Multicast Routing and Wavelength Assignment in Multi-Hop Optical Networks*

Ran Libeskind-Hadas[1] and Rami Melhem[2]

[1] Dept. of Computer Science, Harvey Mudd College, Claremont CA 91711, USA
hadas@cs.hmc.edu
[2] Dept. of Computer Science, University of Pittsburgh, Pittsburgh PA 15260, USA
melhem@cs.pitt.edu

Abstract. This paper addresses multicast routing in multi-hop optical networks employing wavelength-division multiplexing (WDM). 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 although the *routing and wavelength assignment* in this model is NP-complete, the *wavelength assignment* problem can be solved in linear time.

1 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 use of enormous fiber bandwidth by providing data channels whose individual bandwidths more closely match those of the electronic devices at their endpoints.

WDM networks can be classified as either *single-hop* or *multi-hop* networks [4, 5]. 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. 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 [9].

* 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. The authors also gratefully acknowledge the assistance of Mr. Adam Fineman who implemented the algorithms described here and performed the computational experiments.

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]. However, a multi-hop 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 multi-hop connections that minimize the number of transmitters and receivers and/or 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 is used in distributed shared memory clusters to support operations such as cache invalidation and is used in wide-area networks for video distribution and teleconferencing among other applications.

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 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.

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 although the RWA problem in this model is, in general, NP-complete, the WA problem can be solved in linear time. In addition, 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.

Various aspects of multicasting in WDM networks have been investigated recently for both packet- and circuit-switched networks [3, 6–8, 10]. In work most closely related to the results described here, Kovačević and Acampora [2] have investigated the WA problem for multi-hop unicast routing in circuit-switched meshes and Sahasrabuddhe and Mukherjee [6] 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 give a linear time algorithm for the wavelength assignment problem and generalize the algorithm to find "optimal" multicasts. Section 4 describes experimental results using these algorithms. Conclusions are given in Section 5.

2 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 Figure 1. Switches with some similar characteristics were described by Kovačević and Acampora [2] and by Sahasrabuddhe and Mukherjee [6]. We note that the results described in this paper can be adapted to a number of other switch models.

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 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. 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

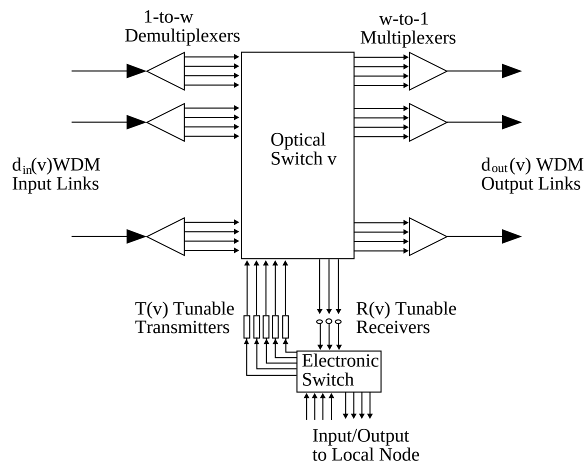

Fig. 1. Schematic of switch model.

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.

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*

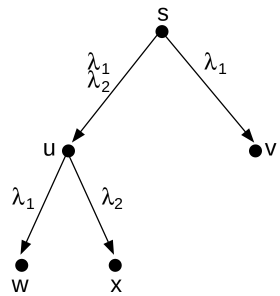

Fig. 2. Node s has two transmitters and all other nodes have no transmitters.

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.

The following two theorems show that although the RWA problem can be solved efficiently for some special cases, the RWA problem is, in general, NP-complete. The proofs of these results are omitted in the interest of space.

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 $O(n)$.*

Theorem 2. *For any $\ell \geq 1$, if $t(v) < W(v)$ for some nodes $v \in V$ then the problem of determining if there exists a RWA is NP-complete.*

3 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.)

We begin in Subsection 3.1 by examining the case that $\ell = 1$. In Subsection 3.2 we show how the algorithm can be adapted to find wavelength assignments that minimize the maximum number of hops from the source to all destinations.

