Efficient Routing and Wavelength Assignment for Reconfigurable WDM Networks

Aradhana Narula-Tam, *Member, IEEE*, Philip J. Lin, *Member, IEEE*, and Eytan Modiano, *Senior Member, IEEE*

***Abstract*—Through the use of configurable wavelength-division-multiplexing (WDM) technology including tunable optical transceivers and frequency selective switches, next-generation WDM networks will allow multiple virtual topologies to be dynamically established on a given physical topology. For N node P port networks, we determine the number of wavelengths required to support *all possible* virtual topologies (PN lightpaths) on a bidirectional ring physical topology. We show that if shortest path routing is used, approximately N wavelengths are needed to map N lightpaths. We then present novel adaptive lightpath routing and wavelength assignment strategies that reduce the wavelength requirements to $\lceil (N/2) \rceil$ working wavelengths per port for protected networks and $\lceil (N/3) \rceil$ wavelengths in each direction per port for unprotected networks. We show that this reduced wavelength requirement is optimal in the sense that it is the minimum required to support the worst case logical topology. Furthermore, we prove that a significant number of logical topologies require this minimum number of wavelengths. We also develop joint routing and wavelength assignment strategies that not only minimize the number of wavelengths required to implement the worst case logical topologies but also reduce average wavelength requirements. Finally, methods for extending these routing and wavelength assignment results to general two-connected and three-connected physical topologies are presented.**

***Index Terms*—Joint routing and wavelength assignment (RWA), logical topology reconfiguration, wavelength division multiplexing, wavelength requirements.**

I. Introduction

N wavelength-division-multiplexing (WDM) systems, multiple signals, separated by wavelength, are carried concurrently on an optical fiber. Each wavelength (channel) operates at peak electronic speeds of 1–10 Gb/s per channel. Configurable optical add/drop multiplexers (ADMs) and cross-connects may be used to allow individual wavelength signals either to be *dropped* to the electronic routers at each node or to pass through the node optically. There are two general classes of WDM network architectures: single-hop and multihop. A single-hop WDM network is an all-optical network in which network traffic is converted to electronics only at the source and destination nodes [1]. In multihop networks, most of the traffic is electronically processed at intermediate node routers between the source and destination [2].

In both single-hop and reconfigurable multihop networks, nodes are typically equipped with a small number of tunable transmitters and receivers. A lightpath between two nodes is formed by tuning the transmitter of one node and the receiver of another node to the same wavelength, and configuring the ADM or cross-connect switches appropriately. Thus a lightpath is unidirectional. The *physical topology* of the network consists of optical nodes and their fiber connections. The *logical topology* describes the lightpaths between the nodes and is determined by the configuration of the transmitters, receivers, and switches on each node. In single-hop networks, extremely rapidly tunable transceivers are required to efficiently time-share the network transceiver ports. Multihop networks may not need to be reconfigured as rapidly since in a connected logical topology, each node can transmit packets to every other node via store and forward or similar mechanisms. Logical topology reconfiguration in multihop networks may be used to reduce network delay and electronic processing loads, as proposed in [3]–[7].

An important characteristic of both single-hop and multihop WDM networks is the independence between the logical and physical topologies. Any logical topology may be implemented on a given connected physical topology if enough wavelengths are available. A network with N nodes and P transceiver ports per node can have up to PN lightpaths. If each lightpath must be routed on a different wavelength, PN wavelengths are required. Building networks with PN wavelengths may not be feasible, since the number of wavelengths available is technology-limited. Furthermore, implementing PN wavelengths may be an expensive and inefficient use of resources. In a ring physical topology, for example, it may be possible to route multiple lightpaths on the same wavelength by intelligent routing and wavelength assignment. Through wavelength reuse, the network wavelength requirements can be significantly reduced.

In this paper, our primary objectives are to determine the minimum number of wavelengths required to allow *all possible* logical topologies to be embedded on a ring physical topology and to obtain routing and wavelength assignment strategies that achieve this minimum wavelength requirement. Supporting all logical topologies provides maximal flexibility in the network design and is preferable since traffic demands are unknown *a priori*. We consider the static or off-line problem for N node P transceiver networks where all PN lightpath requests are received simultaneously; equivalently we design a rearrangeably

Manuscript received February 21, 2001; revised July 29, 2001. The work of A. Narula-Tam was supported by the Defense Advanced Research Projects Agency under AF Contract F19628-00-C-0002. The work of E. Modiano was supported by the National Science Foundation under Grant ANI-0073730.

A. Narula-Tam is with the Massachusetts Institute of Technology, Lincoln Laboratory, Lexington, MA 02420-9108 USA (e-mail: arad@ll.mit.edu).

P. J. Lin is with the Tellabs Research Center and MIT Laboratory for Information and Decision Systems, Cambridge, MA 02139 USA (e-mail: pjlin@ieee.org).

E. Modiano is with the Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

Publisher Item Identifier S 0733-8716(02)00152-X.

nonblocking network. A similar problem has been addressed by Gerstel *et al.* in [8], where minimum wavelength requirements are determined for dynamic lightpath requests (lightpaths are assigned and released on demand) with a maximum of L lightpaths per link. In [8], the worst case wavelength requirements are calculated to ensure a wide-sense nonblocking network. Gerstel *et al.* assume fixed routing schemes, whereas in our work, we determine optimal *joint* routing and wavelength assignment strategies. These routing and wavelength assignment strategies are optimal in the sense that they minimize the network wavelength requirements for networks that must support *all* logical topologies. Although the wavelength requirements are not minimized for each logical topology instance, we show that a significant number of logical topologies require this minimum number of wavelengths; thus the wavelength requirement is not overengineered to support a small number of pathological topologies. We also develop heuristic joint routing and wavelength assignment algorithms that reduce the number of wavelengths required to map logical topologies on *average*, in addition to minimizing the number of wavelengths required for the *worst case* logical topology.

There has been a considerable amount of work in the area of efficient routing and wavelength assignment (RWA). Most prior work decouples the two problems of assigning routes and assigning wavelengths. For example, [9]–[15] assume fixed routing schemes. Given a routing that results in a maximum of L lightpaths on any link, [9] shows that a maximum of $2L-1$ wavelengths are needed to establish all lightpaths. In [10], the static wavelength assignment problem alone is shown to be NP-complete. Wavelength assignment for fixed alternate path routing, where each lightpath route may be selected from a subset of all possible routes, is investigated in [11], [16], and [13]. In [16], the routing and wavelength assignment problem is formulated as an integer linear program (ILP), where given a fixed number of available wavelengths, the goal is to maximize the number of lightpaths supported. The computational complexity of the ILP is somewhat reduced when considering bidirectional ring physical topologies, where the number of alternate paths is two [13]. The ILP is generalized to the case of unconstrained routing (consider all paths between lightpath source and destination) in [17]. Heuristics for unconstrained routing and wavelength assignment are also evaluated in [17]. In [18] and [13], an initial routing and wavelength assignment based on shortest path routing is subsequently modified to reduce wavelength utilization. In this paper, we propose new schemes for joint routing and wavelength assignment that take advantage of the characteristics a logical topology with P ports per node to reduce *both* average and worst case wavelength requirements.

We present our network model and assumptions in Section II. Wavelength requirements are calculated in Section III, where Section III-A focuses on networks using deterministic shortest path routing to route lightpaths and Section III-B develops adaptive *joint* routing and wavelength assignment strategies that minimize network wavelength requirements. In Section IV, we show that a significant number of logical topologies require the maximum number of wavelengths. In Section V, routing and wavelength assignment algorithms are developed that both minimize the worst case wavelength requirements and reduce the average wavelength requirements over fixed routing. Finally, in Section VI, extension of the proposed joint routing and wavelength assignment principles to two-connected and three-connected physical topologies is discussed.

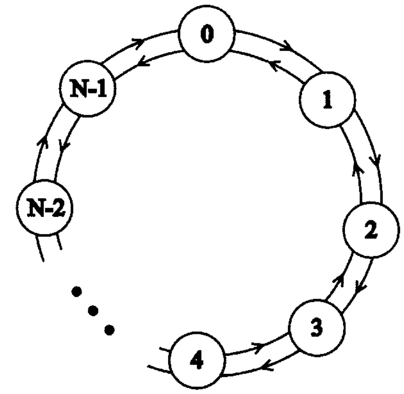

Fig. 1. A bidirectional ring physical topology with a single fiber propagating in each direction.

II. Network Model

We consider networks with N nodes, each equipped with P transceiver ports; thus each logical topology has PN lightpaths. If the physical topology is a unidirectional ring, there is only one possible route between every source and destination pair. Thus, there exist worst case logical topologies that require PN wavelengths. Consider, for example, a logical topology consisting of P rings, where the nodes in each logical ring are ordered in direction opposite to the physical topology. In this case, each lightpath requires a separate wavelength and a total of PN wavelengths are needed. We therefore focus on wavelength provisioning for bidirectional ring physical topologies where route selection can impact the number of required wavelengths.

We consider a bidirectional ring physical topology, shown in Fig. 1, consisting of a minimum of two fibers where half the fibers have wavelengths propagating in the clockwise direction and half the fibers propagate wavelengths in the counterclockwise direction. We assume throughout that the nodes are labeled in increasing order in the clockwise direction. In determining wavelength provisioning requirements, we employ the following accounting method: a set of lightpaths requires one wavelength if the set of lightpaths can be routed on a single wavelength on the same fiber. If a set of lightpaths uses the same color wavelength on both the clockwise and the counterclockwise fibers, we say that the set of lightpaths utilizes *two* wavelengths.

We consider two types of networks: protected and unprotected. For the protected network case, we assume loop-back protection [19] so that half of the total capacity is reserved for protection. If a lightpath is routed on a wavelength in the clockwise direction fiber, a wavelength on the counterclockwise direction fiber is reserved for protection and vice versa. On each fiber, the number of wavelengths used for working traffic changes with different logical topologies. However, the total number of working traffic wavelengths is always equal to the total number of protection wavelengths. When assessing wavelength requirements, we determine the number of wavelengths needed for working traffic. In the protected network, the working traffic wavelength requirements are not restricted by direction, since one always allocates an opposite direction wavelength for protection. In an unprotected network, wavelengths on the bidirectional ring network should be allocated in counterpropagating pairs since there is no benefit in reducing the wavelength requirements in only one direction. If a logical topology requires only x wavelengths in the clockwise direction but $x + \Delta$ wavelengths in the counterclockwise direction, then there also exists a logical topology that requires x counterclockwise wavelengths and $x + \Delta$ clockwise wavelengths. To see this, simply reverse the lightpath directions of the first logical topology. Thus, in order to accommodate both topologies, the network must provide $x + \Delta$ wavelengths in both directions. These differences in the protected and unprotected network cases lead to differing routing and wavelength assignment strategies as well as different wavelength requirements for the two types of networks.

