An On-Line Routing and Wavelength Assignment Algorithm for Dynamic Traffic in a WDM Bidirectional Ring

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

Abstract

We develop an on-line routing and wavelength assignment (RWA) algorithm for a WDM bidirectional ring with N nodes. The algorithm dynamically supports all $\mathbf{k}$-allowable traffic matrices, where $\mathbf{k}$ denotes an arbitrary integer vector $(k_1, ..., k_N)$, and node i, $1 \le i \le N$, can transmit at most k_i wavelengths and receive at most k_i wavelengths. Our algorithm uses $\lceil \sum_{i=1}^{N} k_i/3 \rceil$ wavelengths in each ring direction to support all $\mathbf{k}$-allowable traffic matrices in a rearrangeably nonblocking fashion. Furthermore, our algorithm requires at most three lightpath rearrangements per new session request regardless of the number of nodes N and the amount of traffic $\mathbf{k}$. In a special case with $k_i = k$ for all i, $1 \le i \le N$, the algorithm uses $\lceil kN/3 \rceil$ wavelengths in each ring direction, which is shown in [1] to be the minimum number of wavelengths required for any off-line nonblocking RWA algorithm.

1 Introduction

In a WDM network, the fiber bandwidth is divided into multiple frequency bands often called wavelengths. Without 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 the links from the source node to the destination node. Using wavelength converters allows a network to support a larger set of traffic. However, such converters are presently not commercially available and are likely to be very expensive. Hence, many researchers have focused on the problem of routing and wavelength assignment (RWA) assuming no wavelength conversion. We also focus on the same problem in this paper.

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. It is desirable to have an on-line RWA algorithm which requires few rearrangements of existing lightpaths, if any, in order to support each new session request. In this paper, we design one such algorithm for WDM bidirectional rings.

In section 2, we define a set of $\mathbf{k}$-allowable traffic, formulate the on-line RWA problem for $\mathbf{k}$-allowable traffic in a WDM bidirectional ring, and point out some known results. In section 3, we describe our on-line RWA algorithm in detail and prove its correctness. Finally, we summarize the results in section 4.

2 Problem Formulation and Related Works

We consider traffic sessions each of which takes up a full wavelength and therefore need not consider the problem of traffic grooming [2, 3]. We concentrate on a WDM bidirectional ring with N nodes. Adjacent nodes are connected by two fibers, one in each direction. Let $\mathbf{k}$ denote an arbitrary N-dimensional nonnegative integer vector $(k_1, ..., k_N)$. We assume that node i, $1 \le i \le N$, is equipped with k_i tunable transmitters and k_i tunable receivers. Therefore, at any time, each node i can transmit at most k_i wavelengths and can receive at most k_i wavelengths. Such a traffic matrix is said to belong to a set of $\mathbf{k}$-allowable traffic.

We model the traffic as a session-by-session arrival and departure process in which sessions arrive and depart one at a time. Since the rate of each session is fixed to a full wavelength, a session is fully described by its source-and-destination pair. In addition, since each session can be supported by either a clockwise lightpath or a counter-clockwise lightpath on some wavelength, a lightpath is described by its source/destination pair, its wavelength, and its ring direction.

A new session request is said to be allowable if there is a free transmitter at the source node and a free receiver at the destination node. In other words, a new session request is allowable if the resultant traffic matrix is still in the set of **k**-allowable traffic. An on-line RWA algorithm is said to be rearrangeably nonblocking if any allowable session request can be supported by a lightpath after possibly some rearrangements of existing lightpaths.

Our goal is to develop an on-line RWA algorithm for an N-node bidirectional ring to support **k**-allowable traffic in a rearrangeably nonblocking fashion. It is desirable to keep the number of wavelengths in each ring direction and the number of lightpath rearrangements per new session request to their minimum values.

Let $W_{\mathbf{k},N}$ denote the minimum number of wavelengths in each ring direction for an N-node bidirectional ring to support **k**-allowable traffic in a rearrangeably nonblocking fashion. In [1], it was shown that if the k_i's are all equal to k, then $W_{\mathbf{k},N} = \lceil kN/3 \rceil$ for $N \geq 5$. In addition, an off-line RWA algorithm that uses $\lceil kN/3 \rceil$ wavelengths in each ring direction was developed.

