

A Sequence of Bounds for the Problem of Minimizing Electronic Routing in Wavelength Routed Optical Rings *

Rudra Dutta George N. Rouskas

TR-2000-11

September 27, 2000

Abstract

We consider the problem of designing a virtual topology to minimize electronic routing 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.

Keywords: Virtual Topology, Logical Topology, Wavelength Division Multiplexing (WDM), Wavelength Routed Network, Lightpath, All Optical Network, Optical Ring, Bounds, Heuristics

Department of Computer Science
North Carolina State University
Raleigh, NC 27695-7534

*This work was supported by the NSF under grant NCR-9701113.

1 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 trade off some of the bandwidth for elimination of electro-optic processing delays using *wavelength routing* [5]. It has also been noted in literature that, at least in the short term, physical topologies in the forms of rings are of more interest because of available higher layer protocols such as SONET/SDH [7, 6, 13].

Of late, two concerns have clearly emerged in the treatment of optical ring networks using WDM and wavelength routing. First, it has been recognized that the the cost of network components, specifically electro-optic equipment and SONET add/drop multiplexers (ADMs), is a more meaningful metric for the network or topology rather than the number of wavelengths. Second, earlier studies of wavelength routed networks had assumed that individual traffic demands are comparable to the wavelength bandwidth, and this assumption was reflected in the design of logical topologies [4, 9, 12, 1, 5]. However, it is now evident that the independent 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 that minimize cost as measured by the amount of electro-optic equipment has recently received much attention in the literature. In [11], the problem of wavelength assignment to a given set of lightpaths is considered. The discussion focuses on how limited wavelength conversion affects the capability of the network, and specific cases are discussed with constructive proofs. Several different virtual topologies for ring networks are discussed in [7] in the light of the twin issues of traffic grooming and network cost, and characteristic metrics are derived. The study in [6] complements [7] 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 non-uniform 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

[8], a problem similar to that in [6] 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 [6] and performs rerouting of lightpaths as well, based on the Euler-trail decomposition, is presented. Several ring architectures for dynamic traffic, conforming to standard SONET self-healing rings, are reviewed in [13]. The characteristics of the architectures are discussed both for the case of limited and unlimited number of available wavelengths.

As has been noted in literature [5, 14], the problem of logical topology design is NP-hard even for a physical ring topology, 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 literature [14, 15, 3, 6, 8]. To evaluate the performance of such heuristic approaches when the optimal solution is not available, achievability bounds are useful.

We present a new framework for computing bounds for the problem of optimal traffic grooming in physical ring topologies. 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. Furthermore, the upper bounds we derive are based on actual feasible topologies, so our algorithm for obtaining the upper bounds is a heuristic for the general 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 2. In Section 3 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 4 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 5 contains numerical results and Section 6 concludes the paper.




Figure 1: (a) An N -node unidirectional ring, and (b) detail of a node with a WADM

2 Problem Formulation

We consider a unidirectional ring \mathcal{R} with N nodes¹, numbered from 0 to $N - 1$, as shown in Figure 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 link i . Each node in the ring is equipped with a wavelength add/drop multiplexer (WADM) (see Figure 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 without the need for electro-optic conversion of the signal carried on the wavelength. 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 (see the discussion below) consisting of a set of static lightpaths. These estimates may

¹We 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.

change over time, in which case a new logical topology will have to be obtained; however, we assume that such changes take place over human time scales. In this paper we do not consider the dynamic scenario in which requests for lightpaths or traffic components are received continuously during operation. In this case, dynamic grooming or lightpath allocation algorithms are required, and the performance metric of interest is some function of the blocking characteristics of the network. This latter problem has been considered elsewhere [16, 17].

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 $t^{(sd)} \in \{0, 1, 2, \dots\}$ 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 that has received considerable attention in the literature is the design of logical topologies that optimize a certain performance metric. The performance metric of interest in this work is the amount of electronic forwarding (routing) of traffic streams, since such forwarding involves electro-optic conversion and added message delay and processor load at the intermediate nodes. There is also the possibility of increased buffer requirement and queueing delay. 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 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 this paper we have chosen to concentrate on the former. A single traffic unit may not be bifurcated, but different traffic units for the same source-destination traffic component are looked upon as separate traffic and may be assigned different logical routes. 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 the minimizing of the number of SONET ADM's as in [14, 3, 6, 8]. 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 issue we mentioned above.

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 by adding up all traffic components $t^{(sd)}$ such that the path from s to d includes the physical link l (note that, because of our assumption of unidirectional

traffic flow around the ring, there is exactly one path between any pair of nodes). 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. This is reasonable because carrying the traffic component around the ring consumes more bandwidth, and is likely to require more logical hops.

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 $\{0, 1, 2, \dots, W\}$. 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 $B(l) = \{ (i, j) \mid \text{lightpath } (i, j), \text{ if it existed, would pass through link } l \}$.

With the above notation, 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 Integer Linear Problem (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 the number of wavelengths.

Given:

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

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

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

Find:

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

Subject to:

Traffic Constraints

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

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

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

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

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

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

Wavelength Constraints

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

$$\sum_k c_{ij}^{(k)} = b_{ij}, \quad \forall(i, j) \quad (8)$$

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

To minimize:

$$\min \left(\sum_{s,d,i,j \in \{0 \dots (N-1)\}} t_{ij}^{(sd)} - \sum_{s,d \in \{0 \dots (N-1)\}} t^{(sd)} \right) \quad (10)$$

In the above, 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, 5) define the total traffic on a lightpath and relate it to the lightpath count, respectively. Because of the definition of the quantities $t^{(sd)}(l)$, the constraints (1, 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 number of wavelengths available, (8) relates the wavelength indicators to the lightpath counts, and (9) ensures that no wavelength clash can occur.

The full virtual topology problem defined above can be thought of as consisting of the following three subproblems [5]²:

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:** Determine the wavelength each lightpath uses, that is assign a wavelength to each lightpath in the virtual topology so that wavelength restrictions are obeyed for each physical link.

²For the general virtual topology problem, lightpath routing over the physical topology is another subproblem [5]. In a unidirectional ring physical topology, there is only one physical routing possible for each lightpath, so this subproblem is trivial and can be ignored.

3. **Traffic Routing or Grooming Subproblem:** Groom the traffic components on the lightpaths, that is, route the traffic streams between source and destination nodes 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, limited or otherwise, may be available at the nodes [11]. A physical hop limit for lightpaths may be imposed on feasible topologies, due to physical fiber characteristics or OAM issues. The ring may be bidirectional, either with some simple routing strategy (such as shortest-path, as in [14, 6]) that allows us to consider it as two unidirectional rings, or the lightpath routing (either clockwise or counter-clockwise) may be integrated as part of the optimization process (as in [8]). 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 is not an additive function as in our formulation but some other function such as a min-max type (minimize electronic routing at the node with maximum electronic routing) or a quantified version (minimize number of wavelengths added/dropped, number of SONET ADMs, etc.) our framework can also be adapted to obtain bounds for the optimal value of the objective function.

3 Path Decomposition of a Ring Network

3.1 Definition of Decomposition

We consider a ring \mathcal{R} with N nodes labeled $0 \cdots (N - 1)$, in order, and traffic matrix T . We define a *segment* of length $n, 1 \leq n \leq N$, starting at node $i, 0 \leq i < N$, as the part of the ring \mathcal{R} that includes the n consecutive nodes $i, i \oplus 1, i \oplus 2, \dots, i \oplus (n - 1)$, 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 ring \mathcal{R} 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 . Figure 2 shows such a decomposition.

