10 Optical Network Engineering

George N. Rouskas
Department of Computer Science
North Carolina State University
Raleigh, NC 27695-7534
Email: rouskas@csc.ncsu.edu

10.1 INTRODUCTION

Over the last few years we have witnessed a wide deployment of point-to-point wavelength division multiplexing (WDM) transmission technology in the Internet infrastructure. The corresponding massive increase in network bandwidth due to WDM has heightened the need for faster switching at the core of the network. At the same time, there has been a growing effort to enhance the Internet Protocol (IP) to support traffic engineering [1, 2] as well as different levels of Quality of Service (QoS) [3]. Label Switching Routers (LSRs) running Multi-Protocol Label Switching (MPLS) [4, 5] are being deployed to address the issues of faster switching, QoS support, and traffic engineering. On one hand, label switching simplifies the forwarding function, thereby making it possible to operate at higher data rates. On the other hand, MPLS enables the Internet architecture, built upon the connectionless Internet Protocol, to behave in a connection-oriented fashion that is more conducive to supporting QoS and traffic engineering.

The rapid advancement and evolution of optical technologies makes it possible to move beyond point-to-point WDM transmission systems to an all-optical backbone network that can take full advantage of the available bandwidth by eliminating the need for per-hop packet forwarding. Such a network consists of a number of optical cross-connects (OXCs) arranged in some arbitrary topology, and its main function is to provide interconnection to a number of IP/MPLS subnetworks. Each OXC can switch the optical signal coming in on a wavelength of an input fiber link to the same wavelength in an output fiber link. The OXC may also be equipped with converters that permit it to switch the optical signal on an incoming wavelength of an input fiber to some other wavelength on an output fiber link. The main mechanism of transport in such a network is the lightpath (also referred to as λ-channel), an optical communication channel established over the network of OXCs which may span a number of fiber links (physical hops). If no wavelength converters are used, a lightpath is associated with the same wavelength on each hop. This is the well-known wavelength continuity constraint. Using converters, a different wavelength on each hop may be used to create a lightpath. Thus, a lightpath is an end-to-end optical connection established between two subnetworks attached to the optical backbone.

Currently, there is tremendous interest within both the industry and the research community in optical networks in which OXCs provide the switching functionality. The Internet Engineering Task Force (IETF) is investigating the use of Generalized MPLS (GMPLS) [6] and related signaling protocols to set up and tear down lightpaths. GMPLS is an extension of MPLS that supports multiple types of switching, including switching based on wavelengths usually referred to as Multi-Protocol Lambda Switching (MPλS). With GMPLS, the OXC backbone and the IP/MPLS subnetworks will share common functionality in the control plane, making it possible to seamlessly integrate all-optical networks within the overall Internet infrastructure. Also, the Optical Domain Service Interconnection (ODSI) initiative (which has completed its work) and the Optical Internetworking Forum (OIF) are concerned with the interface between an IP/MPLS subnetwork and the OXC to which it is attached as well as the interface between OXCs, and have several activities to address MPLS over WDM issues [7]. Optical networks have also been the subject of extensive research [8] investigating issues such as virtual topology design [9, 10], call blocking performance [11, 12], protection and restoration [13, 14], routing algorithms and wavelength allocation policies [15, 16, 17], and the effect of wavelength conversion [18, 19, 20], among others.

Given the high cost of network resources and the critical nature of the new applications (especially those supporting business operations) that fuel the growth of the Internet, there has been increasing interest among network providers in traffic engineering techniques [2]. Network resource and traffic performance optimization will continue to be important issues as providers introduce all-optical components in their networks. Since optical device technology has not yet reached the maturity level of electronic component technology, the transmission and switching devices that will need to be deployed to realize an all-optical network tend to be more expensive and bulky (i.e., require more storage space) than their electronic counterparts. Due to the extremely high data rates at which these networks are expected to operate (10-40 Gbps or beyond), a network malfunction or failure has the potential to severely impact critical applications. Also, a single fiber or wavelength may carry a large number of independent traffic streams, and a service outage may have wide implications in terms of the number of customers affected. Therefore, the application of network engineering methods to optimize the utilization of resources while meeting strict performance objectives for the carried traffic will be crucial for the emerging optical networks.

In this chapter we discuss the application of network and traffic engineering techniques to the design and operation of optical networks. The chapter is organized as follows. Section 10.2 introduces the basic elements of the optical network architecture and Section 10.3 discusses the concepts of network and traffic engineering as they apply to optical networks. Section 10.4 describes off-line approaches for the design and capacity planning of optical networks, while Section 10.5 studies on-line algorithms for dynamic provisioning of lightpaths. Section 10.6 discusses standardization activities under way for optical networks, with an emphasis on the control plane issues that are important for traffic engineering. Section 10.7 concludes the chapter.

10.2 OPTICAL NETWORK ARCHITECTURE

The architecture for wide-area WDM networks that is widely expected to form the basis for a future all-optical infrastructure is built on the concept of *wavelength routing*. A wavelength routing network, shown in Figure 10.1, consists of *optical cross-connects (OXCs)* connected by a set of fiber links to form an arbitrary mesh topology. The services that a wavelength routed network offers to attached client subnetworks are in the form of *logical* connections implemented using *lightpaths*. Lightpaths are clear optical paths which may traverse a number of fiber links in the optical network. Information transmitted on a lightpath does not undergo any conversion to and from electrical form within the optical network, and thus, the architecture of the OXCs can be very simple because they do not need to do any signal processing. Furthermore, since a lightpath behaves as a literally transparent "clear channel" between the source and destination subnetwork, there is nothing in the signal path to limit the throughput of the fibers.

The OXCs provide the switching and routing functions for supporting the logical data connections between client subnetworks. An OXC takes in an optical signal at each of the wavelengths at an input port, and can switch it to a particular output port, independent of the other wavelengths. An OXC with N input and N output ports capable of handling W wavelengths per port can be thought of as W independent $N \times N$ optical switches. These switches have to be preceded by a wavelength demultiplexer and followed by a wavelength multiplexer to implement an OXC, as shown in Figure 10.2. Thus, an OXC can cross-connect the different wavelengths from the input to the output, where the connection pattern of each wavelength is independent of the others. By appropriately configuring the OXCs along the physical path, logical connections (lightpaths) may be established between any pair of subnetworks.

As Figure 10.1 illustrates, each OXC has an associated *electronic* control unit attached to one of its input/output ports. The control unit is responsible for control and management functions related to setting up and tearing down lightpaths; these functions are discussed in detail in Section 10.6. In particular, the control unit communicates directly with its OXC, and is responsible for issuing configuration commands to the OXC in order to implement a desired set of lightpath connections; this communication takes place over a (possibly proprietary) interface that depends on the OXC technology. The control unit also communicates with the control units of adjacent OXCs or with attached client subnetworks over *single-hop* lightpaths as shown in Figure 10.1. These lightpaths are typically implemented over administratively configured ports at each OXC and use a separate control wavelength at each fiber. Thus, we distinguish between the paths that data and control signals take in the optical network: data lightpaths originate and terminate at client subnetworks

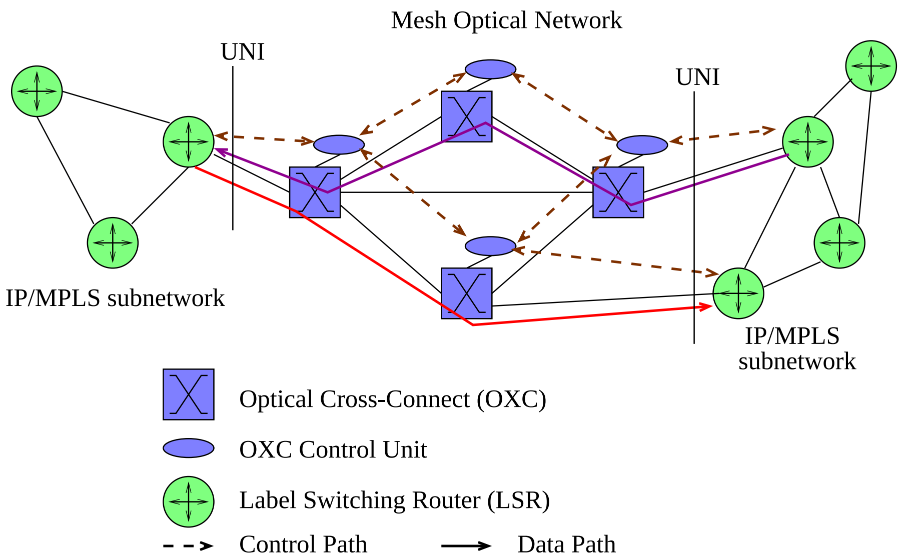

Fig. 10.1 Optical network architecture

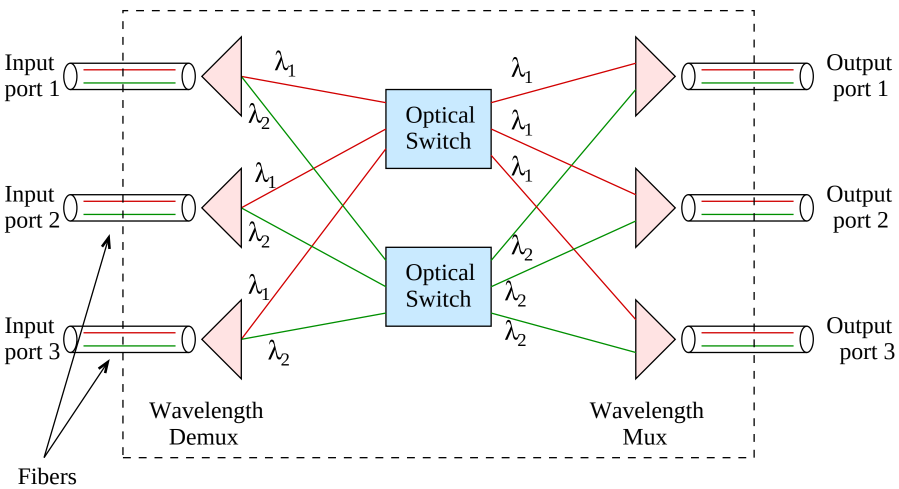

Fig. 10.2 3×3 optical cross-connect (OXC) with two wavelengths per fi ber

and transparently traverse the OXCs, while control lightpaths are electronically terminated at the control unit of each OXC. Communication on the control lightpaths uses a standard signaling protocol (e.g., GMPLS), as well as other standard protocols necessary for carrying out important network functions including label distribution, routing, and network state dissemination. Standardization efforts are crucial to the seamless integration of multi-vendor optical network technology, and are discussed in Section 10.6.

