



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Optical Switching and Networking 1 (2005) 35–49

Optical
Switching and
Networking

www.elsevier.com/locate/osn

Dynamic routing and wavelength assignment with optical bypass using ring embeddings[☆]

Li-Wei Chen, Eytan Modiano*

Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

Received 14 May 2004; received in revised form 15 October 2004; accepted 24 October 2004

Available online 8 December 2004

Abstract

We consider routing and wavelength assignment in ring, torus, and tree topologies with the twin objectives of minimizing wavelength usage and maximizing optical bypass. The P -port dynamic traffic assumption is used, which allows each node to send and receive at most P calls. For rings we show that $\lceil PN/4 \rceil$ wavelengths are necessary and sufficient, and provide a four-hub ring architecture that requires only half of these wavelengths to be locally processed. We extend this approach to develop RWA and bypass algorithms for both tori and trees by embedding virtual rings within these topologies and applying the ring algorithms. For an $R \times C$ torus, we embed $R + C$ rings onto the torus and provide an approach to RWA and banding based on solving disjoint RWA/banding problems for each ring. Our RWA algorithm is more wavelength efficient than any currently known algorithm and uses the minimum number of wavelengths for $R \geq 2C$. Our subsequent banding algorithm allows half of these wavelengths to bypass all but $4R$ hub nodes. Finally, we give a RWA for trees that embeds a single virtual ring and uses the ring to obtain a RWA that requires no more than $\lceil PN/2 \rceil$ total wavelengths; this figure is shown to be optimal for balanced binary trees. A banding algorithm follows that allows half these wavelengths to bypass all non-hub nodes.

© 2004 Elsevier B.V. All rights reserved.

Keywords: WDM; Optical networks; Bypass; Waveband switching; Routing and wavelength assignment

1. Introduction

Optical wavelength division multiplexing (WDM) technology has emerged as an attractive solution for meeting rapidly growing demands for network bandwidth. WDM allows the same fiber to carry many

signals independently as long as each uses a different wavelength. This allows many connections to share the same fiber resources. The *routing and wavelength assignment* (RWA) problem becomes one of assigning routes and wavelengths to each call such that no two calls share the same wavelength on the same link. Calls are additionally subject to the *wavelength continuity constraint*, which requires that a call use the same wavelength on all hops unless wavelength conversion is available at intermediate nodes. If full conversion is available at all nodes, the WDM network

[☆] This work was supported by the National Science Foundation (NSF) under Grant ANI-0073730.

* Corresponding author. Fax: +1 617 452 3414.

E-mail addresses: lwchen@mit.edu (L.-W. Chen), modiano@mit.edu (E. Modiano).

is equivalent to the well-known circuit-switched network; however, the high cost of wavelength converters often makes it desirable to keep the amount of conversion used in the network to a minimum.

There has been considerable work done in the area of finding efficient algorithms for the RWA problem. Here “efficient” implies the attempt to optimize certain figures of merit, which typically depend on the traffic assumptions made. Hence one way to categorize the existing works in the literature is according to their traffic assumptions.

Two broad categories for traffic models are *static traffic* and *dynamic traffic*. In the static traffic model, the set of calls that needs to be supported by the network is fixed and does not change over time. In the dynamic traffic model, the traffic set is allowed to change over time to represent call arrivals and departures.

In the static model, the objective is typically to minimize the number of wavelengths, converters, or other cost parameters [1]. This problem was shown to be NP-complete in [2], and thus the literature has focused on the development of heuristics and bounds. Other approaches include attempting to maximize throughput for fixed capacity [3], to minimize congestion for a fixed traffic set [4], or to maximize the number of calls supported for a fixed number of wavelengths [5]. However, this approach is limited in that it does not allow dynamic call set-up and removals.

In a dynamic model, calls are allowed to arrive and depart over time. One method of modelling call dynamics, known as the *blocking model*, is to adopt a statistical characterization of call arrival rates and holding times and design algorithms to minimize the call blocking probability. Numerous papers have focused on blocking probability analysis under various approximations for simple wavelength assignment algorithms such as the random algorithm [6–11] and first-fit [12]. However, due to the large state-space size of the problem, the blocking probability of a WDM network for more sophisticated algorithms is extremely difficult to analyze. As a result, most statistical algorithms rely on simplifying approximations and heuristics [13].

An alternative approach, the *non-blocking model*, considers designing the network to accommodate any traffic matrix from an admissible set. Call arrivals or departures are equivalent to transitions from one traffic matrix to another. If the transitions can be

accommodated without rearranging any calls, the RWA algorithm is called *wide-sense non-blocking*; algorithms which require call rearrangement are *rearrangeably non-blocking*. Different non-blocking models differ on the choice of admissible set. For example, [14] defines an admissible traffic set containing any traffic set where the worst-case link load is less than a given constant, and attempts to minimize the number of wavelengths required. In [15], the admissible set is similarly defined but the number of wavelengths is fixed and an attempt is made instead to minimize the wavelength converter usage.

We adopt in this paper the P -port model from [16] to define the admissible set. Under these assumptions, the admissible set consists of any traffic set where each node sends and receives at most P calls. The P -port model is very practical since the admissible set is based on actual device limitations in the network. In [16] it is shown that for the case of a bidirectional ring with N nodes and P ports, $\lceil PN/3 \rceil$ wavelengths are necessary and sufficient to support P -port traffic in a rearrangeably non-blocking fashion if no wavelength conversion is employed. An online version which attempts to minimize the number of rearrangements per new call arrival was presented in [17]; this algorithm was later extended from rings to torus networks in [18]. In [19] ring networks with wavelength conversion are considered and it is shown that $\lceil PN/4 \rceil$ wavelengths are both necessary and sufficient under these conditions.

In this paper, we investigate new rearrangeably non-blocking RWA schemes for P -port traffic on a variety of topologies with and without wavelength conversion. The criteria to be optimized will be both the number of wavelengths required and the amount of optical bypass in the network. The details of the traffic model are first discussed; this is followed by a description of the optimization criteria.

1.1. System model

We consider three topologies in this paper: the ring, the two-dimensional torus, and the tree.

In a ring topology, adjacent nodes are connected by two fibers: one supporting wavelengths travelling in the clockwise direction, the other supporting wavelengths in the counterclockwise direction. The two fibers are represented by a single bidirectional

link, where each link can support calls travelling in both directions on every wavelength. If we number the nodes from 0 to $N - 1$, then each node i is connected to nodes $i \oplus 1$ and $i \ominus 1$, where \oplus and \ominus denote addition and subtraction modulo N , respectively.

A two-dimensional (2D) torus is essentially a rectangular topology where each node is a member of two rings. The total number of nodes N must be the product of two integers R and C , and the resultant topology is known as an $R \times C$ torus. Examples of torus topologies are given in Section 3.

A tree topology consists of nodes joined by bidirectional fiber links such that the undirected graph underlying the network must be connected and contain no cycles [20].

