On Optimal Traffic Grooming in WDM Rings

Rudra Dutta, *Student Member, IEEE,* and George N. Rouskas, *Senior Member, IEEE*

***Abstract*—We consider the problem of designing a virtual topology to minimize electronic routing, that is, grooming traffic, in wavelength routed optical rings. The full virtual topology design problem is NP-hard even in the restricted case where the physical topology is a ring, and various heuristics have been proposed in the literature for obtaining good solutions, usually for different classes of problem instances. We present a new framework which can be used to evaluate the performance of heuristics and which requires significantly less computation than evaluating the optimal solution. This framework is based on a general formulation of the virtual topology problem, and it consists of a sequence of bounds, both upper and lower, in which each successive bound is at least as strong as the previous one. The successive bounds take larger amounts of computation to evaluate, and the number of bounds to be evaluated for a given problem instance is only limited by the computational power available. The bounds are based on decomposing the ring into sets of nodes arranged in a path and adopting the locally optimal topology within each set. While we only consider the objective of minimizing electronic routing in this paper, our approach to obtaining the sequence of bounds can be applied to many virtual topology problems on rings. The upper bounds we obtain also provide a useful series of heuristic solutions.**

***Index Terms*—Bounds, electronic grooming, heuristics, lightpath, optical ring, SONET, traffic grooming, wavelength routing.**

I. Introduction

In recent years, wavelength routed optical networks have been seen to be an attractive architecture for the next generation of backbone networks. This is due to the high bandwidth in fibers with *wavelength division multiplexing* (WDM) and the ability to tradeoff some of the bandwidth for elimination of electrooptic processing delays using *wavelength routing* [5]. Of late, two concerns have clearly emerged in the treatment of optical ring networks using WDM and wavelength routing [7], [8]. First, it has been recognized that the cost of network components, specifically electrooptic equipment and SONET add–drop multiplexers (ADMs), is a more meaningful metric for the network or topology than the number of wavelengths. Second, earlier studies of wavelength routed networks had assumed that individual traffic demands are comparable to the wavelength bandwidth [1], [4], [5], [10], [12]. However, it is now evident that the individual traffic streams that wavelength routed networks will carry are likely to have small bandwidth requirements compared even to the bandwidth available in a single wavelength of a WDM system. This assumption is further supported by the fact that the number of different traffic components in a network is likely to be much larger than the number of wavelengths available. These two issues give rise to the concept of *traffic grooming*, which refers to techniques used to combine lower speed traffic components onto available wavelengths in order to meet network design goals such as cost minimization.

The problem of designing logical topologies for rings that minimize cost as measured by the amount of electrooptic equipment has recently received much attention. In [11], the problem of wavelength assignment to a given set of lightpaths is considered, and the focus is on how limited wavelength conversion affects the capability of the network. Several different virtual topologies for ring networks are discussed in [8] in light of the twin issues of traffic grooming and network cost. The study in [7] complements [8] by addressing only the wavelength assignment issue, with the goal of minimizing the number of SONET ADMs. The concept of lightpath splitting in designing a virtual topology is discussed and heuristics for wavelength assignment are developed based on this concept. In [14], the problem of grooming traffic is considered both for unidirectional and bidirectional rings, with a primary goal of either minimizing the number of SONET ADMs or the number of wavelengths. The strategy presented is to first construct circles from the given traffic components and then to groom these circles. Algorithms for exact solutions for uniform traffic and heuristics for the NP-hard nonuniform traffic case are presented. A similar approach of constructing circles first is taken in [15], but now the resultant circles are scheduled in a sequence of virtual topologies. Heuristic algorithms to minimize network cost by grooming are presented in [3], for special traffic patterns such as uniform, certain cases of cross-traffic, and hub. In [2], the concept of a t-allowable traffic pattern is discussed, in which the traffic for each node-pair is constrained to be at most t. A bidirectional ring, but with symmetric traffic, is considered for dynamic traffic. A heuristic algorithm based on a bipartite matching formulation of the problem is presented, and blocking characteristics of the result are discussed. In [9], a problem similar to that in [7] is considered, but the problem of routing the lightpaths is considered as well as the wavelength assignment problem. An Euler-trail decomposition of the ring network is presented, and a heuristic which uses one of the heuristics of [7] and performs rerouting of lightpaths as well, based on the Euler-trail decomposition, is presented.

The problem of logical topology design is NP-hard even for a ring topology [5], [14], and obtaining an exact solution requires significant amount of computation even for modest sized rings. Heuristic approaches are needed for practical purposes and have been reported in [3], [7], [9] [14], [15]. To evaluate the performance of such heuristics when the optimal is not available, achievability bounds are useful.

In this paper, we present a new framework for computing bounds for the problem of traffic grooming in ring topologies.

Manuscript received January 25, 2001; revised July 19, 2001. This work was supported by the National Science Foundation under Grant 9701113.

The authors are with the Department of Computer Science, North Carolina State University, Raleigh, NC 27695-7534 USA (e-mail: rdutta@csc.ncsu.edu; rouskas@csc.ncsu.edu).

Publisher Item Identifier S 0733-8716(02)00147-6.

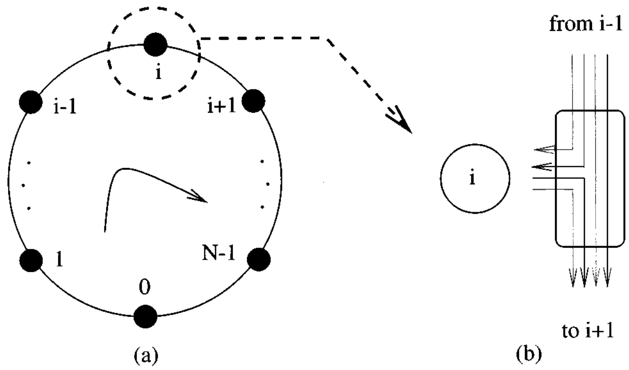

Fig. 1. (a) An N-node unidirectional ring and (b) detail of a node with a WADM.

The framework is based on the idea of decomposing the ring into path segments consisting of successively larger number of nodes. The path segments are solved independently, and the individual results are appropriately combined to obtain bounds on the optimal value of the objective function. In this manner, we obtain a series of bounds, both lower and upper, in which each bound is at least as strong as the previous one. The first few bounds in the sequence require trivial computational effort. Although subsequent bounds take successively larger computational efforts to determine, even several later bounds require significantly less computational effort than the full solution does, depending on the problem instance. We show that this result is due to the fact that solving a path segment exactly takes an amount of time which is orders of magnitude less than that required for solving a ring of the same number of nodes. The problem we consider is very general, as we do not impose any constraints on the traffic patterns. The upper bounds we derive are based on actual feasible topologies, so our algorithm for obtaining the upper bounds is a heuristic for the problem of traffic grooming. Finally, although we illustrate our approach using a specific formulation of the problem, it is straightforward to modify it to apply to a wide range of problem variants with different objective functions and/or constraints.

The rest of the paper is organized as follows. The problem of designing a virtual topology that minimizes electronic routing in a ring network is formulated as an integer linear program (ILP) in Section II. In Section III we describe the decomposition of the ring network into path segments, and we show how to obtain an optimal solution to the path segments from a simplified version of the ILP. Next, in Section IV we describe how to combine efficiently the solutions to individual segments to yield a strong sequence of lower and upper bounds for the original problem. Section V contains numerical results and Section VI concludes the paper.

II. Problem Formulation

We consider a unidirectional ring $\mathcal{R}$ with N nodes,[1] numbered from 0 to $N-1$, as shown in Fig. 1(a). The fiber link between each pair of nodes can support W wavelengths and carries traffic in the clockwise direction; in other words, data flows from a node i to the next node $i \oplus 1$ on the ring, where $\oplus$ denotes addition modulo-N. (Similarly we use $\ominus$ to denote subtraction modulo-N.) The links of $\mathcal{R}$ are numbered from 0 to $N-1$, such that the link from node i to node $i \oplus 1$ is numbered i. Each node in the ring is equipped with a wavelength add–drop multiplexer (WADM) [see Fig. 1(b)]. A WADM can perform three functions. It can optically switch some wavelengths from the incoming link of a node directly to its outgoing link. It can also terminate (drop) some wavelengths from the incoming link to the node; the data carried by the dropped wavelengths is converted to electronic form and undergoes buffering, processing, and possibly, electronic switching at the node. Finally, the WADM can also add some wavelengths to the outgoing link; these wavelengths may carry traffic originating at the node, or they may carry traffic originating at previous nodes in the ring and electronically switched at this node. We assume that estimates of the node-to-node traffic are available, requiring the design of a *virtual* or *logical* topology consisting of a set of static lightpaths. In this paper we do not consider the dynamic scenario in which requests for lightpaths or traffic components are received continuously during operation.