3.1 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. Then $(p(v), v)$ denotes the link from the parent of v to v. Define the predicate $m_v(\Lambda) \rightarrow \{\mathbf{true}, \mathbf{false}\}$ by $m_v(\lambda) = \mathbf{true}$ 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,

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

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 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 $r(v) = 0$,

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

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 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

$$m_v(\lambda) = \bigvee_{A \in \mathcal{A}_{\lambda,t(v)}} \bigwedge_{x \in C(v)} \bigvee_{\lambda' \in A} m_x(\lambda') \tag{3}$$

if $\lambda \in w(p(v), v)$ and otherwise $m_v(\lambda) = \mathbf{false}$. 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 $M = \textbf{true}$ if and only if a WA exists originating at the source node. Then,

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

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. 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 **then** compute $m_v(\lambda)$ using Equation (1)
    if $j > 1$ **and** $r(v) = 0$ **then** compute $m_v(\lambda)$ using Equation (2)
    if $j > 1$ **and** $r(v) > 0$ **then** compute $m_v(\lambda)$ using Equation (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 $O(n+m)$ where $n = |V|$ and $m = |E|$. Since our model assumes 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 **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 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 4.

3.2 Optimal Multicast for $\ell = 1$

In this subsection we show that the dynamic programming solution described in the previous subsection can be adapted to find wavelength assignments which minimize the maximum number of hops required to reach all destination nodes. Similar adaptations can be made for other metrics of optimality.

For each non-root 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 $h_v(\lambda) = \infty$.

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 non-root 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 $r(v) = 0$,

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

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 its 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. Thus, for $r(v) > 0$,

$$h_v(\lambda) = \min_{A \in \mathcal{A}_{\lambda, t(v)}} \max_{x \in C(v)} \min_{\lambda' \in A} \begin{cases} h_x(\lambda') & \text{if } \lambda' = \lambda \\ 1 + h_x(\lambda') & \text{if } \lambda' \neq \lambda \end{cases} \tag{7}$$

if $\lambda \in w(p(v), v)$ and otherwise $h_v(\lambda) = \infty$.

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 Equations (1), (2), (3), and (4) are replaced by Equations (5), (6), (7), and (8), respectively. The asymptotic running time and constants are easily verified to be the same as that of the original algorithm.

3.3 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. The algorithms described above can be easily extended to handle the case that $\ell > 1$. The details are omitted in the interest of space.

4 Experimental Results

In this section we describe experimental results using the algorithms presented in the previous section. In the interest of space, we restrict our attention to the case that $\ell = 1$. The first set of experiments used the wavelength assignment algorithm described in Subsection 3.1 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 random multicast tree with 100 nodes ($n = 100$) was generated in which each node had between 0 and 3 children ($0 \le d_{\text{out}}(v) \le 3$). The generated tree had height 8 and the destination nodes comprised the 53 leaves of the tree. Each link was assumed to carry 10 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 Figure 3 shows the results of these experiments for the two groups of experiments.

We have noted that the dynamic programming algorithms run in time $O(n)$ but the constant term depends on the number of wavelengths, transmitters per node, and degree of the switches. For the experiments described above, the maximum amount of time required by the dynamic program for a multicast request was 0.11 seconds on a 450 MHz Pentium 2. However, for larger values of the parameters the running time was significantly larger. For example, for a problem instance with 100 nodes, 32 wavelengths per link, switches of degree 8, up to 3 transmitters available per node for each multicast request, and an average of half of the 32 wavelengths available on each link, the running time increased to 11.38 seconds.

Therefore, 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 wavelength assignments. 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 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 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 a significantly smaller constant term than that of the dynamic program. In comparison to the 11.38 seconds incurred by the dynamic program for the largest problem instance described above, this heuristic required only 0.01 seconds for the same data. The results of running this greedy heuristic for the data used above are shown in Figure 3 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. Other more sophisticated heuristics could also be considered at the expense of increased running time.

