ISSUES IN WAVELENGTH ROUTED NETWORKS

Some of the important issues in wavelength routed networks include routing and wavelength assignment; minimizing the effect (bandwidth loss) due to wavelength continuity constraint (some possible solutions for this problem include employing wavelength converters, multifibers, and wavelength rerouting); design, reconfiguration, and survivability of virtual topology (optical layer); optical multicasting; control and management; traffic grooming; and IP-over-WDM. We now briefly examine each of these issues.

1 Routing and Wavelength Assignment

In wavelength routed WDM networks, a connection is realized by a lightpath. In order to establish a connection between a source–destination pair, a wavelength continuous route needs to be found between the node pair. An algorithm used for selecting routes and wavelengths to establish lightpaths is known as a *routing and wavelength assignment* (RWA) algorithm. Many problems in wavelength routed WDM networks have RWA as a subproblem. Therefore, it is mandatory to use a good routing and wavelength assignment algorithm to establish lightpaths in an efficient manner. Wavelength assignment is a unique feature in wavelength routed networks that distinguishes them from conventional networks.

Static Versus Dynamic Traffic Demand

The connection requests (traffic demand) can be either static or dynamic. In case of a static traffic demand, connection requests are known *a priori*. The traffic demand may be specified in terms of source–destination pairs. These pairs are chosen based on an estimation of long-term traffic requirements between the node pairs. The objective is to assign routes and wavelengths to all the demands so as to minimize the number of wavelengths used. The dual problem is to assign routes and wavelengths so as to maximize the number of demands satisfied, for a fixed number of wavelengths. The above problems are categorized under the *static lightpath establishment* (SLE) problem. The SLE problem has been shown to be NP-complete (that is, it is computationally intractable or, in other words, the only known algorithms that find an optimal solution require exponential time in the worst case). Therefore, polynomial-time algorithms which produce solutions close to the optimal one are preferred.

In case of a dynamic traffic demand, connection requests arrive to and depart from a network one by one in a random manner. The lightpaths once established remain for a finite time. The dynamic traffic demand models several situations in transport networks. It may become necessary to tear down some existing lightpaths and establish new lightpaths in response to changing traffic patterns or network component failures. Unlike the static RWA problem, any solution to the dynamic RWA problem must be computationally simple, as the requests need to be processed online. When a new request arrives, a route and wavelength need to be assigned to the request with the objective of maximizing the number of connection requests honored (equivalent to minimizing the number of connection requests rejected). Dynamic RWA algorithms perform more poorly than static RWA algorithms because a dynamic RWA algorithm has no knowledge about future connection requests, whereas all the connection requests are known *a priori* to a static RWA algorithm. A dynamic RWA algorithm processes the connection requests strictly in the order in which they arrive, whereas a static RWA algorithm processes the requests in an order decided by some heuristic. One such heuristic is to assign wavelengths to the connections in the nonincreasing order of their hop length, as longer-hop connections are less likely to find the same wavelength free on the entire route. The following example substantiates the above claim.

Centralized Versus Distributed Control

RWA algorithms assume either centralized or distributed control for selecting routes and wavelengths. In the case of centralized control, a central controller is assumed to be available. It keeps track of the state of the network. It is responsible for selecting routes and wavelengths for the requests and sending control signals to appropriate

nodes for establishing and releasing lightpaths. In the case of distributed control, no central controller is used. Up-to-date knowledge of the network state is not known to any node. A node may use distributed breadth-first search to select a route and wavelength to honor a connection request. It may also use precomputed routes and search for free wavelengths on the links of the routes. Control messages are sent to various nodes to reserve wavelengths on the links traversed. Once a route is found and wavelengths are reserved, appropriate control signals are sent to various nodes to configure switches in the routing nodes for establishing the lightpath. Similarly, to release a lightpath, control signals are sent to various nodes by the source node. Centralized algorithms are useful for small networks and are not scalable to large networks. For simplicity and scalability purposes, distributed control protocols are desirable.

Another important problem in wavelength routed networks is the *fairness* between connections with different hop counts. Longer-hop connections are less likely to be accepted than shorter-hop connections. The situation becomes worsened when distributed routing is used, due to the increased possibility of wavelength reservation conflicts between simultaneous attempts of several probes. Therefore, an appropriate mechanism is imperative in order to improve fairness among connections with different hop counts, with a minimum penalty in terms of loss in network-wide performance (for example, throughput).

