

Traffic Grooming in WDM Ring Networks with the Min-Max Objective*

Bensong Chen, George N. Rouskas, and Rudra Dutta

Department of Computer Science, North Carolina State University, Raleigh, NC, USA
{bchen,rouskas,dutta}@eos.ncsu.edu

Abstract. We consider the problem of traffic grooming in WDM ring networks. Previous studies have focused on minimizing *aggregate* representations of the network cost. In this work, we consider a *Min-Max objective*, in which it is desirable to minimize the cost at the node where this cost is maximum. Such an objective is of practical value when dimensioning a network for unknown future traffic demands and/or for dynamic traffic scenarios. We prove that traffic grooming with the Min-Max objective is NP-Complete even when wavelength assignment is not an issue. We also present a new polynomial-time algorithm for Min-Max traffic grooming. Experiments with a wide range of problem instances demonstrate that our algorithm produces solutions which are always close to the optimal and/or the lower bound, and which scale well to large network sizes, large number of wavelengths, and high loads.

1 Introduction

Wavelength division multiplexing (WDM) technology has the potential to satisfy the ever-increasing bandwidth needs of network users on a sustained basis. In WDM networks, nodes are equipped with *optical cross-connects* (OXC) or *optical add-drop multiplexers* (OADMs), devices which can optically switch wavelengths, thus making it possible to establish *lightpath* connections between pairs of network nodes. The set of lightpaths defines a *logical topology*, which can be designed to optimize some performance measure for a given set of traffic demands. The logical topology design problem has been studied extensively in the literature. Typically, the traffic demands have been expressed in terms of whole lightpaths, while the metric of interest has been the number of wavelengths, the congestion (maximum traffic flowing over any link), or a combination of the two. The reader is referred to [1] for a survey and classification of relevant work.

With the deployment of commercial WDM systems, it has become apparent that the cost of network components, especially *line terminating equipment* (LTE), is one of the dominant costs in building optical networks, and is a more meaningful metric to optimize than, say, the number of wavelengths. Furthermore, in order to utilize bandwidth more effectively, new models of optical networks allow several independent traffic streams to *share* the capacity

* This work was supported in part by the NSF under grant ANI-0322107.

of a lightpath. These observations give rise to the concept of *traffic grooming*, a variant of logical topology design, which is concerned with the development of techniques for combining lower speed components onto wavelengths in order to minimize network cost. A comprehensive survey and classification of traffic grooming strategies can be found in [2].

Typically, the goal of traffic grooming studies is to minimize an *aggregate* representation of the LTE cost. While such a metric is important, not imposing a bound on the cost of individual nodes may result in a solution in which some nodes end up with a (very) large amount of LTE while some others with only a small amount of LTE. Such a solution has a number of undesirable properties. First, a node that requires a large amount of LTE may be too expensive or impractical to deploy (e.g., due to high interconnection costs, high power consumption, or space requirements). Second, the network can be highly heterogeneous in terms of the capabilities of individual nodes, making it difficult to operate and manage. Third, a solution minimizing the total LTE cost can be extremely sensitive to the assumptions regarding the traffic pattern, as previous studies [3] have demonstrated. Specifically, a solution that is optimal for a given set of traffic demands may be far away from optimal for a different such set. Since LTE involve expensive hardware devices that cannot be moved from one node to another on demand, an approach that attempts to minimize total LTE cost may not be appropriate for dimensioning a network unless the network operator has a clear picture of traffic demands far into the future *and* these traffic demands are unlikely to change substantially over the life of the network.

In this work, we consider a variant of the traffic grooming problem in which the objective is to minimize the LTE cost at the node where this cost is maximum. We believe that a Min-Max objective is of high practical value to network designers and operators. In particular, an approach that minimizes the maximum LTE cost at any network node is likely to be attractive because all the network nodes are likely to be provisioned with identical equipment. Effectively, all nodes will have a cost that is dictated by the node with the maximum LTE. Such a homogeneous network is easier to operate, manage and maintain, and is likely to be less expensive than a heterogeneous one due to the economies of scale that can be achieved when all nodes are subject to identical specifications. Furthermore, such an optimization approach can be of great importance to dimensioning the network for *unknown and/or dynamic future demands*. Specifically, the network designer may solve the optimization problem for a wide range of traffic scenarios, and equip each node with an amount of LTE equal to the highest solution obtained (plus a certain *fudge factor* for making the solution future-proof). We also note that a similar approach was taken in [4] in a different context, namely for routing and wavelength assignment in the presence of converters. Specifically, an algorithm was developed for distributing a number of converters uniformly across the ring nodes rather than placing them at a single hub node.