We consider both connected and general logical topologies, where a general logical topology can be either connected or unconnected. Connectivity in a multihop network ensures that traffic between every source and destination pair can be continually supported. There may, however, be scenarios where an instance of the logical topology is not necessarily connected, for example, in rapidly tunable single-hop networks.

III. Wavelength Requirements

A. Deterministic Shortest Path Routing

Deterministic shortest path routing (DSPR) is a routing method that assigns a *fixed* shortest path route for each source to destination pair, i.e., if there are multiple shortest path routes between a source and destination node, one of these routes is selected and used exclusively. DSPR schemes are often used because they are simple and because they minimize the resources required to route each lightpath. However, in networks without wavelength converters, the number of wavelengths required to implement a logical topology can be substantially larger than optimal. In this section, we determine wavelength requirements for a network utilizing DSPR. We require the network to be capable of implementing all possible logical topologies. We present two DSPR schemes and then calculate lower and upper bounds on corresponding network wavelength requirements. These bounds will be used to compare the benefits of our adaptive routing and wavelength assignment algorithms to shortest path routing.

When the number of nodes N is even, there are two shortest paths from node i to node $i + (N/2)$. In the deterministic odd even shortest path (DOES) [14] routing scheme, a shortest path between nodes i and $i + (N/2)$ is routed clockwise if i is odd and counterclockwise if i is even. DOES routing was shown to require fewer wavelengths than routing all length $(N/2)$ paths in the same direction. An alternative routing scheme for N that is even preferable to DOES in some cases routes lightpaths from node i to node $i + (N/2)$ and from node $i + (N/2)$ to node i in the clockwise (counterclockwise) direction if i is odd (even) for $0 \le i < (N/2)$. We call this deterministic continuous ring shortest path (DCRS) routing. Note that DOES and DCRS only differ when N is even and $(N/2)$ is odd.

TABLE I
Wavelength Requirements Under Shortest Path Routing for Single Port Per Node Networks. Requirements Are for Working Wavelengths in a Protected Network

		N odd	N even		
			$\frac{N}{2}$ even	$\frac{N}{2}$ odd (DOES)	$\frac{N}{2}$ odd (DCRS)
Connected	W_{LB}	$N-2$	$N-2$	$N-3$	$N-2$
	W_{UB}	$N-2$	$N-1$	$N-1$	$N-1$
General	W_{LB}	$N-1$	$N-2$	N	$N-2$
	W_{UB}	$N-1$	$N-1$	N	$N-1$

Table I shows upper and lower bounds on the number of wavelengths required to implement all possible logical topologies (connected and general) on a network with N nodes and one port per node. A lower bound of W_{LB} indicates that there exists a logical topology that requires at least W_{LB} wavelengths. An upper bound of W_{UB} implies that no logical topology requires more than W_{UB} wavelengths. An upper bound of N is trivial since N lightpaths can require at most N wavelengths. The tighter upper bounds in Table I are obtained by showing that for every logical topology, there exist some lightpaths that are able to share a wavelength. Proofs for the lower bounds are by construction. We derive the lower bound for the N odd connected logical topology case below. The proofs for the other cases are similar.

For N odd, there exists a connected logical topology that requires $N-2$ wavelengths when the lightpaths are routed using shortest path routing. We construct such a connected logical topology as follows. Create lightpaths to connect node i to node $i+\lfloor(N/2)\rfloor$ for $0 \le i < \lfloor(N/2)\rfloor$ and connect node $i+\lfloor(N/2)\rfloor$ to node $i+1$ for $0 \le i < \lfloor(N/2)\rfloor-1$. Each of these $N-2$ lightpaths requires a separate wavelength since each lightpath uses link $(\lfloor(N/2)\rfloor - 1, \lfloor(N/2)\rfloor)$. The final two lightpaths connect node $N-2$ to node $N-1$ and node $N-1$ to node 0. These two lightpaths can share a wavelength with the lightpath from node 0 to node $\lfloor(N/2)\rfloor$. Since $N-2$ of the lightpaths use physical link $(\lfloor(N/2)\rfloor-1, \lfloor(N/2)\rfloor)$, a minimum of $N-2$ wavelengths are required to implement this logical topology. An example for $N = 7$ is shown in Fig. 2. This example illustrates that if routing is deterministic, there will always be logical topologies that require close to N wavelengths.

In calculating the bounds in Table I, no restriction was placed on the directions of the wavelength channels; thus these results correspond to the working traffic wavelength requirements in protected networks. Restricting the wavelength directions can only increase wavelength requirements. Wavelength requirements are shown for both connected and general logical topologies. Trivial upper bounds for P port per node networks can be obtained by multiplying the upper bounds in Table I by P.

Since the lower and upper bounds in Table I are similar and near N, approximately N wavelengths per port are required to

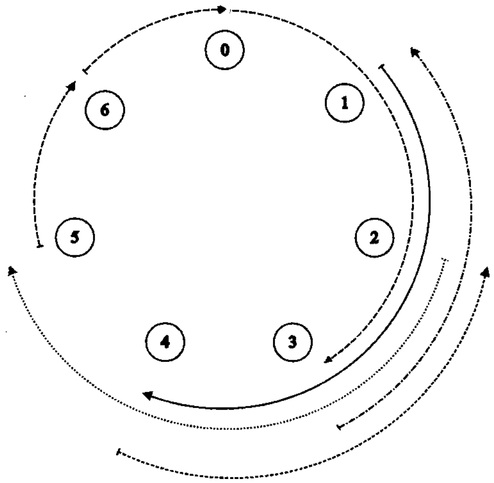

Fig. 2. An example of a connected logical topology that requires a minimum of $W_{\rm LB} = N - 2$ wavelengths when shortest path routing is used.

ensure that *all possible* logical topologies can be established using deterministic shortest path routing. Routing each lightpath on a separate wavelength also requires N wavelengths for each set of N lightpaths. Thus DSPR-based routing and wavelength assignment schemes require nearly the maximum number of wavelengths. In the next section, we develop *adaptive* routing and wavelength assignment schemes that significantly reduce network wavelength requirements.

B. *Adaptive Routing and Wavelength Assignment*

By approaching the routing and wavelength assignment problems jointly, we can reduce the number of network wavelengths required to support all possible logical topologies. Let $W_{\rm req}$ denote the minimum network wavelength requirement. We determine $W_{\rm req}$ by 1) developing adaptive routing and wavelength assignment schemes that can map any set of lightpaths with less than or equal to $W_{\rm req}$ wavelengths and 2) showing that there exist logical topologies that cannot be supported (under any routing strategy) if fewer than $W_{\rm req}$ wavelengths are available. Since the above-mentioned routing and wavelength assignment schemes can implement all logical topologies within the minimum wavelength requirement, the RWA strategies are optimal. Optimal adaptive RWA strategies and wavelength requirements are determined for both protected and unprotected networks. Although these lightpath routing and wavelength assignment algorithms do not minimize the wavelength requirements for *each* logical topology, they do minimize the number of wavelengths required to implement *all possible* logical topologies on the bidirectional ring physical topology.

We consider both connected and general (connected or unconnected) logical topologies. In each section below, we start by considering connected topologies and then generalize the routing and wavelength assignment algorithms and wavelength requirement computations to cover both connected and unconnected topologies.

1) Single Port Per Node Networks: Connectivity in a single port per node network implies that the logical topology forms a ring. Each logical topology can be written as a permutation of the N nodes $(i_0, i_1, \ldots, i_{N-1})$ where there is a lightpath from node i_k to node $i_{(k+1) \bmod N}$, for $0 \le k < N$. Note that when written in ring order, consecutive lightpaths (i_k, i_{k+1}) and (i_{k+1}, i_{k+2}) are adjacent.[1] This fact will be used to develop the RWA algorithms below.

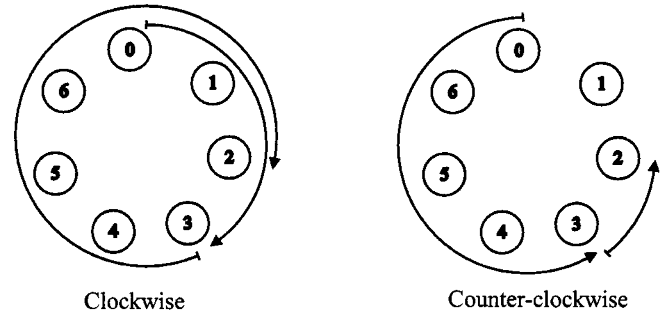

Fig. 3. Two adjacent lightpaths [example (0,3) and (3,2)] can share a wavelength in either the clockwise or counterclockwise direction.

In protected networks, half the wavelengths are used for working traffic and half are used for protection. Each working traffic wavelength can be assigned in either direction since each clockwise (counterclockwise) working traffic wavelength is protected by a counterclockwise (clockwise) protection wavelength. In unprotected networks, wavelengths should be assigned in clockwise/counterclockwise pairs to minimize wavelength requirements.

a) Protected Networks $(P = 1)$*:* The following theorems show that $\lceil (N/2) \rceil$ working traffic wavelengths are necessary and sufficient to implement any connected logical topology. The first theorem relies on the following lemma, which outlines an efficient routing and wavelength assignment strategy.

Lemma 1: In a bidirectional ring physical topology, every pair of adjacent lightpaths can share a wavelength in one of the two directions.

Proof: Let (i_j, i_k) denote a lightpath from source node i_j to destination node i_k. Consider two adjacent lightpaths (i_1, i_2) and (i_2, i_3). If (i_1, i_2) and (i_2, i_3) cannot share a wavelength in the clockwise direction, then i_2 must lie between i_1 and i_3 on the counterclockwise direction fiber; hence, the two lightpaths can share a wavelength in the counterclockwise direction, as shown in Fig. 3. ■

Theorem 1: The maximum number of working wavelengths needed to implement any connected logical topology is equal to $\lceil (N/2) \rceil$.

Proof: By Lemma 1, each pair of adjacent lightpaths can share a wavelength. Since the logical topology forms a ring, the set of lightpaths can be divided into $\lfloor (N/2) \rfloor$ pairs of adjacent lightpaths plus one lightpath if N is odd. Therefore, the maximum number of wavelengths required to route all N lightpaths is $\lceil (N/2) \rceil$. ■

Theorem 2: For $N > 3$, there exists a connected logical topology that requires $\lceil (N/2) \rceil$ wavelengths (regardless of the routing strategy).

