Improved Approaches for Cost-Effective Traffic Grooming in WDM Ring Networks: ILP Formulations and Single-Hop and Multihop Connections

Jian Wang, Wonhong Cho, V. Rao Vemuri, and Biswanath Mukherjee, *Member, IEEE*

***Abstract*—Traffic grooming is the term used to describe how different traffic streams are packed into higher speed streams. In a synchronous optical network–wavelength division multiplexing (SONET–WDM) ring network, each wavelength can carry several lower-rate traffic streams in time division (TDM) fashion. The traffic demand, which is an integer multiple of the timeslot capacity, between any two nodes is established on several TDM virtual connections. A virtual connection needs to be added and dropped only at the two end nodes of the connection; as a result, the electronic add–drop multiplexers (ADMs) at intermediate nodes (if there are any) will electronically bypass this timeslot. Instead of having an ADM on every wavelength at every node, it may be possible to have some nodes on some wavelength where no add–drop is needed on any timeslot; thus, the total number of ADMs in the network (and, hence, the network cost) can be reduced. Under the static traffic pattern, the savings can be maximized by carefully packing the virtual connections into wavelengths. In this work, we allow arbitrary (nonuniform) traffic and we present a formal mathematical definition of the problem, which turns out to be an integer linear program (ILP). Then, we propose a simulated-annealing-based heuristic algorithm for the case where all the traffic is carried on directly connected virtual connections (referred to as the single-hop case). Next, we study the case where a hub node is used to bridge traffic from different wavelengths (referred to as the multihop case). We find the following main results. The simulated-annealing-based approach has been found to achieve the best results, so far, in most cases, relative to other comparable approaches proposed in the literature. In general, a multihop approach can achieve better equipment savings when the traffic-grooming ratio is large, but it consumes more bandwidth.**

***Index Terms*—Linear programming, multihop connection, nonuniform traffic, simulated annealing, single-hop connection, synchronous optical network–wavelength division multiplexing (SONET–WDM) ring, traffic grooming.**

NOMENCLATURE

N	Number of nodes in the network.
W	Number of wavelengths in the network (each wavelength can transmit several circles in time-division fashion).
C	Grooming ratio, which is the number of circles a wavelength can carry.
$T(t_{ij})$	Nonuniform traffic matrix in which t_{ij} represents the traffic from node i to j.
${}^{d}V_{ij}^{cw}$	Virtual connection from node i to node j on circle c, wavelength w. The d represents the direction of a connection and it can be either clockwise or counterclockwise.
O_i, I_i	In the multihop case, O_i represents the virtual connection that starts from the node i and terminates at the hub node. Similarly, I_i represents the virtual connection that starts from the hub node and terminates at node i.
ADM_i^w	Number of ADMs at node i on wavelength w.
e	A link on the physical ring.

Manuscript received April 5, 2000; revised August 9, 2001. This work was supported in part by Intelligent Fiber Optic Systems (IFOS) and the State of California under UC MICRO Program Grant 99-086.

The authors are with the Networks Laboratory, Computer Science Department, University of California, Davis, CA 95616 USA (e-mail: jnwang@ucdavis.edu; cho@cs.ucdavis.edu; rvemuri@das.ucdavis.edu; mukherje@cs.ucdavis.edu).

Publisher Item Identifier S 0733-8724(01)09293-3.

I. INTRODUCTION

SYNCHRONOUS optical network (SONET) ring is a widely deployed optical transport technology because of its high capacity and inherent reliability. In a SONET–wavelength division multiplexing (WDM) ring, each wavelength running at the line rate of OC-N can carry several low-speed OC-M (M $\leq$ N) traffic channels in TDM fashion. Note that, for nonuniform traffic, each connection can have a different OC-M rate. The traffic demand, which is an integer multiple of the timeslot capacity, between any two nodes is established on several TDM virtual connections. A virtual connection needs to be added and dropped only at the two end nodes of the connection; as a result, the electronic add–drop multiplexers (ADMs) at intermediate nodes (if there are any) will electronically bypass this timeslot. It is possible to have some nodes on some wavelength where no add–drop is necessary in any timeslot, so the electronic equipment can be saved. Fig. 1 shows the architecture of a typical node in a SONET–WDM ring network. For some wavelengths ($\lambda 1$ in this example), because there is no need to add or drop any of its timeslots, they can be optically bypassed at the node. For other wavelengths ($\lambda 2$ and $\lambda 3$) where at least one timeslot needs to be added or dropped, an electronic ADM is used.

As shown in Fig. 1, ADMs do not have the timeslot interchange function, and wavelength conversion is not possible without additional equipment, so there are timeslot-continuity and wavelength-continuity constraints at nodes where only ADMs are used. In this study, we refer to these rings that are built only with ADMs as single-hop rings because all the connections are direct connections. A SONET digital cross-connect (DXC) can be used at a node to consolidate or segregate subchannels, but it is a relatively more expensive

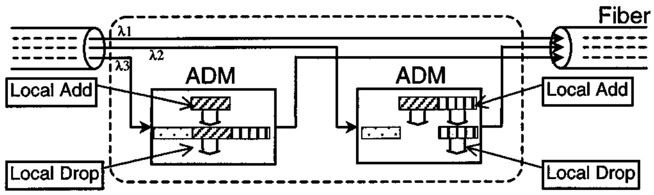

Fig. 1. Architecture of a WDM ring network node.

device. In this study, we consider the case that a DXC is used at only one node of a ring and refer to this kind of network as a multihop ring. For multihop networks, the cost of a DXC is estimated by its add–drop port count.