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




Figure 2: An n -node decomposition: (a) original ring \mathcal{R} with N nodes, and (b) a decomposition $\mathcal{P}_n^{(i)}$ of the ring around a segment of length n starting at node i

$$t_{\mathcal{P}_n^{(i)}}^{(sd)} = \begin{cases} t^{(sd)}, & \text{if } i \prec s \prec d \prec i \oplus (n-1) \\ \sum_{j \notin \{i, i \oplus 1, \dots, d \oplus 1\}} t^{(jd)}, & \text{if } s = S, i \preceq d \preceq i \oplus (n-1) \\ \sum_{j \notin \{s \oplus 1, s \oplus 2, \dots, i \oplus (n-1)\}} t^{(sj)}, & \text{if } d = D, i \prec s \preceq i \oplus (n-1) \\ t_{\text{pass-through}}(i, n), & \text{if } s = S, d = D \\ 0, & \text{if } \begin{cases} s = d & \text{or} \\ s = D & \text{or} \\ d = S & \text{or} \\ i \preceq d \prec s \preceq i \oplus (n-1) \end{cases} \end{cases} \quad (11)$$

where $t_{\text{pass-through}}(i, n)$ 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 , such that $s \prec d$, in the segment is the same as that in the original ring for the corresponding nodes (see the first expression in (11)). 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 the two nodes S

and D together with the links connecting them to the segment. In the decomposed network, node S acts as the source of all traffic components in the original matrix T originated at a node outside the segment and destined to any node in the segment (refer to the second expression in (11)). Node S also acts as the source for all traffic components that pass through the segment, as the fourth expression in (11) indicates. Similarly, node D is the sink for traffic originating at any node in the segment and terminating at a node outside the segment in the original ring (see the third expression in (11)), as well as for pass-through traffic. Finally, any traffic components in the original matrix T that do not originate or terminate at nodes of the segment, or do not traverse any links of the segment, are not included in the traffic matrix for the decomposition. This is captured by the last expression in (11) 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 (11), 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.

3.2 Solving Path Segments in Isolation

Consider the traffic matrix $T_{\mathcal{P}_n^{(i)}} = \begin{bmatrix} t_{\mathcal{P}_n^{(i)}}^{(sd)} \end{bmatrix}$ 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 (11). This matrix can be thought of as representing the traffic demands in a ring network consisting of nodes $S, i, \dots, i \oplus (n - 1), D$, where there is simply no traffic flowing over the link from node D to node S . Consequently, the ILP formulation we presented in Section 2 can be used to obtain a virtual topology that minimizes electronic routing for this ring with traffic matrix $T_{\mathcal{P}_n^{(i)}}$. Since the ILP formulation disallows traffic 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 n -node segment. Such a topology will only be optimal with respect to this n -node 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




Figure 3: A unidirectional n -node path network

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 original ring \mathcal{R} . Our contribution is in proving a looser 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.

Our motivation for using the path decomposition described in this section is two-fold. First, as the number n of nodes in a segment starting at some node i increases, the resulting decomposition $\mathcal{P}_n^{(i)}$ will more closely approximate the original ring. As a result, the bounds we obtain will be tighter with increasing n . Second and more importantly, a path decomposition significantly alleviates the problem of exponential growth in computational requirements for solving the original ILP for an n -node network. This result is a direct consequence of the following lemma.

Lemma 3.1 *A wavelength assignment always exists for a feasible virtual topology on a unidirectional path network, and it can be obtained in time linear in the number of links and the number of wavelengths W per link.*

Proof. Consider a path network with n nodes as shown in Figure 3. Consider any feasible virtual topology on this network, that is any topology for which the number of lightpaths crossing any link is no more than W . If there is any link with less than W lightpaths, we add as many single hop lightpaths on that link as necessary to make the number of lightpaths on that link equal to W . Now we can assign wavelengths to all the lightpaths in a straightforward manner. At the first link from node 0 to node 1, we have W lightpaths to which we can assign W wavelengths in any arbitrary fashion. We then examine these lightpaths in some particular order. Consider a lightpath assigned wavelength w on the first link from node 0 to node 1. If the lightpath continues onto the second link from node 1 to node 2, we assign the wavelength w to this lightpath on the second link. If the lightpath terminates at node 1, we know that another lightpath originates at node 1, since the load on each link is the same. We pick one such lightpath, breaking ties arbitrarily, and assign it wavelength w . We continue in this manner, assigning wavelengths to the lightpaths on each link until we reach the last node and have colored all the lightpaths. We are guaranteed to be able to reach node $n - 1$ on each of the W wavelengths, because no lightpath passes through or originates at node $n - 1$. Now we have assigned a wavelength to all the original lightpaths of the feasible virtual topology we started with and can discard the single hop lightpaths (if any) that were added to bring the number of lightpaths at each link to W . Since all feasible topologies have a wavelength assignment so does the optimal one. ■

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 2. 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, completely solving a six-node ring network using the original formulation (with wavelength assignment) requires between 60 and 90 minutes on a Sun Ultra-10 workstation. However, solving a six-node path network using the simplified formulation (no wavelength assignment) requires only a few seconds. In both cases, we used the LINGO scientific computation package which utilizes branch-and-bound algorithms to solve the ILP.

A further improvement to the ILP formulation that reduces the running time for path networks may be obtained by explicitly forbidding any lightpaths that traverse the link from D to S , that is, by fixing the values of some lightpath counts to zero:

$$b_{ij} = 0, \quad \forall b_{ij} \in B(D, S) \tag{12}$$

where i, j refer to nodes in the decomposition.

We have devised a way to further reduce the computational efforts required to solve the ILP for a path segment; the tradeoff is an increase in the complexity of the formulation. In this approach, we customize the formulation to include information about combinations of lightpath variables that can produce feasible solutions, eliminating combinations we know are not consistent by application of problem-specific logic. We can also prune the search space by a large factor by creating a space that we know contains exactly one optimal solution. The formulation has to be customized for each value of n , where n is the number of nodes in the decomposed segment. The amount of logic and extra information to be included in such a custom formulation increases rapidly with n . We have carried out such customizations for 2- and 3-node decompositions (the case for a single-node decomposition is trivial). The gain in computational effort appears to be more than offset by the complexity of the formulation beyond this point, and as such, this method was not carried out any further.

3.3 Interpretation of the Optimal Value for Decomposed Paths

We denote the optimal objective value for the decomposition $\mathcal{P}_n^{(i)}$ by $\phi_n^{(i)}$. That is, $\phi_n^{(i)}$ is the amount of electronic routing performed by the optimal virtual topology on the decomposition $\mathcal{P}_n^{(i)}$. In the decomposition, n of the $n + 2$ nodes which are the same as in the original ring see the same traffic pattern locally as

they do in the original ring, as we remarked above. 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 in any feasible virtual topology on the decomposition. Thus, the optimal objective value $\phi_n^{(i)}$ for the decomposition $\mathcal{P}_n^{(i)}$ is contributed only by the n nodes abstracted from the ring. Since the traffic pattern seen by these nodes 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. In other words, the electronic routing at this set of nodes is minimized, irrespective of how much electronic routing has to be performed at other nodes of the original ring as a result.