The traffic demands between pairs of nodes in the ring are given in the traffic matrix $T = [t^{(sd)}]$. We assume that the network supports traffic streams at rates that are a multiple of some basic rate (e.g., OC-3). We let C denote the capacity of each wavelength expressed in units of this basic rate. Thus, C denotes the maximum number of traffic units that can be multiplexed on a WDM channel (wavelength). For example, if each wavelength runs at OC-48 rates and the basic rate is OC-3, then $C = 16$. Each quantity $t^{(sd)} \in \{0, 1, 2, \ldots\}$ is also expressed in terms of the basic rate, and it denotes the number of traffic units that originate at node s and terminate at node d.

Given the ring physical topology, a *logical* topology is defined by establishing *lightpaths* between pairs of nodes. A lightpath is a direct optical connection on a certain wavelength. More specifically, if a lightpath spans more than one physical link in the ring, its wavelength is optically passed through by WADMs at intermediate nodes, thus the traffic streams carried by the lightpath travel in optical form throughout the path between the endpoints of the lightpath. We assume that ring nodes are not equipped with wavelength converters, therefore a lightpath must be assigned the same wavelength on all physical links along its path.

An important problem in this context is the design of logical topologies that optimize a certain performance metric. The metric of interest in this work is the amount of electronic forwarding (routing) of traffic streams, since such forwarding involves electrooptic conversion and added message delay and processor load at the intermediate nodes. In a global sense, this means that we want to reduce the number of logical hops taken by traffic components, individually or as a whole. For each node, it also means that we want to reduce the amount of traffic that the node has to store and forward. Thus we have two alternative goals, one is to minimize the total traffic weighted logical hops in the network, and the other is to minimize the maximum number of traffic components electronically routed at a node. In

[1]We describe the working part of the ring. It is assumed that there is a protection part for self-healing or fault-tolerance as dictated by design or required by standards, but our description does not include this.

this paper we have chosen to concentrate on the former. Another possible quantity of interest is the number of wavelengths each WADM is required to add or drop due to lightpath terminations. This would correspond to minimizing of the number of SONET ADMs as in [3], [7], [9], [14]. It should be noted that adding and dropping more wavelengths will result in a larger number of logical hops for traffic components and thus this concern is to some extents incorporated in the electronic forwarding metric we consider.

We let $t(l)$ denote the aggregate traffic load on the physical link l (from node l to node $l \oplus 1$) of the ring. The value of $t(l)$ can be easily computed from the traffic matrix T. The component of the traffic load $t(l)$ due to the traffic from source node s to destination node d is denoted by $t^{(sd)}(l)$. If one or more lightpaths exist from node i to node j in the virtual topology, the traffic carried by those lightpaths is denoted by t_{ij}. The component of this load due to traffic from source node s to destination node d is denoted by $t_{ij}^{(sd)}$. In our formulation, we forbid a traffic component to be carried completely around the ring before being delivered at the destination, thus each traffic component can traverse a given link at most once. Also, a single traffic unit may not be bifurcated, but different traffic units for the same source-destination traffic component are considered as separate traffic and may be assigned different logical routes.

In our formulation we allow for multiple lightpaths with the same source and destination nodes. We denote the *lightpath count* from node i to node j by b_{ij}, taking its value from $\{0, 1, 2, \ldots, W\}$. We also define the *potential lightpath set for a link* to be the set of lightpaths that would pass through a given link and denote it by $B(l) = \{(i,j) \mid \text{lightpath } (i,j), \text{if it existed, would pass through link } l\}$. Finally, we let $c_{ij}^{(k)}$ be the *lightpath wavelength indicator*, i.e., $c_{ij}^{(k)}$ is 1 if a lightpath from node i to node j uses wavelength k, 0 otherwise.

We can now formulate the problem of designing a virtual topology for a ring network such that the total amount of electronic routing at the ring nodes is minimized. The following formulation as an ILP consists of $O(N^4 + N^2W)$ constraints and $O(N^4 + N^2W)$ variables, where N is the number of nodes in the ring, and W is the number of wavelengths.

Given:

The physical topology, a unidirectional ring $\mathcal{R}$ of N nodes.

The traffic matrix $T = [t^{(sd)}], s, d \in \{0 \cdots (N-1)\}, t^{(sd)} \in \{0, 1, 2, \ldots\}, t^{(ss)} = 0, \forall s$.

The wavelength limit W which is the number of distinct wavelengths each link can carry.

Find:

Virtual topology, in terms of lightpath indicators b_{ij}, lightpath wavelength indicators $c_{ij}^{(k)}$, and traffic routing variables $t_{ij}^{(sd)}$.

Subject to:

Traffic Constraints

$$t_{ij}^{(sd)} \le t^{(sd)}(i), \quad \forall (i,j), (s,d) \tag{1}$$

$$t_{ij}^{(sd)} \in \{0, 1, 2, \ldots\}, \quad \forall (i,j) \tag{2}$$

$$\sum_{(i,j) \in B(l)} t_{ij}^{(sd)} = t^{(sd)}(l), \quad \forall (s,d), l \tag{3}$$

$$t_{ij} = \sum_{sd} t_{ij}^{(sd)}, \quad \forall (i,j) \tag{4}$$

$$t_{ij} \le b_{ij} C, \quad \forall (i,j) \tag{5}$$

$$\sum_j t_{ij}^{(sd)} - \sum_j t_{ji}^{(sd)} = \begin{Bmatrix} t^{(sd)}, & s = i \\ -t^{(sd)}, & d = i \\ 0, & s \ne i, d \ne i \end{Bmatrix} \quad \forall i, (s,d). \tag{6}$$

Wavelength Constraints

$$\sum_{(i,j) \in B(l)} b_{ij} \le W, \quad \forall l \tag{7}$$

$$\sum_{k=1}^{W} c_{ij}^{(k)} = b_{ij}, \quad \forall (i,j) \tag{8}$$

$$\sum_{(i,j) \in B(l)} c_{ij}^{(k)} \le 1, \quad \forall l, k. \tag{9}$$

To minimize:

$$\left(\sum_{s,d,i,j \in \{0 \cdots (N-1)\}} t_{ij}^{(sd)} - \sum_{s,d \in \{0 \cdots (N-1)\}} t^{(sd)} \right).$$

The traffic constraint (1) ensures that a lightpath can carry traffic for a source-destination node pair only if it is in the physical route of the traffic component. Constraint (3) states that the physical traffic on a link due to a source-destination node pair must be equal to the sum of the traffic on all lightpaths passing through that link due to that node pair. Constraints (4) and (5) define the total traffic on a lightpath and relate it to the lightpath count, respectively. Because of the definition of the quantities $t^{(sd)}(l)$, constraints (1) and (3) together ensure that no traffic component can be routed completely around the ring before being delivered at the destination node. Constraint (6) is an expression of traffic flow conservation at lightpath endpoints. Among the wavelength constraints, constraint (7) expresses the bound imposed by the number of wavelengths available, (8) relates the wavelength indicators to the lightpath counts, and (9) ensures that no wavelength clash can occur.

The above problem can be thought of as consisting of the following three subproblems [5].

1) **Topology Subproblem:** Determine the virtual topology to be imposed on the physical topology, that is, determine the lightpaths in terms of their source and destination nodes.
2) **Wavelength Assignment Subproblem:** Assign a wavelength to each lightpath in the virtual topology so that wavelength restrictions are obeyed for each physical link.
3) **Traffic Routing or Grooming Subproblem:** Groom the traffic components on the lightpaths, that is, route the traffic streams over the virtual topology obtained.

The framework we present below is based on the above formulation we have chosen, but this formulation is not essential for it. The framework can be adapted to many variations that are possible in the formulation. There may be multiple fiber links between successive nodes, and the nodes may be equipped with wavelength routers instead of WADMs. Hardware for wavelength conversion may be available at the nodes [11]. A phys-

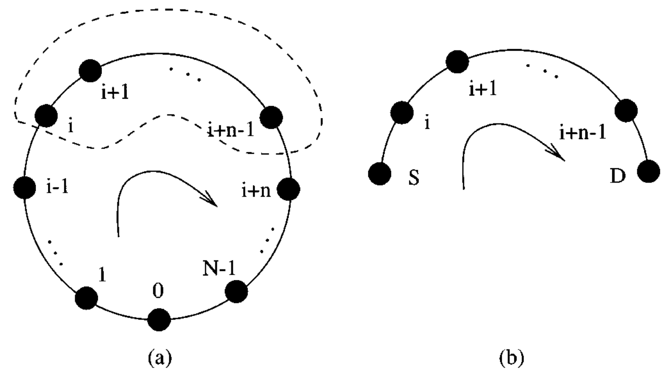

Fig. 2. An n-node decomposition: (a) original ring $\mathcal{R}$ and (b) a decomposition $\mathcal{P}_n^{(i)}$.

ical hop limit for lightpaths may be imposed. The ring may be bidirectional, either with some simple routing strategy (such as shortest path, as in [14] and [7]) that allows us to consider it as two unidirectional rings, or the lightpath routing (either clockwise or counterclockwise) may be integrated as part of the optimization process (as in [9]). In all these cases, the objective may be to minimize electronic routing. The framework we present may be extended, in a straightforward manner for some of these cases, and with some enhancements in others. When the objective function is of a min–max type (e.g., minimize the maximum electronic routing at any node) or a quantized version (e.g., minimize the number of SONET ADMs), our framework can also be adapted to obtain bounds for the optimal value.

