1-1-2001

Multicast Communication in Circuit-Switched Optical Networks

Ran Libeskind-Hadas
Harvey Mudd College

Rami Melhem
University of Pittsburgh - Main Campus

Multicast Communication in Circuit-Switched Optical Networks

Ran Libeskind-Hadas Rami Melhem

Abstract—**In this paper we examine the problem of multicast routing in Wavelength-division multiplexed (WDM) optical networks. In particular, we examine wavelength and routing assignment problems in circuit-switched WDM networks. We show that although the routing and wavelength assignment (RWA) problem is NP-complete in general, the wavelength assignment (WA) problem can be solved in polynomial time.**

Keywords— **Optical networks, WDM, multicast communication, circuit-switching.**

I. Introduction

In Wavelength-division multiplexed (WDM) optical networks the fiber bandwidth is partitioned into multiple data channels which may be transmitted simultaneously on different wavelengths. In *single-hop* (or *all-optical*) WDM 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 [9]. 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 [12].

In *multi-hop* 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. Each conversion of the message from one wavelength to another is called a *hop*. Multi-hop networks have been shown to enjoy higher utilization of bandwidth and lower probability of blocking than single-hop networks [1].

Finally, a network may support *unicast* (or *one-to-one*) communication as well as *multicast* (or *one-to-many*) communication. Multicast communication is used in distributed memory parallel computers to support operations such as barrier synchronization and cache invalidation [10]. In wide-area networks multicast communication is used in video distribution and teleconferencing among other applications [7]. We note that both unicast and broadcast communication operations are special cases of multicast.

In this paper we consider multicast communication in multi-hop circuit-switched networks. We assume 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 communication request 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.

The *routing and wavelength assignment (RWA)* problem 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 *wavelength assignment (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 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.

Various aspects of multicasting in WDM networks have been investigated recently for both packet- and

Ran Libeskind-Hadas is with the Department of Computer Science, Harvey Mudd College, Claremont, CA 91711. This work was supported by the National Science Foundation under grant CCR-9900491 while the author was a visitor at the University of Pittsburgh.

Rami Melhem is with the Department of Computer Science, University of Pittsburgh, Pittsburgh, PA, 15260.

circuit-switched networks [5], [7], [8], [9], [11], [13]. In work most closely related to the results described here, Kovačević and Acampora [4] have investigated the WA problem for multi-hop unicast routing in circuit-switched meshes and Sahasrabuddhe and Mukherjee [8] have formulated the RWA problem for multi-hop multicast routing in packet-switched networks as a mixed-integer linear programming problem.

The remainder of this paper is organized as follows. In Section 2 we formally describe the model under consideration and define notation. In Section 3 we investigate the complexity of the routing and wavelength assignment problem. In Section 4 we give a linear time algorithm for the wavelength assignment problem. Conclusions are given in Section 5.

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.

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

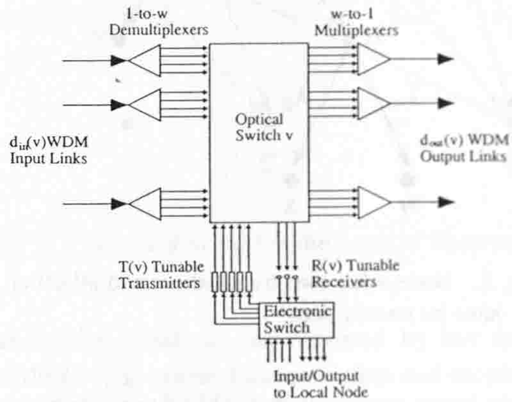

Fig. 1. Schematic of switch model.

Figure 1. Switches with some similar characteristics were described by Kovačević and Acampora [4] and by Sahasrabuddhe and Mukherjee [8]. We note that the results described in this paper can be adapted to a number of other switch models.[1]

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 connection requests are made and released dynamically. At the time that a particular multicast connection 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 and 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 existing connections, or these limits may be imposed in order to reduce cost

[1]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 port can be optically routed directly to a receiver/transmitter pair for retransmission on a new wavelength. This model is similar to the share-per-node switch architecture described in [5].

