Multiwavelength Optical Networks with Limited Wavelength Conversion*

R. Ramaswami

IBM Research
30 Saw Mill River Road
Hawthorne, NY 10532 USA
rajiv@watson.ibm.com

G. H. Sasaki

Dept. Elec. Engg., U. Hawaii
2540 Dole Street
Honolulu, HI 96822 USA
sasaki@spectra.eng.hawaii.edu

Abstract

This paper proposes optical wavelength division multiplexed (WDM) networks with limited wavelength conversion that can efficiently support lightpaths (connections) between nodes. Each lightpath follows a route in the network and must be assigned a channel along each link in its route. The load $\lambda_{\max}$ of a set of lightpath requests is the maximum over all links of the number of lightpaths that use the link. At least $\lambda_{\max}$ wavelengths will be needed to assign channels to the lightpaths. If the network has full wavelength conversion capabilities then $\lambda_{\max}$ wavelengths are sufficient to perform the channel assignment.

We propose ring networks with fixed wavelength conversion capability within the nodes that can support all lightpath request sets with load $\lambda_{\max}$ at most $W-1$, where W is the number of wavelengths in each link. We also propose ring networks with selective pairwise wavelength conversion capability within the nodes that can support all lightpath request sets with load $\lambda_{\max}$ at most W. We also propose a star network with fixed wavelength conversion capability at its hub node that can support all lightpath request sets with load $\lambda_{\max}$ at most W. We extend this result to tree networks and networks with arbitrary topologies. These results show that significant improvements in traffic-carrying capacity can be obtained in WDM networks by providing very limited wavelength conversion capability within the network.

1 Introduction

Wavelength Division Multiplexing (WDM) is an important approach to utilize the large available bandwidth in a single mode optical fiber. WDM is basically frequency division multiplexing in the optical frequency domain, where on a single optical fiber there are multiple communication channels at different *wavelengths* (corresponding to carrier frequencies). There has been a great deal of interest in WDM networks that employ wavelength routing. These networks support *lightpaths*, which are end-to-end circuit-switched communication connections that traverse one or more links and use one WDM channel per link. These lightpaths could serve as the physical communication links for a variety of high-speed networks such as ATM (Asynchronous Transfer Mode) networks.

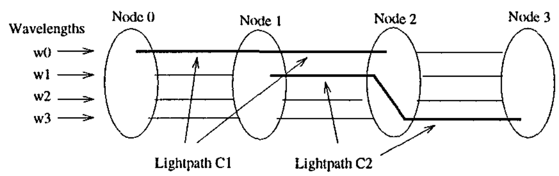

Figure 1: A network with wavelengths $\{\omega_0, \omega_1, \omega_2, \omega_3\}$. Channels are shown as lines between nodes.

An example of a WDM wavelength routing network is shown in Figure 1. It is composed of four nodes with three optical fiber links connecting nodes 0 and 1, nodes 1 and 2, and nodes 2 and 3. Each link has four WDM channels at wavelengths $\{\omega_0, \omega_1, \omega_2, \omega_3\}$. Switching is done at each node so that channels may be connected to form lightpaths. Note that if channels at different wavelengths are to be connected then *wavelength conversion* devices are needed that can shift the wavelength of an optical signal. For example, in Figure 1, lightpath $C1$ is composed of two WDM channels at wavelength ω_0 on links 1 and 2. Hence, it does not need a wavelength converter. However, lightpath $C2$ needs a converter at node 2 because it is composed of two WDM channels at different wavelengths (ω_1 and ω_3). The advantage of wavelength conversion is that WDM channels will be used more efficiently, but the disadvantage is increased cost and complexity.

1.1 Limited Wavelength Conversion

In this paper we will explore circuit-switched wavelength routing WDM network architectures that employ *limited wavelength conversion*, i.e., WDM channels have restrictions on the channels they may be connected to on other links. For example, Figure 2 shows a ring network with seven WDM channels per link, where the channels are at wavelengths $\{\omega_0, \omega_1, ..., \omega_6\}$. The lines within nodes show which pairs of channels

*This work was sponsored in part by grant MDA-972-95-C-0001 by ARPA and grant NCRI-93-96130 by NSF.

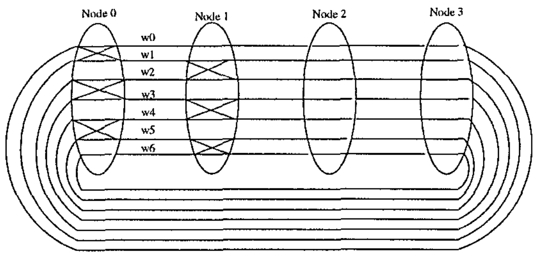

Figure 2: A ring network with wavelengths $\{\omega_0, \omega_1, ..., \omega_6\}$ to illustrate limited wavelength conversion. The lines between nodes represent channels, and lines within nodes indicate which channels may be connected.

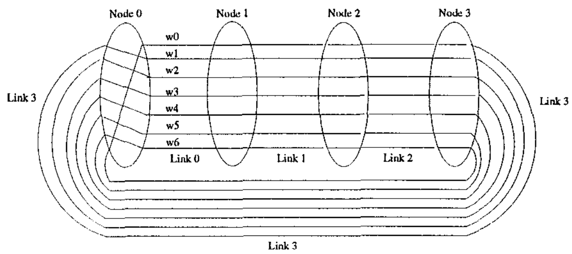

Figure 3: A ring network with wavelengths $\{\omega_0, \omega_1, ..., \omega_6\}$ to illustrate fixed wavelength conversion. The lines between nodes represent channels, and lines within nodes indicate which channels may be connected.

may be connected. Note that the network has some wavelength conversion, but with restrictions. For example, at node 0 channels at wavelengths $\{\omega_0, \omega_1\}$ may only be connected to other channels at wavelengths $\{\omega_0, \omega_1\}$. A special case of limited wavelength conversion is *fixed* wavelength conversion, where at each node, each channel may be connected to exactly one predetermined channel on every other link. Figure 3 shows an example ring network with fixed wavelength conversion. Here, the channel at ω_i in link 3 may be connected only to channel $\omega_{(i+1) \bmod W}$ in link 0 for $i = 0, 1, ..., W-1$. At all other nodes, only channels with the same wavelength are connected.