Route and Wavelength Selection Methods

The important routing methods considered in the literature are *fixed routing, alternate routing*, and *exhaust routing*. In the fixed routing method, only one route is provided for a node pair. Usually this route is chosen to be the shortest route. When a connection request arrives for a node pair, the route fixed for that node pair is searched for the availability of a free wavelength. In the alternate routing method, two or more routes are provided for a node pair. These routes are searched one by one in a predetermined order. Usually these routes are ordered in nondecreasing order of their hop length. In the exhaust method, all possible routes are searched for a node pair. The network state is represented as a graph and a shortest-path-finding algorithm is used on the graph. While the exhaust method yields the best performance when compared to the other two methods, it is computationally more complex. Similarly, the fixed routing method is simpler than the alternate routing method, but it yields poorer performance than the other.

Based on the order in which the wavelengths are searched, the wavelength assignment methods are classified into *most-used, least-used, fixed-order,* and *random-order*. In the most-used method, wavelengths are searched in nonincreasing order of their utilization in the network. This method tries to pack the lightpaths so that more wavelength continuous routes are available for the requests that arrive later. In the least-used method, wavelengths are searched in nondecreasing order of their utilization in the network. This method spreads the lightpaths over different wavelengths. The idea here is that a new request can find a shorter route and a free wavelength on it. The argument is that the most-used method may tend to choose a longer route, as it always prefers the most-used wavelength. In the fixed-order method, the wavelengths are searched in a fixed order. The wavelengths may be indexed and the wavelength with the lowest index is examined first. In the random method, the wavelength is chosen randomly from among the free wavelengths. The most-used and least-used methods are preferred for networks with centralized control. The other two methods are preferred for networks with distributed control. The numerical results reported in the literature show that the most-used method performs better than the least-used method and the fixed-order method performs better than the random method.

RWA algorithms may select routes and wavelengths one after the other. Either routes are searched first or wavelengths are searched first. Alternatively, the routes and wavelengths can be considered jointly. For every route–wavelength pair, a cost value can be associated. Such a method is called as a *dynamic method*. In a *least congested path* routing method, a route with the least congestion is preferred. The least congested path is the one with the maximum number of free wavelengths. This method is expected to leave more wavelength continuous routes for the requests that arrive later.

2 Wavelength-Convertible Networks

One possible way to overcome the bandwidth loss caused by the wavelength continuity constraint is to use wavelength converters at the routing nodes. A wavelength converter is an optical device which is capable of shifting one wavelength to another wavelength. The capability of a wavelength converter is characterized by the degree of conversion. A converter which is capable of shifting a wavelength to any one of D wavelengths is said to have conversion degree D. The cost of a converter grows with increasing conversion degree. A converter is said to have full degree of conversion when the conversion degree equals the number of wavelengths per fiber link. Otherwise, it is said to have partial or limited degree of conversion.

A WXC having one or more wavelength converters is called as a *wavelength interchange crossconnect* (WIXC). A node with wavelength conversion capability is called a wavelength converting (WC) node or a wavelength interchange (WI) node. A WDM network with WC nodes is called a wavelength-convertible network. A node may have a maximum of *Fin x W* converters, where *Fin* is the number of incoming fibers at the node. When every node in a network has a sufficient number of full-degree converters, its performance reaches the best achievable. However, such a network is economically not feasible, as the converters are very expensive. This leads to a number of issues: How many nodes in a network should have conversion capability? How do we choose the converting nodes? How many converters can a node have? How do factors such as traffic demand and network topology affect selection of converting nodes and allocation of converters?

A wavelength-convertible network (a network with WIXCs) performs better than a wavelength-selective network (a network without WIXCs). Wavelength converters relax the continuity constraint at a node. Therefore, they help to reduce the bandwidth (wavelength) loss, resulting in better bandwidth utilization. This advantage is demonstrated in the following example.

3 Multifiber Networks

Since only a few wavelengths are economically and technologically feasible and traffic demand is very high, multifiber networks have been receiving much attention. In a multifiber network, a link between two nodes consists of a bundle of (sometimes more than 100) fibers. Using a bundle of fibers for a link is economical, as laying fiber in the ground is done once (The cost to deploy a new or additional fiber cable has been estimated to be about $70,000 per mile!). At any instant of time, depending upon the current need, only a few fibers are used ("lit"). As and when new demand arrives, additional fibers can be used, if required.

