Topology Design in WDM Rings to Minimize Electronic Routing: Efficient Computation of Tight Bounds *

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. The full virtual topology design problem is NP-hard even in the restricted case where the physical topology is a ring, and heuristics have been proposed in the literature. We present a new framework of bounds which can be used to evaluate the performance of heuristics, and which requires significantly less computation than evaluating the optimal solution. This framework 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 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.

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.

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

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

[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 performances 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 require 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 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 such as multiple fiber links between nodes, 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) (working ring only). 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 *add* or *drop* optical traffic from/into electronic form, for electronic processing. The processing is traffic termination or possibly electronic routing. It can also 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. We assume that estimates of the aggregate 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. (We do not consider the *dynamic* case of requests for lightpaths or traffic components arriving continuously during operation, this latter problem has been considered elsewhere [13, 14, references thereof].) The traffic demands between pairs of nodes in the ring are given in the traffic matrix $T = [t^{(sd)}]$. All traffic is expressed in terms of some base rate (*e.g.* OC3) and the bandwidth available in a single wavelength is denoted by C in this unit.

Given the ring physical topology, a *logical* topology is defined by establishing *lightpaths* between pairs of nodes. A lightpath is a direct optical connection. 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. 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.

In this paper we have chosen to concentrate on the total electronic forwarding (routing) performed by a virtual topology as the performance metric to be optimized. This is of interest because 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. 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 the alternative consumes more bandwidth, and is likely to require more electronic routing. We can 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 as an Integer Linear Problem (ILP), following the formulation in [4] for the general topology case. The specific details and the mathematical formulation 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.

3 Path Decomposition of a Ring Network

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 $i, i\oplus 1, i\oplus 2, \cdots, i\oplus(n-1)$, 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 T, the original traffic matrix, as follows:

$$t^{(sd)}_{\mathcal{P}_n^{(i)}} = \begin{cases} t^{(sd)}, & \text{if } \; i \preceq s \prec d \preceq i \oplus (n-1) \\ \sum_{j \notin \{i, i\oplus 1, \cdots, d\ominus 1\}} t^{(jd)}, & \text{if } \; s = S, i \preceq d \preceq i \oplus (n-1) \\ \sum_{j \notin \{s\oplus 1, s\oplus 2, \cdots, i\oplus(n-1)\}} t^{(sj)}, & \text{if } \; d = D, i \preceq s \preceq i \oplus (n-1) \\ t_{\text{pass-thru}}(i,n), & \text{if } \; s = S, d = D \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

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. 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.

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. The new node S acts as the source of all traffic flowing into the segment in the real ring. Similarly, the new node D acts as the destination for all traffic seen to flow out of the segment. The details of the flow of these traffic components in the rest of the ring is hidden from the segment. 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.

Solving Path Segments in Isolation: Consider the traffic matrix $T_{\mathcal{P}_n^{(i)}} = \left[t_{\mathcal{P}_n^{(i)}}^{(sd)} \right]$ 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 mentioned 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 so obtained achieves the locally best value for electronic routing for the set of nodes in the decomposition, but does not take into account the rest of the ring. This topology may not form the subtopology for these nodes for any optimal solution, or even any feasible solution, of the complete ring. We denote the optimal objective value for the decomposition $\mathcal{P}_n^{(i)}$ by $\phi_n^{(i)}$. Since the new nodes S and D do not contribute any electronic routing, and the traffic pattern seen by the n nodes abstracted from the real 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.

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 is straightforward and is based on assigning arbitrary wavelengths on the first link and extending them over the path. It 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 network, and it 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, we can eliminate the variables and constraints related to wavelength assignment from our formulation. The order of the numbers of variables and constraints remain the same, but a significant number of them are removed. 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.

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.

4.1 Lower Bound

Recall that $\phi_1^{(i)}$ represents the locally best amount of electronic routing that can be achieved at node i considered in isolation and therefore in any toplogy, 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 denote this sum by $\Phi_1 = \sum_{i=0}^{N-1} \phi_1^{(i)}$. 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 $\mathcal{R}$.