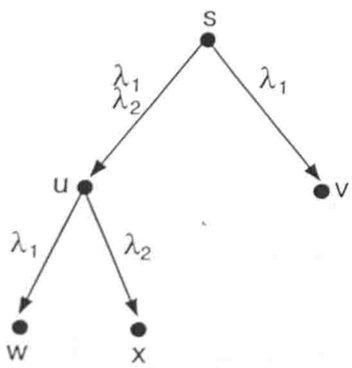

Fig. 2. Node s has two transmitters and all other nodes have no transmitters.

or to leave resources available for subsequent connection requests.

Due to these resource constraints, it may not be possible to realize a multicast connection 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. Figure 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 over each link. 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 may also increase the probability that no routing can be found.

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. A *routing and wavelength assignment (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 that $t(v)$ 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 subtree of G with root s and containing all vertices in D. A *wavelength assignment (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 in (v, x) in τ, and no more that $t(v)$ transmitters and $r(v)$ receivers are used at each node $v \in V$.

In many cases, we wish to find a RWA or WA that is optimal with respect to a given measure. For example, we may wish to find a routing that minimizes the maximum number of hops from the source to all destinations. Alternatively, to minimize use of resources, we may wish to find a routing that minimizes the total number of receivers and transmitters used. These types of "optimal" multicast problems are considered in [6].

In the next section we show that the problem of finding a RWA can be solved in polynomial-time for a special case but, in general, is NP-complete. We then show that the problem of finding a WA can be solved in linear time.

III. The 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 $t(v) \geq W(v)$ and $r(v) \geq 1$ for all $v \in V$ then a 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 $t(v) \geq W(v)$, 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. Thus, without loss of generality we may restrict attention to the case that $\ell = 1$.

Apply 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 a RWA exists. Otherwise, no RWA exists. The running time of 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, a 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 a RWA is NP-complete.

Theorem 2: For any $\ell \geq 1$, if $t(v) < W(v)$ for some $v \in V$ then the problem of determining if there exists a 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 5 clauses. This restricted version of 3SAT is known to be NP-complete [3]. 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 Figure 3.

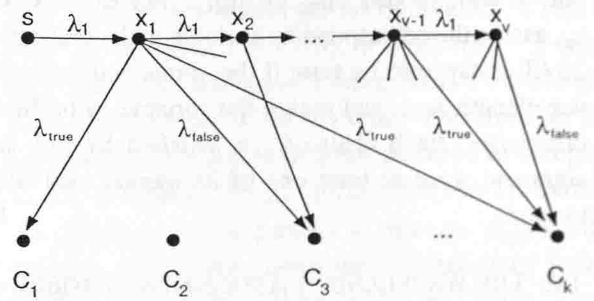

Fig. 3. Construction used in the proof of Theorem 2.

We claim that a 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 which 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.

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 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. The Wavelength Assignment Problem

In this section we show that the wavelength assignment 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 non-destination 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, $r(v) > 0$. (Otherwise no wavelength assignment exists.)
4. For each node v in the multicast tree, $t(v) \leq \min\{w, d_{\text{out}}(v) \times \ell\}$. (A node cannot transmit a total of more than w distinct wavelengths nor can it transmit more than ℓ distinct wavelengths to each of its children, of which there are at most $d_{\text{out}}(v)$.)

We begin by examining the case that $\ell = 1$. We then generalize this result to the case $\ell > 1$.

A. Wavelength Assignment for $\ell = 1$

The algorithm is based on dynamic programming. For each non-root node v, let $p(v)$ denote the parent of v in the given multicast tree. For each non-root node v, define the predicate $m_v(\Lambda) \rightarrow \{\textbf{true}, \textbf{false}\}$ by $m_v(\lambda) = \textbf{true}$ if and only if node v can deliver the message to all destinations in its subtree if it receives the message on wavelength λ. From this definition it follows that for each leaf v in the tree,

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

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

Next, consider an internal non-root node v which has no receivers available. Since $r(v) = 0$, node v may forward the message on the incoming wavelength to its children 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 $r(v) = 0$,

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

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, upon receipt of the message on wavelength λ, all children of v can deliver the message to all destinations in their respective subtrees.