Networks with limited/fixed wavelength conversion will be less costly to implement than networks without restrictions on wavelength conversion (i.e., having full wavelength conversion capability), but may still provide enough conversion to use channels efficiently. The different types of conversion possible within the node are illustrated in Figure 4. Within each node the wavelength conversion can be done all-optically or by receiving the signal, switching it electronically and retransmitting it on another wavelength (O-E-O). The all-optical approach uses optical wavelength converter devices. In some of these devices, such as those based on four-wave mixing [27], the conversion efficiency is a strong function of the input and output wavelengths, naturally leading to limited conversion capability. Even otherwise we can save on the num-

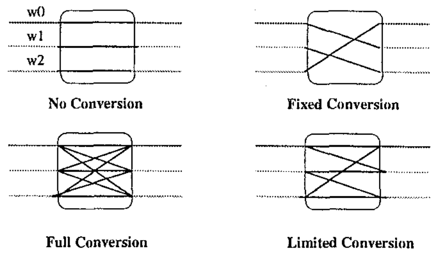

Figure 4: Different types of wavelength conversion.

ber of such devices required in the node. In the O-E-O approach we can implement limited conversion using much fewer electronic switches than would be needed for full conversion.

1.2 Network Model

We assume that the links and WDM channels are bidirectional (or full duplex). Network nodes are connected by fiber optic links, and for simplicity it is assumed that all pairs of nodes have at most one link between them. Each link has W WDM bidirectional channels at wavelengths $\{\omega_0, \omega_1, ..., \omega_{W-1}\}$, where $\omega_0 < \omega_1 < ... < \omega_{W-1}$.

Each node has switching capability to connect WDM channels to form full duplex lightpaths. The switching capability will determine which pairs of channels may be connected to one another. We will refer to two channels that may be connected to one another as being *attached*. For example, at node 0 in Figure 2, channels at wavelengths $\{\omega_0, \omega_1\}$ are attached, channels at $\{\omega_2, \omega_3\}$ are attached, channels at $\{\omega_4, \omega_5\}$ are attached, and channels at ω_6 are attached. A node has *wavelength degree* k (for some integer $k > 0$) if for each pair of incident links, each channel in one link is connected to at most k other channels in the other link. For example, node 0 in Figure 2 has wavelength degree two. A node has *full wavelength conversion* if its wavelength degree is W. A node is said to have *fixed wavelength conversion* if its wavelength degree is one (e.g., see Figure 3). Note that a node with no wavelength conversion has wavelength degree one. (Again these different types of conversion are illustrated in Figure 4.)

The network supports sets of lightpaths. A lightpath is specified by a path in the network that is referred to as a *route*. A lightpath is realized by a set of channels, one on each link along its route so that channels that are incident to a common node are attached at the node. Such a set of channels is referred to as a *channel assignment* for the route. This realization allows communication signals to be sent on a lightpath between the ends of the route by having them transported along attached channels.

A set of lightpaths is specified by a set of routes, one route per lightpath. A set of routes will be referred to as a *request*. A *channel assignment* for a request is a collection of channel assignments, one per route of the request, such that each channel is assigned to at most one route, i.e., no two routes share a channel. Note that a channel assignment for a request realizes the lightpaths corresponding to the request. An important parameter of a request is its *load*, which is the value $\lambda_{\max} = \max_{e \in E} \lambda_e$, where λ_e denotes the number of routes using link e and E denotes the set of links in the network. Clearly at least $\lambda_{\max}$ wavelengths are needed for a channel assignment for a request with load $\lambda_{\max}$.

1.3 Organization

In this paper, we propose ring and star networks with limited wavelength conversion to support sets of lightpaths efficiently. In Section 2, we discuss our results for ring networks. We give a ring network with one node having fixed wavelength conversion and the rest of the nodes with no wavelength conversion such that all requests with load $\lambda_{\max} \leq W - 1$ have channel assignments. We also give a ring network with two nodes with wavelength degree two and the rest of the nodes with no wavelength conversion such that all requests with load $\lambda_{\max} \leq W$ have channel assignments. Note that the networks that have channel assignments for all requests with load at most W utilize the channels as efficiently as networks with full wavelength conversion at all nodes. Also note that the first deployed WDM networks are likely to be rings, as seen from several recent testbeds (see for example [5, 23]).

In Section 3, we discuss our results for star networks as well as extensions to tree networks, and networks with arbitrary topologies where route lengths are at most two. We present a star network that has fixed wavelength conversion and has channel assignments for all requests with load $\lambda_{\max} \leq W$, when W is an even number. In Section 4, we provide conclusions and discuss how our results can be extended when links and channels are *directed.*

Note that we consider the problem of finding channel assignments for sets of lightpaths all at one time. Thus, if a new lightpath is to be included to an existing set of lightpaths (while keeping the same routes for the lightpaths), the channel assignments for all lightpaths may have to be recomputed. In this sense, the channel assignment is done *offline.* There is also the more practical consideration of *online* channel assignment, i.e., setting up new lightpaths without changing the assignment for existing lightpaths. Although we only consider the offline case, we believe that its understanding can lead to fundamental insights to the online case, just as understanding *rearrangeable nonblocking networks* can help to understand efficient *wide-sense* and *strict-sense nonblocking networks* [13]. Also, offline channel assignment will be more efficient in utilizing channels than online channel assignment.

1.4 Related Work

Previous work focuses primarily on networks with either no wavelength conversion or networks with full wavelength conversion (i.e., any pair of WDM channels may be connected).

The joint lightpath routing and channel assignment problem in networks without wavelength conversion is known to be NP-complete [6] and remains NP-complete even for rings [8]. Given a routing already, an algorithm that finds a channel assignment in a ring network without wavelength conversion if $2\lambda_{\max} - 1 \leq W$ is given in [24] as well as [19]. An algorithm that finds channel assignments in a tree network without wavelength conversion if $\frac{3}{2}\lambda_{\max} \leq W$ is given in [19]. Sample requests can easily be constructed for these networks that require $W = 2\lambda_{\max} - 1$ wavelengths and $W = \frac{3}{2}\lambda_{\max}$ wavelengths for rings and stars respectively. [18] gives algorithms that find channel assignments for the case of a directed network without wavelength conversion and directed lightpath requests for trees, if $15\lambda_{\max}/8 \leq W$, and for rings, if $2\lambda_{\max} \leq W$. Several heuristic channel assignment schemes have also been proposed for networks without wavelength conversion [3, 21, 6, 15, 4, 25].

Variants of the limited conversion model are considered in [16, 17, 26, 22]. In [16, 17] it is assumed that each node has a limited number of wavelength converters and that each converter has no restrictions on the wavelengths of the channels it can connect. Here, the restriction is on the *number* of wavelength conversions at a node. In [26], a network with limited wavelength conversion is used to study the performance due to limited wavelength shifting capability of devices based on four wave mixing. The converters allow wavelengths to be shifted within a given range. Also, the work in [22] studies *sparse wavelength conversion*, where networks are comprised of a mix of nodes having full and no wavelength conversion. The channel assignment in these papers [16, 17, 26, 22] are simple heuristics, and their performance analyses are based upon probabilistic models and techniques (i.e., compute blocking probabilities of setting up lightpaths) which may not be as appropriate for networks that require channels to be highly utilized.

There are some recent results on the *online* channel assignment problem for ring and tree networks [9, 11], where lightpath requests arrive and leave the network dynamically. The problem of recovering from link and node faults in ring networks using limited wavelength conversion is addressed in [10].