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

$$\Phi(\sigma_n) = \sum_{k=1}^{|\sigma_n|} \phi_{l_k}^{(i_k)} \tag{2}$$

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

We now define Φ_n as:

$$\Phi_n = \max\{\Phi(\sigma_n)\} \tag{3}$$

where the maximum is taken over all partitions of the ring which contain segments with n or less nodes. Figure 4 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. Note that our definition of Φ_1 before is consistent with this general definition.

It is obvious from definition (3) that the set of partitions we consider in computing Φ_{n+1} is a proper superset of that we consider in computing Φ_n. As a result, we have a strong sequence of bounds in which each is at least as tight as the previous one, that is, $\Phi_{n+1} \geq \Phi_n \, \forall n \in \{1 \cdots (N-1)\}$.

4.2 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 the simplest virtual topology possible, namely, the topology consisting only of W single-hop lightpaths from each node i to the next node $i \oplus 1$. All traffic not terminating at any given node must be electronically routed by that node. 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 [6].) 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 denote the total electronic routing performed under the no-wavelength-routing topology by $\Psi_0 = \sum_{i=0}^{N-1} \psi^{(i)}$. Since this is a feasible topology, Ψ_0 is an upper bound on the optimal electronic routing,

Let us call *concentrator nodes* those nodes which do not perform any optical forwarding (in the no-wavelength-routing topology, every node is a concentrator node). Nodes S and D can be viewed as concentrator nodes in the single node decomposition, thus we see that nodes of any segment preceded and succeeded by a concentrator node each is free to route traffic to achieve the local best case, corresponding to $\phi_n^{(i)}$ for that segment. For example, a topology may be created by alternating concentrator nodes with single nodes which route traffic according to the corresponding single node decomposition solutions. This topology is illustrated in Figure 3, where the even-numbered nodes are concentrator nodes. The objective value for such a topology would be an alternate sum of $\phi_1^{(i)}$ and $\psi^{(i)}$ values. For the strongest possible upper bound, we would choose the topology which minimizes this sum.

It is now straightforward to obtain a strong sequence of upper bounds along the same lines. We define Ψ_n as the lowest objective value we obtain for all topologies which are created by alternating concentrator nodes with segments no larger than n nodes in size. We can once again consider this in light of partitions of the nodes of a ring. Now, however, the partitions are constrained not only to use segments of n-nodes or less, but every alternate segment must contain exactly one node. These alternate single-node segments are used as concentrator nodes in the topology we create, rather than as single-node decompositions.

We 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 $\Psi_{n+1} \leq \Psi_n \, \forall n \in \{0 \cdots (N-2)\}$, giving us a strong sequence of upper bounds.

Because the bounds $\{\Psi_n\}$ are based on actual feasibly topologies, they also provide us with a useful series of heuristic solutions to the ring. Below we mention 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, $\Psi_k, k \geq \lceil N/2 \rceil$ must be as good or better than a double hub design, if the hubs are constrained to be diametrically opposite in the ring.

Both the upper and the lower bounds get progressively tighter with increasing n, and for the final bounds we can assert that:

$$\Psi_{N-1} - \Phi_{N-1} \leq \min_{i=0}^{N-1} (\psi^{(i)} - \phi_1^{(i)}) \tag{4}$$

The derivation is based on considering a partition of the ring into the single node which yields the minimum $\psi^{(i)} - \phi_1^{(i)}$, and the $N-1$ decomposition formed by the rest of the nodes. We then obtain both and upper and lower bound from this partition and compare them to Ψ_{N-1} and Φ_{N-1}. The derivation is omitted here and can be found in [3].

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. The dominating factor in the computational effort is the calculation of the $\phi_n^{(i)}$ values required for a given bound, 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 bounds is limited by this computational effort. The computational effort required to combine the $\phi_n^{(i)}$ values into the best available bounds is $O(2^N)$ for a straightforward algorithm, but the problem can be presented as a dynamic programming problem which takes only $O(n^2N)$ time to solve. We omit this algorithm here. The algorithm can be found in [3], as well as a lemma which can introduce a further constant order reduction in the computation required.