If a network is not properly designed, more ADMs may be needed to carry the same traffic requirement. This problem may be undesirable in metro-area and enterprise networks where there are few users sharing the system cost. In this study, we explore how to minimize the network-wide ADM cost a under static (possibly nonuniform) traffic pattern, which is expected to be the main operation model for the optical layer [1].

A. Example of Traffic Grooming in a SONET–WDM Ring Network

Fig. 2 shows a five-node network with uniform traffic request. In this example, we assume a bidirectional ring with grooming ratio 2 (which may not be practical in a real SONET ring but is used here for illustration purposes). The total number of bidirectional requests is 10, and each request is one unit of subchannel capacity. Fig. 2(a) illustrates all 10 requests. Fig. 2(b) and (c) illustrate two ways of organizing the connections on two wavelengths (one and a half are actually in use). In Fig. 2(b), there is only bypassing traffic at node 2 of the second ring (dashed line), so nine ADMs are needed to carry all the requests. The number of ADMs can be further reduced by simply reconfiguring the rings as Fig. 2(c), where the positions of the connections between nodes 1 and 4 on the second ring (1↔4) and part of the first ring (4↔0, 0↔1) are swapped.

B. Previous Work and Our Contribution

Table I gives a brief review of previous work reported in this area. The various terms used in the table are described.

- Static or dynamic traffic: The traffic pattern will not or will change with time, respectively.
- Uniform (nonuniform) traffic: Traffic demands between any node pair are the same (not the same).
- Single-hop ring: No virtual connection on the ring will be terminated electronically at any intermediate node.
- Multihop ring: Some or all of the virtual connections on the ring may be regenerated electronically at some intermediate nodes.
- Unidirectional (bidirectional) ring: All the traffic on the ring can go along one (both) direction(s).
- Circle: The circle is a virtual ring that is established on one timeslot of a wavelength has the capacity of one unit.
- PPWDM: This represents a point-to-point (opaque) WDM ring, in which signals get cross-connected and regenerated at every node.
- Distance-dependent traffic (used in [4]): The nodes farthest apart exchange one unit of traffic, and the internode traffic demand increases by one unit as the internode distance decreases by one link.
- Egress node: This is a special node, like the central office of an access network, on the ring at which all the traffic terminates or from which it originates .

Papers on this topic were first seen in 1998 [1], [2], [4]. Several network architectures were proposed and compared in [1], [2]. This work also provided an example to show that the number of wavelengths and ADMs cannot be minimized simultaneously for general traffic. Studies [3]–[6] focused on the single-hop case. In [3], the authors showed that the general traffic-grooming problem is nameplate (NP)-complete. For uniform traffic on a special ring with an egress node, the optimal solution is provided. In [4], [5], full mesh connectivity on a bidirectional ring with an odd number of nodes was considered. An elaborate analysis about how to combine different circles into a wavelength channel was provided and the optimal result was obtained. The work in [6] extended previous discussion to arbitrary traffic, arbitrary number of nodes, and both unidirectional and bidirectional rings. A greedy heuristic was proposed. For a recent review of research problems in traffic grooming, please see [7].

In order to improve on previous work, we first provide formal mathematical definitions for the various traffic-grooming problems, which turn out to be integer linear programs (ILPs). Then, we propose a simulated-annealing-based heuristic approach for solving these optimization problems. We believe that this algorithm has achieved the best results, so far, relative to other approaches. Finally, we discuss the multihop case (with a single hub) where a hub node is used to cross-connect traffic between different wavelengths and timeslots. Comparison of the single-hop and multihop approaches is provided.

II. Problem Definition

A. Formulation 1: Single-Hop Connections

The traffic-grooming problem on a single-hop bidirectional ring can be mathematically formulated as shown in Fig. 3.

The traffic-load constraint simply states that the number of links from node i to node j on all circles is equal to the traffic specified in the traffic matrix. In the channel-capacity constraint, ${}^{d}e^{cw}$ represents a d-directional link on wavelength w and subchannel c. If the virtual connection ${}^{d}V_{ij}^{cw}$ uses it, we say that ${}^{d}e^{cw} \in {}^{d}V_{ij}^{cw}$. The channel-capacity constraint requires that a circle carry only one connection on any given link. The last two constraints specify that the number of connections that start and terminate at a node of a ring is bounded by the capacity of the electronic ADM at that node. If there is an ADM, at most C connections can start and terminate there; otherwise, no add–drop can occur.

