Traffic Grooming in an Optical WDM Mesh Network

Keyao Zhu, *Student Member, IEEE*, and Biswanath Mukherjee, *Member, IEEE*

***Abstract*—In wavelength-division multiplexing (WDM) optical networks, the bandwidth request of a traffic stream can be much lower than the capacity of a lightpath. Efficiently grooming low-speed connections onto high-capacity lightpaths will improve the network throughput and reduce the network cost. In WDM/SONET ring networks, it has been shown in the optical network literature that by carefully grooming the low-speed connection and using wavelength-division multiplexer (OADM) to perform the optical bypass at intermediate nodes, electronic ADMs can be saved and network cost will be reduced. In this study, we investigate the traffic-grooming problem in a WDM-based optical mesh topology network. Our objective is to improve the network throughput. We study the node architecture for a WDM mesh network with traffic-grooming capability. A mathematical formulation of the traffic-grooming problem is presented in this study and several fast heuristics are also proposed and evaluated.**

***Index Terms*—Integer linear program, lightpath, mesh network, optical network, traffic grooming, wavelength-division multiplexing.**

I. INTRODUCTION

FIBER OPTICS and wavelength-division multiplexing (WDM) are promising technologies that are expected to satisfy the drastically increasing bandwidth requirements of the Internet. WDM is an approach that can exploit the huge optoelectronic bandwidth mismatch by requiring that each end-user's equipment operate only at electronic rate, but multiple WDM channels from different end-users may be multiplexed on the same fiber [1]. In a wavelength-routed WDM network, a "lightpath" may be established from a source node to a destination node and it may span multiple fiber links [2]. In an all-optical network, the lightpath may remain entirely in the optical domain, optically bypassing the intermediate nodes. Using wavelength-routing switches (WRSs) [1] at intermediate nodes, and via appropriate routing and wavelength assignment (RWA), a lightpath can create logical (or virtual) neighbors out of nodes that are geographically far apart in the network.

Assigning network resources (e.g., wavelengths, transceivers) to successfully carry the connection requests (lightpaths) is well known as the routing and wavelength assignment problem [1], [3]. It is also known as a lightpath-provisioning problem [4]. A number of RWA studies have been reported in the optical networking literature, either based on static traffic demands [1]–[7] or based on dynamic traffic demands [8]–[10]. Most previous studies have assumed that a connection requests bandwidth for an entire lightpath channel. In this study, we assume the bandwidth of connection requests can be some fraction of the lightpath capacity, which makes the problem more practical.

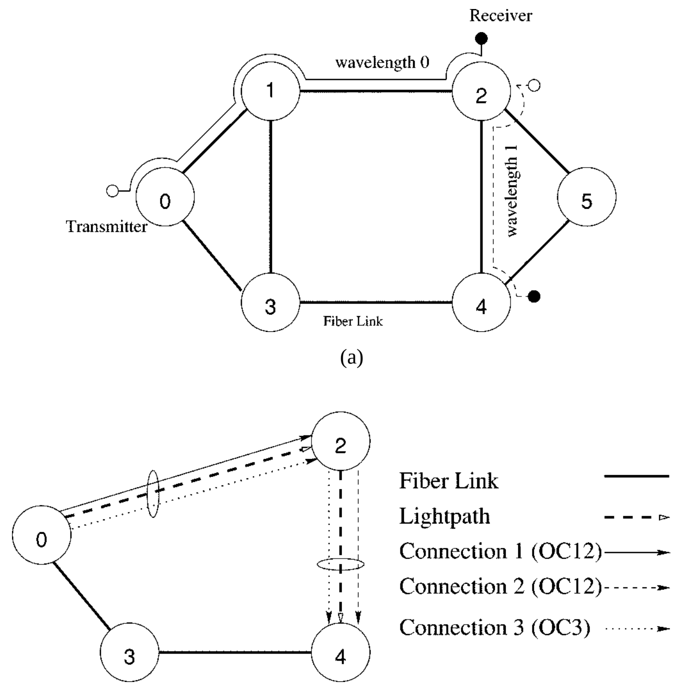

Fig. 1. Illustrative example of traffic grooming.

We investigate the problem of how to "groom" low-speed connection requests to high-capacity lightpaths efficiently. The traffic-grooming problem has been studied on the SONET ring topology. References [11]–[16] reported some previous work on the traffic-grooming problem on the WDM SONET ring networks. The objective function in these studies is to minimize the total network cost, measured in terms of the number of SONET add–drop multiplexers (ADMs). In this paper, we use irregular mesh WDM networks as our study topologies and assume that a connection requests a bandwidth that is a fraction of the wavelength capacity.

Fig. 1 shows an illustrative example of traffic grooming in a WDM mesh network. Fig. 1(a) shows a small six-node network. Each fiber has two wavelength channels. The capacity of each wavelength channel is OC-48, i.e., approx. 2.5 Gb/s.[1] Each node is equipped with a tunable transmitter and a tunable receiver, both of which can be tuned to any wavelength. There are three

[1]Note that the bandwidth of an OC-n channel is approximately $n * 51.84$ Mb/s.

Manuscript received February 15, 2001; revised July 31, 2001. This work was supported in part by the National Science Foundation (NSF) under Grants NCR-9508239 and ANI-9805285. This paper was presented in part at the IEEE International Conference on Communications (ICC'01) at Helsinki, Finland, in June 2001.

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

Publisher Item Identifier S 0733-8716(02)00156-7.

connection requests: (0, 2) with bandwidth requirement OC-12; (2, 4) with bandwidth requirement OC-12; and (0, 4) with bandwidth requirement OC-3. Two lightpaths have already been set up to carry these three connections, as shown in Fig. 1(a). Because of the resource limitations (transmitter in node 0 and receiver in node 4 are busy), we cannot set up a lightpath directly from node 0 to node 4; thus, connection 3 has to be carried by the spare capacity of the two existing lightpaths, as shown in Fig. 1(b). Different connection requests between the same node pair (s, d) can be either groomed on the same lightpath, which directly joins (s, d), using various multiplexing techniques, or routed separately through different virtual paths. A connection may traverse multiple lightpaths if no resources are available to set up a lightpath between the source and destination directly.

We investigate the node architecture for the WDM optical network to support traffic-grooming capability. We study an optical wide-area WDM network which utilizes a grooming-capable optical node architecture, so that a group of lightpaths can be set up to optimally carry the low-speed connection requests.

We formulate the traffic-grooming problem in a mesh network as an optimization problem with the following objective function: for a given traffic matrix set and network resources, maximize the (weighted) network throughput. The mathematical formulation is presented for static traffic demands. Several simple provisioning algorithms, i.e., heuristics, are also proposed and their performance is compared. Finally, we show how to extend the mathematical formulation to accommodate other network optimization criteria.

II. General Problem Statement

The problem of grooming low-speed traffic requests to high-bandwidth wavelength channels on a given physical topology (fiber network) is formally stated below. We are given the following inputs to the problem.

1) A physical topology $G_p = (V, E_p)$ consisting of a weighted unidirectional graph, where V is the set of network nodes and E_p is the set of physical links, connecting the nodes. Nodes correspond to network nodes and links correspond to the fibers between nodes. Though links are unidirectional, we assume that there are an equal number of fibers joining two nodes in different directions. Links are assigned weights, which may correspond to the physical distance between nodes. In this study, we assume that all links have the same weight 1, which corresponds to the fiber hop distance. A network node i is assumed to be equipped with a $D_p(i) \times D_p(i)$ wavelength-routing switch (WRS), where $D_p(i)$ denotes the number of incoming fiber links to node i. [2]
2) Number of wavelength channels carried by each fiber $= W$. Capacity of a wavelength $= C$.
3) A set of $N \times N$ traffic matrices, where $N = |V|$. Each traffic matrix in the traffic-matrix set represents one particular group of low-speed connection requests between the nodes of the network. For example, if C is OC-48, there may exist four traffic matrices: an OC-1 traffic matrix, an OC-3 traffic matrix, an OC-12 traffic matrix, and an OC-48 traffic matrix.
4) The number of lasers (transmitters) (TR_i) and filters (receivers) (RR_i) at each node i. Note that the transceiver can be either wavelength-tunable or part of a fixed-tuned array.