III. Path Decomposition of a Ring Network

A. Definition of Decomposition

We consider a ring $\mathcal{R}$ with N nodes and traffic matrix T. We define a *segment* of length $n, 1 \le n \le N$, starting at node $i, 0 \le i < N$, as the part of the ring that includes the n consecutive nodes $i, i \oplus 1, \ldots, i \oplus (n-1)$, and the links between them. We define a *decomposition of ring $\mathcal{R}$ around a segment of length n starting at node i* as a path $\mathcal{P}_n^{(i)}$ that consists of $n+2$ nodes and $n+1$ links as follows: the n nodes and $n-1$ links of the segment of length n starting at node i, a new node S and a link from S to i, and a new node D and a link from node $i \oplus (n-1)$ to D. We also refer to $\mathcal{P}_n^{(i)}$ as an n-node decomposition of ring $\mathcal{R}$ starting at node i. Fig. 2 shows such a decomposition.

Associated with the decomposition $\mathcal{P}_n^{(i)}$ is a new traffic matrix $T_{\mathcal{P}_n^{(i)}} = [t_{\mathcal{P}_n^{(i)}}^{(sd)}], s, d \in \{i, i \oplus 1, \ldots, i \oplus (n-1), D, S\}$, derived from T, the original traffic matrix, as follows:

$$
\begin{aligned}
t_{\mathcal{P}_n^{(i)}}^{(sd)} &= t^{(sd)}, \quad i \preceq s \prec d \preceq i \oplus (n-1) \\
t_{\mathcal{P}_n^{(i)}}^{(Sd)} &= \sum_{j \notin \{i, \ldots, d \ominus 1\}} t^{(jd)}, \quad i \preceq d \preceq i \oplus (n-1) \\
t_{\mathcal{P}_n^{(i)}}^{(sD)} &= \sum_{j \notin \{s \ominus 1, \ldots, i \ominus (n-1)\}} t^{(sj)}, \quad i \preceq s \preceq i \oplus (n-1) \\
t_{\mathcal{P}_n^{(i)}}^{(SD)} &= t_{\text{pass-thru}}(i, n) \\
t_{\mathcal{P}_n^{(i)}}^{(ss)} &= t_{\mathcal{P}_n^{(i)}}^{(Dd)} = t_{\mathcal{P}_n^{(i)}}^{(sD)} = 0, \quad \forall s, d \\
t_{\mathcal{P}_n^{(i)}}^{(sd)} &= 0, \quad i \preceq d \prec s \preceq i \oplus (n-1)
\end{aligned} \tag{10}
$$

where $t_{\text{pass-thru}}(i, n)$ denotes the traffic of the original matrix T that *passes through* the segment of length n starting at node i, i.e., traffic on ring $\mathcal{R}$ that uses the links of the segment but does not either originate or terminate at any of the nodes in that segment. We call this the *pass-through traffic*. The amount of this traffic can be readily obtained by inspection of traffic matrix T. We have used $s \prec d$ in the above expression to denote that node s precedes node d in the decomposition and $s \preceq d$ to denote that node s precedes and may be the same as node d in the decomposition.

The traffic matrix for the decomposition is defined such that the traffic flowing from a node s to another node $d, s \prec d$, in the segment is the same as that in the original ring [first expression in (10)]. Thus any traffic component, the path of which is entirely in the segment, is unchanged in the decomposition. The decomposition is effected by the introduction of nodes S and D together with the links connecting them to the segment. Node S acts as the source of all traffic components in matrix T originated at a node outside the segment and destined to any node in the segment [second expression in (10)]. Node S also acts as the source for all traffic components that pass through the segment, as the fourth expression in (10) indicates. Similarly, node D is the sink for traffic originating at any node in the segment and terminating at a node outside the segment [third expression in (10)], as well as for pass-through traffic. Finally, any traffic components in matrix T that do not traverse any links of the segment are not included in the traffic matrix for the decomposition. This is captured by the last two expressions in (10) where it is shown that no traffic flows from node D to node S in the decomposition.

Because of the way the traffic matrix for the decomposition is defined in (10), from the point of view of any node $k, i \preceq k \preceq i \oplus (n-1)$, in the segment, the traffic pattern in the new path $\mathcal{P}_n^{(i)}$ is *exactly the same* as in the original ring. The new nodes S and D are introduced in the decomposition to abstract the interaction of traffic components between nodes in and outside the segment. Specifically, the new node S hides the details of how traffic sourced at ring nodes outside the segment and using the links in the segment actually flows over the rest of the ring, by providing a single aggregation point for this traffic. Similarly, the new node D provides a single aggregation point for traffic using the links of the segment and destined to nodes outside the segment, hiding the details of how this traffic flows in the rest of the ring. Finally, the fact that $\mathcal{P}_n^{(i)}$ is a path (i.e., that there is no link from node D to node S) means that the details of traffic in the original ring that does not involve any nodes or links of the segment are hidden in the decomposition.

B. Solving Path Segments in Isolation

Consider the matrix $T_{\mathcal{P}_n^{(i)}} = [t_{\mathcal{P}_n^{(i)}}^{(sd)}]$ of the decomposition $\mathcal{P}_n^{(i)}$ of a segment of length n starting at node i in the ring $\mathcal{R}$, as given in (10). This matrix can be thought of as representing the traffic in a ring network consisting of nodes $S, i, \ldots, i \oplus (n-1), D$, where there is simply no traffic flowing over the link from D to S. Consequently, the ILP formulation in Section II can be used to obtain a virtual topology that minimizes electronic routing for this "ring." Since the formulation disallows routing that carries a traffic component beyond its destination and all around the ring, no lightpaths can be formed to carry traffic over the link from D to S that is absent in the decomposition. Thus, the topology obtained in this manner can be directly applied to the path $\mathcal{P}_n^{(i)}$. When we use the ILP to find an optimal topology for path $\mathcal{P}_n^{(i)}$ we will say that we solve the n-node segment *in isolation*.

The topology obtained by solving an n-node segment in isolation does not take into account the details of the original ring outside of the segment. Such a topology will only be optimal with respect to this segment, in the sense that it will minimize the amount of electronic routing within the segment, but without considering the effects that doing so would have on the amount of electronic routing at nodes of ring $\mathcal{R}$ outside the segment. In fact, this topology may not be optimal for the ring as a whole. In other words, it is possible that, for any optimal topology for the ring as a whole, the subtopology corresponding to the n-node segment will be different than the topology obtained by solving the ILP for $\mathcal{P}_n^{(i)}$ in isolation. Thus it may not be possible to combine optimal solutions to different segments in isolation into a (near-)optimal topology for the ring $\mathcal{R}$. Our contribution is in proving a weaker result: that it is possible to combine optimal solutions to different segments in isolation to obtain lower and upper bounds on the optimal solution to the ring $\mathcal{R}$ as a whole.

As the number n of nodes in a path segment increases, the resulting decomposition will more closely approximate the original ring and the bounds will be tighter. More importantly, a path decomposition significantly alleviates the problem of exponential growth in computational requirements for solving the ILP. This result is a direct consequence of the following lemma.

Lemma III.1: Consider a virtual topology on a unidirectional path, in which the maximum number of lightpaths on any link is L. A wavelength assignment for the virtual topology that uses exactly L wavelengths can be obtained in time linear in L and the number of links of the path.

For the proof, see [13, Sec. 8.5.2].

In solving the decomposed problem, we are merely interested in the optimum value of the objective, since this is the value from which we will obtain the bounds. Since we know that a wavelength assignment is always possible and we are not interested in the details of the wavelength assignment, we can eliminate the wavelength assignment from our formulation altogether. This implies that we can eliminate subproblem (3), the wavelength assignment subproblem, from the list of subproblems we enumerated in Section II. Thus, we eliminate the wavelength variables $c_{ij}^{(k)}$ and the constraints (8) and (9). The number of variables in the formulation remains $O(N^4)$, but N^2W variables are eliminated, and the number of wavelength constraints is reduced from $O(N^2W)$ to $O(N)$, though the total number of constraints remains $O(N^4)$. This creates a formulation that is smaller and requires dramatically less computation to solve. In practice, we have found that eliminating the wavelength assignment subproblem can result in a reduction in computation time by several orders of magnitude. For instance, the LINGO scientific computation package takes between 60 and 90 minutes on a Sun Ultra-10 workstation to solve a six-node ring using the original formulation (with wavelength assignment). However, solving a six-node network using the simplified formulation (no wavelength assignment) requires only a few seconds.