The unidirectional-ring case can be viewed as a special case of the formulation shown in Fig. 3, where d can only be either clockwise or counterclockwise. Sometimes, shortest-path routing is required in the bidirectional ring. This requirement can also be accommodated in the formulation shown in Fig. 3

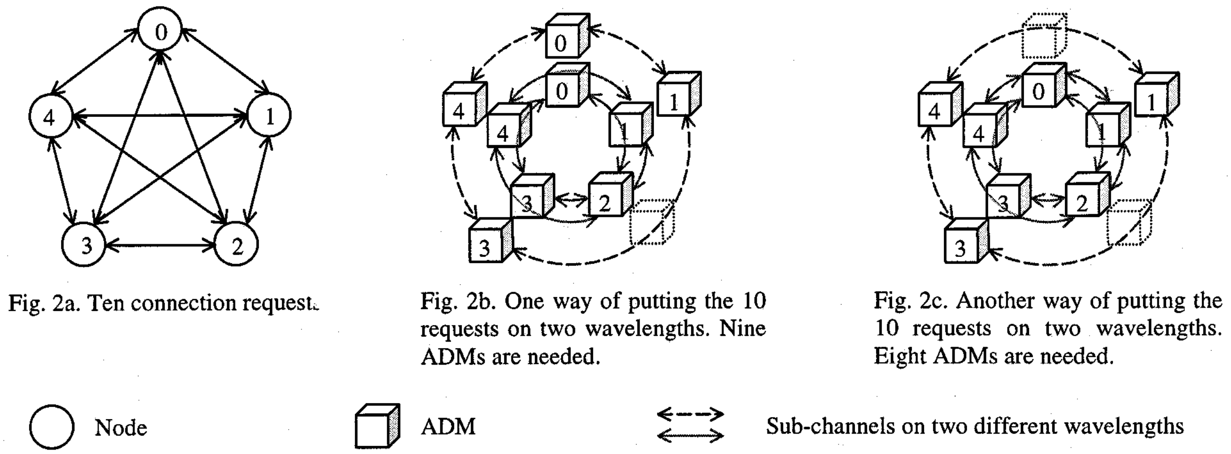

Fig. 2. An example of traffic grooming on a two-wavelength network, showing that different strategies can lead to different network cost. The total number of traffic connections is 10, and each request has a capacity of one unit. (a) Ten connection requests. (b) One way of putting the 10 requests on two wavelengths. Nine ADMs are needed. (c) Another way of putting the 10 requests on two wavelengths. Eight ADMs are needed.

TABLE I
COMPARISION OF PREVIOUS WORK ON TRAFFIC GROOMING ON WDM RING NETWORKS

	Traffic assumptions	Ring architecture	Main result
Gerstel [1][2]	• Static, uniform • Non-statistic dynamic and fixed wavelength	• PPWDM (multihop) • Fully-optical ring (single-hop) • Single-hub (multihop) • Double-hub (multihop) • Hierarchical ring (multihop) • Incremental ring (multihop)	• First paper that tries to minimize transceiver cost. • Study of dynamic traffic and fixed lightpath. • Different architectures are compared.
Modiano [3]	• Egress traffic • Static, uniform	• Unidirectional ring with egress node (single-hop) • Bidirectional ring (single-hop)	• Proof of NP-completeness. • Optimal solution for uniform traffic on egress ring. • Lower bound on uniform all-to-all traffic.
Simmons [4][5]	• Static, uniform • Static, distance-dependent traffic	• Bidirectional ring with odd number of nodes (single-hop)	• How to group timeslots. • Maximal savings for some special cases. • Super node model for distance - dependent traffic.
Zhang & Qiao [6]	• Static, uniform traffic • Static, non-uniform traffic	• Unidirectional and bidirectional ring (single-hop)	• Greedy heuristic for grooming arbitrary traffic. • Heuristic for circle construction for non-uniform traffic.
Our work	• Static, uniform traffic • Static, non-uniform traffic	• Unidirectional and bidirectional ring (single-hop) • Single hub (multihop)	• Formal mathematical problem specification (ILP). • Simulated-annealing-based heuristic for traffic grooming • Greedy heuristic for single-hop and multihop grooming.

by specifying that ${}^{d}V_{ij}^{cw}$ will be zero if the distance from node i to j in the specified direction exceeds $\lceil N/2 \rceil$.

B. Formulation 2: Multihop Method

A multihop ring uses a DXC to do subchannel consolidation or segregation at a hub node. In this study, we simplify the cost calculation by using the following equivalent architecture instead. The hub node has as many ADMs as there are wavelengths. All connections that are passing through it can be terminated and switched to any wavelength and timeslot. A connection can go through the hub node at most once. The formulation for the unidirectional, single-hub ring case is shown in Fig. 4. Bidirectional formulation for the single-hub-based multihop network is also available. We omit it here because it is a straightforward extension of the formulation shown in Fig. 4.

In formulation 2, shown in Fig. 4, O_i^{cw} represents the virtual connection from node i to the hub node on circle c, wavelength w. Similarly, the notation I_i^{cw} represents the connection that starts from the hub node and terminates at node i. The condition $(j > i)$ represents the virtual connection that starts from node

Objective function:

$$\text{Minimize} \sum_i \sum_w ADM_i^w$$

Subject to:

$$\sum_w \sum_c \sum_d {}^dV_{ij}^{cw} = t_{ij} \qquad \forall i, j \qquad \text{(Traffic-load constraint)}$$

$$\sum_{{}^de^{cw} \in {}^dV_{ij}^{cw}} {}^dV_{ij}^{cw} \le 1 \qquad \forall d, e, c, w \qquad \text{(Channel-capacity constraints)}$$

$$\sum_c \sum_j {}^dV_{ij}^{cw} \le C \cdot ADM_i^w \qquad \forall d, i, w \qquad \text{(Transmitter constraints)}$$

$$\sum_c \sum_i {}^dV_{ij}^{cw} \le C \cdot ADM_j^w \qquad \forall d, j, w \qquad \text{(Receiver constraints)}$$

Bounds: ${}^dV_{ij}^{cw}$ *and* ADM_i^w *are both binary numbers.*

Fig. 3. Formulation 1. Mathematical problem formulation for traffic grooming in a single-hop network.