Proof: We can construct logical topologies for N odd and N even that require $\lceil (N/2) \rceil$ wavelengths. Example topologies are shown in Fig. 4.

[1]Two lightpaths are adjacent if the destination node of one lightpath is equal to the source node of the other lightpath.

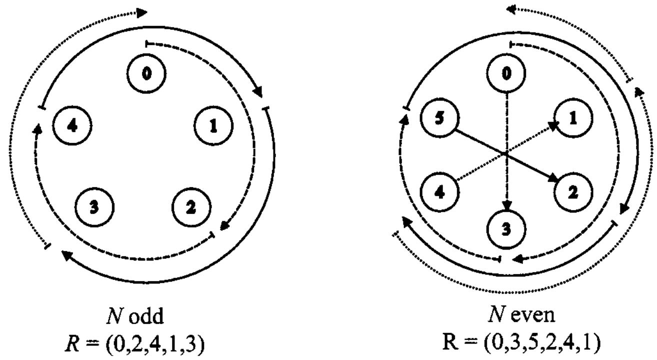

Fig. 4. Example light path topologies that require $\lceil (N/2) \rceil$ wavelengths for N odd and N even. For N even, the $(N/2)$ lightpaths (0,3), (5,2), and (4,1) overlap; thus each overlapping lightpath requires a separate wavelength.

For N *odd*, consider a logical topology connecting node i to node $(i + \lfloor (N/2) \rfloor) \bmod N$. Since each lightpath traverses at least $\lfloor (N/2) \rfloor$ links, at most two lightpaths can share a wavelength. Therefore, at least $\lceil (N/2) \rceil$ wavelengths are required to route all N lightpaths of the logical topology.

For N *even*, the preceding construction does not produce a connected logical topology; therefore, we use a different construction. Construct a connected logical topology by connecting node i_0 to node $i_1 = i_0 + (N/2)$. Next connect node i_1 to node $i_2 = i_1 + (N/2) - 1$. Continue creating lightpaths sequentially in this manner, alternating between adding $(N/2)$ and $(N/2)-1$ until node i_{N-1}. Node i_{N-1} is connected to node i_0. In this logical topology, $(N/2)$ of the lightpaths traverse $(N/2)$ links each, regardless of the routing. Since these $(N/2)$ lightpaths overlap each other, as shown in Fig. 4, each requires a separate wavelength. Therefore, at least $(N/2)$ wavelengths are needed to support this logical topology. ■

These results indicate that by routing pairs of adjacent lightpaths on a single wavelength, any connected logical topology can be supported on a network provisioned with $\lceil (N/2) \rceil$ wavelengths. Furthermore, since $\lceil (N/2) \rceil$ is the minimum number of wavelengths required to support some logical topologies and our objective is to support all connected logical topologies, this adaptive routing strategy is optimal.

We can now generalize the problem to consider both connected and unconnected logical topologies. Recall that a connected logical topology consists of a single directed circuit. In a single port per node network in which all ports are utilized, a disconnected logical topology is the union of multiple edge-disjoint directed circuits[2] [20, Theorem 5.6]. A general logical topology will consist of K directed circuits of size M_i, for $i = 1 \ldots K$, where $\sum_i M_i = N$. If pairs of adjacent lightpaths in each circuit are routed on a single wavelength, then applying Theorem 1 to each circuit shows that at most $\sum_i \lceil (M_i/2) \rceil$ wavelengths will be required. This can be substantially larger than $\lceil (N/2) \rceil$ if the logical topology consists of many odd size circuits. However, the following lemma can be used to show that at most $\lceil (N/2) \rceil + 1$ wavelengths are required to support all (connected or unconnected) logical topologies. Furthermore, we can show that there exist unconnected logical topologies that require $\lceil (N/2) \rceil + 1$ wavelengths.

Lemma 2: Given three directed circuits of odd sizes M_1, M_2, and M_3 that, when routed individually, require $\lceil (M_i/2) \rceil$ wavelengths each, there exists a lightpath from one of the three odd size circuits that can share a wavelength with a lightpath from one of the other two odd size circuits. Thus, the three directed circuits require a total of $\lceil (M_1 + M_2 + M_3)/2 \rceil$ wavelengths.

Proof: The proof of Lemma 2 is in Appendix I. ■

Theorem 3: The maximum number of wavelengths needed to implement any general (connected or unconnected) logical topology is equal to $\lceil (N/2) \rceil + 1$.

Proof: For connected logical topologies, we know from Theorem 1 that a maximum of $\lceil (N/2) \rceil$ wavelengths are required. If the unconnected logical topology contains one odd circuit M_j, then by Theorem 1, the even circuits each require $(M_i/2)$ wavelengths and the odd circuit requires $\lceil (M_j/2) \rceil$ for a total of $\lceil (N/2) \rceil$ wavelengths. If the logical topology contains two odd circuits M_j and M_k, then again by Theorem 1 the even circuits each require $(M_i/2)$ wavelengths and the odd circuits require $\lceil (M_j/2) \rceil$ and $\lceil (M_k/2) \rceil$ for a total of $\sum_i (M_i/2) + (M_j + M_k)/(2) + 1$, or one extra wavelength. Thus assume an unconnected logical topology that contains three odd size circuits of size M_1, M_2, and M_3. Use Lemma 2 and assume without loss of generality that circuits 1 and 2 contain the pair of lightpaths that can share a wavelength. These two circuits require $(M_1 - 1)/(2) + (M_2 - 1)/(2) + 1 = (M_1 + M_2)/(2)$ wavelengths. Thus, the three odd size circuits require $(M_1 + M_2)/(2) + \lceil (M_3/2) \rceil = \lceil (M_1 + M_2 + M_3)/(2) \rceil$ wavelengths. Lemma 2 may also be applied iteratively to logical topologies with larger numbers of odd size circuits. For example, consider a logical topology with five odd size circuits M_1 to M_5. Take any three of these circuits, e.g., circuits 1, 2, and 3, and apply Lemma 2. Two of the three circuits, e.g., circuits 1 and 2, will contain lightpaths that can share a wavelength. These two circuits use $(M_1 + M_2)/(2)$ wavelengths. There remain three odd size circuits 3, 4, and 5. Applying Lemma 2 to these three circuits shows that two of the three circuits, e.g., 3 and 4, contain a pair of lightpaths that can share a wavelength. Therefore, at most $(M_3 + M_4)/(2) + \lceil (M_5/2) \rceil$ wavelengths are needed to establish the three circuits. Consequently, a total of $\lceil (M_1 + M_2 + M_3 + M_4 + M_5)/(2) \rceil$ wavelengths are needed to establish all five circuits. In general, any logical topology contains at most two odd size circuits that require $\lceil (M_i/2) \rceil$ wavelengths each and that do not allow any further sharing of wavelengths between them. In a connected logical topology, these two circuits would require $(M_1 + M_2)/(2)$ wavelengths rather than $\lceil (M_1/2) \rceil + \lceil (M_2/2) \rceil$. Therefore, at most one extra wavelength, for a total of $\lceil (N/2) \rceil + 1$ wavelengths, is required to support unconnected as well as connected logical topologies. ■

b) Unprotected Networks ($P = 1$)*:* In unprotected networks, all wavelengths are used for working traffic. Simply applying the routing algorithms from the protected case when there are no protection wavelengths yields a requirement of $\lceil (N/2) \rceil$ wavelengths in each direction, since some logical topologies may require $\lceil (N/2) \rceil$ clockwise wavelengths while others may require may require $\lceil (N/2) \rceil$ counterclockwise wavelengths. We can do much better. The following theorems

[2]A directed circuit in a directed graph G is defined as a finite sequence of vertices $v_0, v_1, \ldots, v_k$ such that (v_{i-1}, v_i) is an edge in G, $v_0 = v_k$, and all other vertices are unique.

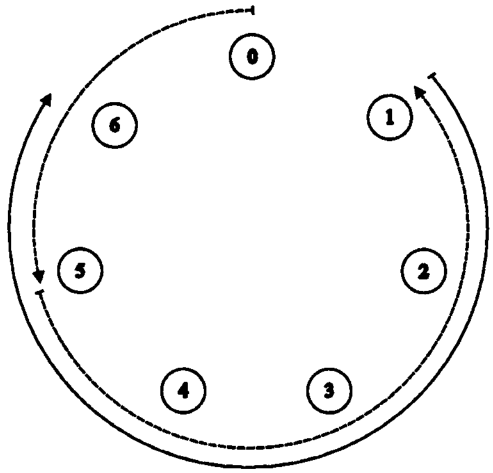

Fig. 5. Any three adjacent lightpaths [example: (0,5), (5,1), and (1,6)] can be routed using one wavelength in each direction.

show that $\lceil (N/3) \rceil$ wavelengths in each direction are necessary and sufficient to implement any connected logical topology. These reduced wavelength requirements are achieved by allocating wavelengths in counterpropagating pairs, as suggested in Section II. ■

Theorem 4: Any connected logical topology can be implemented with $\lceil (N/3) \rceil$ wavelengths in each direction.

Proof: For any given logical topology, use the following routing and wavelength assignment algorithm.

1) Divide the lightpaths into sets of three adjacent lightpaths. If N is not perfectly divisible by three, then there will be one set of lightpaths that has either one or two lightpaths in it.
2) Using Lemma 1, route the first two lightpaths in each set on a single wavelength. Route the third lightpath in each set on a wavelength in the opposite direction.

Since there are $\lceil (N/3) \rceil$ sets, at most $\lceil (N/3) \rceil$ wavelengths are required in each direction. ■

The proof illustrates a method of routing lightpaths to ensure that no more than $\lceil (N/3) \rceil$ wavelengths are needed in each direction. Fig. 5 illustrates the lightpath routing strategy. The next theorem illustrates that at least $\lceil (N/3) \rceil$ wavelengths must be provisioned in each direction, and hence, the optimality of the above routing and wavelength assignment algorithm. ■

Theorem 5: For networks with $N > 4$ nodes, a minimum of $\lceil (N/3) \rceil$ wavelengths in each direction are required to support all possible logical topologies.

Proof: Suppose that only $\lceil (N/3) \rceil - 1$ wavelengths are available in each direction. We can construct logical topologies that *cannot* be supported for the cases of N odd and N even.

