On Optimal Traffic Grooming in WDM Rings*

Rudra Dutta, George N. Rouskas
Department of Computer Science, North Carolina State University, Raleigh, NC 27695-7534
{rdutta,rouskas}@csc.ncsu.edu

ABSTRACT

We consider the problem of designing a virtual topology to minimize electronic routing, that is, grooming traffic, in wavelength routed optical rings. We present a new framework consisting 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. Our approach can be applied to many virtual topology problems on rings. The upper bounds we obtain also provide a useful series of heuristic solutions.

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* [4]. 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 [6, 5].

Two concerns have recently emerged in this area: 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, and 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. These two issues give rise to the concept of *traffic grooming* [6, 5, 11, 12, 2, 1, 7] 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.

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

ACM SIGMETRICS 2001 6/01 Cambridge, MA USA

The problem of designing logical topologies for rings that minimize cost as measured by the amount of electro-optic equipment has recently received much attention in the literature [11, 12, 1, 2, 5, 6, 7]. The problem is addressed in part or full to arrive at heuristic solutions in [11, 12, 2, 7]. A common cost measure used in literature is the number of SONET ADMs [5, 11]. Wavelength assignment to lightpaths has been recognized to be an important part of the problem and several studies focus on this [9, 5], while others consider the lightpath routing problem as well [7]. In [11, 12], a strategy of first grooming traffic components into circles is presented; in [11] these circles are then groomed, in [12] they are scheduled in a sequence of virtual topologies. Heuristic algorithms to minimize network cost by grooming are presented in [2], for special traffic patterns such as uniform, certain cases of cross-traffic, and hub. In [1], a heuristic algorithm based on a bipartite matching formulation of the problem is presented for specific traffic characteristics.

As has been noted in literature [4, 11], the problem of logical topology design is NP-hard even for a physical ring topology, and achievability bounds are useful for evaluating the performance of heuristic algorithms. We present a new framework for computing bounds for the problem of optimal traffic grooming in physical ring topologies. We decompose the ring into path segments consisting of successively larger number of nodes. We show that solving a path segment exactly is much easier than solving a ring of the same number of nodes. We combine the path solutions to obtain a series of bounds, both lower and upper. Computation of the bounds requires less effort than computing the optimal value, and depending on the problem instance, several bounds in the sequence are likely to require significantly less effort.

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 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 such as multiple fiber links, physical hop limit, bidirectional rings, etc.

2. PROBLEM FORMULATION

We consider a unidirectional ring $\mathcal{R}$ with N nodes, 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 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. 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 $t^{(sd)} \in \{0, 1, 2, \cdots\}$ 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

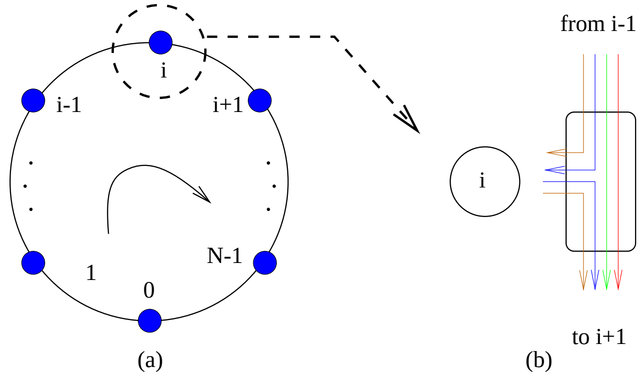

Figure 1: (a) An N-node unidirectional ring, and (b) detail of a node with a WADM

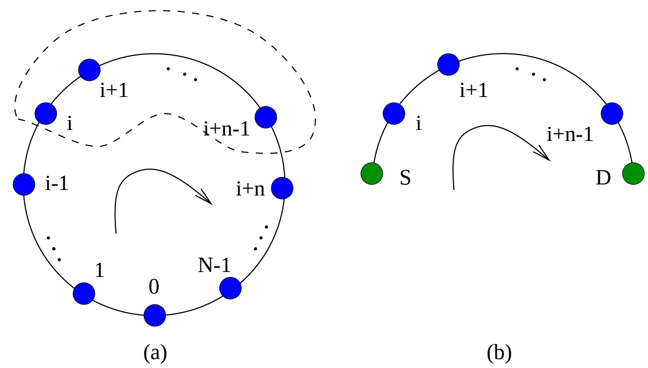

Figure 2: An n-node decomposition: (a) original ring $\mathcal{R}$, and (b) a decomposition $\mathcal{P}_n^{(i)}$

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.