Our goals are to determine the following.

1) A virtual topology $G_v = (V, E_v)$. The nodes of the virtual topology correspond to the nodes in the physical topology. A link between nodes i and j corresponds to a unidirectional lightpath set up between node pair (i, j).
2) Routing connection requests on the virtual topology to either minimize the total network cost or maximize total throughput. In this study, we consider maximizing total throughput.

[2]For any node i, the number of incoming fiber links is equal to the number of outgoing fiber links.

III. Node Architecture

To carry connection requests in a WDM network, lightpath connections may be established between pairs of nodes. A connection request may traverse through one or more lightpaths before it reaches the destination. Two important functionalities must be supported by the WDM network nodes: one is wavelength routing and the other is multiplexing and demultiplexing. A WRS in [1] and [3] provides the wavelength-routing capability to the WDM network nodes. Optical multiplexer/demultiplexer can multiplex/demultiplex several wavelengths to the same fiber link. Low-speed connection requests will be multiplexed on the same lightpath channel by using an electronic-domain TDM-based multiplexing technique. Figs. 2 and 3 show two sample node architectures in a WDM optical network.

The node architecture is composed of two components: WRS and access station. The WRS performs wavelength routing and wavelength multiplexing/demultiplexing. The access station performs local traffic adding/dropping and low-speed traffic-grooming functionalities. WRS is composed of an optical crossconnect (OXC), network control and management unit (NC&M), and optical multiplexer/demultiplexer. In the NC&M, the network-to-network interface (NNI) will configure the OXC and exchange control messages with peer nodes on a dedicated wavelength channel (shown as wavelength 0 in Figs. 2 and 3). The network-to-user interface (NUI) will communicate with the NNI and exchange control information with the user-to-network interface (UNI), the control component of the access station. The OXC provides wavelength-switching functionality. As shown in Fig. 2, each fiber has three wavelengths. Wavelength 0 is used as a control channel for the NC&M to exchange control messages between network nodes. Other wavelengths are used to transmit data traffic.

In Fig. 2, each access station is equipped with some transmitters and receivers (transceivers). Traffic originated from the access station is sent out as an optical signal on one wavelength channel by a transmitter. Traffic destined to the access station is converted from an optical signal to electronic data by a receiver. Both tunable transceivers and fixed transceivers could be used in a WDM network. A tunable transceiver can be tuned between different wavelengths so that it can send out (or receive)

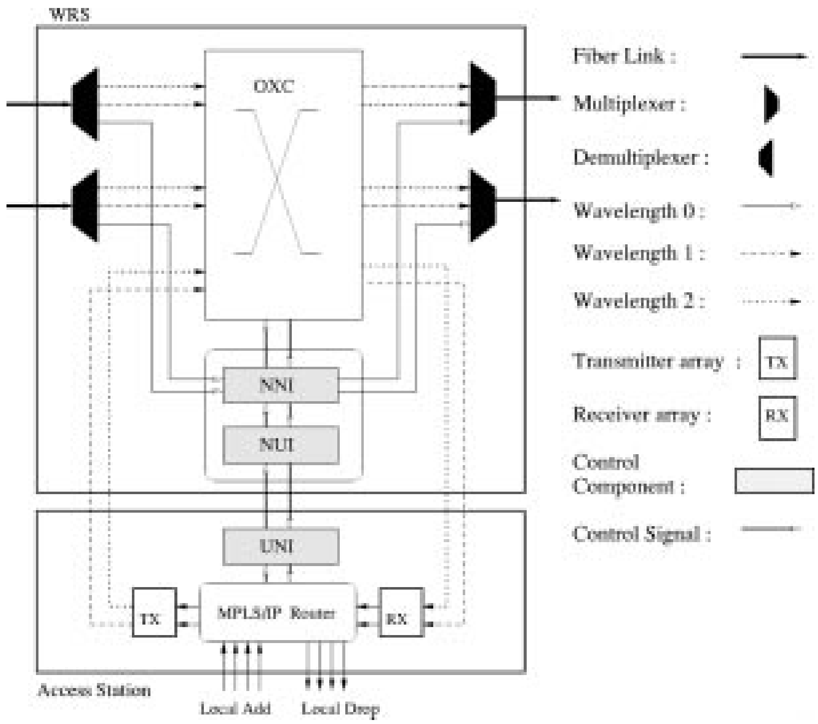

Fig. 2. Node architecture 1: IP over WDM.

an optical signal on any free wavelength in its tuning range. A fixed transceiver can only emit (or receive) an optical signals on one wavelength. To explore all of the wavelength channels on a fiber, a set of fixed transceivers, one per wavelength, is grouped together to form a transceiver array. The size of a fixed transceiver array can be equal to or smaller than the number of wavelengths on a fiber, and the number of tranceiver arrays can be equal to or smaller than the number of fibers joining a node.

The access station in Fig. 2 provides a flexible, software-based bandwidth-provisioning capability to the network. Multiplexing low-speed connections to high-capacity lightpaths is done by the MPLS/IP router using a software-based queuing scheme. The advantages of this model are that: 1) it provides flexible bandwidth granularity for the traffic requests and 2) this MPLS/IP-over-WDM model has much less overhead than the SONET-over-WDM model, shown in Fig. 3. A potential disadvantage of this model is that the processing speed of the MPLS/IP router may not be fast enough compared with the vast amount of the bandwidth provided by the optical fiber link.

In Fig. 3, each access station is equipped with several SONET add–drop multiplexers (ADMs). Each SONET ADM has the ability to separate a high-rate SONET signal into lower rate components [13]. In order for a node to transmit or receive traffic on a wavelength, the wavelength must be added or dropped at the node and a SONET ADM must be used. Generally, each SONET ADM is equipped with a fixed transceiver and operates only on one wavelength as shown in Fig. 3. The digital crossconnect (DCS) can interconnect the low-speed traffic streams between the access station and the ADMs. A low-speed traffic stream on one wavelength can be either dropped to the local client (IP router, ATM switch, etc.) or switched to another ADM and sent out on another wavelength. Fig. 3 presents a SONET-over-WDM node architecture. SONET components (ADM, DCS, etc.) and SONET framing schemes can provide TDM-based fast multiplexing/demulitplexing capability, compared with the software-based scheme in Fig. 3. The disadvantage of this approach is the high cost of SONET components, such as ADM and DCS. In reality, both kinds of access stations may be used together to be connected with an OXC in order to have a multiservice platform for accessing an OXC in a carrier's network.

IV. Mathematical (ILP) Formulation of the Traffic-Grooming Problem

The traffic-grooming problem in a mesh network with static traffic pattern turns out to be an integer linear program (ILP). We make the following assumptions in our study.

