Chapter 1

WDM TECHNOLOGY AND ISSUES IN WDM OPTICAL NETWORKS

1.1 INTRODUCTION

The influence of "networking" on the organization of computer systems has been tremendous, especially in the last 20 years. The old model of a single computer catering to the computational needs of an organization (company or university) has been replaced by one in which a number of separate but interconnected computers carry out the job. Broadly speaking, a computer network is an interconnected (via copper cable, fiber optics, microwaves, or satellites) collection of independent computers that aids communication in numerous ways. Apart from providing a good communication medium, sharing of available resources (programs and data on a computer are available to anyone on the network, without regard to the physical location of the computer and user), improved reliability of service (because of the presence of multiple computers), and cost-effectiveness (as small computers have better price/performance ratio than large ones) are some of the advantages of networking. From the time the ARPANET[1] was conceived to the current-day high-speed networks, the design and technology associated with this (computer networks) field have come a long way. The need for error-free, high-bandwidth communication channels has been on the rise. The services provided by computer networks include remote information access (communication between a person and

[1] The ARPANET, sponsored by U.S. Department of Defense, was the first major effort at developing a network to interconnect computers over a wide geographical area in the late 1960s.

a remote database—for example, World Wide Web browsing) and electronic mail (person-to-person communication) used by millions of people around the globe. The explosive growth of the Internet and bandwidth-intensive applications such as video-on-demand (for example, selecting a movie located at some remote site and watching it online) and multimedia conferencing (which requires setting up high-bandwidth connections among different people, for a virtual meeting, and guaranteeing the desired quality-of- service [QoS] levels—high bandwidth, low latency, and reasonable packet loss rate—for the virtual meeting) require high-bandwidth transport networks whose capacity (bandwidth) is much beyond what current high-speed networks such as asynchronous transfer mode (ATM)[2] networks can provide. Thus, a continuous demand for networks of high capacities at low costs is seen now. This can be achieved with the help of optical networks, as the optical fiber provides an excellent medium for transfer of huge amounts of data (nearly 50 terabits per second [Tb/s]). Apart from providing such huge bandwidth, optical fiber has low cost (approximately $0.30 per yard), extremely low bit error rates (fractions of bits that are received in error, typically 10^{-12}, compared to 10^{-6} in copper cable), low signal attenuation (0.2 decibels per kilometer [dB/km]), low signal distortion, low power requirement, low material use, and small space requirement. In addition, optical fibers are more secure, compared to copper cables, from tapping (as light does not radiate from the fiber, it is nearly impossible to tap into it secretly without detection) and are also immune to interference and crosstalk. Optical networks, employing wavelength division multiplexing (WDM), is seen as the technology of the future for a variety of other reasons which we will mention in Section 1.3.

1.2 OPTICAL NETWORKS

Optical networks (in which data is converted to bits of light called *photons* and then transmitted over fiber) are faster than traditional networks (in which data is converted to *electrons* that travel through copper cable) because photons weigh less than electrons, and further, unlike electrons, photons do not affect one another when they move in a fiber (because they have no electric charge) and are not affected by stray photons outside the fiber. Light has higher frequencies and hence shorter wavelengths, and therefore more "bits" of information can be contained in a length of fiber versus the same length of copper.

[2]ATM networks are connection-oriented and require a connection set up prior to transfer of information from a source to a destination. All information to be transmitted—voice, data, image, and video—is first fragmented into small, fixed-size packets known as cells. These cells are then switched and routed using packet switching principles (see Appendix B).

Optical glass fibers based on the principle of *total internal reflection*, which was well known in the 1850s, were developed for endoscopes early in the 1900s. The use of fiber glass for communication was first proposed by Kao and Hockham in 1966 [76]. The manufacture of optical fiber began in 1970s. A variety of optical networks came into existence in the late 1980s and early 1990s which used optical fiber as a replacement for copper cable to achieve higher speeds. Computer interconnects such as ESCON (Enterprise Serial Connection), Fiber Channel, and HiPPI (High Performance Parallel Interface), for interconnecting computers to other computers or peripheral systems, use low bit-rate optical components which are inexpensive. FDDI (Fiber Distributed Data Interface) uses dual, fiber optic token rings to provide 100–200 megabits per second (Mb/s) transmission between workstations. SONET/SDH (Synchronous Optical NETwork in North America, Synchronous Digital Hierarchy in Europe and Asia)[3]— which forms the basis for current high-speed backbone networks and one of the most successful standards in the entire networking industry—allows seamless interworking of fibers up to OC-192 rate of about 10 gigabits per second (Gb/s). (OC-n [Optical Carrier-n] specifies an electronic data rate of $n \times 51.84$ Mb/s approximately; so OC-48 and OC-192 correspond to approximate data rates of 2.5 Gb/s and 10 Gb/s, respectively. OC-768 [40 Gb/s] is the next milestone in highest realizable electronic communication speed.)

