TRAFFIC GROOMING IN WDM NETWORKS

A THESIS SUBMITTED TO THE GRADUATE DIVISION OF THE UNIVERSITY OF HAWAI'I IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE

IN

ELECTRICAL ENGINEERING

DECEMBER 2003

By
Ramesh Kandula

Thesis Committee:

Galen Sasaki, Chairperson
N T Gaarder
Anders Host-Madsen

To my brother Bhagat Singh.

Acknowledgements

I would like to thank Dr. Galen Sasaki for his enthusiastic supervision and patience. I thank Dr. N.T. Gaarder and Dr. Anders Host-Madsen for their suggestions. Finally, I thank my family members for their patience and support.

ABSTRACT

Wavelength Division Multiplexing (WDM) is the promising technology, which provides very wide bandwidths to support the exponential growth of Internet traffic. A major part of the WDM network cost is the cost of opto-electronic switching components. By grooming low speed traffic connections efficiently onto the lightpaths, the cost of the WDM network can be reduced significantly. In the first part of the thesis, different WDM ring network architectures under a non-statistical traffic model are considered. The average number of Line Terminating Equipments (LTEs) required per node in these networks is significantly lower when compared to a ring network with opaque nodes. To utilize the bandwidth efficiently, the low speed traffic connections have to be rearranged by bridge and roll. In the second part of the thesis, grooming in WDM mesh networks under a static traffic model is considered. The main objective here is to minimize the total lightpaths required. A greedy heuristic algorithm for the lightpath minimization problem is presented.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Wavelength Division Multiplexing (WDM) is the current transmission technology of choice for deployment in back-bone networks to satisfy the exponentially increasing bandwidth requirements. WDM is essentially multiplexing multiple optical signals at different wavelengths onto the same fiber. WDM exploits the huge opto-electronic bandwidth mismatch by requiring that each end-users equipment operate only at electronic rate, but multiple WDM channels from different end-users may be multiplexed on the same fiber [17]. The major advantages of WDM are its low cost per unit of bandwidth, good quality of service and very wide bandwidth. In WDM networks the major component of the cost is no longer the cost of fiber, but the cost of switching components.

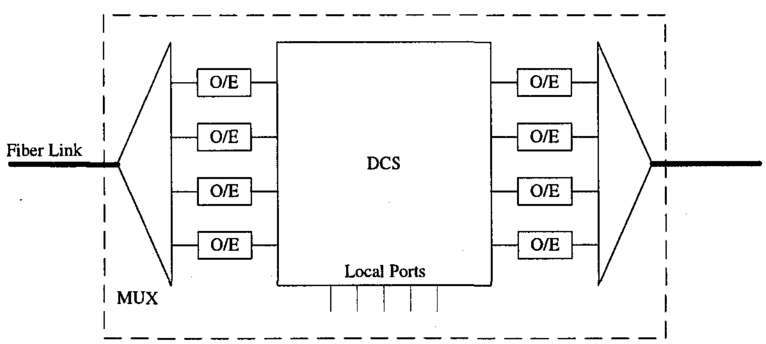

Figure 1.1: Opaque node.

A fiber-link is composed of a pair of counter-propagating fibers that carry optical signals at wavelengths $\lambda_1, \lambda_2, ..., \lambda_n$, where n is the number of wavelengths. The wavelengths function as virtual fibers each carrying an optical communication signal. The straight forward approach to implement switching at a node in a WDM network is shown in Figure 1.1. Such a node is referred to as an *opaque node*, since all the wavelengths are terminated at the node. Each wavelength is terminated by opto-electrical (O/E) equipment known as *line terminating equipment* (LTE). Traffic may be switched electronically between fiber-links by *digital cross-connect* systems (DCS) or packet routers. A drawback of opaque nodes is that they require many LTEs, which are a major contributor to the equipment cost.

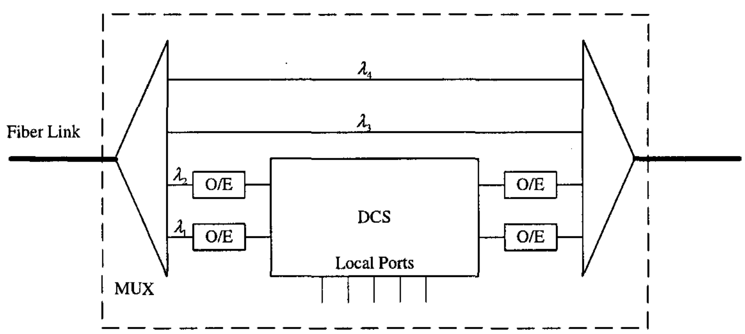

Figure 1.2: Optical node.

We refer to the node shown in Figure 1.2 as an *optical node*. Note that only a fraction of the wavelengths are terminated by the LTEs at this node, and the remaining wavelengths are pass-through. We refer to the wavelengths that pass-through as the *transit* wavelengths for the node, while the wavelengths that terminate at LTEs are referred to as the *local* wavelengths for the node. Transit traffic can use the transit wavelengths, while the traffic that terminates at the node must use local wavelengths. We refer to the traffic that terminate at the node as local traffic. In the Figure 1.2, λ_1 and λ_2 are local wavelengths and λ_3 and λ_4 are transit wavelengths for the node. The optical node shown in Figure 1.3 has a different architecture. Here, the transit wavelengths are terminated at an *optical cross-connect* (OXC) instead of bypassing the node. Depending on the functionality available at the OXC, the wavelength switching at the OXC is either static or reconfigurable. The OXC at an optical node could provide no wavelength conversion, fixed wavelength conversion, limited wavelength conversion or full wavelength conversion. The cost of the optical nodes shown in Figures 1.2 and 1.3 is significantly lower than the cost of an opaque node since they require fewer LTEs.

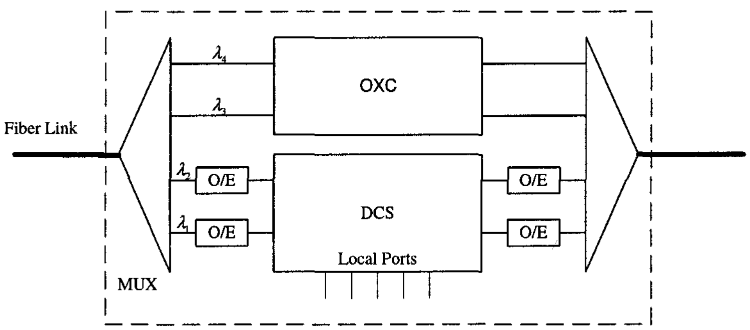

Figure 1.3: Optical node with an OXC.

The optical layer in the WDM network provides *lightpaths* to the higher layer. A lightpath is an end-to-end connection established across the WDM network, and uses a wavelength on each fiber link in a path between the source and destination nodes. In a WDM network with wavelength conversion, a lightpath could occupy different wavelengths on different fiber links. The presence of OXCs at the optical nodes in the WDM network determines if wavelength conversion is possible. The signal traversing a lightpath remains completely in the optical domain. A lightpath via appropriate routing and wavelength assignment creates logical neighbors out of nodes that are physically apart. The lightpath topology is usually called a *logical* (*virtual*) topology.

Given the lightpath demands and the physical network, assigning the lightpaths to the wavelengths in the physical network that minimizes the number of wavelengths is called the *routing and wavelength assignment* (RWA) problem. When the routes for the lightpaths are fixed, the RWA problem becomes a *graph-coloring* problem which is a well-known NP-Complete problem and thus, in general the RWA problem is NP-complete [8]. A number of RWA studies have been reported in the optical networking literature, e.g., [1], [8], [12], and [13].