5 Numerical Results

We define a traffic pattern as *statistically symmetric* in which the traffic components of the form $t^{(s\oplus x, d\oplus x)}$ for all x and for any given s and d are all drawn from the same distribution. In such a case the traffic pattern looks roughly the same from all nodes of the ring. We concentrate on such traffic patterns for producing our numerical results because for highly asymmetric traffic patterns the difference in performance between the best and worst performing topologies is likely to be less. We further categorize statistically symmetric traffic matrices as rising, falling or uniform depending on whether the traffic components sourced by a node to other nodes increase, decrease or remain the same with distance to destination node. This concept is also statistical in nature. We characterize a traffic pattern by the average physical traffic load on the links of the ring (the loads on individual links are likely to be close to each other because of the symmetrical nature of the traffic) expressed as a fraction of WC, the maximum possible, for comparison purposes.

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; for more numerical results see [3].

The quantity Ψ_0 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 [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*. We express all quantities plotted below by normalizing them to Ψ_0 to express grooming effectiveness values accordingly.

Figures 5 and 6 show detailed results for 8-node rings, statistically uniform traffic with 50% load, and statistically rising traffic with 90% load respectively. Figures 7 and 8 both show detailed results for 16 node rings, statistically falling traffic with 90% load. Figure 7 represents a case where the traffic from a node falls to zero at the farthest node, whereas Figure 8 represents one in which the traffic falls to zero at a node halfway around the ring, and is zero to the farther nodes. All the results show similar characteristics. There is a sharp decrease from Ψ_0 to Ψ_1 and more moderate decrease thereafter. The quantity Ψ_7 is between 0.1 and 0.2 in all cases. 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. We also plot the values of two lower bounds computed after the fashion of the Moore bound following [10]. These 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 7), the first bound has a comparable value to the largest Φ_n we have obtained.

The other two figures represent results of 30 traffic matrices of the same traffic characteristics each. Figure 9 represents an ensemble of 30 traffic matrices of 8-node rings, each of uniform traffic with 90% load; in this case the absolute electronic routing values are plotted. Figure 10 represents an ensemble of 30 traffic matrices of 16-node rings, uniform traffic with 50% load; normalized data as with the previous results are plotted. The 2-hop lower bound is also plotted. For the latter ensemble 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.

We also 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, a heuristic attractive for its simplicity and potential for easy OAM. A lower bound on the performance of such a heuristic is easy to obtain by considering the minimum number of times each traffic component must be electronically routed. 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. However, in both detailed and ensemble results we see that most 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.

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 successive bound is at least as strong as the last one.

We present numerical results of the computation of these bounds by computing the bounds on different families of traffic matrices. Numerical results indicate that 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 upper bounds are based on constructing actual feasible topologies on the network, and thus also provide us with a sequence of increasingly good heuristics. The framework can be adapted to other formulations of the problem on the ring network, and we feel this is a useful framework in a broad context of ring problems for wavelength routing optical networks.

References

[1] 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.

[2] 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.

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

[4] 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.

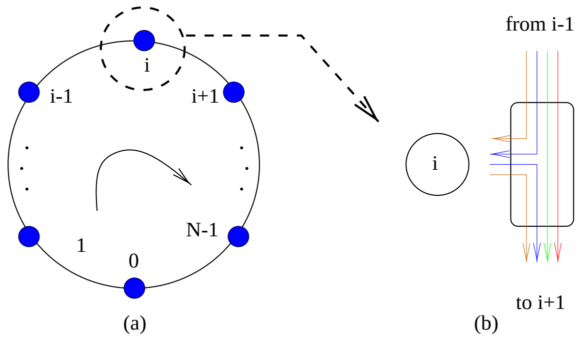

Figure 1: (a) An N-node unidirectional ring, and (b) detail of a node with a WADM

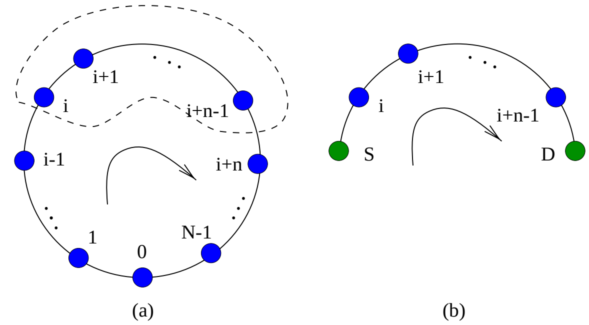

Figure 2: An n-node decomposition: (a) original ring $\mathcal{R}$, and (b) a decomposition $\mathcal{P}_n^{(i)}$