For N *odd*, construct a logical topology by connecting node i to node $(i + \lfloor (N/2) \rfloor) \bmod N$. Since each lightpath traverses a minimum of $\lfloor (N/2) \rfloor$ links, at most two lightpaths can share a wavelength. Furthermore, since each lightpath traverses $\lceil (N/2) \rceil$ links in the counterclockwise direction, lightpaths can only share wavelengths in the clockwise direction. Using the $\lceil (N/3) \rceil - 1$ clockwise wavelengths, we can support $2\lceil (N/3) \rceil - 2$ lightpaths. Each of the $\lceil (N/3) \rceil - 1$ counterclockwise wavelengths supports only one lightpath. Thus the total number of lightpaths supported is $3(\lceil (N/3) \rceil - 1)$, which is always less than N. Hence, this logical topology cannot be implemented with less than $\lceil (N/3) \rceil$ wavelengths in each direction.

For N *even*, construct a connected logical topology as follows. Sequentially, starting with node i_0 and ending at node i_{N-1}, establish a lightpath between node i_j and node $i_{j+1} = (i_j + (N/2) - 1) \bmod N$ if node $(i_j + (N/2) - 1) \bmod N$ is not yet included in the logical ring topology. If it is already included, connect node i_j to node $i_{j+1} = (i_j + (N/2)) \bmod N$. Finally, connect node i_{N-1} to node i_0. When $(N/2)$ is even, each lightpath traverses at least $(N/2) - 1$ physical links. When $(N/2)$ is odd, the connection from i_{N-1} to i_0 traverses $(N/2) - 2$ links in the clockwise direction and $(N/2) + 2$ links in the counterclockwise direction. All other lightpaths traverse at least $(N/2) - 1$ physical links. We consider the two cases of $N = 6$ and $N \geq 8$ separately. When $N = 6$, three of the lightpaths can share a wavelength in the clockwise direction. However, only one lightpath can be established on each counterclockwise wavelength. Therefore, one clockwise and one counterclockwise wavelength can carry at most four of the six lightpaths. For $N \geq 8$, the lightpaths can only share wavelengths in the clockwise direction since each lightpath traverses more than $(N/2)$ links in the counterclockwise direction. Furthermore, at most two lightpaths can share each clockwise wavelength since any three lightpaths require a minimum of $2((N/2) - 1) + (N/2) - 2 > N$ physical links. The $\lceil (N/3) \rceil - 1$ clockwise wavelengths can support at most two lightpaths per wavelength for a total of $2\lceil (N/3) \rceil - 2$ lightpaths. The $\lceil (N/3) \rceil - 1$ counterclockwise wavelengths can each support at most one lightpath. Thus the total number of lightpaths supported is $3(\lceil (N/3) \rceil - 1)$, which is always less than N. Hence, this logical topology cannot be supported with less than $\lceil (N/3) \rceil$ wavelengths in each direction. ■

We have shown that by routing sets of three adjacent lightpaths on a single pair of wavelengths, all connected logical topologies can be supported on a network with $\lceil (N/3) \rceil$ wavelengths in each direction. It can also be shown that all unconnected logical topologies can also be supported with $\lceil (N/3) \rceil$ wavelengths in each direction. The proof, which has been omitted for brevity, uses arguments similar to those used in the proof of Theorem 3 to generalize the protected network results to include unconnected logical topologies.

2) Multiple Ports Per Node: A logical topology with P ports per node is a directed graph with nodes of in-degree and out-degree equal to P. If the logical topology is connected, then the directed graph contains a directed Euler trail ([20, Theorem 5.6]), where an Euler trail is a closed directed trail,[3] which contains all the edges of the graph. Therefore, the PN lightpath logical topology (PN edges of the graph) can be divided into $\lfloor (PN/2) \rfloor$ pairs of adjacent lightpaths plus one lightpath if PN is odd. For the unprotected network case, the PN lightpath topology can be divided into $\lfloor (PN/3) \rfloor$ sets of three adjacent lightpaths plus one set of (PN) mod 3 lightpaths. The routing strategies described in Section III-B1 for single port per node networks can thus be directly applied to multiple port

[3]A closed directed trail in a directed graph G is a finite sequence of vertices $v_0, v_1, \ldots, v_k$, such that (v_{i-1}, v_i) is an edge in G, all edges are distinct, and $v_0 = v_k$. Note that a trail can repeatedly visit the same node.

TABLE II
LOWER BOUND ON WAVELENGTH REQUIREMENTS

Topology Size	Number of Logical Topologies with Lower Bound Wavelength Requirement W_{LB}					
	$W_{LB}=1$	$W_{LB}=2$	$W_{LB}=3$	$W_{LB}=4$	$W_{LB}=5$	Total
$N=4$	2	4				6
$N=5$	2	22				24
$N=6$	2	82	36			120
$N=7$	2	240	478			720
$N=8$	2	616	3,846	576		5,040
$N=9$	2	1,466	24,012	14,840		40,320
$N=10$	2	3,334	126,570	218,574	14,440	362,880

per node networks. The following two theorems summarize the wavelength requirements for connected logical topologies in networks with N nodes and P ports.

Theorem 6: In a protected network, every connected logical topology can be implemented if $\lceil (PN/2) \rceil$ working traffic wavelengths are available.

Proof: Since the logical topology is connected, it contains an Euler trail. Find the Euler trail and divide the set of lightpaths into $\lfloor (PN/2) \rfloor$ sets of adjacent lightpath pairs plus one lightpath if PN is odd. Use Lemma 1 to route the lightpath pairs. Thus $\lceil (PN/2) \rceil$ wavelengths are sufficient. ∎

Theorem 7: In an unprotected network, every connected logical topology can be established if $\lceil (PN/3) \rceil$ wavelengths in each direction are available.

Proof: Since the logical topology is connected, it contains an Euler trail. Find the Euler trail and divide the set of lightpaths into $\lfloor (PN/3) \rfloor$ sets of three adjacent lightpaths plus one set of zero, one, or two lightpaths. Use Lemma 1 to route a pair of lightpaths from each set on one wavelength. Use the wavelength in the opposite direction to route the third lightpath in each set. Thus $\lceil (PN/3) \rceil$ pairs of wavelengths are sufficient to map the logical topology. ∎

Next consider unconnected logical topologies for N node P port networks. A disconnected logical topology G can be divided into a set of K connected components where each component G_i consists of N_i nodes and M_i lightpaths, where $0 \le i \le K-1$ and $\sum_i M_i = PN$. The set of lightpaths in the ith connected component forms an Euler trail on G_i. Thus each set of M_i lightpaths can be routed on $\lceil (M_i/2) \rceil$ wavelengths in the protected case and $\lceil (M_i/3) \rceil$ wavelengths in each direction in the unprotected case. Directly applying the arguments used in the $P=1$ case (Theorem 3) to the K connected components above yields the following theorems specifying the wavelength requirements for general (connected or unconnected) logical topologies. The proof details are omitted in the interest of brevity. ∎

Theorem 8: In a protected network, every logical topology (connected or unconnected) can be implemented if $\lceil (PN/2) \rceil + 1$ working traffic wavelengths are available.

Theorem 9: In an unprotected network, every logical topology (connected or unconnected) can be established if $\lceil (PN/3) \rceil$ wavelengths in each direction are available.

IV. Limited Logical Topology Networks

Thus far, we have considered networks that support all virtual topologies. However, by limiting the number of topologies that can be established, it may be possible to reduce network wavelength requirements.

To investigate the tradeoff between the fraction of topologies supported and the number of wavelengths required, we compute a lower bound on the wavelength requirements for each logical topology. Since the physical topology is a ring, any bisection of the physical topology corresponds to two physical links. Let m be the maximum number of lightpaths that cross any bisection of the physical topology graph. Each of these "crossing lightpaths," i.e., lightpaths that cross the bisection, must be mapped on one of the two physical links. Thus at most two crossing lightpaths can share a wavelength. Consequently, a minimum of $W_{\text{LB}} = (m/2)$ wavelengths are required to implement the logical topology, independent of the RWA strategy. We use this lower bound to calculate the minimum wavelength requirements for all connected single port per node logical topologies of size N. For logical topologies of size $N=4$ to 10, the minimum wavelength requirements (calculated using the lower bound above) and the number of topologies that require this minimum number of wavelengths are shown in Table II. These results show that many of the logical topologies require near the maximum number of $\lceil (N/2) \rceil$ wavelengths. Furthermore, a majority of the logical topologies require at least $\lfloor (N/2) \rfloor - 1$ wavelengths.

For larger values of N, we can use the following analytical result to show that as N increases, a significant number of logical topologies continue to require $\lceil (N/2) \rceil$ wavelengths.

Theorem 10: For N even, a minimum of $(N/2)$ wavelengths are required by at least $((N/2)!)^2$ connected logical topologies with one port per node.

Proof: We compute the number of logical topologies for which each of the N lightpaths crosses a bisection of the physical topology. Since at most two crossing lightpaths can fit on one wavelength, each of these logical topologies requires a minimum of $(N/2)$ wavelengths. Consider any bisection of the physical topology that evenly divides the N nodes. Start with any node i_0. There are $(N/2)$ ways to choose the next node i_1 such that lightpath (i_0, i_1) crosses the bisection. There are

$(N/2) - i$ ways to choose nodes i_{2j} and i_{2j+1}, for $j = 1$ to $(N/2) - 1$, such that all lightpaths cross the physical topology bisection. There are also $(N/2)$ ways to bisect the physical topology. Thus, there are $((N/2)!)^2$ logical topologies with a maximum number of crossings.

The $W_{\mathrm{LB}} = (N/2)$ entries in Table II correspond to $((N/2)!)^2$. This shows that the $\lceil (N/2) \rceil$ wavelength requirement is not artificially high to support rare pathological cases, but rather is necessary to support a significant number of the logical topologies. We note that although the *fraction* of logical topologies that require $\lceil (N/2) \rceil$ wavelengths becomes small as N becomes large, for all practical cases, a large *number* of logical topologies continue to require the maximum number of wavelengths.

V. Reducing Average Wavelength Requirements

Above, we developed a routing and wavelength assignment strategy that minimizes the number of wavelengths required to implement the worst case logical topology. Furthermore, we showed that this RWA strategy reduces the wavelength requirements approximately by one-half over shortest path routing for protected networks and by one-third for unprotected networks. It is also beneficial to reduce wavelength requirements on average. In some networks, for example, extra wavelengths may be used to provide all-optical connections and bypass the electronics altogether. An algorithm that minimizes the average wavelengths will provide more wavelengths for alternate services. In this section, we utilize the routing principles developed in the previous sections to design RWA algorithms that in addition to minimizing the *worst case* wavelength requirement also reduce *average* wavelength requirements. The algorithms are designed for N node P port networks. For simplicity, we assume that all logical topologies are connected.

A. Protected Networks

We describe two heuristic algorithms for routing and wavelength assignment. Recall that in a protected network, the direction of each working wavelength may be selected independently. The first algorithm is a direct implementation of the routing and wavelength assignment strategies suggested by the proofs in the previous sections. The second algorithm is a slight modification of the first that provides a substantial improvement.

