Delay Constrained Traffic Grooming in WDM Ring Networks

Arun Vishwanath, Weifa Liang
Department of Computer Science
The Australian National University
Canberra, ACT-0200, Australia
Email: {arunv, wliang}@cs.anu.edu.au

Abstract

In this paper we study the end-to-end delay constrained traffic grooming problem in WDM ring networks. Our aim is to incorporate Quality of Service (QoS) routing constraints into traffic grooming and address them jointly with the objective of maximizing the network throughput. It is well known that many real-time multimedia traffic not only make use of a fraction of the total wavelength capacity, but also have stringent end-to-end delay requirements. Consequently, while provisioning delay-bounded sub-wavelength traffic, it is of paramount importance to take traffic grooming and QoS routing constraints into consideration simultaneously to reduce the total network cost and improve the overall network performance. In this paper we first present an Integer Linear Program (ILP) formulation for the problem, which is applicable when the problem size is small. We then propose three scalable heuristic algorithms. We finally conduct experiments by simulation to evaluate the performance of the proposed algorithms. The experimental results show that, among the three proposed heuristics, the one based on ILP relaxation offers the best performance.

1. Introduction

Optical networks employing Wavelength Division Multiplexing (WDM) techniques have clearly emerged as the most promising candidate capable of meeting the increasing bandwidth demand of the Internet [1]. In WDM optical networks, different data can be concurrently transmitted on different wavelengths which are multiplexed within a single optical fiber. A routing path between a source node and a destination node is established using *lightpaths*, where a lightpath is an all-optical transmission medium that uses the same wavelength along all the links in its path. The set of lightpaths in the network constitutes the virtual topology of the network. The available bandwidth on a single wavelength is in the order of OC-48/OC-192 (2.48Gbps/10Gbps). It is expected that wavelength capacities of the order of OC-768 (40Gbps) and beyond are likely to be commercially available in the foreseeable future. However, there is a huge disparity between the bandwidth requirement of user connection requests and the available bandwidth on wavelengths. For example, High Definition Television usually requires *20Mbps* only. As wavelengths are critical resources in WDM optical networks, the available bandwidth on each wavelength needs to be used efficiently. Clearly, using a lightpath to realize just one sub-wavelength connection request results in poor network utilization and the degradation of network throughput. To alleviate this disparity, *traffic grooming* techniques have been proposed in the literature [2], which involve grooming several sub-wavelength connection requests onto a single high-capacity wavelength channel, thereby reducing the network cost and improving the network throughput.

In many real-time multimedia applications, various QoS parameters such as bandwidth, jitter and delay play important roles in the establishment of different service level agreements (SLA) between the customer and the network service provider. One such important QoS metric is the *end-to-end delay bound*, which requires that the total delay experienced by any successfully established request be no greater than the delay bound specified a priori. Support of sustained multimedia streams over long periods requires the establishment of dedicated, delay bounded routing paths between source-destination pairs.

It is well known that grooming of sub-wavelength traffic onto a lightpath can only be performed at grooming nodes where the lightpath has been converted from the optical domain into the electronic domain. In other words, grooming involves optical-electronic-optical processing at the nodes. When a connection request is routed on a *multi-hop* routing path (using two or more lightpaths) between a source node and a destination node, then at each intermediate hop in the path[1], the sub-wavelength traffic on an incoming lightpath is dropped and subsequently regenerated and groomed on an outgoing lightpath. As a result, the connection request incurs considerable time delays at each hop due to various O-E-O processing associated with grooming sub-wavelength traffic. These overheads cannot be ignored and have to be taken into account when provisioning connection requests that accommodate highly time dependent multimedia data streams. For large networks in particular, the delay incurred at each intermediate hop is of practical importance since the accumulated delay can be significant. Consequently, SLAs are not fully satisfied and users may not tolerate such delays. Thus, to offer desirable and consistent QoS service, it is of paramount importance for network service providers to incorporate grooming delays into the traffic grooming problem, especially when connection requests with guaranteed end-to-end delay requirements have to be provisioned. This motivates the research in this paper.

[1]Each intermediate hop along a multi-hop routing path refers to the terminating/originating nodes of lightpaths carrying the connection request.

1-4244-0419-3/06/$20.00 ©2006 IEEE

2. Related Work

Traffic grooming in WDM/SONET ring networks has been extensively investigated [2–9]. A brief overview of the traffic grooming problem in WDM/SONET rings was given by Modiano and Lin in [2]. Gerstel et al. [3] consider traffic grooming in WDM ring networks with the objective to reduce the overall SONET transmission equipment cost in the network. They solve the problem in a single step by integrating traffic grooming and lightpath assignment. Zhang and Qiao [4] propose algorithms for traffic grooming and wavelength assignment in SONET/WDM rings to reduce the number of wavelengths and the number of ADMs. Chiu and Modiano [5] focus on unidirectional SONET/WDM ring networks to minimize the number of ADMs. They show that the general traffic grooming problem is NP-Complete, and propose algorithms for special cases that have good performance. Dutta and Rouskas [6] approach the problem of designing a virtual topology to minimize electronic routing in a WDM ring network by presenting a framework to obtain lower and upper bounds on the amount of traffic electronically routed in the network. Gerstel et al. [7] study network designs for WDM optical rings that minimize the overall network cost rather than just the number of wavelengths. The authors include the cost of transceivers required at the nodes and the number of wavelengths on the links as metrics for the total network cost. Wang et al. [8] provide ILP formulations along with greedy algorithms and simulated-annealing techniques to solve the traffic grooming problem in WDM ring networks. Chen et al. [9] give a Min-Max objective to minimize the maximum number of lightpaths originating or terminating at a node. The problem is shown to be NP-Complete and a polynomial time algorithm is also presented.