1) The network is a single-fiber irregular mesh network, i.e., there is at most one fiber link between each node pair.
2) The wavelength-routing switches in the network nodes do not have wavelength conversion capability. A lightpath connection must be set up on the same wavelength channel if it traverses through several fibers. An extension of this problem to include wavelength conversion is straightforward and it actually makes the problem simpler in terms of the number of variables and equations.
3) The transceivers in a network node are tunable to any wavelength on the fiber.
4) A connection request cannot be divided into several lower speed connections and routed separately from the source to the destination. The data traffic on a connection request should always follow the same route.
5) Each node has unlimited multiplexing/demultiplexing capability and time-slot interchange capability. This means

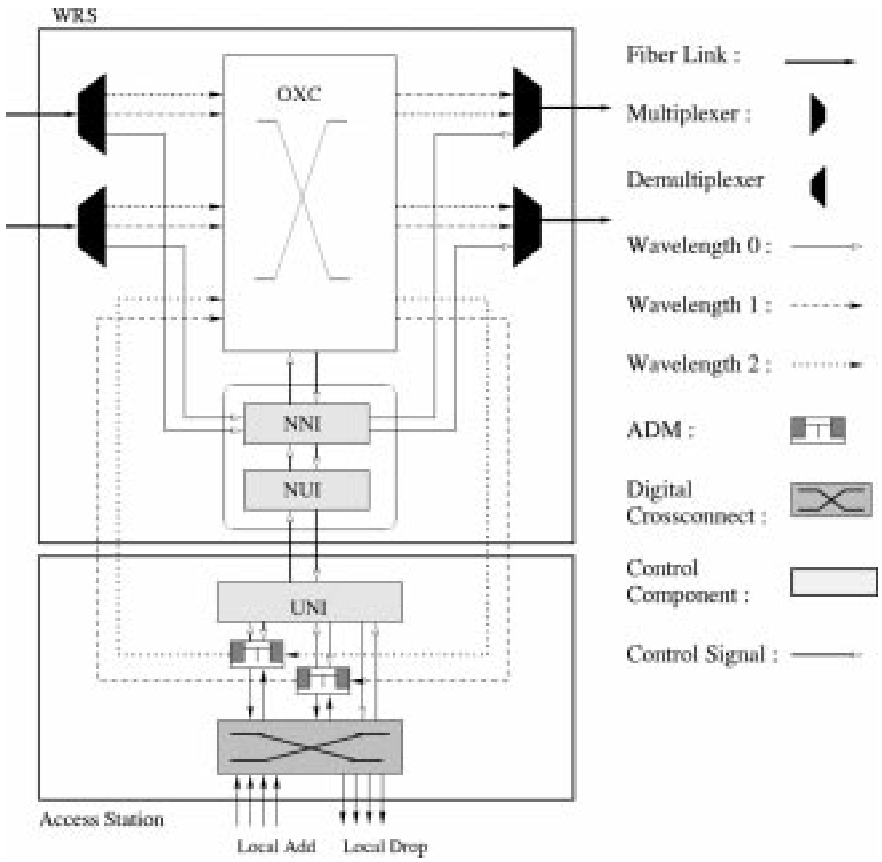

Fig. 3. Node architecture 2: SONET over WDM.

that the access station of a network node can multiplex/demultiplex as many low-speed traffic streams to a lightpath as needed, as long as the aggregated traffic does not exceed the lightpath capacity.[3]

A. Multihop Traffic Grooming

In this section, we assume that a connection can traverse multiple lightpaths before it reaches the destination. So, a connection may be groomed with different connections on different lightpaths. By extending the work in [1] and [6], we formulate the problem as an optimization problem. We will use the following notation in our mathematical formulation.

m and n endpoints of a physical fiber link that might occur in a lightpath.

i and j originating and terminating nodes for a lightpath. A lightpath may traverse single or multiple physical fiber links.

s and d source and destination of the end-to-end traffic request. The end-to-end traffic may traverse through a single or multiple lightpaths. Fig. 4 shows how an end-to-end connection request may be carried.

y granularity of low-speed traffic requests. We assume $y \in \{1, 3, 12, 48\}$, which means that traffic demands between node pairs can be any of OC-1, OC-3, OC-12 and OC-48.

[3]This may only be true for the software-based provisioning scheme in Fig. 2, which may support virtual-circuit connections. The grooming capability of the node architecture in Fig. 3 is limited by the number of output ports of SONET ADM and the size and the functionality of DCS.

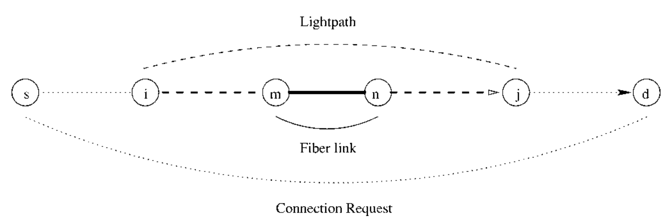

Fig. 4. Illustrative example of a fiber link, a lightpath, and a connection request.

t index of OC-y traffic request for any given node pair (s, d). For example, if there are ten OC-1 requests between node pair (s, d), then $t \in [1, 10]$.

- Given:

N number of nodes in the network.

W number of wavelengths per fiber. We assume all of the fibers in the network carry the same number of wavelengths.

P_{mn} number of fibers interconnecting node m and node n. $P_{mn} = 0$ for node pair which is not physically adjacent to each other. $P_{mn} = P_{nm} = 1$ if and only if there exists a direct physical fiber link between nodes m and n.

P_{mn}^w wavelength w on fiber P_{mn}. $P_{mn}^w = P_{mn}$.

TR_i number of transmitters at node i.

RR_i number of receivers at node i. Note that, in this set of ILP formulation, we assume all the nodes are equipped with tunable transceivers, which can be tuned to any of W wavelengths.

C capacity of each channel (wavelength).

Λ traffic matrix set. $\Lambda = \{\Lambda_y\}$, where y can be any allowed low-speed streams, 1, 3, 12, etc. In our study, $y \in \{1, 3, 12, 48\}$. $\Lambda_{y,sd}$ is the number of OC-y connection requests between node pair (s, d).

- Variables:
  - — Virtual topology:
    - V_{ij} number of lightpaths from node i to node j in virtual topology. $V_{ij} = 0$ does not imply that $V_{ji} = 0$.
    - V_{ij}^w number of lightpaths from node i to node j on wavelength w. Note that, if $V_{ij}^w > 1$, the lightpaths between node i and j on wavelength w may take different paths.
  - — Physical topology route:
    - $P_{mn}^{ij,w}$ number of lightpaths between nodes (i, j) routed through fiber link (m, n) on wavelength w.
  - — Traffic route:
    - $\lambda_{ij,y}^{sd,t}$: The tth OC-y low-speed traffic request from node s to node d employing lightpath (i, j) as an intermediate virtual link.
    - $S_{sd}^{y,t}$ $S_{sd}^{y,t} = 1$ if the tth OC-y low-speed connection request from node s to node d has been successfully routed; otherwise, $S_{sd}^{y,t} = 0$.
- Optimize: Maximize the total successfully-routed low-speed traffic.

$$\text{Maximize}: \sum_{y,s,d,t} y * S_{sd}^{y,t}. \tag{1}$$

- Constraints:
  - — On virtual-topology connection variables

$$\sum_j V_{ij} \le \mathrm{TR}_i \quad \forall i \tag{2}$$

$$\sum_i V_{ij} \le \mathrm{RR}_j \quad \forall j \tag{3}$$

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

$$\mathrm{int}\, V_{ij}, V_{ij}^w. \tag{5}$$

  - — On physical route variables

$$\sum_m P_{mk}^{ij,w} = \sum_n P_{kn}^{ij,w} \quad \text{if } k \ne i, j \quad \forall i, j, w, k \tag{6}$$