In a multifiber wavelength-selective network, a link with *F* fibers each with *W* wavelengths can carry *F* messages on the same wavelength, each on a different fiber. Therefore, the chance of finding a wavelength-continuous route for a request is higher in multifiber networks than in single-fiber networks. As a result, improved performance is achieved by multifiber networks. This means that a multifiber network with *F* fibers per link and *W* wavelengths per fiber performs better than a single-fiber network with *FW* wavelengths per fiber. The reason for the better performance can be understood from the equivalence between a multifiber network and a wavelength-convertible network.

4 Wavelength Rerouting

Apart from wavelength conversion and space division multiplexing, there is yet another way, called *wavelength rerouting*, to reduce the bandwidth loss caused by the wavelength continuity constraint in wavelength routed networks. With wavelength converters employed in a network, a lightpath need to be wavelength-continuous between two consecutive converting nodes only. With space division multiplexing, the chance of finding a wavelength-continuous route is enhanced, as the same wavelength is available on every fiber on a link. Wavelength rerouting creates a wavelength-continuous route by migrating a few existing lightpaths to new wavelengths without changing their route. However, it incurs control overhead and, more important, the service in the rerouted lightpaths needs to be disrupted. Therefore, it is imperative for any algorithm employing wavelength rerouting to migrate as few lightpaths as possible. The usefulness of wavelength rerouting is illustrated in the following example.

The RWA together with the use of wavelength converters, space division multiplexing, and wavelength rerouting is an important and also challenging problem in wavelength routed WDM networks.

5 Virtual Topology Design

Virtual topology (optical layer) in a transport network consists of a set of lightpaths established between a subset of node pairs in the network. The lightpaths can be chosen based on the traffic demand between node pairs. The pattern of connectivity between lightpaths forms a virtual topology or a logical topology. Thus, in a virtual topology, a node corresponds to a routing node in the network and an edge corresponds to a lightpath. If two nodes are connected by a lightpath, then they can communicate in one (light) hop. Due to technological limitations on the number of available wavelengths and the number of available optical transmitters and receivers, it may not be possible to set up lightpaths between all node pairs. If two nodes are not connected directly by a lightpath but are connected by a sequence of lightpaths, the nodes can communicate through the sequence of lightpaths. This type of communication is termed multi-(light)hop communication. In this case, message forwarding between two consecutive lightpaths is performed via electronic processing. For example, two IP routers connected by a lightpath become neighbors in the virtual topology regardless of whether or not they are connected directly by a fiber link in the physical topology. The IP traffic between non-neighbors needs to be processed electronically at every intermediate router.

The traffic between node pairs is routed over the virtual topology in one or more hops. The virtual topology is designed to carry the traffic in such a way as to optimize a certain performance metric, such as the average message delay or network congestion. The message delay can be measured in terms of the number of lightpaths traversed by it. Network congestion is defined as the maximum load on any lightpath. The load on a lightpath can be measured in terms of the amount of traffic carried by it. It can be noted that a lightpath may need to carry traffic that flows between different node pairs.

6 Virtual Topology Reconfiguration

The virtual topology is usually designed based on the estimated average traffic flow between the node pairs in a specific time frame. The length of this time frame depends on whether the planning is long-term or short-term. The traffic flow between the nodes is not constant and is subject to change with time. The underlying virtual topology may not be optimum for all the different patterns of traffic flow. Reconfiguring the virtual topology to be in tune with the changing traffic pattern would help maximize the performance. Reconfiguration requires that a few lightpaths in the existing virtual topology be removed and a few lightpaths be added to form a new virtual topology. The flexibility in wavelength crossconnects to dynamically change the switching patterns of wavelengths from the incoming fibers to the outgoing fibers aids the process of reconfiguration.