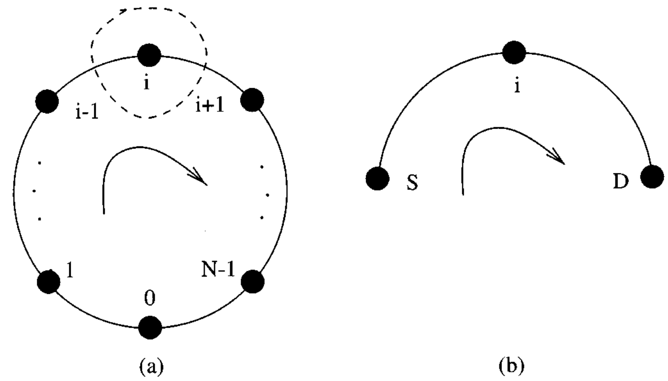

Fig. 3. A single node decomposition of a ring: (a) original ring, and (b) single node decomposition $\mathcal{P}_1^{(i)}$ around node i.

We denote the optimal objective value for the decomposition $\mathcal{P}_n^{(i)}$ by $\phi_n^{(i)}$. The two additional nodes S and D, by the construction of the decomposition, do not have any traffic passing through them at all, and hence they do not contribute any electronic routing. Since the traffic pattern seen by the n nodes abstracted from the ring is the same as when they are included in the ring, $\phi_n^{(i)}$ also represents the locally best (lowest) amount of electronic routing at this set of nodes when considered as part of the ring. That is, the electronic routing at this set of nodes is minimized, irrespective of how much routing has to be performed at other nodes of the original ring as a result.

We note that, because we convert the ring to a path, some information is not taken into account at the point where we "open up" the ring even when considering an N-node path. As a result, solving an N-node decomposition does not provide the optimal solution for the N-node ring.

IV. Bounds

We now describe how to combine the $\phi_n^{(i)}$ values we get from n-node decompositions to obtain lower as well as upper bounds on the total amount of electronic routing performed in the optimal case by a virtual topology on the original ring. We first discuss the case where we only consider single-node decompositions, then move on to the general case where larger decompositions are available.

A. Bounds Based on Single-Node Decompositions

1) Lower Bound: Recall that $\phi_n^{(i)}$ represents the locally best amount of electronic routing at the nodes in the segment of length n starting at node i. In particular, $\phi_1^{(i)}$ represents the locally best amount of electronic routing that can be achieved at node i considered in isolation (see Fig. 3). There may or may not be an optimal (or even feasible) virtual topology for the ring $\mathcal{R}$ that achieves this value of electronic routing at node i, but there can be no topology which achieves an even lower value. Thus, $\phi_1^{(i)}$ is a lower bound on the amount of electronic routing performed at node i for any virtual topology.

Since $\phi_1^{(i)}$ represents contribution to the electronic routing only by node i, we can add the contributions together for each node to obtain a lower bound on the total electronic routing performed in any virtual topology. We call this lower bound Φ_1

$$\Phi_1 = \sum_{i=0}^{N-1} \phi_1^{(i)}. \tag{11}$$

2) Upper Bound: We first consider an upper bound we can obtain directly from the traffic matrix, without recourse to decompositions. This bound corresponds to the virtual topology consisting only of single-hop lightpaths. Consider node i. In this topology, single-hop lightpaths from node $i \ominus 1$ carrying all traffic to node i and beyond terminate at node i. Node i electronically switches all traffic for which it is not the destination, combines it with its own outgoing traffic, and sources a number of single-hop lightpaths (up to W) that carry this traffic to node $i \oplus 1$. We will call this the *no-wavelength-routing* topology, since no wavelengths are optically routed at any node. In such a topology, each node i performs the maximum possible amount of electronic routing, which we denote by $\psi^{(i)}$. Quantities $\psi^{(i)}, i = 0, \ldots, N-1$, can be readily obtained from the traffic matrix T. We let Ψ_0 denote the total electronic routing performed under the no-wavelength-routing topology

$$\Psi_0 = \sum_{i=0}^{N-1} \psi^{(i)}. \tag{12}$$

Since this is a feasible topology, Ψ_0 is an upper bound on the optimal electronic routing.

In general, Ψ_0 is a rather loose upper bound. We now show how to utilize single-node decompositions to improve upon the no-wavelength-routing topology to obtain a tighter bound. Let us refer to Fig. 3(b) which shows a decomposition around node i. Recall now that, in deriving the best local electronic routing $\phi_1^{(i)}$ at node i, we made the assumption that all traffic passing through node i is originated at S and terminated at D. Let us call *concentrator nodes* those nodes which do not perform any optical forwarding (i.e., they terminate and originate all lightpaths). In the no-wavelength-routing topology, every node is a concentrator node. In the single node decomposition nodes S and D can be viewed as concentrator nodes. Thus, we are led to consider a virtual topology such that every other node is a concentrator node (performing the maximum electronic routing possible, $\psi^{(i)}$) while the remaining nodes perform the minimum electronic routing possible, $\phi_1^{(i)}$. This topology is illustrated in Fig. 4, where even-numbered nodes are concentrator nodes.

Depending on which nodes in the ring we select to be the concentrator nodes, we obtain different virtual topologies which yield potentially different values for the total amount of electronic routing. If the number of nodes N in the ring is even, we have two possible topologies, depending on whether even-numbered or odd-numbered nodes are concentrators. If N is odd, any virtual topology constructed in the manner described above

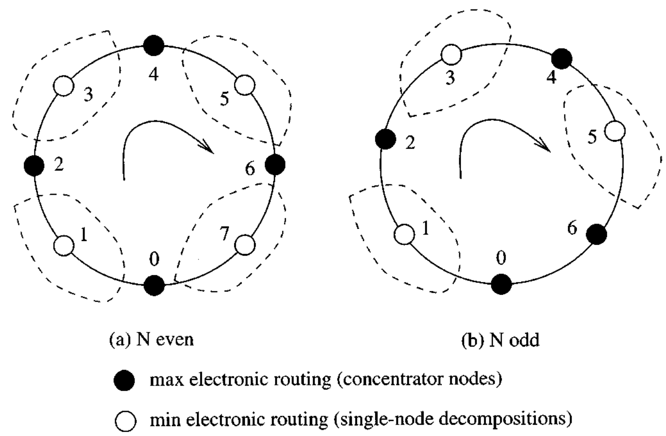

Fig. 4. Virtual topology with alternating single-node decompositions and concentrator nodes.

will have two concentrator nodes next to each other at one position in the ring, as illustrated in Fig. 4(b). Since there are N ways of selecting the position of these two concentrator nodes, there are N possible virtual topologies when N is odd. We take the smallest value of total electronic routing we can obtain from these topologies as the upper bound. This bound is indicated by Ψ_1, and in the general case, it can be expressed as (13), shown at the bottom of the page. Since this is a feasible topology which incurs the maximum electronic routing only at the concentrator nodes while it incurs less at the others, the upper bound set by the objective value of this topology must be at least as strong as Ψ_0; thus we also have that $\Psi_1 \leq \Psi_0$.

B. Bounds Based on Larger Decompositions

1) Lower Bound: In obtaining the bound Φ_1 above, we remarked that we can add the various $\phi_1^{(i)}$ quantities together since they each represented electronic routing at node i only. Consider the quantity $\phi_2^{(i)}$; it represents the minimum possible amount of electronic routing (best case) at node i *and* node $i \oplus 1$ taken together. We cannot add $\phi_1^{(i)}$ or $\phi_1^{(i \ominus 1)}$ to this quantity and still have something that is guaranteed to be a lower bound on the amount of electronic routing these nodes together perform in any feasible topology, because we are potentially counting the traffic routed by a single node twice. However, we can add $\phi_2^{(i)}$ and $\phi_1^{(i \ominus 2)}$, since the two quantities involve sets of nodes that are disjoint, and, therefore, there is no potential double counting of electronic routing. Generalizing this notion, we can add the $\phi_n^{(i)}$ values for any set of decompositions that involve disjoint segments, and we are still guaranteed to obtain a lower bound on the objective value for any feasible topology. We formalize this in the following lemma. The proof follows obviously from the arguments above.

Lemma IV.1: Let σ_n be a set of segments of ring $\mathcal{R}$ which partition the nodes of the ring in segments of length n or smaller.

$$\Psi_1 = \min_{j \in [0, N-1]} \left(\sum_{k \in \{0,2,4,\ldots,2(\lfloor (N-2)/2 \rfloor)\}} \phi_1^{(j+k)} + \sum_{k \notin \{0,2,4,\ldots,2(\lfloor (N-2)/2 \rfloor)\}} \psi^{(j+k)} \right) \tag{13}$$

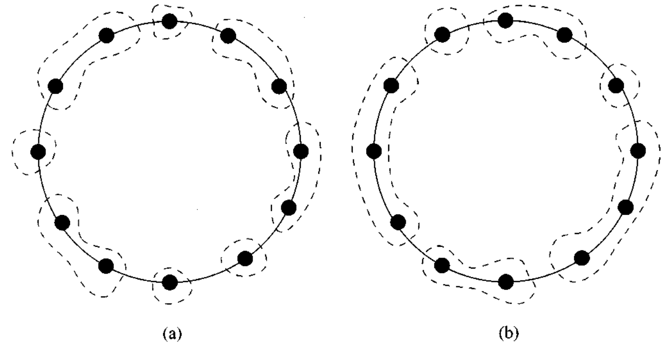

Fig. 5. Partitions of the nodes of a ring into (a) segments of no more than two nodes and (b) segments of no more than three nodes.