In practice, however the bandwidth requirement for a connection request is rarely the capacity of an entire wavelength, which is a common assumption in the RWA studies. Furthermore, minimizing the wavelengths does not necessarily minimize the LTEs as well. In this thesis, we consider the connection requests to be a fraction of the lightpath capacity. The traffic is composed of *tributary* traffic streams, which we refer to as tributary connections. The connections are bidirectional and have *constant bit rates* (CBR). All connections have a common bit rate b. We assume that b divides B, and let g denote B/b. Thus, g tributary connections can be carried on a wavelength. For example, if $B = 2.5$ Gbps (OC-48) and $b = 155$ Mbps (OC-3) then $g = 16$.

In WDM network design one of the main objectives is to minimize the overall cost of the network. The major component of the cost is the opto-electronic switching equipment cost. Packing the low-speed connection requests unto the high-capacity lightpaths efficiently reduces the switching equipment required. This is called the traffic *grooming* problem. In general, given the traffic demands, traffic grooming involves designing a virtual (lightpath) topology, assign the lightpaths to wavelengths on the physical topology and route each traffic demand on the virtual topology. The goal is to minimize a well defined cost function, which in most scenarios is the total number of LTEs required.

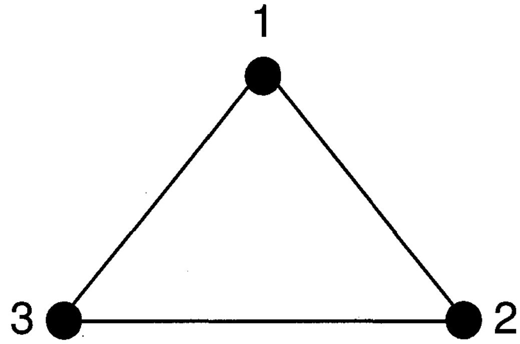

Figure 1.4: A 3 node physical network.

To illustrate traffic grooming, consider a network with three nodes. Let us assume that lightpaths are OC-48 and the connection requests are OC-3. Thus, a lightpath can carry 16 low speed connection requests. The nodes are numbered 1, 2 and 3 in the clockwise direction as shown in Figure 1.4. There are 8 connection requests between each pair of nodes. If the low speed connection requests are not groomed in the network, a lightpath is needed to be set up between each pair of nodes and thus, a total of 6 LTEs will be required. The routing of the connection requests uses the shortest path routing. Since there are lightpaths between every pair of nodes, the traffic between a pair of nodes is routed on the lightpath between that pair of nodes.

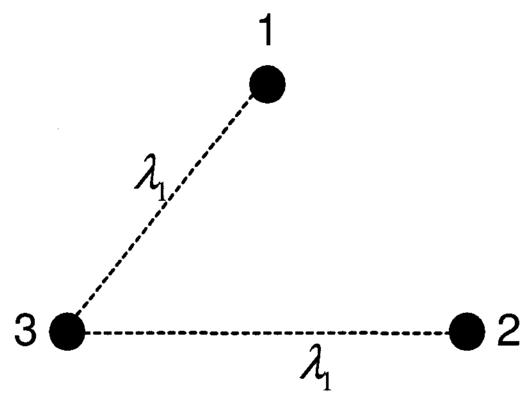

Figure 1.5: A logical topology that grooms traffic streams.

However, if the connection requests are groomed efficiently two lightpaths would suffice to support the traffic as shown in Figure 1.5 and only 4 LTEs will be required. The traffic between nodes (1,3) and (2,3) is routed through the lightpaths between nodes (1,3) and (2,3) respectively. Since a lightpath can support 16 low-speed connections, the traffic between nodes (1,2) can be routed through the lightpaths between nodes (1,3) and (2,3). For reasonably large networks, traffic grooming would result in substantial reduction of LTEs.

The majority of the traffic grooming research deals with the ring networks, e.g., [6], [14], [15], [3], [16], [11], and [7]. This is because the optical networking standard SONET is ring oriented. Both static and dynamic traffic models were used in these studies. A survey of the research in traffic grooming can be found in [4].

In the second chapter, we consider the problem of *dynamic* traffic grooming through traffic *rearrangements* in ring networks. Two optical network architectures are proposed for a nonstatistical traffic model which are as bandwidth efficiency as an opaque network. Since the traffic is dynamic, low speed connections have to be rearranged. In the third chapter, we consider the traffic grooming in mesh networks. Static traffic model is used in this chapter. The objective here is to minimize the total lightpaths needed to support the given traffic demands. Towards the end of this chapter, we present a novel greedy heuristic algorithm. In the last chapter, we conclude this thesis.

Chapter 2

Grooming in WDM Ring Networks

The advantage of an optical node is that it requires lesser number of LTEs and therefore lesser equipment cost than an opaque node. The disadvantage is a loss of switching capability, and potentially less efficient use of bandwidth. As an example, consider the cases where a traffic stream is to be set up between two pairs of nodes (A, B) and (A, C). Suppose the nodes A and C are opaque and the node B is optical as shown in Figure 2.1. The traffic stream between nodes A and C could set up if there is bandwidth available on any wavelength in the link (A, B). For the traffic stream between nodes A and B, consider the case where the wavelengths in (A, B) that terminate at the optical node (λ_1, λ_2, and λ_3) is completely full of traffic, while the other wavelengths (λ_4, and λ_5) are empty. Thus, though there is bandwidth available on (A, B), the traffic stream between nodes A and C cannot be set up because there is no way to switch the traffic.

Although optical nodes have less switching capability, bandwidth inefficiencies may be reduced by intelligently grooming traffic streams at nodes. In this chapter we investigate network architectures and traffic grooming algorithms. There are three types of network cost that will be considered: transmission cost, bandwidth cost, and traffic rearrangement cost. Transmission cost is the number of LTEs in the network. Bandwidth cost is loss in bandwidth utilization. Optical nodes may be less bandwidth efficient than opaque nodes. This will be discussed in Section 2.1. We

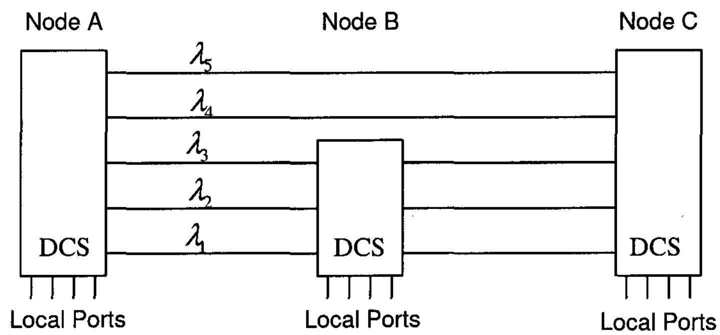

Figure 2.1: A 3 node optical network.

will rely on a non-statistical traffic model that allows arbitrary arrival and departure times of traffic streams. This is also described in the section.

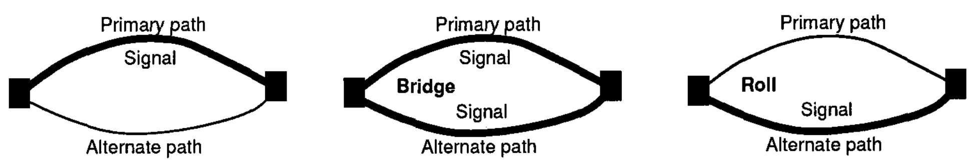

Figure 2.2: Three steps of bridge and roll.