Algorithm 1: Adjacent Routing:

1) Find an Euler trail in the desired logical topology. Order the lightpaths according to the Euler trail. In this way, consecutive lightpaths will be adjacent. Start with the first two lightpaths.
2) Start a new wavelength. Choose wavelength direction so that both lightpaths fit on one wavelength, i.e., use Lemma 1.
3) Try to map the next lightpath on the same wavelength. Continue routing lightpaths until you get to the first lightpath that does not fit on the current wavelength.
4) Take the current (unrouted) lightpath and the next lightpath and go to step 2).

The first algorithm, referred to as *adjacent routing*, fills wavelengths sequentially with adjacent lightpaths. When starting a new wavelength, the direction of lightpath routing is chosen to ensure a minimum of two lightpaths are accommodated. In the second algorithm, termed *first fit adjacent routing*, routing of a lightpath that does not fit on the current wavelength is attempted on wavelengths that already carry lightpaths. A new wavelength is begun only if the lightpath cannot be mapped on any previously used wavelengths.

Algorithm 2: First Fit Adjacent Routing:

1) Find an Euler trail in the desired logical topology. Order the lightpaths according to the Euler trail. In this way, consecutive lightpaths will be adjacent. Start with the first two lightpaths.
2) Start a new wavelength. Using Lemma 1, map the two lightpaths in the appropriate direction so that they fit on the single wavelength.
3) Try to map the next lightpath on the same wavelength. Continue routing lightpaths until you reach the first lightpath that does not fit on the current wavelength.
4) Try to map this lightpath on the previous wavelengths, starting with wavelength 1. Continue routing lightpaths on previous wavelengths until you reach the first lightpath that does not fit on any previously used wavelengths.
5) Take the current lightpath and next lightpath and go to step 2).

We use simulations to compare the average wavelength requirements of the two heuristic algorithms to shortest path routing using first fit wavelength assignment (SPR w/FF). We consider both DOES and DCRS shortest path routing, which differ only when N is even and $(N/2)$ is odd. Fig.6 shows the average wavelength requirements for $N = 6$ to $N = 10$ single port per node logical topologies. The wavelength requirements of all $(N-1)!$ permutations for each N node, $P = 1$ logical topology, are evaluated. Fig. 7 shows the average wavelength requirements for logical topologies with two ports per node. For $P = 2$, we generate 10 000 random logical topologies for $N = 6$ to $N = 10$. With $P = 1$, both adjacent routing and first fit adjacent routing require fewer wavelengths on average (and worst case) than SPR w/FF. First fit adjacent routing reduces

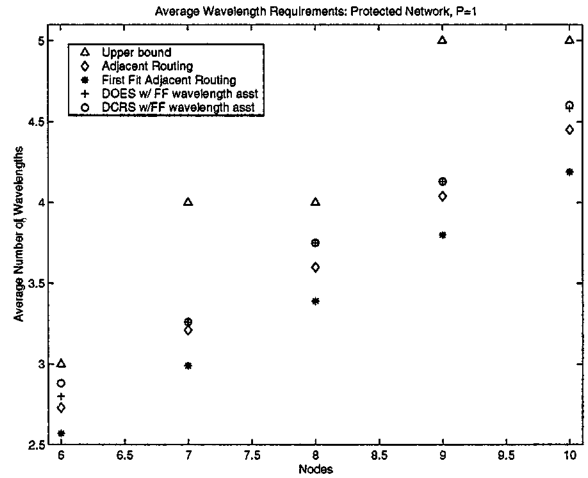

Fig. 6. Comparison of average wavelength requirements using heuristic joint routing and wavelength assignment algorithms to DOES and DCRS shortest path routing using first fit wavelength assignment. All logical topologies with $P = 1$ port per node are routed on a protected network.

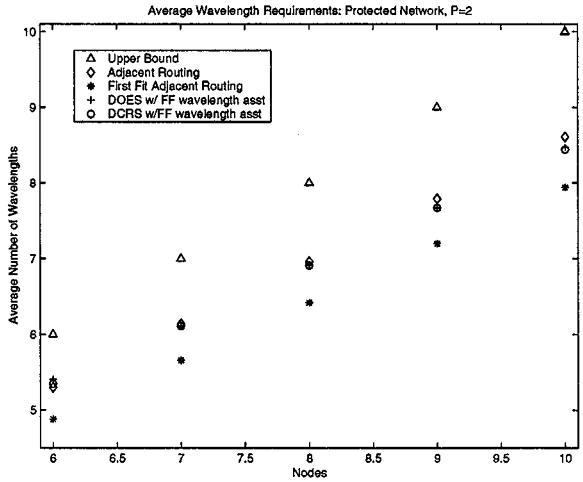

Fig. 7. Comparison of average wavelength requirements for heuristic joint routing and wavelength assignment algorithms to DOES and DCRS shortest path routing using first fit wavelength assignment. Ten thousand random logical topologies with $P = 2$ ports per node are mapped on a protected network.

average wavelength requirements between 8%–11% over SPR w/FF. For $P = 2$ and $N \geq 7$, the average wavelength requirements for adjacent routing, are slightly higher than SPR w/FF. Adjacent routing performs slightly worse than SPR w/FF, since adjacent routing prevents subsequent lightpaths from being mapped on any wavelengths on which a previous lightpath did not fit. First fit adjacent routing, which removes this restriction, provides a significant improvement over SPR w/FF, reducing average wavelength requirements by 6%–10%. Unlike SPR w/FF, adjacent routing and first fit adjacent routing ensure that a minimum of two lightpaths are routed on each wavelength. Although it may be possible to improve these heuristics slightly, these results illustrate that a significant improvement in both worst case and average wavelength requirements is possible by lightpath mapping algorithms that jointly assign routes and wavelengths using the adjacent lightpath routing principle outlined in the proof of Lemma 1. Also plotted in Figs. 6 and 7 are upper bounds on wavelength requirements derived from the results in Section III. The upper bounds correspond to the number of wavelengths required to map the worst case logical topology using our adaptive lightpath routing strategy. As expected, the wavelength requirements for the average case are much less than the worst case. This demonstrates the importance of improving the wavelength requirements for both the worst case and average case.

The average wavelength requirements of adjacent routing for single port per node logical topologies can also be approximated analytically. Note that K adjacent lightpaths are specified by $K+1$ nodes. Mapping K lightpaths on a single wavelength in one direction is equivalent to choosing the first node arbitrarily and then ensuring that the remaining K nodes are in the same order as the physical topology, i.e., in either clockwise (CW) or counterclockwise (CCW) order. Note that for connected logical topologies, if $N-1$ lightpaths fit on one wavelength, then all N lightpaths fit on the wavelength. Let X be the number of lightpaths that fit on a wavelength in either direction. Then, given that at least $K < N$ lightpaths remain to be mapped, we have (1), as shown at the bottom of the page. The number of ways to arrange K nodes in $N-1$ positions is simply $(N-1)(N-2)\ldots(N-K)$. The number of ways to arrange K nodes in R positions so that they are in order (e.g., clockwise) is given by the following recursive expression:

$$L[K,R] = \sum_{i=1}^{R} L[K-1, R-i] \qquad (2)$$

$$L[1,R] = R. \qquad (3)$$

Equation (2) is obtained by considering the position of the first of the K nodes. If the first node is in position i, then the remaining $K-1$ nodes must be placed in $R-i$ positions. The sum is taken over all R possible positions for the first node. The probability that more than K lightpaths fit on a wavelength (when K or more lightpaths need to mapped) is

$$P(X \geq K) = \frac{2L[K, N-1]}{(N-1)(N-2)\ldots(N-K)}. \qquad (4)$$

The average number of lightpaths needed to map a single port per node connected logical topology (ring) can also be computed recursively as follows:

$$A[K] = \begin{cases} 0, & \text{if } K = 0 \\ 1, & \text{if } 1 \leq K \leq 2 \\ 1 + \sum_{i=2}^{K} A[K-i]P(X=i), & \text{if } K > 2, \end{cases} \qquad (5)$$

where K is the number of lightpaths mapped and $P(X = i) = P(X \geq i) - P(X \geq i+1)$. Clearly, when only one or two lightpaths remain to be routed, they can fit on one wavelength. When $K > 2$, exactly i lightpaths fit on one wavelength with probability $P(X = i)$ and the remaining $K - i$ wavelengths use an average of $A[K - i]$ wavelengths. The recursion terminates when less than three lightpaths remain to be routed. The average wavelength requirements can be numerically computed using (5). The computations are approximate since the port and connectivity restrictions cause lightpaths to be correlated. For example, if the first wavelength carries only two lightpaths, this gives us information about the remaining lightpaths. The analytical approximations correspond well to the simulation results in Fig. 6 for $N = 6$ to $N = 10$.

It is found through numeric computation that as N increases, the average number of lightpaths that can fit on a wavelength approaches 2.43. Hence, the average number of wavelengths needed to map an N node ring logical topology approaches $N/2.43$.

$$P(X \geq K) = \frac{\#\text{ of ways in which } K \text{ nodes can be arranged in } N-1 \text{ positions in CW or CCW order}}{\#\text{ of ways to arrange } K \text{ nodes in } N-1 \text{ positions}} \qquad (1)$$

B. Unprotected Networks

In this section, we present two heuristic algorithms for routing and wavelength assignment on unprotected networks. Recall that in unprotected networks, the wavelengths must be allocated in clockwise/counterclockwise pairs. Both algorithms use Lemma 1 to map a pair of lightpaths in the appropriate direction to fit on one wavelength. The next lightpath that does not fit on the current wavelength is mapped on a wavelength in the opposite direction. This ensures that a minimum of three lightpaths are embedded on each pair of wavelengths. Routing of subsequent lightpaths is continued on the current wavelength until a lightpath that does not fit on the current wavelength is encountered. In adjacent routing, a new wavelength pair is begun. In first fit adjacent routing, the current lightpath and subsequent lightpaths are mapped on previously used wavelengths until the first lightpath that does not fit on any lightpath carrying wavelengths is found. At this point, a new wavelength pair is initiated.

Algorithm 1: Adjacent Routing:

1) Find an Euler trail in the desired logical topology. Order the lightpaths according to the Euler trail. In this way, consecutive lightpaths will be adjacent. Start with the first two lightpaths.
2) Start a new wavelength. Choose the direction of the wavelength so that both lightpaths fit on one wavelength.
3) Try to fit additional adjacent lightpaths on this wavelength in the same direction as the first two lightpaths. When the first lightpath that does not fit is encountered, route this lightpath on a new wavelength in the opposite direction.
4) Try to map the next lightpath on the same wavelength. Continue routing lightpaths until you get to the first lightpath that does not fit on the current wavelength.
5) Take the current (unrouted) lightpath and the next lightpath and go to step 2).

