Multiwavelength Optical Networks with Limited Wavelength Conversion*

Rajiv Ramaswami (*corresponding author*)
IBM Research, 30 Saw Mill River Road, Hawthorne, NY 10532
rajiv@watson.ibm.com, 914-784-7356, 914-784-6225 (fax)
and
Galen H. Sasaki
Dept. of Electrical Engg., Univ. of Hawaii, Honolulu, HI 96822
sasaki@spectra.eng.hawaii.edu, 808-956-6103, 808-956-3427 (fax)

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 pairwise 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 also 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.

Keywords: Optical networks, routing, wavelength division multiplexing, wavelength conversion.

*This work was sponsored in part by grant MDA-972-95-C-0001 by ARPA and grant NCRI-93-6130 by NSF.

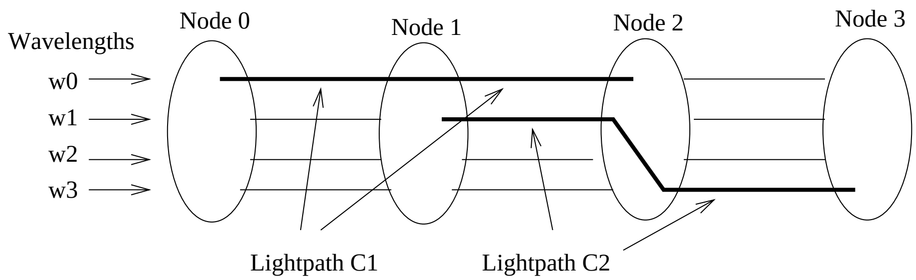

Figure 1: A network with wavelengths $\{\omega_0, \omega_1, \omega_2, \omega_3\}$. Channels are shown as lines between nodes.

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.

An example of a WDM wavelength routing network is shown in Figure 1. It is composed of four nodes with optical fiber links, each having 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

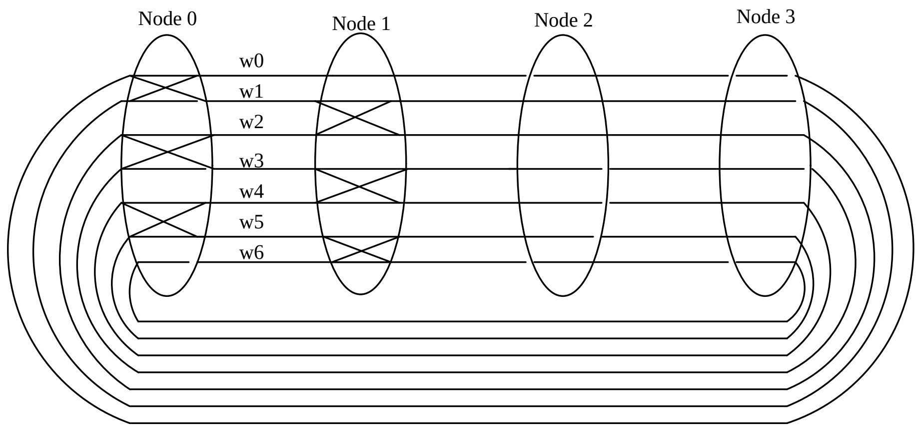

Figure 2: A network with wavelengths $\{\omega_0, \omega_1, ..., \omega_6\}$ to illustrate limited wavelength conversion. The lines between nodes represent channels, and lines within nodes indicate which channels may be connected.

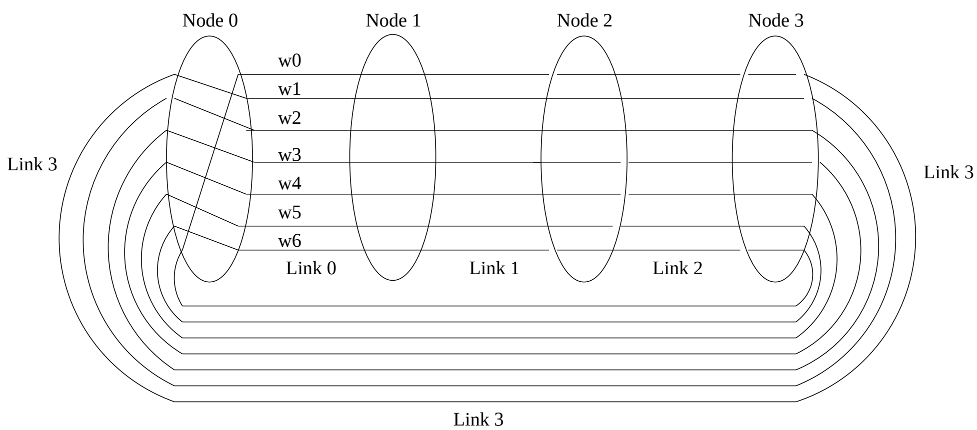

Figure 3: A network with wavelengths $\{\omega_0, \omega_1, ..., \omega_6\}$ to illustrate fixed wavelength conversion. The lines between nodes represent channels, and lines within nodes indicate which channels may be connected.

the channels they may be connected to on other links. For example, Figure 2 shows a network with seven WDM channels between nodes where the channels are at wavelengths $\{\omega_0, \omega_1, ..., \omega_6\}$. The lines within nodes show which pairs of channels 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 which is illustrated in Figure 3. Here, at each node, each channel may be connected to exactly one predetermined channel on every other link. For example, in Figure 3, channel at ω_i in link 3 may be connected only to channel $\omega_{(i+1) \bmod W}$ in link 0 for $i = 0, 1, ..., W-1$.