We consider networks that allow traffic rearrangement. In practice, rearranging existing traffic streams are avoided to reduce the chance of disrupting communication service. However, rearrangement may lead to better bandwidth efficiencies. In addition, rearrangement does occur in practice though often for maintenance purposes such as when replacing or adding equipment. One technique of rearrangement is *bridge and roll* (BR), which is illustrated in Figure 2.2. The "cost" of using BR is that there is a chance that it may fail. We will describe BR for our networks in Section 2.2. The traffic rearrangement cost is the maximum number of BRs to make way for a new traffic stream.

In this chapter we will focus our attention on WDM ring networks. Ring topologies are relevant because SONET, a telecommunication standard for optical network transmission, is ring oriented. In Section 2.1 we will describe a *simple WDM ring* network that is as bandwidth efficient as a fully opaque network but with significantly less LTEs. The network does rearrange existing traffic, and in the section we present bounds on the number of rearrangements.

In Section 2.2 we describe a *hierarchical WDM ring* network, which is a generalized version of the simple WDM ring. In certain cases, it leads to a further reduction in LTEs but with more traffic rearrangement cost. The hierarchical ring network has been presented in [6] [15] but for a different traffic model than we consider in Section 2.1. In [6] [15], traffic streams could arrive at any time but never depart. This models high speed connections that have long lifetimes. In addition, traffic rearrangement has been considered for WDM networks, but the traffic that was rearranged are lightpaths rather than tributary traffic streams [10]. We believe our contribution is the first to consider bridge and roll rearrangement of dynamic tributary traffic under nonstatistical traffic assumptions.

2.1 Network and Traffic Models

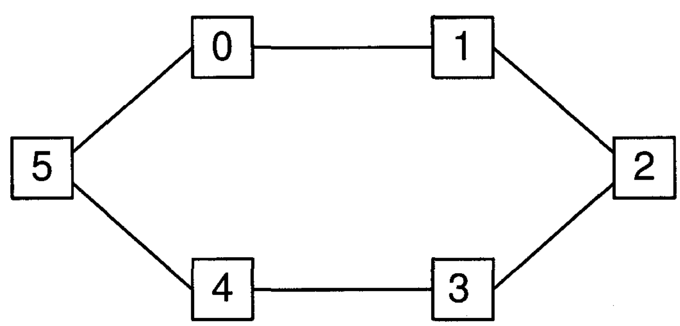

Figure 2.3: A simple ring network with fiber links.

The network has a ring topology as shown in Figure 2.3. The number of nodes is denoted by m, and the nodes are numbered $0, 1, ..., m-1$. The number of wavelengths are denoted by n, and the wavelengths are denoted $1, 2, ..., n$. Each wavelength corresponds to a bidirectional channel. All wavelengths have the same bit rate which we denote by B. For example, B could be 2.5 Gbps (OC-48) or 10 Gbps (OC-192).

Our traffic model is dynamic so that tributary connections can arrive and depart at arbitrary times. We assume that when a tributary connection arrives to be set-up, it already has the route of the links it will traverse. In addition, we assume that for any link, the number of tributary connections for the link at any time is at most gL, where L is an integer parameter. We will refer to this traffic as the *Link Constrained* (LC) Model.

Notice that if the ring network has $n = L$ wavelengths and is fully opaque (i.e., all nodes are opaque) then we can always set up an arriving tributary connection, i.e., there is no blocking. In addition, if there are less than L wavelengths then the fully opaque ring can have blocking. Thus, L is the minimum number of wavelengths for a fully opaque network to support traffic under the LC Model.

Also note that if a ring network with L wavelengths can support the LC Traffic Model then it is necessarily fully opaque. To see this, consider the traffic set where between every pair of adjacent nodes there are tributary connections. Then all fiber-links will be filled with one-hop tributary connections which implies that all their wavelengths must be terminated by LTEs.

Therefore, to make this study more interesting we add the following parameters and constraints to the traffic model. For each node i, we have an integer parameter $t(i)$. We have the constraint that the number of tributary connections that can terminate at node i through any of its incident links is at most $2t(i)$. We refer to the constraint as the node capacity constraint. Note if there are less than $2t(i)$ LTEs at node i then it is possible to get blocking of an arriving tributary connection that terminates at node i. Thus, $2t(i)$ is a lower bound on the number of LTEs to insure no blocking. We refer to the traffic model with the additional node capacity constraint as the *Link and Node Constrained* (LNC) Model.

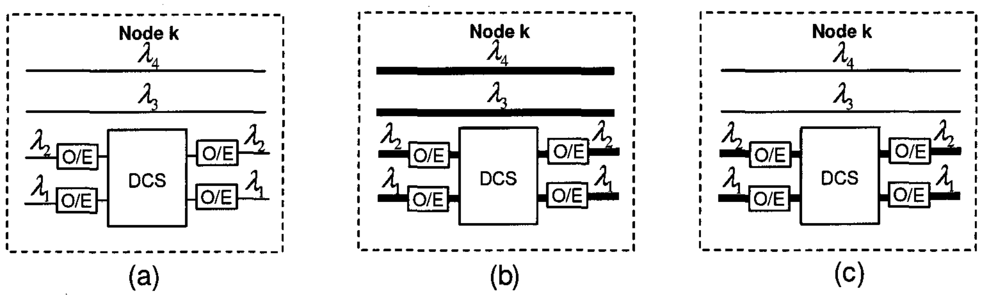

Figure 2.4: (a) An empty optical node k. (b) Transit traffic fills the node. (c) Transit traffic leaves 3 and 4.

In subsequent sections, we will describe WDM ring networks that support traffic under the LNC Model. The networks have $n = L$ wavelengths, so they are as bandwidth efficient as a fully opaque ring. To support traffic under the LNC Model, the networks allow traffic rearrangement. Figure 2.4 illustrates why this is necessary. The figure shows the evolution of tributary connections through an optical node in a ring network. Figure 2.4 (b) shows all wavelengths filled with transit tributary connections. Then the tributary connections that use wavelengths λ_3 and λ_4 terminate leaving the wavelengths empty as shown in Figure 2.4 (c). Now bandwidth is available at these empty wavelengths but a tributary connection cannot be set up from node k because it has no LTEs at those wavelengths. Thus, there is blocking, i.e., the network cannot support LNC traffic. However, we can make room for a new connection if we can move the existing connections in $\{\lambda_1, \lambda_2\}$ to $\{\lambda_3, \lambda_4\}$.

2.2 A Simple WDM Ring With Bridge and Roll

We will describe a ring network which we refer to as the *simple* WDM ring. It has $n = L$ wavelengths. For simplicity, assume that the number of nodes m is even, the odd numbered nodes are opaque, and the even numbered nodes are optical. An example of an optical node is shown in Figure 1.2, where the node k has $t(k) = 2$. At optical node k, its local wavelengths are $\lambda_1, \lambda_2, ..., \lambda_{t(k)}$, and thus they are terminated by LTEs.

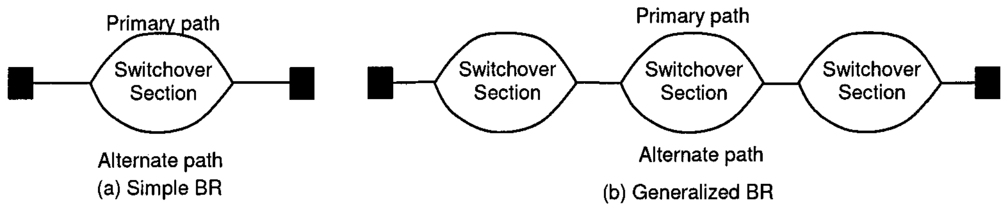

Figure 2.5: Bridge and roll.