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 set $A \subseteq \Lambda$ such that $\lambda \in A$. Let $\mathcal{A}_{\lambda,c}$ denote the set of all wavelength selection sets with respect to λ of size $c + 1$. Thus, every set in $\mathcal{A}_{\lambda,c}$ comprises λ and c additional wavelengths. Then $m_v(\lambda) =$

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

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 set A comprising λ and $t(v)$ 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 c. Define $M = \textbf{true}$ if and only if a WA exists originating at the source node. Then,

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

This rule is analogous to the one in Equation (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 [2]. 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.

Compute a topological ordering $v_{i_1}, \ldots, v_{i_n}$ 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 compute $m_v(\lambda)$ using Eq. (1)
   if $j > 1$ **and** $r(v) = 0$ compute $m_v(\lambda)$ using Eq. (2)
   if $j > 1$ **and** $r(v) > 0$ compute $m_v(\lambda)$ using Eq. (3)
   if $j = 1$ **then** compute M using Equation (4)
  end for (*Comment: End inner for loop*)
end for (*Comment: End outer for loop*)
return (M)

Algorithm 1

Note that the actual WA can be found, if one exists, by recording the wavelength assignments 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 $O(n + m)$ in where $n = |V|$ and $m = |E|$ [2]. Since we are assuming that the degree of each node is upper-bounded by a constant, $m \in O(n)$ and thus the ordering can be computed in time $O(n)$.

There are a total of wn iterations through the nested **for** loops. Among the computations performed inside the **for** loops, the computation in Equation (3) requires the largest number of steps. An upper-bound on the number of steps required to compute $m_v(\lambda)$ in Equation (3) can be derived as follows: For each wavelength λ there are $\binom{w-1}{t(v)}$ distinct wavelength selection sets. For each wavelength selection set A, consider the set of children, $C(v)$, of node v. 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 $\binom{w}{t(v)} d_{\text{out}}(v)(t(v)+1)$. Letting $t = \max_{v \in V} t(v) + 1$ 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 $[w\binom{w}{t}dt]n$. Thus, while the algorithm has $O(n)$ running time, the constant term depends on constants w, t, and d.

B. Wavelength Assignment for $\ell > 1$

As illustrated in the example in Figure 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 subsection we show how the algorithm described above can be generalized for the case that $\ell > 1$.

For a given value of ℓ, let $\mathcal{S}_\ell$ denote the set of all non-empty subsets of Λ of size ℓ or less. We generalize the definition of function m_v as follows: $m_v(\mathcal{S}_\ell) \rightarrow \{\textbf{true}, \textbf{false}\}$ such that for each $S \in \mathcal{S}_\ell$, $m_v(S) =$ **true** if and only if all destinations in the subtree rooted at v can receive the message when v receives the message on all of the wavelengths in set S. Note that $m_v(S) =$ **false** if $S \not\subseteq w(p(v), v)$. From this definition it follows that for each leaf v in the tree,

$$m_v(S) = \begin{cases} \textbf{true} & \text{if } S \subseteq w(p(v), v) \\ \textbf{false} & \text{otherwise} \end{cases} \quad (5)$$

Next, consider an internal non-root node v such that $r(v) = 0$. Then $m_v(S) =$

$$\begin{cases} \bigwedge_{x \in C(v)} m_x(S \cap w(v, x)) & \text{if } S \subseteq w(p(v), v) \\ \textbf{false} & \text{otherwise} \end{cases} \quad (6)$$

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 the wavelengths in $S \cap w(v, x)$. Note that $S \cap w(v, x)$ is the set of wavelengths in S that are available on the link from v to x.