It might appear that as we include more and more nodes in the decomposition, the optimal solution to the original problem will be obtained when all N nodes have been represented in the decomposed network (that is, when we have a decomposition consisting of $N+2$ nodes). However, because we convert the ring to a path, some information is not taken into account at the point where we “open up” the ring when considering an N node decomposed network. As a result, solving an N -node decomposition does not provide the optimal solution for the N -node ring. This is clear from arguments related to this topic in [11, 8], for convenience we present the argument in the current context here. Consider the decomposition of an N -node segment starting at node i . Now note that the decomposed path $\mathcal{P}_N^{(i)}$ contains all nodes and links of the original ring *except* the link from node $i \ominus 1$ to node i ; this is where the ring is “opened up” to introduce the new nodes S and D . Consider a lightpath in the original ring that originates at a node $s \prec i \ominus 1$ and terminates at a node d such that $i \prec d$, and therefore uses the link $i \ominus 1$ (from node $i \ominus 1$ to node i) in the ring. Any such lightpath will be represented by two lightpaths in the decomposition $\mathcal{P}_N^{(i)}$: one from the new node S to node d , and one from the new node D to node s . Since these lightpaths are not constrained to use the same wavelength in the decomposed network, a wavelength assignment for the N -node decomposition may not be feasible for the N -node ring. Thus, the optimal solution to the decomposition may not be possible to translate to a solution to the original ring, and we still have an optimal value $\phi_N^{(i)}$ for the N -node decomposition, rather than an optimal value for the ring \mathcal{R} .

4 Bounds

In this section we describe how we can 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. The method of combination is different for lower and upper bounds. We first discuss the case where we only consider single-node decompositions, then move on to the general case where larger decompositions are available.

4.1 Bounds Based on Single-Node Decompositions

A decomposition of an N -node ring around a single node i is shown in Figure 4. The next two sections show how to obtain lower and upper bounds, respectively, on the optimal ILP solution for the ring as a

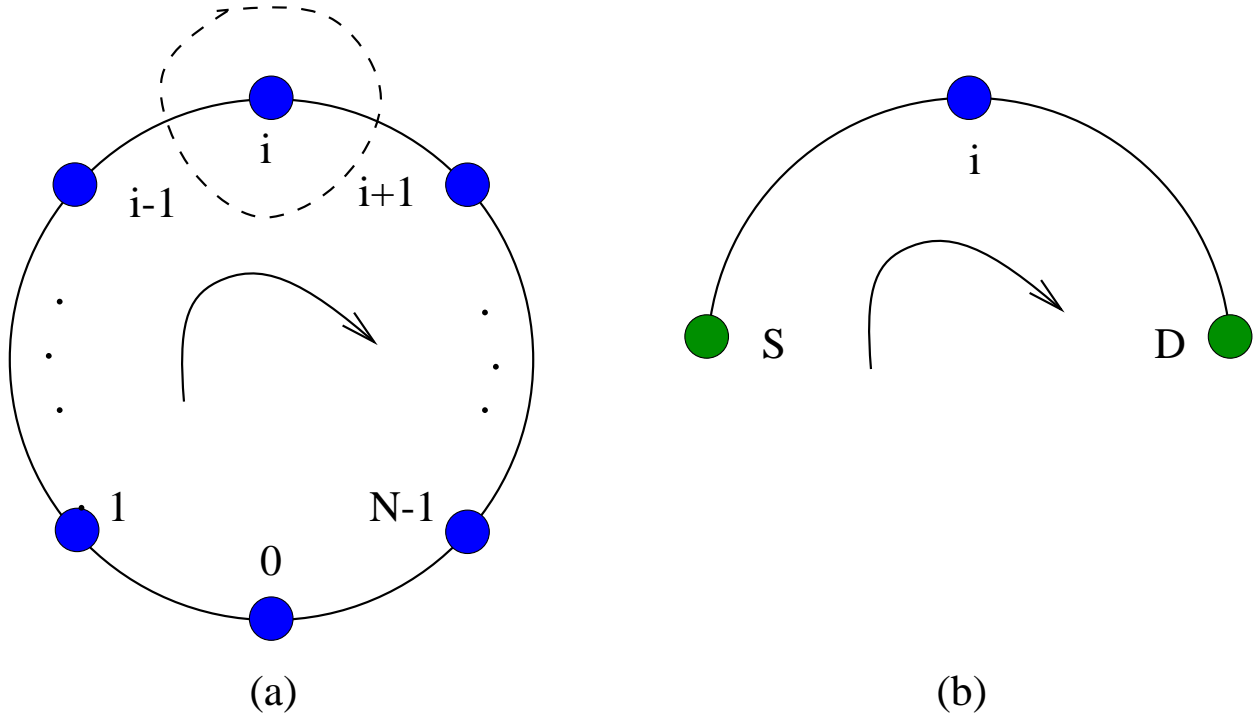


Figure 4: A single node decomposition of a ring: (a) original ring, and (b) single node decomposition $\mathcal{P}_1^{(i)}$ around node i

whole by appropriately combining the optimal ILP solutions $\phi_1^{(i)}$ for the possible single-node decompositions $\mathcal{P}_1^{(i)}$, $i = 0, \dots, (N - 1)$.

4.1.1 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. 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 feasible virtual topology, and in particular, for the optimal 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 for all nodes in a feasible virtual topology. We call this lower bound Φ_1 :

$$\Phi_1 = \sum_{i=0}^{N-1} \phi_1^{(i)} \quad (13)$$

The quantity Φ_1 is a lower bound on the objective value (total electronic routing) for any feasible virtual

topology, and in particular, for the optimal virtual topology for the ring \mathcal{R} .

4.1.2 Upper Bound

We first note that the value of the objective function for any feasible virtual topology sets an upper bound on the optimal value, since it corresponds to an actual solution and the optimal solution can only be better than or as good as this solution. Thus, we consider different achievable topologies and we obtain upper bounds from them.

First we consider an upper bound we can obtain directly from the traffic matrix, without recourse to decompositions. This bound corresponds to the simplest virtual topology possible, namely, the 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 and each lightpath spans exactly one physical link. (This type of topology is called a *PPWDM ring* in [7].) In such a topology, each node i performs the maximum possible amount of electronic routing, which we denote by $\psi^{(i)}$. Quantities $\psi^{(i)}, i = 0, \dots, (N-1)$, 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:

$$\Psi_0 = \sum_{i=0}^{N-1} \psi^{(i)} \quad (14)$$

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 consider how we might utilize single-node decompositions to improve upon the no-wavelength-routing topology and hence obtain a tighter upper bound. To this end, let us refer again to Figure 4(b), which shows a single-node 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 by node S and terminated by node D . Assuming for the moment that nodes S and D are actual nodes in the ring, this is equivalent to saying that no lightpaths are optically routed at either node S or node D , and, consequently, both nodes have to perform the maximum amount of electronic routing possible.

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, and according to our discussion in the last paragraph, nodes S and D can be viewed as concentrator nodes in the single node decomposition. Carrying this line of thought one step further, we are led to consider a new 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 Figure 5, where the even-numbered nodes are concentrator nodes.

