

On-Line Routing and Wavelength Assignment for Dynamic Traffic in WDM Ring and Torus Networks

Poompat Saengudomlert, Eytan H. Modiano, and Robert G. Gallager
 Laboratory for Information and Decision Systems
 Massachusetts Institute of Technology

Abstract— We develop on-line routing and wavelength assignment (RWA) algorithms for WDM bidirectional ring and torus networks with N nodes. The algorithms dynamically support all k -allowable traffic matrices, where k denotes an arbitrary integer vector $[k_1, k_2, \dots, k_N]$, and node i , $1 \leq i \leq N$, can transmit at most k_i wavelengths and receive at most k_i wavelengths. Both algorithms support the changing traffic in a rearrangeably nonblocking fashion. Our first algorithm, for a bidirectional ring, uses $\lceil (\sum_{i=1}^N k_i)/3 \rceil$ wavelengths in each fiber and requires at most three lightpath rearrangements per new session request regardless of the number of nodes N and the amount of traffic k . When all the k_i 's are equal to k , the algorithm uses $\lceil kN/3 \rceil$ wavelengths, which is known to be the minimum for any off-line rearrangeably nonblocking algorithm. Our second algorithm, for a torus topology, is an extension of a known off-line algorithm for the special case with all the k_i 's equal to k . For an $R \times C$ torus network with $R \geq C$ nodes, our on-line algorithm uses $\lceil kR/2 \rceil$ wavelengths in each fiber, which is the same as in the off-line algorithm, and is at most two times a lower bound obtained by assuming full wavelength conversion at all nodes. In addition, the on-line algorithm requires at most $C - 1$ lightpath rearrangements per new session request regardless of the amount of traffic k . Finally, each RWA update requires solving a bipartite matching problem whose time complexity is only $O(R)$, which is much smaller than the time complexity $O(kCR^2)$ of the bipartite matching problem for an off-line algorithm.

Index Terms— WDM networks, routing and wavelength assignment, graph theory.

I. INTRODUCTION

In a wavelength division multiplexed (WDM) network, the fiber bandwidth is divided into multiple frequency bands often called wavelengths. Using reconfigurable optical switches at the network nodes, some wavelengths can be selected at each node for termination and electronic processing, and others selected for optical bypass. In an *all-optical network* architecture, each traffic session optically bypasses electronic processing at each node on its path other than the source node and the destination node. One important benefit of this architecture is the cost saving resulting from using fewer and/or smaller

electronic switches in the network. We focus our attention on all-optical networks in this paper.

Without optical wavelength conversion, routing of traffic sessions is subjected to the *wavelength continuity constraint*, which dictates that the lightpath corresponding to a given session must travel on the same wavelength on all links from the source node to the destination node. Using wavelength converters potentially allows the network to support a larger set of traffic. However, such converters are likely to be expensive. Hence, we focus on the problem of routing and wavelength assignment (RWA) without wavelength converters.

A large body of literature investigates the RWA problem under the wavelength continuity constraint. We can categorize existing results into two groups based on whether static or dynamic provisioning of routes and wavelengths is performed. For static provisioning, the traffic to be supported is assumed known and fixed over time. The goal is often to minimize the number of wavelengths used in the network [1], [2], or to maximize the number of supported traffic sessions for a fixed number of wavelengths [3], [4], [5], [6]. These problems are known to be NP-complete [3]. Consequently, bounds on the optimal costs have been derived [4], [7], and several RWA heuristics have been developed [1], [4], [5], [6], [8], [9].

For dynamic provisioning, we allow the traffic to change over time through session arrivals and session departures. To model dynamic traffic, session arrivals can be assumed to form stochastic processes [10], [11]. In addition, session lifetimes are stochastic. The goal is usually to develop an on-line RWA algorithm which minimizes the average blocking probability for a new session request given a fixed number of wavelengths in the network. We refer to this type of problem formulation as the *blocking formulation*. Due to the complexity in computing blocking probabilities, some approximations are made to simplify the analysis. For example, session arrivals on different links are assumed to be independent [10], [12], or correlated among adjacent links in the same fashion throughout the network [11]. Based on such approximations, several dynamic RWA heuristics have been developed [13], [14].

Another type of problem formulation, referred to as the *nonblocking formulation*, assumes prior knowledge

of the set of all the traffic matrices, or equivalently the traffic demands, to be supported [15], [16], [17], [19], [20]. In [15], the set of traffic matrices is characterized by the maximum link load in the network. In [16], [17], the set of traffic matrices is characterized by the numbers of tunable transmitters and tunable receivers at each node. A new session is said to be *allowable* if its arrival results in a traffic matrix which is still in the set of supportable traffic. The goal is usually to develop an on-line RWA algorithm which does not block any allowable session and uses the minimum number of wavelengths.

If we allow some existing lightpaths to be rearranged in order to support a new session, the corresponding RWA algorithm is said to be *rearrangeably nonblocking*. If we allow no rearrangement of any existing lightpath in order to support a new session, the corresponding RWA algorithm is said to be *wide-sense nonblocking*. Note that if a RWA algorithm is wide-sense nonblocking, it is also rearrangeably nonblocking. Thus, for the same set of traffic matrices, the required number of wavelengths can be no smaller for a wide-sense nonblocking RWA algorithm than for a rearrangeably nonblocking RWA algorithm.

We shall adopt a rearrangeably nonblocking formulation of the RWA problem. As in [16], [17], the supportable traffic set is defined by the number of tunable transmitters and tunable receivers at each end node. We model the traffic as a session-by-session arrival and departure process in which sessions arrive and depart one at a time, and each session utilizes a full wavelength. Our goal is to design an on-line RWA algorithm which is rearrangeably nonblocking, uses the minimum number of wavelengths, and requires few rearrangements of existing lightpaths in order to support each new and allowable session. While our ultimate goal is to perform RWA in an arbitrary mesh topology, we shall focus on a bidirectional ring topology and a torus topology in this paper. In the future, we aim to extend our analytical techniques to obtain a tractable RWA algorithm for arbitrary mesh topologies.

The main contribution of our work is the development of on-line RWA algorithms for supporting dynamic traffic. While there are several results on efficient static RWA algorithms, less is known about efficient dynamic RWA algorithms. For each topology, our on-line algorithm uses the same number of wavelengths as the best-known off-line algorithm and is advantageous in two ways. First, the on-line algorithm guarantees that, for each RWA update due to a traffic change, only a small fraction of existing lightpaths are rearranged. Second, for each RWA update, applying the on-line algorithm instead of the off-line algorithm yields lower computational complexity.

This paper is organized as follows. In section II, we define the set of \mathbf{k} -allowable traffic based on the number of tunable transmitters and tunable receivers at each end node, and formulate the RWA problem for \mathbf{k} -allowable traffic. In section III, we describe our on-line RWA algorithm for a bidirectional ring topology. Section IV contains our on-line RWA algorithm for a torus topology. Finally, we summarize the results and point out future research directions in section V.

II. PROBLEM FORMULATION

Consider an all-optical WDM network with no wavelength conversion. Adjacent nodes are connected by two fibers, one in each direction. In addition, all fibers contain the same number of wavelengths. Assume that each traffic session has a rate of one wavelength. At a given time, only one session can use a specific wavelength in a fiber, but multiple sessions can travel through the same node. Let N be the number of nodes in the network. Node i , $1 \leq i \leq N$, is equipped with k_i fully tunable transmitters and k_i fully tunable receivers. Consequently, at any time, node i can transmit at most k_i wavelengths and receive at most k_i wavelengths. Such a traffic matrix is said to belong to a set of *\mathbf{k} -allowable traffic*, where $\mathbf{k} = [k_1, k_2, \dots, k_N]$.

We model dynamic traffic as a session-by-session arrival and departure process in which sessions arrive and depart one at a time. In other words, a transition from one traffic matrix to another is a result of either a single arrival or a single departure. A new session request is allowable if the resultant traffic matrix is still in the set of \mathbf{k} -allowable traffic. The definition implies that, for each new and allowable session request, there is a free transmitter at the source node and a free receiver at the destination node. For convenience, throughout the paper, a new session is assumed to be allowable unless it is explicitly stated otherwise.