It must be mentioned that none of the above work explicitly considered the end-to-end delay constrained traffic grooming problem in WDM ring networks. The only work that we are aware of is by Yoon [10], who considers traffic grooming in WDM ring networks along with lightpath hop-count constraints simultaneously. The end-to-end delay bound is represented by the hop-count constraint, which limits the number of lightpaths a particular connection request can use. The objective is to minimize the number of wavelengths needed to realize all the connection requests. The problem is formulated as a mixed integer linear program (MILP) and a heuristic algorithm is proposed. The work in this paper essentially differs from theirs in the following way. Our objective, unlike the former, is to maximize the network throughput, given the set of traffic matrices and the number of wavelengths on each link. Furthermore, we also incorporate the constraint on the number of transceivers at the nodes into our formulation, which was not considered by his work. This makes our formulation more robust. In addition, the performance results described in their paper are based on relatively small sized ring networks (4 to 10 nodes), and the maximum number of connection requests is limited to only 45. In contrast, we consider a moderately large-sized ring network (16 nodes) with as many as 768 connection requests. Each traffic matrix we consider has 256 connection requests since the traffic is uniformly distributed between all 16 node pairs. We consider three different traffic matrices at the same time, thus the total number of connection requests is $3 \times 256 = 768$. Our algorithms are thus more suitable for practical sized networks.

The rest of the paper is organized as follows. In Section 3, we introduce the node architecture for traffic grooming. In Section 4, we define the end-to-end delay constrained traffic grooming problem in WDM ring networks. We present the ILP formulation in Section 5 and propose several heuristics in Section 6. We evaluate the performance of the proposed algorithms in Section 7. We conclude the paper in Section 8.

3. Node Architecture

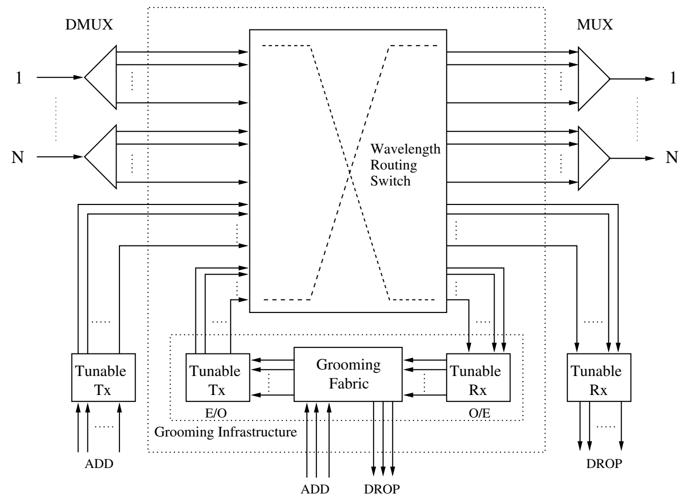

Figure 1. Node architecture for traffic grooming

A WDM ring network consists of optical switching nodes with communication fiber links interconnecting the nodes. Each fiber link carries a certain number of wavelengths, and each wavelength carries traffic from one or more sub-wavelength connection requests. The node architecture for traffic grooming in a WDM ring network is shown in Figure 1. The wavelength routing switch, consisting of all-optical switches, serves as an optical bypass to all traffic that are not destined to the node. Traffic that do not terminate at the node are switched in the optical domain from different input ports to different output ports. Since all-optical wavelength conversion is prohibitively expensive, these switches do not have any wavelength conversion capability. As a result, when a wavelength is switched from an input port to a possibly different output port, the *wavelength continuity constraint* is obeyed, i.e. the traffic on an incoming wavelength is switched on the same wavelength to a corresponding output port.

The node component responsible for grooming sub-wavelength traffic is the grooming infrastructure. It consists of tunable transmitters and receivers along with the grooming fabric. The O-E-O processing associated with traffic grooming takes place in the grooming infrastructure. Tunable receivers first convert the incoming optical signal (wavelength) from the optical domain into the electronic domain. The grooming fabric, with all the electronic processing functionality, performs multiplexing and demultiplexing of sub-wavelength traffic onto or from high capacity wavelengths. Sub-wavelength traffic on incoming wavelengths destined to the node is demultiplexed and dropped locally at the node, while sub-wavelength traffic originating at the node is added and multiplexed onto high capacity wavelengths. The tunable transmitters convert these multiplexed signals from the electronic domain back into the optical domain, following which, the wavelength routing switches switch the resulting wavelength traffic to the respective outgoing ports.

4. Problem Definition

The end-to-end delay constrained traffic grooming problem is formally stated as follows.

