1

Introduction to Optical Networks

Computer communication started with copper wire as the medium for carrying electrical signals encoding the data to be communicated from one computer to another. Copper as a medium of communication has a number of limitations and, in the last two decades, enormous progress has been made in using alternative media for communication. This section reviews the main developments in the area of communication using optical signals with optical fibers as the medium of communication [30, 88, 89, 172, 222, 286, 324, 370].

1.1 What Is an Optical Network

Optical fibers are essentially very thin glass cylinders or filaments which carry signals in the form of light (optical signals). An optical network connects computers (or any other device which can generate or store data in electronic form) using optical fibers. To facilitate data communication, an optical network also includes other optical devices to generate optical (electrical) signals from electrical (respectively optical) data, to restore optical signals after it propagates through fibers, and to route optical signals through the network. These devices will be briefly reviewed later on in Chapter 2. Optical networks have found widespread use because the bandwidth of such networks using current technology is 50 tera-bits per second. In other words, it is theoretically possible to send 50×10^{12} bits per second using a single fiber. First-generation optical networks simply replaced copper wires with optical fibers. However, there are the following important differences between copper and fiber as communication media:

- optical devices are much more expensive compared to electronic devices, so it is important to optimize the use of optical network resources,
- a number of optical signals at different carrier wavelengths may be simultaneously carried by the same fiber,

- the speed at which optical signals may be communicated is far greater than the speed at which data can be processed by electronic circuits.

Second-generation optical networks [80, 286, 324, 370] take into account these differences and the recent developments in optical devices and network technologies. The technology of using multiple optical signals on the same fiber is called *wavelength division multiplexing* (WDM). In this book, the focus is on second-generation WDM networks, with arbitrary topologies, also known as *mesh* networks.

1.1.1 Important Advantages of WDM Optical Networks

As mentioned already, a key advantage of optical technology is speed. Internet traffic has been increasing rapidly since the early 1990s, doubling every six months. The increasing volume of business-to-business communication is another reason for the rapid increase in the demand for additional speed. New applications for data transfer are also a notable development. An example is *video-on-demand*, where users can select a video to be downloaded from some central server. Other important widespread applications include downloading music and exchanging digital pictures. There are other important advantages of optical technology, as follows [286]:

- *Low signal attenuation.* As a signal propagates through fibers, the signal strength goes down at a low rate (0.2 db/km). This means that the number of optical amplifiers needed is relatively small.
- *Low signal distortion.* As a signal is sent along a fiber optic network, it degrades with respect to shape and phase. Signal regenerators are needed to restore the shape and timing. Low signal distortion means that signal regeneration is needed infrequently.
- *Low power requirement.*
- *Low material usage.*
- *Small space requirements.*
- *Low cost.*

1.1.2 Key Terminology in WDM Optical Networks

In an optical network, all possible sources or destinations of data (typically computers) are called *end nodes*. Another important component of an optical network is the *optical router*. Every router node has a number of input (output) fibers, each carrying one or more incoming (outgoing) optical signals. The purpose of a router is to direct each incoming optical signal to an appropriate outgoing fiber. An example of a simplified model of a second-generation optical network is shown in Figure 1.1, where a circle represents an end node, a rectangle represents a router node and a directed line represents a fiber. These fibers are unidirectional and the arrow on the line gives the direction in which optical signals can flow. Such a diagram is called a *physical topology* since it shows the major physical components of the network. A *lightpath* is an optical connection from one end node to another, used to carry data in the form of encoded optical signals. Such a lightpath always starts from an end node, traverses a number of fibers and router nodes, and ends in another end node. Figure 1.2 shows a number of lightpaths using the physical topology in Figure 1.1. The lightpath L_1, shown using a dashed line, starts from end node E_1, passes through router nodes R_1, R_2, R_3, and terminates in end node E_3. The lightpath L_2, shown using a dotted line, starts from end node E_2, passes through router nodes R_2, R_3, R_4, and terminates in end node E_4. The other lightpaths have routes through the physical topology, as shown in Figure 1.2. Since the lightpaths determine which end nodes can directly communicate with each other, once the lightpaths are set up, the physical topology is irrelevant for determining a strategy for communication. It is convenient to view the lightpaths as edges of a directed graph[1] G_L where the nodes of G_L are the end nodes of the physical topology. Such a graph is called the *logical topology* (some researchers use the term *virtual topology*) of an optical network and the edges of such a graph are called *logical edges*. The logical topology corresponding to the lightpaths shown in Figure 1.2 is shown in Figure 1.3.