Objective function:

$$\text{Minimize} \sum_i \sum_w ADM_i^w$$

Subject to:

$$\sum_w \sum_c \Big(\sum_{j(j>i)} V_{ij}^{cw} + O_i^{cw}\Big) = \sum_j t_{ij} \qquad \cdots\cdots \forall i \qquad \text{(Traffic-load constraint)}$$

$$\sum_w \sum_c \Big(\sum_{i(j>i)} V_{ij}^{cw} + I_j^{cw}\Big) = \sum_i t_{ij} \qquad \cdots\cdots \forall j \qquad \text{(Traffic-load constraint)}$$

$$\sum_{e^{cw} \in V_{ij}^{cw}} V_{ij}^{cw} + \sum_{i<e^{cw}} O_i^{cw} + \sum_{j>e^{cw}} I_j^{cw} \le 1 \qquad \cdots\cdots \forall e, c, w \qquad \text{(Channel-capacity constraint)}$$

$$\sum_c \sum_j V_{ij}^{cw} + \sum_c O_i^{cw} \le C \cdot ADM_i^w \qquad \cdots \forall i, w \qquad \textit{(Transmitter constraint)}$$

$$\sum_c \sum_j V_{ji}^{cw} + \sum_c I_i^{cw} \le C \cdot ADM_i^w \qquad \cdots \forall i, w \qquad \textit{(Receiver constraint)}$$

Bounds: *All* V_{ij}^{cw}, O_i^{cw}, I_i^{cw} *and* ADM_i^{cw} *are binary numbers.*

Fig. 4. Formulation 2. Mathematical problem formulation for traffic grooming in a single-hub mutihop network.

i and ends at node j without going through the hub node. The condition $i < e^{cw}$ means that if s is the start node of link e^{cw} and t is the end node, then i is upstream of t. Similarly, $j > e^{cw}$ means that j is downstream of s. The traffic-load constraint needs to be broken into two parts for the multihop case. The first part of the multihop formulation, shown in Fig. 4, specifies that any virtual connection that starts from node i will either terminate before it reaches the hub node or will terminate at the hub node. Similarly, the second traffic-load constraint states that any virtual connection that terminates at node i is either coming from the hub node or from some node downstream of the hub node. The explanations of the other constraints are the same as in the single-hop case.

III. Solving the ILPs Directly

We attempted to solve the ILPs directly under the current restriction that a SONET ring has a maximal size of 16 nodes. Uniform traffic on a single-hop unidirectional ring was chosen for demonstration to give the audience a feeling of how fast the ILPs can be solved, although our formulation is capable of solving nonuniform traffic cases (examples of which will be given in Section V-A2). A commercially available ILP solver (CPLEX, ILOG, Inc., Mountain View, CA 94043) was used. Table II shows the computational time and the optimal solution. For small networks, i.e., six nodes or less, the ILP solver can find the optimal solution in a reasonably short time of a few seconds to a few hours. When the network size grows beyond six nodes, the solver was found to take more than 6 h to discover the optimal solution for some cases (shown as question marks in Table II). The results shown in Table II were obtained from a Hewlett Packard Visualize B132l machine (Hewlett Packard Co., Palo Alto, CA 94304) running a UNIX operating system. For the cases $N = 4, 5$ and $C = 12$ (shaded cells), no grooming is needed because one wavelength can carry all of the traffic. When the network size is larger than seven or eight nodes, we need to turn to heuristics. Table III provides some additional information. When we give the ILP solver a time limit of half an hour for each problem, it usually fails to find even one feasible solution when the network size is larger than eight nodes. Notice that when the grooming ratio is large enough that one wavelength is enough to carry all of the traffic, traffic grooming does not have practical meaning anymore, because an ADM is needed at each node. Also, in Table III, note that the simulated-annealing algorithm reaches better results than the greedy algorithm, sometimes even reaching the lower bound. Question

TABLE II
COMPUTATIONAL TIME (IN SECONDS) AND OPTIMAL SOLUTION FOR THE SINGLE-HOP CASE

	C=3		C=4		C=12	
	T(s)	N_{ADM}	T(s)	N_{ADM}	T(s)	N_{ADM}
N=4	1	7	2	7	0	4
N=5	471	12	45	10	0	5
N=6	?	?	9783	15	191	9
N=7	?	?	?	?	549	12

TABLE III
RESULTS FROM DIFFERENT APPROACHES TO SOLVE THE TRAFFIC-GROOMING PROBLEM IN AN UNIDIRECTIONAL RING WITH UNIFORM TRAFFIC AND SINGLE-HOP CONNECTIONS