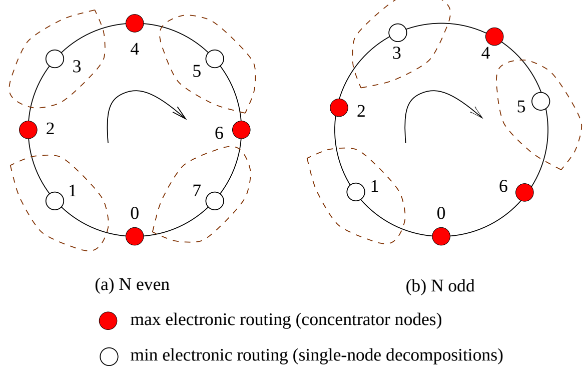

Figure 3: Virtual topology with alternating single-node decompositions and concentrator nodes

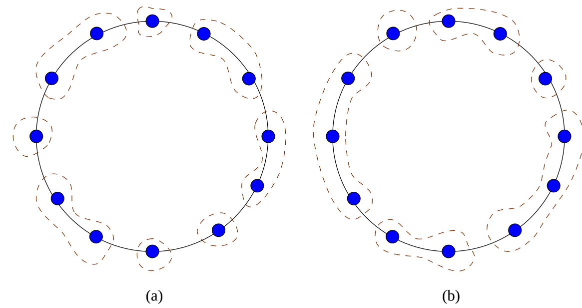

Figure 4: Partitions of the nodes of a ring into (a) segments of no more than 2 nodes, and (b) segments of no more than 3 nodes

[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. *Proceedings of the 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.H. Sasaki. Multiwavelength optical networks with limited wavelength conversion. *Proc. IEEE INFOCOM*, pages 489–498, 1997.

[10] R. Ramaswami and K.N. 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. *SPIE Conference on All-Optical Networking:Architecture, Control and management Issues*, 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.

[13] 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.

[14] 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.