Client subnetworks attach to the optical network via edge nodes which provide the interface between non-optical devices and the optical core. This interface is denoted as UNI (user-to-network interface) in Figure 10.1. The edge nodes act as the terminating points (sources and destinations) for the optical signal paths; the communication paths may continue outside the optical network in electrical form. In Figure 10.1, only the label switching routers (LSRs) of the two IP/MPLS subnetworks which are directly attached to an OXC implement the UNI and may originate or terminate lightpaths. For the remainder of this chapter we will make the assumption that client subnetworks run the IP/MPLS protocols. This assumption reflects the IP-centric nature of the emerging control architecture for optical networks [21]. However, edge nodes supporting any network technology (including ATM switches and SONET/SDH devices) may connect to the optical network as long as an appropriate UNI is defined and implemented.

In addition to simply supporting logical connections between subnetworks, the optical backbone must protect clients from network impairments and the failure of any resources (including fiber links, optical transceivers, and OXCs) in order to ensure reliable operation and continuity of service. In this chapter we assume that the optical network can support the following levels of lightpaths protection [22]:

1. **Dedicated protection**. A protected lightpath consists of a primary (working) route and a dedicated, diversely routed *backup* (protection) route. Under 1+1 protection, the source transmits simultaneously on both the primary and backup path; upon a failure on the primary path, the destination simply switches to the backup path to receive the signal. Under 1:1 protection, transmission occurs on the primary path only, while the backup path may carry low-priority traffic; upon a failure on the primary path, the source and destination of the protected traffic switch to the backup path, and the low-priority traffic is preempted. A 1+1 or 1:1 protected lightpath can recover from any failure on its working route.

2. **Shared protection**. A protected lightpath has a primary route and a diversely routed *shared* backup route. The backup route is shared among a set of protected lightpaths in such a way that any single failure on the primary route of any lightpath in the set can be restored by having the source and destination of the failed lightpath switch to the shared backup path.

3. **No protection**. The lightpath is not protected, i.e., no spare capacity (backup route and wavelength) is set aside to protect from failures along the lightpath route. In the event of a failure the logical connection is lost, however, the

network may attempt to dynamically restore the lightpath by utilizing any available resources.

In [23, 24], the concept of a lightpath was generalized into that of a *light-tree*, which, like a lightpath, is a clear channel originating at a given source node and implemented with a single wavelength. But unlike a lightpath, a light-tree has multiple destination nodes, hence it is a point-to-multipoint channel. The physical links implementing a light-tree form a tree, rooted at the source node, rather than a path in the physical topology, hence the name. Light-trees may be implemented by employing optical devices known as *power splitters* [25] at the OXCs. A power splitter has the ability to split an incoming signal, arriving at some wavelength λ, into up to m outgoing signals, $m \geq 2$; m is referred to as the *fanout* of the power splitter. Each of these m signals is then independently switched to a different output port of the OXC. Note that due to the splitting operation and associated losses, the optical signals resulting from the splitting of the original incoming signal must be amplified before leaving the OXC. Also, to ensure the quality of each outgoing signal, the fanout m of the power splitter may have to be limited to a small number. If the OXC is also capable of wavelength conversion, each of the m outgoing signal may be shifted, independently of the others, to a wavelength different than the incoming wavelength λ. Otherwise, all m outgoing signals must be on the same wavelength λ.

An attractive feature of light-trees is the inherent capability for performing multicasting in the optical domain (as opposed to performing multicasting at a higher layer, e.g., the network layer, which requires electro-optic conversion). Such wavelength routed light-trees are useful for transporting high-bandwidth, real-time applications such as high-definition TV (HDTV). Therefore, OXCs equipped with power splitters will be referred to as *multicast-capable* OXCs (MC-OXCs). Note that, just like with converter devices, incorporating power splitters within an OXC is expected to increase the network cost because of the need for power amplification and the difficulty of fabrication.

10.3 OPTICAL NETWORK AND TRAFFIC ENGINEERING

While the deployment of WDM technology networks in the Internet infrastructure is primarily motivated by the urgent need to accommodate the geometric growth rates of Internet traffic, simply expanding the information carrying capacity of the network is not sufficient to provide high quality service. As the Internet evolves into a mission-critical infrastructure and becomes increasingly commercial and competitive in nature, users will depend on (and indeed expect or demand) reliable operation, expeditious movement of traffic, and survivability, while service providers will need to ensure efficient utilization of network resources and cost-effective operation. Consequently, emerging optical networks must be carefully architected and engineered so as to achieve both user- and provider-specific performance objectives.

In abstract terms, we can think of a network in general, and an optical network in particular, as having three components:

1. a *traffic demand*, either measured or estimated, which is usually expressed as a traffic matrix,

2. a *set of constraints*, which represents the physical fiber layout, the link capacity, the OXCs, and other optical devices (including fiber amplifiers, wavelength converters, etc.) deployed, and

3. a *control policy*, which consists of the network protocols, policies, and mechanisms implemented at the OXC control modules.

Building an optical network that efficiently and reliably satisfies a diverse range of service requests involves the application of *network and traffic engineering* techniques to determine optimal (or near-optimal) operating parameters for each of the above three components.

In essence, therefore, network and traffic engineering is a control problem. In general, depending on the nature of the network technology, control may be exercised at multiple time scales, and different mechanisms are applicable to different time scales [26]. Since optical WDM networks are circuit-switched in nature (i.e., lightpath holding times are assumed to be significantly longer than the round-trip time across the diameter of the optical network), we distinguish two time scales at which it is appropriate to apply control mechanisms:

1. The *connection-level* time scale is the time over which client subnetworks request, use, and tear-down logical connections (lightpaths). The optical network nodes and edge LSRs must implement a *signaling* mechanism for clients to make their requests and declare their traffic requirements. To ensure that the traffic requirements are met, the network must exercise *admission control*, i.e., it may be necessary to deny some lightpath requests. For requests that are allowed, *routing and wavelength assignment* must be performed at connection time. Also, *restoration* mechanisms to recover from network failures operate at this time scale.

2. If the network is persistently overloaded, the admission control algorithm will have to frequently deny lightpath requests and the clients will experience high blocking. The only solution in this situation is to increase network capacity. Since installing new OXCs and fiber trunks may take several months, network engineering techniques that deal with capacity shortage operate at this time scale. These techniques fall under the broad areas of *network design* and *capacity planning*. Also, *protection* techniques must be employed to set aside adequate spare capacity for lightpaths that must be protected from network failures.

Table 10.1 summarizes the control mechanisms appropriate for each of the two time scales.

We note that network and traffic engineering is by nature an adaptive process. Initially the network topology is determined and a control policy is formulated; the policy depends not only on the constraints imposed by the network topology but

Table 10.1 Control time scales in optical networks, and associated mechanisms

Time Scale	Control Mechanism	Point Exercised	
		OXC Control Unit	Edge LSR
Connection-level (network operation phase)	Signaling	Y	Y
	Admission control	Y	N
	RWA	Y	N
	Restoration	Y	N
Months or longer (network design phase)	Capacity planning	Y	N
	Protection	Y	N

also on factors such as the cost structure, the revenue model, and the performance objectives. During network operation, the traffic and network state are continuously observed and analyzed. The observation phase requires that a set of on-line performance and fault monitoring functions be implemented in the network. In the analysis phase, various techniques are applied to identify existing or potential bottlenecks that (may) affect network performance. The feedback from the analysis stage is used to update the network topology and/or reformulate the control policy in order to drive the network to a desirable operating range. The network and traffic engineering cycle then begins anew.

While capacity planning and dynamic provisioning of lightpaths operate at different time scales, the underlying control problem that must be addressed at both the network design and network operation phases is the routing and wavelength assignment (RWA) problem. Because of the prominent role of the RWA problem in the design and operation of optical WDM networks, it is discussed in detail in the following subsection.