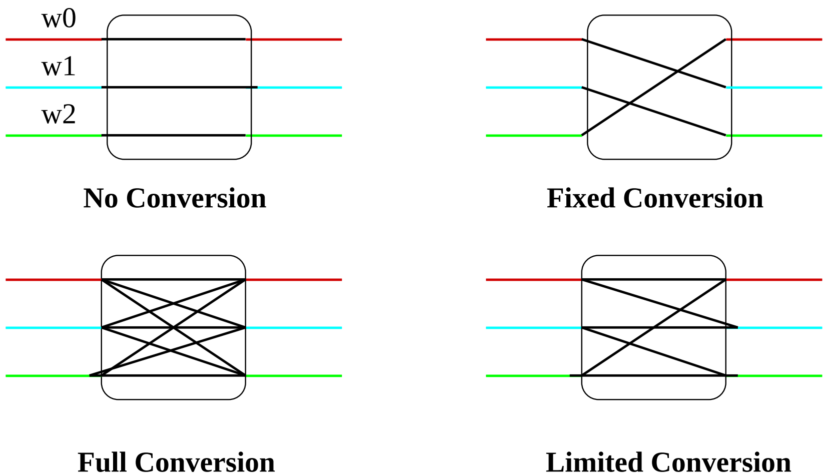

Figure 4: Different types of wavelength conversion.

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 [1], 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 number 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, (for example, 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 a 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 to satisfy 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 first deployed WDM networks are likely to be rings, as seen from several recent testbeds (see for example [2, 3]).

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 $\lambda_{\max} \le W$, when W is an even number. 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. 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* [4]. 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 [5] and remains NP-complete even for rings [6]. Given a routing already, an algorithm that finds a channel assignment in a ring network without wavelength conversion if $2\lambda_{\max} - 1 \le W$ is given in [7] as well as [8]. An algorithm that finds channel assignments in a tree network without wavelength conversion if $\frac{3}{2}\lambda_{\max} \le W$ is given in [8]. 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. [9] gives algorithms that find channel assignments for the case of a directed network without wavelength conversion and directed lightpath requests for trees, if $15\lambda_{\max}/8 \le W$, and for rings, if $2\lambda_{\max} \le W$. Several heuristic channel assignment schemes have also been proposed for networks without wavelength conversion [10, 11, 5, 12, 13, 14].

Variants of the limited conversion model are considered in [15, 16, 17, 18]. In [16, 15] 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 [17], 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 [18] 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 [15, 16, 17, 18] 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 networks [19], 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 [20].

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 *counter-clockwise 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 paths (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 [21]. 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*

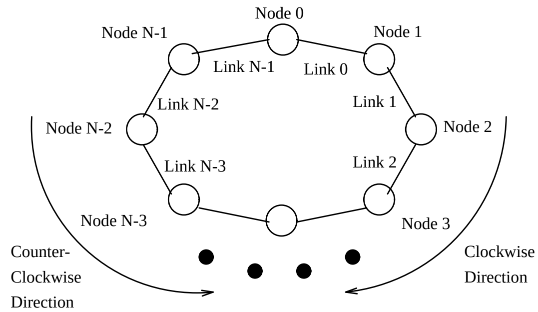

Figure 5: A ring network.

possible load for 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 on length $\leq \lfloor N/2 \rfloor$ initially, their longer routes on the ring will all use the link $\lfloor N/2 \rfloor + i$, 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 a structure for the requests called a *cut-and-color partition*, which we define next.

A cut-and-color partition for a request $\{p_0, ..., p_{m-1}\}$, where m is the number of routes in the request, may be computed as follows. Pick an arbitrary node, say node 0, called the *primary* node. This will be used to "cut" routes in two as explained below. First, refer to routes that pass through node 0 as *cut routes* and the rest of the routes as *uncut.* A set P of routes is generated as follows. Include each uncut route in P. For each cut route p_i, cut (or split) it at node 0 into a pair of paths $\{a_i, b_i\}$ called *residual routes* such that each has node 0 as a terminating node. Let a_i denote the residual route that traverses link $N-1$, and let b_i denote the residual route that traverses link 0. Refer to a_i as the *left* residual route, and b_i as the *right* residual route. Include the residual routes in P. The resulting set P will be referred to as the *uncut and residual routes* of the request. (Figure 6(a) shows a request $\{p_0, p_1, ..., p_6\}$. Note that p_2 and p_4 are cut routes since they pass through node 0. Also note that the routes of the request do not have to be distinct, for example p_1 and p_5 are the same. Figure 6(b) shows the uncut and residual routes in P. Note that a_2 and a_4

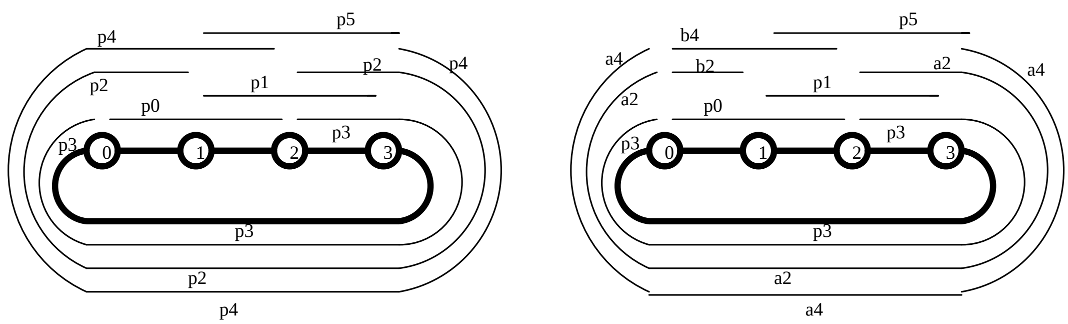

(a) Routes {p0, p1, p2, p3, p4, p5}

(b) The collection of routes P which is composed of uncut routes {p0, p1, p3, p5}, left residual routes {a2, a4}, and right residual routes {b2, b4}.

Figure 6: Shown in (a) is a request $\{p_0, p_1, ..., p_5\}$ for a four node ring network. Shown in (b) is the collection P of uncut and residual routes for the request.