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

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

$$\sum_n P_{in}^{ij,w} = V_{ij}^w \quad \forall i, j, w \tag{9}$$

$$\sum_m P_{mj}^{ij,w} = V_{ij}^w \quad \forall i, j, w \tag{10}$$

$$\sum_{i,j} P_{mn}^{ij,w} \le P_{mn}^w \quad \forall m, n, w \tag{11}$$

$$P_{mn}^{ij,w} \in \{0, 1\}. \tag{12}$$

  - — On virtual-topology traffic variables

$$\sum_i \lambda_{id,y}^{sd,t} = S_{sd}^{y,t} \quad \forall s, d \quad y \in \{1, 3, 12, 48\} \quad t \in [1, \Lambda_{y,sd}] \tag{13}$$

$$\sum_j \lambda_{sj,y}^{sd,t} = S_{sd}^{y,t} \quad \forall s, d, t \quad y \in \{1, 3, 12, 48\} \quad t \in [1, \Lambda_{y,sd}] \tag{14}$$

$$\sum_i \lambda_{ik,y}^{sd,t} = \sum_j \lambda_{kj,n}^{sd,t} \quad \text{if } k \ne s, d \quad \forall s, d, k, t \tag{15}$$

$$\sum_i \lambda_{is,y}^{sd,t} = 0 \quad \forall s, d \quad y \in \{1, 3, 12, 48\} \quad t \in [1, \Lambda_{y,sd}] \tag{16}$$

$$\sum_j \lambda_{dj,y}^{sd,t} = 0 \quad \forall s, d \quad y \in \{1, 3, 12, 48\} \quad t \in [1, \Lambda_{y,sd}] \tag{17}$$

$$\sum_{y,t} \sum_{s,d} y \times \lambda_{ij,y}^{sd,t} \le V_{ij} \times C \quad \forall i, j \tag{18}$$

$$S_{sd}^{y,t} \in \{0, 1\}. \tag{19}$$

- Explanation of equations: The above equations are based on principles of conservation of flows and resources (transceivers, wavelengths, etc.).
  - — Equation (1) shows the optimization objective function.
  - — Equations (2), (3) ensure that the number of lightpaths between node pair (i, j) is less than or equal to the number of transmitters at node i and the number of receivers at node j.
  - — Equation (4) shows that the lightpaths between (i, j) are composed of lightpaths on different wavelengths between node pair (i, j). Note that the value of V_{ij}^w can be greater than 1. For example, in Fig. 1, two lightpaths on the same wavelength w can be set up between node 0 and node 5. One follows route (0, 1, 2, 5), while the other follows route (0, 3, 4, 5).
  - — Equations (6)–(10) are the multicommodity equations (flow conservation) that account for the routing of a lightpath from its origin to its termination. The flow-conservation equations have been formulated in two different ways [5]: i) disaggregate formulation and ii) aggregate formulation. In the disaggregate formulation, every i-j (or s-d) pair corresponds to a commodity, while in the aggregate formulation, all the traffic that is "sourced" from node i (or node s) which corresponds to the same commodity, regardless of the traffic's destination. We employ the disaggregate formulation for the flow-conservation equations since it properly describes the traffic requests between different node pairs. Note that (6)–(10) employ the wavelength-continuity constraint on the lightpath route.
    - Equation (6) ensures that, for an intermediate node k of lightpath (i, j) on wavelength w, the number of incoming lightpath streams is equal to the number of outgoing lightpath streams.

- Equation (7) ensures that, for the origin node i of lightpath (i, j) on wavelength w, the number of incoming lightpath streams is 0.
- Equation (8) ensures that, for the termination node j of lightpath (i, j) on wavelength w, the number of outgoing lightpath streams is 0.
- Equation (9) ensures that, for the origin node i of lightpath (i, j) on wavelength w, the number of outgoing lightpath streams is equal to the total number of lightpaths between node pair (i, j) on wavelength w.
- Equation (10) ensures that, for the termination node j of lightpath (i, j) on wavelength w, the number of incoming lightpath streams is equal to the total number of lightpaths between node pair (i, j) on wavelength w.

— Equations (11), (12) ensure that wavelength w on one fiber link (m, n) can only be present in at most one lightpath in the virtual topology.

— Equations (13)–(19) are responsible for the routing of low-speed traffic requests on the virtual topology, and they take into account the fact that the aggregate traffic flowing through lightpaths cannot exceed the overal wavelength (channel) capacity.

B. Single-Hop Traffic Grooming

In this section, we assume that a connection can only traverse a single lightpath, i.e., only end-to-end traffic grooming is allowed. The formulation of the single-hop traffic grooming problem is similar to the formulation of the multihop traffic-grooming problem, which was presented in the previous section, except for routing of connection requests on the virtual topology. We only present the different part as follows.

- On virtual-topology traffic variables

$$\sum_{y,t} y \times S_{sd}^{y,t} \leq V_{sd} \times C \quad \forall s, d \tag{20}$$

$$S_{sd}^{y,t} \in \{0, 1\}. \tag{21}$$

C. Formulation Extension for Fixed-Transceiver Array

The mathematical formulations in the previous two sections are based on the assumption that the transceivers in a network node are tunable to any wavelength. If fixed-transceiver arrays are used at every network node and if M denotes the number of fixed-transceiver arrays used at each node, we can easily extend our formulation as follows.

- On virtual-topology connection variables

$$\sum_{j} V_{ij}^{w} \leq M \quad \forall i, w \tag{22}$$

$$\sum_{i} V_{ij}^{w} \leq M \quad \forall j, w \tag{23}$$

$$\sum_{w} V_{ij}^{w} = V_{ij} \quad \forall i, j \tag{24}$$

$$\text{int}\, V_{ij}, V_{ij}^{w}. \tag{25}$$

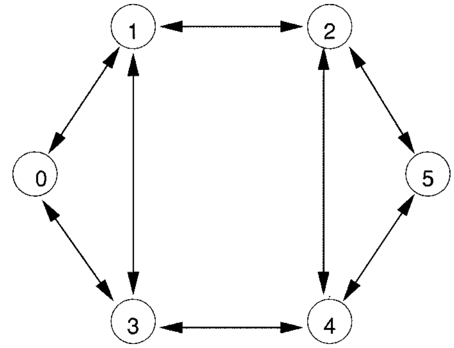

(a)

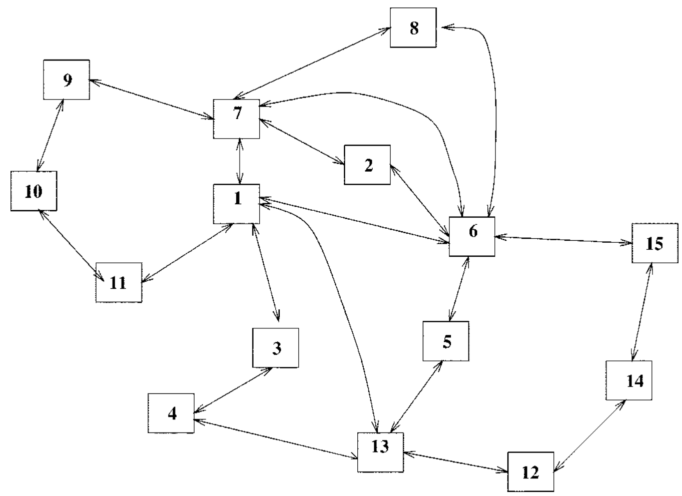

(b)

Fig. 5. (a) A six-node network and (b) a 15-node network.

The other parts of the formulations in the previous two sections are still the same. Equations (22), (23) ensure that the number of lightpaths between node pair (i, j) on wavelength w is less than or equal to the number of transmitters at node i and the number of receivers at node j on the wavelength w (every fixed-transceiver array only has one transceiver on each wavelength).

D. Computational Complexity

It is well known that the RWA optimization problem is NP-complete [1]. If we assume that each connection request requires the full capacity of a lightpath, the traffic grooming problem we are studying becomes the standard RWA optimization problem. It is easy to see that the traffic-grooming problem in a mesh network is also a NP-complete problem since the RWA optimization problem is NP-complete. As the number of variables and equations increases exponentially with the size of network and the number of wavelengths on each fiber, we use a small network topology as an example for obtaining ILP result. For large networks, we will use heuristic approaches.