The paper is organized as follows. In Section 2, we review the general traffic grooming problem, we introduce the Min-Max objective we consider, and we present new theoretical results regarding the complexity of the problem. In

Section 3, we present a polynomial-time algorithm for rings to minimize the maximum LTE cost at any ring node. We present numerical results in Section 4, and we conclude the paper in Section 5.

2 Problem Definition and Complexity Results

New models of optical networks allow several independent traffic streams to share the bandwidth of a lightpath. If the multiplexing and demultiplexing of lower-rate traffic components is performed at the boundaries of the network only (i.e., at edge routers), and the aggregate traffic transparently traverses the optical network, this problem is equivalent to the *routing and wavelength assignment* (RWA) problem. It is known that wavelength assignment to minimize the number of wavelengths can be solved in polynomial time in paths and stars, but that it is NP-hard in rings, trees, and networks of a general topology. However, it is in general impossible to set up lightpaths between every pair of edge routers. Therefore, it is natural to consider optical networks in which nodes have both optical and electronic switching capabilities. Such nodes let some lightpaths pass through transparently, while they may terminate others. Traffic on terminating lightpaths may then be electronically switched (*groomed*) onto new lightpaths towards the destination node. Introducing some amount of electronic switching within the optical network has several advantages: it enhances the degree of *virtual connectivity* among the edge routers, which otherwise is limited by the number of optical interfaces at each router; it increases bandwidth utilization; and it may also reduce the wavelength requirements within the optical network for a given traffic demand. The trade-off is an increase in network cost due to the introduction of expensive active components (i.e., optical transceivers and electronic switches). These observations motivate us to define the following traffic grooming problem.

Let C be the capacity of each wavelength, expressed in units of some arbitrary rate (e.g., OC3); we will refer to parameter C as the *grooming factor*. Let W be the number of wavelengths that each fiber link in the network can support. We represent a traffic pattern by a demand matrix $T = [t^{(sd)}]$, where integer $t^{(sd)}$ denotes the number of traffic streams (each of unit demand) from node s to node d . Given matrix T , the traffic grooming problem involves the following conceptual subproblems (SPs):

1. *logical topology SP*: find a set R of lightpaths,
2. *lightpath routing and wavelength assignment SP*: solve RWA on R , and
3. *traffic routing SP*: route each traffic stream through the lightpaths in R .

The first and third subproblems together constitute the grooming aspect of the problem. Also, the number W of wavelengths per fiber link is taken into consideration as a constraint rather than as a parameter to be minimized.

The goal we consider in this paper is to minimize the *maximum number of lightpaths originating from or terminating at any node*. In our cost model, one unit of cost is incurred for each lightpath that terminates at, or originates

from, a network node. Thus, this cost metric accurately reflects the amount of LTE needed at each network node. Note also that this objective is equivalent to minimizing the *maximum nodal degree* in the logical topology.

In this work, we study the traffic grooming problem in a unidirectional ring \mathcal{R} with N nodes; this problem is NP-Complete since the RWA subproblem in rings is NP-Complete. In this section, we prove that the traffic grooming problem in path networks is also NP-Complete. Since the RWA problem can be solved in linear time in path networks, our results demonstrate that traffic grooming with the Min-Max objective is itself an inherently difficult problem, and that it remains NP-Complete in rings or other general topologies *even when* full wavelength conversion is available at the network nodes.