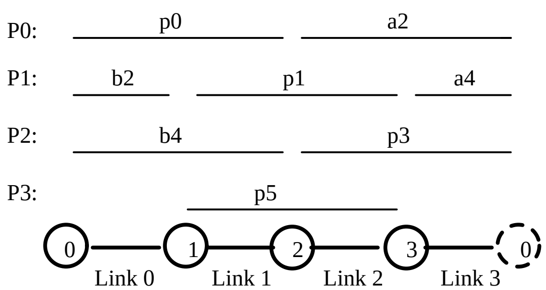

(a) A cut-and-color partition (P0, P1, P2, P3).

Channels \ Links	0	1	2	3
w0	p0	p0	p2	p2
w1	p2	p1	p1	p4
w2	p4	p4	p3	p3
w3	-	p5	p5	-

(b) A channel assignment for the request {p0, p1,..., p5}. Here, the channels are denoted by their wavelengths {w0, w1, w2, w3}.

Figure 7: A cut-and-color partition and a channel assignment.

are the left residual routes of cut routes p_2 and p_4, respectively; and note that b_2 and b_4 are the right residual routes of p_2 and p_4, respectively.)

Next, partition the routes in P into W subsets $(P_0, P_1, ..., P_{W-1})$ such that routes in the same subset do not traverse common links of the ring network. We will refer to the partition $(P_0, P_1, ..., P_{W-1})$ as a *cut-and-color* partition for the request. One way to find a cut-and-color partition is to assign numbers $\{0, ..., W-1\}$ to the routes in P such that routes with a common link have distinct numbers. This is like *coloring* paths in an *interval graph* [22, Sec.16.5] because no route of P crosses through node 0. Hence, we can use a *greedy algorithm* assignment that requires $\lambda_{\max}$ numbers [22, Sec.16.5]. Then for $i = 0, 1, ..., W-1$, all routes that have been assigned to number i are in subset P_i. (Figure 7(a) shows a cut-and-color partition for the request in Figure 6(a). The cut-and-color partition is (P_0, P_1, P_2, P_3), where $P_0 = \{p_0, a_2\}$, $P_1 = \{b_2, p_1, a_4\}$, $P_2 = \{b_4, p_3\}$, $P_3 = \{p_5\}$, $a_2 = (2, 3, 0)$, $b_2 = (0, 1)$,

$a_4 = (3, 0)$, and $b_4 = (0, 1, 2)$.)

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 at most W. Let $(P_0, ..., P_{W-1})$ denote the cut-and-color partition for the request and let P be the set of uncut and residual routes. Note that P is also a request of routes with load at most W. We will first find a channel assignment for P. For $i = 0, 1, ..., W - 1$, assign the wavelength ω_i to each route in P_i. For each route in P, assign it the channels at its wavelength at each link it traverses. The result is a channel assignment for P since the channels assigned to any route are attached (because the channels are at the same wavelength) and no channel is assigned to more than one route (because each P_i has at most one route traversing any link). A channel assignment for R can be modified from P as follows: for each cut route p_i assign it the channels that are assigned to its residual routes. Note that the resulting channel assignment for cut routes are attached because node 0 has full wavelength conversion. (Figure 7(b) shows a channel assignment for the request in Figure 6(a) that results from the cut-and-color partition in Figure 7(a).) □

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 $W - 1$.

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 in the sequence. 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.

An *MCR partition* for a request $R = \{p_0, p_1, ..., p_{m-1}\}$ is a collection of MCRs $(M_0, M_1, ..., M_{k-1})$, where k is the number of MCRs, such that the MCRs are composed of the routes of R and each route is in exactly one MCR. Thus, the MCRs form a partition of R. The next lemma

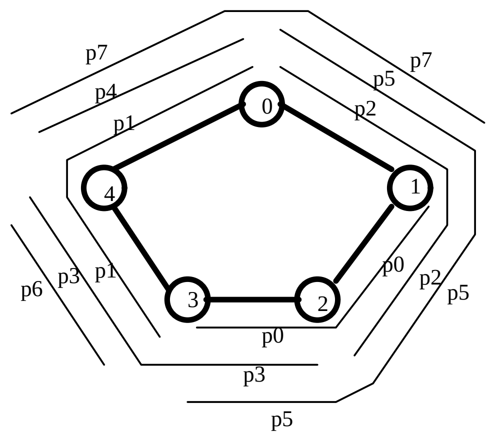

Figure 8: An MCR $(p_0, p_1, ..., p_7)$ that starts and ends at node 1.

states that there is an MCR partition for every *full request*, which is defined to be a request where each link has exactly W routes traversing it.

Lemma 1 *Consider a ring network with a full request* $R = \{p_0, p_1, ..., p_{m-1}\}$. *An MCR partition can be found for the request.*

Proof. We describe a way to construct an MCR partition of R. Find a cut-and-color partition $(P_0, ..., P_{W-1})$ for R. Let P denote the set of uncut and residual routes of R. Note that P is also a full request. Note also that for each $i = 0, 1, ..., W-1$, there can be at most one route of P_i that traverses any link, since P_i belongs to a cut-and-color partition. However since R is a full request, there will be exactly one route of P_i that traverses any link.

For $i = 0, 1, ..., W-1$, let M_i' be the sequence of routes of P_i visited when going around the ring in the clockwise direction starting from node 0. Note that M_i' is an MCR since each link is traversed by exactly one route of P_i. Since the $(P_0, P_1, .., P_{W-1})$ is a partition for P, $(M_0', M_1', ..., M_{W-1}')$ is an MCR partition for P.

Let F denote the MCR partition $(M_0', M_1', ..., M_{W-1}')$. Next is a procedure to iteratively modify F and P so that the residual routes in P are replaced by cut routes, and F remains an MCR partition of P.

Step 1. If P has no residual routes then STOP. Otherwise, in P, find the pair of residual routes of some cut route p_i. Let a_i and b_i denote the left and right residual routes of p_i, respectively.