We will consider networks with and without wavelength converters. A wavelength converter, if available at a given node, can be used to switch a call arriving to that node on one wavelength onto a different wavelength departing from the node. If no conversion is employed, a call passing through a node on one wavelength must exit from the node on the same wavelength.

A traffic matrix or traffic set consists of a set of calls that need to be set up in the network. Each call consists of a source and destination pair and is assumed to occupy a full wavelength. A traffic set is *connected* if the directed graph corresponding to the set of source–destination pairs is connected. In the P -port model, the number and type of calls which the traffic set may contain are based on the port limitations at each node. Each node i is considered to have an equal number of ports P_i for transmitting and receiving, and hence may be the source and destination of at most P_i calls within any given traffic set. This traffic model is natural because it places constraints on the calls which are based on the actual equipment constraints at each node. A traffic set which obeys these constraints is said to be *admissible*.

The dynamic nature of the traffic is modelled by transitions from one admissible traffic set to another over time. The RWA algorithm must provision enough resources to support *any* admissible traffic set, so that regardless of the evolution of the network state over time, no calls are blocked. For this reason, the model is called *non-blocking*. We further allow existing calls to be rearranged between transitions; hence the term *rearrangeably non-blocking*.

In general, each node i can have a different number of ports P_i . The special case where all nodes have the same number of ports $P_i = P$ is called *symmetric P -port traffic*. A special case of symmetric P -port traffic is the *single-port traffic* case, where $P = 1$ for every node.

This paper considers algorithms for assigning routes and wavelengths to every traffic matrix from the admissible set.

1.2. Objective function

In this paper, we consider two approaches to designing RWA algorithms: constructing algorithms that minimize the number of wavelengths in the network, and algorithms that maximize the amount of bypass at each node using banding.

The approach of minimizing the total number of wavelengths is common in the literature for non-blocking networks. Minimizing wavelengths is sensible for two reasons. First, adding additional wavelengths to a network is costly, so wavelengths should be used efficiently. Second, if each wavelength needs to be switched at each node, then optimizing the number of wavelengths is also equivalent to optimizing the number of switch ports at each node, an important property since switching is expensive and switching costs can rapidly dominate in the network as switch size increases.

The second approach is to use *banding* to reduce the number of wavelengths dropped at each node. Banding refers to the grouping of wavelengths into frequency bands; each band contains multiple adjacent wavelengths. Node complexity can be significantly reduced by allowing some bands to completely bypass each node. This is permissible if the RWA algorithm can guarantee that wavelengths within the bypassing bands are never dropped at those nodes.

For the purposes of the RWA problem, we can group the wavelengths into two bands: a *local* band, consisting of wavelengths that can be accessed by all nodes, and a *bypass* band, consisting of wavelengths that can be accessed only by a few designated *hub nodes*. The bypass band can therefore bypass the majority of the nodes in the network. The bypass band can be viewed as a set of “highways” that can only be accessed via the hub nodes, which serve as access ramps by using wavelength conversion to translate




Fig. 1. An OADM architecture. Note that a $(W + P) \times (W + P)$ switch is required, and that the size of this switch increases with the number of wavelengths.

incoming calls on a local wavelength to a bypass wavelength and vice versa, as necessary.

There are several advantages to a banding approach. One is cost savings. Fig. 1 gives an example of an optical add-drop multiplexer (OADM) in a network with no banding. In a system with W total wavelengths, all W wavelengths would have to be switched. Fig. 2 shows a system where the wavelengths are divided into a local band of k wavelengths, and a bypass band of $W - k$ wavelengths. In this case, only the smaller local band of k wavelengths is switched. Another benefit is that the wavelength demultiplexers can be simpler: the first, coarse DMUX need only separate out two large bands, while the second, finer DMUX has a smaller band to work with (only the local wavelengths). Finally, by allowing wavelengths in the bypass band to avoid processing at non-hub nodes altogether, the bypass band can either avoid the switch (thus not suffering power losses due to switching which would reduce the reach of the lightpaths), or be placed in a separate fiber entirely. Such a separate fiber would need to be connected only to the hub nodes and could physically bypass all other nodes entirely.

The potential disadvantage, of course, is that by restricting the access of some of the wavelengths to just the hub nodes, we are imposing additional constraints onto the RWA problem that may decrease the efficiency of the algorithms and necessitate the provisioning of more overall wavelengths.

In this paper, we provide novel RWA algorithms with and without bypass for a variety of topologies. We demonstrate for each case optimal or efficient wavelength usage through comparisons with




Fig. 2. An OADM architecture with bypass. Note that the switch size depends only on the number of wavelengths k entering the node, not the total number W .

theoretical lower bounds, and demonstrate ways for using banding to further simplify node complexity. In Section 2 we cover RWA and banding algorithms for ring networks. Sections 3 and 4 describe a method for using ring embeddings to extend these results to torus and tree topologies, respectively. Table 1 provides a summary of the wavelength usage of the algorithms to be presented. Typically, in each section, a RWA algorithm which achieves the lower bound on the total number of wavelengths used is presented, followed by a banding algorithm which reduces the number of local wavelengths at the expense of the overall total number of wavelengths or wavelength converters.

Table 1
Algorithm performance summary

Type	Algorithm	λ 's (local)	λ 's (total)
Ring	Lower bound	$\lceil PN/8 \rceil$	$\lceil PN/4 \rceil$
	$\lceil PN/4 \rceil$ alg.	$\lceil PN/4 \rceil$	$\lceil PN/4 \rceil$
	Ring banding	$\lceil PN/8 \rceil$	$\lceil 3PN/8 \rceil$
Torus	Lower bound ($R \geq C$)	–	$\lceil PR/4 \rceil$
	TERA ($R \geq 2C$)	$\lceil PR/4 \rceil$	$\lceil PR/4 \rceil$
	Banding ($R \geq 2C$)	$\lceil PR/8 \rceil$	$\lceil 3PR/4 \rceil$
Tree	Lower bound (binary)	–	$\lceil PN/2 \rceil$
	Banding	$\lceil PN/4 \rceil$	$\lceil PN/2 \rceil$

2. Ring networks

2.1. Rings without conversion

In this section, RWA algorithms which minimize the total number of wavelengths required to support

P -port traffic in an N -node ring are considered. We first establish a bound on the minimum number of wavelengths required without conversion. Recall that under the P -port traffic model, any traffic set in which each node i sends and receives no more than P_i calls is admissible; as such, no call belonging to such a set may be blocked. Therefore sufficiently many wavelengths must be provisioned to support all calls within any given admissible traffic set. In [16] it was shown that for a ring with $N > 7$ nodes, there exists an admissible traffic set which requires at least $\lceil PN/3 \rceil$ wavelengths to support it if no wavelength conversion is available. A RWA algorithm was also described that always uses no more than this minimum number of wavelengths for any admissible traffic set. The algorithm was based on the following observation:

Lemma 1. *Define a set of calls to be adjacent if the destination of the first call is the source of the second call, the destination of the second call is the source of the third call, and so forth. Given any two adjacent calls, if the calls cannot fit on a single bidirectional wavelength in the clockwise direction, they must fit onto a single wavelength in the counterclockwise direction.*

A graphical proof is shown in Fig. 3. The algorithm in [16] divides the calls into sets of three adjacent calls each, and observes that as a consequence of Lemma 1 each set could fit on a single wavelength. To do this, determine the direction in which the first two adjacent calls can share a wavelength, and route them in that direction using a wavelength. Route the third call in the opposite direction using the same wavelength.