10.3.1 Routing and Wavelength Assignment (RWA)

A unique feature of optical WDM networks is the tight coupling between routing and wavelength selection. As can be seen in Figure 10.1, a lightpath is implemented by selecting a path of physical links between the source and destination edge nodes, and reserving a particular wavelength on each of these links for the lightpath. Thus, in establishing an optical connection we must deal with both routing (selecting a suitable path) and wavelength assignment (allocating an available wavelength for the connection). The resulting problem is referred to as the *routing and wavelength assignment (RWA)* problem [17], and is significantly more difficult than the routing problem in electronic networks. The additional complexity arises from the fact that routing and wavelength assignment are subject to the following two constraints:

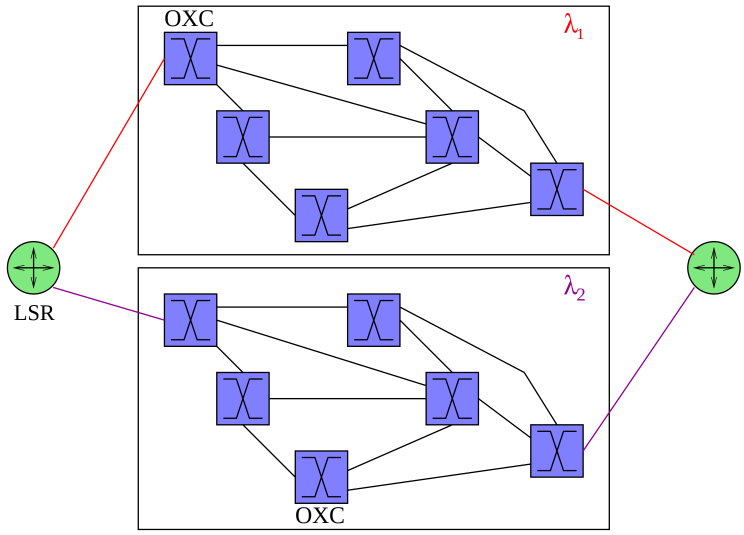

Fig. 10.3 The RWA problem with two wavelengths per fi ber

1. *Wavelength continuity constraint:* a lightpath must use the same wavelength on all the links along its path from source to destination edge node.

2. *Distinct wavelength constraint:* all lightpaths using the same link (fiber) must be allocated distinct wavelengths.

The RWA problem in optical networks is illustrated in Figure 10.3, where it is assumed that each fiber supports two wavelengths. The effect of the wavelength continuity constraint is represented by replicating the network into as many copies as the number of wavelengths (in this case, two). If wavelength i is selected for a lightpath, the source and destination edge node communicate over the i-th copy of the network. Thus, finding a path for a connection may potentially involve solving W routing problems for a network with W wavelengths, one for each copy of the network.

The wavelength continuity constraint may be relaxed if the OXCs are equipped with *wavelength converters* [18]. A wavelength converter is a single input/output device that converts the wavelength of an optical signal arriving at its input port to a different wavelength as the signal departs from its output port, but otherwise leaves the optical signal unchanged. In OXCs without a wavelength conversion capability, an incoming signal at port p_i on wavelength λ can be optically switched to any port p_j, but must leave the OXC on the same wavelength λ. With wavelength converters, this signal could be optically switched to any port p_j on some other wavelength λ'. That is, wavelength conversion allows a lightpath to use different wavelengths along different physical links.

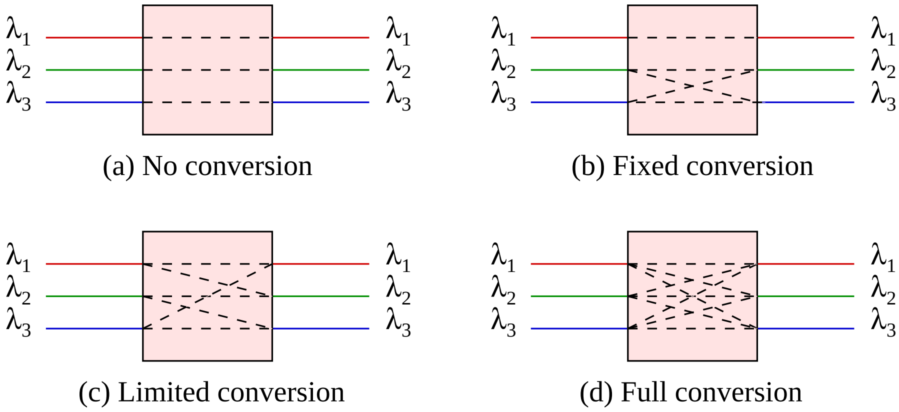

Fig. 10.4 Wavelength conversion

Different levels of wavelength conversion capability are possible. Figure 10.4 illustrates the differences for a single input and single output port situation; the case for multiple ports is more complicated but similar. *Full wavelength conversion* capability implies that any input wavelength may be converted to any other wavelength. *Limited wavelength conversion* [27] denotes that each input wavelength may be converted to any of a specific set of wavelengths, which is not the set of all wavelengths for at least one input wavelength. A special case of this is *fixed wavelength conversion*, where each input wavelength can be converted to exactly one other wavelength. If each wavelength is "converted" only to itself, then we have no conversion.

The advantage of full wavelength conversion is that it removes the wavelength continuity constraint, making it possible to establish a lightpath as long as each link along the path from source to destination has a free wavelength (which could be different for different links). As a result, the RWA problem reduces to the classical routing problem, that is, finding a suitable path for each connection in the network. Referring to Figure 10.3, full wavelength conversion collapses the W copies of the network into a single copy on which the routing problem is solved. On the other hand, with limited conversion, the RWA problem becomes more complex than with no conversion. To see why, note that employing limited conversion at the OXCs introduces links between *some* of the network copies of Figure 10.3. For example, if wavelength λ_1 can be converted to wavelength λ_2 but not to wavelength λ_3, then links must be introduced from each OXC in copy 1 of the network to the corresponding OXC in copy 2, but not to the corresponding OXC in copy 3. When selecting a path for the connection, at each OXC there is the option of remaining at the same network copy or moving to another one, depending on the conversion capability of the OXC. Since the number of alternatives increases exponentially with the number of OXCs that need to be traversed, the complexity of the RWA problem increases accordingly.

Wavelength conversion (full or limited) increases the routing choices for a given lightpath (i.e., makes more efficient use of wavelengths), resulting in better performance. Since converter devices increase network cost, a possible middle ground is to use *sparse conversion*, that is, to employ converters in some, but not all, OXCs in the network. In this case, a lightpath must use the same wavelength along each link in a segment of its path between OXCs equipped with converters, but it may use a different wavelength along the links of another such segment. It has been shown that implementing full conversion at a relatively small fraction of the OXCs in the network is sufficient to achieve almost all the benefits of conversion [11, 19].

With the availability of MC-OXCs and the existence of multicast traffic demands, the problem of establishing light-trees to satisfy these demands arises. We will call this problem the *multicast routing and wavelength assignment (MC-RWA)* problem. MC-RWA bears many similarities to the RWA problem discussed above. Specifically, the tight coupling between routing and wavelength assignment remains, and even becomes stronger: in the absence of wavelength conversion the same wavelength must be used by the multicast connection not just along the links of a single path but along all the links of the light-tree. Since the construction of optimal trees for routing multicast connections is by itself a hard problem [28], the combined MC-RWA problem becomes even harder.

Routing and wavelength assignment is the fundamental control problem in optical WDM networks. Since the performance of a network depends not only on its physical resources (e.g., OXCs, converters, fibers links, number of wavelengths per fiber, etc.) but also on how it is controlled, the objective of an RWA algorithm is to achieve the best possible performance within the limits of physical constraints. The RWA (and MC-RWA) problem can be cast in numerous forms. The different variants of the problem, however, can be classified under one of two broad versions: a static RWA, whereby the traffic requirements are known in advance, and a dynamic RWA, in which a sequence of lightpath requests arrive in some random fashion. The static RWA problem arises naturally in the design and capacity planning phase of architecting an optical network, and is discussed in Section 10.4. The dynamic RWA problem is encountered during the real-time network operation phase and involves the dynamic provisioning of lightpaths; this issue is addressed in Section 10.5.

10.4 OPTICAL NETWORK DESIGN AND CAPACITY PLANNING

If the traffic patterns in the network are reasonably well-known in advance and any traffic variations take place over long time scales, the most effective technique for establishing optical connections (lightpaths) between client subnetworks is by formulating and solving a static RWA problem. Therefore, static RWA is appropriate for provisioning a set of semipermanent connections. Since these connections are assumed to remain in place for relatively long periods of time, it is worthwhile to attempt to optimize the way in which network resources (e.g., physical links and wavelengths) are assigned to each connection, even though optimization may require a considerable computational effort. Because off-line algorithms have knowledge of the entire set of demands (as opposed to on-line algorithms that have no knowledge of future demands), they make more efficient use of network resources and project a lower overall capacity requirement.

In general, the optical network design and planning phase consists of three steps:

1. **Determination of user demand.** If a network already exists, user demand can be measured; otherwise, the demand must be estimated from the expected user population and the expected usage patterns.

2. **Physical topology design.** Given the traffic demand, the physical topology of OXCs and fiber links to connect client subnetworks is determined.

3. **Virtual topology design.** At this step, a static RWA problem is formulated and solved in order to create lightpaths between client subnetworks to satisfy the traffic demand.

For the remainder of this section we will assume that the user demand is known, and we will concentrate on the physical and virtual topology design problems and related issues.

10.4.1 Physical Topology Design

In this phase the network operator has a demand forecast and must decide on a topology to connect client subnetworks through OXCs. This step includes the sizing of links (e.g., determining the number of wavelength channels and the capacity of each channel) and OXCs (e.g, determining the number of ports), as well as the placement of resources such as amplifiers, wavelength converters, and power splitters. Moreover, to deal with link or OXC failures, it is desirable to ensure that there are at least two (or three) paths between any pair of OXCs in the network, i.e., that the graph corresponding to the physical topology of the optical network is two- or three-connected. Often, geographical or administrative considerations may impose further constraints on the physical topology.