- Inputs to the problem

1. A ring network $G = (V, E)$ represents the physical topology, consists of $|V| = n$ nodes and $|E| = m$ links interconnecting the nodes. Each link in G is bidirectional and modeled as a pair of unidirectional links. We assume that all the nodes have grooming capability.
2. $W = \{\lambda_1, \lambda_2, \lambda_3, \ldots\ldots, \lambda_q\}$ is the set of available wavelengths in the network. We assume that each link in the network is assigned the same number of wavelengths.
3. Tx_v and Rx_v denote the number of transceivers (transmitters and receivers) available at node $v \in V$. We assume that all the nodes in G are equipped with the same number of transceivers that can be tuned to any wavelength among the set of wavelengths in W.
4. C is the capacity of each wavelength channel. In this work we assume C to be OC-48.
5. OC_y traffic matrices, where $y \in \{1, 3, 12\}$, representing the OC-1, OC-3 and OC-12 sub-wavelength connection requests between source-destination pairs.
6. Delay matrices Δ_y^{sd}, where $y \in \{1, 3, 12\}$, representing the end-to-end delay bound for each OC_y sub-wavelength connection request between source node s and destination node d.
7. δ_v is the grooming delay incurred at node $v \in V$ due to O-E-O conversion and other associated overheads.

Given the above inputs, we aim to determine a virtual topology and route the connection requests on the virtual topology so that the network throughput is maximized, provided that all the established connection requests meet their individual end-to-end delay bounds. Here the *network throughput* is defined as the sum of the bandwidth of all established connection requests, given the traffic matrices.

5. Integer Linear Program Formulation

In this section, we present an integer linear program formulation for the problem by extending the work in [12]. Following are the notations that will be used in the formulation.

- Integer variables

1. v_{ij} denotes the number of lightpaths between nodes i and j.
2. v_{ij}^{w} denotes the number of lightpaths between nodes i and j on wavelength $\lambda_w \in W$.
3. $p_{mn,w}^{ij} = 1$ if a lightpath between nodes i and j is routed through an intermediate physical link (m, n) on wavelength w; otherwise $p_{mn,w}^{ij} = 0$.
4. $l_{ij}^{sd,y} = 1$ if the OC_y traffic request between nodes s and d uses lightpath between nodes i and j as an intermediate virtual link; otherwise $l_{ij}^{sd,y} = 0$.
5. $s_y^{sd} = 1$ if the OC_y request between nodes s and d is successfully established; otherwise $s_y^{sd} = 0$.

The objective is to maximize the network throughput, which can be expressed as follows.

$$Maximize \quad \sum_{y,s,d} y \cdot s_y^{sd} \tag{1}$$

Subject to the following constraints:

1. Transceiver and virtual topology constraints

$$\forall i: \quad \sum_j v_{ij} \le Tx_i \tag{2}$$

$$\forall j: \quad \sum_i v_{ij} \le Rx_j \tag{3}$$

$$\forall i, j: \quad \sum_w v_{ij}^{w} = v_{ij} \tag{4}$$

$$\forall i, j: \quad v_{ij} \ge 0 \tag{5}$$

$$\forall i, j, w: \quad 0 \le v_{ij}^{w} \le 2 \tag{6}$$

2. Physical topology routing

$$\forall i, j, w: \quad \sum_m p_{mi,w}^{ij} = 0 \tag{7}$$

$$\forall i, j, w: \quad \sum_n p_{jn,w}^{ij} = 0 \tag{8}$$

$$\forall i, j, w, k; k \ne i, j: \quad \sum_m p_{mk,w}^{ij} = \sum_n p_{kn,w}^{ij} \tag{9}$$

$$\forall i, j, w: \quad \sum_n p_{in,w}^{ij} = v_{ij}^{w} \tag{10}$$

$$\forall i, j, w: \quad \sum_m p_{mj,w}^{ij} = v_{ij}^{w} \tag{11}$$

$$\forall m, n, w; (m, n) \in E: \quad \sum_{i,j} p_{mn,w}^{ij} \le 1 \tag{12}$$

3. Virtual topology routing

$$\forall s, d, y \in \{1, 3, 12\}: \quad \sum_j l_{sj}^{sd,y} = s_y^{sd} \tag{13}$$

$$\forall s, d, y \in \{1, 3, 12\}: \quad \sum_i l_{id}^{sd,y} = s_y^{sd} \tag{14}$$

$$\forall s,d,k,y \in \{1,3,12\}; k \neq s,d: \quad \sum_i l_{ik}^{sd,y} = \sum_j l_{kj}^{sd,y} \tag{15}$$

$$\forall s,d,y \in \{1,3,12\}: \quad \sum_i l_{is}^{sd,y} = 0 \tag{16}$$

$$\forall s,d,y \in \{1,3,12\}: \quad \sum_j l_{dj}^{sd,y} = 0 \tag{17}$$

$$\forall i,j: \quad \sum_y \sum_{s,d} y \times l_{ij}^{sd,y} \leq C \times v_{ij} \tag{18}$$

4. End-to-end delay bounds

$$\forall s,d,y \in \{1,3,12\}: \quad \sum_{j,i \neq s} \delta_i \cdot l_{ij}^{sd,y} \leq \Delta_y^{sd} \tag{19}$$

Equations (2) and (3) ensure that the total number of lightpaths between a source node i and a destination node j cannot exceed the number of transmitters at node i and the number of receivers at node j.

Equation (4) shows that the number of lightpaths established between source-destination pair (i,j) is the sum of individual lightpaths between i and j on different wavelengths.

Equation (5) represents the condition that there can be zero or more lightpaths established between a pair of nodes i and j.

Equation (6) ensures that, at most two lightpaths can exist between any pair of nodes i and j on a given wavelength w. One lightpath can be setup in the clockwise direction, while the other can be setup in the anti-clockwise direction.

Equation (7) ensures that, at source node i of lightpath (i,j), the total number of incoming lightpaths on wavelength w is 0.