2.2. Rings with conversion

We next consider P -port bidirectional rings with conversion, again under the objective of minimizing the number of wavelengths. The cut-set bound can be used to provide a lower bound for symmetric P -port traffic. Consider a cut of two links which divides the ring into equal halves of $N/2$ nodes each. A worst-case admissible P -port traffic set can be constructed where each node on the left half of the cut sends all P units of traffic to some node on the right half. This means that $PN/2$ units of traffic would traverse two links, requiring a minimum of $\lceil PN/4 \rceil$ wavelengths.




Fig. 3. (a) Two calls, (n_1, n_2) and (n_2, n_3) , which cannot fit on a single counterclockwise wavelength since the second call overlaps from n_1 to n_3 with the first call. (b) The same two calls can use a single wavelength in the clockwise direction, because the second call must reach destination n_3 in the clockwise direction before encountering the source n_1 of the first call.

An algorithm was provided in [19] called the $\lceil PN/4 \rceil$ algorithm which provides a RWA for any P -port traffic set using no more than $\lceil PN/4 \rceil$ wavelengths. This is optimal in the sense that it meets the lower bound on wavelength usage. It requires a total of $\lceil PN/2 \rceil$ wavelength converters, which can be arbitrarily located within the ring. A detailed description of the algorithm is given in [19]; for the purposes of our discussion here it suffices to know that such an algorithm exists. This algorithm will be used later in the paper to assist in performing RWA on other topologies.

A corollary is that for the special case of $N = 4$ and $P = 1$, only a single wavelength is required, and no wavelength converters are used. This can be extended to the case of arbitrary P by noting that any P -port set can be decomposed into P single-port sets, and each set routed individually. Hence for $N = 4$ the lower bound of P wavelengths can be achieved without conversion.

2.3. The ring banding bound

In this section, we consider maximizing the amount of bypass using banding. Recall that wavelengths are divided into two sets of adjacent wavelengths known as the local band and the bypass band. Wavelengths in the local band, known as local wavelengths, can be accessed by any node; wavelengths in the bypass band, known as bypass wavelengths, can be accessed only by a few special nodes known as hub nodes. Any non-hub node is termed a local node. As described




Fig. 4. A dual ring topology, equivalent to a twenty-node, four-hub ring with some local and some bypass wavelengths. The shaded nodes are hubs.

in Section 1.2, maximizing the size of the bypass band corresponds to reducing the overall cost of the network.

In this section, we assume symmetric P -port traffic and derive a lower bound on the size of the local band for a fixed number of hubs h within an N -node ring. In general, the number of local wavelengths required decreases as the number of hubs increases. It is possible to lower-bound the number of local wavelengths required for a fixed number of hubs using a cut-set bound.

The hub nodes divide the ring into sections, each consisting of a number of local nodes located between two consecutive hub nodes. Suppose the smallest such section consists of N_S local nodes. A traffic set exists where each node within that section sends all P calls to nodes outside that section. Imagine a cut consisting of the two links at the edge of that section. There are $P \cdot N_S$ units of traffic travelling across two links, and so a minimum of $\lceil P \cdot N_S / 2 \rceil$ wavelengths are required. Furthermore, these must all be local wavelengths, since none of the local nodes can access bypass wavelengths. To obtain the tightest such bound, we maximize N_S by distributing the hubs symmetrically within the ring, resulting in $N_S = \lceil N/h \rceil - 1$ and a lower bound of $\lceil P(\lceil N/h \rceil - 1)/2 \rceil$. This situation is illustrated in Fig. 4 for the case of $h = 4$.

2.4. Banding and bypass on rings

In this section, an algorithm called the *ring banding algorithm* that minimizes the number of wavelengths in the local band is described. In general, there exist an entire class of ring banding algorithms depending on the number of hubs in the ring; however, from a

single instance of the algorithm it is straightforward to extrapolate to any other variation since the algorithms are all very similar. Therefore in this paper we focus primarily on the special case of a four-hub ring banding algorithm. A four-hub architecture has the added advantage that routing traffic between the four hubs does not require wavelength conversion to achieve optimal wavelength efficiency, as noted in the preceding section.

For $h = 4$ hubs, the lower bound is $\lceil P(\lceil N/h \rceil - 1)/2 \rceil = \lceil P(\lceil N/4 \rceil - 1)/2 \rceil$ wavelengths. To obtain some intuition about how large this local band is, consider the case where N is a multiple of 4; under this assumption, this minimum number becomes $\lceil P(N - 4)/8 \rceil$. This establishes that at least half of all the wavelengths must be local wavelengths. (Recall that $\lceil PN/4 \rceil$ total wavelengths are required.)

To reduce the number of local wavelengths required, consider a topology where the four hubs are distributed symmetrically within the ring. The ring banding algorithm gives each call a route and wavelength assignment using a three-step process, as follows:

- (1) Starting from the source node, the call travels to the nearest hub. This route uses a local wavelength and is static since the nearest hub node for any given source node is fixed.
- (2) From that hub, it travels via a bypass wavelength to the hub closest to the destination node. This routing is dynamic since the source and destination of the call are variable.
- (3) Finally, the call proceeds from that hub to the destination node via a local wavelength, again by a static route.

We first prove that Steps (1) and (3) use no more than the minimum number of local wavelengths. Then we provide a RWA to dynamically route all the calls in Step (2) using as few bypass wavelengths as possible.

Consider the local wavelength usage. The hubs partition local nodes in the ring into four quarters, each of which contains no more than $\lceil N/4 \rceil - 1$ local nodes. Assign each node within the quarter P bidirectional wavelengths for communication with the hub closest to it. The forward direction, from the node to the hub, is used in Step (1); the reverse direction, from the hub to the node, is used in Step (3). Since half the nodes in each quarter communicate with the hub node on




Fig. 5. Assigning local wavelengths for a single-port sixteen-node ring.

one side and half with the other, each bidirectional wavelength can be shared by two local nodes. This is illustrated for a 16-node ring in Fig. 5. In total, $\lceil P(\lceil N/4 \rceil - 1)/2 \rceil$ local wavelengths are required by this scheme. Note that this meets the lower bound on local wavelengths, and corresponds to roughly half of the total number of wavelengths used by the $\lceil PN/4 \rceil$ algorithm.