To make a distinction between an edge in a physical topology, representing a fiber, and an edge in the logical topology, representing a lightpath, the following notation will be used.

An edge from node N_x to node N_y in the physical topology will be shown by a single arrow $N_x \rightarrow N_y$. Here N_x (N_y) may be either an end node or a router node. An edge from end node E_i to end node E_j, representing a lightpath from E_i to E_j, in the logical topology will be depicted by a double arrow $E_i \Rightarrow E_j$.

It should be noted that the actual route of a lightpath through the physical topology is irrelevant as far as the logical topology is concerned. For instance, using the same physical topology shown in Figure 1.1, in Figure 1.4 lightpaths L_1 and L_2 have routes $E_1 \rightarrow R_1 \rightarrow R_4 \rightarrow R_3 \rightarrow E_3$ and $E_2 \rightarrow R_2 \rightarrow R_1 \rightarrow R_4 \rightarrow E_4$ respectively. The logical topology remains the same as that shown in Figure 1.3.

For a given set of end nodes, different physical topologies, representing different connections between end nodes and router nodes could give the same logical topology. For instance, if the physical topology, shown in Figure 1.5, is used, it is quite simple to set up lightpaths L_1 from end node E_1 to end

[1] In many networks, whenever there is a directed edge from end node E_i to end node E_j, there is also a directed edge from end node E_j to end node E_i. In such networks, an undirected edge may be used between end nodes, replacing a pair of directed edges, for simplicity.

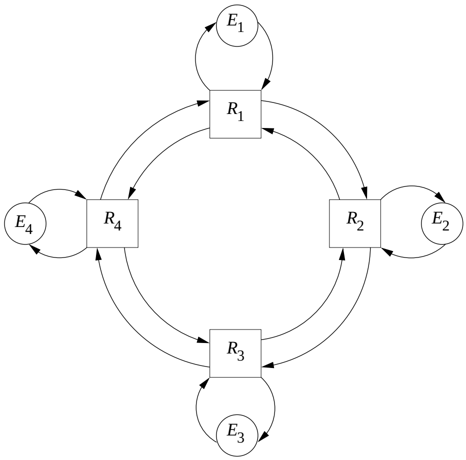

Fig. 1.1. The physical topology of a typical WDM network with four end nodes $E_1, \ldots, E_4$ and four routers $R_1, \ldots, R_4$