We now note that, based on which nodes in the ring we select to be the concentrator nodes, we obtain

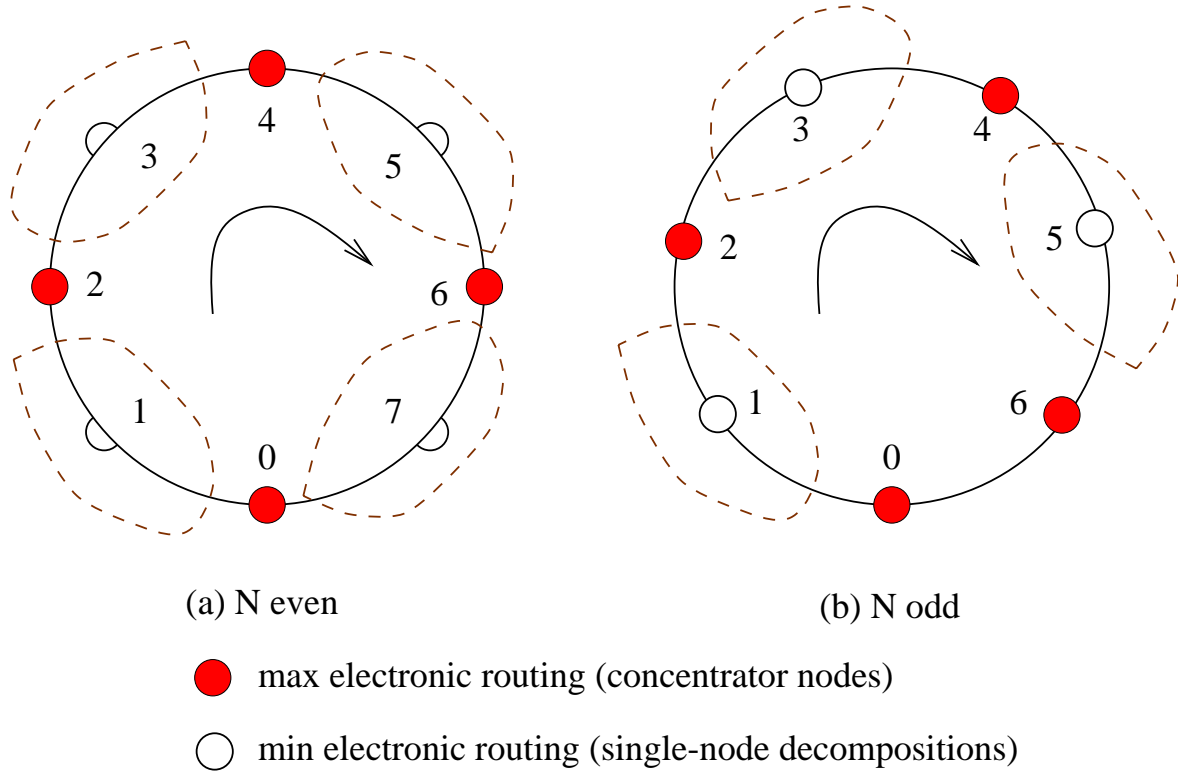


Figure 5: Virtual topology with alternating single-node decompositions and concentrator nodes

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 will have two concentrator nodes next to each other at one position in the ring, as illustrated in Figure 5(b). Since there 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:

$$\Psi_1 = \min_{j \in [0, N-1]} \left(\sum_{k \in \{0, 2, 4, \dots, 2\lfloor (N-2)/2 \rfloor\}} \phi_1^{(j+k)} + \sum_{k \notin \{0, 2, 4, \dots, 2\lfloor (N-2)/2 \rfloor\}} \psi^{(j+k)} \right) \quad (15)$$

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

$$\Psi_1 \leq \Psi_0 \quad (16)$$

4.2 Bounds Based on Larger Decompositions

In this section we consider how we may combine decompositions containing more than a single node from the ring to obtain a sequence of bounds similar to those obtained in the last section.

4.2.1 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 $\phi_1^{(i \oplus 1)}$ 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 $\phi_1^{(i \oplus 2)}$, 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 find that we can add the $\phi_n^{(i)}$ values for any set of decompositions that involve segments that are disjoint in the ring, 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:

Lemma 4.1 *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. Let $l_k, l_k \leq n$, denote the length (number of nodes) of segment $k, k = 1, \dots, |\sigma_n|$, and let i_k denote its starting node. The quantity*

$$\Phi(\sigma_n) = \sum_{k=1}^{|\sigma_n|} \phi_{l_k}^{(i_k)} \quad (17)$$

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:

$$\Phi_n = \max\{\Phi(\sigma_n)\} \quad (18)$$

where the maximum is taken over all partitions of the ring which contain segments with n or less nodes. Figure 6 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 the definition (18), 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 Equation 18, this allows us to draw the general conclusion that:

$$\Phi_{n+1} \geq \Phi_n \quad \forall n \in \{1 \dots (N-1)\} \quad (19)$$

As a result, the sequence $\Phi_1, \Phi_2, \dots, \Phi_N$, is a strong sequence of bounds in which each is at least as tight as the previous one. We note that our definition of Φ_1 in Section 4.1.1 is consistent with our general definition




Figure 6: Partitions of the nodes of a ring into (a) segments of no more than 2 nodes, and (b) segments of no more than 3 nodes

above. We were able to express Φ_1 in a simpler and more explicit form because there is only one possible partition of the nodes of a ring into single-node segments.

We discuss some details of the computation of Φ_n after we describe the upper bounds in the next section.

4.2.2 Upper Bound

It is now straightforward to obtain a strong sequence of upper bounds along the same lines. In Section 4.1.2 we obtained the upper bound Ψ_1 by creating a topology in which single-node segments (where electronic routing is minimized) alternate with concentrator nodes. Similarly, we now 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 (15) 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 . Figure 7 shows two ways we can partition a ring using no larger than 3-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 4.1.2 are consistent with this framework.




Figure 7: Two partitions of a ring into alternating concentrator nodes and segments of no more than 3 nodes

We also note that since every decomposed segment has to alternate with a concentrator node, we can only use up to $N - 1$ node decompositions, and cannot use N -node decompositions as for the lower bounds. As before, the set of all topologies we consider in obtaining Ψ_{n+1} is a superset of the set of all topologies we consider in obtaining Ψ_n , therefore we may assert that:

$$\Psi_{n+1} \leq \Psi_n \quad \forall n \in [0, N - 2] \quad (20)$$

giving us a strong sequence of upper bounds.

Because the bounds $\{\Psi_n\}$ are based on actual 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 5 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 [7, 3, 2, 13], because it considers all topologies with a single concentrator node (which is equivalent to a hub node). For a similar reason, Ψ_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.