If a network does not already exist, the physical topology must be designed from scratch. Obviously, the outcome of this step strongly depends on the accuracy of the demand forecast, and the potential for error is significant when designers have to guess the load in a new network. Therefore, many providers take a cautious approach by initially building a skeleton network and adding new resources as necessary by actual user demand. In this *incremental* network design, it is assumed that sets of user demands arrive over multiple time periods. Resources (e.g., OXCs, fiber links, wavelength channels) are added incrementally to satisfy each new set of demands, in a way that the additional capacity required is minimized.

A physical topology design problem was considered in [29]. Given a number of LSRs and a set of lightpaths to be set up among pairs of LSRs, the objective was to determine the two-connected physical topology with the minimum number of OXCs to establish all the lightpaths (this is a combined physical/virtual topology design problem in that the routing and wavelength assignment for the lightpaths is also determined). An iterative solution approach was considered, whereby a genetic algorithm was used to iterate over the space of physical topologies, and heuristics were employed for routing and wavelength assignment on a given physical topology (refer to the next subsection for details on RWA heuristics). The algorithm was applied to networks with up to 1000 LSRs and tens of thousands of lightpaths, and provided insight into the capacity requirements for realistic optical networks. For example, it was shown that the number of OXCs increases much slower than the number of LSRs, and also that the number of OXCs increases only moderately as the number of lightpaths increases by a factor of two or three. These results indicate that optical networks to interconnect a large number of LSRs can be built to provide rich connectivity with moderate cost.

Other studies related to capacity planning have looked into the problem of optimally placing network resources such as converters or power splitters (for multicast). The problem of converter placement was addressed in [11, 30], and optimal [30] (for uniform traffic only) or near-optimal greedy [11] algorithms (for general traffic patterns) were developed. While both studies established that a small number of converters (approximately 30% of the number of OXCs) is sufficient, the results in [11] demonstrate that (a) the optimal placement of converters is extremely sensitive to the actual traffic pattern, and (b) an incremental approach to deploying converters may not lead to optimal (or near-optimal) results. The work in [31] considered the problem of optimally allocating the multicast-capable OXCs (MC-OXCs) to establish light-trees, and a greedy heuristic was proposed. It was found that there is little performance improvement if more than 50% of the OXCs in the network are multicast-capable, and that the optimal location of MC-OXCs depends on the traffic pattern.

Overall, the physical topology design problem is quite complex because the topology, the link and OXC capacities, and the number and location of optical devices such as converters and amplifiers strongly depends on the routing of lightpaths and the wavelength assignment strategy. If we make the problem less constrained, allowing the topology, routing, wavelength assignment, link capacity, etc., to change, the problem becomes very hard because these parameters are coupled in complicated ways. In practice, the topology may be constrained by external factors making the problem easier to deal with; for instance the existence of a deployed fiber infrastructure may dictate the location of OXCs and the links between them. However, the area of physical topology design for optical networks remains a rich area for future research.

10.4.2 Virtual Topology Design

A solution to the static RWA problem consists of a set of long-lived lightpaths which create a *logical* (or *virtual*) topology among the edge nodes. This virtual topology is embedded onto the physical topology of optical fiber links and OXCs. Accordingly, the static RWA problem is often referred to as the *virtual topology design* problem [9]. In the virtual topology, there is a directed link from edge node s to edge node d if a lightpath originating at s and terminating at d is set up (refer also to Figure 10.1), and edge node s is said to be "one hop away" from edge node d in the virtual topology, although the two nodes may be separated by a number of physical links. The type of virtual topology that can be created is usually constrained by the underlying physical topology. In particular, it is generally not possible to implement fully connected virtual topologies: for N edge nodes this would require each edge node to maintain $N-1$ lightpaths and the optical network to support a total of $N(N-1)$ lightpaths. Even for modest values of N, this degree of connectivity is beyond the reach of current optical technology, both in terms of the number of wavelengths that can be supported and in terms of the optical hardware (transmitters and receivers) required at each edge node.

In its most general form, the RWA problem is specified by providing the physical topology of the network and the traffic requirements. The physical topology corresponds to the deployment of cables in some existing fiber infrastructure, and is given as a graph $G_p(V, E_p)$, where V is the set of OXCs and E_p is the set of fibers that interconnect them. The traffic requirements are specified in a traffic matrix **T** = $[\rho p_{sd}]$, where ρp_{sd} is a measure of the long-term traffic flowing from source edge node s to destination edge node d [32]. Quantity ρ represents the (deterministic) total offered load to the network, while the p_{sd} parameters define the distribution of the offered traffic.

Routing and wavelength assignment are considered together as an optimization problem using integer programming formulations. Usually, the objective of the formulation is to minimize the maximum congestion level in the network subject to network resource constraints [9, 10]. While other objective functions are possible, such as minimizing the average weighted number of hops or minimizing the average packet delay, minimizing network congestion is preferable since it can lead to linear programming (ILP) formulations. While we do not present the RWA problem formulation here, the interested reader may refer to [32, 9, 10]. These formulations turn out to have extremely large numbers of variables, and are intractable for large networks. This fact has motivated the development of heuristic approaches for finding good solutions efficiently.

Before we describe the various heuristic approaches, we note that the static RWA problem can be logically decomposed into four subproblems. The decomposition is approximate or inexact, in the sense that solving the subproblems in sequence and combining the solutions may not result in the optimal solution for the fully integrated problem, or some later subproblem may have no solution given the solution obtained for an earlier subproblem, so no solution to the original problem may be obtained. However, the decomposition provides insight into the structure of the RWA problem and is a first step towards the design of effective heuristics. Assuming no wavelength conversion, the subproblems are as follows.

1. **Topology Subproblem:** Determine the logical topology to be imposed on the physical topology, that is, determine the lightpaths in terms of their source and destination edge nodes.

2. **Lightpath Routing Subproblem:** Determine the physical links which each lightpath consists of, that is, route the lightpaths over the physical topology.

3. **Wavelength Assignment Subproblem:** Determine the wavelength each lightpath uses, that is, assign a wavelength to each lightpath in the logical topology so that wavelength restrictions are obeyed for each physical link.

4. **Traffic Routing Subproblem:** Route packet traffic between source and destination edge nodes over the logical topology obtained.

A large number of heuristic algorithms have been developed in the literature to solve the general static RWA problem discussed here or its many variants. Overall, however, the different heuristics can be classified into three broad categories: (1) algorithms which solve the overall ILP problem sub-optimally, (2) algorithms which tackle only a subset of the four subproblems, and (3) algorithms which address the problem of embedding regular logical topologies onto the physical topology.

Suboptimal solutions can be obtained by applying classical tools developed for complex optimization problems directly to the ILP problem. One technique is to use LP-relaxation followed by rounding [33]. In this case, the integer constraints are relaxed creating a non-integral problem which can be solved by some linear programming method, and then a rounding algorithm is applied to obtain a new solution which obeys the integer constraints. Alternatively, genetic algorithms or simulated annealing [34] can be applied to obtain locally optimal solutions. The main drawback of these approaches is that it is difficult to control the quality of the final solution for large networks: simulated annealing is computationally expensive and thus, it may not be possible to adequately explore the state space, while LP-relaxation may lead to solutions from which it is difficult to apply rounding algorithms.

Another class of algorithms tackles the RWA problem by initially solving the first three subproblems listed above; traffic routing is then performed by employing well-known routing algorithms on the logical topology. One approach for solving the three subproblems is to maximize the amount of traffic that is carried on one-hop lightpaths, i.e., traffic that is routed from source to destination edge node directly on a lightpath. A greedy approach taken in [35] is to create lightpaths between edge nodes in order of decreasing traffic demands as long as the wavelength continuity and distinct wavelength constraints are satisfied. This algorithm starts with a logical topology with no links (lightpaths) and sequentially adds lightpaths as long as doing so does not violate any of the problem constraints. The reverse approach is also possible [36]: starting with a fully connected logical topology, an algorithm sequentially removes the lightpath carrying the smallest traffic flows until no constraint is violated. At each step (i.e., after removing a lightpath), the traffic routing subproblem is solved in order to find the lightpath with the smallest flow.

The third approach to RWA is to start with a given logical topology, thus avoiding to directly solve the first of the four subproblems listed above. Regular topologies are good candidates as logical topologies since they are well understood and results regarding bounds and averages (e.g., for hop lengths) are easier to derive. Algorithms for routing traffic on a regular topology are usually simple, so the traffic routing subproblem can be trivially solved. Also, regular topologies possess inherent load balancing characteristics which are important when the objective is to minimize the maximum congestion.

Once a regular topology is decided on as the one to implement the logical topology, it remains to decide which physical node will realize each given node in the regular topology (this is usually referred to as the *node mapping* subproblem), and which sequence of physical links will be used to realize each given edge (lightpath) in the regular topology (this *path mapping* subproblem is equivalent to the lightpath routing and wavelength assignment subproblems discussed earlier). This procedure is usually referred to embedding a regular topology in the physical topology. Both the node and path mapping subproblems are intractable, and heuristics have been proposed in the literature [36, 37]. For instance, a heuristic for mapping the nodes of shuffle topologies based on the gradient algorithm was developed in [37].