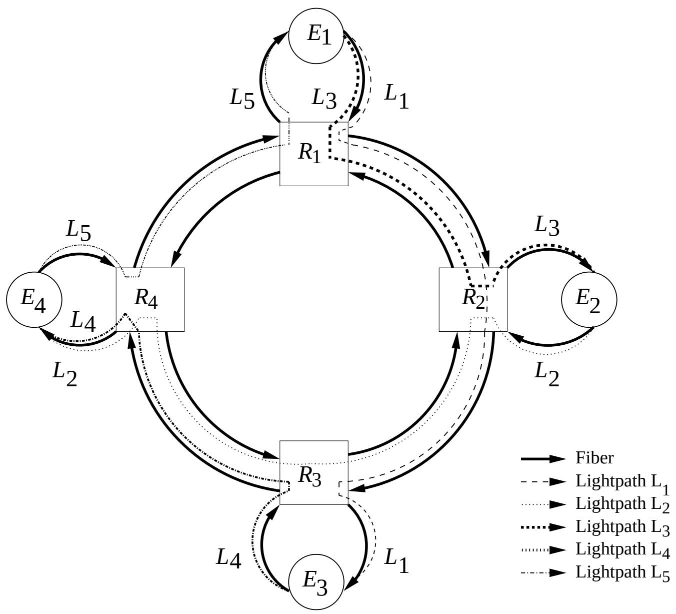

Fig. 1.2. Some lightpaths on the physical topology shown in Figure 1.1

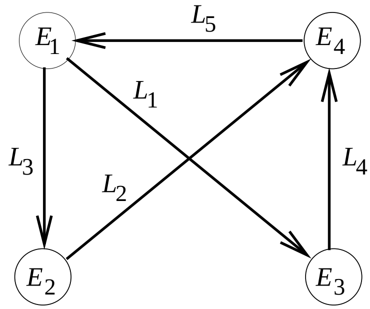

Fig. 1.3. Logical topology G_L corresponding to the lightpaths shown in Figure 1.2

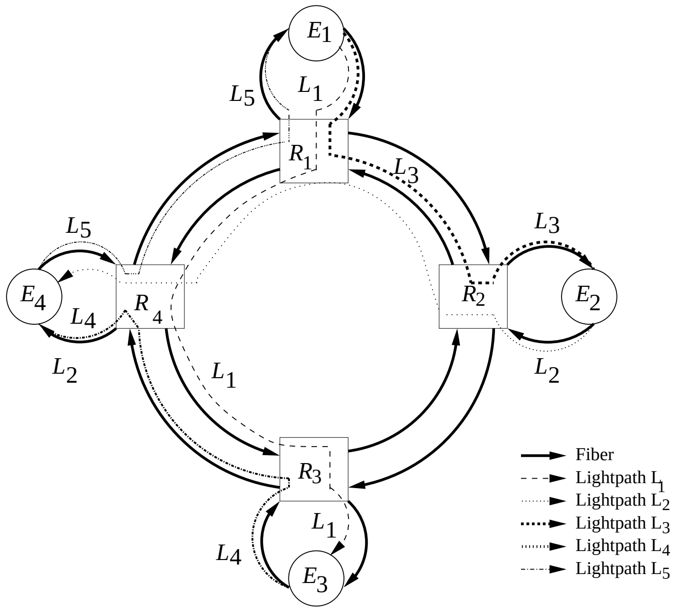

Fig. 1.4. New routes for lightpaths L_1 and L_2 on the physical topology shown in Figure 1.1

node E_3, ... , lightpath L_5 from end node E_4 to end node E_1 giving exactly the logical topology shown in Figure 1.3.

Exercise 1.1. Show a routing for the lightpaths $L_1, \ldots, L_5$ through the physical topology shown in Figure 1.5. □

1.1.3 Data Communication in a WDM Optical Network

If there is a lightpath from end node E_x to end node E_y, then this lightpath is an obvious candidate (but not the only candidate) for data communication