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 set $A \subseteq \Lambda$ such that $S \subseteq A$. Let $\mathcal{A}_{S,c}$ denote the set of all wavelength selection sets with respect to S of size $c + |S|$. Thus, every set in $\mathcal{A}_{\lambda,S}$ comprises the elements of S and c additional wavelengths. Then, $m_v(S) =$

$$\bigvee_{A \in \mathcal{A}_{S,t(v)}} \bigwedge_{x \in C(v)} \bigvee_{S' \subseteq A, |S'| = \ell} m_x(S' \cap w(v,x)) \quad (7)$$

if $S \subseteq w(p(v), v)$ and **false** otherwise. This rule asserts that $m_v(S)$ is **true** if and only if all of the wavelengths in S are available on the link entering v from its parent and there exists some set A comprising the wavelengths in S and $t(v)$ additional wavelengths (to be transmitted at v) with the following property: For every child x of v there exists a subset $S' \subseteq A$ of size ℓ such that x can deliver the message to all descendant destinations if it receives the message on the wavelengths in $S' \cap w(v,x)$.

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 c. Define $M =$ **true** if and only if a WA exists originating at the source node. Then,

$$M = \bigvee_{B \in \mathcal{B}_{t(s)}} \bigwedge_{x \in C(s)} \bigvee_{S' \in B, |S'| = \ell} m_x(S' \cap w(v,x)) \quad (8)$$

The dynamic program shown in Algorithm 1 can now be applied to this case by replacing Equations (1), (2), (3), and (4) with equations (5), (6), (7), and (8), respectively, and by replacing each loop which iterates over single wavelengths λ by a loop which iterates over all non-empty subsets $S \subseteq \Lambda$ such that $|S| \leq \ell$. The running time of the algorithm remains $O(n)$ with a the constant term of at most $2^\ell \binom{w}{t} d \binom{t+\ell}{\ell}$.

V. Conclusion

In this paper we have investigated the routing and wavelength assignment (RWA) and wavelength assignment (WA) problems for multicast in circuit-switched optical networks. We have shown that although the RWA problem is, in general, NP-complete, the WA problem can be solved in linear time using dynamic programming. The reader is referred to [6] where we have investigated variants of this algorithm for finding "optimal" multicasts and report preliminary experimental results using these algorithms.

References

[1] R. A. Barry and P. Humblet. Models of blocking probability. *IEEE Journal on Selected Areas in Communication*, 14(5):858–867, June 1996.

[2] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms*. McGraw-Hill, New York, 1990.

[3] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, New York, 1979.

[4] M. Kovačević and A. S. Acampora. Electronic wavelength translation in optical networks. *Journal of Lightwave Technology*, 14(6):1161–1169, June 1996.

[5] K.-C. Lee and V. O. K. Li. A wavelength-convertible optical network. *Journal of Lightwave Technology*, 11(5/6):962–970, May/June 1993.

[6] R. Libeskind-Hadas and R. Melhem. Multicast routing and wavelength assignment in multi-hop optical networks. In *Proceedings of the 2001 International Conference on Networking*, 2001.

[7] R. K. Pankaj. Wavelength requirements for multicasting in all-optical networks. *IEEE/ACM Transactions on Networking*, 7(3):414–424, June 1999.

[8] L. H. Sahasrabudde and B. Mukherjee. Light-trees: Optical multicasting for improved performance in wavelength-routed networks. *IEEE Communications Magazine*, 37(2):67–73, February 1999.

[9] H. Shen, F. Chin, and Y. Pan. Efficient fault-tolerant routing in multihop optical WDM networks. *IEEE Transactions on Parallel and Distributed Systems*, 10(10):1012–1025, October 1999.

[10] R. Sivaram, R. Kesavan, D. Panda, and C. Stunkel. Where to provide support for efficient multicasting in irregular networks: Network interface or switch? In *Proc. of the 1998 Int. Conf. on Parallel Processing*, pages 452–459, August 1998.

[11] S. B. Tridandapani and B. Mukherjee. Channel sharing in multi-hop WDM lightwave networks: Realization and performance of multicast traffic. *IEEE Journal on Selected Areas in Communications*, 15(3):488–500, April 1997.

[12] J. M. Yates, M. P. Rumsewicz, and J. P. R. Lacey. Wavelength converters in dynamically reconfigurable WDM networks. *IEEE Communications Surveys*, 1999.

[13] Z. Zhang and A. S. Acampora. A heuristic wavelength assignment algorithm for multihop WDM networks with wavelength routing and wavelength reuse. *IEEE Journal of Networking*, 3:281–288, 1995.