4.2.3 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:

$$\zeta^{(j)} = \min_{i=0}^{N-1} (\psi^{(i)} - \phi_1^{(i)}) \quad (21)$$

Consider the topology in which node j is a concentrator node and the rest of the nodes follow the optimal topology for the $(N-1)$ -node decomposition $\mathcal{P}_{N-1}^{(j\oplus 1)}$. The amount of electronic routing performed by this topology is an upper bound on the optimal objective value for the ring. We denote it by Ψ' and note that $\Psi' = \psi^{(j)} + \phi_{N-1}^{(j\oplus 1)}$. Since this is a topology using a segment of $N-1$ nodes alternating with a concentrator node, we must have

$$\Psi' \geq \Psi_{N-1} \quad (22)$$

Now we consider the partition of the ring into the single node j and the $(N-1)$ -node segment comprising all the other nodes again, but this time with a view to obtaining a lower bound. We denote this lower bound by Φ' and it is found by simply adding the objective values for the two corresponding decompositions, so that $\Phi' = \phi_1^{(j)} + \phi_{N-1}^{(j\oplus 1)}$. Once again, this is one of the partitions in the set of all partitions of the nodes of the ring into segments of $N-1$ nodes or less, so that we have

$$\Phi' \leq \Phi_{N-1} \quad (23)$$

However, from the definition of Ψ' and Φ' above, we know that

$$\Psi' - \Phi' = \zeta^{(j)} \quad (24)$$

Combining the above results (22, 23, 24), and considering the definition of $\zeta^{(j)}$, we can assert that:

$$\Psi_{N-1} - \Phi_{N-1} \leq \min_{i=0}^{N-1} (\psi^{(i)} - \phi_1^{(i)}) \quad (25)$$

Of course, depending on the value of N and the computational power available, it may or may not be practical to compute Ψ_{N-1} and Φ_{N-1} ; however, this is the theoretical limitation on the tightness of the framework of bounds we present.

4.2.4 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, and
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 $i \in \{0 \cdots (N-1)\}$, and all values of $n \in \{1 \cdots x\}$. 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 increase 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 the bound Φ_n , we note that a straightforward approach would require us to enumerate all possible 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 $\phi_x^{(i)}$, $x \in \{1 \cdots n\}$, values are available), the number of possible partitions increases with n . The total number of partitions is maximum when $n = N$, and this number is easily seen to be $2^N - 1$ (because each link of the ring can be broken to form partitions or not, with the single case where no link is broken being excluded). For smaller values of n , the total number is smaller but still very large. We note that assuming node i is the first of a segment in the partition (and thus excluding some partitions), the total number of segments in such a partition must always be at least $\lceil N/n \rceil$ for a given value of n , and since each segment can consist of $1, 2, \dots, n$ nodes the total number of such partitions is at least $n^{\lceil N/n \rceil}$. Considering $n = 2$, we can set a lower limit on this value, and thus say that for $2 \leq n \leq N$, the total number of partitions is between $2^{\lceil N/2 \rceil}$ and $2^N - 1$, and is thus exponential in N . Thus the incremental number of partitions to consider for a given value of n is also exponential in N in the worst case. 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 B, and requires $O(n^2N)$ time to find Φ_n given all $\phi_x^{(i)}$ values for $x = 1, \dots, n$. 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 improvement also by using the properties of $\phi_n^{(i)}$ values formalized in the following lemma. These properties introduce a constant factor of improvement to the dynamic programming algorithm described in Appendix B.

Lemma 4.2 *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 4.1 *An x -node decomposition yields at least as large an objective value for the decomposed network as the sum of objective values of any combination of smaller decompositions it can be partitioned into.*

That is, $\phi_x^{(i)} \geq \phi_{y_1}^{(i)} + \phi_{y_2}^{(i+y_1)} + \phi_{y_3}^{(i+y_1+y_2)} + \dots + \phi_{y_n}^{(i+y_1+y_2+\dots+y_{n-1})}$, if $x = y_1 + y_2 + \dots + y_n$.

Proof of Lemma 4.2. Let us first consider the special case $x = y = 1$. Consider a two-node decomposition for nodes i and $r = i \oplus 1$ in the original ring network \mathcal{R} . This yields a value of $\phi_2^{(i)}$ for the node pair decomposed. Now consider a new ring network \mathcal{R}' with four nodes, which has the same structure as the decomposed network above and has a traffic matrix identical to the traffic matrix $T_{\mathcal{P}_2^{(i)}}$ in the above decomposition. It is easy to see that the optimum value of total electronic routing for this new network is $\phi_2^{(i)}$, since this is the best value for nodes i and r , and nodes S and D do not have any pass-through traffic and hence cannot route any traffic electronically. Now consider a single-node decomposition of this new network \mathcal{R}' , yielding $\phi_1^{(S)'}$, $\phi_1^{(i)'}$, $\phi_1^{(r)'}$ and $\phi_1^{(D)'}$ as the best case electronic routing for the four nodes. Once again, since nodes S and D have no pass-through traffic, $\phi_1^{(S)'}$ and $\phi_1^{(D)'}$ are bound to be zero. Now the single-node decomposition matrix for node i , $T_{\mathcal{P}_1^{(i)}}$, is the same no matter whether the decomposition is made from the original network or the new network. Thus, the single node best case electronic routings of the node i is the same in either case, in other words $\phi_1^{(i)'} = \phi_1^{(i)}$. Similarly, $\phi_1^{(r)'} = \phi_1^{(r)}$. Since the sum $\phi_1^{(S)'} + \phi_1^{(i)'} + \phi_1^{(r)'} + \phi_1^{(D)'}$ is known to form a lower bound on the optimum value of the electronic routing for the new network, we can say that $\phi_1^{(i)} + \phi_1^{(r)} \leq \phi_2^{(i)}$.

Although in the above argument we have used special values of 1 for each of x and y , and hence 2 for $x + y$, the arguments retain their validity if any other values are used. Thus, the lemma is true. ■

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 $y_1 + y_2 \leq x$, because the partition we obtain by replacing these two segments by a single $(y_1 + y_2)$ -node segment must yield a higher bound, and we are only 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 found that computing the $\phi_n^{(i)}$ values took minutes and hours for increasing n , while combining them to form the $\{\Phi_n\}$ bounds took milliseconds. We conclude that the limiting factor in determining how many of these 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.

Similar observations apply to the sequence of bounds $\{\Psi_n\}$ as the value of n increases, and a similar algorithm can be used to compute efficiently each bound Ψ_n given the $\phi_n^{(i)}$ values.

5 Numerical Results

In this section, we present the results of using our framework for different traffic matrices. First, we discuss the criteria we used to construct traffic matrices and to group results by these criteria; the results are presented next.

5.1 Traffic Patterns

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. In other words, for a symmetric traffic matrix T we have

$$t^{(sd)} = t^{(s \oplus x, d \oplus x)} \quad \forall i, j \in \{0 \cdots (N - 1)\}, \forall 1 \leq x \leq N - 1 \quad (26)$$

This type of traffic pattern is of interest since the traffic pattern looks similar from different nodes on the ring. If the above relation holds exactly, the optimal topology is likely to be a regular topology. Rather than confine our attention to those traffic matrices in which (26) holds strictly, we consider the general case where traffic components of the form $t^{(s \oplus x, d \oplus x)}$ for a given s and d are all drawn from the same distribution. If the variance of this distribution is zero, then (26) holds exactly, and we call the resulting traffic pattern *strictly symmetric*, otherwise we call it *statistically symmetric*. If a traffic matrix is highly asymmetric, so that the traffic patterns seen by different nodes of the ring are very different, then the optimal topology is likely to perform well in the sections of the ring where there is less congestion and poorly in the sections of high congestion. This is also likely to be the case for any feasible topology, in other words, the difference between the best and worst may be comparatively less. For this reason, we have chosen to concentrate on statistically symmetric traffic matrices for all our results.

Having said that the traffic pattern looks similar from each node, we turn our attention to what it looks like from a given node. We consider three simple cases which are represented schematically in Figure 8. First, the traffic from a given node i to all other nodes may be the same, we refer to this as *uniform* traffic. For a strictly symmetric matrix, this results in each element of the matrix having the same value (other than the diagonal elements, which must be zero). 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 in Figure 8.