Equation (8) ensures that, at the destination node j of lightpath (i,j), the total number of outgoing lightpaths on wavelength w is 0.

Equation (9) ensures that, at any intermediate node k along the physical route of the lightpath between nodes i and j, the total number of incoming lightpaths on wavelength w must be equal to the total number of outgoing lightpaths.

Equation (10) ensures that, at the source node i of lightpath (i,j), the total number of outgoing lightpaths on wavelength w must be equal to the total number of lightpaths between nodes i and j on wavelength w.

Equation (11) ensures that, at the destination node j of lightpath (i,j), the total number of incoming lightpaths on wavelength w must be equal to the total number of lightpaths between nodes i and j on wavelength w.

Equations (7) to (11) above represent the flow conservation equations for routing the lightpath on the physical topology.

Equation (12) ensures that no two lightpaths can be routed on the same wavelength w along the same physical link $(m,n) \in E$. In other words, if any two lightpaths share the same physical link $(m,n) \in E$, this constraint ensures that they will be routed on two distinct wavelengths.

Constraints in Equations (13) to (18) below denote routing on the virtual topology.

Equation (13) ensures that, for a given OC_y request between source node s and destination node d, if the request is successfully established, then there must be exactly one outgoing lightpath edge at s in the virtual topology.

Equation (14) ensures that, for a given OC_y request between source node s and destination node d, if the request is successfully established, then there must be exactly one incoming lightpath edge at d in the virtual topology.

Equation (15) represents the constraint for multi-hop routing on the virtual topology. To route an OC_y request between source node s and destination node d on multiple-hops, this equation ensures that, if there is an incoming lightpath edge at any node $k \neq s,d$ in the virtual topology, there must be a corresponding outgoing edge at k as well.

Equations (16) and (17) ensure that, for an OC_y request to be routed between source node s and destination node d, there must not be any incoming or outgoing lightpath edges at nodes s and d respectively.

Equation (18) represents the wavelength capacity constraint. It ensures that the net traffic flowing on a lightpath between nodes i and j cannot exceed the total wavelength capacity.

Equation (19) denotes the constraint for the end-to-end delay bound of the OC_y request between nodes s and d. This constraint effectively ensures that, if the OC_y request between node pair (s,d) is established, the total delay incurred by it will be less than or equal to the end-to-end delay bound specified a priori.

6. Heuristic Algorithms

The traffic grooming problem without end-to-end delay constraints in WDM ring networks is shown to be NP-Complete [5]. It follows that the problem we consider is NP-Complete as well. The ILP solution in the previous section is computationally intractable even for moderate size networks. Hence, we focus on devising efficient scalable heuristics. In this section, we propose three heuristic algorithms. In what follows, we use the term *demand* to refer to a connection request between a source node and a destination node. The terms demand and connection request will be used interchangeably.

6.1. Algorithm Greedy_OrderBy_Demands

Since different connection requests between the same source-destination pair can have potentially different end-to-end delay bounds, this heuristic attempts to establish lightpaths by considering them one by one. The basic idea of this heuristic lies in the way the demands are sorted. We adopt a greedy approach by sorting demands according to their bandwidth requirement in decreasing order. If the end-to-end delay bound for a demand is 0, it can only be routed on a single lightpath between its two endpoints. These demands are of high priority, thus, we sort all demands with the same bandwidth requirement in increasing order of the end-to-end delay requirements. Subsequently, demands with the same bandwidth and delay requirements are sorted according to their shortest-path hop count in increasing order. The notion behind sorting the demands in this manner is to ensure that the available wavelengths are prudently used to establish demands with high bandwidth requirement and stringent end-to-end delay bounds. Once a lightpath is established, the other demands between the endpoints of the lightpath are groomed on it such that its residual capacity is minimized. Demands that cannot be established on a single lightpath will be routed on the virtual topology, subject to meeting the end-to-end delay bounds. The detailed algorithm is given in Algorithm 1.

Algorithm 1 Greedy_OrderBy_Demands

1. Store all the individual demands from the traffic matrices (OC-1, OC-3 and OC-12) in a new traffic matrix array TM.
2. Sort all the demands in TM according to their bandwidth requirement in descending order.
3. Next, sort all the demands in TM that have the same bandwidth requirement according to their end-to-end delay bounds in ascending order.
4. For each demand in TM, determine the number of hops in the shortest path between the source node and the corresponding destination node. Finally, sort the demands that have the same bandwidth requirement and delay bound according to their hop counts in ascending order.
5. Identify the first source-destination node pair (s, d) in matrix TM. If there is at least one free transmitter at s and at least one free receiver at d, try to establish a lightpath between these two endpoints using shortest path routing and first-fit wavelength assignment policy.
6. If the lightpath can be successfully established, try to groom other demands between node pair (s, d) on this lightpath so as to maximize the lightpath capacity utilization. In other words, pack the lightpath with as much traffic as possible, subject to the available wavelength capacity. Delete the established demands from matrix TM and go to step 5.
7. If the lightpath establishment in step 5 fails, remove all the demands between node pair (s, d) from TM and put them into a secondary traffic matrix array STM. Repeat steps 5 to 7 until matrix TM is empty.
8. Make use of the unused transceivers and wavelengths at the end of step 7 by attempting to establish arbitrary lightpaths between node pairs in increasing order of their hop counts.
9. The complete virtual topology is generated at the end of step 8. Now, sort all the demands in matrix STM according to their bandwidth requirement in descending order.
10. Finally, attempt to route the demands in matrix STM over the virtual topology without violating their end-to-end delay bounds.