On-line RWA algorithms for dynamic traffic are provided in [4] for line, ring, and tree networks. However, the set of traffic considered in [4] is defined by the maximum link load. In addition, it is worth noting that [5] provides some bounds on the minimum number of wavelengths in each fiber for any physical network topology to support **k**-allowable traffic where the k_i's are all equal.

3 On-Line RWA Algorithm for k-Allowable Traffic

In this section, we present our on-line RWA algorithm. Define a directed wavelength as a wavelength in either the clockwise or the counterclockwise ring direction. Given w wavelengths in each ring direction, there are w directed wavelengths in the clockwise ring direction, and w directed wavelengths in the counterclockwise ring direction. Two sessions are said to be adjacent if the destination node of one session is the source node of the other session. The main idea behind our algorithm is to share a directed wavelength between two adjacent sessions, as suggested by the following known lemma [1].

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

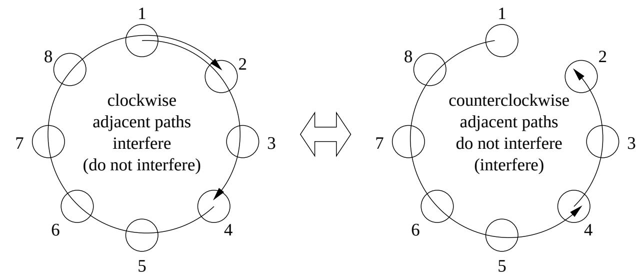

Figure 1: Lightpaths corresponding to a pair of adjacent sessions (1,4) and (4,2) interfere in the clockwise ring direction, but do not interfere in the counterclockwise ring direction.

The proof of lemma 1 is immediate from figure 1, where if two lightpaths overlap in one direction, they do not overlap in the other direction. In what follows, when lightpaths associated with a pair of adjacent sessions share a directed wavelength, we simply say that the adjacent session pair share a directed wavelength.

In our algorithm, we maintain the following two RWA conditions at all time: (i) only adjacent sessions share a directed wavelength, and (ii) at most two adjacent sessions share a directed wavelength. We now describe in detail the on-line RWA algorithm which uses $\lceil (\sum_{i=1}^{N} k_i)/3 \rceil$ wavelengths in each ring direction and at most three lightpath rearrangements per new session request.

Algorithm 1 Provide $\lceil (\sum_{i=1}^{N} k_i)/3 \rceil$ wavelengths in each ring direction.

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 nonsharing sessions and the new session, we claim 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 involving node 1 to sessions involving node N. Once an adjacent session pair is found, there are two possibilities for supporting the new session.

1. 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 session had to be rearranged to form the new adjacent session pair. Sharing of the directed wavelength by the adjacent session pair will free one directed wavelength on which the new session can be supported with only one lightpath rearrangement.
2. If the adjacent session pair cannot share the directed wavelength of either session in the pair, we claim 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 by 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. 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.

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. Before doing so, we prove one useful fact. 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. Let w be the number of wavelengths in use before the new session request. Note that $w = p + q$. For convenience, define $K = \sum_{i=1}^{N} k_i$. We now show that, in step 3, $p < \lfloor K/3 \rfloor$. Since the total number of sessions is at most K in **k**-allowable traffic, we have 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, we have that $K - p > w = 2\lceil K/3 \rceil$, yielding the desired relation

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

We now prove the first claim in step 3 that there always exists a new adjacent session pair. We proceed by contradiction. Suppose that no new adjacent session pair can be formed among nonsharing sessions and the new session. We argue that $q \leq \lfloor (K-p)/2 \rfloor$. To see this, observe 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 at another node, and one receiver at yet another node.

Let p_i be the number of adjacent sessions which share a directed wavelength at node i. Let $k'_i = k_i - p_i$ denote the number of transmitter/receiver pairs which are not used by those p_i adjacent sessions 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 $k_i^r \leq k'_i$.

In step 3, since no new adjacent session pair can be formed among the nonsharing sessions, we have 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

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

Since q is an integer, we have shown that $q \leq \lfloor (K-p)/2 \rfloor$.