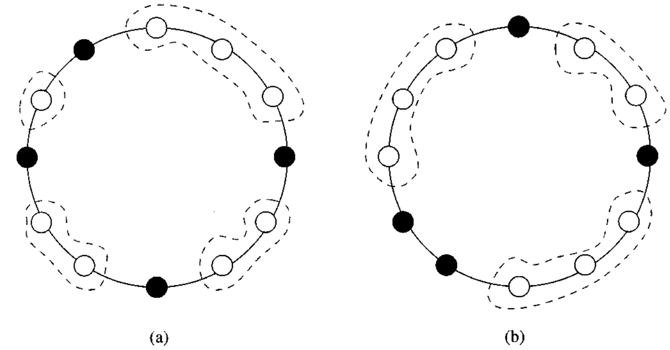

Fig. 6. Two partitions of a ring into alternating concentrator nodes and segments of no more than three nodes.

Let $l_k, l_k \leq n$ denote the length (number of nodes) of segment $k, k = 1, \ldots, |\sigma_n|$, and let i_k denote its starting node. The quantity

$$\Phi(\sigma_n) = \sum_{k=1}^{|\sigma_n|} \phi_{l_k}^{(i_k)} \tag{14}$$

is a lower bound on the objective value for any feasible virtual topology on the ring $\mathcal{R}$ and therefore on the optimal objective value.

We now define Φ_n as

$$\Phi_n = \max\{\Phi(\sigma_n)\} \tag{15}$$

where the maximum is taken over all partitions of the ring which contain segments with n or less nodes. Fig. 5 shows two partitions of the same ring, the first containing only 1- and 2-node segments, and the second containing only 1-, 2-, and 3-node segments.

Because of (15), in computing bound Φ_{n+1} we must consider all partitions (and bounds derived therefrom) considered to compute Φ_n, and, additionally, all partitions which include one or more $(n+1)$-node segments. Specifically, the set of partitions σ_{n+1} we consider for Φ_{n+1} is a proper superset of the set of partitions σ_n we consider for Φ_n. Since we draw the maximum bound from each set as per the definition in (15), we have that

$$\Phi_{n+1} \geq \Phi_n \quad \forall n \in \{1, \ldots, N-1\}. \tag{16}$$

As a result, the sequence $\Phi_1, \Phi_2, \ldots, \Phi_N$ is a strong sequence of bounds. We also note that the definition of Φ_1 in Section IV-A.1 is consistent with the definition above.

2) Upper Bound: It is now straightforward to obtain a strong sequence of upper bounds along the same lines. We define Ψ_n as the lowest objective value we obtain for all topologies which are created by alternating concentrator nodes with segments no larger than n nodes in size. We can once again consider this in light of partitions of the nodes of a ring. Now, however, the partitions are constrained not only to use segments of n-nodes or less, but every alternate segment must contain exactly one node. These alternate single-node segments are used as concentrator nodes in the topology we create, rather than as single-node decompositions. Once again, the form of this upper bound is an alternate sum of $\phi_x^{(i)}$ and $\psi^{(i)}$ values, similar to expression (13) for Ψ_1; but for Ψ_n the value of x is not restricted to 1 as for Ψ_1, instead it can take on any value from 1 to n. Fig. 6 shows two ways we can partition a ring using no larger than three-node segments, thus creating two topologies among the ones we would consider in computing Ψ_3.

We note that the bounds Ψ_0 and Ψ_1 we obtained in Section IV-A.2 are consistent with this framework. As before, the set of topologies we consider in obtaining Ψ_{n+1} is a superset of the set of topologies considered in Ψ_n, therefore we have a strong sequence of upper bounds

$$\Psi_{n+1} \leq \Psi_n \quad \forall n \in \{0, \ldots, N-2\}. \tag{17}$$

Because the bounds $\{\Psi_n\}$ are based on feasible topologies, they also provide us with a useful series of heuristic solutions to the ring. In the next section we derive a result which shows the tightness of the bounds and thus the goodness of the heuristics, and we see in Section V that even the first few solutions of the series can outperform a simplistic heuristic. The later solutions in the series can compare favorably with some heuristics reported in literature. Specifically, Ψ_{N-1} must be as good or better than a *single-hub* architecture [8], [3], [2] because it considers all topologies with a single concentrator (hub) node. For a similar reason, $\Psi_k, k \geq \lceil N/2 \rceil$ must be as good or better than a double hub design, if the hubs are constrained to be diametrically opposite in the ring.

3) Tightness of Bounds: Consider the value $\psi^{(i)} - \phi_1^{(i)}$ for each node of a ring, which is the difference between the minimum and maximum traffic the node can route in any virtual topology. Let the node for which this difference is minimum be node j, and let the corresponding difference be $\zeta^{(j)}$, so that

$$\zeta^{(j)} = \min_{i=0}^{N-1} \left(\psi^{(i)} - \phi_1^{(i)} \right). \tag{18}$$

The final upper and lower bounds are guaranteed to be no further apart than this quantity. We state this in the following lemma, the proof of which is omitted here and can be found in [6].

Lemma IV.2: The guarantee on the final values in the sequences of upper and lower bounds is

$$\Psi_{N-1} - \Phi_{N-1} \leq \zeta^{(j)}. \tag{19}$$

4) Computational Considerations: The bounds Φ_n (and Ψ_n) for successive values of n incorporate progressively more information about the problem and as such require progressively more computational effort to determine. This increase in computational effort manifests itself in two ways:

1) the calculation of the $\phi_n^{(i)}$ values required for a given bound;
2) the evaluation of all partitions of the ring in segments of length at most n by an appropriate combination of the $\phi_n^{(i)}$ values.

In the discussion that follows we focus on the sequence $\{\Phi_n\}$ of the lower bounds, but the observations we make are equally applicable to the sequence of upper bounds.

The computation of a bound utilizing a certain size of decompositions requires knowledge of decompositions of all lower sizes as well. Thus, computing Φ_x would require us to compute $\phi_n^{(i)}$ for all values of $i \in \{0, \ldots, N-1\}$ and all values of $n \in \{1, \ldots, x\}$. However, the *incremental* computation of decompositions required to determine Φ_x consists only of determining $\phi_x^{(i)}$ for all nodes i, since $\phi_n^{(i)}$ for $n < x$ would already have been computed when determining Φ_{x-1}. Naturally, as x increases, this incremental effort required also increases; as we have noted before the number of variables and constraints increases as $O(n^4)$. Thus the maximum value of n for which we can determine the corresponding bound is limited by this computational effort.

Regarding the second factor that affects the computation time required to obtain Φ_n, we note that a straightforward approach would require us to enumerate all partitions of an N-node ring into segments of length at most n. While *evaluating* each partition (i.e., computing the lower bound for it) takes time linear in the number of segments of the partition (assuming that the individual $\phi_x^{(i)}, x \in \{1, \ldots, n\}$, values are available), the number of possible partitions increases very rapidly with n. We show in [6] that the incremental number of partitions to consider for a given value of n is exponential in N. Thus, a straightforward approach to combine decompositions into bounds would severely limit the maximum value of N for which we can determine the corresponding bounds.

However, by exploiting the particular characteristics of $\phi_n^{(i)}$, we have developed a polynomial-time algorithm to compute Φ_n, assuming that the appropriate $\phi_n^{(i)}$ values are available. The algorithm is based on incrementally building the best sum of $\phi_n^{(i)}$ around the ring and following only the best partial sum at an intermediate node. This algorithm is presented as a dynamic programming problem in Appendix A and requires $O(n^2 N)$ time to find Φ_n given all $\phi_x^{(i)}$ values for $x = 1, \ldots, n$. For the largest number of total partitions in the case $n = N$, this corresponds to $O(N^3)$ instead of $O(2^N)$ time, and in fact becomes linear in N for a small given value of n.

We can achieve an additional constant factor of improvement by using the properties of $\phi_n^{(i)}$ values formalized in the following lemma; its proof is omitted but can be found in [6].

Lemma IV.3: An $(x+y)$-node decomposition yields at least as large an objective value for the decomposed network as the sum of objective values of the x-node and y-node decompositions it exactly contains. That is, $\phi_{x+y}^{(i)} \geq \phi_x^{(i)} + \phi_y^{(i+x)}$, if x and y are positive and $x + y < N$.