We want to design an on-line RWA algorithm which supports \mathbf{k} -allowable traffic in a rearrangeably nonblocking fashion, uses the minimum number of wavelengths, and requires few rearrangements of existing lightpaths in order to support each new session request. Our algorithm will be centralized in nature. We assume that traffic does not change too frequently and the algorithm always has correct knowledge of the current RWA in the network. In addition, we assume there is sufficient time for lightpath rearrangements between consecutive transitions of the traffic matrix.

We shall consider two regular topologies, a bidirectional ring topology and a torus topology. Fig. 1 illustrates the two topologies. In either topology, each node is considered an end node, i.e. it sources and/or

sinks traffic as well as passes intermediate traffic. Since a bidirectional ring topology is widely used, its investigation is an important practical problem. Although the torus topology may not be implemented in practice, its investigation should give us better understanding of the RWA problem for dynamic traffic in a more densely connected network.

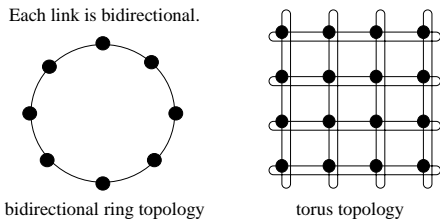


Fig. 1. Bidirectional ring and torus topologies.

III. BIDIRECTIONAL RING TOPOLOGY

In this section, we present an on-line RWA algorithm for k -allowable traffic in an N -node bidirectional ring. Let W_k denote the minimum size of a set of wavelengths such that if each wavelength is provided in each fiber, we can support k -allowable traffic with no wavelength conversion. Note that W_k is the number of wavelengths used to support *any* traffic matrix in the k -allowable set. Thus, for a specific traffic matrix, we may need fewer wavelengths than in the worst case. In [17], it was shown that, if all the k_i 's are equal to k , then $W_k = \lceil kN/3 \rceil$ for $N \geq 7$. In addition, an off-line RWA algorithm that uses at most $\lceil kN/3 \rceil$ wavelengths in each fiber, or equivalently in each ring direction, was developed. In general, we can show that, if all the k_i 's are equal to k , then [18]

$$W_k = \begin{cases} \lceil 3k/4 \rceil, & N = 3, \\ k, & N = 4, \\ \lceil 5k/3 \rceil, & N = 5, 6, \\ \lceil Nk/3 \rceil, & N \geq 7. \end{cases}$$

We shall present an on-line RWA algorithm that uses $\lceil (\sum_{i=1}^N k_i)/3 \rceil$ wavelengths in each fiber to support k -allowable traffic. Note that, for $N \geq 7$, when all the k_i 's are equal to k , the algorithm uses the minimum number of wavelengths found in [17]. In all the other cases, the algorithm yields the upper bound $W_k \leq \lceil (\sum_{i=1}^N k_i)/3 \rceil$.

Define a *directed wavelength* as a wavelength in either the clockwise or the counterclockwise ring direction. Given w wavelengths in each fiber, there are w directed wavelengths in the clockwise ring direction, and w directed wavelengths in the counterclockwise ring direction. Note that any traffic session can be supported on a directed wavelength in either ring direction. Two sessions are said to be *adjacent* if the destination node

of one session is the source node of the other. The main idea behind our algorithm involves sharing a directed wavelength between two adjacent sessions, as suggested by the following known lemma in [17].

Lemma 1: In a bidirectional ring, lightpaths corresponding to any pair of adjacent sessions can either share a directed wavelength in the clockwise ring direction or the counterclockwise ring direction.

The proof of lemma 1 is immediate from Fig. 2, where if two lightpaths overlap in one direction, they do not overlap in the other. In particular, the lightpaths corresponding to sessions (1,4) and (4,2) overlap in the clockwise direction, but do not overlap in the counterclockwise direction.¹ In what follows, when the lightpaths associated with a pair of adjacent sessions share a directed wavelength, we simply say that the adjacent session pair share a directed wavelength.

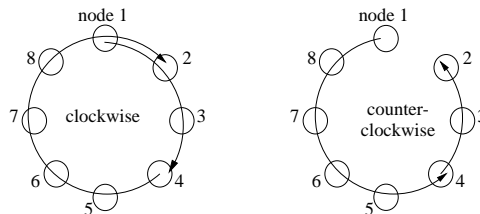


Fig. 2. Adjacent sessions share a directed wavelength.

The main idea of our algorithm is to maintain the following two RWA conditions at all times: (i) only adjacent sessions share a directed wavelength, and (ii) at most two adjacent sessions share a directed wavelength.

To give some intuition on the main idea of our algorithm, consider the special case with all the k_i 's equal to 1. In this case, our algorithm uses $\lceil N/3 \rceil$ wavelengths in each fiber. We next describe informally how to use $\lceil N/3 \rceil$ wavelengths to support the traffic. We ignore integer rounding in the informal discussion below.

Given a traffic matrix, form as many adjacent session pairs as possible up to $N/3$ pairs in a greedy fashion, i.e. it does not matter if we end up with less than the maximum possible number of pairs. Let p denote the number of adjacent session pairs formed. Consider two cases.

Case 1: $p = N/3$. In this case, we support $N/3$ adjacent session pairs containing $2N/3$ sessions on $N/3$ directed wavelengths in the required ring directions. This is always possible since there are $N/3$ directed wavelengths available in each ring direction. Having done so, there are at most $N - 2N/3 = N/3$ remaining sessions each of which we support on one directed wavelength in

¹A session from node i to node j is denoted by session (i, j) .

either ring direction. Thus, the total number of directed wavelengths required is at most $N/3 + N/3 = 2N/3$. It follows that $N/3$ wavelengths in each fiber, i.e. each ring direction, are sufficient.

Case 2: $p < N/3$. In this case, we support p adjacent session pairs containing $2p$ sessions on p directed wavelengths in the required ring directions. This is always possible since there are $N/3$ directed wavelengths available in each ring direction. Note that we cannot form any more adjacent session pairs in this case.

Consider only the sessions in the p adjacent session pairs formed above. Define a *common node* to be a node which transmits a wavelength and receives a wavelength. Observe that each adjacent session pair has at least one common node. For example, Fig. 3 shows two ($p = 2$) adjacent session pairs (7,4) and (4,3) together with (1,8) and (8,7). The pair (7,4) and (4,3) has node 4 as a common node, while the pair (1,8) and (8,7) has node 8 as a common node. In addition, sessions (8,7) and (7,4) make node 7 a common node. In general, given p adjacent session pairs, there are at least p common nodes.

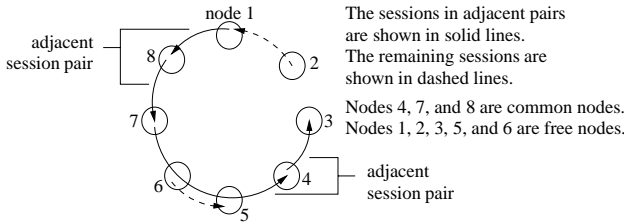


Fig. 3. Adjacent session pairs, common nodes, and free nodes.

Define a node which is not a common node as a *free node*. A free node still has a free transmitter and/or a free receiver. For example, in Fig. 3, nodes 1, 2, 3, 5, and 6 are free nodes. Since there are at least p common nodes, there are at most $N - p$ free nodes.

Consider the remaining sessions which are not in the p adjacent session pairs formed above. Observe that each free node terminates, i.e. either transmits or receives, at most one remaining session. To see this, note that each free node cannot transmit more than one remaining session since it only has one transmitter. By the same argument, each free node cannot receive more than one remaining session. Moreover, each free node cannot transmit a remaining session and receive a remaining session simultaneously, or else we could form another new adjacent session pair, i.e. have more than p pairs. Thus, each remaining session is terminated at two distinct free nodes. For example, in Fig. 3, the remaining session (2,1) is terminated at free nodes 1 and 2. No other remaining session is terminated at either node 1 or

node 2. Since there are at most $N - p$ free nodes, there are at most $(N - p)/2$ remaining sessions. We support each remaining session on one directed wavelength in either ring direction. Thus, the total number of directed wavelengths required is $p + (N - p)/2 = N/2 + p/2 < N/2 + N/6 = 2N/3$. It follows that $N/3$ wavelengths in each fiber are sufficient.