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. We also allow for multiple lightpaths with the same source and destination nodes.

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, following the formulation in [4] for the general topology case. The specific details and the mathematical formulation as an Integer Linear Problem (ILP) for the ring network is omitted here and can be found in [3]. It 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.

The framework we present below is based on the 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 [9]. 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 [11, 5]) 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 [7]). 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 type 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 bound the optimal value of the objective function.

3. PATH DECOMPOSITION OF A RING

3.1 Definition of Decomposition

We consider a ring network $\mathcal{R}$ with N nodes labeled $0\cdots(N-1)$, in order, 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 $\mathcal{R}$ that includes the n consecutive nodes $i, i\oplus 1, i\oplus 2,\cdots, i\oplus(n-1)$, 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 $T_{\mathcal{P}_n^{(i)}}=\left[t^{(sd)}_{\mathcal{P}_n^{(i)}}\right]$, $s,d\in\{i,i\oplus 1,\cdots,i\oplus(n-1),D,S\}$, derived from the original matrix T, where $t^{(sd)}_{\mathcal{P}_n^{(i)}}$ is given by:

$$\begin{array}{ll} t^{(sd)}, & \text{if} \quad i\preceq s\prec d\preceq i\oplus(n-1) \\ \sum_{j\notin\{i,i\oplus 1,\cdots,d\ominus 1\}} t^{(jd)}, & \text{if} \quad s=S, i\preceq d\preceq i\oplus(n-1) \\ \sum_{j\notin\{s\oplus 1,\cdots,i\oplus(n-1)\}} t^{(sj)}, & \text{if} \quad d=D, i\preceq s\preceq i\oplus(n-1) \\ t_{\text{pass-thru}}(i,n), & \text{if} \quad s=S, d=D \\ 0, & \text{if} \quad \left\{\begin{array}{l} s=d \quad \text{or} \\ s=D \quad \text{or} \\ d=S \quad \text{or} \\ i\preceq d\prec s\preceq i\oplus(n-1) \end{array}\right. \end{array} \qquad (1)$$

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, 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 (1)). 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. The 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 (1)). Node S also acts as the source for all traffic components that pass through the segment, as the fourth expression in (1) 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 (1)), 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 (1) 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 (1), 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 $T_{\mathcal{P}_n^{(i)}} = \left[t^{(sd)}_{\mathcal{P}_n^{(i)}}\right]$ 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 (1). This matrix can be thought of as representing the traffic demands 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 node D to node S. Consequently, the ILP formulation we described in Section 2 can be used to obtain a virtual topology that minimizes electronic routing for this "ring". 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 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, the proof of which is omitted here and can be found in [3].

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

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 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.

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 $\mathcal{P}_n^{(i)}$. 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. 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.

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. 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 From Single-Node Decompositions

A decomposition of an N-node ring around a single node i is shown in Figure 3. The next two sections show how to obtain lower and upper bounds, respectively, on the optimal ILP solution for the ring as a whole by appropriately combining the optimal ILP solutions $\phi_1^{(i)}$ for the possible single-node decompositions $\mathcal{P}_1^{(i)}, i = 0, \cdots, (N-1)$.

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:

$$\Phi_1 = \sum_{i=0}^{N-1} \phi_1^{(i)} \tag{2}$$

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 $\mathcal{R}$.

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. In such a topology, each node i performs the maximum possible amount of electronic routing, which we denote by $\psi^{(i)}$. Quantities $\psi^{(i)}, i = 0, \cdots, (N-1)$, 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:

$$\Psi_0 = \sum_{i=0}^{N-1} \psi^{(i)} \tag{3}$$

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 to obtain

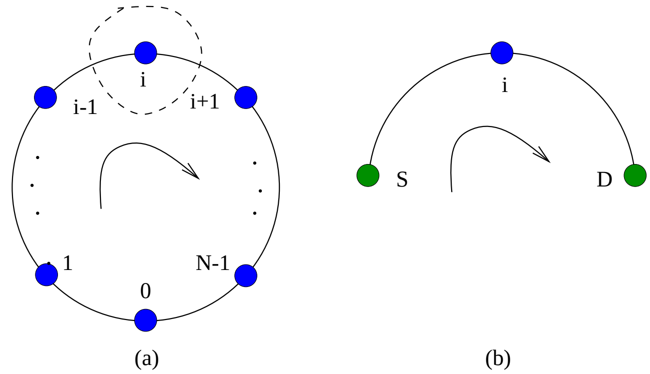

Figure 3: A single node decomposition of a ring: (a) original ring, and (b) single node decomposition $\mathcal{P}_1^{(i)}$ around node i