		N=4	N=5	N=6	N=7	N=8	N=9	N=10	N=11	N=12	N=13	N=14	N=15	N=16
C=3	Lower Bound	6	11	15	21	29	36	45	56	66	78	92	105	120
	ILP	7	12	18	28	46	?	?	?	?	?	?	?	?
	Simulated Annealing	7	12	17	21	31	36	48	57	69	78	95	105	124
	Greedy	7	12	17	26	35	44	56	67	81	98	113	131	152
C=4	Lower Bound	6	10	15	21	28	36	45	55	66	78	91	105	120
	ILP	7	10	15	27	31	?	?	?	?	?	?	?	?
	Simulated Annealing	7	10	15	21	28	36	45	55	66	78	91	105	120
	Greedy	7	11	17	23	30	37	48	60	72	84	99	112	130
C=12	Lower Bound	4	5	9	11	15	18	23	28	33	39	46	53	60
	ILP	4	5	9	13	18	?	?	?	?	?	?	?	?
	Simulated Annealing	4	5	9	12	16	18	24	30	36	39	49	57	64
	Greedy	4	5	10	13	18	19	27	33	41	49	56	69	73
C=16	Lower Bound	4	5	6	10	12	16	18	23	28	33	37	42	48
	ILP	4	5	6	11	14	18	?	?	?	?	?	?	?
	Simulated Annealing	4	5	6	11	14	18	20	26	33	37	42	46	57
	Greedy	4	5	6	13	15	20	21	30	35	42	48	58	65
C=48	Lower Bound	4	5	6	7	8	9	10	15	17	19	21	26	29
	ILP	4	5	6	7	8	9	10	?	?	?	?	?	?
	Simulated Annealing	4	5	6	7	8	9	10	16	19	22	24	31	37
	Greedy	4	5	6	7	8	9	10	19	23	24	25	31	34
C=64	Lower Bound	4	5	6	7	8	9	10	11	15	18	20	22	24
	ILP	4	5	6	7	8	9	10	11	?	?	?	?	?
	Simulated Annealing	4	5	6	7	8	9	10	11	15	19	22	25	28
	Greedy	4	5	6	7	8	9	10	11	15	25	26	27	28

marks represent the cases where the ILP solver could not find a feasible solution within a half-hour time limit. Shaded data means that all traffic can be carried on one wavelength, so traffic grooming is necessary. Lightly shaded data indicates the only situation where the greedy algorithm was found to outperform the simulated-annealing algorithm in our experiments.

IV. HEURISTICS

The general traffic-grooming problem is NP-complete [3]. Solving the ILP directly is not practical even for moderate-size networks because of the long solution time. Although a simple greedy heuristic has been provided in [6], we propose alternative approaches to discover improved results. Specifically, we propose a simulated-annealing algorithm for single-hop connections and a simple heuristic for multihop connections.

A. Simulated Annealing for Single-Hop Connections

To reduce the complexity, the traffic-grooming problem is usually divided into two components (see [4]–[6]). In the first step, the traffic demands are assigned to circles. In the second step, a traffic-grooming algorithm is employed to reorganize the circles or connections on wavelengths. We adopted the same strategy, and our first-step heuristic is built upon the wavelength-assignment algorithm proposed in [6] for reducing the wavelength consumption. An important practical concern is whether the heuristic is a good foundation for our second-step heuristic. It has been shown in [1] that the number of wavelengths and the number of ADMs cannot always be minimized simultaneously. Both of our ILPs (discussed in Section II) and our heuristic (to be discussed in this section) have the potential to find the true optimal number of ADMs when there are more than enough wavelengths. Although there is a chance that the heuristic gives

```
do { // iterate around all states
        do {//accept ANN_CONST*C times
                //randomly pick two circles, swap part or all of them
                dcost=perturb(); //
                if (Δcost<0 || (Δcost>0 && exp(-Δcost/control)>rand[0,1))){
                        accept_change();// accepted the change
                        chain++;
                }
                else reject_change();
        }while (chain<ANN_CONST*C);
        control = control*DEC_CONST;
} while (control>END);
```

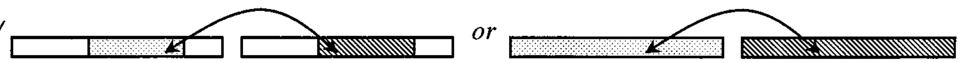

Fig. 5. Simulated annealing-based traffic-grooming algorithm for single-hope method.

us fewer wavelengths than necessary to minimize the ADM usage, the chance of this occurring is low. Furthermore, the emphasis of this study is on comparing different second-step heuristics, so the heuristic is adequate. Our second-step heuristic randomly chooses some virtual connections in the network and changes their position, by using the simulated-annealing technique to help accelerate the process of branch-and-bound to find a good solution. The implementation of simulated annealing follows the Monte Carlo method [8] referred to as the *Metropolis algorithm*. The algorithm is depicted in Fig. 5.

In this implementation, "perturb" means to randomly pick two circles on different wavelengths, swapping part of them or the whole circles (illustrated in the comment of the pseudocode). The chance for doing partial swapping is selected to be very small. Whenever the criterion for doing partial swapping is satisfied, we go ahead to check if there are segments in the two circles that are swappable (if swapped, they do not break any existing connections). If the result is negative, we simply swap the whole circles. The consequence of the perturbation may or may not be helpful to bring down the ADM usage. If the perturbation helps, we will accept the perturbation; otherwise, we will calculate $\exp(-\Delta$ *cost–control*) and compare it with a random number. The perturbation will still be accepted if the random number is smaller. After repeating the above process for a certain number of times, we consider that the system has reached the equilibrium state, and then we go on to decrease the *control* variable (the temperature). The process will be terminated when the *control* variable satisfies some predefined criterion.