We shall later prove by similar arguments that $\lceil (\sum_{i=1}^N k_i)/3 \rceil$ wavelengths are sufficient to support k -allowable traffic. We now describe our on-line RWA algorithm which is rearrangeably nonblocking, uses $\lceil (\sum_{i=1}^N k_i)/3 \rceil$ wavelengths in each fiber, and requires at most three lightpath rearrangements per new session request. We shall refer to this algorithm as the *ring RWA algorithm*.

Ring RWA Algorithm: (Use $\lceil (\sum_{i=1}^N k_i)/3 \rceil$ wavelengths in each fiber.)

Session termination: When a session terminates, simply remove its associated lightpath from the ring without any further lightpath rearrangement.

Session arrival: When a session arrives and the resultant traffic matrix is still k -allowable, proceed as follows.

Step 1: If there is a nonsharing session, i.e. a session which does not share its directed wavelength with any session, and it is adjacent to and can share its directed wavelength with the new session, assign the two sessions to share that directed wavelength. In this case, no lightpath rearrangement is required. Otherwise, proceed to step 2.

Step 2: If there is a free directed wavelength in either ring direction, assign a free directed wavelength to the new session. In this case, no lightpath rearrangement is required. Otherwise, proceed to step 3.

Step 3: Among the nonsharing sessions and the new session, we claim and shall prove shortly that there must exist a pair of adjacent sessions. Form such an adjacent session pair by searching through all pairs of sessions in some order, e.g. from sessions terminating at node 1 to sessions terminating at node N . Once an adjacent session pair is found, there are two possibilities.

(3a) If the adjacent session pair can share the directed wavelength of one session in the pair, assign the adjacent session pair to share that directed wavelength. In this case, the adjacent session pair does not include the new session since step 1 would have otherwise applied. Therefore, one existing lightpath must be rearranged. Sharing of the directed wavelength by the adjacent session pair will free one directed wavelength on which

the new session can be supported. Fig. 4a illustrates this scenario. In particular, existing sessions (1,5) and (5,2) form an adjacent session pair which can be supported on the directed wavelength of session (5,2). After the lightpath of session (1,5) is rearranged, the new session (1,4) is supported on the directed wavelength previously used by session (1,5).

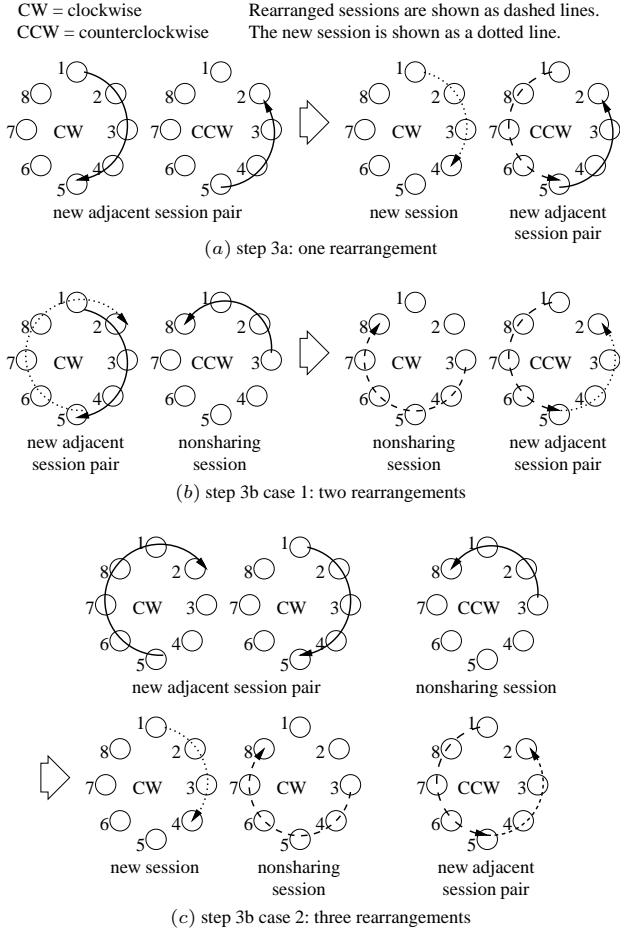


Fig. 4. Step 3 of the ring RWA algorithm.

(3b) If the adjacent session pair cannot share the directed wavelength of either session in the pair, we claim and shall prove shortly that there must exist a directed wavelength with a nonsharing session in the opposite ring direction, i.e. the ring direction in which the adjacent session pair can share a directed wavelength. Remove the lightpath of that nonsharing session from its directed wavelength, and assign the adjacent session pair to share that directed wavelength. When the adjacent session pair includes the new session, the new session will now be supported, and sharing of the directed wavelength by the adjacent session pair will free one directed wavelength on which the removed nonsharing session can be supported. In this case, a total of two lightpath rearrangements are made. Fig. 4b illustrates this scenario. In particular, existing session (1,5) and the new

session (5,2) form an adjacent session pair which can be supported on the directed wavelength of existing session (3,8). After the lightpaths of sessions (1,5) and (3,8) are rearranged, the new session (5,2) shares a directed wavelength with session (1,5) on the directed wavelength previously used by session (3,8), while session (3,8) is supported on the directed wavelength previously used by session (1,5).

When the adjacent session pair does not include the new session, sharing of the directed wavelength by the adjacent session pair will free two directed wavelengths on which the removed nonsharing session and the new session can be supported. In this case, a total of three lightpath rearrangements are made. Fig. 4c illustrates this scenario. In particular, existing sessions (1,5) and (5,2) form an adjacent session pair which can be supported on the directed wavelength of existing session (3,8). After the lightpaths of sessions (1,5), (5,2), and (3,8) are rearranged, the adjacent session pair (1,5) and (5,2) are supported on the directed wavelength previously used by session (3,8), session (3,8) is supported on the directed wavelength previously used by session (1,5), and the new session (1,4) is supported on the directed wavelength previously used by session (5,2).

Before proving the correctness of the ring RWA algorithm, we establish two useful lemmas related to step 3 of the algorithm. In what follows, let p be the number of adjacent session pairs which share a directed wavelength before the new session request. Let q be the number of nonsharing sessions before the new session request. The two lemmas give upper bounds on p and q respectively. In addition, let w be the number of wavelengths in use before the new session request. Note that $w = p + q$. For convenience, define $K = \sum_{i=1}^N k_i$.

Lemma 2: In step 3 of the ring RWA algorithm, $p < \lfloor K/3 \rfloor$.

Proof: Since the total number of sessions is at most K in \mathbf{k} -allowable traffic, it follows that $2p + q < K$ before the new session request. Thus, w is bounded by

$$w = p + q < p + (K - 2p) = K - p.$$

In step 3, since there is no free directed wavelength for the new session, it follows that the number of wavelengths in use w is equal to the total number of directed wavelengths $2\lceil K/3 \rceil$. Therefore, $K - p > w = 2\lceil K/3 \rceil$, yielding the desired relation

$$p < K - 2\lceil K/3 \rceil \leq \lfloor K/3 \rfloor. \quad \square$$

Lemma 3: In step 3 of the ring RWA algorithm, if no adjacent session pair can be formed among the nonsharing sessions and the new session, then $q \leq \lfloor (K - p)/2 \rfloor$.

Proof: Note that node i , $1 \leq i \leq N$, is equipped with k_i tunable transmitter/receiver pairs. Overall, we have a total of K transmitter/receiver pairs. Each pair of adjacent sessions which share a directed wavelength utilizes one transmitter/receiver pair at some node, one transmitter, and one receiver elsewhere.

Let p_i be the number of adjacent session pairs which share a directed wavelength and have node i as a common node. Since an adjacent session pair may have more than one common node, $\sum_{i=1}^N p_i \geq p$. Let $k'_i = k_i - p_i$ denote the number of transmitter/receiver pairs which are not used by those p_i adjacent session pairs at node i . In addition, let k_i^t and k_i^r denote the numbers of nonsharing sessions transmitted and received at node i respectively. It is clear that $k_i^t \leq k'_i$ and $k_i^r \leq k'_i$.

Since no new adjacent session pair can be formed among the nonsharing sessions, it follows that, at each node i , either $k_i^t = 0$ or $k_i^r = 0$. Thus, $k_i^t + k_i^r \leq k'_i$. Because each nonsharing session uses one transmitter and one receiver, it follows that

$$2q = \sum_{i=1}^N (k_i^t + k_i^r) \leq \sum_{i=1}^N k'_i = K - \sum_{i=1}^N p_i \leq K - p.$$

Since q is an integer, $q \leq \lfloor (K - p)/2 \rfloor$. \square

Proof of algorithm correctness: From the algorithm description, it is clear that we always keep the two desired RWA conditions, i.e. (i) only adjacent sessions share a directed wavelength, and (ii) at most two adjacent sessions share a directed wavelength. In addition, it is clear that at most three lightpath rearrangements are made to support each new session request.