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, we have that

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

It follows that

$$p \;\geq\; 2\lceil K/3 \rceil - \lfloor (K-p)/2 \rfloor \;\geq\; 2K/3 - (K-p)/2,$$

or quivalently, $p \geq K/3$, which contradicts the above established fact that $p < \lfloor K/3 \rfloor$ in step 3. Hence, we have shown that a new adjacent session pair always exists in step 3.

Finally, we 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 apparent from the fact that $p < \lfloor K/3 \rfloor$ 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 paths. It follows that, in either ring direction, a directed wavelength with a nonsharing session exists. □

The construction of our on-line RWA algorithm implies the following theorem.

Theorem 1 *For a bidirectional ring with N nodes and $\mathbf{k}$-allowable traffic, the required number of wavelengths in each ring direction for rearrangeably nonblocking $W_{\mathbf{k},N}$ is bounded by*

$$W_{\mathbf{k},N} \quad \leq \quad \left\lceil \frac{\sum_{i=1}^{N} k_i}{3} \right\rceil .$$

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 ring direction and requires at most three lightpath rearrangements per new session request.

When $N \geq 5$ and $k_i = k$ for all i, $1 \leq i \leq N$, we have from [1] that $W_{\mathbf{k},N} \geq \lceil kN/3 \rceil$. In this case, the above upper bound is tight. Otherwise, the above upper bound is not necessarily tight and algorithm 1 may use more than $W_{\mathbf{k},N}$ wavelengths. An interesting example 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 $k_i = 1$, $2 \leq i \leq N$. It can be shown that $W_{\mathbf{k},N} = \lceil (N-1)/2 \rceil$, which is less than the upper bound $\lceil 2(N-1)/3 \rceil$ from theorem 1. In addition, there exists an on-line RWA algorithm which uses $\lceil (N-1)/2 \rceil$ wavelengths and requires at most four lightpath rearrangements per new session request. Therefore, for a fixed value of $\sum_{i=1}^{N} k_i$ equal to kN for some positive integer k, the case in which the k_i's are all equal yields the maximum value of $W_{\mathbf{k},N}$.

4 Conclusion

We developed an on-line RWA algorithm for dynamic $\mathbf{k}$-allowable traffic in an N-node WDM bidirectional ring. The algorithm uses $\lceil (\sum_{i=1}^{N} k_i)/3 \rceil$ wavelengths in each ring direction, is rearrangeably nonblocking, and requires at most three lightpath rearrangements per new session request regardless of the number of nodes N and the amount of traffic $\mathbf{k}$. Our algorithm also provides an alternative derivation on the value of $W_{\mathbf{k},N}$ given in [1] when the k_i's are all equal.

The developed algorithm implies the upper bound on the minimum number of wavelengths for rearrangeable nonblocking: $W_{\mathbf{k},N} \leq \lceil (\sum_{i=1}^{N} k_i)/3 \rceil$. The bound is tight for the case in which the k_i's are all equal. In addition, we observed that, for $N \geq 5$ and a fixed value of $\sum_{i=1}^{N} k_i$ equal to kN for some positive integer k, the case in which the k_i's are all equal yields the maximum value of $W_{\mathbf{k},N}$.

References

[1] A. Narula-Tam, P.J. Lin, and E.H. Modiano, "Efficient Routing and Wavelength Assignment for Reconfigurable WDM Networks," *IEEE Journal of Selected Areas in Communications*, January 2001.

[2] O. Gerstel, R. Ramaswami, and G. Sasaki, "Cost-Effective Traffic Grooming in WDM Rings," *IEEE/ACM Transactions on Networking*, Vol.8, No.5, October 2000.

[3] A.L. Chiu and E.H. Modiano, "Traffic Grooming Algorithms for Reducing Electronic Multiplexing Costs in WDM Ring Networks," *Journal of Lightwave Technology*, Vol.18, No.1, January 2000.

[4] O. Gerstel, G. Sasaki, S. Kutten, and R. Ramaswami, "Worst-case Analysis of Dynamic Wavelength Allocation in Optical Networks," *IEEE/ACM Transactions on Networking*, Vol.7, No.6, December 1999.

[5] R.K. Pankaj, *Architectures for Linear Lightwave Networks*, MIT Ph.D. Thesis, September 1992.