There are a number of issues concerning the virtual topology reconfiguration. Migrating one topology to another topology not only incurs control overhead but also introduces service disruption, which is very expensive. It is desirable that the new topology be as close to the current topology as possible. At the same time, the performance metric needs to be optimized for the new topology. There are different approaches to address this problem. One approach is to design an optimal (from the perspective of the performance metric) topology for the new traffic demand and determine a sequence of steps to migrate the current topology to the new topology in such a way as to minimize the service disruption at each step. This approach ignores the current topology information while designing the new optimal topology. Therefore it may choose a new topology which is not close enough to the current topology. Another approach is to choose an optimal topology that requires minimum changes to the current topology. This approach, however, requires complex mathematical models and algorithms. The third approach is to strike a trade-off between the performance metric and the number of changes. Computationally simpler algorithms can be used to choose a new topology that requires only a few changes while keeping the deviation of the performance metric from the optimum one within an acceptable level. The process of reconfiguration is illustrated in the following example.

Now, assume that the traffic demand between nodes changes from T_1 to T_2 as depicted in Fig. 1.18(c). If the same virtual topology V_1 is used to carry the new

7 Survivable Networks

Another important issue in WDM networks is how network component failures are dealt with. Since a huge amount of traffic is carried in WDM networks, it is mandatory that the service recovery be very fast and the recovery time be of the order of milliseconds. Failure recovery can be done either at the optical layer or at the client layers. SONET and ATM systems may employ their own failure recovery techniques. However, handling failures at the optical layer has some advantages. First, failures can be recovered at the lightpath level faster than at the client layer. Second, when a component such as a node or link fails, the number of lightpaths failed (and thus need to be recovered) is much smaller when compared to the number of failed connections at the client layer. This will not only help restore service quickly but will also result in lesser traffic and control overhead.

There are different approaches to handle failures at the lightpath level in an optical layer. Every working (primary) lightpath can be protected by preassigning resources (wavelengths) to its backup (secondary) lightpath. Upon detecting a failure, service can be switched from the working lightpath to the backup lightpath. Here, the service recovery is almost immediate, as the backup lightpath is readily available. However, it requires excessive resources to be reserved. To overcome this shortcoming, instead of preassigning resources to a backup lightpath, it can be dynamically searched after a failure actually occurs. However, this will result in longer service recovery time and also resources are not guaranteed to be available. Thus, any solution to the survivability problem needs to optimize a certain metric such as resource (wavelength, fiber) requirement, connection acceptance rate, and failure recovery time.

8 Optical Multicast Routing

So far we have assumed that a communication (connection establishment) is between only two nodes—a source and a destination. This (one-to-one) communication is referred to as *unicasting*. However, in several new applications a source needs to communicate to a set of destinations. This (one-to-many) communication is referred to as *multicasting*. Applications which require multicasting include video conferencing, distance learning (where geographically dispersed students listen to the same lecture), distributed databases (all copies of a replicated file are updated at the same time), real-time work groups (files, graphics, and messages are exchanged among active group members in real time). A simple way to realize a multicast session is to establish a unicast path between the source and each of the destinations. However, multicasting using such unicast paths requires many resources (number of distinct wavelengths and wavelength channels— a wavelength channel refers to a wavelength on a link) and also the source of the multicast session has to transmit the same data a number of times which equals the number of destinations. These drawbacks can be overcome if the nodes in a WDM network have the capability to tap and split an optical signal. A wavelength routing node with tapping capability taps a small amount of optical signal (power) from a wavelength channel while forwarding the data on that channel to an output link. The tapped optical power is used by the local node if it is a destination. A node with optical power tapping capability is called a *drop and continue node* or a *DaC node*. A node with splitting capability makes copies of data in optical domain via optical power splitting and thus can forward incoming data to more than one output link.

If a network has splitting capability at all nodes, then it is referred to as a network with *full splitting capability*. The construction of a multicast tree in a network with full splitting capability and full wavelength conversion capability is similar to the construction of a multicast tree in a conventional electronic network, except that bandwidth is replaced by a wavelength. However, a split-capable node (*split node*) is very expensive due to its complex architecture. Hence, only a subset of nodes in a network is assumed to be split-capable, and such a network is referred to as a network with *sparse splitting*.

A node with splitting capability is useful in expanding the multicast tree. However, a node with only splitting capability cannot support more than one connection on the same outgoing link, whereas a split-capable node with wavelength conversion capability can do so by using different wavelengths. This phenomenon can be viewed as a source transmitting multiple messages on any wavelength to as many outgoing links as needed. Hence, a node with splitting and wavelength conversion capabilities is termed a *virtual source* or simply a *VS node*.