6.2. Algorithm Greedy_OrderBy_Net_Traffic

Unlike the previous heuristic, this algorithm does not take into account the end-to-end delay bounds of individual demands when attempting to establish lightpaths. Instead, the demands are sorted according to the total uncarried traffic between source-destination pairs in descending order (greedy policy), where uncarried traffic refers to the sum of the bandwidth of individual demands between source-destination pairs that are yet to be established. Once a lightpath is established, the demands between the endpoints of the lightpath are groomed on it such that its residual capacity is minimized. Demands that cannot be groomed on the lightpath are routed on the virtual topology, subject to meeting the end-to-end delay bounds. The detailed algorithm is given in Algorithm 2.

Algorithm 2 Greedy_OrderBy_Net_Traffic

1. Sort all the demands according to the net (total) uncarried traffic between source-destination pairs in descending order. Store these demands in a traffic matrix array TM.
2. Identify the first source-destination node pair (s, d) in matrix TM. If there is at least one free transmitter at s and at least one free receiver at d, try to establish a lightpath between these two endpoints using shortest path routing and first-fit wavelength assignment policy.
3. If the lightpath can be successfully established, route all the demands between node pair (s, d) directly on this lightpath so as to maximize the lightpath capacity utilization. Determine the remaining uncarried traffic between node pair (s, d) and update its corresponding value in matrix TM. Go to step 1.
4. If the lightpath establishment in step 2 fails, delete the corresponding (s, d) entry from matrix TM. Identify all the individual demands between node pair (s, d) and store them in a secondary traffic matrix array STM. Repeat steps 1 to 4 until matrix TM is empty.
5. Make use of the unused transceivers and wavelengths at the end of step 4 by attempting to establish arbitrary lightpaths between node pairs in increasing order of their hop counts.
6. The complete virtual topology is generated at the end of step 5. Now, sort all the demands in matrix STM according to their bandwidth requirement in descending order.
7. Finally, attempt to route the demands in matrix STM over the virtual topology without violating their end-to-end delay bounds.

6.3. Algorithm ILP_Relaxed_Solution

Rather than relying purely on greedy approaches to establish lightpaths as described in the previous two heuristics, this algorithm starts by establishing lightpaths using the results obtained by solving the corresponding Linear Program (LP). Once a lightpath is established, the demands between the endpoints of the lightpath are groomed on it such that its residual capacity is minimized. All demands that cannot be routed directly on a single lightpath are routed on the virtual topology using its spare capacity, subject to meeting the end-to-end delay bounds. The detailed algorithm is given in Algorithm 3.

7. Performance Evaluation

In this section, we first describe the simulation environment, and then present the experimental results of the proposed algo-

Algorithm 3 ILP_Relaxed_Solution

1. Formulate the ILP for the problem.
2. Solve the Linear Program of the problem by relaxing the integer constraints on the variables used in the ILP formulation.
3. Obtain the values for all the v_{ij}^{w} variables from the resulting LP solution and store them in matrix L. The variables i and j in v_{ij}^{w} indicate the source node and the destination node of the lightpath (i, j) on wavelength w. Sort matrix L according to the v_{ij}^{w} values in descending order.
4. Execute steps 1 to 4 of the algorithm Greedy_OrderBy_Demands.
5. Using shortest path routing and first-fit wavelength assignment policy, try to establish a lightpath between the first source-destination node pair (i, j) in matrix L, subject to the transceiver constraints.
6. If the lightpath can be successfully established, identify all the demands in matrix TM that are between the endpoints (i, j). Route these demands directly on the lightpath so as to maximize the lightpath capacity utilization. Delete the demands from matrix TM that were successfully routed on the lightpath. Also, delete node pair (i, j) from matrix L and go to step 5.
7. If the lightpath establishment in step 5 fails, delete the node pair (i, j) from matrix L. Identify all the individual demands in matrix TM that are between the endpoints (i, j). Delete these demands from matrix TM and store them in a secondary traffic matrix array STM. Repeat steps 5 to 7 until matrix L is empty.
8. Next, scan matrix TM for any unestablished demands between source node s and destination node d. Try to establish a lightpath between node pair (s, d) using shortest path routing and first-fit wavelength assignment policy, subject to the transceiver constraints.
9. If the lightpath can be successfully established, identify the remaining demands in matrix TM that are between the endpoints (s, d). Route these demands directly on the lightpath so as to maximize the lightpath capacity utilization. Delete all the demands from matrix TM that were successfully routed on the lightpath. Go to step 8.
10. If the lightpath establishment in step 8 fails, delete all the demands between the two endpoints (s, d) from matrix TM and store them in a secondary traffic matrix array STM. Repeat steps 8 to 10 until matrix TM is empty.
11. Make use of the unused transceivers and wavelengths at the end of step 10 by attempting to establish arbitrary lightpaths between node pairs in increasing order of their hop counts.
12. The complete virtual topology is generated at the end of step 11. Now, sort all the demands in matrix STM according to their bandwidth requirement in descending order.
13. Finally, attempt to route the demands in matrix STM over the virtual topology without violating their end-to-end delay bounds.

