "Power-efficient design of multicast wavelength-routed networks," *IEEE J. Select. Areas Commun.*, vol. 18, pp. 1852–1862, Oct. 2000.
[2] R. A. Barry and P. Humblet, "Models of blocking probability," *IEEE J. Select. Areas Commun.*, vol. 14, pp. 858–867, June 1996.
[3] J.-C. Bermond, L. Gargano, S. Perennes, A. Rescigno, and U. Vaccaro, "Efficient collective communication in optical networks," presented at the Annu. Int. Colloq. Automata, Languages and Programming (ICALP), Paderborn, Germany, July 1996.
[4] T. H. Cormen, C. E. Leiserson, and R. L. Rivest, *Introduction to Algorithms*. New York: McGraw-Hill, 1990.
[5] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York: Freeman, 1979.
[6] M. Kovačević and A. S. Acampora, "Electronic wavelength translation in optical networks," *J. Lightwave Technol.*, vol. 14, pp. 1161–1169, June 1996.
[7] K.-C. Lee and V. O. K. Li, "A wavelength-convertible optical network," *J. Lightwave Technol.*, vol. 11, pp. 962–970, May/June 1993.
[8] W. Liang and H. Shen, "Multicasting and broadcasting in large WDM networks," in *Proc. 12th Int. Parallel Processing Symp. (IPPS/SPDP)*, Orlando, FL, 1998, pp. 516–523.
[9] R. Malli, X. Zhang, and C. Qiao, "Benefit of multicasting in all-optical networks," in *Proc. SPIE Conf. All-Optical Networks*, vol. 3531, Boston, MA, Nov. 1998, pp. 209–220.
[10] B. Mukherjee, "WDM-based local lightwave networks, Part II: Multihop systems," *IEEE Network*, vol. 6, pp. 20–32, July 1992.
[11] ——, "WDM-based local lightwave networks, Part I: Single-hop systems," *IEEE Network*, vol. 6, pp. 12–27, May 1992.
[12] R. K. Pankaj, "Wavelength requirements for multicasting in all-optical networks," *IEEE/ACM Trans. Networking*, vol. 7, pp. 414–424, June 1999.
[13] 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.
[14] G. Sahin and M. Azizoglu, "Multicast routing and wavelength assignment in wide-area networks," in *Proc. SPIE Conf. All-Optical Networks*, vol. 3531, Boston, MA, Nov. 1998, pp. 196–208.
[15] H. Shen, F. Chin, and Y. Pan, "Efficient fault-tolerant routing in multihop optical WDM networks," *IEEE Trans. Parallel Distrib. Syst.*, vol. 10, pp. 1012–1025, Oct. 1999.
[16] S. B. Tridandapani and J. S. Meditch, "Supporting multipoint connections in multihop WDM optical networks," *J. High-Speed Networks*, vol. 4, no. 2, pp. 169–188, 1995.
[17] S. B. Tridandapani and B. Mukherjee, "Channel sharing in multihop WDM lightwave networks: Realization and performance of multicast traffic," *IEEE J. Select. Areas Commun.*, vol. 15, pp. 488–500, Apr. 1997.
[18] J. Yates, M. Rumsewicz, and J. Lacey. (1999) "Wavelength converters in dynamically reconfigurable WDM networks". *IEEE Commun. Surveys* [Online]. Available: http://www.comsoc.org/livepubs/surveys/index.html
[19] X. Zhang, J. Wei, and C. Qiao, "Constrained multicast routing in WDM networks with sparse light splitting," in *Proc. INFOCOM*, Tel Aviv, Israel, Mar. 2000, pp. 1781–1790.
[20] Z. Zhang and A. S. Acampora, "A heuristic wavelength assignment algorithm for multihop WDM networks with wavelength routing and wavelength reuse," *IEEE Trans. Networking*, vol. 3, pp. 281–288, June 1995.
[21] T. Znati, T. Alrabiah, and R. Melhem, "Point to multipoint path establishment schemes to support multicasting in WDM networks," in *Proc. 3rd IFIP Working Conference on Optical Network Design and Modeling (ONDM'99)*, Paris, France, Feb. 1999.

Ran Libeskind-Hadas (S'91–M'93) received the B.A. degree in applied mathematics from Harvard University, Cambridge, MA, in 1987 and the M.S. and Ph.D. degrees in computer science from the University of Illinois at Urbana-Champaign in 1989 and 1993, respectively.

Since 1993, he has been with the Department of Computer Science, Harvey Mudd College, Claremont, CA. He held the Iris and Howard Critchell Assistant Professorship and is currently an Associate Professor. His research interests are in multicast routing algorithms and optical networking.

Dr. Libeskind-Hadas is a Member of the Editorial Board of the IEEE TRANSACTIONS ON COMPUTERS.

Rami Melhem (S'82–M'83–SM'97–F'00) received the B.E. degree in electrical engineering from Cairo University, Cairo, Egypt, in 1976, the M.A. degree in mathematics and the M.S. degree in computer science from the University of Pittsburgh, Pittsburgh, PA, in 1981, and the Ph.D. degree in computer science from the University of Pittsburgh in 1983.

He was an Assistant Professor at Purdue University, West Lafayette, IN, prior to joining the faculty of the University of Pittsburgh in 1986, where he is currently a Professor of computer science and electrical engineering and the Chair of the Computer Science Department. He is the editor for the Plenum Book Series in Computer Science and is on the editorial board of the Computer Architecture Letters. He has served on program committees of numerous conferences and workshops and was the General Chair for the Third International Conference on Massively Parallel Processing Using Optical Interconnections. His research interests include optical interconnection networks, real-time and fault-tolerant systems, high-performance computing and parallel computer architectures.

Dr. Melhem was on the Editorial Boards of the IEEE TRANSACTIONS ON COMPUTERS and the IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS, and currently serves on the advisory boards of the IEEE technical committees on Parallel Processing and on Computer Architecture. Dr. Melhem is a Member of the Association of Computing Machinery.