In the algorithm shown in Fig. 5, as the computation goes on and the *control* variable goes down, the chance of having a good perturbation and the chance of accepting a bad perturbation will decrease, so the time spent in lower temperature is much longer than in higher temperature. The constants ANN_CONST and DEC_CONST are critical for the algorithm's performance. ANN_CONST decides how long we have to wait before we consider that the system has reached its equilibrium. DEC_CONST decides how quickly we decrease the temperature. After experimenting with these parameter values, we finally adopted ANN_CONST to be between 4 and 20, depending on the size of the search space, and DEC_CONST $= 0.95$ for best results for our numerical examples. The start temperature is found to be important for some cases. We do not recommend a very high starting temperature. The circle-assignment heuristic in the first step serves as a pregrooming heuristic, which puts similar connections as close to one another as possible; thus, a very high starting temperature can counterbalance the effort of the first step and prolong the convergence process. Some special techniques are also needed when implementing the *perturb()* function, details of which are skipped here. These techniques are mainly used to increase the convergence speed of the algorithm.

Although the simulated-annealing-based heuristic adopted the two-step strategy, it is intrinsically different from the greedy heuristic [6] in that it always has the potential to find the true optimal solution. For the greedy heuristic, the second step can only regroup the circles but not change the existing circles. This is one of the major reasons for getting a suboptimal solution in [6]. The simulated-annealing-based second-step heuristic provides a chance to change the circles so that it can jump out of traps.

B. Heuristic for the Multihop Method

A greedy heuristic is proposed for the multihop method [9]. It puts all the traffic on circles sequentially, and then it applies the algorithm shown in Fig. 6.

In the above pseudocode, the Wavelength_Combining function checks two wavelengths link by link. If the total load on any link will not exceed the wavelength capacity, then combine them. The Segment_Swapping function finds the underutilized links in different wavelengths and combines them into one wavelength through segment swapping.

V. Illustrative Numerical Results and Comparisons

A. Single-Hop Approach

1) Uniform Traffic: The uniform-traffic case has been well studied in the literature, so it provides a good reference to eval-

```
while (the number of ADMs and the number of wavelengths can be reduced) {
        Establish connection on the shortest path;
        Wavelength_Combining();
        Segmen_Swapping();
}
```

Fig. 6. Heuristic traffic-grooming algorithm for multihop method.

uate our algorithms. Table III shows the required number of ADMs from the different algorithms for the single-hop case defined in Formulation 1 (see Fig. 3) for an unidirectional ring. For each value of the grooming ratio C, the first row shows the lower bound calculated by the algorithm proposed in [6]. The second row shows the result from solving our ILP formulation, given half an hour for each value; a question mark means that the ILP solver failed to obtain a feasible solution within half an hour. The third row shows the result from our simulated-annealing heuristic, and the fourth row shows the result from the greedy heuristic in [6]. The simulated-annealing algorithm was run for 30 trials and the best result was chosen.

We noticed that, most of the time, the simulated-annealing approach achieves better results than the greedy heuristic. Sometimes, it even reaches the lower bound. Even for the one case where the greedy heuristic showed a better result ($N = 16$, $C = 48$), we believe that the simulated-annealing algorithm might have done a better job if it was given more time. In this study, all algorithms started from the same network configurations (step-1 heuristic results), but the simulated-annealing algorithm can always be started from the greedy result if we choose to do so, and this will guarantee equal or better results. To calculate all of the simulated-annealing results in Table III, the computational time for all 30 trials was a little less than 100 min on a 200-MHz Pentium machine running Windows NT (Microsoft Corporation, Redmond, WA 98052-6399). The fact that the simulated-annealing algorithm outperforms the simple greedy algorithm is not surprising because it tried more possibilities and consumed more computational time. In [3] and [4], the authors provided a more elaborate algorithm for the special case of uniform traffic on bidirectional rings. Although this algorithm has been claimed to be optimal, it does need one more ADM than our simulated-annealing-based heuristic for the case $N = 9$ and $C = 4$. The reason for this difference is that the algorithm in [3] and [4] will always try to fill partially occupied wavelengths first before it examines a new wavelength. For this special case ($N = 9$, $C = 4$), three wavelengths are needed to carry all the traffic and the optimal happens to occur when two wavelengths are filled with only three circles and one with four circles, which is out of the search space for the algorithm in [3], [4]. For our algorithm, this is a natural result.

2) Nonuniform Traffic: For handling nonuniform traffic, we first show an example that is small enough to be solved by an ILP solver in a reasonable amount of time. We picked a four-node unidirectional-ring network with a random traffic matrix $\{\{0, 1, 8, 4\}, \{12, 0, 3, 9\}, \{1, 2, 0, 2\}, \{4, 1, 7, 0\}\}$. In this traffic matrix, if nodes are numbered 0 through 3, then the traffic demand from node 1 to node 3 is nine units. The grooming ratio C (or wavelength capacity) of this network is chosen to be 3. We run all the three programs, i.e., ILP, simulated annealing, and greedy. After a 6-h time limit, the ILP algorithm does not reach its final conclusion and its best result was 15 wavelengths and 31 ADMs. The simulated-annealing algorithm produces the same result in about 2 s, whereas the greedy algorithm yields 33 ADMs in negligible time. Our experiment indicates that the ILP usually cannot handle networks larger than six nodes when given nonuniform traffic. Results from the simulated-annealing algorithm are shown in Fig. 7. In this figure, each wavelength is drawn as three lines (recall that $C = 3$) representing the three circles. The four nodes on the ring are denoted as nodes 0–3. As an example, the eight units of traffic from node 0 to node 2 are carried by wavelength 4, wavelength 9, and part of wavelength 7 (timeslots 2 and 3), as shown in bold in Fig. 7.