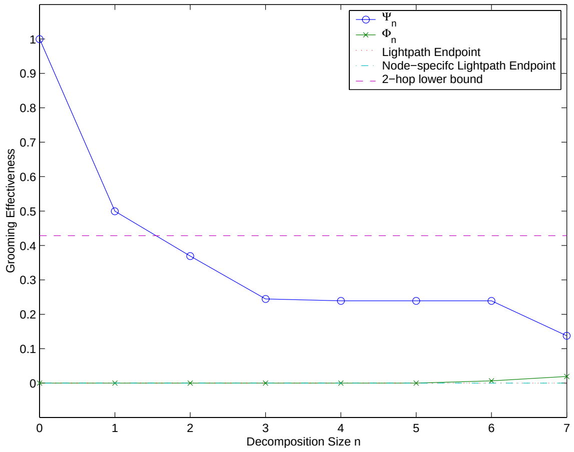

Figure 5: Detailed results for $N = 8$, Statistically uniform pattern, 50% load

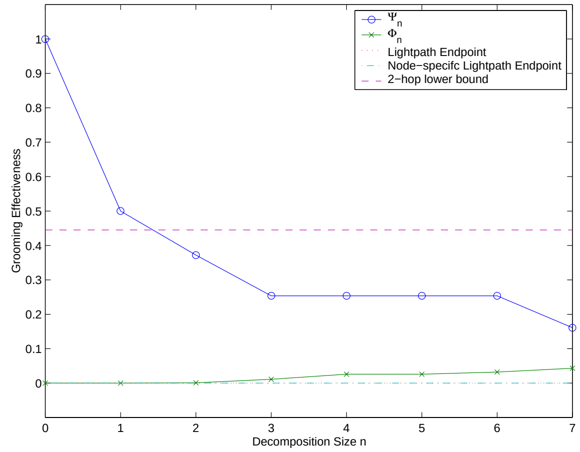

Figure 6: Detailed results for $N = 8$, Statistically rising pattern,90% load

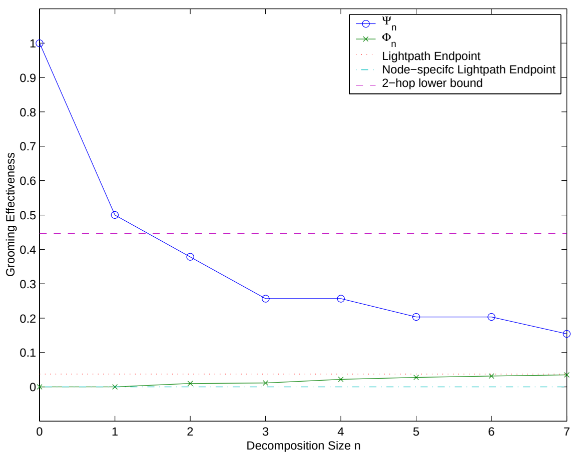

Figure 7: Detailed results for $N = 16$, Statistically falling pattern, 90% load, falling to end

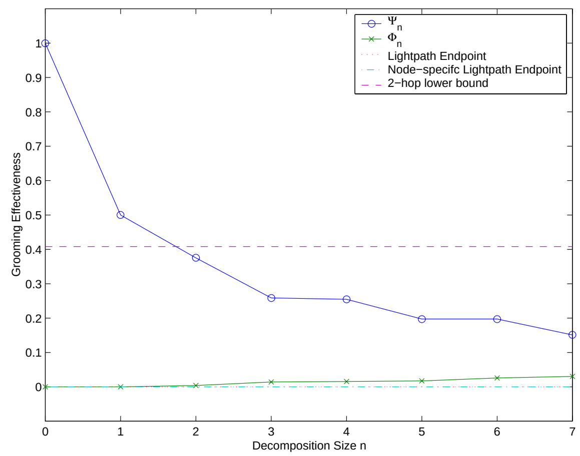

Figure 8: Detailed results for $N = 16$, Statistically falling pattern, 90% load, falling to $N/2$

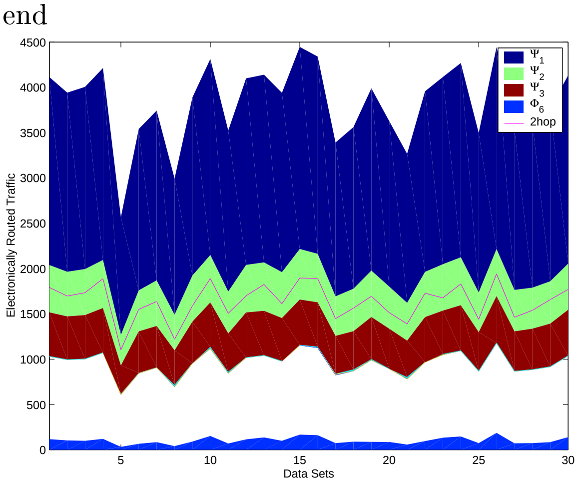

Figure 9: Ensemble results for $N = 8$, Statistically uniform pattern, 90% load, electronic routing

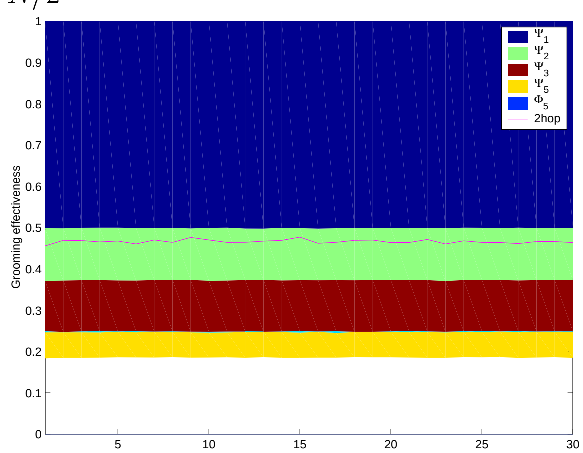

Figure 10: Ensemble results for $N = 16$, Statistically uniform pattern, 50% load, (normalized)