Let us consider a unidirectional path network \mathcal{P} with N nodes. There is a single directed fiber link from node i to node $i + 1$, for each $i \in \{1, 2, \dots, N - 1\}$. An instance of the traffic grooming problem is provided by specifying a number N of nodes in the path, a traffic matrix $T = [t^{(sd)}]$, $1 \leq s < d \leq N$, a grooming factor C , a number of wavelengths W , and a goal F . The problem asks whether a valid logical topology may be formed on the path and all traffic in T routed over the lightpaths of the logical topology so that the number of incoming or outgoing lightpaths at any node in the path is less than or equal to F .

We distinguish two problem variants depending on whether bifurcated routing of traffic is allowed or not. In the former variant, a traffic component $t^{(sd)}$ is allowed to be split into various subcomponents which may follow different routes (i.e., different lightpath sequences). The bifurcation is restricted to integer subcomponents. In the latter variant, for any pair (s, d) such that $t^{(sd)} \leq C$, we require that all $t^{(sd)}$ traffic units be carried on the *same* sequence of lightpaths from s to d . On the other hand, if $t^{(sd)} > C$, it is not possible to carry all the traffic on the same lightpath. In this case, we allow the traffic demand to be split into $\lfloor \frac{t^{(sd)}}{C} \rfloor$ subcomponents of magnitude C and at most one subcomponent of magnitude less than C , and the no-bifurcation requirement applies to each subcomponent independently. We then have the following two theorems; their proofs are omitted due to space constraints, and can be found in [5].

Theorem 1. *The decision version of the grooming problem in unidirectional paths with the Min-Max objective (bifurcated routing not allowed) is NP-complete.*

Theorem 2. *The decision version of the grooming problem in unidirectional paths with the Min-Max objective (bifurcated routing allowed) is NP-complete.*

Because of the construction of the proofs, we have the following corollary, which demonstrates that, even when solutions to the first two subproblems of the traffic grooming problem (see Section 2) are provided, the problem remains NP-Complete by virtue of the third subproblem (traffic routing). Therefore, traffic grooming is inherently more difficult than the well-known RWA problem.

Corollary 1. *The decision version of the traffic grooming problem in unidirectional paths with the Min-Max objective (bifurcated routing allowed or not) is NP-Complete even when a logical topology is provided.*

3 Min-Max Traffic Grooming Algorithm for Rings

We now present an algorithm for traffic grooming in WDM rings to minimize the maximum nodal degree (indegree or outdegree) in the logical topology. Rather than solving all three subproblems simultaneously (see Section 2), we decouple the logical topology and traffic routing subproblems from the RWA subproblem and tackle them independently. Specifically, the algorithm consists of the following steps which are explained in more detail in the following subsections:

- **Step 1.** Solve the logical topology and traffic routing subproblems on the ring network using the algorithm in Section 3.1. The result of this step is a set R of lightpaths and a routing of the traffic demands over the lightpaths in R that minimize the maximum nodal degree.
- **Step 2.** Use the algorithm in Section 3.2 to color the lightpaths of set R . The result of this step is a wavelength assignment that does not use more than W wavelengths. However, at the end of this step, the degree at one node, say, node i , of the ring may increase beyond the corresponding value after Step 1, by an amount equal to some value Δ .
- **Step 3.** Use the algorithm presented in [4] to distribute the additional Δ LTE at node i to other nodes in the ring network.

3.1 Min-Max Traffic Grooming Algorithm

We present a polynomial-time algorithm for the logical topology and traffic routing subproblems of the traffic grooming problem. Our algorithm attempts to minimize the maximum nodal degree by creating long lightpaths that bypass intermediate nodes whenever possible. Because of the results in Section 2, the traffic grooming subproblem is itself NP-Complete, and hence our polynomial-time algorithm will terminate without necessarily finding an optimal solution. However, numerical results to be presented later indicate that the solutions we obtain are close to the optimal and/or the lower bound.

Before we proceed, we introduce the concept of *reduction* of a traffic matrix. Specifically, we reduce the matrix T so that all elements are less than the capacity C of a single wavelength, by assigning a whole lightpath to traffic between a given source-destination pair that can fill it up completely. The available wavelengths on the links of the path segment from the source to the destination node are also decremented by the number of lightpaths thus assigned. Since breaking such lightpaths would increase the amount of LTE at some intermediate nodes of the path, this procedure does not preclude us from reaching an optimal solution, nor does it make the problem inherently easier or more difficult.