It remains to prove the two claims in step 3. The first claim states that there always exists a new adjacent session pair. We proceed by contradiction. Suppose that no new adjacent session pair can be formed among the nonsharing sessions and the new session. From lemma 3, $q \leq \lfloor (K - p)/2 \rfloor$. Since there is no free directed wavelength for the new session in step 3, it follows that the number of wavelengths in use w is equal to the total number of directed wavelengths $2\lceil K/3 \rceil$. Therefore,

$$p + \lfloor (K - p)/2 \rfloor \geq p + q = w = 2\lceil K/3 \rceil.$$

It follows that

$p \geq 2\lceil K/3 \rceil - \lfloor (K - p)/2 \rfloor \geq 2K/3 - (K - p)/2$, or equivalently, $p \geq K/3$, which contradicts the fact that $p < \lceil K/3 \rceil$ in step 3 from lemma 2. Hence, a new adjacent session pair always exists in step 3.

We now prove the second claim in step 3 that if we need to find a nonsharing session in the opposite ring direction, i.e. the ring direction in which the new adjacent session pair can share a directed wavelength, one always

exists. The claim is a direct consequence of lemma 2, i.e. $p < \lceil K/3 \rceil$ in step 3. In other words, the number of sharing session pairs is less than the number of directed wavelengths in each ring direction. Since step 2 was not taken, all the other $2\lceil K/3 \rceil - p$ directed wavelengths are taken by nonsharing sessions. Therefore, in either ring direction, a directed wavelength with a nonsharing session exists. \square

The construction of the ring RWA algorithm yields the upper bound $W_{\mathbf{k}} \leq \lceil (\sum_{i=1}^N k_i)/3 \rceil$. We now derive a simple lower bound on $W_{\mathbf{k}}$. Let $L_{\mathbf{k}}$ denote the minimum number of wavelengths which, if provided in each fiber, can support \mathbf{k} -allowable traffic given full wavelength conversion at all nodes. For a ring network, any pair of links correspond to a cut, denoted by \mathcal{C} , which separates the N nodes into two sets, denoted by $\mathcal{N}_{\mathcal{C},1}$ and $\mathcal{N}_{\mathcal{C},2}$. The maximum possible traffic, in wavelength units, across the cut from one set of nodes to the other is equal to $\min(\sum_{i \in \mathcal{N}_{\mathcal{C},1}} k_i, \sum_{i \in \mathcal{N}_{\mathcal{C},2}} k_i)$. Since there are two fibers across the cut from one set of nodes to the other, one fiber must contain at least $\lceil \min(\sum_{i \in \mathcal{N}_{\mathcal{C},1}} k_i, \sum_{i \in \mathcal{N}_{\mathcal{C},2}} k_i)/2 \rceil$ wavelengths. The maximum possible traffic over all the cuts yields the lower bound on $L_{\mathbf{k}}$ given below.

Lemma 4: For a bidirectional ring with N nodes and \mathbf{k} -allowable traffic, $L_{\mathbf{k}}$ is bounded by $L_{\mathbf{k}} \geq \max_{\mathcal{C}} \lceil \min(\sum_{i \in \mathcal{N}_{\mathcal{C},1}} k_i, \sum_{i \in \mathcal{N}_{\mathcal{C},2}} k_i)/2 \rceil$.

Since $L_{\mathbf{k}} \leq W_{\mathbf{k}}$, we can use the lower bound on $L_{\mathbf{k}}$ in lemma 4 as a lower bound on $W_{\mathbf{k}}$. However, this lower bound is not tight in general. For example, when $N \geq 7$ and all the k_i 's are equal to k , we know from [17] that $W_{\mathbf{k}} = \lceil kN/3 \rceil$, but the lower bound from lemma 4 is $W_{\mathbf{k}} \geq \lceil k \lfloor N/2 \rfloor / 2 \rceil$. Note that this lower bound is obtained from the cut which separates any consecutive $\lfloor N/2 \rfloor$ nodes from the other $\lceil N/2 \rceil$ nodes.

We summarize the results of this section in the theorem below.

Theorem 1: For a bidirectional ring with N nodes and \mathbf{k} -allowable traffic,

$$\begin{aligned} \max_{\mathcal{C}} \left\lceil \frac{\min(\sum_{i \in \mathcal{N}_{\mathcal{C},1}} k_i, \sum_{i \in \mathcal{N}_{\mathcal{C},2}} k_i)}{2} \right\rceil &\leq L_{\mathbf{k}} \\ &\leq W_{\mathbf{k}} \leq \left\lceil \frac{\sum_{i=1}^N k_i}{3} \right\rceil. \end{aligned}$$

In addition, there exists, by construction, an on-line RWA algorithm which uses $\lceil (\sum_{i=1}^N k_i)/3 \rceil$ wavelengths in each fiber and requires at most three lightpath rearrangements per new session request.

We have pointed out that, when $N \geq 7$ and all the k_i 's are equal to k , the upper bound in theorem 1 is tight but the lower bound is not. An interesting example in which the upper bound is not tight but the lower bound is tight is an N -node bidirectional ring which contains one hub node, say node 1, with $k_1 = N - 1$, and the other $N - 1$ nodes each with $k_i = 1$. In this case, the lower bound in theorem 1 corresponds to the cut which separates the hub node from all the other nodes, and is given by $W_{\mathbf{k}} \geq \lceil (N - 1)/2 \rceil$. In the next subsection, we construct a specialized on-line RWA algorithm to show that $W_{\mathbf{k}} \leq \lceil (N - 1)/2 \rceil$, yielding the following lemma.

Lemma 5: For an N -node bidirectional ring with $k_1 = N - 1$ and $k_2 = k_3 = \dots = k_N = 1$, $W_{\mathbf{k}} = \lceil (N - 1)/2 \rceil$.

Proof: The above discussion shows that $W_{\mathbf{k}} \geq \lceil (N - 1)/2 \rceil$. The proof of $W_{\mathbf{k}} \leq \lceil (N - 1)/2 \rceil$ is in the next subsection. \square

A. Single-Hub Ring RWA Algorithm

Consider a bidirectional ring with N nodes. In particular, node 1 acts as a hub node with $k_1 = N - 1$. In addition, for $2 \leq i \leq N$, $k_i = 1$. Note that the non-hub nodes can directly transmit and/or receive wavelengths among themselves.

We shall first give an informal argument to show that $W_{\mathbf{k}} \leq \lceil (N - 1)/2 \rceil$. Afterwards, we shall present a formal proof based on an on-line RWA algorithm which uses $\lceil (N - 1)/2 \rceil$ wavelengths.

As in the ring RWA algorithm, the main idea in this section involves sharing of a directed wavelength by an adjacent session pair. As a reminder, two sessions are said to be adjacent if the destination node of one session is the source node of the other. Two sessions are said to be mutually adjacent if the destination node of one session is the source node of the other and vice versa. For convenience, we shall call two sessions which are adjacent but not mutually adjacent a *nonmutual adjacent* session pair. While a nonmutual adjacent session pair can share a directed wavelength in only one ring direction, a mutual adjacent session pair can share a directed wavelength in either ring direction.

We shall refer to an adjacent session pair which has the hub node as a common node as an adjacent session pair *at the hub*. Our RWA is based on the following two RWA conditions: (i) only adjacent session pairs at the hub share a directed wavelength, and (ii) all mutual adjacent session pairs at the hub share a directed wavelength. We first give an informal proof that $\lceil (N - 1)/2 \rceil$ wavelengths are sufficient to support the traffic. We ignore integer rounding in the informal discussion below.

Given a traffic matrix, form all the mutual adjacent session pairs at the hub, but do not assign directed wavelengths for them at this point. Then form all the nonmutual adjacent session pairs at the hub in a greedy fashion. Let r and s denote the numbers of mutual and nonmutual adjacent session pairs at the hub respectively. Let t be the number of the remaining sessions. Note that we cannot form another adjacent session pair at the hub among these t sessions.