Corollary IV.1: An x-node decomposition yields at least as large an objective value as the sum of objective values of any combination of smaller decompositions it can be partitioned into. That is, $\phi_x^{(i)} \geq \phi_{y_1}^{(i)} + \phi_{y_2}^{(i+y_1)} + \phi_{y_3}^{(i+y_1+y_2)} + \cdots + \phi_{y_n}^{(i+y_1+y_2+\cdots+y_{n-1})}$, if $x = y_1 + y_2 + \cdots + y_n$.

The corollary follows immediately from the lemma by repeated application within the same decomposition and allows us to discard partitions in which a small segment follows another. Specifically, if we are computing Φ_x, we can discard a partition in which a y_1-node segment is followed by a y_2-node segment if $y_1 + y_2 \leq x$, since the partition we obtain by replacing these segments by a $(y_1 + y_2)$-node segment must yield a higher bound, and we are interested in the maximum bound.

The above methods allow us to compute the $\{\Phi_n\}$ bounds by combining the $\phi_n^{(i)}$ values in an insignificant fraction of the time taken to compute the $\phi_n^{(i)}$ values themselves. In practice, we have found that computing the $\phi_n^{(i)}$ values takes minutes (even hours) for increasing n, while combining them to form the $\{\Phi_n\}$ bounds takes milliseconds. We conclude that the limiting factor in determining how many bounds can be computed in a reasonable amount of time is the effort required to solve the ILP for n-node decomposition in order to compute each of the $N \phi_n^{(i)}$ values.

V. Numerical Results

In this section, we present the results of using our framework for different traffic matrices. We first create the distinction between *symmetric* and *asymmetric* traffic patterns. The term symmetric applies to the ring, rather than the traffic matrix itself. We call a traffic pattern symmetric if the traffic pattern from any node to the others is repeated for all the nodes. This type of traffic pattern is of interest since the traffic pattern looks similar from different nodes on the ring. In the general case, traffic components of the form $t^{(s \ominus x, d \ominus x)}$ for every given s and d, and for all values of x, are all drawn from the same distribution. If the variance of this distribution is zero, we call the resulting traffic pattern *strictly symmetric*; otherwise, we call it *statistically symmetric*. With respect to what traffic looks like from a given node, we consider three simple cases. First, the traffic from a given node i to all other nodes may be the same, we refer to this as *uniform* traffic. When the traffic components to all the other nodes are not the same, we can speak of a *falling* traffic pattern in which the traffic from node s to node $s \oplus x$ decreases linearly as x increases. Similarly, we speak of a *rising* pattern. Again, we introduce the concept of statistical variation so that actual matrix elements vary from these patterns statistically and do not vary strictly linearly as described above.

In order to have a good basis for comparison for the above three types of traffic patterns, we focus on the concept of *characteristic physical load* of the traffic matrix. For the problem instance to be feasible, the traffic flowing over each physical link must be less than or equal to WC. If the matrix is statistically symmetric, the loads on the links will all be close to some value, because the traffic pattern is the same looking from any node or link. We call this value the characteristic physical load of the matrix and obtain it by taking the average of the physical load on each link and express it as a percentage of WC. For the

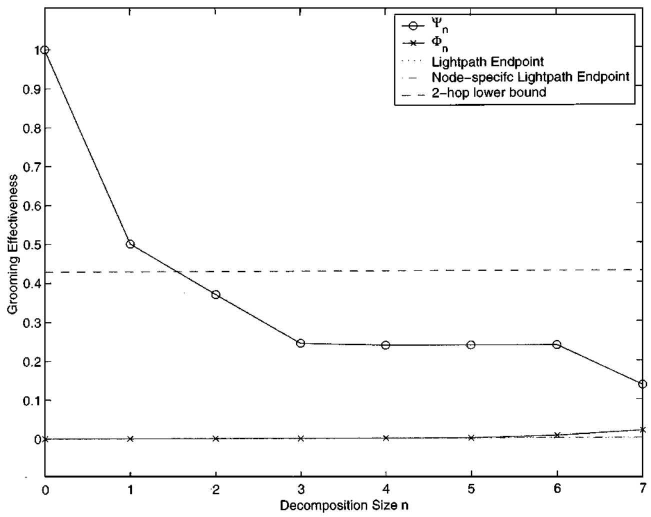

Fig. 7. $N = 8$, statistically uniform, 50% load.

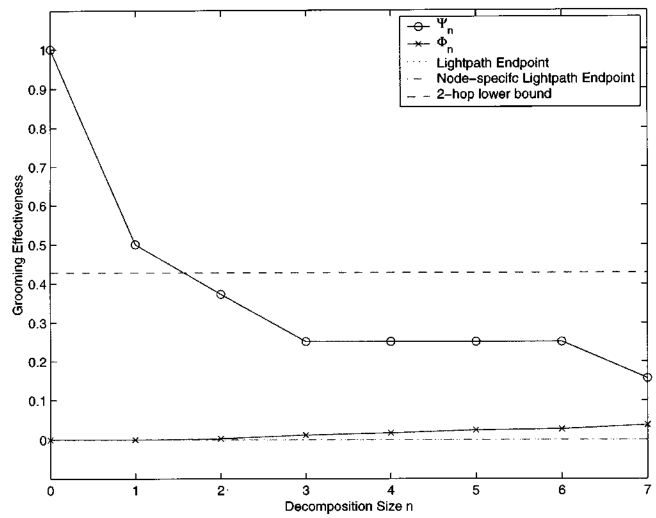

Fig. 8. $N = 8$, statistically uniform, 90% load.

same pattern, the characteristic load scales with the matrix elements. For example, by multiplying each element of a matrix of characteristic load 50% by a factor of 1.5, it is converted into a matrix of characteristic load 75% but of the same pattern.

We present results pertaining to 8-node and 16-node rings. For most of our results, the value of W was taken to be between 16 and 20 and the value of C was 48. We used randomly generated statistically symmetric traffic matrices for all the runs. A discrete uniform probability distribution was used for all traffic generation. We focus on characteristic physical load values of 50% and 90%. Only a sampling of the results obtained are presented here.

Because of different traffic patterns and different characteristic load values, the absolute values of the different bounds plotted are not easy to compare. It is necessary to express them in terms of some characteristic of the problem that makes it sensible to compare them. Recall that Ψ_0 denotes the amount of electronic routing performed by a topology that does not employ optical forwarding at all and corresponds to *no grooming*. The other extreme (not necessarily achievable) is *complete grooming*, in which all traffic is groomed into lightpaths and no electronic routing is performed. The actual amount of electronic routing performed by any feasible topology falls between these limits and may be expressed as a fraction of Ψ_0 to indicate the *effectiveness of grooming*. Thus, 1 indicates that all the traffic has been left ungroomed, while 0 corresponds to the best situation in which no traffic is left to be groomed. The values $\{\Psi_n\}$ expressed as such ratios indicate the upper bounds on the optimal grooming effectiveness, while the values $\{\Phi_n\}$ represent lower bounds on the grooming effectiveness that can be reached in the optimal case. In our plots, we normalize each set of results to the corresponding value of Ψ_0 and plot the grooming effectiveness ratios as above. Other quantities in the plot which we discuss below are similarly normalized.

We present two broad sets of data. In the first, or detailed section, we present Φ_n and Ψ_n for values of n up to 7, for $N = 8, 16$, for loads of 50% and 90%, and statistically uniform, rising and falling patterns. Figs. 7 and 8 show the results for eight-node rings, statistically uniform traffic, 50% and 90% characteristic load, respectively. Figs. 9 and 10 similarly show the results for statistically falling traffic, 50% load, and rising traffic, 90% load, respectively. Figs. 11 and 12 both show detailed results for 16-node rings with falling data and 90% load; in the first case the traffic from a given node falls to zero at the furthest destination node along the ring, while in the second case it falls to zero at a destination node halfway around the ring and is zero for all nodes further along the ring. The latter pattern could be of interest if a bidirectional ring is decomposed into two unidirectional rings by adopting shortest path routing for all traffic components *a priori*.

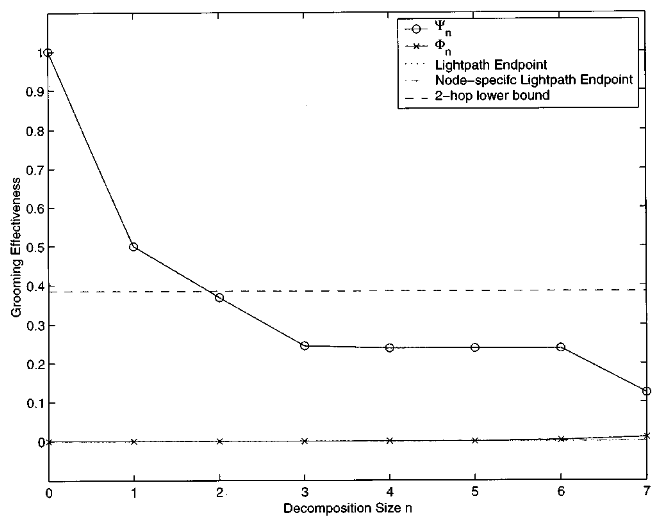

Fig. 9. $N = 8$, statistically falling, 50% load.