After the reduction, we initialize the logical topology to one in which a sufficient number of single-hop lightpaths is formed on each link of the ring network to carry the traffic using this link. We note that this initial solution is a feasible solution to the logical topology and traffic routing subproblems, in that it does not use more than W lightpaths on any link. However, this initial topology yields a large value of the maximum nodal degree F , which is equal to the number of

single-hop lightpaths in the most congested link. Our approach, then, is to improve on this initial solution by joining short lightpaths to form longer ones, thus lowering the degrees at intermediate nodes. In the following, we summarize our algorithm for joining short lightpaths.

Let us define the relationship $i \prec j$ between ring nodes to denote that node i “precedes” node j in the direction of traffic flow; similarly, we will use the notation $i \preceq j$ to denote that node i precedes, or may be the same as, node j . The main idea of our algorithm is to consider the node with the maximum degree, and to attempt to decrease its degree by one at each iteration; this process repeats until no more improvement is possible. Let m be the node with the maximum degree. The algorithm searches for a pair of nodes (i, j) , $i \prec m \prec j$, such that there exist lightpaths (i, m) and (m, j) . The objective is to shift all the traffic from lightpaths (i, m) and/or (m, j) to either an existing or a new lightpath (i, j) in order to decrease the maximum of the indegree and outdegree of node m by one. If the traffic can be shifted entirely to an existing lightpath (i, j) , then this procedure is always possible, since no wavelength limit constraints are violated, and also the degree of nodes i and/or j may also decrease in the process. However, if a new lightpath (i, j) must be created, the above procedure is carried out only if the wavelength limit constraint is not violated and the degrees of nodes i and j do not increase above the current maximum degree minus one (the minus one is necessary to ensure that the algorithm will not get into an infinite loop). A more detailed description is provided in Figure 1.

We now argue that at the end of each iteration of the **repeat** loop, the algorithm produces a solution to the logical topology and traffic routing subproblems that is feasible (i.e., no link carries more than W lightpaths – recall that we are not concerned with wavelength assignment at this point), and the maximum nodal degree is no larger (but possibly smaller) than that of the logical topology at the beginning of the iteration. At each iteration of the **repeat** loop, the algorithm tries to replace the lightpaths (i, m) and (m, j) for some nodes $i \prec m \prec j$ with a (possibly) new and longer lightpath (i, j) so as to decrease the nodal degree of node m . However, no action is taken if replacing the two lightpaths with the longer one would violate any wavelength constraints or would increase the degrees of i or j to more than the maximum at the start of the iteration, as we can see at Step 15. Since the initial topology at Step 2 of the algorithm is feasible, we conclude that the topology at the end of each iteration will be feasible and will not increase the maximum nodal degree.

The running time complexity of our algorithm is determined by the main iteration between Steps 6 and 21. In turn, the complexity of the iteration is determined by the two **for** loops, the inner **for** loop from Step 11 to 14, and the outer loop from Step 8 to 20. Each of these loops takes time $O(N^2)$ in the worst case, where N is the number of nodes in the ring. Therefore, each iteration through the **repeat** loop from Step 6 to 21 takes $O(N^4)$ time in the worst case. The main iteration of the algorithm (i.e., the **repeat** loop) will be executed at most $N\delta$ times, where δ is the maximum decrease in the degree of any node. Since $\delta < W$ always, the worst-case complexity of the algorithm is $O(WN^5)$.

Min-Max Traffic Grooming Algorithm

Input: A ring network with N nodes, W wavelengths each of capacity C , and reduced traffic matrix $T = [t^{(sd)}]$.

Output: The number of lightpaths b_{ij} from i to j , and the traffic routing quantities $t_{ij}^{(sd)}$ (which indicate the amount of the component $t^{(sd)}$ routed over a lightpath from i to j), so that the solution does not violate any wavelength constraints and minimizes the maximum number of lightpaths terminating at, or originating from, any node.