We first support the s nonmutual adjacent session pairs at the hub on s directed wavelength in the required ring directions. We now show this is always possible. Observe that each non-hub node terminates, i.e. transmits or receives, at most one session in these s adjacent pairs. To see this, note that each non-hub node cannot transmit more than one session since it only has one transmitter. By the same argument, each non-hub node cannot receive more than one session. Moreover, each non-hub node cannot transmit a session and receive a session in these s adjacent pairs simultaneously, or else we can form another mutual adjacent session pair at the hub. It follows that each nonmutual adjacent session pair at the hub is terminated at two non-hub nodes, and no other nonmutual adjacent session pair at the hub is terminated at any of these two nodes. Since there are $N - 1$ non-hub nodes, it follows that $s \leq (N - 1)/2$. Since there are $(N - 1)/2$ directed wavelengths available in each ring direction, there are enough wavelengths to support the s session pairs.

We next support the r mutual adjacent session pairs at the hub on any r unused directed wavelengths. We now show this is always possible. Note that each mutual adjacent session pair at the hub is terminated at one distinct non-hub node. From the above discussion, each nonmutual adjacent session pair at the hub is terminated at two distinct non-hub nodes. Since there are $N - 1$ non-hub nodes, it follows that $r + 2s \leq N - 1$, or equivalently $r \leq (N - 1) - 2s$. Since there are $(N - 1) - s$ unused directed wavelengths left for this step, the inequality $r \leq (N - 1) - 2s$ implies that there are enough directed wavelengths to support the r session pairs.

In the final step, we support the t remaining sessions on any t unused directed wavelengths. We now show this is always possible. Since we cannot form any adjacent session pair at the hub from these t sessions, the hub node can either transmit or receive some or all of these t sessions but not both. Without loss of generality, assume that the hub node transmits none of these t sessions. Consider the transmitters at the non-hub nodes. Each of the r mutual adjacent session pairs at the hub uses one transmitter at some non-hub node. Similarly, each of the s nonmutual adjacent session pairs at the hub uses one

transmitter at some non-hub node. Since the hub node does not transmit any of the t remaining sessions, each of the t sessions uses one transmitter at some non-hub node. Since there are $N-1$ non-hub nodes, it follows that $r+s+t \leq N-1$, or equivalently $t \leq (N-1)-r-s$. Since there are $(N-1)-r-s$ unused directed wavelengths left for this step, there are enough directed wavelengths to support the remaining t sessions.

Based on the above discussion, we now formally present an on-line RWA algorithm which uses $\lceil (N-1)/2 \rceil$ wavelengths in each fiber, is rearrangeably non-blocking, and requires at most four lightpath rearrangements per new session request. We shall refer to this algorithm as the *single-hub ring RWA algorithm*.

Single-Hub Ring RWA Algorithm: (Use $\lceil (N-1)/2 \rceil$ wavelengths in each fiber.)

Session termination: When a session terminates, simply remove its associated lightpath from the ring without any further lightpath rearrangement.

Session arrival: When a session arrives and the resultant traffic matrix is still k -allowable, proceed as follows.

Step 1: If the new session, denoted by u , can form a mutual adjacent session pair at the hub with some existing session, denoted by x , there are two possibilities.

(1a) If x is not sharing its directed wavelength, assign the mutual adjacent session pair u and x to share this directed wavelength. In this case, no lightpath rearrangement is required.

(1b) If x is sharing a directed wavelength with another existing session, denoted by y , then x and y are not mutually adjacent at the hub, or else u and x cannot be mutually adjacent at the hub. Remove y from its directed wavelength and assign the mutual adjacent session pair u and x to share the directed wavelength of y .

If there is a free directed wavelength, use it to support y . In this case, one lightpath rearrangement is made. Otherwise, we claim and shall prove shortly that y can form another adjacent session pair at the hub with some nonsharing session, denoted by z . Note that y and z cannot be mutually adjacent at the hub, or else they would have shared a directed wavelength.

If the directed wavelength of z can support y , assign y and z to share this directed wavelength. In this case, one lightpath rearrangement is made. Otherwise, we claim and shall prove shortly that there must exist either a nonsharing session or a mutual adjacent session pair in the opposite ring direction. In the case of a nonsharing session in the opposite ring direction, we remove that nonsharing session and support y and z on its directed wavelength. The removed nonsharing session can then be

supported on the directed wavelength of z . In this case, a total of three lightpath rearrangements are made. In the case of a mutual adjacent session pair in the opposite ring direction, we remove that mutual adjacent session pair and support y and z on their directed wavelength. The removed mutual adjacent session pair can then be supported on the directed wavelength of z . In this case, a total of four lightpath rearrangements are made.

Step 2: If u cannot form a mutual adjacent session pair at the hub with any existing session and there is a free directed wavelength, use a free directed wavelength to support u . In this case, no lightpath rearrangement is made.

Step 3: If u cannot form a mutual adjacent session pair at the hub with any existing session and there is no free directed wavelength, we claim and shall prove shortly that, among nonsharing sessions and u , a nonmutual adjacent session pair at the hub can be formed. Denote this session pair by y and z . There are two possibilities.

(3a) If u is in the session pair, i.e. $y = u$ or $z = u$, assume without loss of generality that $y = u$. If the directed wavelength of z can support y , assign y and z to share this directed wavelength. In this case, no lightpath rearrangement is required. Otherwise, we claim and shall prove shortly that there must exist either a nonsharing session or a mutual adjacent session pair in the opposite ring direction. In the case of a nonsharing session in the opposite ring direction, we remove that nonsharing session and support y and z on its directed wavelength. The removed nonsharing session can then be supported on the directed wavelength of z . In this case, a total of two lightpath rearrangements are made. In the case of a mutual adjacent session pair in the opposite ring direction, we remove that mutual adjacent session pair and support y and z on their directed wavelength. The removed mutual adjacent session pair can then be supported on the directed wavelength of z . In this case, a total of three lightpath rearrangements are made.

(3b) If u is not in the session pair, then $y \neq u$ and $z \neq u$. If the directed wavelength of either y or z can support the session pair, assign y and z to share this directed wavelength. This sharing frees one directed wavelength on which u can be supported. In this case, one lightpath rearrangement is made. Otherwise, we claim and shall prove shortly that there must exist either a nonsharing session or a mutual adjacent session pair in the opposite ring direction. In the case of a nonsharing session in the opposite ring direction, we remove that nonsharing session and support y and z on its directed wavelength. The removed nonsharing session and the new session can then be supported on the directed wavelengths of y and

z . In this case, a total of three lightpath rearrangements are made. In the case of a mutual adjacent session pair in the opposite ring direction, we remove that mutual adjacent session pair and support y and z on their directed wavelength. The removed mutual adjacent session pair and the new session can then be supported on the directed wavelengths of y and z . In this case, a total of four lightpath rearrangements are made.

Proof of algorithm correctness: From the algorithm description, it is clear that we always keep the two desired RWA conditions, i.e. (i) only adjacent sessions at the hub share a directed wavelength, and (ii) all mutual adjacent sessions at the hub share a directed wavelength. In addition, it is clear that at most four lightpath rearrangements are made to support each new session request. We shall prove the two claims in step 1, and the other three claims in step 3.

The first claim in step 1 and the first claim in step 3 are essentially the same. We shall prove the two claims at the same time. The claims state that if a session to be supported, denoted by w , is not mutually adjacent to any existing session at the hub and there is no free directed wavelength to support it, then there exists among nonsharing sessions and w an adjacent session pair at the hub, denoted by y and z .

We proceed by contradiction. Assume that an adjacent session pair at the hub cannot be found. Let p be the number of mutual adjacent session pairs at the hub. Let q be the number of nonmutual adjacent session pairs at the hub which share a directed wavelength. Let r be the number of nonsharing sessions including session w . We argue that $r \leq N - 1 - p - q$. To see this, define r_i^t and r_i^r , $1 \leq i \leq N$, to be the number of nonsharing sessions transmitted and received at node i respectively. Since there is no adjacent session pair at the hub (node 1) among these r sessions, we have that either $r_1^t = 0$ or $r_1^r = 0$. Without loss of generality, assume $r_1^t = 0$. Note that each of the $p + q$ sharing session pairs which are adjacent at the hub uses one transmitter at a nonhub node. There are in total $N - 1$ transmitters at nonhub nodes. Thus, the number of transmitters used for nonsharing sessions at nonhub nodes are bounded by $\sum_{i=2}^N r_i^t \leq N - 1 - p - q$. It follows that

$$r = \sum_{i=1}^N r_i^t = r_1^t + \sum_{i=2}^N r_i^t \leq N - 1 - p - q.$$

Since we have a total of $2\lceil(N-1)/2\rceil$ directed wavelengths, the number of directed wavelengths available to support nonsharing paths is $2\lceil(N-1)/2\rceil - p - q$, which is at least the number of nonsharing paths $N - 1 - p - q$. This contradicts the assumption that there is no free directed

wavelength to support w . Thus, we have shown that an adjacent session pair at the hub must exist.