Step 2. Let M_a and M_b denote the MCRs that contain a_i and b_i, respectively. (Note that M_a and M_b could be the same if a_i and b_i are in the same MCR.)

Step 3. Form a new sequence M of routes as follows. First, let $M = M_a$. If $M_a \neq M_b$ (i.e., a_i and b_i are not in the same MCR of F) then append the sequence M_b to M. Note that M is an MCR of P and contains both a_i and b_i.

Step 4. Modify M as follows, so that residual routes $\{a_i, b_i\}$ are replaced by their corresponding cut route p_i:

(a) If $M_a = M_b$ (i.e., a_i and b_i are in the same MCR of F) then modify M by removing b_i and replacing a_i with p_i.

(b) Otherwise, if $M_a \neq M_b$ then modify M by replacing the pair (a_i, b_i) with p_i. (Note that this can be done because $\{a_i, b_i\}$ are consecutive routes in M, which in turn is due to a_i being the last route in M_a, and b_i being the first route in M_b.)

Step 5. Modify P by removing $\{a_i, b_i\}$ and adding p_i. Modify F by removing $\{M_a, M_b\}$ and adding M. Note that F is still an MCR partition for P. Go to Step 1.

Note that the procedure stops when $P = R$. Since F is an MCR partition for P, the procedure computes an MCR partition for R. $\square$

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 channels called *multi-cycles of channels* (MCC), which will be defined next. An MCC is a sequence of distinct channels $(c_0, c_1, ..., c_{k-1})$ 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. Note that k is a multiple of N and that for $i = 0, 1, ..., k-1$, channel c_i is in link $(i + j) \bmod N$. 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(a) shows an MCC with multiplicity 3 that starts and ends at node 1. Note that the starting (and ending) node could be any node in the ring, but for convenience we will choose specific starting and ending nodes.

Note that for any $m = 0, 1, ..., W-1$, the set of channels in an MCC with multiplicity m is a channel assignment for the set of routes in an MCR with multiplicity m. To illustrate this, let $(p_0, p_1, ..., p_{k-1})$ denote the MCR, and for $i = 0, 1, ..., k-1$, let n_i denote the number

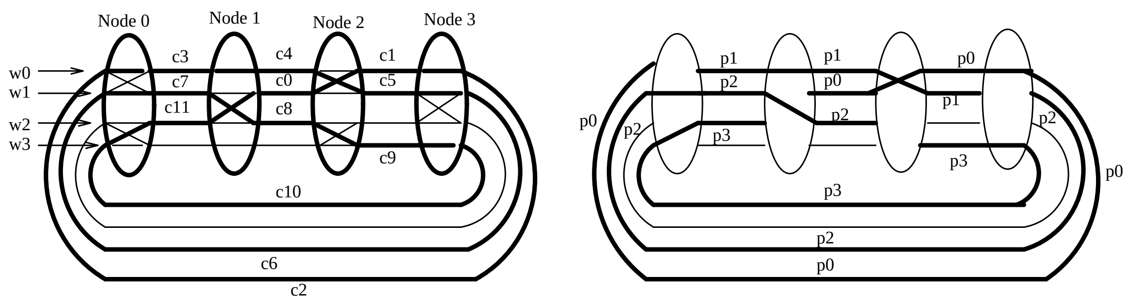

(a) An MCC {c0, c1, ..., c11} with multiplicity 3 in a ring network with wavelengths {w0, w1, w2, w3}. The lines between nodes are WDM channels, and lines within nodes indicate the attached channels.

(b) The channels of the MCC of (a) is a channel assignment for the routes of an MCR (p0, p1, p2, p3).

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

of links traversed by p_i. Note that the MCR goes around the ring in the clockwise direction, and that each route goes in the clockwise direction. Let j denote the first link that p_0 traverses. Assume without loss of generality that the MCC also starts at link j. Then a channel assignment for the routes of the MCR is as follows: the first n_0 channels of the MCC are assigned to p_0, the next n_1 channels of the MCC are assigned to p_1, and so forth. Figure 9(b) shows how the channels of the MCC in Figure 9(a) can be a channel assignment for an MCR with multiplicity 3.

Now consider ring networks with fixed conversion such that their channels form a single MCC, i.e., the MCC has multiplicity W. We will refer to such networks as a *single-MCC* ring network. An example of such a network is the network shown in Figure 3. Here 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). Another example is the network shown in Figure 10. 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 with load at most $W-1$ has a channel assignment.*

Proof. Let R be an arbitrary request with load at most $W-1$. We will first show that

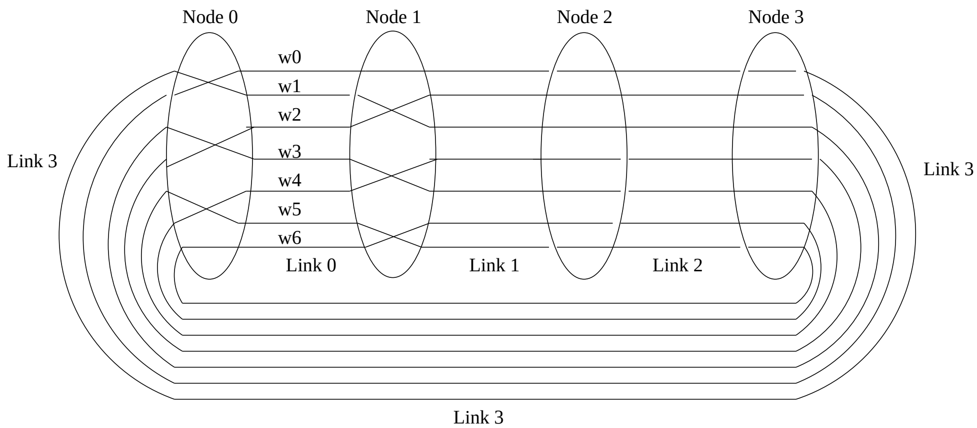

Figure 10: An example of a single-MCC ring network with wavelengths $\{\omega_0, \omega_1, ..., \omega_6\}$.

an MCR M with multiplicity W can be formed from the routes of R and additional *dummy routes*.