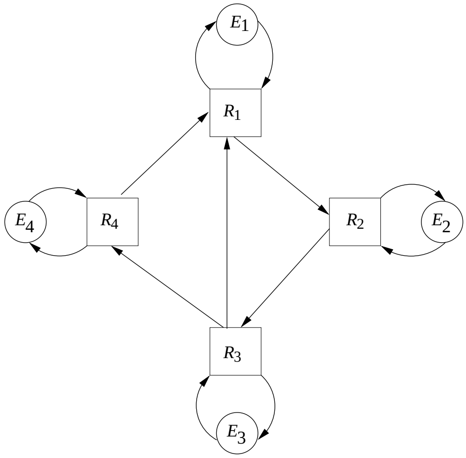

Fig. 1.5. A different physical topology

from E_x to E_y. If there is no lightpath from E_x to E_y, then a communication from E_x to E_y is possible only if there are end nodes E_{i_1}, E_{i_2}, ..., E_{i_m}, such that there is a lightpath from E_x to E_{i_1}, a lightpath from E_{i_1} to E_{i_2}, ..., a lightpath from E_{i_m} to E_y. End node E_x has to communicate the data to E_{i_1}, E_{i_1} then communicates the data to $E_{i_2}, \ldots, E_{i_m}$ finally communicates the data to E_y. It is convenient to use the notion of a directed path (or dipath), in the logical topology G_L, from one end node to another, for such a communication. A directed path through a logical topology is called a *logical path*. The logical path used by the above communication from E_x to E_y is $E_x \Rightarrow E_{i_1} \Rightarrow E_{i_2} \Rightarrow \ldots \Rightarrow E_{i_m} \Rightarrow E_y$. In general, there may be a number of logical paths between any pair of end nodes. If data has to be communicated through an optical network from source E_x to destination E_y, one or more appropriate logical paths from E_x to E_y have to be selected for this communication.

For example, in the logical topology shown in Figure 1.3, to send data from end node E_1 to end node E_3, lightpath L_1 corresponding to the logical edge $E_1 \Rightarrow E_3$ is the only choice. To send data from end node E_1 to end node E_4, one possible logical path is $E_1 \Rightarrow E_2 \Rightarrow E_4$. If this logical path is used, the lightpath L_3 will be used to send the data from the source end node E_1 to the intermediate end node E_2. The lightpath L_2 will be used to send data from the intermediate end node E_2 to the destination end node E_4. In this process, the data in the source end node E_1 is converted into optical form and is sent to end node E_2 using lightpath L_3. At end node E_2 the data is extracted from the lightpath L_3 so that the data is again in electronic form. Then the data is again converted into optical form and is communicated from end node E_2 to end node E_4 using lightpath L_2 where the data is again converted back to electronic form.

Exercise 1.2. Consider the logical topology shown in Figure 1.6.

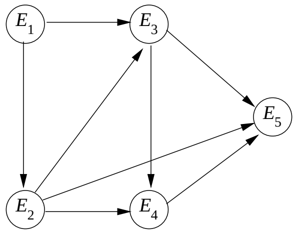

Fig. 1.6. The logical topology of a typical WDM network

i) Find all the logical paths from end node E_1 to end node E_5.
ii) In this logical topology, is it possible to communicate from any end node to any other end node? If not, then is it possible to add one lightpath so that any end node can communicate with any other end node?

□

1.2 Categorizations of WDM Networks

WDM networks may be categorized in a number of different ways. Three important ways of categorization of WDM networks are described below.

1.2.1 Broadcasting Networks and Wavelength-Routed Networks

WDM networks may be broadly classified as *wavelength-routed* networks and *broadcast-and-select* networks.

In a *wavelength-routed network*, the wavelength of the optical signal and the fiber it is using determine the subsequent path used (hence the name wavelength-routed) by the signal. Since each optical signal is sent along a specified path and not broadcast to all nodes in the network, the power requirement of such a network is lower than that of a broadcast-and-select network. This type of network may contain a large number of end nodes but is more complex and expensive than a broadcast-and-select network. The network shown in Figure 1.2 is a wavelength-routed network since the end nodes communicate using lightpaths which are routed from their sources to their respective destinations based on their wavelengths. For instance, the routers R_1, R_2, or R_3 have been set up in such a way that, when signals using the wavelength of lightpath L_1 are received by router nodes R_1, R_2, and R_3, they are sent forward to router node R_2, router node R_3, and end node E_3 respectively. The same is true for all the lightpaths in this network.