V. Illustrative Numerical Result From ILP Formulations

This section presents numerical examples of the traffic-grooming problem using Fig. 5(a) as our physical topology. The traffic matrices are randomly generated. As an

TABLE I
TRAFFIC MATRIX OF OC-1 CONNECTION REQUESTS

	Node 0	Node 1	Node 2	Node 3	Node 4	Node 5
Node 0	0	5	4	11	12	9
Node 1	0	0	8	5	16	6
Node 2	14	12	0	9	6	16
Node 3	4	11	15	0	1	5
Node 4	10	2	3	3	0	9
Node 5	2	1	8	15	13	0

TABLE II
TRAFFIC MATRIX OF OC-3 CONNECTION REQUESTS

	Node 0	Node 1	Node 2	Node 3	Node 4	Node 5
Node 0	0	6	2	1	5	4
Node 1	8	0	8	6	7	8
Node 2	1	3	0	0	2	7
Node 3	5	7	3	0	2	6
Node 4	6	4	5	0	0	2
Node 5	5	4	4	2	0	0

TABLE III
TRAFFIC MATRIX OF OC-12 CONNECTION REQUESTS

	Node 0	Node 1	Node 2	Node 3	Node 4	Node 5
Node 0	0	1	1	1	0	0
Node 1	1	0	1	1	0	2
Node 2	0	1	0	2	1	0
Node 3	2	0	2	0	2	0
Node 4	1	2	0	2	0	1
Node 5	1	1	2	2	2	0

example, we allow the traffic demand to be any one of OC-1, OC-3, and OC-12. The traffic matrices are generated as follows: 1) the number of OC-1 connection requests between each node pair is generated as a uniformly distributed random number between 0 and 16; 2) the number of OC-3 connection requests between each node pair is generated as a uniformly distributed random number between 0 and 8; and 3) the number of OC-12 connection requests between each node pair is generated as a uniformly distributed random number between 0 and 2. These three traffic matrices are shown in Tables I–III, and the total traffic demand turns out to be the equivalent of OC-988. The capacity of each wavelength (channel) is OC-48.

Table IV shows the corresponding result for the network throughput obtained from a commercial ILP solver, "CPLEX", based on different network resource parameters. In Table IV, T denotes the number of transceivers and W denotes the number of wavelengths. In the single-hop case, we only allow a connection to transverse a single lightpath, which means that only end-to-end traffic-grooming (multiplexing) is allowed. In the multihop case, a connection is allowed to traverse multiple lightpaths, i.e., a connection can be dropped at intermediate nodes and groomed with other low-speed connections on different lightpaths before it reaches its destination. Fig. 1(b) shows a multihop grooming case, where connection 3 traversed two lightpaths; it was groomed with connection 1 on lightpath (0, 2) and with connection 2 on lightpath (2, 4). As expected, the multihop case leads to higher throughput than the single-hop case.

We can see from Table IV that, when the number of tunable transceivers at each node is increased from 3 to 5, the network throughput increases significantly, both in the multihop case and in the single-hop case. But when the number of tunable transceivers at each node increases from 5 to 7, network throughput does not improve. This is because there are not enough wavelengths to setup more lightpaths to carry the blocked connection requests. Some illustrative results of transceiver and wavelength utilization for the multihop case are shown in Tables V and VI.

In multihop case, when the transceiver is not a limited resource compared with wavelength, more short lightpaths may be set up to carry connections through multiple lightpaths, but this scenario is less likely in the single-hop case. This is shown in Table IV where $T = 5, W = 3$ and $W = 4$. When $T = 5, W = 4$, if multihop grooming is allowed, the network throughput is 100%; otherwise, some connections get blocked. In the multihop case, 29 lightpaths are established compared with 28 lightpaths in the single-hop case.

Tables V and VI show some results for the node transceiver utilization and link wavelength utilization for the multihop case. When the number of transceivers is increased (from 3 to 5), the overall wavelength utilization is increased, as shown in Table VI. This is because more lightpaths have been established to carry the connection requests, shown in Table IV. As we mentioned, when most of the links have fully utilized the available wavelengths, increasing the number of transceivers (from 5 to 7) will not help to improve the network throughput and will result in lower transceiver utilization, shown in Table V ($T = 7$ and $W = 3$).

Table VII shows the virtual topology and the lightpath capacity utilization for the multihop case, when $T = 5$ and $W = 3$. As we can see, most of the lightpaths have high capacity utilization (above 90%). There are some node pairs ((0, 1), (1, 3), etc.) which have multiple lightpaths set up, though the aggregate traffic between them can be carried by a single lightpath. The extra lightpaths are used to carry multihop connection traffic.

In the ILP formulation, we treat the low-speed connection requests separately. The results from the ILP solutions show that, if there is a lightpath set up between (s, d), the low-speed connections between (s, d) tend to be packed on this lightpath channel directly. Based on this observation, we propose two simple heuristic algorithms for solving the traffic-grooming problem in large networks.

VI. HEURISTIC APPROACH

The optimization problem of traffic grooming is NP-complete. It can be partitioned into the following four subproblems, which are not necessarily independent.

1) Determine a virtual topology, i.e., determine the number of lightpaths between any node pair.
2) Route the lightpaths over the physical topology.
3) Assign wavelengths optimally to the lightpaths.
4) Route the low-speed connection requests on the virtual topology.

TABLE IV
THROUGHPUT AND NUMBER OF LIGHTPATHS ESTABLISHED (TOTAL TRAFFIC DEMAND IS OC-988)

	Multi-hop		Single-hop	
	Throughput	Lightpath #	Throughput	Lightpath #
T=3, W=3	74.7%(OC-738)	18	68.0%(OC-672)	18
T=4, W=3	93.8%(OC-927)	24	84.1%(OC-831)	24
T=5, W=3	97.9%(OC-967)	28	85.7%(OC-847)	24
T=7, W=3	97.9%(OC-967)	28	85.7%(OC-847)	24
T=3, W=4	74.7%(OC-738)	18	68.0%(OC-672)	18
T=4, W=4	94.4%(OC-933)	24	84.7%(OC-837)	24
T=5, W=4	100% (OC-988)	29	95.5%(OC-944)	28

TABLE V
RESULTS: TRANSCEIVER UTILIZATION (MULTIHOP CASE)

		T=3, W=3	T=5, W=3	T=7, W=3
Node 0	Transmitter	100%	100%	71.4%
	Receiver	100%	100%	71.4%
Node 1	Transmitter	100%	100%	71.4%
	Receiver	100%	100%	71.4%
Node 2	Transmitter	100%	100%	71.4%
	Receiver	100%	100%	71.4%
Node 3	Transmitter	100%	100%	71.4%
	Receiver	100%	100%	71.4%
Node 4	Transmitter	100%	80%	57.4%
	Receiver	100%	80%	57.4%
Node 5	Transmitter	100%	80%	57.4%
	Receiver	100%	80%	57.4%

TABLE VI
RESULTS: WAVELENGTH UTILIZATION (MULTIHOP CASE)

	T=3, W=3	T=5, W=3	T=7, W=3
Link (0,1)	33.3%	100%	100%
Link (0,3)	100%	100%	100%
Link (1,0)	100%	100%	100%
Link (1,2)	100%	100%	100%
Link (1,3)	33.3%	66.7%	66.7%
Link (2,1)	100%	100%	100%
Link (2,4)	66.7%	100%	100%
Link (2,5)	66.7%	100%	100%
Link (3,0)	33.3%	100%	100%
Link (3,1)	100%	66.7%	66.7%
Link (3,4)	66.7%	100%	100%
Link (4,2)	66.7%	100%	100%
Link (4,3)	66.7%	100%	100%
Line (4,5)	66.7%	66.7%	66.7%
Link (5,2)	66.7%	100%	100%
Link (5,4)	66.7%	66.7%	66.7%

A. Routing

The routing and wavelength assignment problem (RWA) has received a lot of attention in the WDM networking literature. The current well-known routing approaches are fixed routing, fixed-alternate routing, and adaptive routing [10].