A network with sparse splitting and wavelength conversion contain VS nodes, split nodes, and DaC nodes. The main issue in constructing a multicast tree in such a network is to minimize the resources (wavelengths and wavelength channels) and also the (setup) time required for the construction of the tree, exploiting the capabilities of each of these nodes.

9 Network Control and Management

In a wavelength routed WDM network, a control mechanism is needed to set up and tear down lightpaths (all-optical connections). When a connection request arrives, this mechanism must be able to select a route, assign a wavelength to the connection, configure the appropriate switches (WXCs/WIXCs) along the route, and provide information such as what are the existing lightpaths and which wavelengths are currently being used on each fiber link. The control mechanism can be either centralized or distributed. A distributed control mechanism is generally preferred, as it is more robust. The objectives here are to maximize the number of connections established (throughput), minimize the connection setup times, and minimize the bandwidth used for control signals.

The functions performed by a network management system include performance management (monitoring and managing the various parameters, such as throughput, resource [wavelength] utilization, and bit error rate, which measure the performance of the network), fault management (detecting and isolating failed components, and restoring the disrupted traffic), security management (protecting data belonging to network users from being tapped or corrupted by unauthorized users), and accounting management (tracking the usage of network components and charging/billing).

The increasing deployment of WDM technology has presented service providers with a new problem of managing wavelengths (optical channels - OChs) to provide fast and reliable services to their end customers. In order to cost-effectively manage the increasing number of wavelengths, the WDM optical network should support per-wavelength or OCh-level operations, administration, and maintenance (OAM) functions. The OAM functions can be realized using SONET/SDH overhead bytes. But this requires that the signals on each of the wavelengths (client signals) be in SONET/SDH format. *Digital wrapper* technology (adding additional bytes to signal in the optical layer) provides OCh-level OAM functionality similar to SONET/SDH, but the client signals can be in any format (legacy SONET/SDH signals or ATM cells or IP packets).

10 Transmission Impairment

Developing network-layer solutions to counter physical-layer impairments, such as laser shift, dispersion in fiber, and also impairment that affects optical components such as amplifiers, switches, and wavelength converters, is another important issue. For example, most previous networking solutions for the RWA problem assume an ideal physical layer. However, in practice, a signal degrades in quality due to physical-layer impairment as it travels through switches (picking up crosstalk) and EDFAs (picking up noise). This may cause a high bit error rate (BER) at the receiving end of a lightpath. The work in [134] estimates the *online BER* on candidate routes and wavelengths before establishing a connection between a source–destination pair. Thus, one approach is to establish a connection with minimum BER. Another is to establish a connection with BER lower than a certain threshold (for example, 10^{-12}); if no such connection is found, the connection request is rejected. Another networking study which considers physical-layer device characteristics while attempting to solve a network-layer problem is amplifier placement in WDM optical networks.

Optimizing Amplifier Placements

In wavelength routed networks, optical amplification is required to combat various power losses such as fiber attenuation and coupling loss in wavelength routers. Since optical amplifiers are costly, their total number in the network should be minimized, apart from determining their exact placements in the network. However, optical amplifiers have constraints on the maximum gain and the maximum output power they can supply. When optical signals on different wavelengths originating at various nodes at locations separated by large distances arrive at an amplifier, their power levels may be very different. This phenomenon, known as *near–far effect*, can limit the amount of amplification available since the higher-powered wavelengths could saturate the amplifier and limit the gain seen by the lower-powered wavelengths. The amplifier placement problem considering the limitations of the devices (for example, maximum power of a transmitter, fiber attenuation, minimum power required on a wavelength for detection [this represents both the receiver sensitivity level and the amplifier sensitivity level], maximum power available from an amplifier, and maximum [small-signal] amplifier gain) is studied in [135]. The general problem of minimizing the total amplifier count is a mixed-integer nonlinear optimization problem.

11 Ring Networks and Traffic Grooming

WDM optical ring networks have several attractive features over mesh networks, such as easy planning and management, simple control, faster failure restoration, and lower cost. Ring networks are more attractive in metropolitan areas covering a small geographic area, as they can be built from WADMs, which are less expensive than WXCs. Also, they do not require sophisticated amplifiers and repeaters, leading to significant cost savings.