Next we consider the dynamic routing of calls between hubs in Step (2). Each hub is responsible for sending and receiving all calls belonging to the $\lceil N/4 \rceil$ nodes closest to it (including itself) to and from other hubs. Therefore, for Step (2) each of the four hubs acts as though it has $P' = P\lceil N/4 \rceil$ ports. From the results earlier in this section, we know that $P\lceil N/4 \rceil$ bypass wavelengths are necessary and sufficient to support such traffic, and furthermore that we can do so without the use of wavelength converters. (To see this, note that in the four-hub ring $N = 4$ and $P' = P\lceil N/4 \rceil$.)

A summary of the properties of the ring banding algorithm follows:

Local node requirements: Each local node can use fixed routes and wavelengths to communicate with its nearest hub node. This allows the local node architecture to be extremely simple. In fact, using this algorithm, not only do local nodes never need to access bypass wavelengths, they do not need to switch local wavelengths either since these wavelengths are statically assigned. Therefore, the

switching block shown in Fig. 2 can be replaced with static connections from the P ports to the P assigned wavelengths.

Hub node requirements: Each of the four hub nodes needs to be able to switch between any of the local wavelengths and any of the bypass wavelengths. Bypass wavelengths never need to be switched onto other bypass wavelengths—i.e. no conversion is required between bypass wavelengths. $\lceil PN/2 \rceil$ converters between bypass and local wavelengths are required at each hub node.

Call rearrangements: Call departures never require existing calls to be rearranged since the wavelengths used can just be removed. It can be shown that call arrivals may require the rearrangement of at most seven bypass wavelengths [17] (two sets of single-port traffic for the four hub nodes, less one for the newly arriving call). Local wavelengths never need rearrangement, as they are statically assigned.

We illustrate the operation of the four-hub ring banding algorithm with an example.

Example 1. Fig. 5 considers single-port traffic on a 16-node ring, and shows the wavelengths that would be assigned in Steps (1) and (3) of the algorithm to handle the traffic between the local nodes and the hubs. Note that only the minimum number $\lceil (16/4 - 1)/2 \rceil = 2$ of local wavelengths are required. For any fixed traffic set, the four hubs would then provide a RWA using the bypass wavelengths via the $\lceil PN/4 \rceil$ algorithm.

To illustrate the full RWA for a given call, consider a call travelling from node 6 to node 12. The closest hub for node 6 is 5, and the closest hub for node 12 is 13. Therefore a route is assigned as follows: (1) from node 6 to node 5 via a local wavelength, (2) from node 5 to node 13 via a bypass wavelength, and (3) from node 13 to node 12 via a local wavelength again. Fig. 6 shows an example of a possible RWA for the call using this approach.

3. Torus networks

3.1. The torus lower bound

Torus networks with no wavelength conversion were considered in [18], which presented a RWA algorithm based on column-first routing requiring




Fig. 6. The route for call (6, 12). Hubs 5 and 13 are used to access the bypass wavelengths.

twice the minimum number of wavelengths. We investigate the use of wavelength conversion to reduce the number of wavelengths required. We will first show a lower bound on wavelengths required; subsequently, we give a novel algorithm for RWA on a torus that makes efficient use of wavelengths and achieves the bound in certain cases.

Define an $R \times C$ torus to be a network consisting of RC nodes, each of which is uniquely assigned to a row r , $1 \leq r \leq R$, and a column c , $1 \leq c \leq C$. Each node is connected via four bidirectional links to four other nodes: the two adjacent nodes occupying the same row, and the two adjacent nodes in the same column. More specifically, denote the node in row r and column c by $n_{r,c}$. Then the node $n_{r,c}$ is connected via bidirectional links to the nodes $n_{r \oplus 1, c}$, $n_{r \ominus 1, c}$, $n_{r, c \oplus 1}$, and $n_{r, c \ominus 1}$, where operations on the rows and columns are modulo R and C , respectively. Fig. 7 shows an example of a 6×3 torus. Note that each row contains C nodes, and each column contains R nodes.

We again use the P -port traffic model, where each node is assumed to have P ports, and hence can send and receive at most P calls. We next investigate how many wavelengths must be provisioned in the torus to support any admissible traffic set under these assumptions.

Suppose without loss of generality that $R \geq C$. (If the opposite is true, rotate the picture of the torus 90 degrees and relabel the rows as the columns and vice versa.) Let R be even. Consider a horizontal cut across the columns which removes $2C$ links and

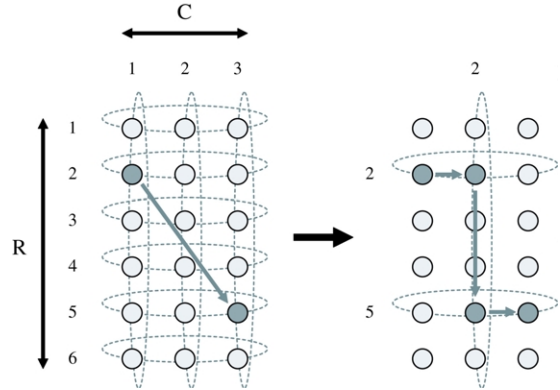


Fig. 7. Breaking up a call into three sub-calls using a bridging column. The single call on the left, from $n_{2,1}$ to $n_{5,3}$, is routed as two row-ring calls and a column-ring call using the bridging column 2. Each of the sub-calls can be routed independently on their respective rings using the $\lceil PN/4 \rceil$ algorithm. Additional wavelength conversion may be required at nodes $n_{2,2}$ and $n_{5,2}$ if the sub-calls are not assigned to the same wavelength.

divides the torus into two equal sets of $RC/2$ nodes each. Consider the first set. Under the P -port model, there exists a worst-case traffic set in which each node in that set sends all P calls to some node in the other set. In this case, there are $PRC/2$ calls traversing $2C$ links, which means that a minimum of $\lceil (PRC/2)/(2C) \rceil = \lceil PR/4 \rceil$ wavelengths are required. A similar argument for R odd yields a bound of $\lceil P(R - 1)/4 \rceil$ wavelengths.

3.2. The TERA algorithm—overview

In this section we describe an algorithm based on the Torus Embedded-Ring Approach (TERA) for routing and wavelength assignment. We will show that by judicious use of wavelength conversion, the TERA algorithm will use no more than $\max\{\lceil PC/2 \rceil, \lceil PR/4 \rceil\}$ wavelengths. For tori where $R \geq 2C$, this achieves the lower bound of $\lceil PR/4 \rceil$ wavelengths; in the worst case ($R = C$) it uses $\lceil PR/2 \rceil$.

We will describe the algorithm in detail later in this section, but the general idea is as follows. For any given call going from n_{r_1, c_1} to n_{r_2, c_2} , instead of considering all possible route assignments (of which there are many), we break the problem down into finding a route from the source n_{r_1, c_1} to some intermediate node n_{r_1, c_b} in the same row, from n_{r_1, c_b}

to some n_{r_2, c_b} in the same column, and finally from n_{r_2, c_b} to the destination n_{r_2, c_2} in the same row. The advantage to this approach is that instead of having a single call travelling through a torus, the call has been subdivided into three smaller calls, each on a different ring, and the results for rings in the preceding section can be used to do the routing and wavelength assignment for each sub-call. Fig. 7 gives an example of routing a call using this approach.