1. **begin**
 2. Initialize the logical topology to one with only single-hop lightpaths and initialize all b_{ij} and $t_{ij}^{(sd)}$ accordingly
 3. **for all** i, j **do** $r_{ij} \leftarrow$ capacity unused on the (i, j) direct lightpath in the current topology (residual capacity)
 4. **for all** j **do** $I_j \leftarrow$ indegree of j in the current topology
 5. **for all** j **do** $O_j \leftarrow$ outdegree of j in the current topology
 6. **repeat** // Main iteration
 - $m \leftarrow$ some node s. t. $\max\{I_m, O_m\}$ is maximum in the ring
 - // Our objective is to replace a lightpath (i, m) and/or (m, j) with a // direct lightpath (i, j) in order to reduce $\max\{I_m, O_m\}$ by one
 8. **for each pair** (i, j) such that $i < m < j$ **do**
 9. **if** $I_m == O_m$ **then** $TotalToShift \leftarrow \max\{C - r_{im}, C - r_{mj}\}$
else if $I_m > O_m$ **then** $TotalToShift \leftarrow C - r_{im}$
else $TotalToShift \leftarrow C - r_{mj}$
// $TotalToShift$ is the amount of traffic to be shifted to other
// wavelengths in order to reduce $\max\{I_m, O_m\}$ by one
 10. $TrafficToShift \leftarrow 0$
 11. **for each pair** (s, d) , $s \preceq i, j \preceq d$ **do**
 12. $TrafficToShift \leftarrow TrafficToShift + \min\{t_{im}^{(sd)}, t_{mj}^{(sd)}\}$
 13. **if** $TrafficToShift > TotalToShift$ **then break**
// No more (s, d) pairs needed
 14. **endfor** // of the (s, d) loop
 15. **if** $TrafficToShift < TotalToShift$ **then break**
// Cannot replace lightpaths (i, m) or (m, j) ; continue w/ next (i, j) pair
else if replacing lightpaths (i, m) or (m, j) would violate any
wavelength limit constraints **or** would create new maximum nodal
degrees at i or j **then break**
else
 16. Remove lightpaths (i, m) , (m, j) // $\max\{I_m, O_m\}$ decreases by 1
 17. **for all pairs** (s, d) contributing to $TrafficToShift$ **do**
 18. Reduce $t_{im}^{(sd)}$, $t_{mj}^{(sd)}$ by contributing amount, add equal amount to $t_{ij}^{(sd)}$
// This step reflects the new routing over the lightpath (i, j)
 19. **endif**
 20. **endfor** // of the (i, j) loop
 21. **until** no decrease in the max degree at any node is possible
 22. **end** // of the algorithm
-

Fig. 1. Algorithm for logical topology and traffic routing

In practice, however, the algorithm runs much faster than the above worst-case analysis indicates, taking on the order of few tens of milliseconds for any problem instance with $N = 16$ nodes and $W = 128$ wavelengths.

3.2 An Algorithm for Wavelength Assignment

The output of the above algorithm is a set of lightpaths R between pairs of ring nodes, and a routing of the traffic elements $\{t^{(sd)}\}$ over these lightpaths. While the algorithm guarantees that the resulting logical topology is such that no link carries more than W wavelengths, it may not be possible to color the lightpaths in R using no more than W wavelengths. We now present a polynomial-time algorithm to perform wavelength assignment with at most W colors; the tradeoff in ensuring that the number of wavelengths does not exceed W is a modification of the logical topology (i.e., the set R) which may result in an increase in the degree of some node in the ring. Consequently, the objective F of our optimization problem may increase. Therefore, we then refine the new logical topology to decrease the objective F .