We consider a method of rearranging existing connections called bridge-and-roll (BR). As shown in Figure 2.2, BR moves a connection by first finding a path of free bandwidth which we call the alternate path. Then the signal is copied to both the primary and alternate paths. Finally, the signal on the primary path is terminated, and now the connection has completed its migration to the alternate path. This is similar to $1 + 1$ protection switching. More generally, BR is shown in Figure 2.5 (a). Here, the alternate path overlaps the primary path except at one section, which we refer to as the switchover section. Thus, the bridge and roll is really done over the switchover section. Since each BR has some risk of failing, they should be minimized to reduce the likelihood of service disruption. In addition, note that the BR we consider have primary and alternate "paths" that use the same fiber links. These "paths" are paths (or trails) of WDM channels over the same route.

Figure 2.5 (b) shows a more general form of switching from a primary path to an alternate path. There are a multiple k switchover sections, where in this case $k = 3$. We do not consider this further since such a BR can be implemented with k ordinary BRs, where the first BR is over the first switchover section, the second BR is over the second switchover section, and so forth.

The next lemma has a bound on the number of BRs required to set up a tributary connection.

Proposition 2.2.1 *Consider the simple WDM ring and traffic model LNC. Setting up a tributary connection c requires at most q BRs, where q is the number of end nodes of the connection that are optical. Since q is at most 2, the number of BRs to accommodate a new tributary connection is at most 2.*

Proof. Let the route of the connection be the sequence of nodes $(u_1, u_2, ..., u_p)$, where p denotes the number of nodes on the path. First, we will prove the proposition for the case when $q = 0$. Note that the route is made up of sections, where a section is an optical node u_{k+1} surrounded by opaque nodes u_k and u_{k+2}. We will show that we can set up c through such a section. Since we are about to set up c through the section and c can contribute to traffic load on links, the LNC Model implies there must be free bandwidth on links of the section. If there is free bandwidth on a transit wavelength (for u_{k+1}) then we can set up c through that wavelength. If all transit wavelengths are full then there must be free bandwidth on the local wavelengths (for u_{k+1}). Then we can set up c through the free bandwidth by cross-connecting through u_{k+1}. Since we can set up c through each section and the sections are connected by opaque nodes, connection c can be set up. Thus, the proposition is true when $q = 0$.

Consider the case when $q = 1$, and without loss of generality let u_1 be the optical node. Since c adds to the load, the LNC Model implies there is free bandwidth on link (u_1, u_2). If there is free bandwidth on a local wavelength (for u_1) then we can set up c through it. Then to set up c from the opaque node u_2 along the rest of the route of c to up is the same as for the case $q = 0$. Thus, we can set up c along the rest of the route. If there is no free bandwidth on the $t(u_1)$ local wavelengths on link (u_1, u_2) then because of the node capacity constraint of the LNC Model there must be transit traffic on the local wavelengths. In addition, there must be a transit wavelength with free bandwidth. Then we can BR an existing transit connection as shown in Figure 2.6 from the local to transit wavelengths. Then we can set up c.

The last case is $q = 2$. It is similar to case 1 except that you may have to BR at both end nodes. ∎

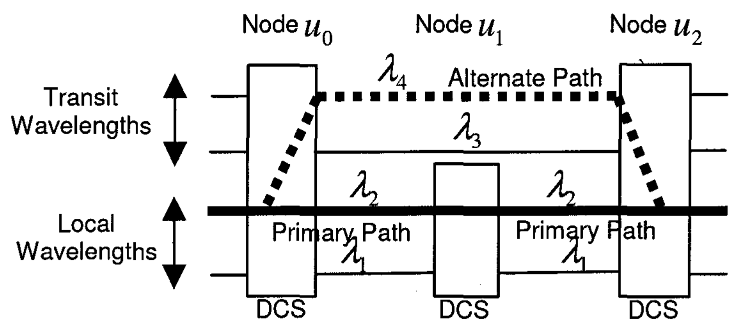

Figure 2.6: Bridge and roll to free bandwidth on a local wavelength for node u_1.

Let T be a collection of tributary connections that arrive and depart at arbitrary times but satisfy the LNC constraint. Let M_T denote the average number of BRs per new tributary connection for T. Let M_{max} denote the maximum (more precisely supremum) of M_T over all possible traffic sets T.

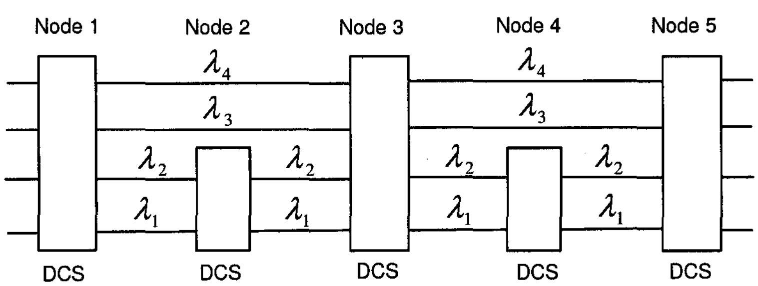

Figure 2.7: A section of the ring of 5 nodes.

Proposition 2.2.2 *For a simple WDM ring network with at least* 6 *nodes,* M_{max} *is at least* $2/3$.

Proof. Since the network has at least six nodes, it has a section of 5 nodes as shown in Figure 2.7. We will describe a traffic that achieves the lower bound of M_{max}. First we have an arrival of tributary connections for node pair $\{1, 3\}$ and node pair $\{3, 5\}$ until the wavelengths are filled.

The resulting traffic configuration will be called the "full state". Then there is the following set of arrivals and deletions of connections. First, a connection is deleted between the pair $\{1, 3\}$ at λ_3, and another connection is deleted between node pair $\{3, 5\}$ at λ_3. Then a connection arrives between node pair $\{2, 4\}$. This causes two BRs to make way for the new connection. One of the BRs moves a connection between node pair $\{1, 3\}$ on local wavelength λ_1 to wavelength λ_3, and the other BR moves a connection between node pair $\{3, 5\}$ on local wavelength λ_1 to wavelength λ_3. Next the connection that just arrived is deleted. Then two connections arrive, one between $\{1, 3\}$ and the other between $\{3, 5\}$. This leaves the wavelengths completely filled, and we are now back to the "full state". The set of arrivals and deletions are then repeated forever.

Notice that each set of arrivals and deletions has 3 arrivals and 2 BRs. Thus the traffic set has on average 2 BRs for every 3 arrivals. ∎

We now compare costs with the fully opaque ring. First we consider bandwidth cost. Notice that a fully opaque ring network requires L wavelengths to support traffic under the LNC Traffic Model. Since the simple WDM ring uses the same number of wavelengths and does not block any new connection, it is as bandwidth efficient as a fully opaque ring.

Next, we consider the traffic rearrangement cost. Proposition 2.2.1 implies the simple WDM ring has at most 2 BRs per tributary connection, and Proposition 2.2.2 implies that on average the number of BRs per tributary connection can be at least $2/3$. On the other hand, there are no traffic rearrangement costs for the fully opaque ring.

Finally, we consider transmission cost. To determine the number of LTEs for the ring networks, we needed to choose values for L and $t()$ of the LNC traffic model. Our values for these parameters were chosen under the assumption of a uniform traffic matrix and that tributary connections follow shortest hop routes. In particular, each node pair was assumed to have the same number y of tributary connections. To determine the value of L, we first compute the average number of tributary connections over a link. This is $\frac{1}{2}(m-1)yh$, where h is the average number of hops per connection. Thus, $L = \lceil\frac{(m-1)yh}{2g}\rceil$. For each node k, the number of connections terminating at the node is $(m-1)y$, where half come through one incident link and the other half come through the other incident link. Thus, we choose $t(k) = \tau$, where $\tau = \lceil \frac{my}{2g} \rceil$. We shall refer to these parameters for the LNC traffic model as the uniform traffic parameters. We consider two values for the number of nodes, $m = 8$ and 16. In addition, we are interested in values of y that result in the number of wavelengths being 32 and 64, i.e., $L = 32$ or 64. For $m = 8$, this corresponds to $\tau = 16$ and 32, respectively. For $m = 16$, this corresponds to $\tau = 8$ and 16, respectively.