In a *broadcast-and-select network* for unicast[2] communication, the source end node selects an appropriate wavelength λ_p and broadcasts the data to be transmitted to all end nodes in the network using the wavelength λ_p. The receiver at the destination end node must be tuned to the wavelength λ_p while the receivers at all other end nodes are tuned to wavelengths different from λ_p. The net result is that the data is detected and processed only at the destination node. A typical broadcast-and-select network is shown in Figure 1.7. In this figure, each end node has one transmitter and one receiver. For convenience, the transmitters in each end node are shown on one side and the receivers on the other side. Each end node has a connection to a device called a *passive star* (to be reviewed in Section 2.3.2). The i^{th} end node E_i has a transmitter t_i and a receiver r_i, and generates a signal using wavelength λ_i, for all $i, 1 \leq i \leq 4$. The receivers in the end nodes are tuned, as shown by the wavelength above the rectangle depicting each end node. As shown in the figure, the passive star receives signals using wavelengths $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ from end nodes E_1 to E_4 and broadcasts all these signals to all the receivers, r_1, r_2, r_3, and r_4 in the corresponding end nodes. Only the receiver at end node E_1, being tuned to wavelength λ_3, selects this signal. Similarly, each of the remaining receivers only selects one signal. A broadcast-and-select network is simple and easy to implement but the size of the network is limited due to the requirement that the signal has to be broadcast to all end nodes. This type of network will be discussed in more detail in Chapter 6.

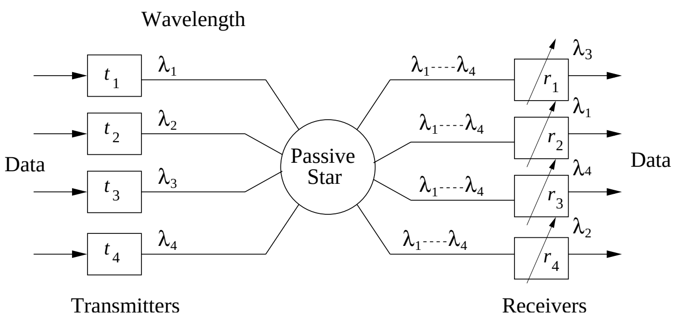

Fig. 1.7. A broadcast-and-select network

1.2.2 Static and Dynamic Lightpath Allocation

There are two approaches for deciding a strategy for data communication in a wavelength-routed network. The more standard approach is to set up lightpaths on a semipermanent basis so that, once the lightpaths are set up to

[2] Unicast communication means communication from one source to only one destination.

handle the expected volume of data between the ordered pairs of end nodes, the lightpaths will continue to exist for a relatively long period of time (weeks or months). This approach is called *static lightpath allocation*. When the communication pattern changes sufficiently, the existing lightpaths will be taken down and new lightpaths will be set up to handle the changes in traffic. In the second approach, called *dynamic lightpath allocation*, lightpaths are set up on demand and, when a communication is over, the corresponding lightpath is taken down (i.e., the lightpath no longer remains operational).

1.2.3 Single-hop and Multi-hop WDM Networks

In order to make the most economical use of a network, maximizing the throughput of the network is important. The cost of devices for optoelectronic conversion is high, and the speed of processing in electronic circuits is much less than the speed at which optical signals propagate through fibers. Minimizing the number of stages of electronic processing of data in WDM networks is therefore a crucial objective. It is highly desirable that the data be kept in optical form from the time it leaves the source end node until it reaches its destination end node. In a *single-hop* network, all data communication involves a path length of one logical edge. In other words, exactly one lightpath is involved in each communication. Single-hop networks are also called *all-optical networks*, since the communication is always in the optical domain [81, 178, 286, 322, 324].

In a network with $\mathcal{N}_E$ end nodes, the number of end node pairs is $\mathcal{N}_E \cdot (\mathcal{N}_E - 1)$ so that the number of lightpaths, and hence the amount of optical resources needed in a single-hop optical network, becomes uneconomical when $\mathcal{N}_E$ is relatively large. In a *multi-hop* network, [1, 157, 236, 286, 324], some data communication involves more than one lightpath. In the network shown in Figure 1.1, with lightpaths shown in Figure 1.2, the communication from end node E_1 to end node E_4 involves two lightpaths (either L_3 and L_2 or L_1 and L_4). The network is therefore a multi-hop network. If, in the network shown in Figure 1.1, there were lightpaths between every ordered pair of end nodes, then it would become a single-hop network.