Without loss of generality, assume that each link of the ring has exactly $W-1$ routes of R crossing it (otherwise, we can add dummy one-hop routes). Then R is a full request for a ring network with $W-1$ wavelengths. Thus, we can apply Lemma 1 to find an MCR partition $(M_0, M_1, ..., M_{k-1})$, where $k \leq W$ is the number of MCRs, such that $\sum_{i=0}^{k-1} m_i = W-1$, where m_i denotes the multiplicity of M_i for $i = 0, 1, ..., k-1$. Now for $i = 0, 1, ..., k-1$, let x_i be the node where MCR M_i starts and ends. Without loss of generality, let $x_0 \leq x_1 \leq ... \leq x_{k-1}$. 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 $x_0 \leq x_1 \leq ... \leq x_{k-1}$, i.e., p' has load at most one.

An MCR M can be formed by combining the MCRs $(M_0, M_1, ..., M_{k-1})$ and dummy routes $\{p'_0, p'_1, ..., p'_{k-1}\}$ as follows: $M = (M_0, p'_0, M_1, p'_1, ..., M_{k-1}, p'_{k-1})$. 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 p' has 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 exactly 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. $\square$

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.

Theorem 4 *For any value of W, there is a value n such that the following is true. For any ring network with W wavelengths, fixed wavelength conversion at every node, and $N \geq n$ (recall that N is the number of nodes), there is a request with load W that does not have a channel assignment.*

Outline of Proof: We will show that for any ring network with fixed wavelength conversion and a sufficient number of nodes, there is a full request that does not have a channel assignment. We will allow requests to contain routes that start and end at a common node and circumvent the ring clockwise exactly once. Such routes will be referred to as *full cycle* routes, and their starting (and ending) nodes will be referred to as their *terminating nodes.* (At the end of this outline, we will discuss requests that disallow full cycle routes.) Also, let $n = W$.

First, notice that a ring network with fixed conversion at every node has its channels attached to form a single collection of channel disjoint MCCs (i.e., they do not share any channels). Consider two cases:

Case 1, one of the MCCs has multiplicity $m > 1$: Consider a full request consisting of W full cycle routes with distinct terminating nodes (which is possible since $N \geq W$). Let us suppose there is a channel assignment for the request. We proceed to show that this is false. Let H be an MCC with multiplicity m, and let p_0 be one of the routes that has its channels from H. Without loss of generality assume that p_0 has terminating node 0. Note that the terminating node 0 has the incident links 0 and $N-1$. Now let c_0 denote the channel in link 0 assigned to p_0. Let c_{N-1} denote the channel in link $N-1$ that is attached to c_0. Note that c_{N-1} cannot be assigned to p_0 because H has multiplicity $m > 1$. Note that c_{N-1} cannot be assigned to any other route because the other routes do not have node 0 as a terminating node and c_0 (the channel attached to c_{N-1} through node 0) is assigned to p_0. Since c_{N-1} is not assigned to a route and the request is full, the channel assignment is invalid.

Case 2, all of the MCCs have multiplicity one: Then we have W MCCs. Consider the following request which is composed of $W-2$ full cycle routes and three routes (M_0, M_1, M_2), which form an MCR with multiplicity two. The full cycle routes have terminating node 0. Route M_0 starts at node 3 and ends at node 2, route M_1 starts at node 2 and ends at node 1, and route M_2 starts at node 1 and ends at node 3. Note that the routes of the request overlap one another. Since the routes overlap, each route requires channels from a different MCC. Since there are $W+1$ routes, there can be no channel assignment.

Note that the above argument can be modified when full cycle routes are disallowed. This can be done by replacing full cycle routes with MCRs with multiplicity one. However, n must be large enough to insure that terminating nodes are distinct. □

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. The ring network of the theorem is a special case of a class of ring networks we refer to as *multi-MCC ring networks.* A *multi-MCC ring network* is one where for any $k \geq 1$ and collection $(m_0, m_1, ..., m_{k-1})$ of positive integers that satisfies $\sum_{i=0}^{k-1} m_i = W$, there is 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 full. Otherwise, we can make it full by adding *dummy* one-hop routes. From Lemma 1, there is an MCR partition $(M_0, M_1, ..., M_{k-1})$ 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 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 $i = 0, 1, ..., k-1$, we can construct a channel assignment for the routes of M_i from the channels in H_i. $\Box$

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 with wavelength conversion capability called the *primary* and *secondary* nodes. 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. Now at the primary and secondary nodes, WDM channels at the same wavelength are attached. 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.

Lemma 2 *Consider the PW ring network. Let j and k be arbitrary integers satisfying $0 \le j \le j+k \le W-1$. Then there is an MCC with multiplicity $k+1$ that only uses channels with wavelengths $\{\omega_j, \omega_{j+1}, ..., \omega_{j+k}\}$.*

Proof. The proof will be by induction on k. For the case when $k = 0$, the lemma is true because a sequence of channels at ω_j that goes around the ring once is an MCC described by the lemma.

Now suppose the lemma is true for the case $k-1$, i.e., there is an MCC M with multiplicity k that only uses channels at wavelengths $\{\omega_j, \omega_{j+1}, ..., \omega_{j+k-1}\}$. We will show that the lemma is true for k.

Now consider the case when $j+k-1$ is odd (the case when $j+k-1$ is even will be discussed later). Without loss of generality, assume that M starts and ends at the primary node, and that its first channel is at ω_{j+k-1}. At the primary node, the first channel is attached to channels at ω_{j+k-1} and ω_{j+k} because $j+k-1$ is odd. Since M has attached first and last channels and only occupies wavelengths $\{\omega_j, \omega_{j+1}, ..., \omega_{j+k-1}\}$, the last channel of M must be at ω_{j+k-1}. Therefore, we have the following

- The MCC M starts and ends at the primary node, and its first and last channels are at ω_{j+k-1}.
- At the primary node, the first and last channels of M are attached to the channels at ω_{j+k}.