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. Given a traffic matrix, we can compute the total traffic flowing over a given link of the ring in a straightforward manner. For the problem instance to be feasible, this traffic must be less than or equal to WC , where W is the number of wavelengths and C is the capacity in units of traffic of each wavelength. If the matrix is statistically symmetric, the load on each link will all be close to some value, because the traffic pattern is the same looking from any node or link.

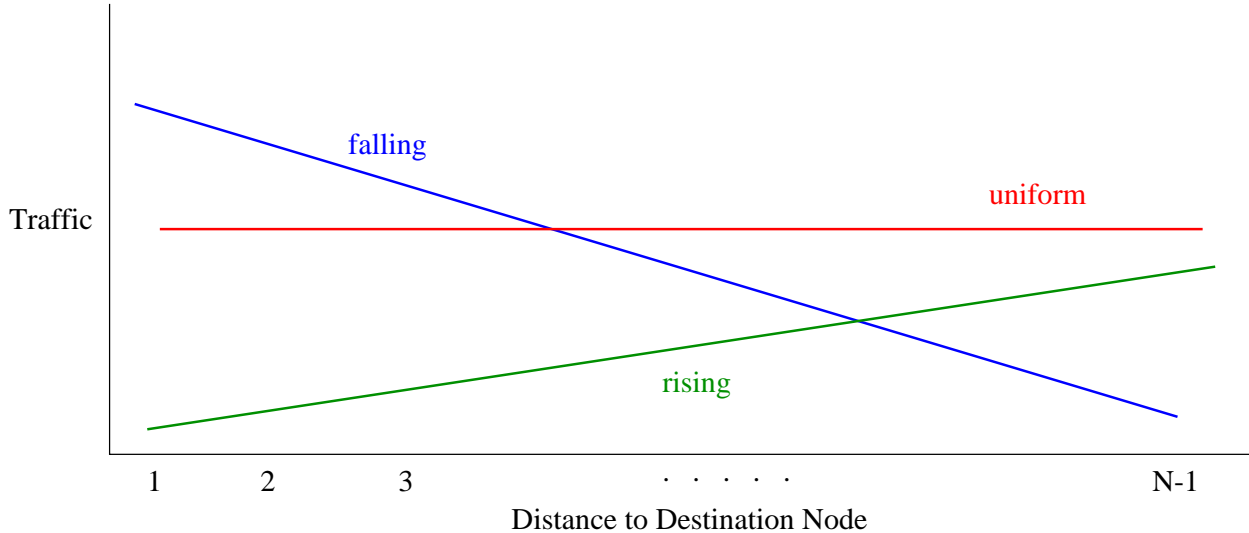


Figure 8: Different traffic patterns.

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 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.

5.2 Results

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 around 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%.

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. We concentrate on the quantity Ψ_0 , which denotes the amount of electronic routing performed by a topology that does not employ optical forwarding at all. This is often actually used in networks at transitory stages in which the fiber medium is employed with WDM, but no wavelength routing is employed [5, 10]. We can consider this case to correspond to *no grooming*, that is, no effort has been made to groom individual traffic components into lightpaths. 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 and themselves represent the grooming effectiveness of the heuristic solutions created. The values $\{\Phi_n\}$ represent lower bounds on the grooming effectiveness that can be reached in principle, that is, 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 upto 7, for $N = 8, 16$, for the two characteristic loads and statistically uniform, rising and falling patterns. Figures 9 - 14 present the data for 8-node rings and Figures 15 - 18 present the data for 16-node rings. For $N = 8$, the rising and falling matrices were so generated that the lowest traffic component in a row (traffic to nearest node for rising pattern, traffic to farthest node for falling pattern) are close to zero. For $N = 16$, we present data only for two types of falling traffic patterns. One is a statistically falling matrix in which the traffic falls to zero at the farthest node as with the 8-node case. We also have a data set in which the traffic from node s falls to zero at node $s \oplus 8$, that is, halfway around the ring. Such a 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*, as we mentioned in Section 2. No data set for a rising pattern or uniform pattern is included for $N = 16$.

Figures 9, 10 show the results for 8-node rings, statistically uniform traffic, 50% and 90% characteristic load respectively. Figures 11, 12 similarly show the results for statistically falling traffic, 50% and 90% load, and figures 13, 14 the results for statistically falling traffic, 50% and 90% load, respectively. 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, 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. The growth of Φ_n is comparatively less. Figures 15, 16 show detailed results for 16-node rings, statistically falling traffic, falling to zero at the end of the ring, 50% and 90% characteristic load respectively. Figures 17, 18 show similar results for statistically falling traffic, falling to zero at a point halfway around the ring. These show similar characteristics as the results for 8-node rings. Even though the largest value of n for which Ψ_n is obtained, 7, is now less than half of N , we observe that Ψ_7 once again represents grooming effectiveness between 0.1 and 0.2 in all the cases. Thus we generally observe that we get good grooming effectiveness and that the lower bounds are comparatively less in magnitude. This validates our approach of describing the values of the bounds with respect to the no-optical-forwarding case rather than the optimal value, because it indicates that a high value of electronic routing for some 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 this detailed section, we also plot three other quantities (these do not vary with number of nodes like Φ_n and Ψ_n). The first is a lower bound computed after the fashion of the Moore bound found in literature (for example, [12]) by considering the number of lightpath endpoints available to traffic, while the second is based on a per-node consideration of the lightpath endpoints, derived after the fashion of [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, in most cases, we obtain only the




Figure 9: Detailed results for $N = 8$, Statistically uniform pattern, 50% load




Figure 10: Detailed results for $N = 8$, Statistically uniform pattern, 90% load




Figure 11: Detailed results for $N = 8$, Statistically falling pattern, 50% load




Figure 12: Detailed results for $N = 8$, Statistically falling pattern, 90% load




Figure 13: Detailed results for $N = 8$, Statistically rising pattern, 50% load




Figure 14: Detailed results for $N = 8$, Statistically rising pattern, 90% load




Figure 15: Detailed results for $N = 16$, Statistically falling pattern, 50% load, falling to end




Figure 16: Detailed results for $N = 16$, Statistically falling pattern, 90% load, falling to end




Figure 17: Detailed results for $N = 16$, Statistically falling pattern, 50% load, falling to $N/2$




Figure 18: Detailed results for $N = 16$, Statistically falling pattern, 90% load, falling to $N/2$




Figure 19: Ensemble results for $N = 8$, Statistically uniform pattern, 90% load, electronic routing