The second claim in step 1 and the last two claims in step 3 are essentially the same. We shall prove them all at the same time. The claim states that if a nonmutual adjacent session pair at the hub, denoted by y and z , cannot fit on a directed wavelength of either y or z and there is no free directed wavelength in the opposite ring direction, then there exists either a nonsharing session or a mutual adjacent session pair on a directed wavelength in the opposite ring direction. As defined above, let p be the number of mutual adjacent session pairs at the hub. Let \hat{q} be the number of nonmutual adjacent session pairs at the hub including sessions y and z . Note that each of these \hat{q} session pairs may or may not share a directed wavelength. We first show that $\hat{q} \leq \lfloor(N-1)/2\rfloor$. Define the following quantities for node i , $2 \leq i \leq N$. Let \hat{q}_i^t and \hat{q}_i^r denote the number of sessions in those \hat{q} session pairs which are transmitted and received at node i respectively. It is clear that $\hat{q}_i^t \leq k_i$ and $\hat{q}_i^r \leq k_i$. By definition, each of these \hat{q} session pairs is not a mutual adjacent session pair at the hub. Thus, at each nonhub node i , either $\hat{q}_i^t = 0$ or $\hat{q}_i^r = 0$. It follows that $\hat{q}_i^t + \hat{q}_i^r \leq k_i$. Because each of the \hat{q} session pairs uses one transmitter and one receiver at nonhub nodes, it follows that

$$2\hat{q} = \sum_{i=2}^N (\hat{q}_i^t + \hat{q}_i^r) \leq \sum_{i=2}^N k_i = N - 1.$$

Since \hat{q} is an integer, we have shown that $\hat{q} \leq \lfloor(N-1)/2\rfloor$.

The claim is now apparent from the fact that $\hat{q} \leq \lfloor(N-1)/2\rfloor$. In other words, the number of supported nonmutual adjacent session pairs at the hub $\hat{q} - 1$ is strictly less than the number of directed wavelengths in each ring direction $\lceil(N-1)/2\rceil$. Given that there is no free directed wavelength, it follows that, in either ring direction, either a nonsharing session or a mutual adjacent session pair exists. \square

B. On-Line RWA for Bidirectional Sessions

In this subsection, we discuss the special case in which sessions are bidirectional. More precisely, for each node pair i and j , if there are m sessions from i to j , there are also m sessions from j to i .

We define a special kind of adjacent session pairs as follows. Two sessions form a *mutual adjacent* session pair if the source node of one session is the destination node of the other and vice versa. Note that each pair of bidirectional sessions form a mutual adjacent session pair. In addition, a mutual adjacent session pair can share

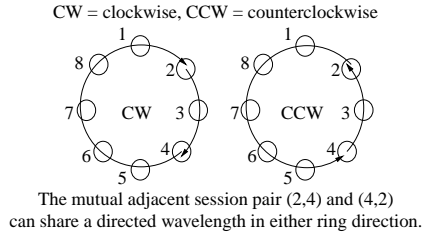


Fig. 5. Supporting a mutual adjacent session pair on a directed wavelength.

a directed wavelength in either ring direction, as shown in Fig. 5 for the pair (2,4) and (4,2).

The number of mutual adjacent session pairs in k -allowable traffic with bidirectional sessions is at most $\lceil (\sum_{i=1}^N k_i)/2 \rceil$. Since each mutual adjacent pair can be supported on one directed wavelength in any ring direction, it follows that $\lceil \lceil (\sum_{i=1}^N k_i)/2 \rceil / 2 \rceil$ wavelengths are sufficient to support k -allowable traffic with bidirectional sessions. In addition, if we assume that each traffic matrix change is a result of arrivals or departures of a pair of bidirectional sessions, then, with $\lceil \lceil (\sum_{i=1}^N k_i)/2 \rceil / 2 \rceil$ wavelengths, new arrivals can use any free directed wavelength and require no rearrangement of existing sessions.

Notice that lemma 4 is still valid for bidirectional sessions since the cut that corresponds to the lower bound can have bidirectional sessions travel across it. Therefore, we have the following bounds on L_k and W_k .

Lemma 6: For a bidirectional ring with N nodes and k -allowable traffic with bidirectional sessions,

$$\max_c \left\lceil \frac{\min \left(\sum_{i \in \mathcal{N}_{c,1}} k_i, \sum_{i \in \mathcal{N}_{c,2}} k_i \right)}{2} \right\rceil \leq L_k$$

$$\leq W_k \leq \left\lceil \left\lceil \frac{\sum_{i=1}^N k_i}{2} \right\rceil / 2 \right\rceil.$$

In addition, if sessions arrive and depart in bidirectional pairs, there exists a wide-sense nonblocking on-line algorithm that uses $\lceil \lceil (\sum_{i=1}^N k_i)/2 \rceil / 2 \rceil$ wavelengths in each fiber.

Note that, when all the k_i 's are equal to k , the bounds in theorem 6 are rather tight, i.e. $\lceil k \lfloor N/2 \rfloor / 2 \rceil \leq W_k \leq \lceil \lceil kN/2 \rceil / 2 \rceil$. For the one-hub ring with $k_1 = N - 1$ and $k_i = 1$ for all other i , the bounds in theorem 6 coincide and $W_k = \lceil (N - 1)/2 \rceil$.

IV. TORUS TOPOLOGY

In this section, we present an on-line RWA algorithm for k -allowable traffic in a torus network. We shall consider only the cases in which all the k_i 's are equal to some integer k . For convenience, we refer to the k -allowable traffic in which all the k_i 's are equal to k as *symmetric k -allowable traffic*.

Consider an $R \times C$ torus network with N nodes, where $N = RC$ and $R \geq C$. If $R < C$, we can reverse the roles of columns and rows and the following discussion remains valid. Let L_k and W_k be the minimum number of wavelengths which, if provided in each fiber, can support symmetric k -allowable traffic with and without wavelength conversion respectively. Note that L_k and W_k are the numbers of wavelengths used to support *any* traffic matrix in the symmetric k -allowable set. Thus, for a specific traffic matrix, we may need fewer wavelengths than in the worst case. In [20], it was shown that L_k and W_k are bounded by

$$\left\lceil \frac{kR}{4} \right\rceil \leq L_k \leq W_k \leq \left\lceil \frac{kR}{2} \right\rceil.$$

The lower bound on L_k was derived by finding the cut with the maximum traffic per fiber across it. The upper bound on W_k was derived by constructing an off-line RWA algorithm based on a bipartite matching problem formulated in [19], [20].² We shall extend the results in [19], [20] to develop an on-line RWA algorithm that uses the same number of wavelengths and provides two advantages. First, for each RWA update due to a traffic change, our algorithm guarantees that at most $C - 1$ existing lightpaths are rearranged. Second, for each RWA update, our algorithm needs to solve a simple bipartite matching that requires the running time $O(R)$. As will be seen shortly, if the off-line algorithm is used for each RWA update, the associated bipartite matching problem requires the running time $O(kCR^2)$.

We first explain the off-line algorithm in [19], [20]. Define a *directed wavelength* in a torus network as follows. Each wavelength consists of an upward directed wavelength and a downward directed wavelength. An upward directed wavelength is directed upwards along any column and to the right along any row. A downward directed wavelength is directed downwards along any column and to the left along any row.

We apply *column-first routing* where each lightpath travels along the source column and then along the destination row. In addition, each lightpath is supported by no more than one directed wavelength, i.e. if it travels upwards along the source column, then it must travel to the right along the destination row. Under column-first routing, a set of sessions from distinct source columns to distinct destination rows can all be supported on a single directed wavelength, which can be either upward or downward directed [20], as shown in Fig. 6.