Let M^* be the MCC that starts and ends at the primary node and has channels only at ω_{j+k}. Note that the last and first channels of M are attached to the first and last channels of M^*, respectively, at the primary node. Appending the sequence M^* to M results in an MCC sequence as described in the Lemma.

Now consider the case when $j+k-1$ is even. Then an MCC for the lemma can be found in a similar way as when $j+k-1$ were odd except that M and M^* start and end at the secondary node. □

Lemma 3 *The PW ring network is a multi-MCC ring network.*

Proof. Suppose $k > 0$ is some integer, and $(m_0, m_1, ..., m_{k-1})$ is a collection of positive integers that satisfy $\sum_{i=0}^{k-1} m_i = W$. We will show that there exists a collection of MCCs $(M_0, M_1, ..., M_{k-1})$ such that for $i = 0, 1, ..., k-1$, the multiplicity of M_i is m_i.

Note that the wavelengths may be partitioned into subsets $(\Omega_0, \Omega_1, ..., \Omega_{k-1})$, where $\Omega_0 = \{\omega_0, \omega_1, ..., \omega_{m_0-1}\}$, $\Omega_1 = \{\omega_{m_0}, \omega_{m_0+1}, ..., \omega_{m_0+m_1-1}\}$, and so forth. Note that for $i = 0, 1, ..., k-1$, the number of wavelengths in Ω_i is m_i. From Lemma 2, we know that for $i = 0, 1, ..., k-1$, there is an MCC M_i with multiplicity m_i that only uses the wavelengths in Ω_i. The resulting collection $(M_0, M_1, ..., M_{k-1})$ completes the proof. □

Proof of Theorem 5: The theorem follows from Theorem 6, Lemma 3, and the fact that a PW ring network has wavelength degree two at two nodes and no wavelength conversion at the other nodes. □

Next we describe another class of multi-MCC ring networks based upon *permutation interconnection networks* such as the Benes network [23]. We will refer to these as *permutation ring networks*. The permutation interconnection network corresponding to a permutation ring network will be referred to as a *template permutation network*. The template permutation network for a ring network of N nodes has $N+1$ stages of links, each stage having W links and the stages are labeled $0, 1, ..., N$. The links in the first stage are called the *inputs*, and the links in the last stage are called the *outputs*. The set of input (resp., output) links

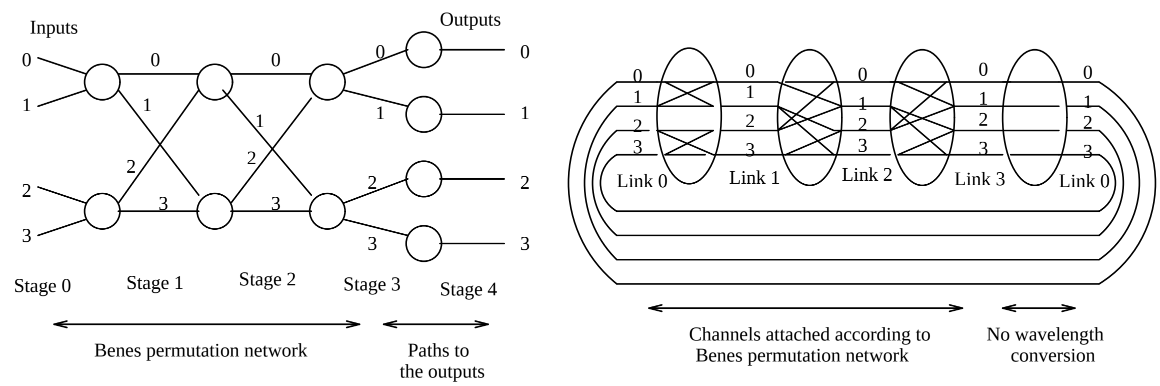

(a) Template permutation network

(b) Permutation ring network

Figure 11: Shown in (a) is template permutation network composed of a Benes permutation network followed by a set of paths to the output links. Shown in (b) is a permutation ring network with four nodes, $W = 4$, and based on the template permutation network in (a).

are labeled $\{0, 1, ..., W-1\}$. The template network has the property that for any permutation $(\pi(0), \pi(1), ..., \pi(W-1))$ of $(0, 1, ..., W-1)$, there is a set of W link-disjoint paths between the inputs and outputs such that for $i = 0, 1, ..., W-1$, the path that starts from input i ends at output $\pi(i)$. An example template network is shown in Figure 11(a) which is composed of a Benes permutation network for the first $2 \log_2 W$ stages of links and then a collection of paths to the output links.

From the template permutation network G_{temp}, a permutation ring network G_{ring} can be defined as follows. For $i = 1, 2, ..., N-1$, the W links at stage i of G_{temp} are assigned to the W channels at link i of G_{ring}. The $2W$ links at stages 0 and N of G_{temp} are assigned to the W channels at link 0 of G_{ring} so that the both the input and output links labeled i are assigned to the channel at wavelength ω_i, for $i = 0, 1, ..., W-1$. The permutation ring network has its channels attached according to the interconnection of G_{temp}, i.e., if two links of G_{temp} are incident to a common node then their corresponding channels in G_{ring} are attached. Figure 11(b) shows a permutation ring network based on the template network of Figure 11(a). Note that the ring network has wavelength degree at most two at each node.