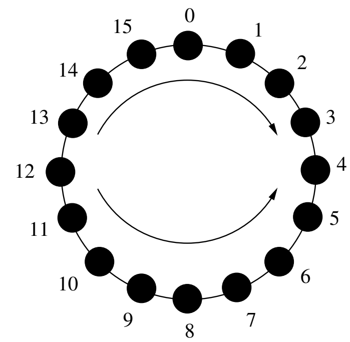

Figure 2. 16 node bidirectional ring network

rithms using the 16 node bidirectional ring network shown in Figure 2 as the physical topology. We consider three different traffic matrices - OC-1, OC-3 and OC-12, with the total wavelength capacity being OC-48. Following the same simulation environment setting in [12], the traffic matrices are generated as follows. The number of OC-1 connection requests between each node pair is a uniformly distributed random number between 0 and 16. The number of OC-3 connection requests between each node pair is a uniformly distributed random number between 0 and 8. Finally, the number of OC-12 connection requests between each node pair is a uniformly distributed random number between 0 and 2. The grooming delay at each node is randomly assumed to be 3 units. The end-to-end delay bound for the connection requests is a uniformly distributed random number from the set $\{0, 3, 6, 9, 12\}$. A connection request with a delay bound of 0 indicates that it can only be routed on a single lightpath (one hop) directly between its two endpoints, while the value 12 indicates that the connection request can be routed in at most 5 lightpath-hops. The total bandwidth of all the connection requests in our experiment is equivalent to OC-7668. Further, assume that the connection requests cannot be divided into several smaller connection requests and routed separately. Each connection request has to be routed along a single path between the source node and the destination node.

7.1. Experimental Results

In Figures 3 to 6, we plot the network throughput curves by varying the number of wavelengths on each fiber link and fixing the number of transceivers at each node at 10, 11, 12 and 13 respectively. To estimate the upper bound on the network throughput, we relax the original ILP to a corresponding LP. We then solve this LP directly using the CPLEX [15] linear optimizer (version 10). Following the theory of Linear Programming, the solution of the LP formulation specifies an upper bound on the solution of the original ILP. In addition, if this relaxed solution is a feasible solution to the problem, it is also the optimal solution to the problem. The upper bound curves are indicated by ILP_Relaxed_Upper_Bound.

Figure 3 shows the network throughput curves when the number of transceivers at each node is fixed at 10 and the number of available wavelengths on each link is varied from 5 to 15. It can be clearly observed from the figure that the ILP_Relaxed_Solution outperforms the other two heuristic algo-

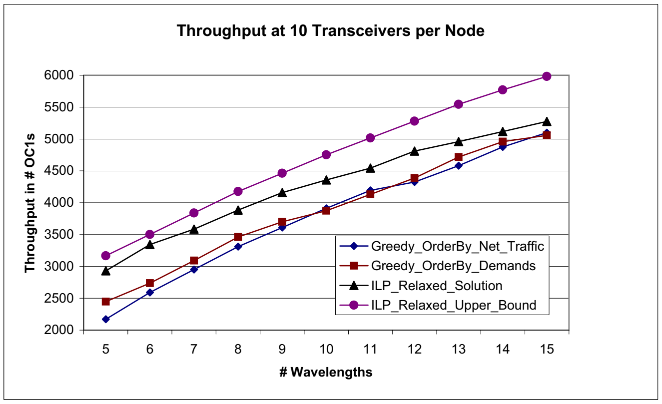

Figure 3. Throughput when number of transceivers is fixed at 10 and number of wavelengths varied from 5-15

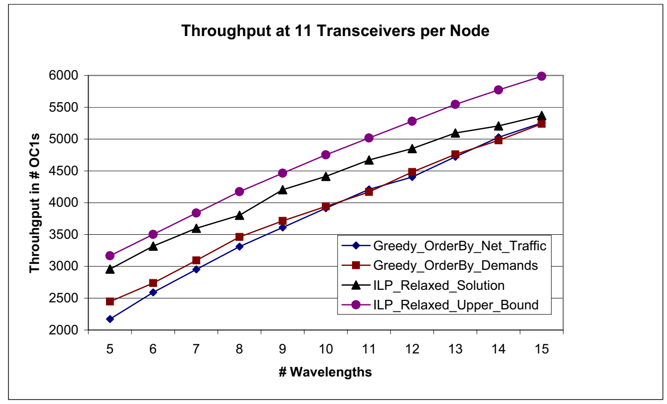

Figure 4. Throughput when number of transceivers is fixed at 11 and number of wavelengths varied from 5-15

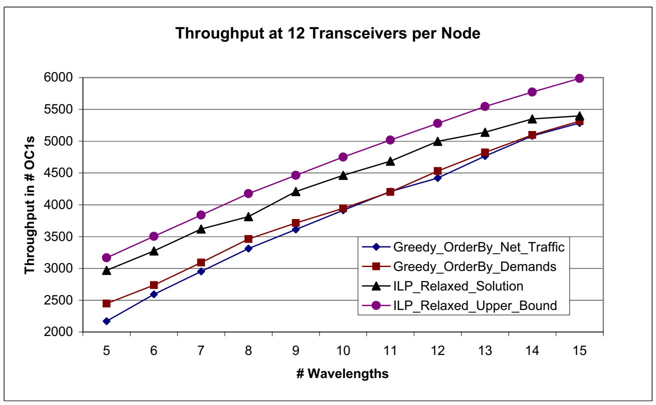

Figure 5. Throughput when number of transceivers is fixed at 12 and number of wavelengths varied from 5-15