2 Rings

In this section we will consider ring networks. Without loss of generality, it will be assumed that a ring network has a *clockwise direction* (and *counterclockwise direction*) as shown in Figure 5. Its nodes are numbered $0, 1, ..., N-1$ consecutively in the clockwise direction, where N denotes the number of nodes. The links are also numbered $0, 1, ..., N-1$ in the clockwise direction such that for each $i = 0, 1, ..., N-1$, the link between node i and node $(i+1) \bmod N$ is numbered i (see Figure 5).

Most of the results of this section assume that there is a collection of lightpaths to be set up, and their set of routes (i.e., a *request*) is already given. However, we should note that for the ring network, there is an algorithm that can compute minimum load requests for sets of lightpaths, specified by their terminating nodes [7]. However, simple shortest path routing does not perform very poorly as shown below.

Theorem 1 *Suppose we are given a request of source-destination pairs and the minimum possible load for*

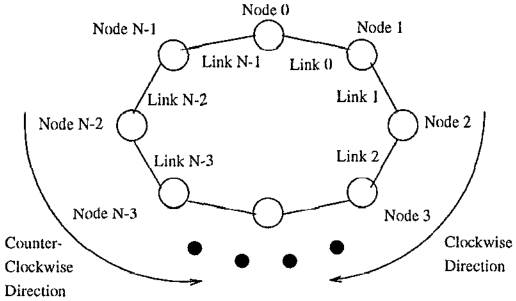

Figure 5: A ring network.

satisfying this request is $\lambda_{\max}$. *Then shortest-path routing yields a load of at most* $2\lambda_{\max}$.

Proof. Suppose shortest-path routing yields a load λ_{sp}. Consider a link i with load λ_{sp}. Rerouting k lightpaths using link i using their longer routes on the ring can reduce the load on link i to at most $\lambda_{sp} - k$. Note that since all these lightpaths are routed on paths of length $\leq \lfloor N/2 \rfloor$ initially, their longer routes on the ring will all use the link $(\lfloor N/2 \rfloor + i)$ mod N, increasing its load by k. Therefore an optimal routing algorithm would have a load given by $\lambda_{\max} \geq \min_k(\lambda_{sp} - k, k)$, or $\lambda_{\max} \geq \lceil \lambda_{sp}/2 \rceil$. □

For the rest of the section, we describe ring networks that lead to efficient channel assignments for requests. To find channel assignments, we will use structures for the requests. In particular, for each request R, the structures are another request R^c and a set of labels for the routes in R^c. The request R^c will be referred to as the *cut request* for R because it is generated from R by "cutting" the routes of R at node 0. The label for each route $p \in R^c$ will be denoted by $\gamma(p)$, which is an integer from the set $\{0, 1, ..., \lambda_{\max} - 1\}$, where $\lambda_{\max}$ is the load for R.

We will now define R^c and then $\gamma(\cdot)$. For the sake of discussion, denote the routes of R by $\{p_0, ..., p_{m-1}\}$, where m is the number of routes. Also, refer to routes that pass through node 0 as *cut routes*, and ones that do not as *uncut routes*. R^c is generated from R as follows. For each uncut route $p_i \in R$, include p_i in R^c. For each cut route $p_i \in R$, cut (or split) it at node 0 into a pair of routes $\{a_i, b_i\}$ such that each has node 0 as a terminating node. Include a_i and b_i into R^c. We will refer to a_i and b_i as the *residual routes* for p_i. Route a_i is assumed to be the one that intersects link $N-1$ and is referred to as the *left* residual route, while b_i is the one that intersects link 0 and is referred to as the *right* residual route. (Figure 6 shows a request $\{p_0, p_1, ..., p_6\}$. Note that p_2, p_4, and p_6 are cut routes since they pass through node 0. Figure 7 shows the routes of R^c.)

Next we define the labels $\gamma(\cdot)$ for R^c. The labels for the routes in R^c are such that routes traversing a common link have distinct labels. This is like *coloring* paths in an *interval graph* [2, Sec.16.5] because no route of R^c crosses through node 0. Hence, we can use a *greedy algorithm* assignment that requires $\lambda_{\max}$

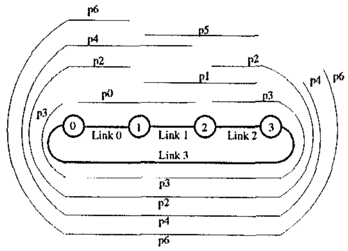

Figure 6: A request $R = \{p_0, p_1, ..., p_6\}$ for a four node ring network.

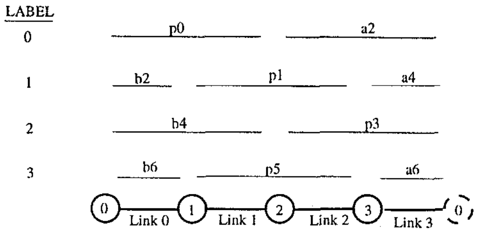

Figure 7: Cut request R^c and label $\gamma(\cdot)$.

numbers [2, Sec.16.5]. Figure 7 shows the labels for R^c.

Theorem 2 *Suppose the ring network has full wavelength conversion at one node and no wavelength conversion at the other nodes. Then any request with load at most* W *has a channel assignment.*

Proof. Without loss of generality, assume that node 0 has the full wavelength conversion. Suppose we are given an arbitrary request R with load $\lambda_{\max}$ at most W. Let R^c be the cut request for R, and let $\gamma(\cdot)$ be the corresponding label. Consider the following channel assignment for R. For each uncut route $p_i \in R$, assign it the channels at wavelength $\omega_{\gamma(p_i)}$ on all its links. For each cut route $p_i \in R$, assign it the channels at wavelength $\omega_{\gamma(a_i)}$ on all links corresponding to its left residual route a_i, and assign it channels at wavelength $\omega_{\gamma(b_i)}$ on all links corresponding to its right residual route b_i.

The channel assignment is valid which can be checked as follows. Note that channels assigned to a route in R are attached because they have the same wavelength, except at node 0, which has full wavelength conversion. Note that different routes of R are assigned different channels because of the way $\gamma(\cdot)$ is defined. Finally, note that only W wavelengths are used because $\gamma(\cdot)$ takes values from $\{0, ..., \lambda_{\max} - 1\}$ and $\lambda_{\max} \leq W$. □

The ring network of Theorem 2 has a channel assignment for every request with load at most W. However, one of its nodes has wavelength degree W. The next theorem states that there is a ring network with wavelength degree one (i.e., fixed wavelength conversion) that has channel assignments for every request with load at most $W - 1$.