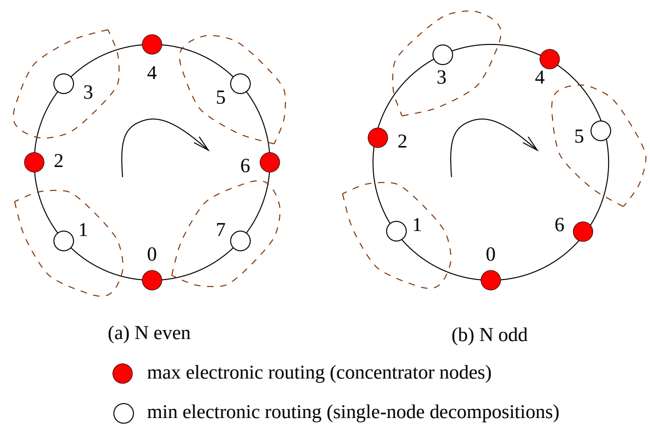

Figure 4: Virtual topology with alternating single-node decompositions and concentrator nodes

a tighter upper bound. To this end, let us refer again to Figure 3(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.

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. Looking back at the single node decomposition, we see that the nodes S and D can be viewed as concentrator nodes. Carrying this line of thought one step further, we are led to consider a new virtual topology where every other node is a concentrator node (performing the maximum electronic routing, $\psi^{(i)}$) while the remaining nodes perform the minimum electronic routing possible, $\phi_1^{(i)}$. This topology is illustrated in Figure 4, 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 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 4(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:

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

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 $\Psi_1 \leq \Psi_0$.

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. The proof follows obviously from the arguments above and is omitted.

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 $k, k = 1, \cdots, |\sigma_n|$, and let i_k denote its starting node. The quantity*

$$\Phi(\sigma_n) = \sum_{k=1}^{|\sigma_n|} \phi_{l_k}^{(i_k)} \quad (5)$$

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:

$$\Phi_n = \max\{\Phi(\sigma_n)\} \quad (6)$$

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

$$\Phi_{n+1} \geq \Phi_n \qquad \forall n \in \{1 \cdots (N-1)\} \quad (7)$$

As a result, the sequence $\Phi_1, \Phi_2, \cdots, \Phi_N$, 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 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 (4) 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 6 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. 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 any N-node decompositions. 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

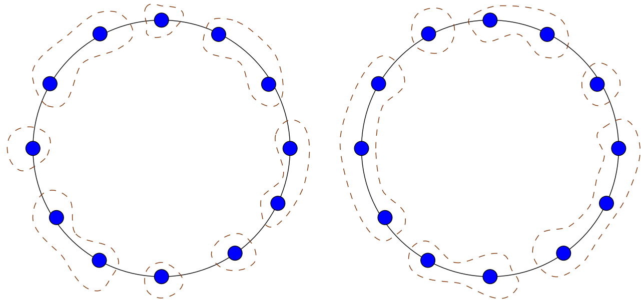

(a) (b)

Figure 5: 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

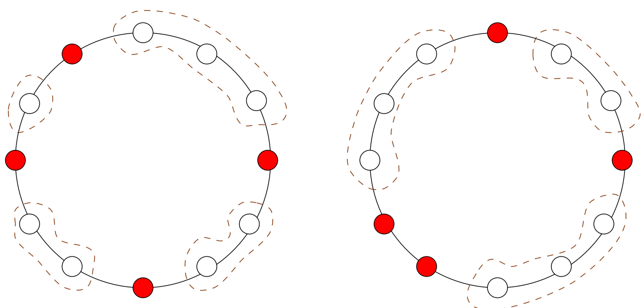

(a) (b)

Figure 6: Two partitions of a ring into alternating concentrator nodes and segments of no more than 3 nodes

obtaining Ψ_n, therefore we may assert that:

$$\Psi_{n+1} \le \Psi_n \qquad \forall n \in \{0 \cdots (N-2)\} \tag{8}$$

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 [6, 2, 1], because it considers all topologies with a single concentrator node (which is equivalent to a hub node). For a similar reason, $\Psi_k, k \ge \lceil N/2 \rceil$ 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:

$$\zeta^{(j)} = \min_{i=0}^{N-1} (\psi^{(i)} - \phi_1^{(i)}) \tag{9}$$

The final upper and lower bound in our framework 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 [3].