1.3 Important Problems in WDM Networks and Solution Approaches

Some important problems in the area of WDM network design that have been investigated are as follows:

- How to set up a lightpath, or a set of lightpaths, to make optimum use of the network resources.
- How to best define the logical topology of a multi-hop network.

- Given a logical topology, what is the optimum strategy to handle all the requests for data communication in the network.
- How to handle faults in the network.

The problems mentioned above are (mathematical) optimization problems [377], and are often tackled using *linear program* formulations or using *combinatorial optimization*. The idea of a *linear program*[3] (LP) is to optimize some objective function, subject to linear constraints of a number of decision variables [377] appearing in the problem formulation. If, in a problem formulation, some of the variables are constrained to have nonnegative integer values and the others are *continuous* variables, each capable of having any nonnegative value, the formulation is called a *mixed integer linear program* (MILP) [377]. In most of the optical network design problems considered in this book, a mixed integer linear program is needed for an exact solution. If some of these integer variables are required to have a value of either 0 or 1, such variables are called *binary* variables. A major problem with the MILP formulations is the fact that they are NP-complete. This means that, for optical networks, as the size of the network grows, the time required to solve the MILP formulation grows exponentially. Even for small problems, the actual time to solve them can be unacceptably large. For this reason, heuristics are often used to solve the problems within a reasonable amount of time, giving solutions that are "good" but not necessarily optimal. In this book, heuristics as well as MILP formulations will be covered.

1.4 A Typical Problem in Multi-hop Wavelength-Routed Network Design

As an informal introduction to the above topics, the problems in designing a multi-hop wavelength-routed network using static lightpath allocation are given below. The design takes into account the following:

- the physical topology of the network,
- the optical hardware available at each end node which determines how many lightpaths may originate from or end at that end node,
- the characteristics of the fibers, which, for example, determine how many lightpaths may be allowed in a single fiber,
- the traffic requirements between each pair of end nodes.

It is convenient to represent the traffic requirements in the form of a matrix $T = [t(i,j)]$, often called a *traffic matrix*. The entry $t(i,j)$ in row i and column j of traffic matrix T denotes the amount of traffic from end node E_i to E_j, $i \neq j$. The unit may be specified in a number of ways. Some popular units for specifying data communication are

[3] A review of the simplex method to solve an LP appears in Appendix 1.

- megabits/second (Mbps),
- gigabits/second (Gbps),
- the amount of data that may be carried by a single lightpath,
- the signal rate, using the *Optical Carrier level* notation (OC-n), where the base rate (OC-1) is 51.84 Mbps and OC-n means $n \times 51.84$ Mbps.

A typical traffic matrix T, for the network in Figure 1.1, is shown below. Here the capacity of a lightpath is used as the unit.

$$T = \begin{bmatrix} 0.00 & 0.30 & 0.10 & 0.30 \\ 0.20 & 0.00 & 0.30 & 0.20 \\ 0.35 & 0.10 & 0.00 & 0.30 \\ 0.00 & 0.20 & 0.10 & 0.00 \end{bmatrix}$$

The entry $t(1,4)$ is 0.30, meaning that the amount of data to be communicated from end node E_1 to end node E_4 is 0.30 units. If a single lightpath can carry data at the rate of 10.0 Gbps[4] (i.e., OC-192), the expected data communication rate from end node E_1 to end node E_4 is 3.0 Gbps.