The average number of LTEs per node for a fully opaque ring is $2L$. The average number of LTEs per node for the simple WDM ring is $L + \tau$. The following table has the average number of LTEs per node for $m = 8$ and 16 and $L = 32$ and 64. Note that for $m = 8$, the simple optical ring has 25 % less LTEs than fully opaque. For $m = 16$, simple optical ring has 37.5 % less LTEs than fully opaque.

Table 2.1: Comparison of LTEs required in WDM opaque and WDM simple rings.

Parameters	Avg No.LTEs/node fully opaque ring	Avg No. LTEs/node simple WDM ring
m = 8, L = 32	64	48
m = 8, L = 64	128	96
m = 16, L = 32	64	40
m = 16, L = 64	128	80

2.3 A Hierarchical WDM Ring

We consider a generalization of the simple WDM ring that has less LTEs in some cases, but at the expense of increasing the number of BRs. We will refer to this as the hierarchical WDM ring.

Before describing the hierarchical WDM ring, we will revisit the simple WDM ring. The simple ring is composed of sections, where the boundaries of the sections are the opaque nodes. A

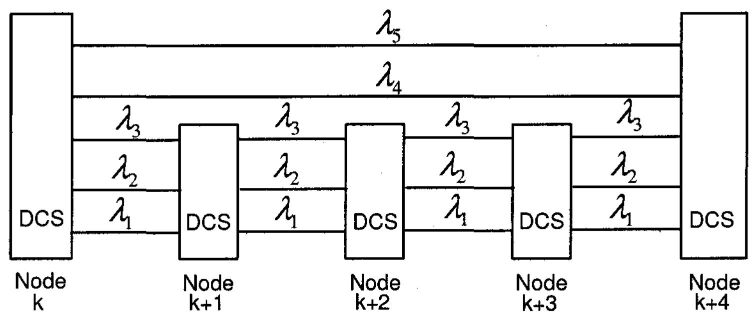

Figure 2.8: A section of a ring.

section consists of an optical node between two opaque nodes. We can generalize a section so that it is composed of any number of consecutive optical nodes between two opaque nodes as shown in Figure 2.8. The figure shows a section with three optical nodes $k+1, k+2$, and $k+3$ between two opaque nodes k and $k+4$. The optical nodes are also referred to as intermediate nodes. For the section, the local wavelengths are defined to be $\lambda_1, \lambda_2, ..., \lambda_r$, where $r = \sum_{j \in I} t(j)$ and I is the set of intermediate nodes. In the Figure 2.8 $I = \{k+1, k+2, k+3\}$, $t(k+1) = t(k+2) = t(k+3) = 1$, and thus, $r = 3$. The other wavelengths are the transit wavelengths. At the optical nodes, there are LTEs at local wavelengths, but not at transit wavelengths. All tributary connections that terminate at an intermediate node is referred to as local for the section, and all other connections are referred to as transit.

We have the following properties for a section. Suppose we are to set up a new tributary connection c through the section. Note that due to the LNC traffic model, there must be free bandwidth along the links. There are two cases to consider.

- **Case 1.** The connection c goes through the section, from opaque node to opaque node. Since the connection traverses the section, there must be free bandwidth on all links of the section. If there is free bandwidth on some transit wavelength then the connection can use the wavelength. If not, then there must be free bandwidth on local wavelengths on all links. The free bandwidth on the local wavelengths can be cross-connected to support the connection c.

- **Case 2.** The connection c terminates at one or more of the intermediate optical nodes. We consider two subcases. The first subcase is when there is free bandwidth on local wavelengths along the route of connection c. Then the connection can be set up on this free bandwidth. The other subcase is when some link e along the route does not have free bandwidth on a local wavelength. There are two facts to consider.
  - The link e must have free bandwidth on a transit wavelength λ_j. Since the wavelength is transit (i.e., no cross-connection at the intermediate optical nodes), there must be free bandwidth on this wavelength across the section.
  - Note that the number of local wavelengths on e is r. Also note that this is sufficient to carry the maximum number of local connections on e. Since e has no free bandwidth on a local wavelength, it must have at least one transit connection c' on a local wavelength. The connection c' can be BR'd to the transit wavelength λ_j. This will free up bandwidth at the local wavelengths for c, and then c can be set up.

We just described how to set up a connection through individual sections. Essentially, this describes how to set up a connection through the network because the sections are connected by opaque nodes, that can cross-connect bandwidth between any pair of wavelengths. The two cases imply the following proposition.

Proposition 2.3.1 *Consider the simple WDM ring but with the generalization of sections. Assume traffic model LNC. Then setting up a tributary connection c requires at most q BRs, where q is the number sections that has an end node of c as an intermediate (optical) node. Since q can be at most 2, the number of BRs for c is at most 2.*

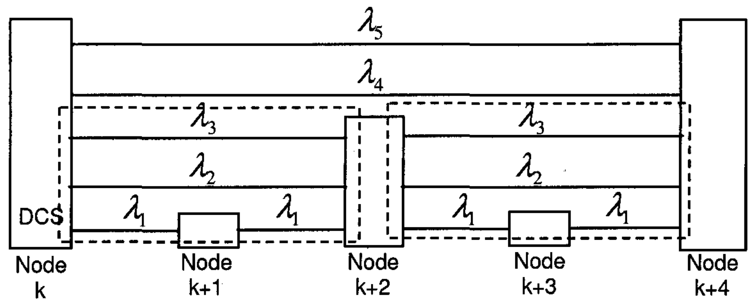

Figure 2.9: A level-1 section with two level-2 sections.

Consider the section of Figure 2.8 again. Ignore the transit wavelengths and only consider the r local wavelengths. They support all traffic under the LNC Model except that the link capacity constraint is now rg. Notice that the intermediate nodes are all opaque. Thus, we can reduce the LTEs by replacing some of the opaque nodes with optical nodes just as we did for the simple WDM ring. This is shown in Figure 2.9. Notice that the section is made up of two subsections. The subsections have a border node that behaves like an optical node if we ignore the transit wavelengths. In the figure, the border node is $k+2$. Thus, the border node has LTEs at all the r local wavelengths for the section.

We will refer to a section as a level-1 section, and a subsection as a level-2 section. We will refer to the ring network with these sections as a hierarchical WDM ring.

Proposition 2.3.2 *Consider the hierarchical WDM ring. Assume the traffic model is LNC. Then setting up a tributary connection c requires at most q BRs, where q is the number of level-1 and level-2 sections that have an end node of c as an intermediate node. Since q is at most* 4, *the number of BRs to accommodate the tributary connection is at most* 4.

Proof. Earlier in this section, we described the properties of a section in hierarchical ring network. From Case 1 of the properties, a BR will not occur in a level-1 or level-2 section if the connection c is passing through it. From Case 2 of the properties, at most one BR occurs for a level-1 section if the connection c terminates in one of its intermediate nodes. In addition, if the connection also terminates in a level-2 section then at most one additional BR occurs for the level-2 section. Thus, the proposition is true. ∎

We will now compare LTE cost for the fully opaque ring, simple WDM ring, and hierarchical WDM ring when the number of nodes $m = 16$. For this case, the hierarchical WDM ring was divided into four level-1 sections, each with 3 intermediate nodes. Each level-1 section was composed of two level-2 sections, each with 1 intermediate node. Thus, a level-1 section was structured like Figure 2.9.