LEMMA 4.2. *The guarantee on the final values in the sequences of upper and lower bounds is:*

$$\Psi_{N-1} - \Phi_{N-1} \quad \le \quad \zeta^{(j)} \tag{10}$$

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 guarantee on the tightness of the framework 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 appropriate combinations of $\phi_n^{(i)}$ 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 $1, 2, \cdots, n$ 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 [3], and requires $O(n^2 N)$ time to find Φ_n given all $\phi_x^{(i)}$ values for $x = 1, \cdots, n$. For the largest number of total partitions (when $n = N$), this corresponds to $O(N^3)$ instead of $O(2^N)$ time, and 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, the proof of which is omitted here and can be found in [3]. These properties introduce a constant factor of improvement to the dynamic programming algorithm mentioned above.

LEMMA 4.3. *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 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, $\phi_x^{(i)} \geq \phi_{y_1}^{(i)} + \phi_{y_2}^{(i+y_1)} + \phi_{y_3}^{(i+y_1+y_2)} + \cdots + \phi_{y_n}^{(i+y_1+y_2+\cdots+y_{n-1})}$, 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 $y_1 + y_2 \leq x$, 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 $\{\Psi_n\}$.

5. 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 $t^{(s \oplus x, d \oplus x)}$ 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*. 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.

We now turn our attention to the different traffic components originated at 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 same pattern, the characteristic load scales with the matrix elements.

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%. 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. 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 [4, 8]. 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 7, 8 show the results for 8-node rings, statistically uniform traffic, 50% and 90% characteristic load respectively. Figures 9, 10 similarly show the results for statistically falling traffic, 50% load, and rising traffic, 90% load, respectively. Figures 11, 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. 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.

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 8-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 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