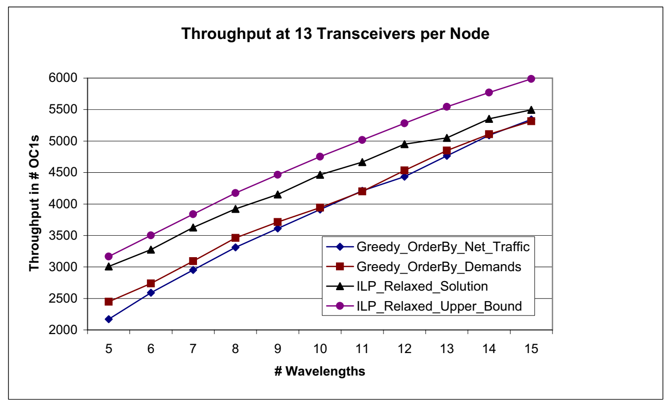

Figure 6. Throughput when number of transceivers is fixed at 13 and number of wavelengths varied from 5-15

rithms. The upper bound curves indicated in the figure are not feasible solutions to the problem because the LP solutions are not integer solutions. In fact, the actual upper bound will be *lesser* than the upper bounds shown in the figure. From this, we can infer that the difference in network throughput between the optimal solution and the ILP relaxed solution is actually *smaller* than that seen in the figure. This renders the network throughput from the ILP relaxed solution much closer to the optimal solution. It can also be seen that, when the number of available wavelengths is relatively small, Greedy_OrderBy_Demands heuristic performs slightly better than the Greedy_OrderBy_Net_Traffic heuristic. This suggests that, when the available wavelength resources are limited, it might be better to order and route demands individually (ILP_Relaxed_Solution, Greedy_OrderBy_Demands) as opposed to ordering them according to the net uncarried traffic (Greedy_OrderBy_Net_Traffic).

In the above network setting, the gain in the network throughput using the ILP_Relaxed_Solution is between 4% and 35% when compared with the Greedy_OrderBy_Net_Traffic heuristic. In comparison with the performance offered by the Greedy_OrderBy_Demands heuristic, the measured gain is between 4% and 22%. Similar trend can be observed in the other network configurations as well. When the nodes are equipped with 13 transceivers (see Figure 6), the gain in the network throughput varies between 3% and 39% when compared with the former heuristic, and between 3% and 23% when compared with the latter. These performance results clearly demonstrate the effectiveness of the relaxation technique.

Next, we plot the network throughput curves versus the number of transceivers when the number of wavelengths on each link is fixed at 9, 11, 13 and 15 respectively. These curves are shown in Figures 7 to 10. These results also confirm the high efficiency of the ILP_Relaxed_Solution over the other two algorithms. The network throughput gain achieved varying between 2% and 17%.

8. Conclusion

In this paper, we investigated the end-to-end delay constrained traffic grooming problem in WDM ring networks. Given individual traffic matrices, end-to-end delay bounds, wavelength capacity, the number of transceivers and the number of wavelengths, we first proposed an ILP formulation for the problem, which is applicable when the network size is small. We then proposed three scalable heuristics and finally evaluated the performance of the proposed algorithms by experimental simulations using a 16 node representative sized ring network. The experimental

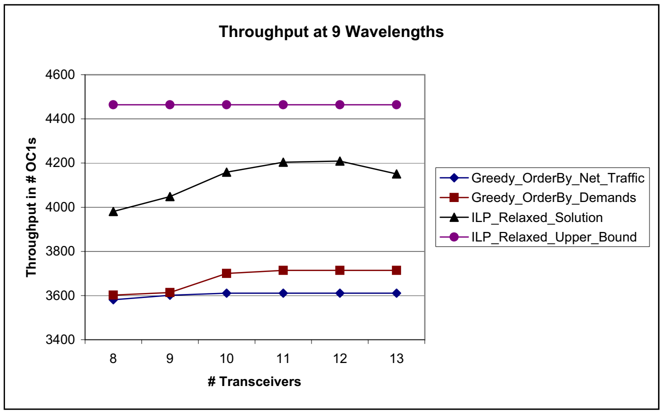

Figure 7. Throughput when number of wavelengths is fixed at 9 and number of transceivers varied from 8-13

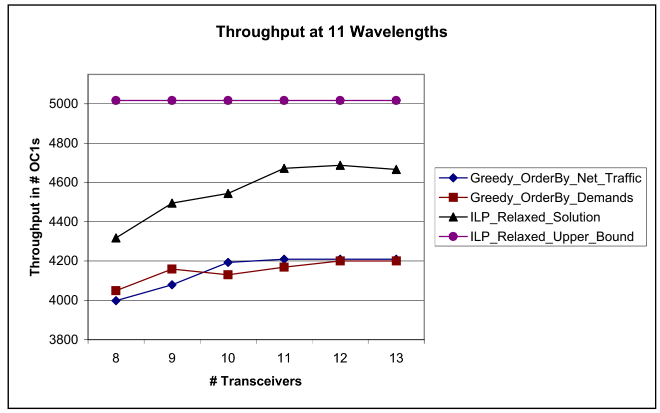

Figure 8. Throughput when number of wavelengths is fixed at 11 and number of transceivers varied from 8-13