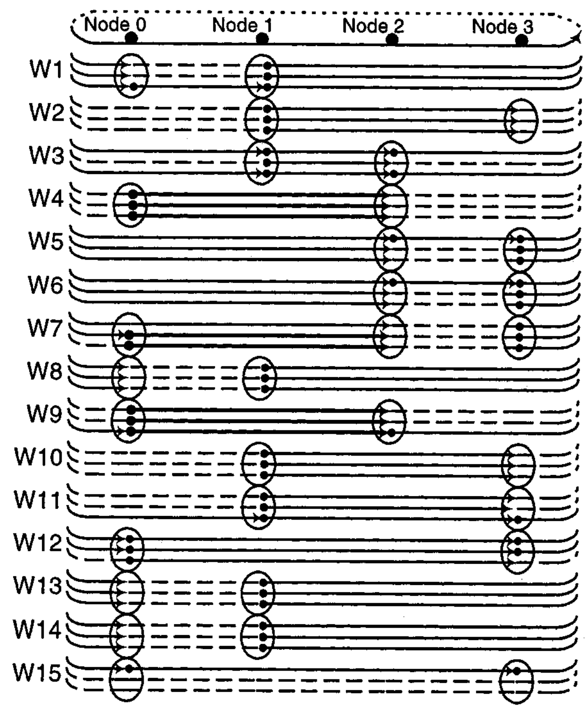

Fig. 7. Nonuniform traffic example. Dots and arrows represent the start and the end of a connection. Circles represent ADMs.

Notice that the traffic matrix is highly asymmetric, which is very practical in today's data networks and is very different from most previous work on traffic grooming. In the above example, suppose the wavelength capacity is OC-3, then traffic demand from node 0 to node 2 equals OC-8, which can be two OC-3 and two OC-1 connections (as was shown in Fig. 7). One demand may have to be transmitted on different wavelengths if needed,

TABLE IV
COMPARISON OF SINGLE-HOP AND MULTIHOP APPROACHES ON NONUNIFORM TRAFFIC ($N = 4$)

	Single-hop		Multihop	
	W	ADM	W	ADM
C=3	15	31	19	38
C=12	4	14	5	11
C=48	1	4	2	5

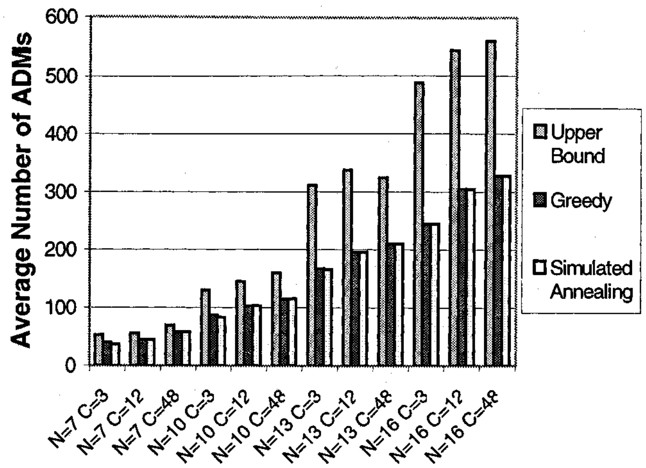

Fig. 8. Statistical result comparison for nonuniform traffic on a bidirectional ring.

e.g., the OC-8 traffic from a single user may be transmitted over three OC-3 wavelengths.

Fig. 8 shows statistical result comparisons for the bidirectional rings. The results are calculated in the following way. For each ring scenario (described by its size and wavelength capacity), 30 random traffic matrices were generated. The traffic demand between any node pair is uniformly distributed between 0 and twice the wavelength capacity ($t_{ij} = \text{rand}(0, 2 \times C)$). Three values are calculated for each traffic matrix, namely, the upper bound (no grooming), the result from the greedy approach, and the result from the simulated-annealing approach. Then, the average number of ADMs is shown in the figure for each approach.

The results show that both traffic-grooming algorithms achieve very good savings. Although still visible in some cases, the benefit of the simulated-annealing-based approach is not that significant for this example. The reason is that the traffic load given to the networks is heavy and unregulated. Unregulated traffic causes bubbles (unusable segments) in single-hop networks. The two reasons combined lead to the number of circles to tens and even hundreds, which in turn increases the search space dramatically. Even after increasing the computation time, the simulated-annealing result does not improve significantly because the search space is very large.

B. Multihop Approach

We compared the result obtained by multihop approach (Section IV–B) with the one from the single-hop approach under arbitrary traffic. Table IV shows an example using the same traffic matrix as the one used to obtain Fig. 7. Our results show that, in order to get the most benefit from using the switching (hub) node, the grooming ratio cannot be too small or too large. When the grooming ratio is 3 or 4, the single-hop approach usuually results in fewer ADMs. The reason is that the DXC cost is estimated by its port count, so more wavelengths mean higher DXC cost. When the grooming ratio is large enough that one or two wavelengths can accommodate all the traffic, the two approaches do not have much difference. When the grooming ratio has a moderate value, the cross-connect function at the hub node usually helps to reduce the number of ADMs. Notice that the ADM savings achieved by multihopping is usually at the cost of increased wavelength usage in our algorithm. Intuitively, the explanation is that, for single-hop networks, all of the connections are established on the shortest path; for the multihop case, some connections may travel more than one hop and traverse the ring more than once in order to take advantage of the switching function of the hub node. Although switching may help to fill the bubbles around the switching node, with just one switching node, it is not easy to counterbalance the waste caused by the longer connections. Currently, our multihop heuristic is only designed for minimizing the number of ADMs. An open problem is to incorporate minimizing the number of wavelengths, as well.