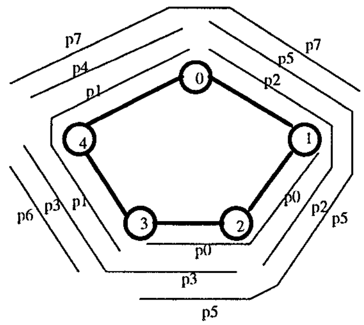

Figure 8: An MCR $(p_0, p_1, ..., p_7)$ that starts and ends at node 1.

First we will define additional structure for a request that will help determine a channel assignment. Consider a sequence of routes $(p_0, p_1, ..., p_{k-1})$, where k is the number of routes. The sequence is referred to as a *multi-cycle of routes* (MCR) if for $i = 0, 1, ..., k-1$, the last node of p_i is the first node of $p_{(i+1) \bmod k}$ when we consider the routes to be going in the clockwise direction. Hence, an MCR is just a sequence of routes that circumvents the ring in the clockwise direction one or more times. The number of times an MCR goes around the ring is called its *multiplicity*. Figure 8 shows an MCR with multiplicity three.

Now we will consider requests such that the number of routes traversing any link is the same. Such requests will be referred to as *uniform* requests. Given a uniform request R, we can define a collection of MCRs $\{M_0, M_1, ..., M_{k-1}\}$ (for some k) with the property that each route of R is in exactly one MCR, and the MCRs are composed only of the routes of R. The collection is called an *MCR partition* for R.

We now describe how to construct an MCR partition from a uniform request $R = \{p_0, p_1, ..., p_{m-1}\}$, with load denoted by $\lambda_{\max}$. First, find the cut request R^c for R and its corresponding label $\gamma(\cdot)$. Note that R^c is also a uniform request with load $\lambda_{\max}$. Second, construct an MCR partition for R^c. This MCR partition will be denoted by $\{M_0, M_1, ..., M_{\lambda_{\max}-1}\}$, where M_j (for $0 \leq j < \lambda_{\max}$) is the MCR with multiplicity one formed by the routes $p \in R^c$ that have label $\gamma(p) = j$. Notice that the MCRs each have at most one right residual route and at most one left residual route. Without loss of generality, assume that MCRs with residual routes start and end at node 0, and go clockwise around the ring. Thus, a right residual route would be the first route in an MCR, while a left residual route would be the last route.

Third, pair the MCRs that have exactly one residual route such that each pair has a left and a right residual route. (Note such a pairing is possible since the number of left residual routes is equal to the number of right residual routes.) Then merge each pair of MCRs to form new MCRs. To be more specific, consider a pair (M_l, M_r), where M_l and M_r holds the left and right residual routes, respectively. Merge the pair by attaching M_l after M_r.

Note that the MCRs are now of two types. The first type has no residual routes, and these are referred to as being *unmergeable*. The second type starts at node 0 with a right residual route, goes clockwise, and ends at node 0 with a left residual route. We refer to these MCRs as being *mergeable*. Now as long as there are mergeable MCRs do the following. Form a sequence of mergeable MCRs $(M'_0, M'_1, ..., M'_{j-1})$ (for some j) such that for $i = 0, 1, ..., j-1$, the last route of M'_i and the first route of $M'_{(i+1) \bmod j}$ are the right and left residual routes, respectively, of some cut route in R. Merge the MCRs $M'_0, M'_1, ..., M'_{j-1}$ by attaching them in succession to form a new MCR. Then in the new MCR, for $i = 0, 1, ..., j-1$, replace the last route of M'_i and the first route of $M'_{(i+1) \bmod j}$ with their corresponding cut route. The new MCR is now unmergeable, i.e., it has no residual routes.

After the merging, the set of MCRs only contain the routes of R. Thus, it is an MCR partition for R.

Example 1: The following is an example of constructing a MCR partition from a request, which is shown in Figure 6. Note that its cut request is shown in Figure 7. The MCR partition of the cut request is $M_0 = (p_0, a_2)$, $M_1 = (b_2, p_1, a_4)$, $M_2 = (b_4, p_3)$, and $M_3 = (b_6, p_5, a_6)$. Note that M_0 and M_2 are the MCRs with only one residual route. These are merged together to form a new MCR $M_{20} = (b_4, p_3, p_0, a_2)$. The remaining MCRs $\{M_1, M_{20}, M_3\}$ are all mergeable. Merge M_1 with M_{20} to form the MCR $M' = (b_2, p_1, a_4, b_4, p_3, p_0, a_2)$. Replace (a_4, b_4) with the cut route p_4, and replace (a_2, b_2) with cut route p_2. This leaves the MCR $M' = (p_1, p_4, p_3, p_0, p_2)$, which is unmergeable. This leaves M_3 as the only mergeable MCR. Replace (a_6, b_6) with p_6 which leaves $M_3 = (p_5, p_6)$. Then M_3 is unmergeable. Note that $\{M_3, M'\}$ is an MCR partition for the request in Figure 6. □

We will define ring networks that take advantage of the MCR partition structure. The channels of these ring networks can be organized into sequences of attached channels called *multi-cycles of channels* (MCC). An MCC is a sequence of distinct channels $(c_0, c_1, ..., c_{k-1})$, where k is a multiple of N, that starts at some node j goes around the ring in the clockwise direction one or more times and ends at node j. In addition, for $i = 0, 1, ..., k-1$, the channels c_i and $c_{(i+1) \bmod k}$ must be attached. The *multiplicity* of an MCC is the number of times it goes around the ring, i.e., it is equal to $\frac{k}{N}$. Figure 9 shows an MCC with multiplicity 3 that starts and ends at node 1.

Note that an MCC with multiplicity m (for $1 \leq m \leq W$) will lead to a channel assignment for the routes in an MCR with multiplicity m. In particular, the channel assignment can be computed as follows. First assume without loss of generality that the MCC and MCR start at the same node. Starting from the node, go clockwise around the ring m times one link at a time. While going around the ring, keep track of the corresponding route of the MCR and the channel of the MCC, and assign the channel to the route.

Now consider ring networks with fixed conversion

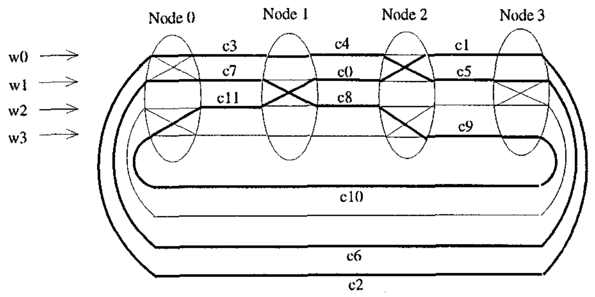

Figure 9: An example MCC with multiplicity 3 that starts and ends at node 1.