In the above approach, an intelligent choice of a column c_b for each call is required so that subsequent routing of the resulting sub-calls uses as few wavelengths as possible.

For notational purposes, define the ring formed by the nodes $\{n_{i,j} \mid j = 1, \dots, C\}$ to be the row-ring i , and the ring formed by the nodes $\{n_{i,j} \mid i = 1, \dots, R\}$ to be the column-ring j . Under this nomenclature, there are R row-rings and C column-rings. Let us call the columns $\{c_b\}$ used to generate the sub-calls the *bridging columns*, since calls will use these columns to travel between row-rings. In an $R \times C$ torus, there are C bridging columns. We will use a bipartite matching approach to associate each call with a bridging column in such a way that the resulting sub-calls will form a $2P$ -port traffic set on each row-ring, and a P -port traffic set on each column-ring. Once this is done, it is evident that by using the $\lceil PN/4 \rceil$ algorithm on each row-ring and column-ring, the total number of wavelengths required in the torus will be $\max\{\lceil 2PC/4 \rceil, \lceil PR/4 \rceil\} = \max\{\lceil PC/2 \rceil, \lceil PR/4 \rceil\}$, as claimed.

3.3. Bridging column assignment

In this section, we will describe a method for using matchings to assign a bridging column to each call such that the resulting sub-calls form a $2P$ -port traffic set on each row-ring, and a P -port traffic set on each column-ring.

Consider a call $(n_{r_1, c_1}, n_{r_2, c_2})$ that is divided into three sub-calls $(n_{r_1, c_1}, n_{r_1, c_b})$, $(n_{r_1, c_b}, n_{r_2, c_b})$, and $(n_{r_2, c_b}, n_{r_2, c_2})$. We will call the $(n_{r_1, c_1}, n_{r_1, c_b})$ the *starting sub-call*, $(n_{r_1, c_b}, n_{r_2, c_b})$ the *bridging sub-call*, and $(n_{r_2, c_b}, n_{r_2, c_2})$ the *ending sub-call*. We wish to determine the conditions that the bridging column assignment is subject to.

1. Row-ring conditions: For each row-ring of size C , there are $2PC$ sub-calls to be routed on it, P of

which are starting sub-calls and the remaining P of which are ending sub-calls. Each node already uses P ports for sending the starting sub-calls, and P ports for receiving the ending sub-calls. We wish to choose bridging columns such that the starting sub-calls on each row-ring use no more than P additional destination ports per node, and the ending sub-calls use no more than P additional source ports per node. If these conditions are met, then each row-ring needs no more than $2P$ ports per node.

2. Column-ring conditions: For each column-ring, we wish to choose bridging columns such that the P -port assumption holds true for each column-ring; i.e. the set of PR bridging calls for each column-ring should use no more than P source and destination ports per node.

Lemma 2. *If the bridging conditions imposed by the column-rings are satisfied, then the bridging conditions for the row-rings are also satisfied.*

Proof. For each row-ring, at least P ports are required regardless of the choice of column-rings since the sources of each starting sub-call and the destinations of each ending sub-call are fixed. Thus it remains to show only that each node is the destination of no more than P starting sub-calls, and the source of no more than P ending sub-calls, to prove that only another P ports per node are required (for a total of $2P$ per node).

To see that this is true, note that the destination of each starting sub-call is the source of a corresponding bridging sub-call. Therefore, if a node in a row-ring were the destination of more than P starting sub-calls, that node would also be the source of more than P bridging sub-calls on its column-ring. Since we assumed that the column-ring conditions were satisfied, this cannot be true.

Similarly, the source of each ending sub-call is the destination of a corresponding bridging sub-call. If a node in a row-ring were the source of more than P ending sub-calls, that node would also be the destination of more than P bridging sub-calls on its column-ring, which cannot be true. \square

Lemma 2 tells us that it suffices to assign calls to bridging columns so that the column-ring conditions are satisfied. We achieve this by using a bipartite matching approach, as follows.

Consider the traffic set for the $R \times C$ torus. We construct a bipartite graph consisting of two sets of R vertices, which we will call $\{S_i\}$ and $\{D_j\}$, where $i = 1, \dots, R$ and $j = 1, \dots, R$. In a given set, each vertex corresponds uniquely to one of the R row-rings.

Each call in the traffic set will correspond to an edge in the graph. A call from a node n_{r_1, c_1} to n_{r_2, c_2} will be represented by an edge from vertex S_{r_1} to D_{r_2} ; in other words, a call in the torus is represented by an edge in the graph connecting the vertex in $\{S_i\}$ representing its source row-ring, and the vertex in $\{D_j\}$ representing its destination row-ring. Since we have a P -port traffic set, and each row-ring contains C nodes, each vertex in the graph will have degree PC .

We will call the bipartite graph thus constructed the *bridging graph*. In the next step, we will make use of the following theorem for bipartite matchings with equal nodal degree.

Theorem 1. *Define a perfect matching to be a set of edges where exactly one edge is incident on every vertex. Then, in a bipartite graph (V_1, V_2, \mathcal{E}) in which each vertex in V_1 and in V_2 has degree m , the set \mathcal{E} can be partitioned into m disjoint perfect matchings.*

Proof. The proof is basically by induction using Hall's theorem [21] and was given in [22]. It is omitted here for brevity. \square

In the context of our constructed bipartite matching, Theorem 1 guarantees that we can obtain a set of PC disjoint perfect matchings. Each perfect matching corresponds to a set of R calls where exactly one call originates from each row-ring, and one call is destined for each row-ring. Since for any given call the bridging sub-call has source node equal to the source row of the original call, and destination node equal to the destination row of the original call, this means that if all calls in a matching use the same bridging column, the set of resultant bridging sub-calls will correspond to a single-port traffic matrix for that column.

The preceding idea forms the basis for the assignment of bridging columns. Recall that Theorem 1 guarantees that we will have PC disjoint perfect matchings. Divide these matchings into C sets of P disjoint perfect matchings. Assign each set of matchings to one of the C columns. All calls in a matching

source		destination		graph edge	
row	column	row	column	S_i	D_j
1	1	3	2	1	3
1	2	2	1	1	2
2	1	1	1	2	1
2	2	4	1	2	4
3	1	2	2	3	2
3	2	1	2	3	1
4	1	4	2	4	4
4	2	3	1	4	3

Fig. 8. A traffic set for the single-port 4×2 torus considered in Example 2. The first two pairs of columns give the row-column pairs for the source and destination nodes, while the last two columns give the edges that represent each respective call in the bridging graph.

assigned to a given column will use that column as its bridging column. Since each matching requires only a single port per node, the P matchings in each column will require no more than P ports per node. Thus the column-ring conditions (and subsequently the row-ring conditions, by Lemma 2) are satisfied.