Next, the dynamic programming formulation from Subsection 3.2 was used to measure the number of hops required for the same parameters used in the above experiments. The results are shown in Figure 4 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 2 hops

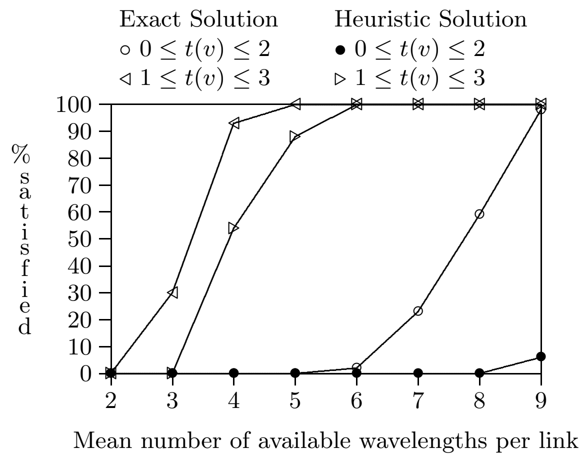

Fig. 3. Percentage of multicast requests satisfied as a function of number of available wavelengths.

and no multicast request required more than 7 hops. These results indicate the relationship between hop counts and percentage of satisfied requests.

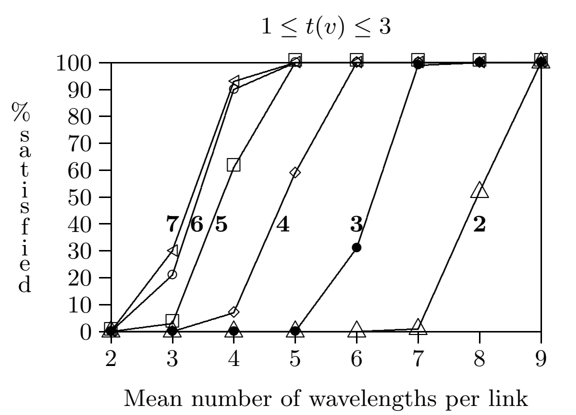

Fig. 4. Percentage of multicast requests satisfied using at most h hops as a function of the number of available wavelengths. Curves are labeled with h.

5 Conclusion and Future Research

In this paper we have investigated the problems of multicast routing and wavelength assignment. We have shown that the wavelength assignment problem for any fixed multicast tree can be solved in time linear 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 wavelength assignment problem can be adapted to find wavelength assignments that minimize the maximum number of hops from the source to all destinations. Similar adaptations can be made to find solutions that are optimal with respect to other metrics.

The algorithms described in this paper can be used either to find exact solutions to the wavelength assignment problem or to evaluate solutions found by faster and simpler heuristics. Heuristics for minimizing the maximum number of hops, transmitter and receiver usage, and other measures of optimality are currently under investigation.

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. M. Kovačević and A. S. Acampora. Electronic wavelength translation in optical networks. *Journal of Lightwave Technology*, 14(6):1161–1169, June 1996.
3. W. Liang and H. Shen. Multicasting and broadcasting in large WDM networks. In *Proceedings of the 12th International Parallel Processing Symposium (IPPS/SPDP)*, pages 516–523, 1998.
4. B. Mukherjee. WDM-based local lightwave networks, part II: Multi-hop systems. *IEEE Network*, 6(4):20–32, July 1992.
5. B. Mukherjee. WDM-based local lightwave networks, part I: Single-hop systems. *IEEE Network*, 6(3):12–27, May 1992.
6. 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.
7. S. B. Tridandapani and J. S. Meditch. Supporting multipoint connections in multi-hop WDM optical networks. *Journal of High-Speed Networks*, 4(2):169–188, 1995.
8. 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.
9. J. M. Yates, M. P. Rumsewicz, and J. P. R. Lacey. Wavelength converters in dynamically reconfigurable WDM networks. *IEEE Communications Surveys*, 1999.
10. 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.