Permutation ring networks are multi-MCC ring networks because the template permutation network has the property that there is a set of link-disjoint paths between its inputs and outputs according to any permutation. For example, an MCC of multiplicity m can be determined as follows. Note that there is a set of m link disjoint paths $\{p_0, p_1, ..., p_{m-1}\}$ in G_{temp} such that for $i = 0, 1, ..., m-1$, p_i starts at input i and ends at output $i + 1 \bmod m$. Note that for $i = 0, 1, ..., m-1$, the links traversed by path p_i corresponds to a sequence

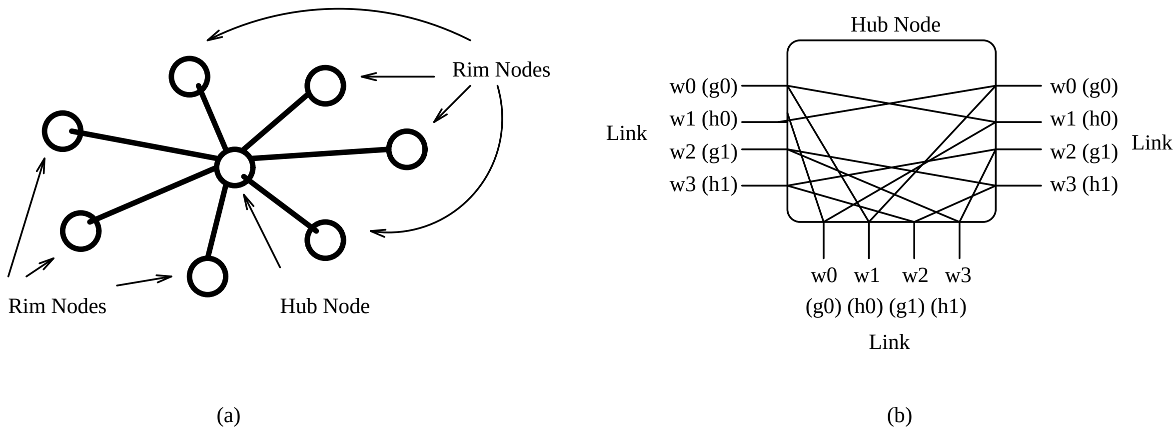

Figure 12: Shown in (a) is a star network with a hub node at center and rim nodes connected to the hub. Shown in (b) is how channels are attached at a hub node that has FCWP and when $W = 4$.

of attached channels that start at ω_i at link 0, goes clockwise around the ring network, and ends at $\omega_{i+1 \bmod m}$ at link 0. Thus, the channels of the ring network that correspond to the links of G_{temp} traversed by $\{p_0, p_1, ..., p_{m-1}\}$ can form an MCC of multiplicity m.

Since a permutation ring network is a multi-MCC ring network, Theorem 6 implies the following theorem.

Theorem 7 *For a permutation ring network, any request with load at most W has a channel assignment.*

While the PW-network is a better choice than a permutation network for efficient wavelength assignment, since it requires fewer switches/wavelength converters, permutation networks will turn out to be useful for handling failures in the network. This topic is explored in detail in [20].

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 12(a). There are links between the rim nodes and the hub only. The star network we

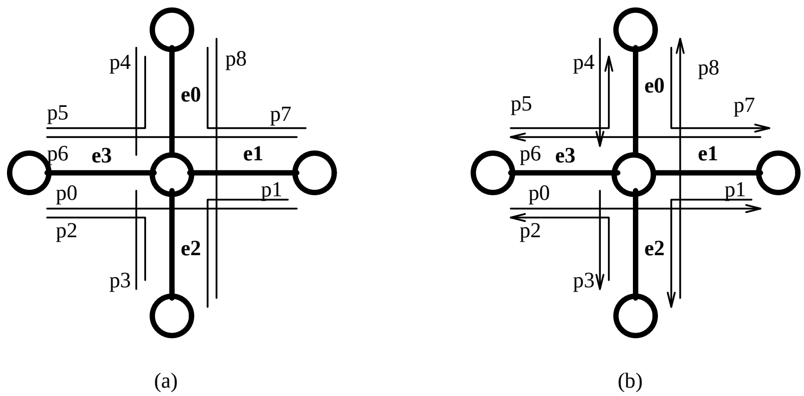

Figure 13: Shown in (a) is a request $\{p_0, p_1, ..., p_8\}$ for a star network of four links. Shown in (b) is a balanced set of directions for the request.

consider has a hub node with fixed channel 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 12(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 assignments. 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 13(a) shows a request for a star network, and Figure 13(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 *full request* (recall, a full request is one where every link has W routes traversing it).

Lemma 4 *Consider a tree network T with W even. Every full request $R = \{p_0, p_1, ..., p_{m-1}\}$ 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. Now note that since R is full, 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, we will show that the routes of R may be put into a sequence $\mathbf{r} = (p_{\pi(0)}, p_{\pi(1)}, ..., p_{\pi(m-1)})$, where $(\pi(0), \pi(1), ..., \pi(m-1))$ is a permutation of $(0, 1, ..., m-1)$, and may be given directions such that for $i = 0, 1, ..., m-1$, the last node of $p_{\pi(i)}$ is the first node of $p_{\pi((i+1) \bmod m)}$. To show that the sequence $\mathbf{r}$ and the directions exist, consider the following *multigraph* G, i.e., a graph which can have multiple edges between nodes. The graph G has the same nodes as T. Its edges are represented by the set $\{p_0, p_1, ..., p_{m-1}\}$, where the p_i is an edge in G between two nodes u and v if u and v are the two terminating nodes of route p_i in T. Note that G is a multigraph where each of its nodes has 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 [24, Chap. 7], and the tour corresponds to $\mathbf{r}$. The way in which the Euler tour can be traversed will give a set of directions for the routes in $\mathbf{r}$.