We will use the uniform traffic parameter values for LNC traffic. In particular, the parameters are $L = 32$ $(\tau = 8)$ and 64 $(\tau = 16)$. The following table shows the average number of LTEs per node. Simple WDM rings have 37.5 % less LTEs than fully opaque. Hierarchical WDM rings have 44 % less LTEs than fully opaque.

Table 2.2: Comparison of LTEs required in WDM opaque, simple and hierarchical rings.

Parameters	Avg No. LTEs/node fully opaque ring	Avg No. LTEs/node simple WDM ring	Avg No. LTEs/node hierarchical WDM ring
L = 32	64	40	36
L = 64	128	80	72

Chapter 3

Grooming in General WDM Mesh Networks

In this chapter we consider the problem of grooming in general WDM mesh networks under a static traffic model. As stated in Chapter 1, traffic grooming involves packing the low speed connection requests onto the lighpaths efficiently and assigning the lightpaths to the wavelengths in the physical network. Here, we only consider the grooming of traffic in the logical topology. It means that we are not concerned about the assignment of lightpaths to wavelengths in the physical topology. We assume that all logical topologies are implementable on a physical topology. Although this assumption is restrictive, it is satisfied in the following situation. In wide-area WDM networks, if the number of physical links (fibers) available between the adjacent nodes is large, then the number of wavelengths in the system is not important and two lightpaths can always be routed on two different physical links even though they use the same wavelength [9]. The cost of the logical topology is dominated by the number of transceivers, which is twice the total number of lightpaths. Thus our goal is to minimize the total number of lightpaths in the logical topology, which can completely route the given traffic demands. Throughout this chapter we assume that every traffic demand is some integer multiple of an unit demand, where the unit demand is a fraction of the capacity of the lightpath. This exact problem was considered in [9] and [2]. In Section 4.1 the mathematical formulation of the problem is presented. In Section 4.2 the lower and upper bounds for the optimal number of lightpaths is presented. In Section 4.3 a brief overview of the known heuristics for this problem is given and towards the end of this section a novel greedy heuristic is presented.

3.1 Problem Statement

The demand traffic is represented by matrix $T = [t_{ij}]$ where t_{ij} denotes the sum of traffic streams (each of unit demand) from node i to node j and from node j to node i. Without loss of generality, we can assume that $t_{ij} = 0 \ \forall \ i \leq j$. Let c (an integer multiple of unit demand) be the capacity of each lightpath. Further, we can assume that $t_{ij} < c$. If k $(k > c)$ low speed connections are to be set up between nodes i and j, then one could set up a lightpath between the nodes i and j and route c of these connections on this lightpath and then consider the problem with $k - c$ connections between nodes i and j. Let N be the number of nodes. The logical topology is represented by the lightpath matrix $L = [l_{ij}]$, where l_{ij} denotes the number of lightpaths between node i and node j. Let f_{ijkl} denote the number of traffic streams between node k and node l that are routed through the lightpath between node i and node j. Note that $f_{ijkl} \geq 0 \ \forall \ i, \ j, \ k,$ and l. The goal is to minimize $\sum_{i,j} l_{ij}$ subject to the following constraints.

1. **Capacity constraint.** This constraint implies that the number of traffic streams routed through a link in the logical topology cannot exceed the total capacity of the lightpaths that constitute the link.

$$\sum_{k,l} f_{ijkl} \leq c\, l_{ij} \ \forall \ i, \ j \, .$$

2. **Flow Conservation constraint.** This constraint implies that the traffic flowing into a node is equal to the sum of traffic flowing out of the node and the traffic dropped at that node. Thus, for each $i\,, j\,, k\,, l\,,$

$$\sum_i f_{ijkl} - \sum_i f_{jikl} = \begin{cases} 0 & \text{if } j \neq k \text{ } and \text{ } j \neq l \\ t_{ij} & \text{if } j = k \\ -t_{ij} & \text{if } j = l. \end{cases}$$

The lightpath minimization problem could be interpreted as follows. Consider a $N \times N$ complete graph with the capacity of each edge $E_{ij} \gg c$, where c is the capacity of lightpath. Consider each traffic demand t_{ij} in T as a different commodity that has to be routed from node i to node j in this complete graph. Let the cost of routing through the edge E_{ij} be a step function denoted by $\lceil \frac{x}{c} \rceil$, where x is number of unit traffic demands routed through edge E_{ij}. Lighpath minimization problem is equivalent to finding the minimum cost multicommodity network flow in this complete graph. In general graphs, the minimum cost multicommodity flow problem is NP-Complete for both linear and non-linear cost functions. However, note that in the above interpretation the graph is complete. We do not know whether the lightpath minimization problem (or minimum cost multicommodity flow problem in complete graphs with non-linear cost function) is NP-Complete.

3.2 Bounds

In this section we derive bounds on the number of lightpaths under the static traffic model. For lower bounds, we first obtain a lower bound on the total path length required to support all the traffic demands on the lightpath topology and then find the corresponding lower bound on the number of lightpaths. The path length of a unit traffic demand is the number of lightpaths that it is routed through.

3.2.1 Lower Bounds

The path length of every traffic demand is at least one. Thus, the total path length required to support the given traffic on any logical topology is

$$\sum_{i=1}^{N}\sum_{j=1}^{N} t_{ij} \,.$$

The minimum number of lightpaths required to support this total path length is

$$LB_1 = \left\lceil \frac{\sum_{i=1}^{N}\sum_{j=1}^{N} t_{ij}}{c} \right\rceil . \tag{3.2.1}$$

The lower bound in Equation (3.2.1) [2] can be improved by observing the fact that no lightpath carries c traffic streams whose path lengths are one. Note that $t_{ij} < c \forall\, i,\, j$. This implies that a lightpath between nodes i and j could carry at most $c - 1$ traffic streams between nodes i and j. The c^{th} connection on this lightpath could either be empty or be used to route a traffic demand whose path length is at least two. Thus,

$$\sum_{i=1}^{N}\sum_{j=1}^{N} t_{ij} + \frac{L}{2} \leq Lc \,,$$

where L is the number of lightpaths.

$$LB_2 = \left\lceil \frac{\sum_{i=1}^{N}\sum_{j=1}^{N} t_{ij}}{c - \frac{1}{2}} \right\rceil . \tag{3.2.2}$$

Next, we generalize a lower bound found in [5], which only considered uniform traffic. If there are L lightpaths in the logical topology, at most the sum of first L largest traffic demands can have a path length of one. The remaining traffic demands must have path lengths greater than or equal to 2. Let t^k denote the k^{th} largest integer in the matrix $[t_{ij}]$. Then the total path length that needs to be supported by any logical topology is

$$2\sum_{i=1}^{N}\sum_{j=1}^{N} t_{ij} - \sum_{k=1}^{L} t^k \,.$$

By replacing $\sum_{k=1}^{L} t^k$ with Lt^1 we obtain the following lower bound.

$$
\begin{aligned}
2\sum_{i=1}^{N}\sum_{j=1}^{N} t_{ij} - Lt^1 &\leq \sum_{i=1}^{N}\sum_{j=1}^{N} t_{ij} - \sum_{k=1}^{L} t^k \\
&\leq Lc\,.
\end{aligned}
$$

Thus,

$$
LB_3 = \left\lceil \frac{2\sum_{i=1}^{N}\sum_{j=1}^{N} t_{ij}}{c+t^1} \right\rceil . \tag{3.2.3}
$$

For the uniform traffic case with $t_{ij} = t \;\forall\; i,\, j$ and $i \neq j$, the lower bound in Equation (3.2.3) becomes $\left\lceil \frac{N(N-1)}{c+t} \right\rceil$ [5].