Let us start by describing how to assign wavelengths to the lightpaths of set R . Our approach is based on the observation that, while the wavelength assignment problem is hard for ring networks, it is solvable in linear time in paths. Consider some node m of the ring. Let R_1 denote the lightpaths in R which *optically bypass* node m , and let $R_2 = R - R_1$ be the set of remaining lightpaths. The lightpaths in set R_2 can be viewed as the logical topology on a path network, and thus, can be colored using no more than W wavelengths. Now consider all the lightpaths in set R_1 . It may be possible to color some of them without violating any wavelength continuity constraints; in general, however, there may be some lightpaths in this set that cannot be colored without the need for additional wavelengths. In this case, we break such a lightpath (x, y) , $x \prec m \prec y$ into two lightpaths (x, m) and (m, y) . The new lightpaths do *not* bypass node m , and thus, can be colored along with the lightpaths in set R_2 using no more than W wavelengths. While breaking such a lightpath will increase the indegree and outdegree of node m by one, this approach guarantees a coloring of the new set of lightpaths that satisfies the wavelength constraints. The following steps describe our algorithm in more detail.

1. Let $R_1 \subset R$ be the set of lightpaths that optically bypass node m . Let $R_2 = R - R_1$ be the subset of remaining lightpaths.
2. Sort the lightpaths in R_2 in increasing order of their length.
3. Use the first-fit policy to color the lightpaths in R_2 . This step is always possible since it corresponds to a first-fit wavelength assignment for a path.
4. Sort the lightpaths in R_1 in decreasing order of their length.
5. Use the first-fit policy to color the lightpaths in R_1 . If lightpath $l = (x, y)$, $y \prec x$, cannot be colored, then: break l into two lightpaths, $l_1 = (x, m)$ and $l_2 = (m, y)$ which do not bypass node m ; increment the indegree and outdegree of node m to accommodate the new lightpaths; and repeat from Step 1 with $R_2 \leftarrow R_2 \cup \{l_1, l_2\}$ and $R_1 \leftarrow R_1 - \{l\}$.

Note that we arbitrarily select the node m above. An improvement would be to consider all N nodes, run the algorithm for each of them, and then select the solution with the least maximum nodal degree. However, we have found that the incremental improvement from such an approach does not justify the increase in running time, and thus we run the above algorithm just once.

Let Δ denote the increase in the nodal degree of node m after the termination of the above wavelength assignment algorithm. This increase is due to the fact that Δ lightpaths which optically bypassed node m under the initial logical topology defined by the set R , have now been broken into two lightpaths each. Of the 2Δ new lightpaths, Δ terminate at node m and Δ originate from it. Therefore, node m needs an additional Δ pairs of LTE, one for each of the original Δ lightpaths that used to bypass the node. Consequently, this increase of the objective F by Δ at node m increases the LTE cost in the ring network.

We now refine the new logical topology at the end of the wavelength assignment to improve on the objective F . Our approach is based on the observation that each additional pair of LTE at node m , one for an incoming and one for an outgoing lightpath, can be thought of as a *wavelength converter*. Indeed, consider, one of the Δ lightpaths that initially bypassed the node. This lightpath was broken by the algorithm into two shorter lightpaths that terminate at and originate from the node, respectively. This action was taken in Step 6 of the algorithm because it was not possible to assign the two shorter lightpaths on the two links in either side of node m the same color. Therefore, the additional pair of LTE at node m acts as a converter, changing the wavelength of the new incoming short lightpath to the wavelength of the new outgoing lightpath.

Consider a logical topology and corresponding feasible wavelength assignment on a ring network that requires a number Δ of converters at some node m . The recent study in [4] showed that it is possible to modify the wavelength assignment such that 2Δ converters are uniformly distributed across all N ring nodes (i.e., each node has at most $\lceil 2\Delta/N \rceil$ converters). Therefore, we use this algorithm to distribute the Δ pairs of LTE (i.e., “converters”) at node m to the other ring nodes. As a result, the maximum degree at all the ring nodes will increase, but the maximum nodal degree of the network (i.e., of node m) will decrease, resulting in a new logical topology with a smaller value for the objective F . For the details of this algorithm, the reader is referred to [4].