Note that the directions for the routes of $\mathbf{r}$ is also a set of directions for the routes of R. We will verify that the directions are balanced by checking the routes that cross an arbitrary link e. Now consider traversing the sequence of routes in $\mathbf{r}$ by following their directions. Notice that the traversal follows the links of T, and that the tree T can be considered as 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 the request is full and W is even. Hence, there are $\frac{W}{2}$ routes crossing e in either direction. Since e is an arbitrary link, the set of directions for the request is balanced. $\square$

Theorem 8 *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 full (otherwise, we can make it full by adding *dummy* one-hop routes). Let $R_2 = \{p_0, p_1, ..., p_{m-1}\}$ denote the two-hop routes of R. Note that the other routes of R are one-hop routes, and denote these by R_1. From

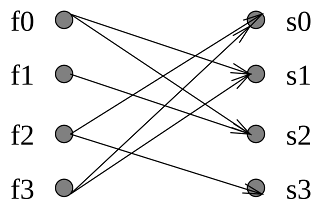

Figure 14: The bipartite graph G corresponding to the star network and directed routes of Figure 13(b).

Lemma 4, 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 14) 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 to when e_i is the second link.) Bipartite graph G has edges $b_0, ..., b_{m-1}$, where for $i = 0, 1, ..., m-1$, b_i is between f_j and s_k if e_j and e_k are the first and second links, respectively, that route p_i traverses. Note that each vertex of G has at most $\frac{W}{2}$ incident edges because R is balanced, i.e., at most $\frac{W}{2}$ routes in R_2 use a link as their first (resp., second) link. 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 of any 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 [25]. 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. $\Box$

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 8 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. □

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 $\{g_0, g_1, ..., g_{W/2-1}\}$ attached to channels at $\{h_0, h_1, ..., h_{W/2-1}\}$ as *patch nodes*. Note that patch nodes have wavelength degree $\frac{W}{2}$.

Theorem 9 *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.*

Outline of Proof: For each request R, a channel assignment can be determined as follows. First, find a balanced set of directions for R, which is possible by Lemma 4. Next, "cut" routes into smaller "residual routes" at the patch nodes. This leaves a collection of uncut and residual routes that traverse at most two links. In addition, all of the routes that traverse two links have an intermediate node that has FCWP. This leaves a collection of "star networks" with hubs at nodes having FCWP. The proof of Theorem 9 implies that channel assignments can be found each set of routes corresponding to a "star network" (here, we assume that the directions of the residual and uncut routes that are used for the channel assignment are the same ones computed earlier for R). The resulting channel assignments can be "patched"

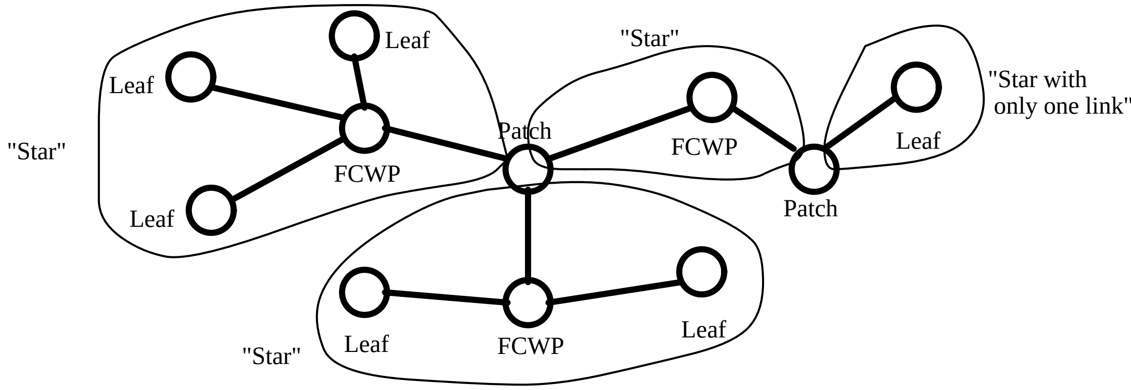

Figure 15: A tree network with leaf nodes, patch nodes, and nodes with FCWP. The neighbor nodes of nodes with FCWP are leaf and patch nodes.

together at the patch nodes to form a channel assignment for R because the patch nodes have the necessary channel attachments. □

See Figure 15 for an example tree network described in Theorem 9. The figure also shows the "star networks" described in the proof of the theorem. It should be noted that since the tree is a bipartite graph, a tree network requires at most half of its nodes to be patch nodes for it to satisfy the conditions of the theorem.

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 8 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 8 (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] 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.

[2] G. K. Chang, G. Ellinas, J. K. Gamelin, M. Z. Iqbal, and C. A. Brackett, "Multiwavelength reconfigurable WDM/ATM/SONET network testbed," *IEEE/OSA J. Lightwave Tech./IEEE JSAC*, vol. 14, June 1996.

[3] 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, pp. 800–813, June 1996.

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

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

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

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

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

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

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

[11] 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 Infocom'94.

[12] 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*, June 1996. An earlier version appeared in Proc. Infocom'95.

[13] A. Birman, "Computing approximate blocking probabilities for a class of optical networks," *IEEE JSAC/JLT Special Issue on Optical Networks*, June 1996. An earlier version appeared in Proc. Infocom'95.

[14] N. Waters and P. Demeester, "Design of the optical path layer in multiwavelength cross-connected networks," *IEEE JSAC/JLT Special Issue on Optical Networks*, June 1996.

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

[16] 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.

[17] J. Yates, J. Lacey, D. Everitt, and M. Summerfield, "Limited-range wavelength translation in all-optical networks," in *Proc. Infocom*, pp. 954–961, 1996.

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

[19] O. Gerstel and S. Kutten, "Dynamic wavelengh allocation in WDM ring networks," in *Submitted to Infocom'97*, 1997.

[20] O. Gerstel, R. Ramaswami, and G. H. Sasaki, "Fault tolerant WDM rings with limited wavelength conversion," in *Submitted to Infocom'97*, 1996.

[21] 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.

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

[23] V. Beneš, *Mathematical Theory of Connecting Networks.* Academic Press, 1965.

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

[25] T. Inukai, "An efficient SS/TDMA time slot assignment algorithm," *IEEE Trans. on Commun.*, vol. 27, pp. 1449–1455, Oct. 1979.