Example 2. In this example, we consider the problem of assigning bridging columns to a traffic set on a single-port 4×2 torus. The traffic set is given in Fig. 8. The corresponding bridging graph is shown in Fig. 9. Recall that Theorem 1 states that since the graph has vertex degree $PC = 1 \cdot 2 = 2$, we can find two disjoint matchings. One possible such choice is given.

Under the choice of matchings given in Fig. 9, we assign all calls in the first matching to column 1, and all calls in the second matching to column 2. This uniquely specifies a bridging column for each call in the traffic set, as shown in Fig. 10. For example, the first call in the set, from $n_{1,1}$ to $n_{3,2}$, corresponds to the graph edge S_1 to D_3 . This edge is in the first matching, so the call from $n_{1,1}$ to $n_{3,2}$ is assigned to the bridging column of 1. The resulting sub-calls are $(n_{1,1}, n_{1,1})$, $(n_{1,1}, n_{3,1})$, and $(n_{3,1}, n_{3,2})$. (Note that the $(n_{1,1}, n_{1,1})$ sub-call happened to be degenerate, and in practice would not require a wavelength.)

3.4. The TERA algorithm—operation

We can now formally state the TERA algorithm.




Fig. 9. The bridging graph for Example 2. As expected, each vertex in the bridging graph has vertex degree $PC = 1 \cdot 2 = 2$. Using Theorem 1, this can be divided into $C = 2$ disjoint perfect matchings.

call		bridging column
src row	src col	
1	1	1
1	2	2
2	1	1
2	2	2
3	1	1
3	2	2
4	1	1
4	2	2

Fig. 10. The resultant assignment of bridging columns to calls for Example 2.

The TERA Algorithm.

- (1) Given a P -port traffic set for the torus, construct the corresponding bipartite bridging graph, as described in Section 3.3.
- (2) Divide the edges on the bridging graph into C sets of P disjoint bipartite matchings. Assign each set to a different bridging column.
- (3) Now that each call has a bridging column, divide each call into a starting sub-call, a bridging sub-call, and an ending sub-call.
- (4) For each row-ring, use the $\lceil PN/4 \rceil$ algorithm to perform RWA on all starting and ending sub-calls within that row-ring. This requires $\lceil PC/2 \rceil$ wavelengths on each row-ring.
- (5) For each column-ring, use the $\lceil PN/4 \rceil$ algorithm to perform RWA of all bridging sub-calls within that column-ring. This requires $\lceil PR/4 \rceil$ wavelengths on each column-ring.

- (6) Give each original call in the torus the route and wavelength assignment formed by the concatenation of the RWA of the starting, bridging, and ending sub-calls. Up to two converters may be needed to change between sub-calls.

3.5. Banding and bypass on tori

In this section, we give an algorithm for banding on a torus which reduces size of the band of local wavelengths dropped at non-hub nodes.

Since the TERA algorithm essentially results in a problem of disjointly routing traffic on different rings in the rows and columns, we can apply an approach similar to the ring banding algorithm. Rather than using the $\lceil PN/4 \rceil$ algorithm to route the sub-calls on the row-rings and column-rings, the ring banding algorithm is used. In this discussion, we focus on using the four-hub ring banding algorithm, but the results could be extended to using different numbers of hubs.

We assume again by convention that $R \geq C$. Using a four-hub architecture on each row, we can reduce the number of local wavelengths required on those rings to $P\lceil C/8 \rceil$. A minimum of $4R$ hubs are required to do this, since no two row-rings can share a hub node. In order for rings along the columns to also require a local band of no more than $P\lceil C/8 \rceil$ wavelengths, enough hubs must be allocated along the columns such that no node in a column-ring is further than $\lceil C/8 \rceil$ hops from a hub node. This requires at least $\lceil R/\lceil C/4 \rceil \rceil$ hub nodes along each column. An upper bound on the total number of hubs required is therefore $4R + \lceil R/\lceil C/4 \rceil \rceil \cdot C$, or approximately $8R$ hubs. Therefore the number of hubs h^* is bounded by $4R \leq h^* \leq 8R$. We can achieve the lower bound by using clever hub designation to allow the same nodes to serve as hubs for both a row-ring and a column-ring, reducing the total number of hubs required.

We describe a hub allocation scheme that uses the minimum number of hubs. For the first row, designate nodes $n_{1,h_1}, n_{1,h_2}, n_{1,h_3}$, and n_{1,h_4} to be hubs, where the column numbers h_i are given in Table 2. In the second row, nodes $n_{2,h_1 \oplus 1}, n_{2,h_2 \oplus 1}, n_{2,h_3 \oplus 1}$, and $n_{2,h_4 \oplus 1}$ are hubs; note that this is a cyclic shift in the rightward direction (modulo C) of the hub allocation for the previous row. This pattern repeats for all subsequent rows. In general, $n_{r,c}$ is a hub if $c \oplus$

Table 2
Hub column numbers for the first row-ring

if $C \bmod 4 = \dots$	h_1	h_2	h_3	h_4
0	0	$C/4$	$2C/4$	$3C/4$
1	0	$\lfloor C/4 \rfloor$	$2\lfloor C/4 \rfloor$	$2\lfloor C/4 \rfloor + \lceil C/4 \rceil$
2	0	$\lceil C/4 \rceil$	$\lfloor C/4 \rfloor + \lceil C/4 \rceil$	$\lfloor C/4 \rfloor + 2\lceil C/4 \rceil$
3	0	$\lceil C/4 \rceil$	$2\lceil C/4 \rceil$	$3\lceil C/4 \rceil$

$(r - 1) = h_i$ for some $i = 1, 2, 3$, or 4 ; i is called the *hub index* of that hub.

It is trivial to note that each row now has four hubs, and hence only $P\lceil C/8 \rceil$ local wavelengths are required along the rows. The following lemma claims that this hub allocation also requires no more than $P\lceil C/8 \rceil$ local wavelengths along the columns.

Lemma 3. Designating hub nodes along the row-rings as described results in a hub allocation with the property that along each column-ring, no local node is more than $\lceil C/8 \rceil$ hops away from the nearest hub node.

Proof. The proof will show that along any column, no two hubs are separated by more than $\lceil C/4 \rceil$ nodes, from which it follows that no node can be more than $\lceil C/8 \rceil$ hops away from a hub.

Recall that a node $n_{r,c}$ is a hub if $c \oplus (r - 1) = h_i$; equivalently, $n_{r,c}$ is a hub iff $r = h_i \ominus (c - 1)$ for some $i = 1, 2, 3$, or 4 . Consider two adjacent hubs in the same column $n_{r_1,c}, n_{r_2,c}$ with hub indices i and j ; the distance between them is $r_1 - r_2 = h_i - h_j$. From Table 2, for any two adjacent h_i and h_j , $h_i - h_j \leq \lceil C/4 \rceil$, completing the proof. \square

The consequence of the lemma is that any local node is no more than $\lceil C/8 \rceil$ nodes away from a hub; hence no more than $P\lceil C/8 \rceil$ wavelengths are required. Fig. 11 gives an example of a 21×21 torus, and illustrates the hub allocation obtained from this construction.