Given that all the algorithms for the RWA problem are based on heuristics, it is important to be able to characterize the quality of the solutions obtained. To this end, one must resort to comparing the solutions to known bounds on the optimal solution. A comprehensive discussion of bounds for the RWA problem and the theoretical considerations involved in deriving them can be found in [9]. A simulation-based comparison of the relative performance of the three classes of heuristic for the RWA problem is presented in [10]. The results indicate that the second class of algorithms discussed earlier achieve the best performance.

The study in [23] also focused on virtual topology design (i.e., static RWA) for point-to-point traffic but observed that, since a light-tree is a more general representation of a lightpath, the set of virtual topologies that can be implemented using light-trees is a superset of the virtual topologies that can be implemented only using lightpaths. Thus, for any given virtual topology problem, an optimal solution using light-trees is guaranteed to be at least as good and possibly an improvement over the optimal solution obtained using only lightpaths. Furthermore, it was demonstrated that by extending the lightpath concept to a light-tree, the network performance (in terms of average packet hops) can be improved while the network cost (in terms of the number of optical transmitters/receivers required) decreases.

The static MC-RWA problem has been studied in [38, 39]. The study in [38] focused on demonstrating the benefits of multicasting in wavelength routed optical networks. Specifically, it was shown that using light-trees (spanning the source and destination nodes) rather than individual parallel lightpaths (each connecting the source to an individual destination) requires fewer wavelengths and consumes a significantly lower amount of bandwidth. In [39] an ILP formulation that maximizes the total number of multicast connections was presented for the static MC-RWA problem. Rather than providing heuristic algorithms for solving the ILP, bounds on the objective function were presented by relaxing the integer constraints.

10.4.3 Design of Survivable Optical Networks

In the previous section we considered the RWA problem under the assumption that lightpaths are not protected. To ensure that protected lightpaths survive network failures, it is important to develop solutions to the static RWA problem that take into account the requirements of lightpaths in terms of dedicated or shared protection (refer to Section 10.2). In other words, for each protected lightpath a diversely routed backup path (dedicated or shared) must also be computed and reserved during the network design phase. Clearly, the provisioning of backup lightpaths will increase the capacity requirements compared to a network that provides no protection. Therefore, a common objective in the design of survivable optical networks is to provide adequate protection while minimizing the total network capacity.

Optical network architectures are exposed to a wide range of risks due to either human activities (e.g., accidental fiber cuts or operational errors), or equipment malfunctions such as laser and OXC failures. To ensure survivability, the primary and backup route for a lightpath must be diverse so that a single failure will not affect both paths. Since several network links may pass though the same conduit or share the same right-of-way, two link-disjoint paths may not be immune to a single failure. The concept of a *shared risk link group* (SRLG) [40] can be used to express the risk relationship of optical channels with respect to a single failure. Specifically, each point of failure in the network (e.g., a fiber, a cable, or a conduit) is associated with an SRLG which is the set of channels that would be affected by the failure. A particular link can belong to several SRLGs, and diverse routing of lightpaths implies that the primary and backup routes are SRLG-disjoint. Therefore, the availability of SRLG information is crucial to the design of survivable networks. In general, the different types of SRLGs are defined by the network operator, and the mapping of links into SRLGs must be manually configured since it may be impossible for network elements to self-discover this information. We also note that SRLG information does not change dynamically, i.e, it need not be updated unless changes are made to the network capacity or topology.

Protection from failures can be classified as either *local span (link) protection* or *path (end-to-end) protection*. In path protection, the ingress and egress OXCs of the failed lightpath coordinate to restore the signal on a predefined backup path which is SRLG-disjoint from the primary path. In local span protection, the OXCs closest to the failure coordinate to reroute the lightpath through alternate channels around the failure. These alternate channels are different for each failure, and must be configured in advance and set aside for protection.

Given the SRLG information, the type of protection (local span or path protection), and the protection requirements of each lightpath (dedicated, shared, or no protection), the problem of designing a virtual topology that meets the lightpaths protection requirements while minimizing total network capacity can be be modeled as an ILP and has been studied in [41, 42, 43]. While we do not provide the formulation here, the reader is referred to [42] for three different formulations corresponding to dedicated path protection, shared path protection, and shared local span protection. The formulation is similar to the one discussed in the previous section for unprotected lightpaths, with additional constraints to account for the fact that the primary and backup paths for a given lightpath are SRLG-disjoint, that two primary paths can share a link (channel) on their backup path only if they are SRLG-disjoint, and that a backup path consumes capacity if and only if there is a failure on the primary path. Consequently, similar solution approaches can be employed to obtain a near-optimal set of primary and backup paths for the given set of lightpaths.

Finally, we note that the techniques described in this section can also be used in designing virtual private networks (VPNs) over the optical network [44]. A VPN is an overlay network implemented over a public network infrastructure using tunneling, such that the client nodes at the endpoints of a tunnel appear to be connected by a point-to-point link of a specified bandwidth. VPN services are expected to be key drivers of new network technologies since they allow corporations to extend their services over geographically dispersed regions in a cost-effective manner. Providing VPN service over optical networks has the additional advantage of enhancing privacy and allaying security concerns due to the data transparency in the optical domain. From the point of view of the optical network provider, a VPN is simply a set of client edge nodes that need to be interconnected by lightpaths with certain protection requirements. The provider's objective is to maximize the total amount of VPN traffic by optimally utilizing the capacity of the optical network. This problem is similar to the virtual topology design problem we discussed in this section and can be addressed by the methodologies we have already described.

10.5 DYNAMIC LIGHTPATH PROVISIONING AND RESTORATION

During real-time network operation, edge nodes submit to the network requests for lightpaths to be set up as needed. Thus, connection requests are initiated in some random fashion. Depending on the state of the network at the time of a request, the available resources may or may not be sufficient to establish a lightpath between the corresponding source-destination edge node pair. The network state consists of the physical path (route) and wavelength assignment for all active lightpaths. The state evolves randomly in time as new lightpaths are admitted and existing lightpaths are released. Thus, each time a request is made, an algorithm must be executed in real time to determine whether it is feasible to accommodate the request, and, if so, to perform routing and wavelength assignment. If a request for a lightpath cannot be accepted because of lack of resources, it is blocked.

Because of the real-time nature of the problem, RWA algorithms in a dynamic traffic environment must be very simple. Since combined routing and wavelength assignment is a hard problem, a typical approach to designing efficient algorithms is to decouple the problem into two separate subproblems: the routing problem and the wavelength assignment problem. Consequently, most dynamic RWA algorithms for wavelength routed networks consist of the following general steps:

1. Compute a number of candidate physical paths for each source-destination edge node pair and arrange them in a path list.

2. Order all wavelengths in a wavelength list.

3. Starting with the path and wavelength at the top of the corresponding list, search for a feasible path and wavelength for the requested lightpath.

The specific nature of a dynamic RWA algorithm is determined by the number of candidate paths and how they are computed, the order in which paths and wavelengths are listed, and the order in which the path and wavelength lists are accessed.

10.5.1 Route Computation

Let us first discuss the routing subproblem. If a *static* algorithm is used, the paths are computed and ordered independently of the network state. With an *adaptive* algorithm, on the other hand, the paths computed and their order may vary according to the current state of the network. A static algorithm is executed off-line and the computed paths are stored for later use, resulting in low latency during lightpath establishment. Adaptive algorithms are executed at the time a lightpath request arrives and require network nodes to exchange information regarding the network state. Lightpath set up delay may also increase, but in general, adaptive algorithms improve network performance.

The number of path choices for establishing an optical connection is another important parameter. A *fixed* routing algorithm is a static algorithm in which every source-destination edge node pair is assigned a single path. With this scheme, a connection is blocked if there is no wavelength available on the designated path at the time of the request. In *fixed-alternate* routing, a number $k, k > 1$, of paths are computed and ordered off-line for each source-destination edge node pair. When a request arrives, these paths are examined in the specified order and the first one with a free wavelength is used to establish the lightpath. The request is blocked if no wavelength is available in any of the k paths. Similarly, an adaptive routing algorithm may compute a single path, or a number of alternate paths at the time of the request. A hybrid approach is to compute k paths off-line, however, the order in which the paths are considered is determined according to the network state at the time the connection request is made (e.g., least to most congested).

In most practical cases, the candidate paths for a request are considered in increasing order of *path length* (or *path cost*). Path length is typically defined as the sum of the weights (costs) assigned to each physical link along the path, and the weights are chosen according to some desirable routing criterion. Since weights can be assigned arbitrarily, they offer a wide range of possibilities for selecting path priorities. For example, in a static (fixed-alternate) routing algorithm, the weight of each link could be set to 1, or to the physical distance of the link. In the former case, the path list consists of the k minimum-hop paths, while in the latter the candidate paths are the k minimum-distance paths (where distance is defined as the geographic length). In an adaptive routing algorithm, link weights may reflect the load or "interference" on a link (i.e., the number of active lightpaths sharing the link). By assigning small weights to least loaded links, paths with larger number of free channels on their links rise to the head of the path list, resulting in a *least loaded* routing algorithm. Paths that are congested become "longer" and are moved further down the list; this tends to avoid heavily loaded bottleneck links. Many other weighting functions are possible.

When path lengths are sums of of link weights, the k-shortest path algorithm [45] can be used to compute candidate paths. Each path is checked in order of increasing length, and the first that is feasible is assigned the first free wavelength in the wavelength list. However, the k shortest paths constructed by this algorithm usually share links. Therefore, if one path in the list is not feasible, it is likely that other paths in the list with which it shares a link will also be infeasible. To reduce the risk of blocking, the k shortest paths can be computed so as to be pairwise link-disjoint. This can be accomplished as follows: when computing the i-th shortest path, $i = 1, \cdots, k$, the links used by the first $i-1$ paths are removed from the original network topology and Dijkstra's shortest path algorithm [46] is applied to the resulting topology. This approach increases the chances of finding a feasible path for a connection request.