We observe that all the figures look similar. There is a sharp decrease from Ψ_0 to Ψ_1 and more moderate decrease thereafter. In all cases of eight-node rings, there is a marked decrease for Ψ_{N-1}, and the grooming effectiveness for Ψ_{N-1} is between 0.1 and 0.2 in all cases. For 16-node cases, Ψ_7 (which is no longer the final value Ψ_{N-1}) is again between 0.1 and 0.2. Thus we generally observe that we get good grooming effectiveness and that the lower bounds are comparatively small in magnitude. This indicates that a high value of electronic routing for some

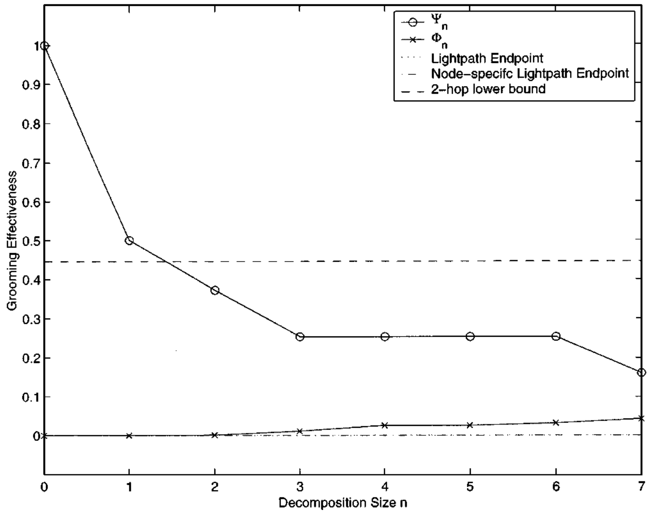

Fig. 10. $N = 8$, statistically rising, 90% load.

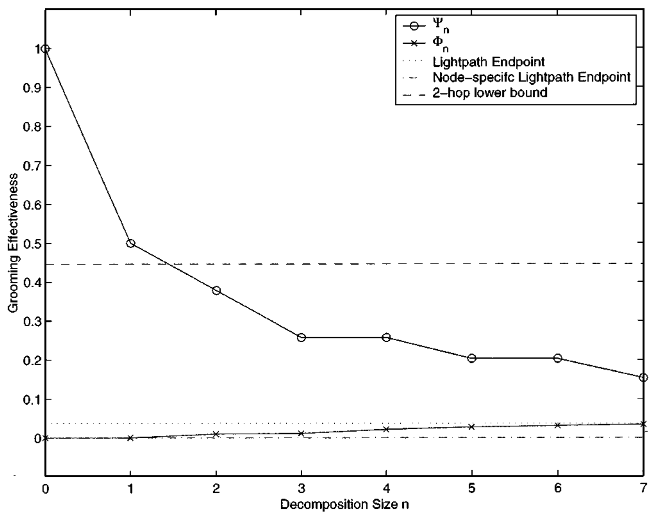

Fig. 11. $N = 16$, statistically falling to end, 90% load.

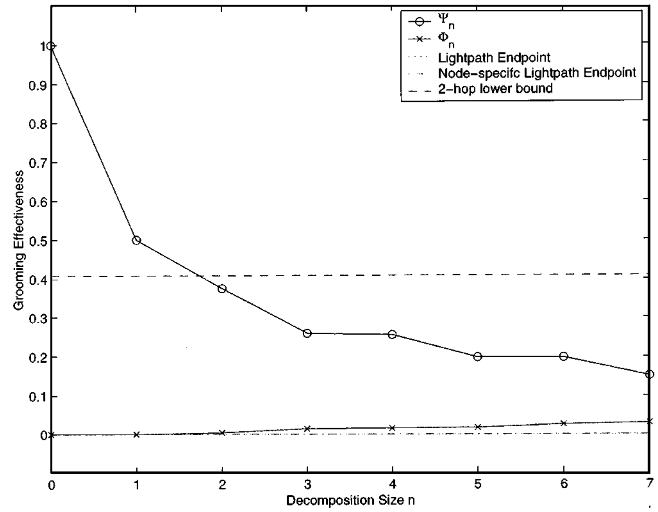

Fig. 12. $N = 16$, statistically falling to $N/2$, 90% load.

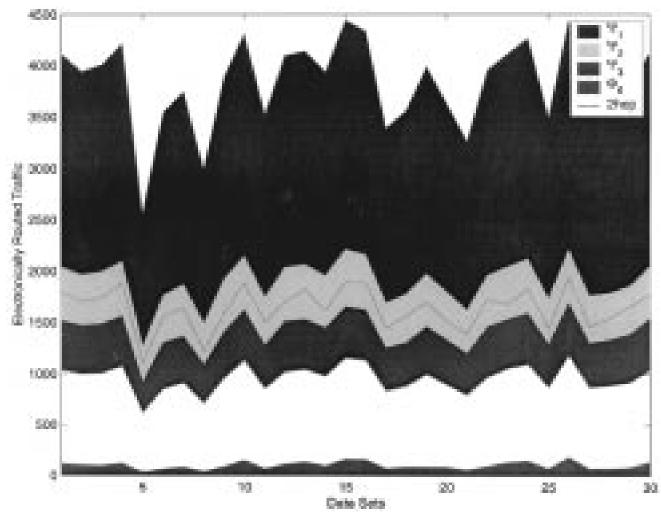

Fig. 13. Ensemble, $N = 8$, statistically uniform, 90% load, electronic routing.

feasible topology is more likely to result from lack of grooming (and can be corrected by proper grooming) than being the inevitable consequence of a high optimal value.

In all the above graphs, we also plot three other quantities (these do not vary with the number of nodes as Φ_n and Ψ_n do). The first is a lower bound computed after the fashion of the Moore bound [12], while the second is based on a per-node consideration of the lightpath endpoints, derived in [12]. Bounds of this type have been developed by consideration of general topologies and it is expected that our bounds, derived for the special case of the ring, will be tighter. In fact we see that, except in Fig. 11, we obtain only the trivial value of zero for these bounds. We also plot an easy to compute lower bound on the performance of a heuristic which is based on solving the problem optimally but using only single-hop and two-hop lightpaths. For even values of N, this heuristic is especially attractive since a wavelength assignment is always possible and thus need not be performed, a result for which we omit the proof here. A lower bound on the performance of such a heuristic is easy to obtain by considering that a traffic component from node s to node $s \oplus m$ must be electronically routed at least $\lfloor (m-1)/2 \rfloor$ times, for $m > 2$. We call this bound the *two-hop lower bound*. Since Ψ_0 and Ψ_1 are obtained from topologies that contain no lightpaths longer than two hops, the objective value of the optimal two-hop topology will by definition be equal to or less than these. This is borne out by the results. However, in each case we see that all Ψ_n values for $n > 1$ are lower than the two-hop lower bound. Thus, even the first few solutions of our framework can outperform simplistic heuristics such as the two-hop optimal topology.

In the second set of data, we present different runs in each of which results are plotted for 30 traffic matrices of the same pattern and same value of N (either 8 or 16). Fig. 13 presents the results for an ensemble of eight-node statistically uniform traffic matrices with characteristic load around 90%; the actual values of electronic routing are plotted. Fig. 14 presents an ensemble of 16-node statistically uniform traffic matrices with characteristics load around 50%; the normalized grooming effectiveness values as described above are plotted. Only the Ψ_n values which produce appreciable improvements over the previous ones are labeled for improved readability. Similarly, only the highest Φ_n value obtained is plotted. The two-hop lower bound is also plotted.

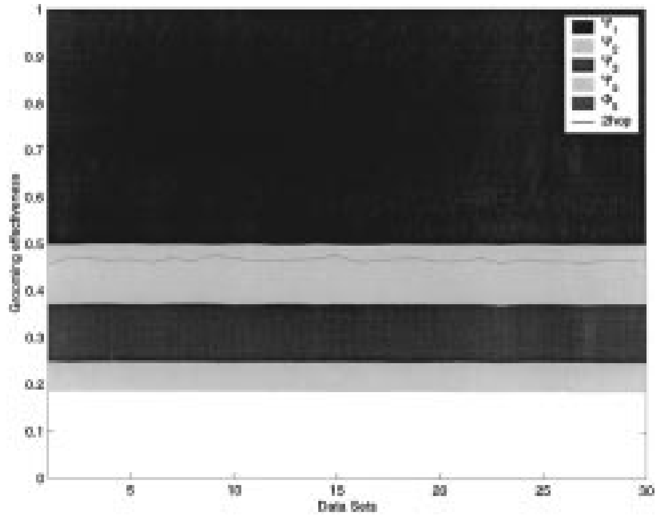

Fig. 14. Ensemble, $N = 16$, statistically uniform, 50% load, (normalized).