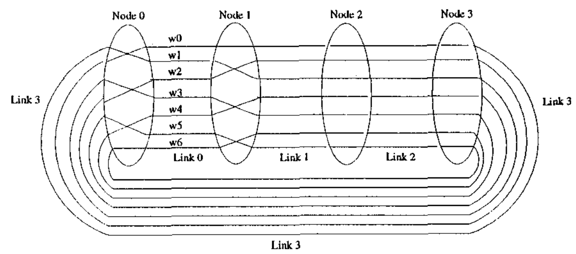

Figure 10: An example of a single-MCC ring network with wavelengths $\{\omega_0, \omega_1, ..., \omega_6\}$.

such that their channels form a single MCC, i.e., the MCC has multiplicity W. We will refer to such networks as *single-MCC* ring networks, and we give two examples.

Example 2: The network shown in Figure 3 is a single-MCC ring network. At node 0, the channel at ω_i in link 3 is attached to the channel at $\omega_{(i+1) \bmod W}$ in link 0 for $i = 0, 1, ..., W-1$. At the other nodes, for $i = 0, 1, ..., W-1$, channels at ω_i are attached to each other (no wavelength conversion). □

Example 3: The network shown in Figure 10 is a single-MCC ring network. At node 0, channels at ω_i are attached to channels at ω_{i+1} for even values of $i < W-1$. If W is odd then channels at ω_{W-1} are attached to each other. At node 1, channel ω_i is attached to channel ω_{i+1} for odd values of $i < W-1$. Channels at ω_0 are attached to each other, and if W is even then channels at ω_{W-1} are attached to each other. At other nodes there is no wavelength conversion. □

Theorem 3 *For a single-MCC ring network, any request R with load at most $W-1$ has a channel assignment.*

Proof. Without loss of generality, assume that R has uniform load (otherwise, we can add one-hop *dummy routes*). There is an MCR partition $\{M_0, M_1, ..., M_{k-1}\}$ for R, where k is the number of MCRs in the partition. We will transform the MCR partition into a single MCR with multiplicity W by adding dummy routes as follows.

For $i = 0, 1, ..., k-1$, let x_i be the node where MCR M_i starts and ends. Without loss of generality, let $x_0 \le x_1 \le ... \le x_{k-1}$. Let $p' = \{p'_0, p'_1, ..., p'_{k-1}\}$ be a collection of dummy routes such that for $i = 0, 1, ..., k-1$, p'_i starts at node x_i, ends at node $x_{(i+1) \bmod k}$, and goes clockwise around the ring. However, if $x_i = x_{(i+1) \bmod k}$ then the route p'_i has zero length, i.e., it is a path that starts and ends at node $x_i = x_{(i+1) \bmod k}$ but does not traverse any links. Note that each link of the ring network has at most one route of p' traversing it because $x_0 \le x_1 \le ... \le x_{k-1}$, i.e., p' has load at most one. Create a new MCR $M = (M_0, p'_0, M_1, p'_1, ..., M_{k-1}, p'_{k-1})$. Note that M is an MCR because for $i = 0, 1, ..., k-1$, the dummy route p'_i starts at the end of M_i and ends at the beginning of $M_{(i+1) \bmod k}$.

The routes of M have load that is either $W-1$ or W because R has load $W-1$ and the dummy routes have load at most one. If the routes of M have load $W-1$ then another dummy route p_k can be appended to M that starts and ends at node x_0 (the starting and ending point for M) and circumvents the ring exactly once. Then M will be an MCR with multiplicity W. Since M is an MCR with multiplicity W and the single-MCC ring network has an MCC with multiplicity W, there is a channel assignment for the routes of M. Thus, there is a channel assignment for R. □

The following theorem states that with fixed conversion, all requests with load at most W are not guaranteed to have a channel assignment, proving that Theorem 3 provides the best possible construction and channel assignment for a fixed-conversion ring network. For brevity, the proof is omitted. Its outline can be found in [20].

Theorem 4 *For any ring network with W wavelengths, fixed wavelength conversion at every node, and a sufficiently large number of nodes, there is a request with load W that does not have a channel assignment.*

Theorem 3 illustrates that providing fixed wavelength conversion is sufficient to obtain efficient channel assignments for the offline case, as long as the load of the request is at most $W-1$. By allowing a bit more wavelength conversion, a ring network may be designed to have channel assignments for all requests with load at most W. This is stated in the next theorem.

Theorem 5 *There is a ring network that has wavelength degree two at two nodes and no wavelength conversion at the other nodes such that every request with load at most W has a channel assignment.*

The proof of the theorem will be given after preliminary definitions and results. Consider a ring network with W wavelengths. A vector $(m_0, m_1, ..., m_{k-1})$ is be referred to as a *multiplicity vector* if its elements are positive integers, $k \in \{1, 2, ..., W\}$, and $\sum_{i=0}^{k-1} m_i = W$. The ring network is referred to as a *multi-MCC ring network* if for each multiplicity vector $(m_0, ..., m_{k-1})$, it has a collection of MCCs $(H_0, H_1, ..., H_{k-1})$ that are channel disjoint and, for $i = 0, 1, ..., k-1$, H_i has multiplicity m_i.

Theorem 6 *For a multi-MCC ring network, any request with load at most W has a channel assignment.*

Proof. Without loss of generality, we may assume that the request is uniform and has load W. (Otherwise, we can make it so by adding *dummy* one-hop routes.) Let $\{M_0, M_1, ..., M_{k-1}\}$ be an MCR partition for the request, where k is the number of MCRs in the partition. For $i = 0, 1, ..., k-1$, let m_i be the multiplicity of M_i. Note that $\sum_{i=0}^{k-1} m_i = W$. Therefore, from the definition of a multi-MCC ring network, there is a collection of channel disjoint MCCs $(H_0, H_1, ..., H_{k-1})$ such that for $i = 0, 1, ..., k-1$, m_i is the multiplicity of H_i. It is now straight forward to construct a channel assignment for the request since for each $i = 0, 1, ..., k-1$, a channel assignment for the routes of M_i can be constructed from the channels in H_i. □

We now describe a particular multi-MCC ring network which we refer to as a *paired wavelengths* (PW) ring network. The network has two nodes, called the *primary* and *secondary* nodes, that have wavelength conversion capability. All other nodes have no wavelength conversion.

To describe the wavelength conversion capability at the primary and secondary nodes we will use the following terminology. We say that a pair of wavelengths (ω, ω') form a *switching pair* at a node if WDM channels at $\{\omega, \omega'\}$ are attached to channels at $\{\omega, \omega'\}$. At the primary node, the following pairs of wavelengths form switching pairs: (ω_0, ω_1), (ω_2, ω_3), and so forth. At the secondary node the following pairs of wavelengths form switching pairs: (ω_1, ω_2), (ω_3, ω_4), and so forth. Figure 2 shows how the channels are attached for the case $W = 7$, and where the primary and secondary nodes are nodes 0 and 1, respectively. Note that the network has channel degree two at the primary and secondary nodes.