In fixed routing, the connections are always routed through a predefined fixed route for a given source-destination pair. One example of such an approach is fixed shortest path routing. The shortest path route for each source-destination pair is calculated offline using standard shortest path algorithms, such as Dijkstra's algorithm. If there are not enough resources to satisfy a connection request, the connection gets blocked.

In fixed-alternate routing, multiple fixed routes are considered when a connection request comes. In this approach, each node in the network is required to maintain a routing table that contains an ordered list of a number of fixed routes to each destination node. For example, these routes can be the first shortest path, the second shortest path, etc. When a connection request comes, the source node attempts to establish the connection on each of the routes from the routing table in sequence, until the connection is successfully established. Since fixed-alternate routing provides simplicity of control for setting up and tearing down connections, it is also widely used in the dynamic connection-provisioning case. It has been shown that, for certain networks, having as few as two alternate routes provides significantly lower blocking than having full wavelength conversion at each node with fixed routing [17].

In adaptive routing, the route from a source node to a destination node is chosen dynamically, depending on the current network state. The current network state is determined by the set of all connections that are currently in progress [10]. For example, when a connection request arrives, the current shortest path between the source and the destination is calculated based on the available resources in the network; then, the connection is established through the route. If a connection gets blocked in the adaptive-routing approach, it will also be blocked in the fixed-alternate routing approach. Since each time a new connection request comes to a node, the route needs to be calculated based on the current network state, adaptive routing will require more computation and a longer response time than fixed-alternate routing, but it is also more flexible than fixed-alternate routing.

In our heuristics, we will use adaptive routing.

B. Wavelength Assignment

Once the route has been chosen for each lightpath, the number of lightpaths going through a physical fiber link defines the congestion on that particular link. With the wavelength-continuity constraint, we need to assign wavelengths to each lightpath such that any two lightpaths passing through the same physical link are assigned different wavelengths.[4] Ten wavelength-assignment approaches have been compared in [10], and all of them were found to perform similarly. We will use one simple approach, first-fit (FF). In FF, all wavelengths are numbered. When searching for an available wavelength, a lower numbered

[4]We assume a single-fiber network system. There is only one fiber in each direction if two nodes are connected.

TABLE VII
RESULT: VIRTUAL TOPOLOGY AND LIGHTPATH UTILIZATION (MULTIHOP CASE WITH $T = 5$ AND $W = 3$)

	Node 0	Node 1	Node 2	Node 3	Node 4	Node 5
Node 0	0	2 (70%)	0 (100%)	1 (89%)	1 (100%)	1 (100%)
Node 1	1 (100%)	0	1 (100%)	2 (100%)	1 (100%)	0
Node 2	1 (100%)	1 (95%)	0	1 (100%)	1 (100%)	1 (70%)
Node 3	2 (100%)	1 (100%)	1 (100%)	0	0	1 (100%)
Node 4	1 (100%)	1 (100%)	0	0	0	1 (91%)
Node 5	0 (100%)	0	2 (98%)	1 (100%)	1 (100%)	0

wavelength is considered before a higher numbered wavelength. The first available wavelength is then selected. The idea behind this simple scheme is that it tries to pack all of the in-use wavelengths toward the low end of the wavelength space.

C. Heuristics

We propose two heuristic algorithms for the traffic-grooming problem. Let $T(s,d)$ denote the aggregate traffic between node pair s and d, which has not been successfully carried. Let $t(s,d)$ denote one connection request between s and d, which has not been successfully carried yet. Let C denote the wavelength capacity.

- *Maximizing Single-Hop Traffic* (MST). The basic idea of this heuristic is introduced in [1] for the traditional virtual-topology design problem. This simple heuristic attempts to establish lightpaths between source-destination pairs with the highest $T(s,d)$ values, subject to constraints on the number of transceivers at the two end nodes and the availability of a wavelength in the path connecting the two end nodes. The connection requests between s and d will be carried on the new established lightpath as much as possible. If there is enough capacity in the network, every connection will traverse a single lightpath hop. If there are not enough resources to establish enough lightpaths, the algorithm will try to carry the blocked connection requests using currently available spare capacity of the virtual topology. The pseudocode for this heuristic follows.

```
Step 1: Construct virtual topology:
1.1: Sort all of the node pairs (s,d)
     according to the sum of uncarried
     traffic request T(s,d) between (s,d)
     and put them into a list L in de-
     cending order.
1.2: Try to setup a lightpath between
     the first node pair (s',d') in L using
     first-fit wavelength assignment
     and shortest-path routing, sub-
     ject to the wavelength and tran-
     ceiver constraints. If it fails,
     delete (s',d') from L; otherwise, let
     T(s,d) = Max[T(s,d)-C, 0] and go to Step
     1.1 until L is empty.
Step 2: Route the low-speed connections on
     the virtual topology constructed
     in Step 1.
2.1: Statisfy all of the connection re-
     quests which can be carried through
     single lightpath hop, and update the
     virtual topology network state.
2.2: Route the remaining connection re-
     quests based on the current virtual
     topology network state, in the de-
     cending order of the connections'
     bandwidth requirement.
```

- *Maximizing Resource Utilization* (MRU). Let $H(s,d)$ denote the hop distance on physical topology between node pair (s,d). Define $T(s,d)/H(s,d)$ as the connection resource utilization value, which represents the average traffic per wavelength link. This quantity shows how efficiently the resources have been used to carry the traffic requests. This heuristic tries to establish the lightpaths between the node pairs with the maximum resource utilization values. When no lightpath can be set up, the remaining blocked traffic requests will be routed on the virtual topology based on their connection resource utilization value $t(s,d)/H'(s,d)$, where $t(s,d)$ denotes a blocked connection request, and $H'(s,d)$ denotes the hop distance between s and d on the virtual topology. The pseudocode for this heuristic follows the same steps as the pseudocode of MST, except that the node pairs and blocked connections are sorted based on their resource utilization values.

Both heuristic algorithms have two stages. Based on our observations from the ILP results, we find that packing different connections between the same node pair within the same existing lightpath, which directly joins the end points, is a very efficient grooming scheme. In the first stage, the algorithms try to establish lightpaths as much as possible to satisfy the aggregate end-to-end connection requests. If there are enough resources in the network, every connection request will be successfully routed through a single lightpath hop. This will minimize the traffic delay. In the second stage, the spare capacity of the currently established lightpath channels is used to carry the connection requests blocked in the first stage, and the algorithms give single-hop groomable connections high priority to be satisfied.

D. Heuristic Results and Comparison

Table VIII shows a comparison between the results obtained from ILP solver and the heuristic algorithms for the six-node network in Fig. 5(a). We can observe that the MST and MRU heuristic algorithms show reasonable performance when compared with the results obtained from the ILP solver. The heuristic approaches have much less computation complexity than the ILP approach. The two proposed algorithms are

TABLE VIII
THROUGHPUT RESULTS COMPARISON BETWEEN ILP AND HEURISTIC ALGORITHMS (TOTAL TRAFFIC DEMAND IS OC-988)

	Multi-hop (ILP)	Single-hop (ILP)	Heuristic (MST)	Heuristic (MRU)
T=3, W=3	74.7%(OC-738)	68.0%(OC-672)	71.0%(OC-701)	67.4%(OC-666)
T=4, W=3	93.8%(OC-927)	84.1%(OC-831)	89.4%(OC-883)	93.6%(OC-925)
T=5, W=3	97.9%(OC-967)	85.7%(OC-847)	94.4%(OC-933)	94.4%(OC-933)
T=7, W=3	97.9% (OC-967)	85.7%(OC-847)	94.4%(OC-933)	94.4%(OC-933)
T=3, W=4	74.7%(OC-738)	68.0%(OC-672)	71.0%(OC-701)	67.4%(OC-666)
T=4, W=4	94.4%(OC-933)	84.7%(OC-837)	93.1%(OC-920)	93.6%(OC-925)
T=5, W=4	100% (OC-988)	95.5%(OC-944)	100%(OC-988)	100%(OC-988)