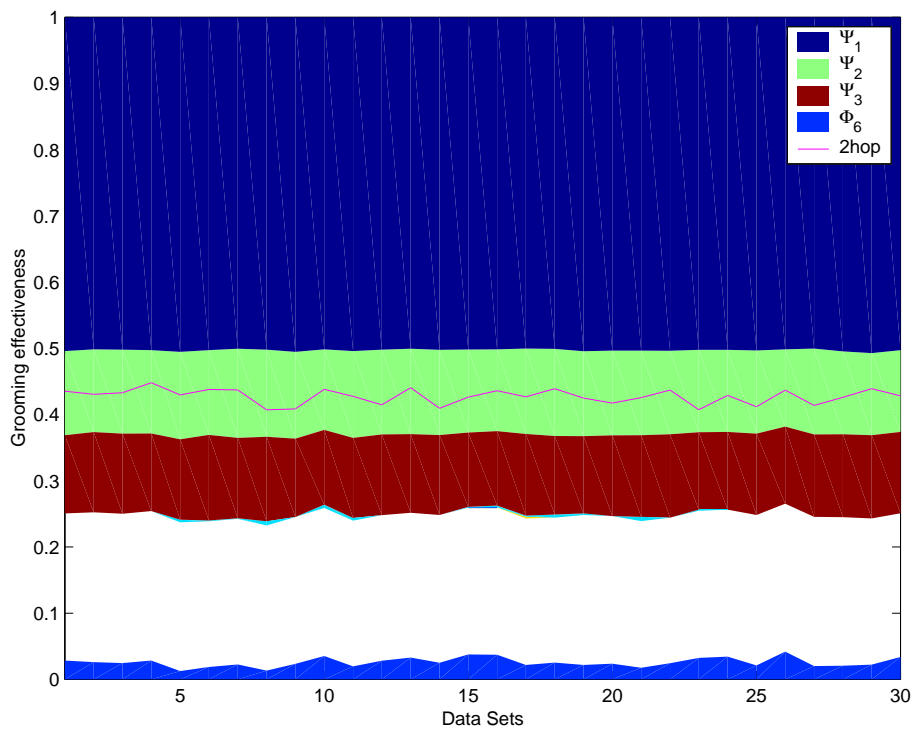


Figure 20: Ensemble results for $N = 8$, Statistically uniform pattern, 90% load, (normalized)




Figure 21: Ensemble results for $N = 8$, Statistically falling pattern, 50% load, (normalized)




Figure 22: Ensemble results for $N = 8$, Statistically falling pattern, 90% load, (normalized)




Figure 23: Ensemble results for $N = 8$, Statistically rising (step) pattern, 50% load, (normalized)




Figure 24: Ensemble results for $N = 16$, Statistically uniform pattern, 50% load, (normalized)

trivial value of 0 for these bounds. However, for the 16-node ring in the case where traffic falls to zero at the end of the ring, (figures 15, 16) the first bound has a comparable value to the largest Φ_n we have obtained (in one case it is nearly equal and in the other case distinctly larger).

Finally, we plot an easy to compute lower bound on the performance of a simple heuristic which is based on solving the problem optimally but using only single-hop and two-hop lightpaths. For even values of N , the simplicity of 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 formal statement and 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 *2-hop lower bound*. Since Ψ_0 and Ψ_1 are obtained from topologies that can by definition 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 2-hop lower bound. Thus even the first few solutions provided by 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). For each matrix, Φ_n and Ψ_n are plotted for values of n upto a maximum, the 2-hop lower bound is also plotted. Figures 19 - 24 present these results. Only the Ψ_n values which produce appreciable improvements over the previous ones are plotted for improved readability. Similarly, only the highest Φ_n value obtained is plotted. The 2-hop lower bound is also plotted.

Figures 19,20 both plot the results for the same ensemble, 8-node rings with statistically uniform traffic of 90% characteristic load. Figure 19 shows the actual values of electronic routing while figure 20 expresses the same data normalized to produce grooming effectiveness values as described above. All the rest of the figures use the normalized values. Figures 21, 22 represent ensembles of 8-node rings, statistically falling traffic matrices of 50% and 90% characteristic load values respectively. Figure 23 shows the results for an ensemble of 8-node rings, 50% load; the traffic pattern is statistically a step rise, traffic to near nodes being the same on average, while the traffic to further nodes is again the same but a higher value. Figure 24 represents an ensemble of 16-node rings, statistically uniform traffic with 50% load. For the last two ensembles the highest value of n for which Φ_n and Ψ_n are plotted is 5, for the earlier one it is 6.

The ensemble results confirm the detailed results we obtained earlier. Because we have obtained bound values upto 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 for most cases.

6 Concluding Remarks

We have considered the problem of virtual topology design for wavelength routed optical networks in the special case where the physical topology is a unidirectional ring. The problem of interest is to groom lower speed traffic streams into the wavelengths available, that is minimize electronic routing of traffic. In this context, 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. We derive a result showing how close the final upper and lower bounds are guaranteed to be to each other. An algorithm is developed which enables us to combine the decompositions into bounds without facing a combinatorial explosion.

We present numerical results of the computation of these bounds by computing the bounds on different families of traffic matrices. For small rings, we can compute nearly all the bounds in the framework. Numerical results indicate that in these cases the expectations from theoretical considerations are fulfilled. For larger rings, the sequence cannot be computed to the end and we are limited by the availability of computing power in how far we can compute this sequence. The numerical results show that we can get good results in either case.

The framework can be adapted to other formulations of the problem on the ring network, such as bidirectional rings or quantized objective functions. The bounds help us in evaluating heuristics which must be used in practical situations because of the NP-hard nature of the problem. The upper bounds are based on constructing actual feasible topologies on the network, and thus also provide us with a sequence of increasingly good heuristics. Thus, we feel this is a useful framework in a broad context of ring problems for wavelength routing optical networks.

References

- [1] S. Banerjee, J. Yoo, and C. Chen. Design of wavelength-routed optical networks for packet switched traffic. *Journal of Lightwave Technology*, 15(9):1636–1646, Sept 1997.
- [2] R. Berry and E. Modiano. Minimizing electronic multiplexing costs for dynamic traffic in unidirectional sonet ring networks. *IEEE International Conference on Communications*, 3:1724–1730, 1999.
- [3] A.L. Chiu and E.H. Modiano. Reducing electronic multiplexing costs in unidirectional sonet/wdm ring networks via efficient traffic grooming. *IEEE Globecom*, 1:322–327, Nov 1998.
- [4] I. Chlamtac, A. Ganz, and G. Karmi. Lightpath communications: An approach to high bandwidth optical wans. *IEEE Trans. Communications*, 40(7):1171–1182, Jul 1992.
- [5] R. Dutta and G.N. Rouskas. A survey of virtual topology design algorithms for wavelength routed optical networks. *Optical Networks Magazine*, 1(1):73–89, Jan 2000.
- [6] O. Gerstel, P. Lin, and G. Sasaki. Wavelength assignment in a wdm ring to minimize cost of embedded sonet rings. *Proc. IEEE INFOCOM*, pages 94–101, 1998.
- [7] O. Gerstel, R. Ramaswami, and G. Sasaki. Cost effective traffic grooming in WDM rings. *Proc. IEEE INFOCOM*, pages 69–77, 1998.
- [8] G. Li and R. Simha. Efficient routing to reduce the cost of add-drop multiplexers in wdm optical ring networks. *Proceedings of the SPIE*, 3843:258–267, 1999.
- [9] B. Mukherjee, S. Ramamurthy, D. Banerjee, and A. Mukherjee. Some principles for designing a wide-area optical network. *Proc. IEEE INFOCOM*, pages 110–119, 1994.
- [10] B. Mukherjee, S. Ramamurthy, D. Banerjee, and A. Mukherjee. Some principles for designing a wide-area optical network. *IEEE/ACM Trans. Networking*, 4(5):684–696, Oct 1996.
- [11] R. Ramaswami and G.H. Sasaki. Multiwavelength optical networks with limited wavelength conversion. *Proc. IEEE INFOCOM*, pages 489–498, 1997.
- [12] R. Ramaswami and K.N. Sivarajan. Design of logical topologies for wavelength-routed optical networks. *IEEE JSAC*, 14(5):840–851, June 1996.
- [13] G.H. Sasaki. Minimal cost wdm sonet ring designs. *Optical Networks Magazine*, (to appear).
- [14] X. Zhang and C. Qiao. An effective and comprehensive solution to traffic grooming and wavelength assignment in sonet/wdm rings. *SPIE Conference on All-Optical Networking: Architecture, Control and management Issues*, 3531:221–232, Nov 1998.
- [15] X. Zhang and C. Qiao. On scheduling all-to-all personalized connections and cost-effective designs in wdm rings. *IEEE/ACM Transactions on Networking*, 7(3):435–445, Jun 1999.