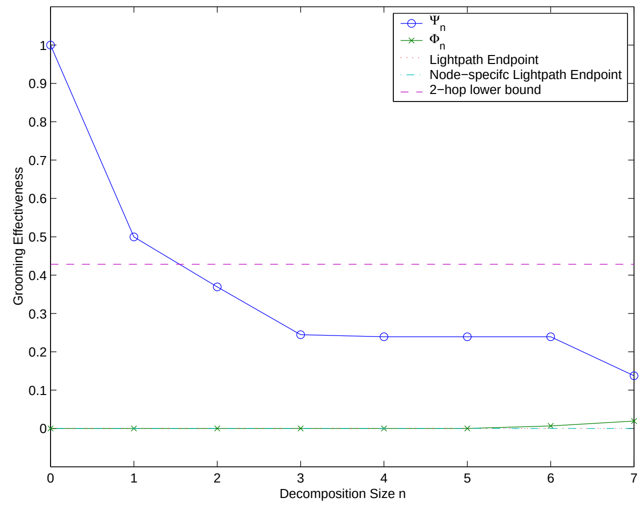

Figure 7: $N = 8$, Statistically uniform, 50% load

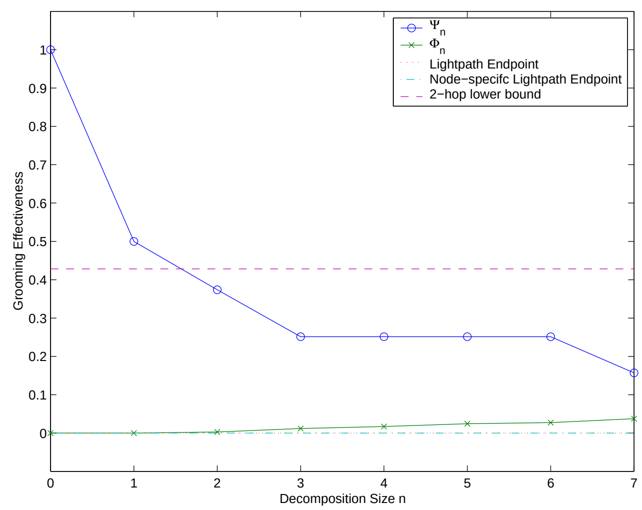

Figure 8: $N = 8$, Statistically uniform, 90% load

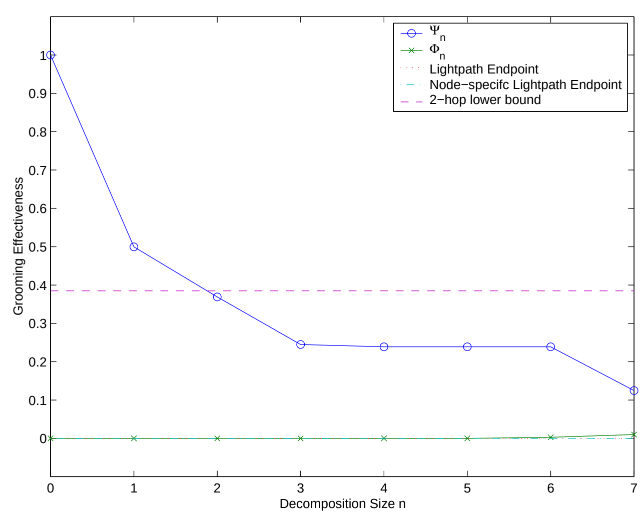

Figure 9: $N = 8$, Statistically falling, 50% load

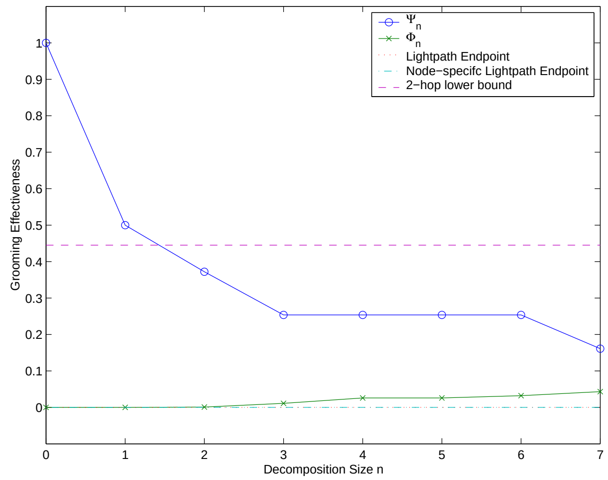

Figure 10: $N = 8$, Statistically rising,90% load

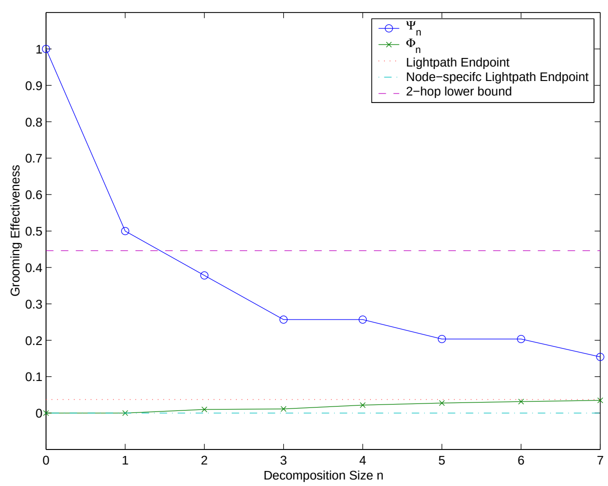

Figure 11: $N = 16$, Statistically falling to end, 90% load

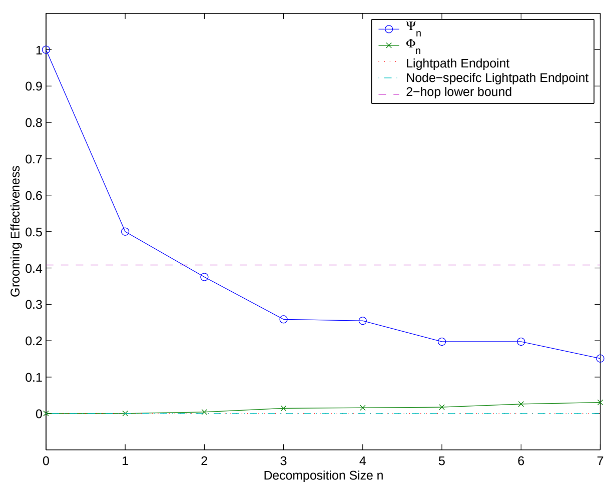

Figure 12: $N = 16$, Statistically falling to $N/2$, 90% load

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, [10]) 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 [10]. 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 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, (figure 11) the first bound has a comparable value to the largest Φ_n we have obtained.