VI. CONCLUSION

In this study, we first provided the formal mathematical specifications of the traffic-grooming problem in several ring networks, i.e., single-hop and multihop (with a single hub) cases of unidirectional and bidirectional rings. Then, we proposed a simulated-annealing-based traffic-grooming algorithm for the single-hop case and a greedy heuristic for the multihop case. The simulated-annealing-based heuristic overcomes the suboptimal problem caused by the two-step strategy in previous greedy heuristic. We believe that the simulated-annealing algorithm has reached the best result so far. It was also shown that, for nonuniform traffic, the greedy approach usually is good enough when compared with the simulated-annealing algorithm. The multihop approach could achieve more ADM savings when the grooming ratio is neither too small nor too large, but it usually results in more wavelength usage due to the prolonged connection length.

REFERENCES

[1] O. Gerstel, P. Lin, and G. Sasaki, "Wavelength assignment in a WDM ring to minimize cost of embedded SONET rings," in *Proc. IEEE INFOCOM '98*, Mar. 1998, pp. 94–101.
[2] O. Gerstel and G. Sasaki, "Cost effective traffic grooming in WDM rings," in *Proc. IEEE INFOCOM '98*, Mar. 1998, pp. 69–77.

[3] E. H. Modiano and A. L. Chiu, “Traffic grooming algorithms for minimizing electronic multiplexing costs in unidirectional SONET/WDM ring networks,” in *Proc. CISS ’98*, Mar. 1998.
[4] J. M. Simmons, E. L. Goldstein, and A. A. M. Saleh, “On the value of wavelength-add/drop in WDM rings with uniform traffic,” in *Proc. OFC ’98*, San Jose, CA, Feb. 22–27, 1998, pp. 361–362.
[5] ——, “Quantifying the benefit of wavelength add–drop in WDM rings with distance-independent and dependent traffic,” *J. Lightwave Technol.*, vol. 17, pp. 48–57, Jan. 1999.
[6] X. Zhang and C. Qiao, “An effective and comprehensive solution to traffic grooming and wavelength assignment in SONET/WDM rings,” in *Proc. SPIE*, vol. 3531, 1998, pp. 221–232.
[7] E. Modiano and P. J. Lin, “Traffic grooming in WDM networks,” *IEEE Commun. Mag.*, vol. 39, pp. 124–129, July 2001.
[8] P. J. M van Laarhoven, *Theoretical and Computational Aspects of Simulated Annealing*. Amsterdam, The Netherlands: Center Mathematics Comput. Sci., 1988.
[9] W. Cho, J. Wang, and B. Mukherjee, “Improved approaches for cost-effective traffic gromming in WDM ring networks: Uniform-traffic case,” *Photon. Network Commun.*, vol. 3, no. 2, pp. 245–254, Apr. 2001.

Jian Wang received the B.S. degree in physics from Shandong University, Shandong, P.R. China, and the M.S. degrees in biomedical engineering and applied science from Shandong Medical University, Shandong, P.R. China, and the University of California, Davis, respectively. He is currently pursuing the Ph.D. degree in optical networking at the University of California, Davis, and conducting research in the Networks Research Laboratory at that university.

Wonhong Cho received the B.S., M.S., and Ph.D. degrees in electronic engineering from Yonsei University, Seoul, Korea, in 1985, 1987, and 1994, respectively.

From March 1994 to May 2000, he was an Assistant Professor of Computer Science, Kunsan National University, Kunsan, Jeollabuk-do, Korea. In January 1999, he joined the University of California, Davis, as Visiting Professor of Computer Science. His research interests include WDM networks.

V. Rao Vemuri holds a joint appointment as a Professor in the Department of Applied Science and the Department of Computer Science, University of California, Davis. He is also a scientist in the Computational Directorate of the Lawrence Livermore National Laboratory, Livermore, CA.

Biswanath Mukherjee (S’82–M’83) received the B.Tech. (Hons.) degree from the Indian Institute of Technology, Kharagpur, India, in 1980 and the Ph.D. degree from University of Washington, Seattle, in June 1987.

At the University of Washington, he held a GTE Teaching Fellowship and a General Electric Foundation Fellowship. In July 1987, he joined the Unversity of California, Davis, where he has been a Professor of Computer Science since July 1995, and Chairman of Computer Science since September 1997. He serves on the editorial boards of the IEEE/ACM TRANSACTIONS ON NETWORKING, IEEE NETWORK, *ACM/Baltzer Wireless Information Networks*, and *Optical Network Magazine*. He also served as Editor-at-Large for Optical Networking and Communications for the IEEE INFOCOM ’96 conference. He is author of the textbook *Optical Communication Networks*, which received the 1997 Association of American Publishers’ Honorable Mention in Computer Science. His research interests include lightwave networks, network security, and wireless networks.

Dr. Mukherjee is a co-winner of awards for papers presented at the 1991 and the 1994 National Computer Security Conferences.