The second algorithm differs from Algorithm 1 at step 5), where mapping of a lightpath that does not fit on the current wavelength is attempted on previous wavelengths.

Algorithm 2: First Fit Adjacent Routing:

1) Find an Euler trail in the desired logical topology. Order the lightpaths according to the Euler trail. In this way, consecutive lightpaths will be adjacent. Start with the first two lightpaths.
2) Start a new wavelength. Using Lemma 1, map the two lightpaths in the appropriate direction so that they fit on a single wavelength.
3) Try to map the next lightpath on the same wavelength. Continue routing lightpaths until you reach the first lightpath that does not fit on the current wavelength.
4) Route this lightpath on a new wavelength in the opposite direction. Continue routing lightpaths until you get to the first lightpath that does not fit on the current wavelength.
5) Try to map this lightpath on the previous wavelengths, starting with wavelength 1. Continue routing lightpaths on previous wavelengths until you reach the first lightpath that does not fit on any previously used wavelengths.
6) Take current (unrouted) lightpath and next lightpath and go to step 2).

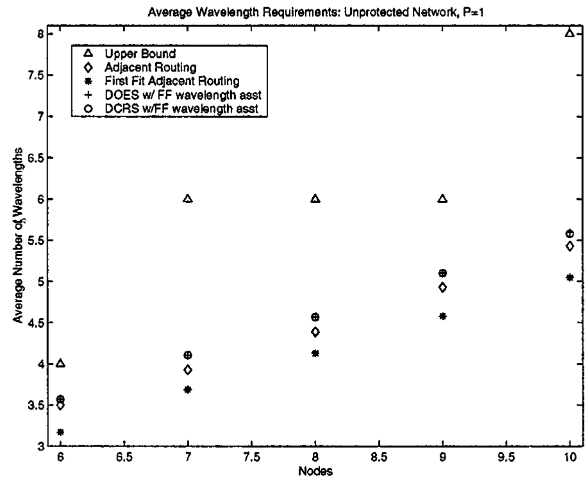

Fig. 8. Comparison of average wavelength requirements using heuristic joint routing and wavelength assignment algorithms to DOES and DCRS shortest path routing using first fit wavelength assignment. An unprotected network is assumed. The average wavelength requirements for all logical topologies with $P = 1$ port per node is determined.

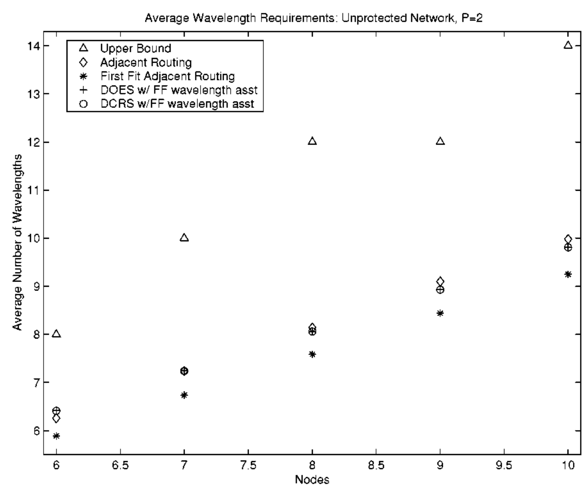

Fig. 9. Comparison of average wavelength requirements for heuristic joint routing and wavelength assignment algorithms to DOES and DCRS shortest path routing using first fit wavelength assignment. An unprotected network is assumed. The average wavelength requirements for 10 000 random logical topologies with $P = 2$ ports per node is determined.

Average wavelength requirements for logical topologies with $P = 1$ and $P = 2$ are shown in Figs. 8 and 9, respectively. For single port per node logical topologies, both adjacent routing and first fit adjacent routing produce significant improvements over SPR w/FF. First fit adjacent routing provides a 9%–11% reduction in average wavelength requirements. With $P = 2$, first fit adjacent routing reduces average wavelength requirements by 5%–8%.

Figs. 6–9 illustrate the improvement in average wavelength requirements achieved from using our adaptive algorithms as compared to fixed routing schemes. These figures also include an upper bound showing the wavelength requirements for worst case logical topologies. The average wavelength requirements are significantly lower than worst case wavelength requirements. These results demonstrate that it is important to reduce both average and worst case wavelength requirements.

VI. Extensions to k-Connected Physical Topologies

Above, we have shown for bidirectional ring physical topologies that a minimum of two lightpaths can be mapped on any wavelength if the wavelength direction is selected appropriately. We used this observation to determine minimum wavelength requirements for supporting worst case logical topologies and developed a routing and wavelength assignment algorithm that achieves this minimum wavelength requirement. We also developed heuristic routing and wavelength assignment algorithms that illustrate how this adjacent lightpath routing principle can be used to reduce average wavelength requirements. In this section, we show how the adjacent lightpath routing principle might be extended for use on two-connected and three-connected physical topologies. We assume that each physical link consists of two fibers, one propagating in each of two directions. Note that a set of k adjacent lightpaths is specified by an ordered set of $k+1$ nodes. A set of k adjacent lightpaths is said to use a single wavelength if there is an edge disjoint trail that passes through all $k+1$ distinct nodes in order.

A bidirectional ring is a specific case of a two-connected physical topology. However, not all two-connected graphs are rings or even contain a cycle that traverses all N nodes. We show that Lemma 1 can be generalized to two-connected physical topologies. Specifically, we show that any pair of adjacent lightpaths can be mapped on one wavelength of a two-connected physical topology by properly choosing the lightpath routes.

Theorem 11: On a two-connected graph, for any pair of adjacent lightpaths (a, b) and (b, c), there exists an edge disjoint trail from a to b to c.

Proof: Consider two adjacent lightpaths (a, b) and (b, c). Add a node t to the graph and connect it to nodes a and c. The new graph is still two-connected. Thus by Menger's theorem [20], there are two edge disjoint paths from node b (the source) to node t (the sink). Furthermore, one of these two paths passes through a and the other through c. Thus there exist paths from a to b and from b to c that are edge disjoint. ■

We can also show the following result for three-connected physical topologies.

Theorem 12: On a three-connected graph, any three adjacent lightpaths (a, b), (b, c), and (c, d) can be routed on a single wavelength.

The proof, given in Appendix II, is by construction; thus it also provides a method for routing the three lightpaths. It is fairly straightforward to generate sets of three adjacent lightpaths that *cannot* be mapped on a two-connected physical topology for any routing strategy. Similarly, one can generate sets of four adjacent lightpaths that *cannot* be mapped on a three-connected physical topology. Thus, the mappings described are maximal in this sense. We conjecture that these results can be generalized to k-connected physical topologies for all k, i.e., we expect that any k adjacent lightpaths can be mapped on one wavelength of a k-connected physical topology. The result for k even has been established in [21]. For k odd, the problem is yet unsolved [22].

Theorems 11 and 12 can be used to devise heuristic routing and wavelength assignment strategies for two-connected and three-connected physical topologies. In fact, slight variations of the adjacent routing and first fit adjacent routing algorithms can be implemented. Here instead of selecting a wavelength direction, routing the lightpaths on one wavelength corresponds to finding a directed trail for the lightpaths through the physical topology. For two-connected physical topologies, the lightpath routes are selected to ensure that two lightpaths fit on one wavelength. One can easily find pairs of adjacent lightpaths that *cannot* be mapped using shortest path routing on one wavelength of a two-connected physical topology. Thus it is also possible to generate logical topologies in which shortest path routing using first fit wavelength assignment would not be able to map a minimum of two lightpaths on each wavelength of a two-connected physical topology. Similarly, for three-connected physical topologies, the lightpath routes are chosen to ensure that three lightpaths fit on one wavelength. One can easily generate sets of three adjacent lightpaths that if routed using shortest path routing will require more than one wavelength on a three-connected physical topology. Thus, SPR w/FF on a three-connected physical topology will not map a minimum of three lightpaths on each wavelength. As the physical topology becomes more connected, and the number of lightpaths (equivalently, ports) increases, we expect shortest path routing to use most wavelengths fairly efficiently since most lightpaths will use a small number of physical links. However, algorithms such as first fit adjacent routing can improve wavelength utilization, since they ensure that a minimum of two lightpaths are routed on *every* wavelength on a two-connected physical topology and that a minimum of three lightpaths are mapped on *every* wavelength on a three-connected physical topology.

VII. Conclusion

The minimum number of wavelengths required W_{req} to implement all virtual topologies on an N node P port network has been determined. For connected logical topologies, $W_{\mathrm{req}} = \lceil (PN/2) \rceil$ working traffic wavelengths are required on a protected network and $W_{\mathrm{req}} = \lceil (PN/3) \rceil$ wavelengths in each direction are required on an unprotected network. A significant fraction of logical topologies require nearly W_{req} wavelengths, thus reducing the wavelength requirement by designing the network to support only a limited number of logical topologies is not a worthwhile proposition. This also indicates that our focus on supporting all possible logical topologies does not result in a substantial overprovisioning of resources.

Adaptive lightpath routing strategies that can embed all logical topologies within the minimum W_{req} wavelength requirement were developed. These adaptive routing schemes required far fewer wavelengths than RWA schemes based on shortest path routing.

Joint routing and wavelength assignment algorithms that improve average wavelength requirements were also developed. These algorithms reduced average wavelength requirements *in addition* to minimizing worst case logical topology wavelength requirements. If networks are provisioned with W_{req} wavelengths, embedding a topology with less than W_{req} wavelengths allows the extra wavelengths to be used for alternative services. The joint routing and wavelength assignment algorithms developed use fewer wavelengths than shortest path routing-based RWA both for average and worst case topologies.

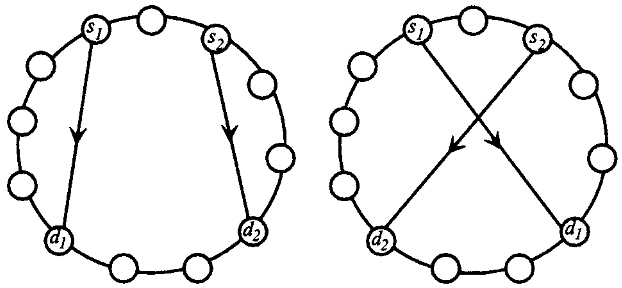

(a) (b)

Fig. 10. (a) Two parallel links. (b) Two intersecting links.