If the lightpath requires protection (dedicated or shared), a similar procedure can be used at connection setup time to determine two SRLG-disjoint paths, a primary (working) path and a backup (protection) one. The path cost for a lightpath with dedicated protection is the sum of the costs of the links along the primary and backup paths. Thus, the objective is to pick a pair of link-disjoint paths with the minimum combined cost. If the lightpath requires shared protection, the links of the backup path may be shared among other lightpaths. Note, however, that when determining the backup path for a given lightpath, selecting any link that is already shared by the backup paths of other lightpaths does not incur any additional cost (since the link has already been set aside for protection). Thus, the above algorithm can be used for determining the primary and backup paths by setting the weight of shared links in the backup path to zero.

An alternative to pre-computed backup paths is to attempt to restore a lightpath after a failure, by dynamically determining a new route for the failed lightpath [14]. As with protection (see Section 10.4.3), restoration techniques can be classified as either *local span (link) restoration*, whereby the OXCs closest to the failure attempt to restore the lightpath on any available channels around the failure, or *end-to-end (path) restoration*, whereby the ingress OXC dynamically computes a new route to the egress OXC of the failed lightpath. This approach has the advantage that it incurs low overhead in the absence of failures. However, dynamic restoration has two disadvantages that make it unsuitable for lightpaths carrying critical traffic: there is no guarantee of successful restoration (e.g., due to lack of available resources at the time of recovery), and it takes significantly longer to restore traffic (e.g., due to the need to compute a new route on the fly) than when protection resources are provisioned in advance.

The problem of determining algorithms for routing multicast optical connections has also been studied in [39, 47]. The problem of constructing trees for routing multicast connections was considered in [47] independently of wavelength assignment, under the assumption that not all OXCs are multicast capable, i.e., that there is a limited number of MC-OXCs in the network. Four new algorithms were developed for routing multicast connections under this *sparse light splitting* scenario. While the algorithms differ slightly from each other, the main idea to accommodate sparse splitting is to start with the assumption that all OXCs in the network are multicast capable and use an existing algorithm to build an initial tree. Such a tree is infeasible if a non-multicast-capable OXC is a branching point. In this case, all but one branches out of this OXC are removed, and destination nodes in the removed branches have to join the tree at a MC-OXC. In [39], on the other hand, the MC-RWA problem was solved by decoupling the routing and wavelength assignment problems. A number of *alternate* trees are constructed for each multicast connection using existing routing algorithms. When a request for a connection arrives, the associated trees are considered in a fixed order. For each tree, wavelengths are also considered in a fixed order (i.e., the first-fit strategy discussed in the next subsection). The connection is blocked if no free wavelength is found in any of the trees associated with the multicast connection.

We note that most of the literature (and the preceding discussion) has focused on the problem of obtaining paths that are optimal with respect to total path cost. In transparent optical networks, however, optical signals may suffer from physical layer impairments including attenuation, chromatic dispersion, polarization mode dispersion (PMD), amplifier spontaneous emission (ASE), cross-talk, and various nonlinearities [40]. These impairments must be taken into account when choosing a physical path. In general, the effect of physical layer impairments may be translated into a set of constraints that the physical path must satisfy; for instance, the total signal attenuation along the physical path must be within a certain power budget to guarantee a minimum level of signal quality at the receiver. Therefore, a simple shortest path first (SPF) algorithm (e.g., Dijkstra's algorithm implemented by protocols such as OSPF [48]) may not be appropriate for computing physical paths within a transparent optical network. Rather, constraint-based routing techniques such as the one employed by the constraint-based shortest path first (CSPF) algorithm [5] are needed. These techniques compute paths by taking into account not only the link cost but also a set of constraints that the path must satisfy. A first step towards the design of constraint-based routing algorithms for optical networks has been taken in [40] where it was shown how to translate the PMD and ASE impairments into a set of linear constraints on the end-to-end physical path. However, additional work is required to advance our understanding of how routing is affected by physical layer considerations, and constraint-based routing remains an open research area [49].

10.5.2 Wavelength Assignment

Let us now discuss the wavelength assignment subproblem which is concerned with the manner in which the wavelength list is ordered. For a given candidate path, wavelengths are considered in the order in which they appear in the list to find a free wavelength for the connection request. Again, we distinguish between the static and adaptive cases. In the static case, the wavelength ordering is fixed (e.g., the list is ordered by wavelength number). The idea behind this scheme, also referred to as *first-fit*, is to pack all the in-use wavelengths towards the top of the list so that wavelengths towards the end of the list will have higher probability of being available over long continuous paths. In the adaptive case, the ordering of wavelengths is typically based on usage. Usage can be defined either as the number of links in the network in which a wavelength is currently used, or as the number of active connections using a wavelength. Under the *most used* method, the most used wavelengths are considered first (i.e., wavelength are considered in order of decreasing usage). The rationale behind this method is to reuse active wavelengths as much as possible before trying others, packing connections into fewer wavelengths and conserving the spare capacity of less-used wavelengths. This in turn makes it more likely to find wavelengths that satisfy the continuity requirement over long paths. Under the *least used* method, wavelengths are tried in the order of increasing usage. This scheme attempts to balance the load as equally as possible among all the available wavelengths. However, least used assignment tends to “fragment” the availability of wavelengths, making it less likely that the same wavelength is available throughout the network for connections that traverse longer paths.

The most used and least used schemes introduce communication overhead because they require global network information in order to compute the usage of each wavelength. The first-fit scheme, on the other hand, requires no global information, and since it does not need to order wavelengths in real-time, it has significantly lower computational requirements than either the most used or least used schemes. Another adaptive scheme that avoids the communication and computational overhead of most used and least used is *random* wavelength assignment. With this scheme, the set of wavelengths that are free on a particular path is first determined. Among the available wavelengths, one is chosen randomly (usually with uniform probability) and assigned to the requested lightpath.

We note that in networks in which all OXCs are capable of wavelength conversion, the wavelength assignment problem is trivial: since a lightpath can be established as long as at least one wavelength is free at each link and different wavelengths can be used in different links, the order in which wavelengths are assigned is not important. On the other hand, when only a fraction of the OXCs employ converters (i.e., a sparse conversion scenario), a wavelength assignment scheme is again required to select a wavelength for each segment of a connection’s path that originates and terminates at an OXC with converters. In this case, the same assignment policies discussed above for selecting a wavelength for the end-to-end path can also be used to select a wavelength for each path segment between OXCs with converters.

10.5.3 Performance of Dynamic RWA Algorithms

The performance of a dynamic RWA algorithm is generally measured in terms of the call blocking probability, that is, the probability that a lightpath cannot be established in the network due to lack of resources (e.g., link capacity or free wavelengths). Even in the case of simple network topologies (such as rings) or simple routing rules (such as fixed routing), the calculation of blocking probabilities in WDM networks is extremely difficult. In networks with arbitrary mesh topologies, and/or when using alternate or adaptive routing algorithms, the problem is even more complex. These complications arise from both the link load dependencies (due to interfering lightpaths) and the dependencies among the sets of active wavelengths in adjacent links (due to the wavelength continuity constraint). Nevertheless, the problem of computing blocking probabilities in wavelength routed networks has been extensively studied in the literature, and approximate analytical techniques which capture the effects of link load and wavelength dependencies have been developed in [11, 19, 16].

A detailed comparison of the performance of various wavelength assignment schemes in terms of call blocking probability can be found in [50].

Though important, average blocking probability (computed over all connection requests) does not always capture the full effect of a particular dynamic RWA algorithm on other aspects of network behavior, in particular, *fairness*. In this context, fairness refers to the variability in blocking probability experienced by lightpath requests between the various edge node pairs, such that lower variability is associated with a higher degree of fairness. In general, any network has the property that longer paths are likely to experience higher blocking than shorter ones. Consequently, the degree of fairness can be quantified by defining the *unfairness factor* as the ratio of the blocking probability on the longest path to that on the shortest path for a given RWA algorithm. Depending on the network topology and the RWA algorithm, this property may have a cascading effect which can result in an unfair treatment of the connections between more distant edge node pairs: blocking of long lightpaths leaves more resources available for short lightpaths, so that the connections established in the network tend to be short ones. These shorter connections "fragment" the availability of wavelengths, and thus, the problem of unfairness is more pronounced in networks without converters, since finding long paths that satisfy the wavelength continuity constraint is more difficult than without this constraint.

Several studies [11, 19, 16] have examined the influence of various parameters on blocking probability and fairness, and some of the general conclusions include the following:

- Wavelength conversion significantly affects fairness. Networks employing converters at all OXCs sometimes exhibit orders of magnitude improvement in fairness (as reflected by the unfairness factor) compared to networks with no conversion capability, despite the fact that the improvement in overall blocking probability is significantly less pronounced. It has also been shown that equipping a relatively small fraction (typically, 20-30%) of all OXCs with converters is sufficient to achieve most of the fairness benefits due to wavelength conversion.

- Alternate routing can significantly improve the network performance in terms of both overall blocking probability and fairness. In fact, having as few as three alternate paths for each connection may in some cases (depending on the network topology) achieve almost all the benefits (in terms of blocking and fairness) of having full wavelength conversion at each OXC with fixed routing.