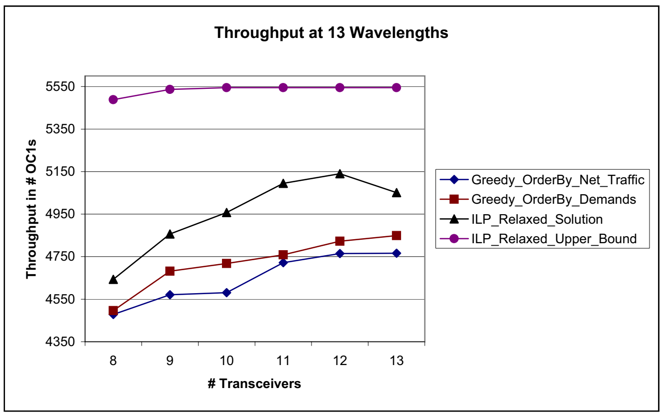

Figure 9. Throughput when number of wavelengths is fixed at 13 and number of transceivers varied from 8-13

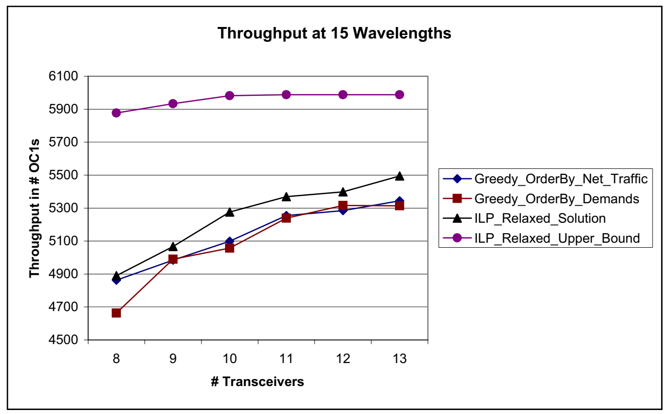

Figure 10. Throughput when number of wavelengths is fixed at 15 and number of transceivers varied from 8-13

results showed that the heuristic based on ILP relaxation outperforms the other two heuristic algorithms significantly. We are currently working on establishing theoretical upper bounds for the network throughput from the ILP formulation. Extending the work to solve the problem in the context of WDM mesh networks is also being investigated.

9. Acknowledgment

It is acknowledged that the work by the authors was supported by research grant #DP0449431, funded by the Australian Research Council under its Discovery Schemes.

References

[1] R. Ramaswami and K. N. Sivarajan, *Optical Networks: A Practical Perspective,* Morgan Kaufmann Publishers, 1998

[2] E. Modiano and P. Lin, "Traffic grooming in WDM networks," *IEEE Communications Magazine,* vol. 39, no. 7, pp. 124-129, July 2001

[3] O. Gerstel, P. Lin and G. Sasaki, "Combined WDM and SONET network design," *Proc. IEEE INFOCOM,* pp. 734-743, Mar/Apr 1999

[4] X. Zhang and C. Qiao, "An effective and comprehensive approach for traffic grooming and wavelength assignment in SONET/WDM rings," *IEEE/ACM Transactions on Networking,* vol. 8, no. 5, pp. 608-617, Oct. 2000

[5] A. L. Chiu and E. Modiano, "Traffic grooming algorithms for reducing electronic multiplexing costs in WDM ring networks," *IEEE Journal of Lightwave Technology,* vol. 18, no. 1, pp. 2-12, Jan. 2000

[6] R. Dutta and G. N. Rouskas, "On optimal traffic grooming in WDM rings," *IEEE Journal on Selected Areas in Communications,* vol. 20, no. 1, pp. 1-12, Jan. 2002

[7] O. Gerstel, R. Ramaswami and G. Sasaki, "Cost-effective traffic grooming in WDM rings," *IEEE/ACM Transactions on Networking,* vol. 8, no. 5, pp. 618-630, Oct. 2000

[8] J. Wang, W. Cho, V. R. Vemuri and B. Mukherjee, "Improved approaches for cost-effective traffic grooming in WDM ring networks: ILP formulations and single-hop and multihop connections," *IEEE Journal of Lightwave Technology,* vol. 19, no. 11, pp. 1645-1653, Nov. 2001

[9] B. Chen, G. N. Rouskas and R. Dutta, "Traffic grooming in WDM ring networks with the min-max objective," *Proc. IFIP NETWORKING,* pp. 174-185, 2004

[10] M. G. Yoon, "Traffic grooming and light-path routing in WDM ring networks with hop-count constraint," *Proc. IEEE ICC,* pp. 731-737, 2001

[11] K. Zhu and B. Mukherjee, "A review of traffic grooming in WDM optical networks: Architectures and challenges," *Optical Networks Magazine,* vol. 4, no. 2, pp. 55-64, Mar/Apr 2003

[12] K. Zhu and B. Mukherjee, "Traffic grooming in an optical WDM mesh network," *IEEE Journal on Selected Areas in Communications,* vol. 20, no. 1, pp. 122-133, Jan. 2002

[13] K. Zhu, H. Zang and B. Mukherjee, "A comprehensive study on next-generation optical grooming switches," *IEEE Journal on Selected Areas in Communications,* vol. 21, no. 7, pp. 1173-1186, Sep. 2003

[14] R. Ramaswami and K. N. Sivarajan, "Design of logical topologies for wavelength-routed optical networks," *IEEE Journal on Selected Areas in Communications,* vol. 14, no. 5, pp. 840-851, June 1996

[15] ILOG CPLEX, http://www.ilog.com/products/cplex