Although we primarily focused on the bidirectional ring physical topology, methods for extending the adjacent lightpath routing principles to general two-connected and three-connected physical topologies were also presented. Determination of the minimum wavelength requirements for k-connected physical topologies topologies is an area for future work. The joint routing and wavelength assignment algorithms developed for bidirectional ring physical topologies were adapted to two-connected and three-connected physical topologies. Performance analysis of the these algorithms as well as development of new heuristic joint RWA strategies for k-connected physical topologies are areas for future work. Investigation of joint routing and wavelength assignment on other physical topologies is another area of interest.

APPENDIX I
PROOF OF LEMMA 2

Denote the three directed circuits of odd sizes M_1, M_2, and M_3 as A, B, and C. We show below that these three circuits can be routed using a total of $\lceil (M_1+M_2+M_3)/(2) \rceil$ wavelengths.

We begin with some preliminary definitions. A *link* in the directed circuit corresponds to a directed logical connection between a source node s and destination node d. We say that two logical links $(s1, d1)$ and $(s2, d2)$ are parallel if nodes $s2$ and $d2$ lie on the same side of the bisection of the physical topology formed by link $(s1, d1)$, as illustrated in Fig. 10(a). We say that two links *intersect* if they are not parallel, as shown in Fig. 10(b). Two parallel links are said to traverse the *same direction* if traversing the ring physical topology to go from $d1$ to $s2$ requires going through node $s1$ in one direction and $d2$ in the other. Note that links $(s1, d1)$ and $(s2, d2)$ in Fig. 10(a) traverse the same direction.

The proof of Lemma 2 uses the following theorem and corollary.

Theorem 13: Given a directed circuit of odd size M, any link l can intersect at most $M-1$ links in the circuit.

Proof: Any link l that cuts a circuit of size M will divide the circuit such that k of the nodes are to the left of the link and $M-k$ of the nodes are to the right of the link. The number of links from the circuit that intersect link l is thus at most $2\min(k, M-k) = 2\lfloor (M/2) \rfloor = M-1$, since M is odd. ■

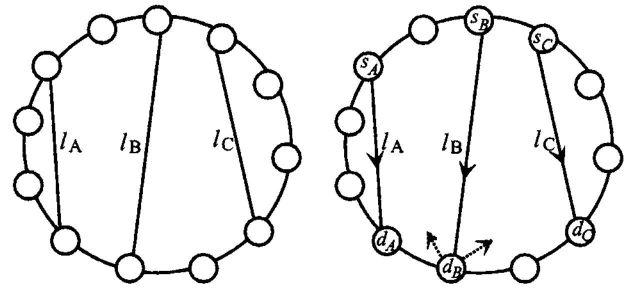

(a) (b)

Fig. 11. Scenario 1): three links from three circuits are parallel.

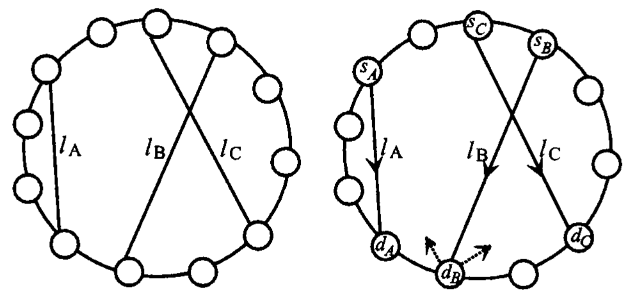

(a) (b)

Fig. 12. Scenario 2): a link from circuit A is parallel to links from circuits B and C, but the links from circuits B and C intersect.

Corollary to Theorem 13: Given three directed circuits A, B, and C of odd size, there exists a link in A denoted l_A that is parallel to a link in B, denoted l_B, and also to a link in C, denoted l_C.

Now we can prove Lemma 2. By applying Lemma 1, we know that each circuit requires at most $\lceil (M_i/2) \rceil$ wavelengths. Thus if we can show that there exist two links, one each from two odd circuits that can share a wavelength, the number of wavelengths needed will be $(M_i-1)/(2)+(M_j-1)/(2)+1+\lceil (M_k/2) \rceil = \lceil (M_i+M_j+M_k)/(2) \rceil$ as required.

Let $a \parallel b$ denote that link a is parallel to link b. From the corollary, there are two possible scenarios: 1) all three links are parallel or 2) links l_B and l_C intersect but are both parallel to link l_A. In scenario 1), shown in Fig. 11(a), links $l_A \parallel l_B$, $l_A \parallel l_C$, and $l_B \parallel l_C$. If no two of the three links can share a wavelength, this implies that the three links must go in the same direction, for if they do not, the two opposite direction links can share a wavelength. Now consider scenario 2), shown in Fig. 12(a), where $l_A \parallel l_B$ and $l_A \parallel l_C$ but l_B is not parallel to l_C. As shown in Fig. 12(b), links l_B and l_C must traverse the same direction as l_A; otherwise they can share a wavelength with l_A. Thus suppose that l_A is in the same direction as l_B and l_C, as shown in Figs. 11(b) for scenario 1) and 12(b) for scenario 2). Consider the link following link l_B in circuit B, i.e., the link with source node equal to the destination node of l_B. If the destination of this link is to the right of link l_A, then this link can share a wavelength with link l_A. Otherwise the destination of this link must be to the left of link l_C, in which case this link can share a wavelength with link l_C. Therefore, there always

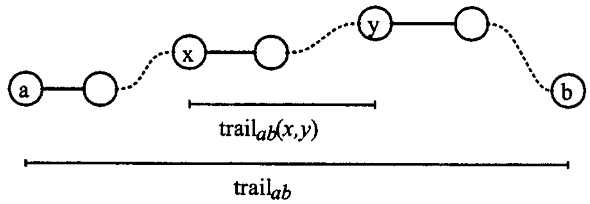

Fig. 13. Definition of a trail from node a to b, trail$_{ab}$, and its segment from node x to node y, trail$_{ab}(x, y)$.

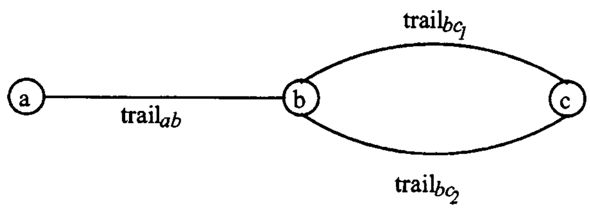

Fig. 14. Construction of three edge disjoint trails using Menger's theorem and the max-flow min-cut theorem.

exists at least two links, one each from two different circuits of A, B, and C, that can share a wavelength.

Appendix II
Proof of Theorem 12

In this section, we prove Theorem 12, which shows that on a three-connected physical topology, any three lightpaths can be mapped on a single wavelength. The proof is by construction; consequently, it also illustrates a method for selecting the routes of the three lightpaths.

We begin with some preliminary definitions. Let trail$_{sd}$ be a trail from nodes s to node d without repeated edges. Although a trail may contain loops, we assume for simplicity that trails do not contain loops. The proof can be easily extended to cases where trails contain loops. Define trail$_{sd}(x, y)$, for all nodes x, y on trail$_{sd}$, as a *segment* of trail$_{sd}$ that starts at node x, follows trail$_{sd}$, and terminates on node y. An example of a trail and its segment is shown in Fig. 13. Note that trail$_{ab}(a, b) =$ trail$_{ab}$ and that trail$_{ab}(y, x)$ has the same edges as trail$_{ab}(x, y)$ except it is traversed in the opposite direction. Two trails are *disjoint* when they do not share an edge. Two trails *touch* at node z if node z belongs to both trails. Two trails *overlap* at edge (x, y) if (x, y) is on both trails. Note that if two trails overlap at (x, y), then they touch at x and y.

To prove Theorem 12, given a three-connected graph, we must establish a trail, trail$_{ad}$, through any four distinct nodes $\{a, b, c, d\}$, such that the trail passes through all four nodes, terminates on nodes a and d, and such that the trail segments trail$_{ad}(a, b)$, trail$_{ad}(b, c)$, and trail$_{ad}(c, d)$ are all edge disjoint.

Proof: Let node b be the source of a flow of capacity 3, node a be a sink of capacity 1, and node c be a sink of capacity 2. Using Menger's theorem and max-flow min-cut arguments [20], we can find three edge disjoint trails: trail$_{ab}$, trail$_{bc_1}$, and trail$_{bc_2}$, where nodes a and b are terminals of trail$_{ab}$ and nodes b and c are terminals of trail$_{bc_1}$ and trail$_{bc_2}$, as shown in Fig. 14. Note that the three trails are only guaranteed to be *edge* disjoint.

Using Menger's theorem again, find three edge disjoint trails from node d to node c. Name these trails trail$_{dc_1}$, trail$_{dc_2}$, and

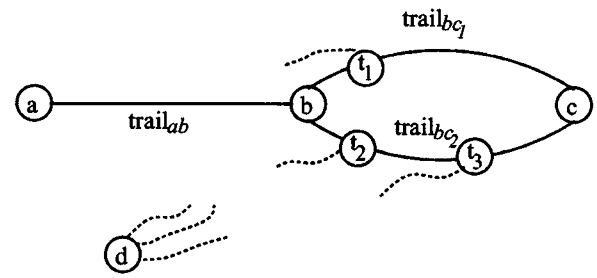

Fig. 15. The three trails from node d to node c will eventually touch trail$_{bc_1}$ ∪ trail$_{bc_2}$. The locations where these trails touch are denoted t_i for $i = 1, 2, 3$. Note that the first place the trails touch may be at node c.

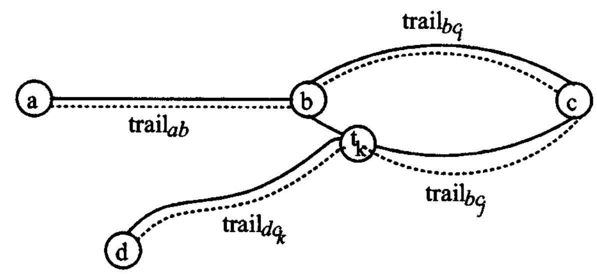

Fig. 16. Construction of the desired trail when at least one of the three trails from node d to node c does not overlap with trail$_{ab}$ before touching trail$_{bc_1}$ ∪ trail$_{bc_2}$.