- Wavelength assignment policies also play an important role, especially in terms of fairness. The random and least used schemes tend to "fragment" the wavelength availability, resulting in large unfairness factors (with least used having the worst performance). On the other hand, the most used assignment policy achieves the best performance in terms of fairness. The first-fit scheme exhibits a behavior very similar to most used in terms of both fairness and overall blocking probability, and has the additional advantage of being easier and less expensive to implement.

10.6 CONTROL PLANE ISSUES AND STANDARDIZATION ACTIVITIES

So far we have focused on the application of network design and traffic engineering principles to the control of traffic in optical networks with a view to achieving specific performance objectives, including efficient utilization of network resources and planning of network capacity. Equally important to an operational network are associated control plane issues involved in automating the process of lightpath establishment and in supporting the network design and traffic engineering functions. Currently, a number of standardization activities addressing the control plane aspects of optical networks are underway [51, 52, 53] within the Internet Engineering Task Force (IETF) [54], the Optical Domain Service Interconnection (ODSI) coalition [55], and the Optical Internetworking Forum (OIF) [56]. In this section we review the relevant standards activities and discuss how they fit within the traffic engineering framework; we note, however, that these are ongoing efforts and will likely evolve as the underlying technology matures and our collective understanding of optical networks advances.

Let us return to Figure 10.1 which illustrates the manner in which client subnetworks (IP/MPLS networks in the figure) attach to the optical network of OXCs. The figure corresponds to the vision of a future optical network which is capable of providing a bandwidth-on-demand service by dynamically creating and tearing down lightpaths between client subnetworks. There are two broad issues that need to be addressed before such a vision is realized. First, a signaling mechanism is required at the user-to-network interface (UNI) between the client subnetworks and the optical network control plane. The signaling channel allows edge nodes to dynamically request bandwidth from the optical network, and supports important functions including service discovery and provisioning capabilities, neighbor discovery and reachability information, address registration, etc. Both the ODSI coalition [57] and the OIF [58] have developed specifications for the UNI; the OIF specifications are based on GMPLS.

Second, a set of signaling and control protocols must be defined within the optical network to support dynamic lightpath establishment and traffic engineering functionality; these protocols are implemented at the control module of each OXC. Currently, most of the work on defining control plane protocols in the optical network takes place under the auspices of IETF, reflecting a convergence of the optical networking and the IP communities to developing technology built around a single common framework, namely, GMPLS, for controlling both IP and optical network elements [59]. There are three components of the control plane that are crucial to setting up lightpaths within the optical network, and thus, relevant to traffic engineering (refer to Figure 10.5):

- **Topology and resource discovery.** The main purpose of discovery mechanisms is to disseminate network state information including resource usage, network connectivity, link capacity availability, and special constraints.

- **Route Computation.** This component employs RWA algorithms and traffic engineering functions to select an appropriate route for a requested lightpath.

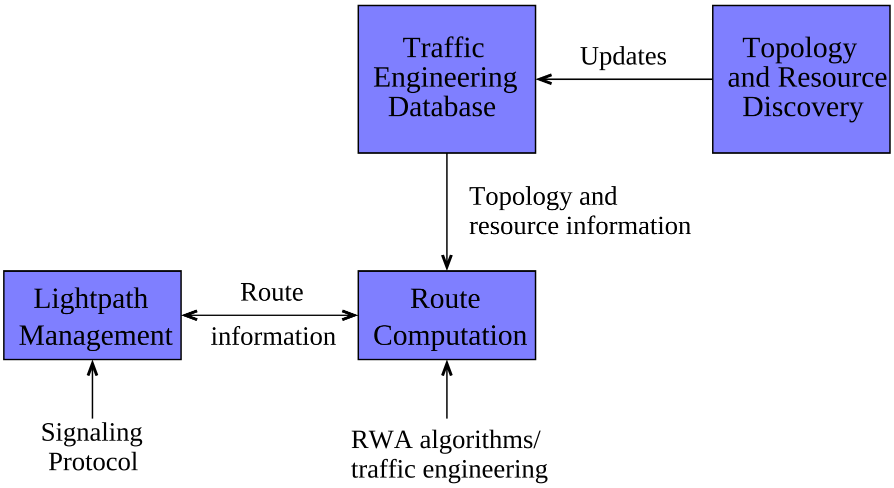

Fig. 10.5 Control plane components

- **Lightpath Management.** Lightpath management is concerned with setup and tear-down of lightpaths, as well as coordination of protection switching in case of failures.

The reader is referred to Chapter 14 for details on the control plane, and to Chapter 16 for details on standards activities. We now briefly discuss some of the standardization activities aimed at defining new protocols or extending existing ones to carry out these three functions.

Topology and resource discovery includes neighbor discovery, link monitoring, and state distribution. The link management protocol (LMP) [60] has been proposed to perform neighbor discovery and link monitoring. LMP is expected to run between neighboring OXC nodes and can be used to establish and maintain control channel connectivity, monitor and verify data link connectivity, and isolate link, fiber, or channel failures. Distribution of state information is typically carried out through link state routing protocols such as OSPF [48]. There are currently several efforts under way to extend OSPF to support GMPLS [61] and traffic engineering [62]. In particular, the link state information that these protocols carry must be augmented to include optical resource information including: wavelength availability and bandwidth, SRLG information (discussed in Section 10.4.3), physical layer constraints (discussed in Section 10.5.1), and link protection information, among others. This information is then used to build and update the optical network traffic engineering database (see Figure 10.5) which guides the route selection algorithm.

Once a lightpath is selected, a signaling protocol must be invoked to set up and manage the connection. Two protocols have currently been defined to signal a lightpath setup: RSVP-TE [63] and CR-LDP [64]. RSVP-TE is based on the resource reservation protocol (RSVP) [65] with appropriate extensions to support traffic engineering, while CR-LDP is an extension of the label distribution protocol (LDP) [66] augmented to handle constraint-based routing. The protocols are currently being extended to support GMPLS [67, 68]. Besides signaling the path at connection time, both protocols can be used to automatically handle the switchover to the protection path once a failure in the working path has occurred.

As a final note, the control plane elements depicted in Figure 10.5 are independent of each other and, thus, separable. This modularity allows each component to evolve independently of others, or to be replaced with a new and improved protocol. As the optical networking and IP communities come together to define standards, the constraints and new realities (e.g., the explosion in the number of channels in the network) imposed by the optical layer and WDM technology will certainly affect our long-held assumptions regarding issues such as routing, control, discovery, etc., which have been developed for mostly opaque electronic networks. As we carefully rethink these issues in the context of transparent (or almost-transparent) optical networks, protocol design will certainly evolve to better accommodate the new technology. Therefore, we expect that the control plane protocols will continue to be refined and/or replaced by new, more appropriate ones. The interested reader should frequently check with the activities within IETF and OIF for the most recent developments.

10.7 SUMMARY

We have discussed the issues arising in the design and operation of optical WDM networks. The emphasis has been on capacity planning and traffic engineering techniques as they apply to different time scales, from resource sizing and topology design to real-time provisioning of lightpaths. The underlying routing and wavelength assignment control problem has been extensively studied and is well-understood within the research community. A number of control plane technologies in support of traffic engineering functions are also being developed at a rapid rate within the standards bodies. The next step is to fully incorporate RWA control algorithms and protocols into optical network elements, initially in field trials and later for deployment in commercial networks, in order to realize a seamless integration of optical networking within the overall Internet infrastructure.

References

1. D. O. Awduche. MPLS and traffic engineering in IP networks. *IEEE Communications*, 37(12):42–47, December 1999.

2. D. Awduche, J. Malcolm, J. Agogbua, M. O'Dell, and J. McManus. Requirements for traffic engineering over MPLS. RFC 2702, September 1999.

3. S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss. An architecture for differentiated services. RFC 2475, December 1998.

4. E. Rosen, A. Viswanathan, and R. Callon. Multiprotocol label switching architecture. RFC 3031, January 2001.

5. B. Davie and Y. Rekhter. *MPLS Technology and Applications*. Morgan Kaufmann Publishers, San Diego, California, 2000.

6. P. Ashwood-Smith *et al.* Generalized MPLS – signaling functional description. IETF Draft <draft-ietf-mpls-generalized-signaling-06.txt>, April 2001. Work in progress.

7. D. H. Su and D. W. Griffith. Standards activities for MPLS over WDM networks. *Optical Networks*, 1(3), July 2000.

8. O. Gerstel, B. Li, A. McGuire, G. N. Rouskas, K. Sivalingam, and Z. Zhang (Eds.). Special issue on protocols and architectures for next generation optical WDM networks. *IEEE Journal Selected Areas in Communications*, 18(10), October 2000.

9. R. Dutta and G. N. Rouskas. A survey of virtual topology design algorithms for wavelength routed optical networks. *Optical Networks*, 1(1):73–89, January 2000.

10. E. Leonardi, M. Mellia, and M. A. Marsan. Algorithms for the logical topology design in WDM all-optical networks. *Optical Networks*, 1(1):35–46, January 2000.

11. Y. Zhu, G. N. Rouskas, and H. G. Perros. A path decomposition approach for computing blocking probabilities in wavelength routing networks. *IEEE/ACM Transactions on Networking*, 8(6):747–762, December 2000.

12. L. Li and A. K. Somani. A new analytical model for multifiber WDM networks. *IEEE Journal Selected Areas in Communications*, 18(10):2138–2145, October 2000.