The ensemble results confirm the detailed results we obtained earlier. Because we have obtained bound values up to a smaller value of n, we do not see the low values of grooming effectiveness we encountered in the detailed results, but the values of the earlier bounds indicate that the same characteristic are likely to emerge. We note from the normalized graphs that Ψ_1 is likely to achieve a grooming effectiveness of around 0.5, and this is likely to be the case irrespective of the characteristic load or traffic pattern at least in the range we have varied them. Later Ψ_n values produce decreasing benefits. We note both from the detailed as well as ensemble results that several Ψ_n values before (but not including) Ψ_{N-1} are likely to produce little incremental benefit. The ensemble results also confirm that Ψ_2 is likely to outperform the two-hop optimal topology heuristic.

VI. Concluding Remarks

We have considered the problem of grooming traffic in virtual topology design for wavelength-routed optical networks. We have created a framework of bounds, both upper and lower, on the optimal value of the amount of traffic electronically routed in the network. The bounds are obtained based on the idea of decomposing the ring network a few nodes at a time. We specify the decomposition method and derive a result showing that solving the decompositions is a considerably more tractable problem than solving the complete problem. We present a method of combining these partial solutions into a sequence of bounds, both upper and lower, in which every successive bound is at least as strong as the last one.

Appendix I
Dynamic Programming Algorithm for Φ_n

Let $D_n(j,k)$ denote the largest sum of $\phi_x^{(i)}$ values obtained by partitioning the segment of the ring $\mathcal{R}$ comprised by the nodes $j, j \oplus 1, \ldots, j \oplus k \ominus 1$ into smaller segments no larger than n nodes, which we refer to as subsegments. That is, we consider the set of partitions of the nodes from j to $j \oplus k \ominus 1$ inclusive of the ring into subsegments no larger than n nodes. For each partition we sum the objective values of the decompositions obtained from the subsegments and use $D_n(j,k)$ to denote the maximum of these over all the partitions we consider. The first subsegment is constrained to start with the node j and the last subsegment is constrained to end with the node $j \oplus k \ominus 1$. For any value of j, valid values for k are $1, 2, \ldots, N$, and in addition we define $D_n(j,0) = 0$ for notational convenience. Thus $D_n(j,k)$ are partial best sums of $\phi_x^{(i)}$ values. We obtain a total bound by extending these partial sums completely around the ring as below.

By definition, $D_n(j,1) = \phi_1^{(j)}$. We now show how to obtain $D_n(j,k+1)$ given all the values $D_n(j,1), D_n(j,2), \ldots, D_n(j,k)$. All the partitions we need to consider to obtain $D_n(j,k+1)$ end with a segment which itself must end with the node $j \oplus k$. This last segment can have only between 1 and n nodes, giving us n classes of partitions we have to consider. The number of classes is smaller if $k < (n-1)$, because we are constrained to start our partition from node i. For each class, we need only consider the partition yielding the best sum of $\phi_x^{(i)}$ values and add a single $\phi_x^{(i)}$ value to it to extend the partition to node $j \oplus k$. This allows us to write the recurrence relation

$$D_n(j,k+1) = \max_{0 \le r \le \min\{n-1,k\}} \left(D_n(j, k \ominus r) + \phi_{r+1}^{(j \ominus k \ominus r)} \right). \quad (20)$$

Now $D_n(j,N)$ denotes the maximum bound that can be obtained from all partitions of the ring that are constrained to have node j as the first node of a segment. To consider other partitions in which node j is not the first of a segment, we note that node j can only be the second, third,..., nth node of a segment. Hence, it suffices to consider the n sums $D_n(j,N)$ for n successive values of j to ensure that all partitions have been considered. Any n successive nodes can be used, we assume without loss of generality that the nodes $N-n, N-n+1, \ldots, N-1$ are used, thus

$$\Phi_n = \max_{i \in \{0,\ldots,n-1\}} D(N-n+i, N). \quad (21)$$

From (20), to compute $D_n(j,k+1)$ we need at most n additions and n comparisons. To obtain $D_n(j,N)$ for any given value of j thus needs a total of $O(nN)$ time. Finally, we need to repeat this n times as per (21), thus needing a total of $O(n^2N)$ time for this algorithm.

References

[1] S. Banerjee, J. Yoo, and C. Chen, "Design of wavelength-routed optical networks for packet switched traffic," *J. Lightwave Technol.*, vol. 15, pp. 1636–1646, Sept. 1997.

[2] R. Barry and E. Modiano, "Minimizing electronic multiplexing costs for dynamic traffic in unidirectional sonet ring networks," *Proc. IEEE ICC*, vol. 3, pp. 1724–1730, 1999.

[3] A. L. Chiu and E. H. Modiano, "Reducing electronic multiplexing costs in unidirectional sonet/wdm ring networks via efficient traffic grooming," in *IEEE GLOBECOM*, vol. 1, Nov. 1998, pp. 322–327.

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

[5] R. Dutta and G. N. Rouskas, "A survey of virtual topology design algorithms for wavelength routed optical networks," *Optical Networks Mag.*, vol. 1, no. 1, pp. 73–89, Jan. 2000.

[6] R. Dutta, "Virtual Topology Design for Traffic Grooming in WDM Networks," Ph.D. dissertation, North Carolina State Univ., Raleigh, NC, Aug. 2001.

[7] O. Gerstel, P. Lin, and G. Sasaki, "Wavelength assignment in a WDM ring to minimize cost of embedded sonet rings," *Proc. IEEE INFOCOM*, pp. 94–101, 1998.
[8] O. Gerstel, R. Ramaswami, and G. Sasaki, "Cost effective traffic grooming in WDM rings," *Proc. IEEE INFOCOM*, pp. 69–77, 1998.
[9] G. Li and R. Simha, "Efficient routing to reduce the cost of add-drop multiplexers in WDM optical ring networks," in *Proc. SPIE*, vol. 3843, 1999, pp. 258–267.
[10] B. Mukherjee, S. Ramamurthy, D. Banerjee, and A. Mukherjee, "Some principles for designing a wide-area optical network," *Proc. IEEE INFOCOM*, pp. 110–119, 1994.
[11] R. Ramaswami and G. H. Sasaki, "Multiwavelength optical networks with limited wavelength conversion," *Proc. IEEE INFOCOM*, pp. 489–498, 1997.
[12] 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.
[13] ——, *Optical Networks*. San Francisco, CA: Morgan Kaufmann, 1998.
[14] X. Zhang and C. Qiao, "An effective and comprehensive solution to traffic grooming and wavelength assignment in SONET/WDM rings," in *Proc. SPIE*, vol. 3531, Nov. 1998, pp. 221–232.
[15] ——, "On scheduling all-to-all personalized connections and cost-effective designs in WDM rings," *IEEE/ACM Trans. Networking*, vol. 7, pp. 435–445, June 1999.

Rudra Dutta (S'00) was born in Calcutta, India, in 1968. After completing elementary schooling in Calcutta, he received the B.E. degree in electricial engineering from Jadavpur University, Calcutta, India, in 1991, an M.E. degree in systems science and automation from Indian Institute of Science, Bangalore, India in 1993, and the Ph.D. degree in computer science from North Carolina State University, Raleigh, in 2001.

From 1993 to 1997, he worked for IBM as a Software Developer and Programmer in various networking related projects. He is currently employed as Assistant Professor in the department of Computer Science at the North Carolina State University, Raleigh.

George N. Rouskas (S'92–M'95–SM'01) received the Diploma in electrical engineering from the National Technical University of Athens (NTUA), Athens, Greece, in 1989, and the M.S. and Ph.D. degrees in computer science from the College of Computing, Georgia Institute of Technology, Atlanta, GA, in 1991 and 1994, respectively.

He joined the Department of Computer Science, North Carolina State University in August 1994, and he has been an Associate Professor since July 1999. During the 2000–2001 academic year, he spent a sabbatical term at Vitesse Semiconductor, Morrisville, NC, and in May and June 2000 he was an Invited Professor at the University of Evry, France. His research interests include network architectures and protocols, optical networks, multicast communication, and performance evaluation.

Dr. Rouskas is a recipient of a 1997 NSF Faculty Early Career Development (CAREER) Award and a co-author of a paper that received the Best Paper Award at the 1998 SPIE conference on All-Optical Networking. He also received the 1995 Outstanding New Teacher Award from the Department of Computer Science, North Carolina State University, and the 1994 Graduate Research Assistant Award from the College of Computing, Georgia Institute of Technology. He was a co-guest editor for the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS, Special Issue on Protocols and Architectures for Next Generation Optical WDM Networks, published in October, 2000, and is on the editorial boards of the IEEE/ACM TRANSACTIONS ON NETWORKING and the *Optical Networks Magazine*. He is a member of the ACM and of the Technical Chamber of Greece.