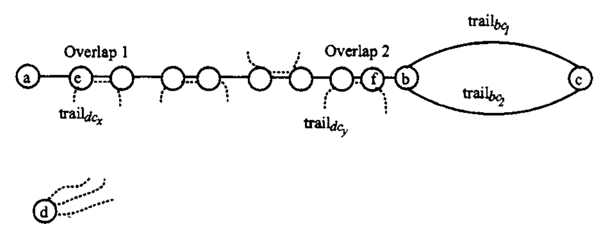

Fig. 17. Only consider overlaps that occur before trail$_{dc_i}$, $i = 1, 2, 3$, touches trail$_{bc_1}$ ∪ trail$_{bc_2}$.

trail$_{dc_3}$. Note that these three trails may overlap with any of the three trails previously defined.

Traverse trail$_{dc_i}$ $(i = 1, 2, 3)$ from node d toward node c. Stop whenever trail$_{dc_i}$ touches either trail$_{bc_1}$ or trail$_{bc_2}$ for the very first time. Let t_i be the point where the trails touch. Note that trail$_{dc_i}(d, t_i)$ is edge disjoint with trail$_{bc_1}$∪trail$_{bc_2}$, as illustrated in Fig. 15.

If one of the three d to c trails does not overlap with trail$_{ab}$, then we have our desired trail$_{ad}$. Assume that trail$_{dc_k}$ is the trail that has this property and that it first touches trail$_{bc_j}$ at node t_k. Let trail$_{bc_i}$ denote the other node b to node c trail. Then the desired trail is trail$_{ad} =$ trail$_{ab}$ + trail$_{bc_i}$ + trail$_{bc_j}(c, t_k)$ + trail$_{dc_k}(t_k, d)$, as illustrated in Fig. 16.

The only case left to consider is when all three d to c trails overlap with trail$_{ab}$ before they first touch trail$_{bc_1}$ ∪ trail$_{bc_2}$. There may also be repeated overlaps, and the ordering of the overlap locations may intertwine. Nevertheless, the number of overlaps is finite. There is an overlap that is closest to node a; call this Overlap1. Label the node closest to node a on Overlap1 as node e. There is also an overlap that is closest to node b; call this Overlap2. Label the node closest to node b on Overlap2 as node f. This is illustrated in Fig. 17.

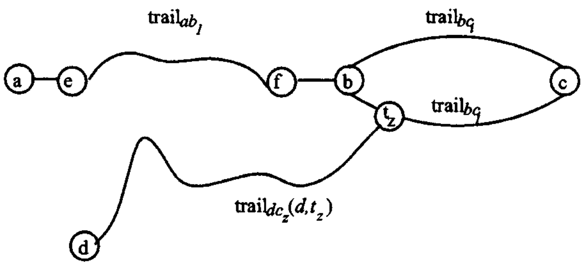

Fig. 18. Construction of the desired trail when all three node d to node c trails overlap with trail_{ab} before they touch $\text{trail}_{bc_1} \cup \text{trail}_{bc_2}$.

Let trail_{dc_x} be the trail that creates Overlap1 and trail_{dc_y} be the trail that creates Overlap2. The last d to c trail is then trail_{dc_z}. Note that trail_{dc_x} may be the same trail as trail_{dc_y}. Assume that trail_{dc_z} first touches $\text{trail}_{bc_1} \cup \text{trail}_{bc_2}$ at node t_z.

Claim: There exists a trail from node a to node b that is edge disjoint from $\text{trail}_{dc_z}(d, t_z)$.

Proof of Claim: In the case that trail_{dc_y} is different from trail_{dc_x}, consider $\text{trail}_{ab}(a, e) + \text{trail}_{dc_x}(e, d) + \text{trail}_{dc_y}(d, f) + \text{trail}_{ab}(f, b)$. This trail is edge disjoint with $\text{trail}_{dc_z}(d, t_z)$ because on trail_{ab}, there is no overlap before node e and after node f, and since trail_{dc_x} and trail_{dc_y} are edge disjoint from trail_{dc_z}.

In the case that trail_{dc_y} is the same as trail_{dc_x}, consider $\text{trail}_{ab}(a, e) + \text{trail}_{dc_y}(e, f) + \text{trail}_{ab}(f, b)$. This trail is edge disjoint with $\text{trail}_{dc_z}(d, t_z)$ because on trail_{ab}, there is no overlap before node e and after node f, and since $\text{trail}_{dc_y} = \text{trail}_{dc_x}$ is edge disjoint from trail_{dc_z}. ■

In whichever case, the above claim is true. Label this new node a to node b as trail_{ab_1}. Note that trail_{ab_1} is also edge disjoint with both of the node b to node c trails.

The proof can now be finished. The desired trail, shown in Fig. 18, is $\text{trail}_{ab_1} + \text{trail}_{bc_i} + \text{trail}_{bc_j}(c, t_z) + \text{trail}_{dc_z}(t_z, d)$, where trail_{dc_z} first touches $\text{trail}_{bc_1} \cup \text{trail}_{bc_2}$ at node t_z of trail_{bc_j}.

REFERENCES

[1] B. Mukherjee, "WDM-based local lightwave networks Part I: Single-hop systems," *IEEE Network*, pp. 12–27, May 1992.

[2] A. S. Acampora, "A multichannel multihop local lightwave network," in *Proc. GLOBECOM*, vol. 3, 1987, pp. 1459–1467.

[3] J. P. Labourdette and A. S. Acampora, "Logically rearrangeable multihop lightwave networks," *IEEE Trans. Commun.*, vol. 39, pp. 1223–1230, Aug. 1991.

[4] S. Banerjee and B. Mukherjee, "The photonic ring: Algorithms for optimized node arrangements," *Fiber Integ. Opt.*, vol. 12, no. 2, pp. 133–171, 1993.

[5] R. Ramaswami and K. N. Sivarajan, "Design of logical topologies for wavelength-routed optical networks," *IEEE J. Select. Areas Commun.*, vol. 14, pp. 840–851, June 1996.

[6] G. N. Rouskas and M. H. Ammar, "Dynamic reconfiguration in multihop WDM networks," *J. High Speed Networks*, vol. 4, no. 3, pp. 221–238, June 1995.

[7] A. Narula-Tam and E. Modiano, "Dynamic load balancing in WDM packet networks with and without wavelength constraints," *IEEE J. Select. Areas Commun.*, vol. 18, pp. 1972–1979, Oct. 2000.

[8] O. Gerstel, G. Sasaki, S. Kutten, and R. Ramaswami, "Worst-case analysis of dynamic wavelength allocation in optical networks," *IEEE Trans. Networking*, vol. 7, pp. 833–845, Dec. 1999.

[9] A. Tucker, "Coloring a family of circular arcs," *Siam J. Appl. Math.*, vol. 29, no. 3, pp. 493–502, Nov. 1975.

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

[11] A. Birman and A. Kershenbaum, "Routing and wavelength assignment methods in single-hop all-optical networks with blocking," in *IEEE Infocom*, vol. 1, 1995, pp. 431–438.

[12] O. Gerstel and S. Kutten, "Dynamic wavelength allocation in all-optical ring networks," in *Proc. IEEE ICC*, vol. 1, 1997, pp. 432–436.

[13] L. Wuttisittikukij, S. Leelanunnukul, S. Arreewanit, and P. Prapinmongkolkarn, "Routing and wavelength allocation in mult-wavelength all-optical ring networks," in *Proc. IEEE ICC*, 1999, pp. 2018–2022.

[14] C. Qiao and Y. Mei, "Off-line permutation embedding and scheduling in multiplexed optical networks with regular topologies," *IEEE Trans. Networking*, vol. 7, pp. 241–250, Apr. 1999.

[15] S. Baroni and P. Bayvel, "Wavelength requirements in arbitrarily connected wavelength-routed optical networks," *J. Lightwave Technol.*, vol. 15, pp. 242–251, Feb. 1997.

[16] R. Ramaswami and K. N. Sivarajan, "Routing and wavelength assignment in all-optical networks," *IEEE Trans. Networking*, vol. 3, pp. 489–500, Oct. 1995.

[17] A. Mokhtar and M. Azizoglu, "Adaptive wavelength routing in all-optical networks," *IEEE Trans. Networking*, vol. 6, pp. 197–206, Apr. 1998.

[18] S. Banerjee and C. Chen, "Design of wavelength-routed optical networks for circuit switched traffic," in *IEEE Infocom*, vol. 1, 1996, pp. 306–310.

[19] R. Ramaswami and K. N. Sivarajan, *Optical Networks: A Practical Perspective*. San Francisco, CA: Morgan Kaufmann, 1998.

[20] M. N. S. Swamy and K. Thulasiraman, *Graphs, Networks, and Algorithms*. New York: Wiley, 1981.

[21] H. Okamura, "Paths in k-edge connected graphs," *J. Combin. Theory Series B*, vol. 45, no. 3, pp. 345–355, Dec. 1988.

[22] A. Narula-Tam and T. Chow, "A linear version of k-ordered graphs," *Discrete Mathematics (Special Issue)*, to be published.

Aradhana Narula-Tam (M'97) received the B.S.E. degree from the University of Pennsylvania, Philadelphia, in 1990 and the S.M. and Ph.D. degrees from the Massachusetts Institute of Technology (MIT), Cambridge, in 1992 and 1997, respectively, all in electrical engineering.

From 1997 to 1998, she was in the Research and Advanced Development Group, Motorola Information Systems Group, Mansfield, MA. Since 1998, she has been with the Advanced Networking Group at MIT Lincoln Laboratory. Her research interests include communication theory and optical networks.

Philip J. Lin (M'95) received the B.S. degree in engineering from the California Institute of Technology, Pasadena, in 1989 and the S.M. and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1991 and 1996, respectively.

He is a Staff Research Engineer at the Tellabs Research Center and a Visiting Scientist at the Massachusetts Institute of Technology. His Ph.D. thesis work was on all-optical backbone networks. His current research areas include optical network architecture, optical system design, and optical access networks.

Eytan Modiano (S'90–M'93–SM'00) received the B.S. degree in electrical engineering and computer science from the University of Connecticut at Storrs in 1986 and the M.S. and Ph.D. degrees in electrical engineering from the University of Maryland, College Park, in 1989 and 1992, respectively.

He was a Naval Research Laboratory Fellow between 1987 and 1992 and a National Research Council Postdoctoral Fellow during 1992–1993, while he was conducting research on security and performance issues in distributed network protocols. Between 1993 and 1999, he was with the Communications Division at the Massachusetts Institute of Technology (MIT) Lincoln Laboratory, where he designed communication protocols for satellite, wireless, and optical networks and was the Project Leader for the laboratory's Next Generation Internet project. Since 1999, he has been a Member of the Faculty of the Aeronautics and Astronautics Department and the Laboratory for Information and Decision Systems at MIT, where he conducts research on communication networks and protocols with emphasis on satellite and hybrid networks and high-speed networks.