13. S. Ramamurthy and B. Mukherjee. Survivable WDM mesh networks, part I – protection. In *Proceedings of INFOCOM '99*, pages 744–751, March 1999.

14. S. Ramamurthy and B. Mukherjee. Survivable WDM mesh networks, part II – restoration. In *Proceedings of ICC '99*, pages 2023–2030, June 1999.

15. A. Mokhtar and M. Azizoglu. Adaptive wavelength routing in all-optical netowrks. *IEEE/ACM Transactions on Networking*, 6(2):197–206, April 1998.

16. E. Karasan and E. Ayanoglu. Effects of wavelength routing and selection algorithms on wavelength conversion gain in WDM optical networks. *IEEE/ACM Transactions on Networking*, 6(2):186–196, April 1998.

17. H. Zang, J. P. Jue, and B. Mukherjee. A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks. *Optical Networks*, 1(1):47–60, January 2000.

18. B. Ramamurty and B. Mukherjee. Wavelength conversion in WDM networking. *IEEE Journal Selected Areas in Communications*, 16(7):1061–1073, September 1998.

19. S. Subramaniam, M. Azizoglu, and A. Somani. All-optical networks with sparse wavelength conversion. *IEEE/ACM Transactions on Networking*, 4(4):544–557, August 1996.

20. T. Tripathi and K. Sivarajan. Computing approximate blocking probabilities in wavelength routed all-optical networks with limited-range wavelength conversion. In *Proceedings of INFOCOM '99*, pages 329–336, March 1999.

21. N. Ghani. Lambda-labeling: A framework for IP-over-WDM using MPLS. *Optical Networks*, 1(2):45–58, April 2000.

22. O. Gerstel and R. Ramswami. Optical layer survivability – a service perspective. *IEEE Communications*, 38(3):104–113, March 2000.

23. L. H. Sahasrabuddhe and B. Mukherje. Light-trees: Optical multicasting for improved performance in wavelength-routed networks. *IEEE Communications*, 37(2):67–73, February 1999.

24. D. Papadimitriou *et al.* Optical multicast in wavelength switched networks – architectural framework. IETF Draft <draft-poj-optical-multicast-01.txt>, July 2001. Work in progress.

25. B. Mukherjee. *Optical Communication Networking*. McGraw-Hill, 1997.

26. S. Keshav. *An Engineering Approach to Computer Networking*. Addison-Wesley, Reading, MA, 1997.

27. V. Sharma and E. A. Varvarigos. Limited wavelength translation in all-optical WDM mesh networks. In *Proceedings of INFOCOM '98*, pages 893–901, March 1999.

28. S. L. Hakimi. Steiner's problem in graphs and its implications. *Networks*, 1:113–133, 1971.

29. Y. Xin, G. N. Rouskas, and H. G. Perros. On the design of MPλS networks. Technical Report TR-01-07, North Carolina State University, Raleigh, NC, July 2001.

30. S. Subramaniam, M. Azizoglu, and A. K. Somani. On the optimal placement of wavelength converters in wavelength-routed networks. In *Proceedings of INFOCOM '98*, pages 902–909, April 1998.

31. M. Ali and J. Deogun. Allocation of multicast nodes in wavelength-routed networks. In *Proceedings of ICC 2001*, pages 614–618, 2001.

32. R. Ramaswami and K. N. Sivarajan. Design of logical topologies for wavelength-routed optical networks. *IEEE Journal Selected Areas in Communications*, 14(5):840–851, June 1996.

33. D. Banerjee and B. Mukherjee. A practical approach for routing and wavelength assignment in large wavelength-routed optical networks. *IEEE Journal Selected Areas in Communications*, 14(5):903–908, June 1996.

34. B. Mukherjee *et al.* Some principles for designing a wide-area WDM optical network. *IEEE/ACM Transactions on Networking*, 4(5):684–696, October 1996.

35. Z. Zhang and A. Acampora. A heuristic wavelength assignment algorithm for multihop WDM networks with wavelength routing and wavelength reuse. *IEEE/ACM Transactions on Networking*, 3(3):281–288, June 1995.

36. I. Chlamtac, A. Ganz, and G. Karmi. Lightnets: Topologies for high-speed optical networks. *Journal of Lightwave Technology*, 11:951–961, May/June 1993.

37. S. Banerjee and B. Mukherjee. Algorithms for optimized node placement in shufflenet-based multihop lightwave networks. In *Proceedings of INFOCOM '93*, March 1993.

38. R. Malli, X. Zhang, and C. Qiao. Benefit of multicasting in all-optical networks. In *Proceedings of SPIE*, volume 3531, pages 209–220, November 1998.

39. G. Sahin and M. Azizoglu. Multicast routing and wavelength assignment in wide-area networks. In *Proceedings of SPIE*, volume 3531, pages 196–208, November 1998.

40. J. Strand, A. L. Chiu, and R. Tkach. Issues for routing in the optical layer. *IEEE Communications*, pages 81–96, February 2001.

41. M. Alanyali and E. Ayanoglu. Provisioning algorithms for WDM optical networks. *IEEE/ACM Transactions on Networking*, 7(5):767–778, October 1999.

42. S. Ramamurthy and B. Mukherjee. Survivable WDM mesh networks, part I - protection. In *Proceedings of INFOCOM '93*, pages 744–751, March 1999.

43. B. T. Doshi, S. Dravida, P. Harshavardhana, O. Hauser, and Y. Wang. Optical network design and restoration. *Bell Labs Technical Journal*, pages 58–84, January-March 1999.

44. Y. Qin, K. Sivalingam, and B. Li. Architecture and analysis for providing virtual private networks (VPNs) with QoS over optical WDM networks. *Optical Networks*, 2(2):57–65, March/April 2001.

45. E. Lawler. *Combinatorial Optimization: Networks and Matroids*. Holt, Rinehart and Winston, 1976.

46. D. Bertsekas and R. Gallager. *Data Networks*. Prentice Hall, Inc., Englewood Cliffs, NJ, 1992.

47. X. Zhang, J. Y. Wei, and C. Qiao. Constrained multicast routing in WDM networks with sparse light splitting. *Journal of Lightwave Technology*, 18(12):1917–1927, December 2000.

48. J. Moy. OSPF version 2. RFC 2328, April 1998.

49. A. Chiu, J. Luciani, A. Banerjee, and A. Fredette. Impairments and other constraints on optical layer routing. IETF Draft <draft-ietf-ipo-impairments-00.txt>, May 2001. Work in progress.

50. Y. Zhu, G. N. Rouskas, and H. G. Perros. A comparison of allocation policies in wavelength routing networks. *Photonic Network Communications*, 2(3):265–293, August 2000.

51. Z. Zhang, J. Fu, D. Guo, and L. Zhang. Lightpath routing for intelligent optical networks. *IEEE Networks*, 15(4):28–35, July/August 2001.

52. C. Assi, M. Ali, R. Kurtz, and D. Guo. Optical networking and real-time provisioning: An integrated vision for the next-generation internet. *IEEE Networks*, 15(4):36–45, July/August 2001.

53. S. Sengupta and R. Ramamurthy. From network design to dynamic provisioning and restoration in optical cross-connect mesh networks: An architectural and algorithmic overview. *IEEE Networks*, 15(4):46–54, July/August 2001.

54. The internet engineering task force. http://www.ietf.org.

55. Optical domain service interconnect. http://www.odsi-coalition.com.

56. The optical internetworking forum. http://www.oiforum.com.

57. G. Bernstein, R. Coltun, J. Moy, A. Sodder, and K. Arvind. ODSI functional specification version 1.4. ODSI Coalition, August 2000.

58. User network interface (UNI) 1.0 signaling specification. OIF2000.125.6, September 2001.

59. B. Rajagopalan *et al.* IP over optical networks – a framework. IETF Draft <draft-many-ip-optical-framework-03.txt>, March 2001. Work in progress.

60. J. P. Lang *et al.* Link management protocol (LMP). IETF Draft <draft-ietf-mpls-lmp-02.txt>, September 2001. Work in progress.

61. K. Kompella *et al.* OSPF extensions in support of generalized MPLS. IETF Draft <draft-ietf-ccamp-ospf-gmpls-extensions-00.txt>, September 2001. Work in progress.

62. D. Katz, D. Yeung, and K. Kompella. Traffic engineering extensions to OSPF. IETF Draft <draft-katz-yeung-ospf-traffic-06.txt>, October 2001. Work in progress.

63. D. Awduche *et al.* RSVP-TE: Extensions to RSVP for LSP tunnels. IETF Draft <draft-ietf-mpls-rsvp-lsp-tunnel-08.txt>, February 2001. Work in progress.

64. O. Aboul-Magd *et al.* Constraint-based LSP setup using LDP. IETF Draft <draft-ietf-mpls-cr-ldp-05.txt>, February 2001. Work in progress.

65. R. Braden *et al.* Resource reservation protocol – version 1. RFC 2205, September 1997.

66. L. Andersson, P. Doolan, N. Feldman, A. Fredette, and B. Thomas. LDP specification. RFC 3036, January 2001.

67. P. Ashwood-Smith *et al.* Generalized MPLS signaling – RSVP-TE extensions. IETF Draft <draft-ietf-mpls-generalized-rsvp-te-05.txt>, October 2001. Work in progress.

68. P. Ashwood-Smith *et al.* Generalized MPLS signaling – CR-LDP extensions. IETF Draft <draft-ietf-mpls-generalized-cr-ldp-04.txt>, July 2001. Work in progress.