4 Numerical Results

We present experiments to demonstrate the performance of our traffic grooming algorithm. The experiments are characterized by the following parameters: the traffic pattern, the number N of nodes in the ring, the number W of wavelengths per link, the capacity C of each wavelength, and the load L on the link carrying the most traffic. The maximum amount of traffic that can flow through a link is WC , hence we express the load L as a percentage of WC . For each experiment (i.e., each set of values for the above parameters), we generate

50 problem instances, and we compare the following values for each instance: the value of the objective F (i.e., the maximum nodal degree) returned by our algorithm; the *lower bound* F^l on the objective F , which can be obtained as: $F^l = \max_s (\max (\lceil \sum_d t^{(sd)} / C \rceil, \lceil \sum_d t^{(ds)} / C \rceil))$; and the value F^e of the objective for a network using *all-electronic routing*, which corresponds to the number of wavelengths needed to carry the traffic on the link with the heaviest traffic load: $F^e = \max_l \lceil \sum_{(x,y) \in B(l)} t^{(sd)} / C \rceil$. We consider three traffic patterns in our study: uniform, random, and locality; for details on how these patterns are generated, refer to [5].

For small rings (i.e., $N \leq 8$), we have obtained the optimal solution by using CPLEX to solve an ILP formulation (see [5]) of the min-max traffic grooming problem. We have found that our algorithm produces solutions at, or very close to, the optimal, for all traffic patterns we considered. Unfortunately, CPLEX fails to terminate for instances with $N > 8$ even after running for several days.

We now present results for rings with $N = 16$ and $W = 128$; such networks are of practical interest because the maximum size of a SONET ring is 16 nodes, and also because WDM links supporting 128 wavelengths are becoming commercially available. We note that our algorithm needed less than a second to find a solution for each problem instance that we present in this section.

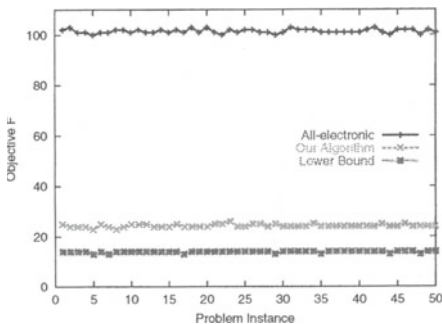


Fig. 2. Uniform pattern, $N = 16$, $W = 128$, $C = 12$, $L = 80\%$

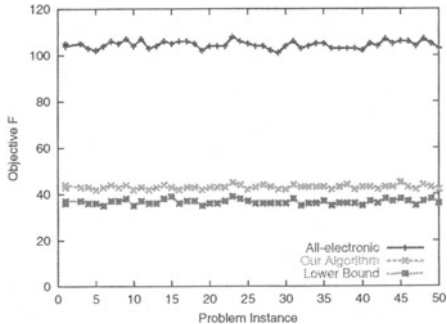


Fig. 3. Locality pattern, $N = 16$, $W = 128$, $C = 12$, $L = 80\%$

Figures 2, 3, and 4 present the results of 50 different problem instances for each of the uniform, locality, and random traffic patterns, respectively. As we can see, the values of the objective F returned by our algorithm are close to the lower bound, and track it well across the different traffic patterns and problem instances within each pattern. The all-electronic solutions, on the other hand, are quite high with values at around LW . This result indicates that our approach of minimizing the maximum LTE cost in the network can produce significant cost savings. Another related observation is that, by using a Min-Max objective, we can ensure that the cost of any individual network node (as well as the overall network cost) is determined by the traffic demands of the node, and will not