SONET rings are a popular architecture used in many existing networks. A SONET ring is built from SONET add drop multiplexers (SADMs). An SADM is capable of multiplexing and demultiplexing a number of low-speed SONET connections to and from a high-speed connection onto an optical ring using time division multiplexing (TDM) technique. For example, four OC-ff connections at 155.52 Mb/s can be multiplexed onto an OC-12 ring at 622.08 Mb/s. Similarly, 16 OC-ff connections at 155.52 Mb/s can be multiplexed onto an OC-48 ring at 2.48832 Gb/s and four OC-12 connections at 622.08 Mb/s can be multiplexed onto an OC-48 ring at 2.48832 Gb/s. By applying WDM technology, several rings each on a different wavelength can be formed on a physical ring network. Similar to SONET ring networks, WDM ring networks can also be classified as two-fiber unidirectional ring (UR-2), two-fiber bidirectional ring (BR-2), and four-fiber bidirectional ring (BR-4) networks. In a UR-2 network, there are two unidirectional rings in opposite directions. One ring carries traffic during normal operation and the other ring is used for service protection against failures. In a BR-2 network, half of the wavelengths in one ring and half of the wavelengths in the other ring are used for routing normal traffic. All the remaining wavelengths are used for protection purposes. In a BR-4 network, one pair of rings in opposite directions is used for carrying normal traffic while the other pair of rings is used for protection purpose. A multifiber ring network uses a number of fibers in multiples of two or four. In a multifiber ring network, WXCs can be used in place of WADMs to provide flexibility in switching a wavelength to any of the fibers. Another useful architecture for backbone networks is an interconnected ring network, wherein several rings are interconnected. The nodes in a ring can employ WADMs except for those nodes connecting different rings. Such bridging nodes can employ WXCs.

A WDM ring network can be designed to meet a static traffic demand. Here, the objective is to minimize the number of wavelengths or to minimize the number of fibers when the number of wavelengths is fixed. The ring can support dynamic traffic wherein maximizing the accepted traffic demand is an important objective. A configurable WADM can be employed to support dynamic traffic. The wavelength assignment problem is predominant in a WDM ring network when compared to the routing problem as there is a limited choice in selecting a route between two nodes. In a unidirectional ring network, there is only one route between any two nodes. In a bidirectional ring network, there are only two routes between any two nodes and usually the shorter route is preferred as it consumes lesser bandwidth. In a multifiber network, choosing both the wavelength and fiber are critically important when compared to mesh networks, because, the same wavelength may be available on more than one fiber and the possible routes are limited to one or two as discussed above.

Traffic Grooming

A ring combining WDM and SONET technologies, referred to as a WDM-SONET ring, can be built from WADMs and SADMs. A WDM-SONET ring can have *W* number of SONET rings each on a different wavelength, where there are *W* wavelengths per fiber. At a node, SADMs are required one for each of the wavelengths added and dropped. An SADM multiplexes (combines) low speed SONET streams to a high speed traffic and transmits it on the ring corresponding to a wavelength. Similarly, it receives a wavelength from the ring and demultiplexes it into a number of low speed streams. In other words, associated with each SADM there is a unique wavelength. A simple way of designing an N -node WDM-SONET ring is to use W number of SADMs at each node, each for one wavelength. However, this design is not efficient as every wavelength is dropped and added at every node and a large number (N W) of SADMs are used. In practice, the traffic added or dropped at a node does not require all the wavelengths. Thus, a cost-effective design should allow transit traffic to optically bypass a node. Since the electronic processing cost at a node is critical, minimizing the number of SADMs required by a WDM-SONET ring is an important problem. Figure 1.20(a) shows a WDM-SONET ring with four nodes and three wavelengths per fiber. There are three SONET rings, each on a different wavelength. A simple way of designing the ring is to use three SADMs, each corresponding to a wavelength, at every node. In this case, the messages (including the transit traffic) on all the wavelengths are dropped at every node. Allowing the transit traffic to optically bypass intermediate nodes by using the WADMs might help reduce the number of SADMs. In Fig. 1.20(a), wavelength w_0 is optically bypassed by none of the WADMs, wavelength w_1 is optically bypassed by one WADM, and wavelength w_2 is optically bypassed by one WADM.