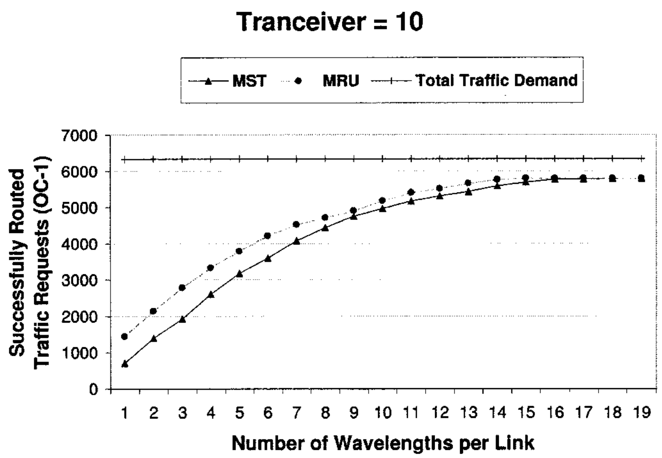

Fig. 6. Network throughput versus number of wavelengths for the network topology in Fig. 5(b) with ten tunable transceivers at each node.

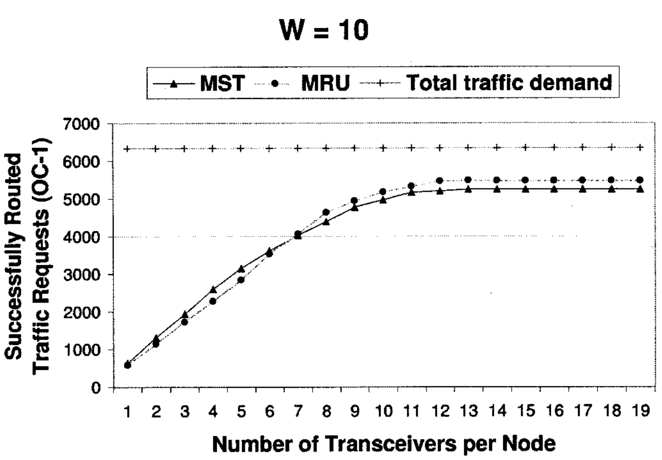

Fig. 7. Network throughput versus number of tunable transceivers for the network topology in Fig. 5(b) with ten wavelengths on each fiber link.

relatively simple and straightforward; by using other RWA algorithms instead of adaptive routing and first-fit wavelength assignment, it may be possible to develop other elaborate heuristic algorithms to achieve better performance.

Figs. 6–8 show the results from the two heuristic algorithms, when applied to the larger network topology in Fig. 5(b). The traffic matrices follow the same distribution as we mentioned in Section V.

In Fig. 6, we plot the network throughput versus the number of wavelengths on every fiber link when each node is

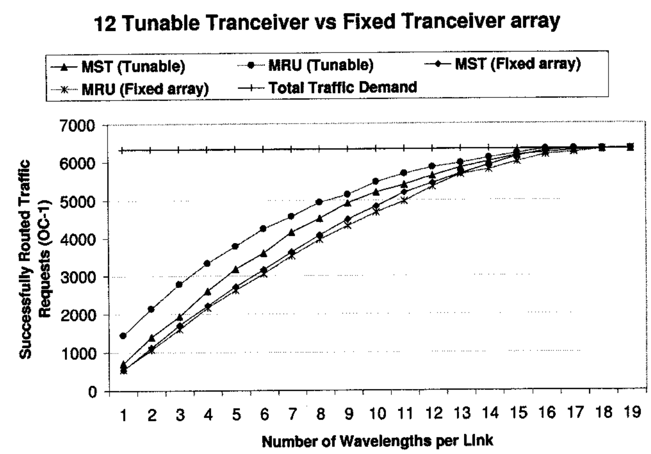

Fig. 8. Network throughput versus number of wavelengths (size of fixed transceiver array) for the network topology in Fig. 5(b) with 12 tunable transceivers at each node.

equipped with ten tunable transceivers. We observe that the MRU heuristic performs better than the MST algorithm with respect to network throughput. Since the number of tunable transceivers at each node is limited (=10 in this case), when the number of wavelengths on each fiber link reaches a certain value (16 in this case), increasing the number of wavelengths does not help to increase the network throughput.

In Fig. 7, we plot the network throughput versus the number of transceivers at every node when each fiber link carries ten wavelengths. We compare the performance of the two heuristic algorithms. The results show that, when the number of transceivers is small (≤ 7 in this case), MST performs better than MRU. This is because MRU tries to utilize wavelengths efficiently. When the number of transceivers is small, wavelength is not the limiting resource in the network any more. So maximizing wavelength utilization will not help to improve the performance.

Fig. 8 compares the performance using tunable transceiver and fixed transceiver in every network node. Each node is equipped with 12 transceivers if we use a tunable transceiver. Each node is equipped with one transceiver array if we use fixed transceivers and the size of the transceiver array is equal to the number of wavelengths supported by every fiber link. The results in Fig. 8 indicate that, when nodes are equipped with the same number of transceivers, the tunable-transceiver architecture has better performance. For the fixed-transceiver case, MST performs better than MRU.

VII. Mathematical Formulation Extension for Other Optimization Criteria

In this section, we show how to extend our ILP formulations to handle different optimization criteria for the traffic-grooming problem.

A. Extension for Network Revenue Model

It has been shown earlier that the low-speed connection requests between the same node pair tend to be packed together on to the same lightpath channel. The connections, which can be carried by a single lightpath channel are more likely to be satisfied than the connections which have to traverse multiple lightpaths, when they have the same bandwidth requirement and the optimization objective is to maximize network throughput. To make the problem more realistic, it is reasonable for us to assume that two connection requests may have different priority, even if they have the same bandwidth requirement. This is because different connection requests may have different end-node distance, quality-of-service requirement, etc. A connection's priority can be represented by a "weight" associated with it. In this section, we assume that the weight is determined by the bandwidth requirement and end-node distance of the connection request. For a given network topology and traffic demand, the objective is to maximize the weighted network throughput. Let W_i denote the weight of connection i, D_i denote the end-node distance of connection i, and C_i denote the bandwidth requirement of connection i. The connection's weight function is defined as

$$W_i = D_i \times C_i^{\alpha} \tag{26}$$

where $0 \leq \alpha \leq 1$ and D_i is measured by the shortest path distance of the connection's end nodes on the physical topology. Equation (26) is called the power-law cost function [18]. It is used to study the actual tariffs demanded by communications services for high-speed telecommunication channels, and there is effectively a "quantity discount" (controlled by α) in that capacity cost (per unit of channel capacity) decreases as the capacity increases. Thus, the network's weighted throughput becomes

$$T = \sum_{i=1}^{H} D_i * C_i^{\alpha} * S_i \tag{27}$$

where $S_i = 1$ if connection request i has been satisfied; otherwise $S_i = 0$, and the total number of connection requests is H. T can also be called "*network revenue*". We can easily modify our ILP formulation to optimize T. The only part of the equation which should be modified is shown as follows:

- Optimize: Maximize network revenue

$$\text{Maximize}: \sum_{y,s,d,t} D_{sd} * y^{\alpha} * S_{sd}^{y,t} \tag{28}$$

where D_{sd} denotes the distance between node pair (s, d).

B. Illustrative Results

In this section, we show some illustrative results to optimize network revenue using our ILP formulation extension. We use the same network topology and traffic matrix set as in Section V. In (28), D_{sd} is measured by the shortest path hop distance between node s and d on the physical topology, and α is equal to 0.8.