Designing a multi-hop network is complex due to the following considerations:

- In general, there are numerous choices for the routes of each of the lightpaths through the physical topology.
- There are restrictions on the properties of lightpaths. For instance, as noted above, the number of lightpaths allowed on a fiber is limited by the characteristics of the fibers.
- In general, a very large number of logical topologies may be mapped to the same physical topology. Given a set of traffic requirements, possibly in the form of a traffic matrix T, some logical topologies are better in the sense of utilizing network resources[5] better.
- There may be multiple logical paths between a given pair of end nodes. The traffic between such a pair of end nodes may be shared by some or all of these logical paths. For example, in the logical topology shown in Figure 1.3 there are two logical paths ($E_1 \Rightarrow E_2 \Rightarrow E_4$ and $E_1 \Rightarrow E_3 \Rightarrow E_4$) between end nodes E_1 and E_4. In traffic matrix T, the traffic between end nodes E_1 and E_4 is 3.0 Gbps. Part of this traffic may be routed using the logical path $E_1 \Rightarrow E_2 \Rightarrow E_4$ and the rest using the logical path $E_1 \Rightarrow E_3 \Rightarrow E_4$.

The objective of the design is to optimize network performance. A very useful metric for this is the traffic on the logical edge carrying the maximum

[4] The amount of data that can be carried by a lightpath depends on the technology used and is likely to increase in the future. Unless otherwise stated, it is assumed that all lightpaths in a network have the same capacity for data communication (i.e., 10.0 Gbps in this example).

[5] The question of what are the "network resources" and what is meant by "better utilization" will be discussed in later chapters.

traffic. For reasons that will be discussed later, an important objective is to minimize this traffic. One technique to solve such optimization problems is *linear programming*. To specify this problem, integer variables are needed. A detailed discussion of this problem appears later on, but, informally, any ordered pair of end nodes may have a lightpath between them and each lightpath may choose any valid path, in the physical topology, from its source to its destination. The presence (or absence) of a lightpath requires a 0/1 variable and the use of a particular edge in the physical topology in a given lightpath also requires a 0/1 variable. In such a representation, for nontrivial networks, the number of integer variables becomes extremely large and the problem becomes computationally intractable.

To address this situation, it may be necessary to use techniques that are faster but may give suboptimal solutions. For example, it is possible to substantially simplify the problem by decoupling the problem of finding a logical topology from the problem of routing through the logical topology to minimize the traffic on the logical edge carrying the maximum traffic in the network, but at the expense of potentially worse solutions. The problem of finding a logical topology may be considerably simplified by using a heuristic.[6] In this case, the problem of routing through the logical topology reduces to that of a *multi-commodity network flow problem* [4] (to be discussed in detail in Section 7.3) that can be specified using straightforward linear programming techniques. Even so, it is interesting to note that each nonzero entry in the traffic matrix corresponds to a commodity in the multi-commodity network flow model. Since most pairs of end nodes are expected to have some traffic between them, the number of such nonzero entries in a traffic matrix is expected to be $O(\mathcal{N}_E{}^2)$, where $\mathcal{N}_E$ is the number of end nodes in the network. In other words, the number of commodities increases rapidly with $\mathcal{N}_E$. For example, if $\mathcal{N}_E = 30$, the number of nonzero entries in the traffic matrix (commodities in the multi-commodity network flow model) is almost 900. This poses major computational challenges requiring appropriate operations research techniques that will be discussed later on in Section 7.3.

Exercise 1.3. Consider the logical topology shown in Figure 1.3 and the traffic matrix T given above. You must show how each entry in the traffic matrix $t(i,j), t(i,j) \neq 0$, for all $i, j, 1 \leq i, j \leq 4$, may be communicated from its source E_i to its destination E_j. When determining the strategy to communicate traffic $t(i,j)$, you are permitted to split the data into smaller parts. For example, for the traffic $t(1,4) = 0.3$ from E_1 to E_4, there are two logical paths to consider: $E_1 \Rightarrow E_2 \Rightarrow E_4$ and $E_1 \Rightarrow E_3 \Rightarrow E_4$. Part of the total traffic of 3 Gbps from E_1 to E_4 may be sent using the logical path $E_1 \Rightarrow E_2 \Rightarrow E_4$ and the rest using the path $E_1 \Rightarrow E_3 \Rightarrow E_4$. In general, each lightpath will be used to carry traffic from multiple source-destination pairs. For instance, the lightpath from E_1 to E_2 will also be used to carry the traffic