The number of bypass wavelengths required can be obtained by examining the number of calls arriving and departing from each hub node. Along the rows, each hub is responsible for $\lceil C/4 \rceil$ nodes, and each such node has $2P$ ports; therefore each row-hub has $P' = 2P\lceil C/4 \rceil$. Using the $\lceil PN/4 \rceil$ algorithm,




Fig. 11. A 21×21 torus. The local nodes are at all intersection points of the grid, while the hub nodes are shown as shaded circles. For $R = C = 9$, we have that $\lfloor N/4 \rfloor = 2$ and $\lceil N/4 \rceil = 3$, so $h_1 = 0, h_2 = 2, h_3 = 4$, and $h_4 = 7$. Adding each of the row numbers modulo 9 gives the hub assignments shown. As a check, note that $g_1 = 0, g_2 = 2, g_3 = 5$, and $g_4 = 7$ also correctly yields the resulting hub allocations down the columns.

$P\lceil C/2 \rceil$ bypass wavelengths are required. A similar argument on the columns shows that $P\lceil R/4 \rceil$ bypass wavelengths are required.

The wavelength assignment also has the following properties:

- (1) *Non-hub nodes have fixed routing and fixed wavelengths.* Conversion between local wavelengths is hence not required, since each node is assigned its own wavelength to send and receive from its hub. Also, local nodes need no knowledge of network state.

- (2) *The only place where conversion is required is at the hubs.* No conversion is required to connect calls continuing from a row-ring onto a column-ring, or from a column-ring onto a row-ring.

4. Tree networks

4.1. The tree lower bound

In this section, a bound on the minimum number of wavelengths required to support P -port traffic in trees is established. We use the cut-set bound to obtain a lower bound on the number of wavelengths. Since a tree contains no cycles, the removal of any single link i disconnects the tree into two disjoint sets of nodes. For each i , we call these sets S_i^1 and S_i^2 , and let $|S_i^1|$ and $|S_i^2|$ denote the number of nodes in each set, respectively. Suppose $|S_i^1| \leq |S_i^2|$. Then there exists a worst-case admissible set where each node in S_i^1 sends all P units of traffic to some node in S_i^2 . Since all this traffic must cross link i , at least $W_i = |S_i^1|$ wavelengths are required to support it. If $|S_i^1| > |S_i^2|$, similar reasoning gives $W_i = |S_i^2|$.

We can obtain the tightest lower bound by maximizing over all links i . Let the greatest lower bound thus obtained be $W = \max_i \{W_i\} = \max_i \{\min\{|S_i^1|, |S_i^2|\}\}$. A link which achieves this lower bound is known as a *bottleneck link*. Fig. 12 gives an example of determining the bottleneck link.

From the preceding discussion, it is clear that the lower bound obtained in this section is dependent on the specific topology of the tree, and not just on the number of nodes in the tree. For example, for a balanced d -ary tree, the links adjacent to the root node are the bottleneck links, and $W = P(N - 1)/d$. For balanced binary trees, where $d = 2$, the bound is $W = P(N - 1)/2$.

4.2. The $\lceil PN/2 \rceil$ embedded-ring approach

In this section, we describe a RWA based on embedding a virtual ring in the tree topology. We will show that for connected P -port traffic sets, this approach requires at most $\lceil PN/2 \rceil$ wavelengths for any tree topology, and hence is optimal for tree topologies where $W = \lceil PN/2 \rceil$. For example, it is

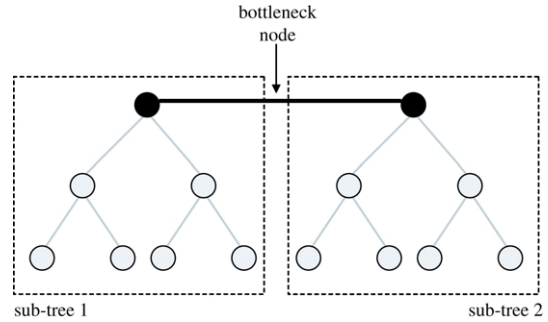


Fig. 12. A 14-node balanced binary tree. The bottleneck link is the link in heavy black; removal of this link from the graph disconnects the graph into the two indicated sub-trees, each containing $|S_i^1| = |S_i^2| = 7$ nodes. By considering all links in turn, it can be shown that this is the bottleneck link, resulting in a lower bound of seven wavelengths.

optimal for balanced binary trees. Furthermore, no wavelength conversion is required.

The ring-embedding idea is intuitively very simple. In any tree, by using depth-first search, we can form a circuit which visits each node in the tree at least once while traversing each link only twice (once in each direction). This circuit is said to form a *virtual ring* in the following sense. Consider a ring topology where the nodes are connected in the order in which each corresponding node in the tree is first visited by the circuit. Then any RWA for this ring has a one-to-one correspondence with a RWA for the tree. Each link between two adjacent nodes on the ring corresponds to the links traversed by the circuit in travelling between those two nodes on the tree. Such a circuit is illustrated in Fig. 13 for a 15-node balanced binary tree.

A single, unidirectional wavelength on the ring corresponds to the use of a single, bidirectional wavelength on the tree. (A bidirectional wavelength is required because a single circuit around the tree used each link once in each direction.) Recall from Lemma 1 that any two adjacent calls on a ring must fit on a single wavelength in one direction or the other. The ring-embedding algorithm simply divides the traffic into adjacent pairs, and determines the single-wavelength direction for each pair. Each pair is then routed on a single directed wavelength on the virtual ring, which corresponds to a RWA on the tree which uses a single bidirectional wavelength per pair. Since there are a total of $\lceil PN/2 \rceil$ pairs, no more than $\lceil PN/2 \rceil$ wavelengths are required.

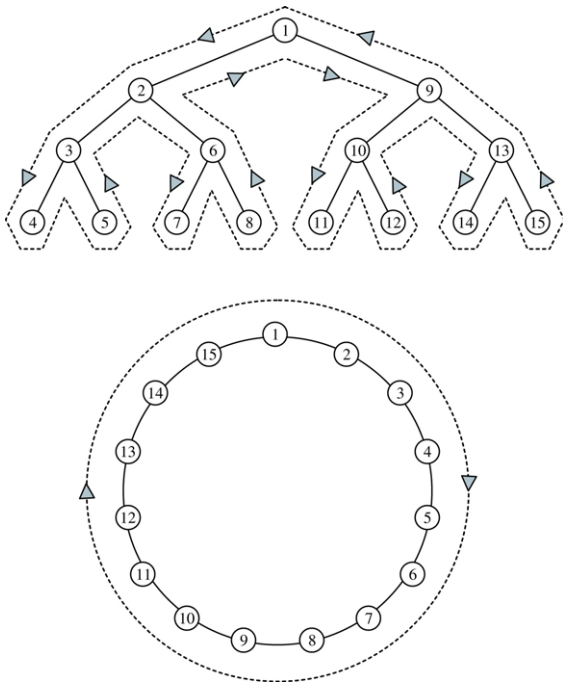


Fig. 13. Embedding a cycle in a 15-node balanced binary tree. The nodes have been numbered so that the cycle visits them in order of increasing index. The corresponding virtual ring topology is shown below the tree.