- [16] Y. Zhu, G.N. Rouskas, and H.G. Perros. Blocking in wavelength routing networks, part i: The single path case. *Proceedings of IEEE Infocom '99, New York, NY*, pages 321–328, March 1999.
- [17] Y. Zhu, G.N. Rouskas, and H.G. Perros. Blocking in wavelength routing networks, part ii: Mesh topologies. *Proceedings of the Sixteenth International Teletraffic Congress (ITC 16), Edinburgh, UK*, pages 1321–1330, Jun 1999.

A Note on Obtaining a Strong Sequence of Bounds

When deriving total upper and lower bounds from single node decompositions in Section 4, we were able to combine them in a straightforward manner without any combinatorial considerations. This is because we had only one kind of decomposed segments at hand, and simply traversed the nodes in order with such segments, alternating with concentrator nodes in the case of the upper bound. It might appear that we could have followed this approach to obtain successively stronger bounds (even if each bound is not as strong as with the combinatorial approach), without recourse to the combinatorial complications, for each value of x resorting to segments with less than x nodes only to complete the ring. We note in this appendix that such an approach would not give us a strong sequence of bounds.

We denote a total lower bound formed using only n node segment decompositions (with at most one smaller segment) as Φ'_n , and a total upper bound formed similarly as Ψ'_n . When two node decomposition values are available, we can partition the ring into segments of two adjacent nodes. If N is even, there are two ways we can form this partition ($0 - 1, 2 - 3, \dots$ or $N - 0, 1 - 2, 3 - 4, \dots$). If N is odd, we can form $\lfloor N/2 \rfloor$ two node decompositions, and use a single node decomposition for the last node. Depending on which node we choose as the singleton one, there are N possible ways to partition the ring. In general, if n -node decomposition values are available, there are n possible partitions of the ring into n -node decompositions if N is a multiple of n , and N possible partitions into $\lfloor N/n \rfloor$ decompositions of n nodes each and one smaller one if N is not a multiple of n . In each case we can sum all the corresponding $\phi_n^{(i)}$ values to obtain a total bound for the ring. For any n , Φ'_n is the largest sum of $\phi_n^{(i)}$ from these possible decompositions, and Ψ'_n is similarly the lowest sum obtained from only n -node segments alternating with concentrator nodes, with at most one smaller segment.

If x is a multiple of y , and N is a multiple of x , we are guaranteed that Φ'_x is at least as strong as Φ'_y . This follows from Lemma 4.2 because we can take the partition of the ring into y -node sets that result in Φ'_y , then obtain a partition into x -node sets simply by combining x/y adjacent y -node sets to obtain each x -node set. Now the sum must be at least as large, because $\phi_x^{(i)}$ for each x -node set must be at least as large as the sum of $\phi_y^{(i)}$ values for the y -node sets it contains. For upper bounds, similar considerations exist except that the characterization of the relationships between bounds is a little more complex. For a bound derived from x -node decompositions to be stronger than one derived from y -node decompositions, x must exactly include a number of y -node sets *as well as* the concentrator nodes between them, and this can be translated into a criterion for the total upper bound to be at least as strong.

But this type of argument cannot be applied in the general case where we move from n -node decompositions to $(n + 1)$ -node decompositions, because $n + 1$ is in general not a multiple of n , and the above criteria for upper bounds is generally not fulfilled. In fact lower and upper bounds formed simply as above do not possess the desirable characteristic that each successive bound is at least as strong as the previous one. For example, consider the two traffic matrices:

$$T_1 = \begin{bmatrix} 0 & 45 & 30 & 20 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 35 & 20 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 25 & 47 & 31 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 45 & 21 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 37 & 55 & 10 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 43 & 10 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 35 & 51 & 20 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 45 & 10 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 53 & 41 & 20 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 45 & 20 \\ 20 & 20 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 40 \\ 30 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

$$T_2 = \begin{bmatrix} 0 & 45 & 30 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 35 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 25 & 47 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 45 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 37 & 55 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 43 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 35 & 51 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 45 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 53 & 41 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 45 & 0 \\ 20 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 40 \\ 30 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

In both cases, $N = 12$, $W = 10$, $C = 16$. For T_1 , we have $\Phi'_3 < \Phi'_2$, and for T_2 we have $\Psi'_2 > \Psi'_1$. Specifically, for T_1 , $\Phi'_2 = 24$ and $\Phi'_3 = 19$; for T_2 , $\Psi'_1 = 0$ and $\Psi'_2 = 75$.

Thus we do not necessarily get a progression of bounds in which each is at least as strong as the last when we use successively larger decompositions, even though we utilize more information in computing the successively larger decompositions. The reason for this is that in deriving Φ'_n and Ψ'_n we constrain ourselves to using *only* the n -node decompositions. This means that in computing Φ'_x and Ψ'_x we are forced to ignore any structure in the given traffic matrix which exists at a scale y smaller than x , if x is not a multiple of y , as in the above examples. However, we can get a strong progression of bounds in which each is at least as strong as the last if we merely increase the scope of the cases we consider when forming the bounds as we use successively larger decompositions, rather than constraining ourselves as above. This is the reason for adopting the combinatorial approach.

B Combining Partitions to Obtain Best Case

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 $j, j \oplus 1, \dots, j \oplus k \ominus 1$ 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 $1, 2, \dots, N$, 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 consider how to obtain $D_n(j, k + 1)$ given all the values $D_n(j, 1), D_n(j, 2), \dots, D_n(j, k)$. We note that all 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 be comprised only of 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:

$$D_n(j, k + 1) = \max_{0 \leq r \leq \min\{n-1, k\}} \left(D_n(j, k \ominus r) + \phi_{r+1}^{(j \oplus k \ominus r)} \right) \quad (27)$$

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, \dots n -th of a segment. Hence it suffices to consider the n sums $D_n(j, N)$ for n successive values of j to be sure that all possible partitions have been considered. Any n successive nodes can be used, we assume without loss of generality that the nodes $N - n, N - n + 1, \dots, N - 1$ are used, and make the statement:

$$\Phi_n = \max_{i \in \{0 \dots (n-1)\}} D(N - n + i, N) \quad (28)$$

To determine the complexity of this method to determine Φ_n , we first assume that we already have the values $\phi_x^{(i)}$ for $x \in \{1 \dots n\}$ and $i \in \{0 \dots (N - 1)\}$. In that case, to find $D_n(j, k + 1)$, we need at most n additions and n comparisons as evident from (27). 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 (28), thus needing a total of $O(n^2 N)$ time for this algorithm.