We can view the set of sessions from distinct source columns to distinct destination rows as a matching in

²Several other resource allocation problems can be casted into bipartite matching problems, e.g. [21].

a bipartite graph [19]. For a given traffic matrix, we can construct the *column-to-row* bipartite graph, denoted by $(\mathcal{V}_1, \mathcal{V}_2, \mathcal{E})$, as follows. The set of abstract nodes \mathcal{V}_1 contains C nodes corresponding to the C source columns. The set of abstract nodes \mathcal{V}_2 contains R nodes corresponding to the R destination rows. In the set of edges \mathcal{E} , an edge between node i in \mathcal{V}_1 and node j in \mathcal{V}_2 corresponds to a session from a source in column i to a destination in row j .

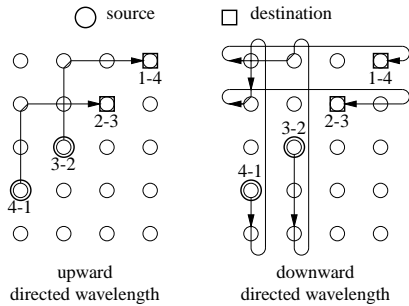


Fig. 6. Sessions from distinct source columns to distinct destination rows can share a directed wavelength under column first routing.

Since the sessions belonging to a matching in the column-to-row bipartite graph are transmitted from distinct source columns to distinct destination rows, these sessions can be supported on one directed wavelength using column-first routing. The off-line algorithm in [19], [20] assigns a single matching to a single directed wavelength. Since the bipartite graph has the maximum node degree kR , the edges can be partitioned into kR disjoint matchings [22]. It follows that, kR directed wavelengths, i.e. $\lceil kR/2 \rceil$ wavelengths, are sufficient.

The complexity of the off-line RWA algorithm is due largely to finding kR disjoint bipartite matchings. The best known algorithm for finding a maximum bipartite matching in [23] has the running time $O(\Delta E)$, where Δ is the maximum node degree and E is the number of edges in the bipartite graph. For our purpose, this algorithm finds kR matching in time $O(kCR^2)$.

We now present our on-line algorithm. The algorithm keeps kR disjoint bipartite matchings of kR directed wavelengths such that each traffic session corresponds to an edge in one matching. When a session departs, we simply remove its corresponding lightpath from the network. When a new session, say $(\mathcal{C}_i, \mathcal{R}_j)$, arrives, we find one directed wavelength which is not used by any source in column i , and one directed wavelength which is not used by any destination in row j . If the two directed wavelengths are the same, we can support the new session without any lightpath rearrangement. Otherwise, we rearrange some existing lightpaths on the two directed wavelengths to support the new session. The

following lemma makes the above discussion concrete and states an upper bound on the number of lightpath rearrangements.

Lemma 7: In a bipartite graph $(\mathcal{V}_1, \mathcal{V}_2, \mathcal{E})$ with $|\mathcal{V}_1| = C \leq |\mathcal{V}_2|$, given a new edge $(\mathcal{C}_i, \mathcal{R}_j)$, $\mathcal{C}_i \in \mathcal{V}_1$, $\mathcal{R}_j \in \mathcal{V}_2$, a matching \mathcal{M}_1 of directed wavelength λ_1 which is not incident on \mathcal{C}_i , and a matching \mathcal{M}_2 of directed wavelength λ_2 which is not incident on \mathcal{R}_j , there exist two disjoint bipartite matchings which cover all the edges in \mathcal{M}_1 and \mathcal{M}_2 as well as the new edge $(\mathcal{C}_i, \mathcal{R}_j)$.

In addition, these two disjoint matchings can be assigned to λ_1 and λ_2 so that the number of lightpath rearrangements is at most $C - 1$.

*Proof:*³ Consider the bipartite graph $(\mathcal{V}_1, \mathcal{V}_2, \mathcal{E}')$ whose set of edges \mathcal{E}' contains all of the edges in \mathcal{M}_1 and \mathcal{M}_2 as well as the new edge $(\mathcal{C}_i, \mathcal{R}_j)$. Observe that each node has degree at most 2. It follows that \mathcal{E}' can be partitioned into two disjoint matchings [22], denoted by \mathcal{M}'_1 and \mathcal{M}'_2 respectively.

Without loss of generality, assume that $(\mathcal{C}_i, \mathcal{R}_j)$ belongs to \mathcal{M}'_1 . Let set \mathcal{P} contain the edges in \mathcal{M}_1 assigned to \mathcal{M}'_2 and the edges in \mathcal{M}_2 assigned to \mathcal{M}'_1 . Let set \mathcal{Q} contain the edges in \mathcal{M}_1 assigned to \mathcal{M}'_1 and the edges in \mathcal{M}_2 assigned to \mathcal{M}'_2 . Notice that \mathcal{P} and \mathcal{Q} cover all the edges in \mathcal{M}_1 and \mathcal{M}_2 . Since there are at most $2C - 2$ edges in \mathcal{M}_1 and \mathcal{M}_2 , it follows that $|\mathcal{P}| + |\mathcal{Q}| \leq 2C - 2$.

If $|\mathcal{P}| \leq C - 1$, assigning \mathcal{M}'_1 to λ_1 and \mathcal{M}'_2 to λ_2 yields the desired result that the number of lightpath rearrangements, which is equal to the sum of the number of edges in \mathcal{M}_1 assigned to \mathcal{M}'_2 and the number of edges in \mathcal{M}_2 assigned to \mathcal{M}'_1 , is at most $C - 1$. Otherwise, it is true that $|\mathcal{Q}| \leq C - 1$. In this case, assigning \mathcal{M}'_1 to λ_2 and \mathcal{M}'_2 to λ_1 yields the desired result. \square

Below is our on-line RWA algorithm which we shall refer to as the *torus RWA algorithm*.

Torus RWA Algorithm: (Use $\lceil kR/2 \rceil$ wavelengths in each fiber.)

Session termination: When a session terminates, simply remove its associated lightpath from the network without any further lightpath rearrangement.

Session arrival: When a session arrives and it is allowable, proceed as follows. Let i and j denote the source column and the destination row of the new session.

³As an alternative proof, we can map this lightpath rearrangement problem into a problem of connection rearrangement in a three-stage Clos network of switches [24]. The switches in the first and third stages correspond to the nodes in \mathcal{C} and \mathcal{R} respectively. The switch in the second stages correspond to wavelengths. It follows that Paull's theorem on the number of rearrangements can be applied to obtain the result of lemma 7. The proof above is simpler and more direct.

Step 1: If there is a directed wavelength, denoted by λ_0 , which is used by neither a source in column i nor a destination in row j , then assign the new session to λ_0 , and use column-first routing. In this case, no lightpath rearrangement is required. Otherwise, proceed to step 2.

Step 2: Find a directed wavelength, denoted by λ_1 , which is not used by any source in column i , and another directed wavelength, denoted by λ_2 , which is not used by any destination in row j . Since the new session is assumed allowable, λ_1 and λ_2 exist.

Modify the RWA of only the sessions on λ_1 and λ_2 . Construct the column-to-row bipartite graph $(\mathcal{V}_1, \mathcal{V}_2, \mathcal{E}')$ in which the set of edges \mathcal{E}' contains the bipartite matchings of λ_1 and λ_2 as well as the new edge $(\mathcal{C}_i, \mathcal{R}_j)$. Notice that $|\mathcal{V}_1| = C \leq R = |\mathcal{V}_2|$ and each abstract node has degree at most 2. From lemma 7, the set \mathcal{E}' can be partitioned into two disjoint matchings. In addition, lemma 7 tells us that the two matchings can be assigned to λ_1 and λ_2 such that at most $C - 1$ existing lightpaths need to be rearranged.

We now argue that each RWA update in the torus RWA algorithm requires solving a bipartite matching problem whose time complexity is only $O(R)$ using the following edge coloring procedure with two colors.

In the beginning of each coloring step, we select any uncolored edge and assign any color to it. We then iteratively proceed to its adjacent uncolored edge and assign the different color. Since we have a bipartite graph with maximum node degree 2, when the step terminates, we either have a path or a cycle of even length that never visits the same node more than once. In either case, it is clear that two colors are sufficient. At this point, we start the next coloring step until all the edges are colored. Since each node belongs to a single path or cycle, the edges colored in different steps are never adjacent and can thus be colored separately. It is clear that the edges with the same color form a matching. Since each edge is visited once and the number of edges is bounded by $2R$ in a bipartite graph with maximum node degree 2, the running time of this algorithm is $O(R)$.