Aggregating low-speed SONET connections onto a high-speed wavelength so as to minimize the number of SADMs required is known as *the traffic grooming problem*. This problem has been shown to be NP-complete. The maximum number of low-speed connections that can be multiplexed onto a wavelength defines the *multiplexing factor*. For example, if the low-speed SONET streams operate at OC-x speed and a wavelength can operate at a speed of OC-y, $y = x$, then the multiplexing degree is given by y/x. The traffic grooming is said to be valid only if the total bandwidth of the SONET connections carried by any wavelength on a fiber does not exceed the wavelength capacity. A bad traffic grooming algorithm may result in an excessive number of SADMs. Therefore, it is imperative that a good algorithm be used for traffic grooming in order to reduce the cost of the WDM-SONET ring design.

In typical telecommunication backbone networks, multiple rings are interconnected together to provide large geographical coverage. In , traffic grooming in interconnected multi-WDM rings is studied addressing two issues: how to interconnect WDM ring networks, and how to groom the traffic in interconnected rings ,traffic grooming in mesh topology WDM networks is examined. Although mesh topology WDM networks will be of greater importance in the future—at least in the near future—ring topology WDM networks are viable because SONET/SDH architectures are ring-oriented.

12 Virtual Private Networks over WDM Optical Networks

A *virtual private network* (VPN) is a communication network between two or more machines or networks, built for the private use of an organization, over a shared public infrastructure such as the Internet. In other words, a VPN turns the public network (Internet) into a simulated WAN by letting an organization securely extend its network services to remote users, branch offices, and partner companies. VPNs require strong security protocols, such as IPSec (IP Security), to be used for data transfer, as they consist of several machines not under the control of the organization—IP routers and the Internet that carries the traffic [126]. VPNs can make use of the concept of a lightpath offered by WDM, to create secure tunnels (channels) of bandwidth across the WDM backbone network.

13 Access Networks

In the recent past, the backbone telecommunications network has undergone phenomenal changes, with a majority of long-distance links being upgraded to utilize fiber optic technology and data transmission becoming almost totally digital. However, the local access network, which connects customers' homes with the backbone network, has remained relatively unchanged over the past several years: the local telephone network is still based mainly on twisted-pair technology and the cable television network is based on coaxial cable for the most part. Today, most residential connections to the Internet use dial-up modems operating at low speeds on twisted pairs. To run video and advanced Internet applications, even residential customers need huge channel capacity that seems hardly realizable with the traditional copper cable-based access networks. The problem of bringing fiber close to residential and small business customers in order to solve the "last mile" (bandwidth or access) bottleneck in the local loop—the two-wire connections between a subscriber's telephone and the local telephone central office—in a cost-effective manner is a challenging one.

An access network consists of a hub (for example, a telephone central office or a cable television *head end*), remote nodes (RNs), and network interface units (NIUs) [147]. Each hub (which itself may be a part of the backbone network) serves several homes via the NIUs (which may be located in individual homes or may serve several homes). In order to avoid running cables from a hub to each individual NIU, the hub may be connected to several RNs (through a network called a *feeder network*) deployed in the field and each RN in turn may be connected to several NIUs (through a network called a *distribution network*). Access networks can be classified based on the type of feeder network and distribution network. The feeder network can either assign each NIU its own *dedicated* bandwidth or can have the entire bandwidth *shared* for short periods by all the NIUs. The distribution network may be either a switched or a broadcast network. The telephone network is a switched network with each NIU getting its own dedicated bandwidth, whereas the cable television network is a broadcast network with all NIUs sharing the total cable bandwidth.

The different solutions for access networks which are currently being developed include digital subscriber loops (DSLs), hybrid fiber coax (HFC), and fiber-in-the-loop (FITL). It is not clear which of these will better meet the needs of the future broadband telecommunication network, known as the B-ISDN (Broadband Integrated Services Digital Network). *Broadband* refers to a data transmission scheme in which multiple signals share the bandwidth of a medium. It enables the transmission of voice, data, and video signals over a single medium. Cable television employs broadband techniques to deliver dozens of channels over one coaxial cable.