Proof of Theorem 5: In light of Theorem 6 it is sufficient to prove that the PW ring network is a multi-MCC ring network. To do this consider an arbitrary multiplicity vector $(m_0, m_1, ..., m_{k-1})$. We now proceed to show that there is a collection $(H_0, H_1, ..., H_{k-1})$ of channel disjoint MCCs such that for $i = 0, ..., k-1$, H_i has multiplicity m_i. We will do this by constructing the MCCs.

For each $i = 0, 1, ..., k-1$, assign m_i wavelengths to construct H_i. In particular, assign wavelengths $\{\omega_0, \omega_1, ..., \omega_{m_0-1}\}$ to H_0, assign wavelengths $\{\omega_{m_0}, \omega_{m_0+1}, ..., \omega_{m_0+m_1-1}\}$ to H_1, and so forth. Since the MCCs use different wavelengths, they will be channel disjoint.

To complete the proof, we define the H_i (for $0 \le i < k$) from its m_i assigned wavelengths $\{\omega_j, \omega_{j+1}, ..., \omega_{j+m_i-1}\}$, where $j = \sum_{s=0}^{i-1} m_s$. The definition is recursive. It starts by noting that there is an MCC with multiplicity one using channels only at wavelength ω_j. Now for $s = 1, 2, ..., m_i - 1$, suppose there is an MCC H' with multiplicity s that only uses channels at wavelengths $\{\omega_j, \omega_{j+1}, ..., \omega_{j+s-1}\}$. Then we can form an MCC H'' with multiplicity $s+1$ that only uses channels at wavelengths $\{\omega_j, \omega_{j+1}, ..., \omega_{j+s}\}$ as follows. Note that there is an MCC H^* with multiplicity one that uses channels only at wavelength ω_{j+s}, and that there is a node x (either the primary or secondary node) that has $(\omega_{j+s-1}, \omega_{j+s})$ as a switching pair. Let H'' be the MCC by joining H' and H^* at node 0 through the switching pair $(\omega_{j+s-1}, \omega_{j+s})$. In other words, H'' is the sequence of channels that starts at node x at wavelength ω_{j+s}, follows the channels of H^* until returning to x at wavelength ω_{j+s}, then follows the channels of H' starting at wavelength ω_{j+s-1} until it returns to x at wavelength ω_{j+s-1}. Finally, H_i is defined to be H'' when $s = m_i - 1$. □

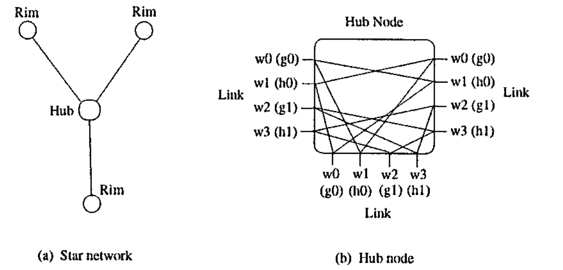

Figure 11: A star network with three rim nodes is shown in (a). A diagram showing how the channels in the hub are attached when $W = 4$ is shown in (b).

Another kind of multi-MCC ring network is described in [20]. It has its channels attached to form a *permutation interconnection network* such as the Benes network [1]. Although the ring network requires more switches/wavelength converters than the PW network, it is useful in handling failures in the network. This topic is explored in detail in [10].

3 Stars, Trees and Meshes

In this section, we will first consider star networks, and then consider more general networks. Throughout this section we will assume that W is even, so that the wavelengths may be paired as follows: $(\omega_0, \omega_1), (\omega_2, \omega_3), ..., (\omega_{W-2}, \omega_{W-1})$. For $i = 0, 1, ..., \frac{W}{2} - 1$, let (g_i, h_i) denote the pair $(\omega_{2i}, \omega_{2i+1})$.

A star network consists of a *hub* node and one or more *rim* nodes, as shown in Figure 11(a). There are links between the rim nodes and the hub only. The star network we consider has a hub node with fixed wavelength conversion such that for $i = 0, 1, ..., \frac{W}{2} - 1$, channels at g_i are only attached to channels at h_i. We call such a node as one having *fixed conversion wavelength pairs* (FCWP). Figure 11(b) shows a hub node having FCWP. Notice that channels at the same wavelength are not attached.

As in the previous section, requests will be given structure that will lead to efficient channel assign-

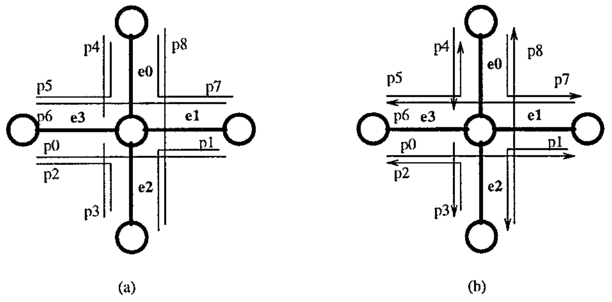

Figure 12: Shown in (a) is a request $\{p_0, p_1, ..., p_8\}$ for a star network of four links $\{e_0, e_1, e_2, e_3\}$. Shown in (b) is a balanced set of directions for the request.

ments. In particular, the routes of a request will be given *directions*, and these directions will be used to determine a channel assignment. To direct a route, one of its terminating nodes is designated as the *first* node, the other terminating node is designated as the *last* node, and the route is assumed to go from the first to the last node. Now consider a request, and suppose the routes of the request are directed. If the routes along each link are directed such that exactly half traverse the link in one direction and the other half traverse the link in the opposite direction then the request is said to have a *balanced set of directions.* Figure 12(a) shows a request for a star network, and Figure 12(b) shows a balanced set of directions for the request.

The following lemma considers tree networks, of which star networks are a special case. It also considers a uniform request with load W (i.e., each link has W routes traversing it).

Lemma 1 *Consider a tree network T with W even. Every uniform request $R = \{p_0, p_1, ..., p_{m-1}\}$ with load W may have its routes directed such that it has a balanced set of directions.*

Proof. First we will show that at each node in T, the number of routes of R that terminate at the node is even. Consider an arbitrary node u in T, and let d denote the number of links that are incident to it. Note that each link has exactly W routes traversing it. Thus, $\sum_{e \in E_u} \rho_e = d \cdot W$, where ρ_e is the number of routes of R that traverse link e, and E_u is the set of links incident to u. Now let n_1 denote the number of routes that terminate at u, and let n_2 denote the number of routes that have u as an intermediate node. Therefore, $n_1 + 2n_2 = \sum_{e \in E_u} \rho_e$. Thus, $n_1 + 2n_2 = d \cdot W$. Since both W and $2n_2$ are even, n_1 must be even.