For an $R \times C$ torus network with $R < C$, we can obtain similar results by reversing the roles of columns and rows. We summarize the results in this section in the following theorem.

Theorem 2: For an $R \times C$ torus network with symmetric k -allowable traffic, there exists, by construction, an on-line RWA algorithm which uses $\lceil k \max(R, C)/2 \rceil$ wavelengths in each fiber and requires at most $\min(R, C) - 1$ lightpath rearrangements per new session request.

V. CONCLUSION

We developed an on-line routing and wavelength assignment (RWA) algorithm for WDM bidirectional ring with N nodes to support k -allowable traffic in a rearrangeably nonblocking fashion. The algorithm uses $\lceil (\sum_{i=1}^N k_i)/3 \rceil$ wavelengths in each ring direction and requires at most three lightpath rearrangements per new session request regardless of the number of nodes N and the amount of traffic \mathbf{k} .

The developed algorithm implies the upper bound on $W_{\mathbf{k}}$, i.e. $W_{\mathbf{k}} \leq \lceil (\sum_{i=1}^N k_i)/3 \rceil$. The bound is tight for the case in which $N \geq 7$ and all the k_i 's are equal to some positive integer k . In addition, we observed that, for $N \geq 7$ and a fixed value of $\sum_{i=1}^N k_i$ equal to kN for some positive integer k , the case in which all the k_i 's are equal yields the maximum value of $W_{\mathbf{k}}$.

We extended the off-line RWA algorithm in [19], [20] to obtain an on-line algorithm for an $R \times C$ torus topology to support k -allowable traffic, where all the k_i 's are equal to some positive integer k , in a rearrangeably nonblocking fashion. Our on-line algorithm uses $\lceil k \max(R, C)/2 \rceil$ wavelengths in each fiber and requires at most $\min(R, C) - 1$ lightpath rearrangements per new session request regardless of the amount of traffic k .

Each of the above on-line RWA algorithms uses the same number of wavelengths as the best-known off-line algorithm, and is advantageous in two ways. First, each on-line algorithm guarantees that, for each RWA update due to a traffic change, only a small fraction of existing lightpaths are rearranged. Second, for each RWA update, applying our on-line algorithm instead of the off-line algorithm yields lower computational complexity.

We observe from the two algorithms that the number of lightpath rearrangements per new session request is related to the number of lightpaths supported on a single directed wavelength. For a bidirectional ring, up to two lightpaths are supported on a single directed wavelength. Since the ring RWA algorithm modifies only the RWA of the sessions on at most three directed wavelengths, it follows that the number of lightpath rearrangements depends on neither the number of nodes N nor the amount of traffic \mathbf{k} . For a torus topology, up to C lightpaths are supported on a single directed wavelength. Since the torus RWA algorithm modifies only the RWA of the sessions on at most two directed wavelengths, it follows that the number of lightpath rearrangements depends on the dimension of the network C but not on the amount of traffic \mathbf{k} . It is interesting to find out whether we can design an on-line RWA algorithm for an arbitrary mesh topology to have a similar property regarding lightpath rearrangements. We hope that our

analytical approaches in this paper can be used in future development of such an algorithm.

REFERENCES

- [1] D. Banerjee and B. Mukherjee, "A practical approach for routing and wavelength assignment in large wavelength-routed optical networks," *IEEE JSAC*, vol. 14, no. 5, pp. 903-908, June 1996.
- [2] R. Ramaswami and K.N. Sivarajan, "Design of logical topologies for wavelength-routed optical networks," *IEEE JSAC*, vol. 14, no. 5, pp. 840-851, June 1996.
- [3] I. Chlamtac, A. Ganz, and G. Karmi, "Lightpath communications: an approach to high bandwidth optical WAN's," *IEEE Trans. on Comm.*, vol. 40, no. 7, pp. 1171-1182, July 1992.
- [4] R. Ramaswami and K.N. Sivarajan, "Routing and wavelength assignment in all-optical networks," *IEEE/ACM Trans. on Networking*, vol. 3, no. 5, pp. 489-500, October 1995.
- [5] Z. Zhang and A.S. Acampora, "A heuristic wavelength assignment algorithm for multihop WDM networks with wavelength routing and wavelength reuse," *IEEE/ACM Trans. on Networking*, vol. 3, no. 3, pp. 281-288, June 1995.
- [6] C. Chen and S. Banerjee, "A new model for optimal routing and wavelength assignment in wavelength division multiplexed optical networks," *IEEE INFOCOM 1996*, pp. 164-171, April 1996.
- [7] S. Banerjee, J. Yoo, and C. Chen, "Design of wavelength-routed optical networks for packet switched traffic," *Jour. of Lightwave Tech.*, vol. 15, no. 9, pp. 1636-1646, September 1997.
- [8] K.C. Lee and V.O.K. Li, "A wavelength rerouting algorithm in wide-area all-optical networks," *Jour. of Lightwave Tech.*, vol. 14, no. 6, pp. 1218-1229, June 1996.
- [9] B. Mukherjee *et al.*, "Some principles for designing a wide-area WDM optical network," *IEEE/ACM Trans. on Networking*, vol. 4, no. 5, pp. 684-696, October 1996.
- [10] A. Birman, "Computing approximate blocking probabilities for a class of all-optical networks," *IEEE JSAC*, vol. 14, no. 5, pp. 852-857, June 1996.
- [11] S. Subramaniam, M. Azizoglu, and A.K. Somani, "All-optical networks with sparse wavelength conversion," *IEEE/ACM Trans. on Networking*, vol. 4, no. 4, pp. 544-557, August 1996.
- [12] R.A. Barry and P.A. Humblet, "Models of blocking probability in all-optical networks with and without wavelength changers," *IEEE JSAC*, vol. 14, no. 5, pp. 858-867, June 1996.
- [13] L. Li and A.K. Somani, "Dynamic wavelength routing using congestion and neighborhood information," *IEEE/ACM Trans. on Networking*, vol. 7, no. 5, pp. 779-786, October 1999.
- [14] Y. Zhu, G.N. Rouskas, and H.G. Perros, "A path decomposition algorithm for computing blocking probabilities in wavelength routing networks," *IEEE/ACM Trans. on Networking*, vol. 8, no. 6, pp. 747-762, December 2000.
- [15] O. Gerstel, G. Sasaki, S. Kutten, and R. Ramaswami, "Worst-case analysis of dynamic wavelength allocation in optical networks," *IEEE/ACM Trans. on Networking*, vol. 7, no. 6, pp. 833-845, December 1999.
- [16] R.K. Pankaj, *Architectures for Linear Lightwave Networks*, MIT Ph.D. Thesis, September 1992.
- [17] A. Narula-Tam, P.J. Lin, and E.H. Modiano, "Efficient routing and wavelength assignment for reconfigurable WDM networks," *IEEE JSAC*, vol. 20, no. 1, pp. 75-88, January 2002.
- [18] P. Saengudomlert, *Architectural Study of High-Speed Networks with Optical Bypassing*, MIT Ph.D. Thesis, September 2002.
- [19] A. Youssef, "Off-line permutation scheduling on circuit-switched fixed routing networks," *Symp. on Frontiers of Massively Parallel Computation*, pp. 389-396, October 1992.
- [20] C. Qiao and Y. Mei, "Off-line permutation embedding and scheduling in multiplexed optical networks with regular topologies," *IEEE/ACM Trans. on Networking*, vol. 7, no. 2, pp. 241-250, April 1999.
- [21] A. Ganz and Y. Gao, "Efficient algorithms for SS/TDMA scheduling," *IEEE Trans. on Communications*, vol. 40, no. 8, pp. 1367-1374, August 1992.
- [22] C. Berge, *Graphs*. North-Holland, 1985.
- [23] A. Schrijver, "Bipartite edge coloring in $O(\Delta m)$ time," *SIAM Journal on Computing*, vol. 28, no. 3, pp. 841-846, 1998.
- [24] J.Y. Hui, *Switching and Traffic Theory for Integrated Broadband Networks*. Kluwer Academic, 1990.