A DSL system uses fiber to connect a central office to a remote node, which performs optical-to-electrical signal conversion and delivers service to several thousand homes by transferring the optical signals onto twisted pairs. A DSL system can be upgraded to higher capacities by using *asymmetric digital subscriber line* (ADSL) technology on the twisted-pairs. ADSL uses sophisticated modulation and coding techniques to provide significantly more bandwidth over the existing twisted-pair infrastructure. The frequency spectrum is divided into two regions, one for conventional analog telephone signals and the other for bidirectional digital transmission. An ADSL system is asymmetric in that the user can transmit upstream into the network at speeds in the range 64–640 kb/s and can receive information downstream from the network at speeds in the range 1.54–6.14 Mb/s, depending on the distance from the central office. This asymmetry in upstream/downstream transmission rates is said to meet the requirements of applications such as upstream requests and downstream page transfers in World Wide Web applications. An ADSL system requires the central office and also the home to have an ADSL modem.

In the past, cable television networks have been used as downstream networks for delivering analog television. Newer cable television systems are HFC networks in which transmission of information from the cable television company head end to remote (fiber) nodes is over fiber optic links and then over a coaxial cable network to subscribers' homes. The current 300- to 450-MHz coax cables will be replaced by 750-MHz coax cables, thereby upgrading the capacity from 50 to 75 6-MHz channels to 125 6-MHz channels (in North America channels are 6 MHz wide). Seventy-five of the 125 channels are for transmitting analog television and the new 50 channels using sophisticated modulation techniques provide a total 2 Gb/s (50_40 Mb/s per channel) of new bandwidth. However, this requires the cable operators to replace all the existing cables with 750-MHz coax and unidirectional analog amplifiers with bidirectional split-band amplifiers for allowing information to flow in both up/down streams. As in the case of ADSL, the upstream/downstream transmission rates are asymmetric here for two-way services (such as high-speed Internet access via cable modem, cable telephony, and pay-per-view video). Note that in DSL the final segment is a point-to-point local loop using twisted pairs, whereas in HFC the final segment is a shared coaxial cable. While DSL is fully switched, HFC uses a shared medium without switching.

FITL is receiving a lot of attention at present because fiber can be economically pushed closer to the customer through the use of a *passive optical network* (PON) for achieving even higher bandwidths [91]. In a PON, a single feeder fiber connects a central office to a passive optical splitter which distributes the transmitted signals to *optical* *network units* (ONUs). The ONUs perform optical-to-electrical signal conversion and deliver service to homes. If an ONU is located in a home, this is called as a *fiber-to-the-home* (FTTH) system; if the ONU is located in the neighborhood (the curb) of its central office, it is called a *fiber-to-the-curb* (FTTC) system. Note that in an FTTC system, an ONU is shared over several homes and hence it requires an additional distribution network from the ONU to the NIUs.

PON technology has several attractive features. A PON allows for a longer distance (over 20 km) between a central office and customer premises compared to that (about 5.5 km) in DSL. A PON provides higher bandwidth due to deeper fiber penetration (FTTH; however, FTTC is the most economical deployment today). PONs use passive optical devices instead of active devices (such as multiplexers and demultiplexers) in the splitting locations, thus relieving network operators from maintaining them and providing power to them. Moreover, the fiber infrastructure being transparent to bit rates and modulation formats, PONs allow easy upgrades to higher bit rates or additional wavelengths without changing the infrastructure itself.

A WDM PON uses a wavelength routing device (for example, arrayed waveguide grating, AWG) at the passive optical splitter to provide a single, dedicated wavelength to every ONU. Different multiplexing techniques can be used to implement bidirectional (up/down stream) transmission. When a single fiber is dedicated to every customer, the downstream traffic can be transmitted employing wavelengths belonging to the 1.55-micron optical window, while for upstream traffic the wavelengths in the 1.3-micron window can be used. When two fibers are dedicated to every customer, each can carry the traffic in one direction (space division bidirectionality). Unfortunately, implementing WDM PONs is a very costly affair. A *time-division multiplexed* (TDM) PON uses an optical power splitter instead of wavelength routing device as in WDM PON, and can be implemented with low-cost components. Since many customers share the fiber, an important issue is how different customers should access the fiber bandwidth. A very simple way is for all customers to operate in the same wavelength band and share the fiber bandwidth in time slots. A TDM PON uses TDM for downstream traffic and *time-division multiple access* (TDMA) for upstream traffic. A large international group of network operators and equipment vendors—the Full Services Access Network (FSAN) Consortium—has been working to standardize PON access systems . By standardizing fiber access systems that will be purchased by a large number of network operators, the FSAN hopes that PON access systems will become more affordable.