4.3. Banding and bypass using the embedded-ring approach

Banding can also be performed using the embedded-ring approach by directly applying the ring banding algorithm on the virtual ring. For a four-hub architecture, for example, the approach of Section 2.4 can be directly used to obtain a RWA for the virtual ring which uses $\lceil PN/8 \rceil$ local wavelengths and $\lceil PN/4 \rceil$ bypass wavelengths. Since a single unidirectional wavelength on the virtual ring requires a bidirectional wavelength on the tree, this implies that $2\lceil PN/8 \rceil$ local wavelengths and $2\lceil PN/4 \rceil$ bypass wavelengths are required for the tree using this approach.

5. Conclusions

We considered routing and wavelength assignment for P -port traffic on ring, torus, and tree topologies with and without wavelength conversion. For each

topology, we first used the objective of minimizing wavelengths, then moved on to maximizing optical bypass via banding.

For rings, we recap the $\lceil PN/3 \rceil$ and $\lceil PN/4 \rceil$ algorithms from [16] and [19] which are optimal in minimizing the number of wavelengths used for networks without and with wavelength conversion, respectively. We then propose a novel approach for banding which minimizes the size of the band of local wavelengths dropped at all nodes, and additionally eliminates the need for switching at all but h hub nodes, where h is a design parameter. Furthermore, only the hubs need to have knowledge of the network state, eliminating the need for a centralized management scheme which needs to coordinate with all nodes.

For tori, we derive a lower bound of $\lceil PR/4 \rceil$ wavelengths required to support P -port traffic, and provide a novel algorithm (the TERA algorithm) which uses ring embeddings and requires $\max\{\lceil PC/2 \rceil, \lceil PR/4 \rceil\}$ wavelengths for an $R \times C$ torus using sparse conversion. We show that this is optimal for $R \geq 2C$. This improves over the algorithm in [18] which requires $\lceil PR/2 \rceil$ wavelengths. We then provide a banding algorithm based on the ring banding approach which eliminates switching at all but $4R$ hubs. Again, only the hubs need information about the network state in this second RWA.

Finally, for trees, we describe a method based on the cut-set bound which obtains a lower bound on the number of wavelengths required to support any P -port tree. We then use a ring-embedding approach to produce the $\lceil PN/2 \rceil$ algorithm, which has low complexity and is optimally efficient in wavelength usage for a certain subset of trees, including balanced binary trees. A method for applying the ring banding algorithm to trees for bypass requiring only half the wavelengths to be dropped at local nodes is then derived.

References

- [1] R. Ramaswami, K.N. Sivarajan, *Optical Networks: A Practical Perspective*, Morgan Kaufmann, Los Altos, CA, 1998.
- [2] I. Chlamtac, A. Ganz, G. Karmi, Lightpath communications: an approach to high bandwidth optical WANS, *IEEE Trans. Commun.* 40 (1992) 1171–1182.
- [3] D. Banerjee, B. Mukherjee, Wavelength-routed optical networks: linear formulation, resource budgeting tradeoffs,

- and a reconfiguration study, *IEEE/ACM Trans. Networking* 8 (2000) 598–607.
- [4] R.M. Krishnaswamy, K.N. Sivarajan, Design of logical topologies: a linear formulation for wavelength-routed optical networks with no wavelength changers, *IEEE/ACM Trans. Networking* 9 (2001) 186–198.
- [5] R. Ramaswami, K.N. Sivarajan, Routing and wavelength assignment in all-optical networks, *IEEE/ACM Trans. Networking* 3 (1995) 489–500.
- [6] A. Birman, Computing approximate blocking probabilities for a class of all-optical networks, *IEEE J. Select. Areas Commun.* 14 (5) (1996) 852–857.
- [7] R.A. Barry, P.A. Humblet, Models of blocking probability in all-optical networks with and without wavelength changers, *IEEE J. Select. Areas Commun.* 14 (1996) 858–867.
- [8] M. Kovacevic, A. Acampora, Benefits of wavelength translation in all-optical clear-channel networks, *IEEE J. Select. Areas Commun.* 14 (1996) 868–880.
- [9] S. Subramaniam, M. Azizoglu, A. Somani, A new analytical model for multifiber WDM networks, *IEEE J. Select. Areas Commun.* 18 (2000) 2138–2145.
- [10] Y. Zho, G.N. Rouskas, H.G. Perros, A path decomposition algorithm for computing blocking probabilities in wavelength routing networks, *IEEE/ACM Trans. Networking* 8 (2000) 747–762.
- [11] S. Ramesh, G.N. Rouskas, H.G. Perros, Computing blocking probabilities in multiclass wavelength-routing networks with multicast calls, *IEEE J. Select. Areas Commun.* 20 (2002) 89–96.
- [12] E. Karasan, E. Ayanoglu, Effects of wavelength routing and selection algorithms on wavelength conversion gain in WDM optical networks, *IEEE/ACM Trans. Networking* 6 (1998) 186–196.
- [13] L. Li, A.K. Somani, Dynamic wavelength routing using congestion and neighborhood information, *IEEE/ACM Trans. Networking* 7 (1999) 779–786.
- [14] O. Gerstel, G. Sasaki, S. Kuttan, R. Ramaswami, Worst-case analysis of dynamic wavelength allocation in optical networks, *IEEE/ACM Trans. Networking* 7 (1999) 833–845.
- [15] R. Ramaswami, G. Sasaki, Multiwavelength optical networks with limited wavelength conversion, *IEEE/ACM Trans. Networking* 6 (1998) 744–754.
- [16] A. Narula-Tam, P.J. Lin, E. Modiano, Efficient routing and wavelength assignment for reconfigurable WDM networks, *IEEE J. Select. Areas Commun.* 20 (2002) 75–88.
- [17] P. Saengudomlert, E. Modiano, R.G. Gallager, An on-line routing and wavelength assignment algorithm for dynamic traffic in a WDM bidirectional ring, in: *Proc. JCIS*, March 2002.
- [18] P. Saengudomlert, E. Modiano, R.G. Gallager, On-line routing and wavelength assignment for dynamic traffic in WDM ring and torus networks, in: *Proc. IEEE INFOCOM*, April 2003.
- [19] L.-W. Chen, E. Modiano, Efficient routing and wavelength assignment for reconfigurable WDM networks with wavelength converters, in: *Proc. IEEE INFOCOM*, April 2003.
- [20] D. Bertsekas, R. Gallager, *Data Networks*, 2nd edition, Prentice-Hall, Englewood Cliffs, NJ, 1992.
- [21] M.N.S. Swamy, K. Thulasiraman, *Graphs, Networks, and Algorithms*, Wiley, New York, 1981.
- [22] P.J. Lin, *Wide area optical backbone networks*, Ph.D. Thesis, MIT, February 1996.