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). Figure 13 presents the results for an ensemble of 8-node statistically uniform traffic matrices with characteristic load around 90%, the actual values of electronic routing are plotted. Figure 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 2-hop lower bound is also plotted.

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 characteristics 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

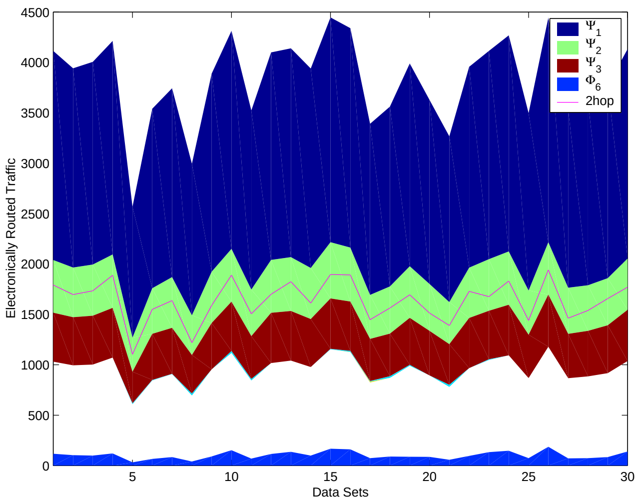

Figure 13: Ensemble, $N = 8$, Statistically uniform, 90% load, electronic routing

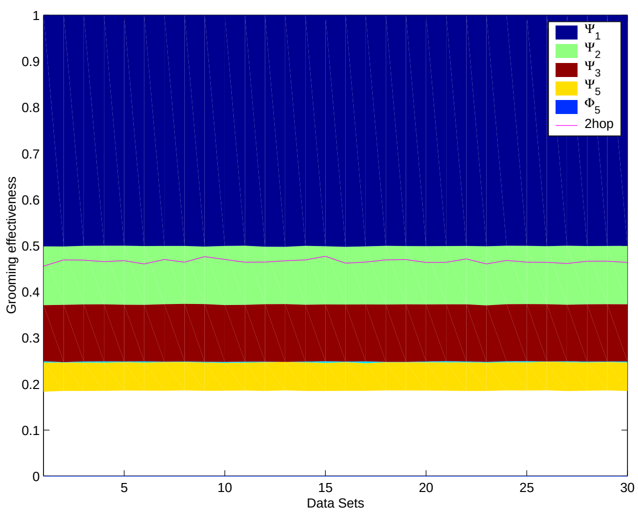

Figure 14: Ensemble, $N = 16$, Statistically uniform, 50% load, (normalized)

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 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 bound is at least as strong as the last one. Numerical results indicate that the expectations from theoretical considerations are fulfilled.

7. REFERENCES

[1] R. Berry and E. Modiano. Minimizing electronic multiplexing costs for dynamic traffic in unidirectional sonet ring networks. *Proc. IEEE ICC*, 3:1724–1730, 1999.

[2] A. Chiu and E. Modiano. Reducing electronic multiplexing costs in unidirectional SONET/WDM ring networks via efficient traffic grooming. *Proc. IEEE Globecom*, 1:322–327, Nov 1998.

[3] R. Dutta and G. Rouskas. A sequence of bounds for the problem of minimizing electronic routing in wavelength routed optical rings. Technical Report TR-2000-11, NCSU, 2000.

[4] R. Dutta and G. Rouskas. A survey of virtual topology design algorithms for wavelength routed optical networks. *Optical Networks Magazine*, 1(1):73–89, Jan 2000.

[5] 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.

[6] O. Gerstel, R. Ramaswami, and G. Sasaki. Cost effective traffic grooming in WDM rings. *Proc. IEEE INFOCOM*, pages 69–77, 1998.

[7] G. Li and R. Simha. Efficient routing to reduce the cost of add-drop multiplexers in WDM optical ring networks. *Proc. of SPIE*, 3843:258–267, 1999.

[8] 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.

[9] R. Ramaswami and G. Sasaki. Multiwavelength optical networks with limited wavelength conversion. *Proc. IEEE INFOCOM*, pages 489–498, 1997.

[10] R. Ramaswami and K. Sivarajan. Design of logical topologies for wavelength-routed optical networks. *IEEE JSAC*, 14(5):840–851, June 1996.

[11] X. Zhang and C. Qiao. An effective and comprehensive solution to traffic grooming and wavelength assignment in SONET/WDM rings. *Proc. SPIE*, 3531:221–232, Nov 1998.

[12] 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.