Second, consider the following *multigraph* G, i.e., a graph which can have multiple edges between vertices. The number of vertices in G is equal to the numbers of nodes in T. Each vertex in G represents a distinct node in T and is labeled with the node it represents. G has m edges with labels $\{p_0, p_1, ..., p_{m-1}\}$. The edge labeled p_i in G is between vertices u and v if the route p_i in T has u and v as terminating nodes. Note that the vertices of G have an even number of incident edges because the number of routes of R that terminate at any node in T is even. Thus, there is an Euler tour for G [12, Chap. 7]. Traverse the edges of G according to the tour, traversing each edge exactly once. Let $(\pi(0), \pi(1), ..., \pi(m-1))$ be the permutation of $(0, 1, ..., m-1)$ such that $(p_{\pi(0)}, p_{\pi(1)}, ..., p_{\pi(m-1)})$ is the sequence of edges traversed. Direct each edge according to how it was traversed. Direct each route p_i in R according to how its corresponding edge in G is directed.

We will verify that the directions are balanced by checking an arbitrary link e in T. Now note that for $i = 0, 1, ..., m-1$, the last node of $p_{\pi(i)}$ is the first node of $p_{\pi((i+1) \bmod m)}$. Then the traversal of T that follows the sequence of directed routes $(p_{\pi(0)}, p_{\pi(1)}, ..., p_{\pi(m-1)})$ goes from node to neighboring node. Now notice that tree T can be considered to be two subtrees connected by e. Thus, if the traversal crosses link e in one direction, then the next time it crosses e it will be in the opposite direction because the only way to get between the two subtrees is across e. Therefore, the traversal crosses link e exactly $\frac{W}{2}$ times in both directions because link e has W routes going through it. Hence, there are $\frac{W}{2}$ routes crossing e in either direction. This verifies that the directions for the request are balanced because e is arbitrary. □

Theorem 7 *Consider a star network with N nodes, W even, and where the hub node has FCWP. Then any request with load at most W has a channel assignment.*

Proof. First let $\{e_0, e_1, ..., e_{N-2}\}$ denote the links of the star network. Next, consider a request R and assume without loss of generality that it is uniform with load W (otherwise, we can add dummy one-hop routes). Let $R_2 = \{p_0, p_1, ..., p_{m-1}\}$ denote the two-hop routes of R. The other routes of R are one-hop routes, and denote these by R_1. From Lemma 1, we can direct the routes of R so that they are balanced. Since the routes of R_2 are now directed, they have *first* and *second* links that they traverse.

The channel assignment for a route p_i in R_2 will be specified by a number n, that satisfies $0 \leq n < \frac{W}{2}$ and which is referred to as the *wavelength pair index* (WPI) for p_i. In particular, the channels for p_i is the channel at g_n on the first link of p_i and the channel at h_n on the second link. Note that the channels are attached because the hub node has FCWP. Now refer to a collection of WPIs, one per route of R_2, as a *feasible* collection if, for each link e, the WPIs of routes that use e as their first link are distinct and the WPIs of routes that use e as their second link are distinct. Note that the resulting channel assignment will have each channel assigned to at most one route. Thus, a feasible collection of WPIs leads to a channel assignment for R_2.

To determine a feasible collection of WPIs, consider a bipartite graph G (see Figure 13) which has two sets of vertices $\{f_0, f_1, ..., f_{N-2}\}$ and $\{s_0, s_1, ..., s_{N-2}\}$. (For $i = 0, 1, ..., N-2$, the vertices f_i and s_i both represent link e_i in the star network, but f_i corresponds to when e_i is the first link of a route and s_i corresponds

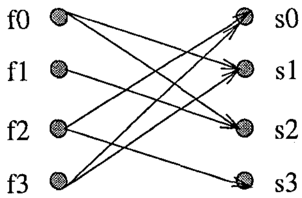

Figure 13: The bipartite graph G corresponding to the star network and directed routes of Figure 12(b).

to when e_i is the second link.) Bipartite graph G has m edges corresponding to the routes of R_2. In fact, they are labeled $p_0, ..., p_{m-1}$. For $i = 0, 1, ..., m-1$, if route p_i has e_j and e_k as its first and second links, respectively, then the edge of G labeled p_i is between vertices f_j and s_k.

Since the directions of the routes of R are balanced, at any link in the star network, at most $\frac{W}{2}$ routes of R_2 are directed towards (resp., away from) the hub node. Therefore, for any link in the star network, at most $\frac{W}{2}$ routes of R_2 use the link as their first (resp., second) link. Hence, each vertex of G has at most $\frac{W}{2}$ incident edges. Then numbers $\{0, 1, ..., \frac{W}{2} - 1\}$ can be assigned to the edges of G such that, at each vertex, the numbers assigned to the edges incident to the vertex are distinct. In particular, the assignment can be accomplished using the scheduling algorithms used for Satellite Switched/Time Division Multiple Access (SS/TDMA) systems [14]. Let these numbers be the WPIs for the routes corresponding to their edges. Note that these WPIs for the routes in R_2 are feasible, and so there is a channel assignment for R_2.

To complete the channel assignment for R, we have to find channel assignments for the routes in R_1. This is trivial since the the routes in R_1 are one-hop. □

Corollary 1 *Consider an arbitrary topology network G with W even and where all nodes with two or more incident links have FCWP. There is a channel assignment for each request that has load at most W and routes traversing at most two links.*

Proof. A star network graph G_{star} is used to represent the links of G. Each link e in G is represented by a distinct edge in G_{star} which is also labeled e. A channel assignment for a request $R = \{p_0, p_1, ..., p_{m-1}\}$ in network G can be determined as follows. Let $R' = \{p'_0, p'_1, ..., p'_{m-1}\}$ be a set of routes for G_{star} such that for $i = 0, 1, ..., m-1$, p'_i traverses the edges of G_{star} with the same labels as the links traversed by p_i. Note that R' has load at most W on G_{star} since it has the same load as R on G. Thus, we can apply Theorem 7 to R' on G_{star} to get a channel assignment for R'. Note that this channel assignment can be translated into a channel assignment for R on G in a straightforward way because the channels assigned to the same route of R are attached. The channels assigned to a route are attached because routes that traverse two links in G have intermediate nodes with FCWP. □

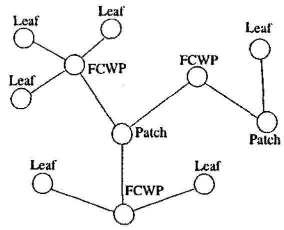

Figure 14: A tree network with leaf nodes, patch nodes, and nodes with FCWP.

In the next theorem, we consider a tree network. We will refer to nodes that have one incident link as a *leaf node*. We will refer to nodes that have channels at wavelengths $\{g_0, g_1, ..., g_{W/2-1}\}$ attached to channels at wavelengths $\{h_0, h_1, ..., h_{W/2-1}\}$ as *patch nodes*. Note that patch nodes have wavelength degree $\frac{W}{2}$.