We could obtain a different lower bound from the following observation. The number of lightpaths incident at a node i has to be at least $\left\lceil \frac{\sum_{j=1}^{N}(t_{ij}+t_{ji})}{c} \right\rceil$ in any logical topology that completely routes the given traffic demands. Let L denote the number of lightpaths and let d_i denote the lightpaths incident at node i. Thus, the number of lightpaths in the logical topology is $L = \frac{1}{2}\sum_{i=1}^{N} d_i$. If some lightpath topology completely routes the given traffic demands, then we have the following inequalities.

$$
\begin{aligned}
L &= \frac{1}{2}\sum_{i=1}^{N} d_i \\
&\geq \frac{1}{2}\left\lceil \frac{\sum_{j=1}^{N}(t_{ij}+t_{ji})}{c} \right\rceil .
\end{aligned}
$$

Thus,

$$
LB_4 = \frac{1}{2}\sum_{i=1}^{N}\left\lceil \frac{\sum_{j=1}^{N}(t_{ij}+t_{ji})}{c} \right\rceil . \tag{3.2.4}
$$

This lower bound is particularly interesting because we use this lower bound to show that the lightpaths used in any star logical topology is never more than twice the optimum number of lightpaths required.

3.2.2 Upper Bounds

If we setup a lightpath between any two nodes i, j with $t_{ij} > 0$, then the number of lightpaths in this logical topology is

$$UB_{trivial} = \sum_{i=1}^{N}\sum_{j=1}^{N}\left\lceil \frac{t_{ij}}{c} \right\rceil . \qquad (3.2.5)$$

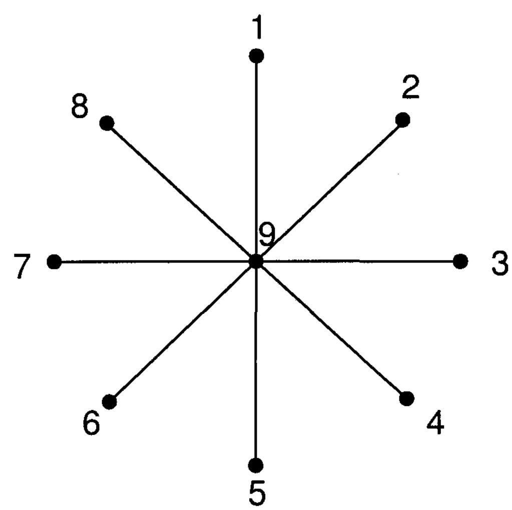

Figure 3.1: Logical star topology.

Consider the logical star topology as shown in Figure 3.1. In this topology, the lightpaths can be setup only between the hub (central node) and the end nodes. Thus all the switching is done only at the hub. The number of lightpaths that need to be setup between an end node i and the hub is $\left\lceil \frac{\sum_{j=1}^{N}(t_{ij}+t_{ji})}{c} \right\rceil$. Let us assume that node h is the hub. Then the total number of lightpaths in this star logical topology is

$$\begin{aligned} UB_{star} &= \sum_{i=1}^{N}\left\lceil \frac{\sum_{j=1}^{N}(t_{ij}+t_{ji})}{c} \right\rceil - \left\lceil \frac{\sum_{j=1}^{N}(t_{hj}+t_{jh})}{c} \right\rceil \\ &= \sum_{i=1,\, i\neq h}^{N}\left\lceil \frac{\sum_{j=1}^{N}(t_{ij}+t_{ji})}{c} \right\rceil . \end{aligned} \qquad (3.2.6)$$

Theorem 3.2.1 *The total number of lightpaths needed in any star logical topology is at most twice the optimum number of lightpaths.*

Proof. It is clear that the ratio of Equation (3.2.6) and Equation (3.2.5) is at most 2. Since the optimum value is greater than the lower bound, the ratio of the number of lightpaths in any star logical topology Equation (3.2.6) and the optimum number of lightpaths is at most 2. ∎

3.3 Greedy Heuristic Algorithm

The mathematical formulation of lightpath minimization problem in Section 3.1 indicates that it is an *Integer Linear Programming* problem with $O(N^4)$ variables. This ILP becomes unmanageable for reasonably large networks. Thus, we need to turn to good heuristics to solve this problem. A good heuristic is characterized by its algorithmic complexity and performance. The heuristic proposed in [2] starts with an empty topology and successively adds lightpaths to route the remaining traffic streams. If a traffic stream cannot be routed between a pair of nodes, then a new lightpath is setup between those nodes. However the choice of the new traffic stream is random. The heuristic presented in [9] starts with an initial topology and successively deletes lightpaths after rerouting the traffic streams the lightpaths carry. In the remainder of this section, we present a novel heuristic that attempts to minimize the lightpaths.

Let the lightpath graph after k iterations be denoted by $D^k = [l_{ij}c]$, where l_{ij} is the number of lightpaths between nodes i and j and c is the capacity of each lightpath. Let R^k denote the unused capacity that remains in D^k. We refer to R^k as the residual graph after the k^{th} iteration. The traffic demand matrix is denoted by $T = [t_{ij}]$, where t_{ij} is the number of traffic streams, each of unit demand that have to be set up between nodes i and j. Let t^k_{ij} denote the number of streams that are yet to be set up between nodes i and j after the k^{th} iteration. Thus, T^k denotes the traffic matrix that is not yet routed after the k^{th} iteration. This implies that $t_{ij} - t^k_{ij}$ traffic streams between nodes

i and j are already routed on the current lightpath graph D^k. The heuristic algorithm is described below.

- **Step 1.** Initially the index of iteration, $k = 0$ and thus $T^k = T$, $R^k = \phi$ and $D^k = \phi$.
- **Step 2.** If $t^k_{ij} > 0$ for some i and j, add a lightpath between some pair of nodes (p, q) in the residual graph, which can route the maximum number of traffic streams from T^k.
- **Step 3.** Update D^{k+1}, T^{k+1} and R^{k+1} and go to Step 2.

The above heuristic begins with an empty lightpath topology. It adds lightpaths successively to the residual topology, so that with each successive addition it can route the maximum number of traffic streams that are not yet routed. We use certain properties of the residual graph to show that this could be solved as a multitude of single commodity network flows.

Definition Consider a residual graph R^k and a traffic demand matrix $T^k = [t^k_{ij}]$. Suppose for each $t^k_{ij} > 0$, the nodes i and j are disconnected in the graph R^k. We call the graph R^k disconnected *with respect to* T^k.

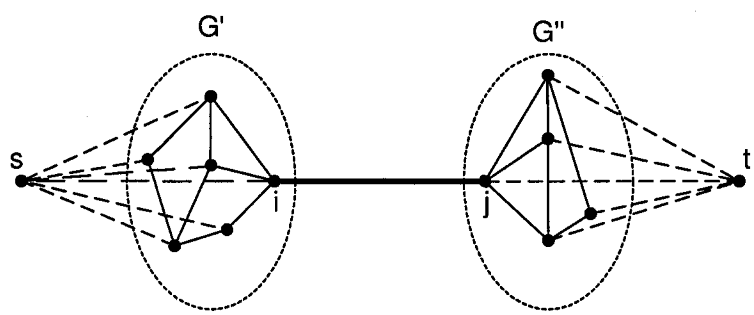

Figure 3.2: A sample residual network.