TABLE IX
Results of Comparison Between Revenue Model and Network Throughput Model

	Optimize Revenue		Optimize Throughput
	Revenue	Throughput	Throughput
T=3, W=3	83.7%	72.4%	74.7%
T=5, W=3	98.5%	97.2%	97.9%
T=7, W=3	98.5%	97.2%	97.9%
T=3, W=4	83.7%	72.4%	74.7%
T=4, W=4	94.3%	91.7%	94.4%
T=5, W=4	100%	100%	100%

Table IX compares the results between the two optimization models. In Table IX, T denotes the number of tunable transceivers per node and W denotes the number of wavelengths per fiber link. Multihop grooming is allowed in both models. It is shown that, in the revenue model, when $T = 3$ and $W = 3$, the maximal achievable revenue is 83.7%, and 72.4% of traffic requests have been satisfied to achieve the revenue, while the maximal achievable traffic load the network can carry is 74.7%. In revenue model, we find that if there is a lightpath set up between (s, d), it may first be used to carry some long multihop connections (with higher weight) which will traverse this lightpath as an intermediate hop. Thus, some connections directly between (s, d) may be blocked. This means that packing different connections between the same node pair within the same existing lightpath, which directly joins the end points, is not a good grooming scheme any more. We find that, because of the quantity-discount parameter α in (26), lower speed connections are more likely to be satisfied than higher speed connection requests. It is obvious that different heuristics are needed based on the different optimization criteria.

VIII. Conclusion

This study was devoted to the traffic-grooming problem in a WDM mesh network. We studied the architecture of a node with grooming capability. We presented the ILP formulation for traffic-grooming in such a WDM mesh network. We compared the performance of the single-hop grooming approach and multihop grooming approach on a small six-node network with randomly generated traffic pattern. Results from ILP showed that the end-to-end aggregate traffic between the same node pair tends to be groomed on to the same lightpath channel, which directly joins the end points, if the optimization objective is to maximize the network throughput. Two heuristic approaches were also proposed for solving the traffic-grooming problem in large networks. We compared the performance of these two heuristic algorithms, MST and MRU, with different network resource parameters. The comparison results showed that MRU performs better if tunable transceivers are used and MST performs better if fixed transceivers are used. We extended the optimization problem to a network-revenue model and found a different grooming scheme, which can be used to design an efficient heuristic algorithm on network-revenue model.

We showed that, with proper extension, our ILP mathematical model can be used to examine good grooming schemes for different models. These schemes can be used to design efficient heuristic algorithms, which are practical for large and realistic networks.

REFERENCES

[1] B. Mukherjee, *Optical Communication Networks*. New York: McGraw-Hill, 1997.

[2] I. Chlamtac, A. Faragó, and T. Zhang, "Lightpath (wavelength) routing in large WDM networks," *IEEE J. Select. Areas Commun.*, vol. 14, pp. 909–913, June 1996.

[3] R. Ramaswami and K. N. Sivarajan, *Optical Networks: A Practical Perspective*. San Francisco, CA: Morgan Kaufmann, 1998.

[4] M. Alanyali and E. Ayanoglu, "Provisioning algorithms for WMD optical networks," *IEEE/ACM Trans. Networking*, vol. 7, pp. 767–778, Oct. 1999.

[5] R. Ramaswami and K. N. Sivarajan, "Design of logical topologies for wavelength-routed optical networks," *IEEE J. Select. Areas Commun.*, vol. 14, pp. 840–851, June 1996.

[6] D. Banerjee and B. Mukherjee, "Wavelength-routed optical networks: Linear formulation, resource budgeting tradeoffs, and a reconfiguration study," *IEEE/ACM Trans. Networking*, vol. 8, pp. 598–607, Oct. 2000.

[7] ——, "A practical approach for routing and wavelength assignment in large wavelength-routed optical networks," *IEEE J. Select. Areas Commun.*, vol. 14, pp. 903–908, June 1996.

[8] X. Zhang and C. Qiao, "Wavelength assignment for dynamic traffic in multi-fiber WDM networks," in *Proc. 7th Int. Conf. Computer Communications and Networks*, Oct. 1998, pp. 479–485.

[9] R. A. Barry and S. Subramaniam, "The MAX-SUM wavelength assignment algorithm for WDM ring networks," presented at the Proc. OFC'97, Feb. 1997.

[10] H. Zang, J. P. Jue, and B. Mukherjee, "A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks," *Optical Network Mag.*, vol. 1, no. 1, pp. 47–60, Jan. 2000.

[11] O. Gerstel, P. Lin, and G. Sasaki, "Combined WDM and SONET network design," in *Proc. INFOCOM'98*, vol. 2, San Francisco, CA, Mar. 1998, pp. 734–743.

[12] X. Zhang and C. Qiao, "On scheduling all-to-all personalized connections and cost-effective designs in WDM rings," *IEEE/ACM Trans. Networking*, vol. 7, pp. 435–443, June 1999.

[13] A. L. Chiu and E. H. Modiano, "Traffic grooming algorithms for reducing electronic multiplexing costs in WDM ring networks," *J. Lightwave Technol.*, vol. 18, pp. 2–12, Jan. 2000.

[14] J. Wang, V. R. Vemuri, W. Cho, and B. Mukherjee, "Improved approaches for cost-effective traffic grooming in WDM ring networks: ILP formulations and single-hop and multihop connections," *IEEE/OSA J. Lightwave Technol.*, vol. 19, pp. 1645–1653, Nov. 2001.

[15] P. J. Wan, L. Liu, and O. Frieder, "Grooming of artitrary traffic in SONET/WDM BLSRs," *Proc. IEEE GLOBECOM'99*, pp. 1012–1016, Nov. 1999.

[16] P. J. Wan, G. Calinescu, and O. Frieder, "Grooming of arbitrary traffic in SONET/WDM BLSRs," *IEEE J. Select. Areas Commun.*, vol. 18, pp. 1995–2003, Oct. 2000.

[17] S. Ramamurthy, "Optimal Design of WDM Network Architectures," Ph.D. dissertation, Dept. Computer Science, Univ. California, Davis, CA, Sept. 1998.

[18] L. Kleinrock, "Analytic and simulation methods in computer network design," in *AFIPS Conf., 1970 Spring Joint Computer Conf.*, vol. 42, 1970, pp. 569–579.

Keyao Zhu (S'98) received the B.S. degree from Peking University, Beijing, China, in 1998, and the M.S. degree from the University of California, Davis, in March 2001. Currently, he is a Ph.D. student in the Computer Science Department, University of California, Davis.

He works as a Research Assistant in the networks laboratory at the University of California, Davis. His research interests include wavelength-routed WDM network design and analysis and WDM network protection and restoration.

Biswanath Mukherjee (S'84–M'87) received the B.Tech. (Hons.) degree from the Indian Institute of Technology, Kharagpur, India, in 1980, and the Ph.D. degree from the 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 University of California, Davis, where he has been a Professor of Computer Science since July 1995 and Chairman of the Computer Science Department from September 1997 to June 2000. He is author of the textbook *Optical Communication Networks* (New York: McGraw-Hill, 1997), a book which received the Association of American Publishers, Inc.'s 1997 Honorable Mention in Computer Science. His research interests include lightwave networks, network security, and wireless networks.

Dr. Mukherjee is a co-winner of paper awards presented at the 1991 and the 1994 National Computer Security Conferences. He serves or has served on the editorial boards of the IEEE/ACM TRANSACTIONS ON NETWORKING, *IEEE Network*, *ACM/Baltzer Wireless Information Networks* (WINET), *Journal of High-Speed Networks*, *Photonic Network Communications*, and *Optical Network Magazine*. He also served as Editor-at-Large for optical networking and communications for the IEEE Communications Society. He served as the Technical Program Chair of the IEEE INFOCOM'96 Conference.