Theorem 8 *Consider a tree network with W even and is composed of nodes that are patch nodes, leaf nodes, and nodes that have FCWP. Suppose that the network is such that all nodes with FCWP have neighboring nodes that are leaf or patch nodes. Then every request with load at most W has a channel assignment.*

The proof of the theorem is omitted for brevity. Its outline can be found in [20].

4 Conclusion

The paper considered a network model with bidirectional links, WDM channels, and lightpaths for ring, star, and tree networks, as well as arbitrary topology networks with restrictions on lightpath route lengths. It was shown that there are ring and star networks with minimal wavelength conversion capabilities that can perform off-line channel assignment as well as networks with full wavelength conversion. In fact, the ring and star networks of Theorems 5 and 7 only required wavelengths to be shifted to their nearest wavelengths, i.e., the required range of wavelength shifting is minimal.

It should also be noted that the results of this paper can be extended to the case when links, WDM channels, and lightpaths are unidirectional. The results on ring networks in Section 2 can be extended in a straightforward way. Theorem 7 (the result on the bidirectional star network) can be extended to the unidirectional case and where the hub does not have any wavelength conversion (here, the bidirectional star has two unidirectional links that go in opposite directions between the hub and any rim node). Extending all these results to arbitrary topologies is an important open problem.

Acknowledgement: We thank Ori Gerstel for his comments on the paper.

References

[1] V. Benes, *Mathematical Theory of Connecting Networks.* Academic Press, 1965.

[2] C. Berge, *Graphs and Hypergraphs.* North Holland 1976.

[3] A. Birman, "Computing approximate blocking probabilities for a class of optical networks," *IEEE JSAC/JLT Special Issue on Optical Networks*, vol. 14, no. 5, pp. 852-857, June 1996. An earlier version appeared in Proc. Infocom '95.

[4] A. Birman and A. Kershenbaum, "Routing and wavelength assignment methods in single-hop all-optical networks with blocking," *Proc. IEEE Infocom '95*, pp. 431-438, 1995.

[5] G. K. Chang, G. Ellinas, J. K. Gamelin, M. Z. Iqbal, and C. A. Brackett, "Multi-wavelength reconfigurable WDM/ATM/SONET network testbed," *IEEE/OSA J. Lightwave Tech./IEEE JSAC*, vol. 14, no. 6, pp. 1320-1340, June 1996.

[6] I. Chlamtac, A. Ganz, and G. Karmi, "Lightpath communications: an approach to high-bandwidth optical WANs" *IEEE Trans. on Commun.*, vol. 40, pp. 1171-1182, July 1992.

[7] A. Frank, T. Nishizeki, N. Saito, H. Suzuki, and E. Tardos, "Algorithms for routing around a rectangle," *Discrete Applied Math*, vol. 40, pp. 363-378, 1992.

[8] M. Garey, D. Johnson, G. Miller, and C. Papadimitrou, "The complexity of coloring circular arcs and chords," *SIAM Journal on Discrete Math*, vol. 1, no. 2, pp. 216-227, 1980.

[9] O. Gerstel and S. Kutten, "Dynamic wavelength allocation in WDM ring networks," IBM Research Report RC 20462, May 1996.

[10] O. Gerstel, Rajiv Ramaswami, and G. Sasaki, "Fault-tolerant multiwavelength optical networks with limited wavelength conversion," to appear in *Proc. IEEE Infocom '97.*

[11] O. Gerstel, G. Sasaki, and R. Ramaswami, "Dynamic channel assignment for WDM optical networks with little or no wavelength conversion," to appear in *Proc. 1996 Allerton Conference*, Monticello, IL, 1996.

[12] F. Harary, *Graph Theory.* Addison-Wesley, 1969.

[13] J. Y. Hui, *Switching and Traffic Theory for Integrated Broadband Networks.* Kluwer Academic Publishers, 1990.

[14] T. Inukai, "An efficient SS/TDMA time slot assignment algorithm," *IEEE Trans. Commun.* vol. 27, pp. 1449-1445, Oct. 1979.

[15] M. Kovacevic and A. S. Acampora, "On the benefits of wavelength translation in all optical clear-channel networks," *IEEE JSAC/JLT Special Issue on Optical Networks*, vol. 14, no. 5, pp. 868-880, June 1996. An earlier version appeared in Proc. Infocom '95.

[16] K.-C. Lee and V. O. K. Li, "Routing and switching in a wavelength convertible lightwave network," *Proc. Infocom '93*, pp. 578-585, 1993.

[17] K.-C. Lee and V. O. K. Li, "A wavelength-convertible optical network," *IEEE/OSA J. Lightwave Tech.*, vol. 11, pp. 962-970, May/June 1993.

[18] M. Mihail, C. Kaklamanis, and S. Rao, "Efficient access to optical bandwidth," *Proc. of IEEE Symp. on Foundations of Computer Science*, pp. 548-557, 1995.

[19] P. Raghavan and E. Upfal, "Efficient routing in all-optical networks," *Proc. 26th ACM Symp. Theory of Computing*, pp. 134-143, May 1994.

[20] R. Ramaswami and G. Sasaki, "Multiwavelength optical networks with limited wavelength conversion," IBM Research Report RC 20503, July 1996.

[21] R. Ramaswami and K. N. Sivarajan, "Routing and wavelength assignment in all-optical networks," *IEEE/ACM Trans. on Networking*, pp. 489-500, Oct. 1995. An earlier version appeared in IEEE Infocom '94.

[22] S. Subramaniam, M Azizoglu, and A. K. Somani, "Connectivity and sparse wavelength conversion in wavelength-routing networks," *Proc. IEEE Infocom '96*, pp. 148-155, 1996.

[23] H. Toba, K. Oda, K. Inoue, K. Nosu, and T. Kitoh, "An optical FDM based self-healing ring network employing arrayed-waveguide-grating ADM filters and EDFAs with level equalizers," *IEEE JSAC/JLT Special Issue on Optical Networks,* vol. 14, no. 5, pp. 800-813, June 1996.

[24] A. Tucker, "Coloring a family of circular arcs," *SIAM Journal on Applied Math*, vol. 29, no. 3, pp. 493-502, 1975.

[25] N. Wauters and P. Demeester, "Design of the optical path layer in multiwavelength cross-connected networks," *IEEE JSAC/JLT Special Issue on Optical Networks*, vol. 14, no. 5, pp. 881-892, June 1996.

[26] J. Yates, J. Lacey, D. Everitt, and M. Summerfield, "limited-range wavelength translation in all-optical networks," *Proc. IEEE Infocom '96*, pp. 954-961, 1996.

[27] J. Zhou, N. Park, K. Vahala, M. Newkirk, and B. Miller, "Four-wave mixing wavelength conversion efficiency in semiconductor traveling-wave amplifiers measured to 65 nm of wavelength shift," *IEEE Photon. Tech. Lett.*, vol. 6, pp. 984-987, Aug. 1994.