Proposition 3.3.1 *Let the residual graph R^k be disconnected with respect to some demand matrix T^k. If an edge of some capacity is added to the graph between a pair of nodes, then the maximum* *number of demands that can be routed on R^k could be solved as a maximum single commodity network flow. Further, if the maximum number of traffic streams are routed on R^k, then R^{k+1} is disconnected with respect to T^{k+1}.*

Proof. Since R^k is disconnected *with respect to* T^k, if two nodes x and y are connected in R^k, then $t^k_{xy} = 0$. This follows from the Definition 3.3. Let the end nodes of the new edge be i and j. Consider the connected subgraph G' that contains node i as shown in Figure 3.2. Since the subgraph G' is connected, $t^k_{xy} = 0 \,\forall\, x$ and $y \in G'$. Create a source node s and directly connect it to all the nodes in the subgraph G' with new edges. Let the capacity of edge E_{sl}, where $l \in G'$, be $\sum_{m \in G''} t_{lm}$. Consider the connected subgraph G'' that contains node j. Note that $t^k_{xy} = 0 \,\forall\, x$ and $y \in G''$. Create a sink node t and directly connect it to all the nodes in G'' with new edges. Let the capacity of edge E_{mt}, where $m \in G''$, be $\sum_{l \in G'} t_{lm}$. Note that every demand that is routed on R^k has to be routed through the new edge $(i,\, j)$. Thus, the maximum number of demands from T^k that can be routed on R^k is the maximum flow from s to t.

We will now prove by contradiction that R^{k+1} is disconnected *with respect to* T^{k+1}. Suppose there is a pair of nodes l and m in R^{k+1}, where $t^{k+1}_{lm} > 0$ and l is connected to m. Since l was not connected to m in R^k, l must be connected to m in R^{k+1} through the added edge $(i,\, j)$. This implies that $l \in G'$, $m \in G''$, and there is unused capacity in R^{k+1} from l to m through the edge $(i,\, j)$. This contradicts the fact that the amount of traffic routed on R^k through the edge $(i,\, j)$ is maximized. Therefore, $t^{k+1}_{lm} > 0$ implies that l and m are disconnected in $R^{k+1} \,\forall\, l$ and m. ∎

Since R^k is disconnected *with respect to* $T^k \,\forall\, k$, it is clear from Proposition 3.3.1 that the Step 2 of the heuristic involves enumerating all pairs of nodes (i, j) and solving a single commodity network flow per pair. We choose a lightpath between nodes i and j that maximizes this flow. We need to solve $O(N^2)$ single commodity network flows in the Step 2 of the heuristic. The number of iterations performed by the algorithm depends on the traffic demand matrix, T. A trivial upper bound on the number of iterations is $\sum_{i=1}^{N} \sum_{j=1}^{N} t_{ij}$. Note that the routes of the traffic streams that are routed at an earlier stage in the heuristic do not use the lightpaths that are added later. Hence, the heuristic may not give an optimal solution. However, we currently do not know how this heuristic performs when compared to the optimal solution. To gain an insight into its performance we need to do computational experiments.

Chapter 4

Conclusion

In this thesis, we considered the traffic grooming in both WDM ring and mesh networks. In the second chapter, we considered WDM ring networks called the simple WDM ring and the hierarchical WDM ring. They are as bandwidth efficient as fully opaque rings for the LNC Traffic Model. They were shown to require significantly less number of LTEs for a particular collection of LNC Traffic Model parameters. We believe the parameters are reasonable since they reflect a uniform traffic pattern. For the case of large rings with 16 nodes, the WDM rings can have 44 % less LTEs than a fully opaque network. The drawback of these ring networks is that they require traffic rearrangement.

We gave an example that showed that rearrangement is necessary because the bandwidth efficiency is high. For the hierarchical WDM rings with 16 nodes, the reduction in LTEs can be 44% but the maximum number of BRs per new tributary connection is 4. For simple WDM rings with 16 nodes, the reduction in LTEs can be 37.5 % but the maximum number of BRs per new tributary connection is 2. Though the hierarchical WDM ring provides better LTE reduction, the simple WDM ring may be the best compromise between reducing LTEs and minimizing BRs.

In the third chapter, we considered traffic grooming in WDM mesh networks. We provided the mathematical formulation of this problem. The objective is to minimize the lightpaths required to support the given traffic demands. A demand could be a fraction of the lightpath capacity. Then, we presented bounds on the lightpaths required to support the given traffic demands. We showed that any star logical lightpath topology never uses more than twice the optimum number of lightpaths required. We presented a novel greedy heuristic algorithm. However we do not know how this heuristic performs when compared to the optimal solution.

A lot of research has been carried out in the area of traffic grooming in WDM networks. However most of the work has concentrated on the ring networks. There has not been much work on traffic grooming in general topologies. It is worthy to develop algorithms that can be flexibly and efficiently applied within a variety of optical network and cost models.

Bibliography

[1] D. Banerjee and B. Mukhrjee. A pratical appproach for routing and wavelength assignment in large wavelength-routed optical networks. *IEEE Journal on Selected Areas in Communications*, 14(5):909–913, June 1996.

[2] Mauro Brunato and Roberto Battiti. A multistart randomized greedy algorithm for traffic grooming on mesh logical topologies. *Proceedings of ONDM*, 2002.

[3] Angela L. Chiu and Eytan H. Modiano. Traffic grooming algorithms for reducing electronic multiplexing costs in wdm ring networks. *Journal of Lightwave Technology*, 18(1), January 2000.

[4] Rudra Dutta and George N. Rouskas. Traffic grooming in wdm networks: Past and future. *IEEE Network*, 16(6):46–56, November/December 2002.

[5] O. Gerstel, P. Lin, and G. Sasaki. Combined wdm and sonet network design. *Proc. IEEE Infocom*, 1999.

[6] O. Gerstel, R. Ramaswami, and G. Sasaki. Cost effective traffic grooming in wdm rings. *IEEE/ACM Trans*, 8(5):608–617, October 2000.

[7] H. Ghafouri-Shiraz, G. Zhu, and F. Yuan. Effective wavelength assignment algorithms for optimizing design costs in sonet/wdm rings. *IEEE Journal of Lightwave Technology*, 19(9), October 2001.

[8] E. Hyytia and J. Virtamo. Wavelength assignment and routing in wdm networks. *Nordic Teletraffic Seminar (NTS)*, 1998.

[9] V. R. Konda and T. Y. Chow. Algorithm for traffic grooming in optical networks to minimize the number of transceivers. *IEEE Workshop on High Performance Switching and Routing*, pages 218–221, 2001.

[10] K. C. Lee and V. O. K. Li. A wavelength rerouting algorithm in wide area all optical networks. *IEEE/OSA J. Lightwave Technology (special issue on Multiwavelength Technology and Networks)*, 14(6):1218–1228, June 1996.

[11] E. Modiano and A. Narula. Mechanisms for providing optical bypass in wdm-based networks. *Optical Networks Magazine*, January 2000.

[12] R. Ramaswami and K. N. Sivarajan. Design of logical topologies for wavelength-routed optical networks. *IEEE Journal on Selected Areas in Communications*, 14:840–851, June 1996.

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

[14] G. Sasaki and O. Gerstel. Minimal cost wdm sonet rings that guarantee no blocking. *Optical Networks Magazine*, October 2000.

[15] G. Sasaki, O. Gerstel, and R. Ramaswami. A wdm ring network for incremental traffic. *Proc. 36th Annual Allerton Conference on Communication, Control, and Computing, Monticello, IL*, pages 662–672, September 1998.

[16] X. Zhang and C. Qiao. An effective and comprehensive approach to traffic grooming and wavelength assignment in sonet/wdm rings. *IEEE/ACM Transactions on Networking*, 8(5), October 2000.

[17] Keyao Zhu and Biswanath Mukherjee. Traffic grooming in an optical wdm mesh network. *IEEE Journal on Selected Areas in Communications*, 20(1), JANUARY 2002.