[6] One problem with many of the heuristics is that there is no guarantee about the quality of the solution.

from E_1 to E_2. Make sure that the traffic on each lightpath is less than 1.0, the capacity of the lightpath. Indicate how much traffic is being carried by each lightpath. □

Exercise 1.4. Consider again the logical topology shown in Figure 1.3 and the traffic matrix $T_1 = [t_1(i,j)]$ given below.

$$T_1 = \begin{bmatrix} 0.00 & 0.30 & 0.10 & 0.50 \\ 0.20 & 0.00 & 0.30 & 0.20 \\ 0.35 & 0.10 & 0.00 & 0.30 \\ 0.30 & 0.20 & 0.10 & 0.00 \end{bmatrix}$$

Is it possible to ensure that each entry in the traffic matrix $t_1(i,j), t_1(i,j) \neq 0$, for all $i, j, 1 \leq i, j \leq 4$, may be communicated from its source E_i to its destination E_j on this logical topology?

If traffic T_1 cannot be handled, can you add one lightpath to the network (in other words, add one logical edge to the logical topology shown in Figure 1.3) so that this traffic may be handled?

The solution is not unique and any logical edge which would give a valid routing such that no logical edge carries a total traffic of more than 1.0 is a valid solution. □

1.5 Structure of the Book

In Chapter 2, major optical devices will be briefly summarized. The purpose of this chapter is to familiarize the readers with the terminology used in describing optical hardware and its purpose. Readers interested in more details on how the optical hardware works will be directed to relevant books and papers.

Chapter 3 gives an overview of different types of WDM networks and some important problems in WDM network design. This chapter gives more details about broadcast-and-select networks, wavelength-routed networks, single-hop and multi-hop networks, the route and wavelength assignment (RWA) problem, networks with static and dynamic lightpath allocation, and networks with wavelength conversion.

Chapters 4 and 5 cover different ways of setting up lightpaths. In Chapter 4, the RWA problem is modeled as a graph-coloring problem, thus allowing tools from graph theory to be applied. The hardness of the problem is established, and bounds for general digraphs and general requests are derived. The RWA problem for specific types of requests and for specific networks is also considered.

In Chapter 5, a number of techniques for route and wavelength assignment (RWA) in different types of networks are considered. This includes networks with and without wavelength converters, networks with static and dynamic lightpath allocation, and ring networks. For static lightpath allocation, since the lightpaths, once established, will persist for a considerable time, it is worthwhile to spend some time in finding the "best" allocation. For this purpose MILPs have been proposed. Heuristic solutions for handling this problem within a reasonable time have also been suggested. It is important to find a quick RWA in a lightpath allocation scenario. Some heuristics and distributed algorithms for dynamic RWA are also covered in this chapter.

Chapter 6 describes a logical topology design suitable for broadcast-and-select networks. This topology is interesting in that additional nodes can be added relatively easily, with only small changes to the network.

Chapter 7 covers the design of the logical topology to handle a set of traffic requirements. This problem has been discussed briefly in Section 1.4. MILP formulations are proposed to solve the problem. Such a formulation is computationally intractable except for very small networks. Heuristic solutions are also proposed to tackle this problem. Traffic routing over large logical topologies is a problem in itself and needs special techniques.

In view of the fact that lightpaths typically carry 10 Gbps each (and may carry even 40 Gbps each in the future) and cover all the continents, it is easy to see that optical networks are quite vulnerable to faults. The most common fault is due to fiber cuts, for instance, a fiber may be accidentally cut during earth-moving operations. Chapter 8 discusses different types of faults and how to handle them. The problems are complicated and solutions have been proposed using MILP and heuristics.

The capacity of a lightpath is much larger than a typical request for data communication from an individual user. Chapter 9 discusses *traffic grooming* techniques for efficiently utilizing the large data communication capability of lightpaths to handle many requests for data communication at relatively lower rates.