Fig. 4. Random pattern, $N = 16$, $W = 128$, $C = 12$, $L = 40 - 60\%$

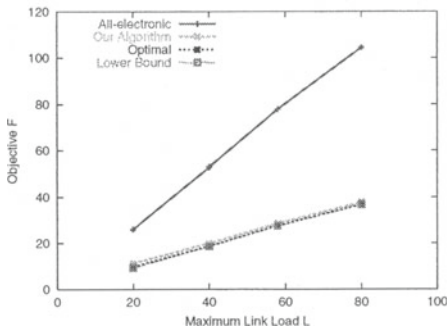


Fig. 5. Uniform pattern, $N = 6$, $W = 128$, $C = 12$, various loads

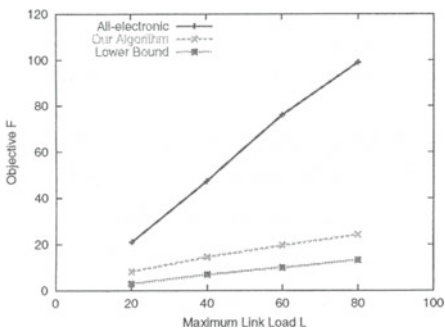


Fig. 6. Uniform pattern, $N = 16$, $W = 128$, $C = 12$, various loads




Fig. 7. Uniform pattern, $N = 16$, $W = 128$, $C = 12$, $L = 95\%$

scale with the number of wavelengths. Finally, from the three figures we see that the optimal value of the objective depends on both the traffic pattern and the maximum link load L . This observation suggests that, for the purposes of dimensioning the network for unknown future demands, the network designer should run our algorithm for a variety of traffic patterns and loads, and then equip each node with an amount of LTE determined by the highest value of F .

In Figures 5 and 6, we plot the values of F^e , F^l , and F against the load L for a uniform traffic pattern; each point in the figures is the average over 50 problem instances for the given value of L (the optimal values F^* are also plotted for the small network in Figure 5). The experiments presented in the two figures are identical in all respects except that Figure 5 is for 6-node rings, while Figure 6 is for 16-node rings. We have obtained similar results for the other two patterns as well. As we can see, F^e , F^l , and F all increase linearly with the load L . However, the curve corresponding to the all-electronic solution has a slope much steeper than that of the curves corresponding to the lower bound and our solution. This behavior demonstrates that the cost benefit of our optimization approach increases with the load of the network. This property is an important

one, and it implies that network operators will be able to operate the network at high loads with only an incremental increase in cost.

Another important property of the Min-Max optimization approach becomes evident once we compare the curves of Figure 5 to those of Figure 6. Specifically, we note that the curves corresponding to the all-electronic solutions have similar slopes in both figures. On the other hand, the curves corresponding to the lower bound and our solution have a steeper slope in Figure 5 (which corresponds to a 6-node network) than in Figure 6 (for a 16-node network). Equally important, for a given load L , the actual values of the lower bound and our solution are *lower* for the larger network than the respective values for the smaller network; this is true despite the fact that link loads of both networks are similar and that the 16-node network carries a much larger aggregate amount of traffic than the 6-node network. Based on this observation, we conclude that our solution approach scales well not only with the load and number of wavelengths, as we explained above, but also with the size of the network.

Finally, in Figure 7, we present results for problem instances with $N = 16$ nodes, a uniform traffic pattern, and a high load $L = 95\%$. Let us compare Figure 7 to Figure 2 which plots results for the same values of the system parameters, except that the load $L = 80\%$ in Figure 2. As we can see, the all-electronic solution F^e is about 20% higher in Figure 7, as expected, and the lower bound is also somewhat higher. However, there is little difference in the values of F returned by our algorithm when the load increases from 80% to 95%, reinforcing our earlier conclusion that the benefit of Min-Max optimization increases with the traffic load, enabling cost-effective network operation at very high loads.

5 Concluding Remarks

We studied a variant of traffic grooming in rings, where the objective is to minimize the maximum LTE cost at any node. We presented an algorithm for grooming and wavelength assignment, and we showed that it performs well under a wide range of traffic patterns and system parameter values. Our algorithm can be used to dimension networks for unknown and/or dynamic traffic demands.

References

1. Dutta, R., Rouskas, G.N.: A survey of virtual topology design algorithms for wavelength routed optical networks. *Optical Networks* **1** (2000) 73–89
2. Dutta, R., Rouskas, G.N.: Traffic grooming in WDM networks: Past and future. *IEEE Network* **16** (2002) 46–56
3. Dutta, R., Rouskas, G.N.: On optimal traffic grooming in WDM rings. *IEEE Journal on Selected Areas in Communications* **20** (2002) 110–121
4. Chen, L.W., Modiano, E.: Efficient routing and wavelength assignment for reconfigurable WDM networks with wavelength converters. In: *IEEE INFOCOM*. (2003)
5. Chen, B., Rouskas, G.N., Dutta, R.: Traffic grooming in WDM rings with the Min-Max objective. Technical report, North Carolina State Univ., Raleigh, NC (2003)