1.2.1 Optical Fiber Principles

Optical fiber consists of a very fine cylinder of glass (core) through which light propagates. The core is surrounded by a concentric layer of glass (cladding) which is protected by a thin plastic jacket as shown in Fig. 1.1(a). The core has a slightly higher index of refraction than the cladding. The ratio of the indices of refraction of the cladding and the core defines a *critical angle*, θ_c. What makes fiber optics work is *total internal reflection*: when a ray of light from the core approaches the core-cladding surface at an angle less than θ_c, the ray is completely reflected back into the core (see Fig. 1.1[b]).

Since any ray of light incident on the core cladding surface at an angle less than θ_c (critical angle) is reflected internally, many different rays of light from the core will be bouncing at different angles. In such a situation, each ray is said to have a different mode and hence a fiber having this property is called a *multimode fiber* (see Fig. 1.2[a]). Multiple modes cause the rays to interfere with each other, thereby

[3] A set of standards for transmitting digital information over optical networks (see Appendix C).

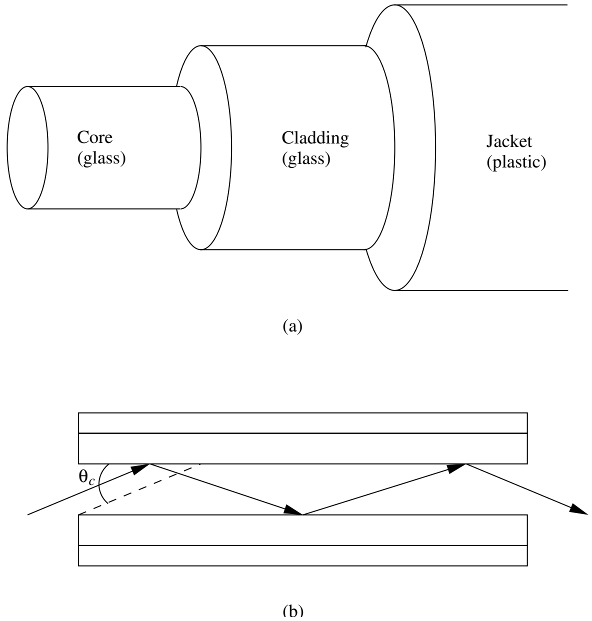

Figure 1.1 (a) Optical fiber. (b) Reflection in fiber.

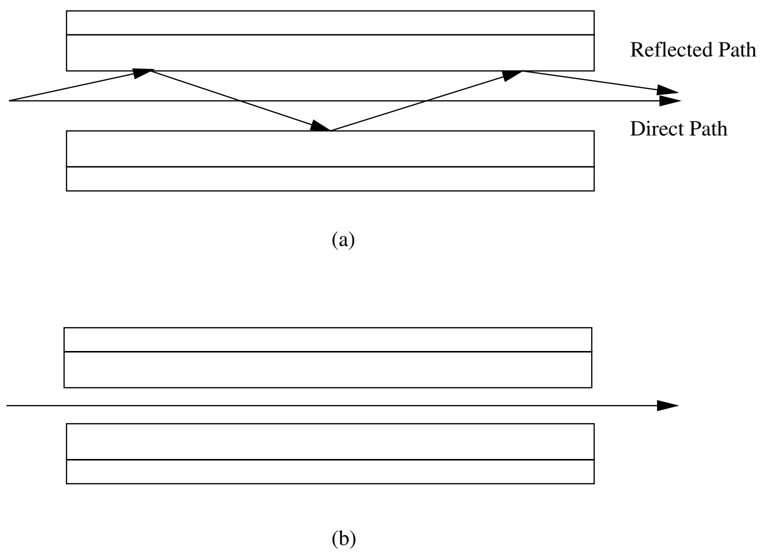

Figure 1.2 (a) Multimode fiber (multiple rays follow different paths). (b) Single-mode fiber (only direct path propagates in fiber).

limiting the maximum bit rates that are achievable using a multimode fiber. If the diameter of the core is made very narrow, the fiber acts like a waveguide, and the light can travel in a straight line along the center axis of the fiber. A fiber having this property is called a *single-mode fiber* (see Fig. 1.2[b]). Single-mode fibers can transmit data at several gigabits per second over hundreds of kilometers and are more expensive. In multimode fibers the core is 50 microns (1 micron is 10^{-6} meter) in diameter, whereas in single-mode fibers the core is 8 to 10 microns.

1.2.2 Optical Transmission System

An optical transmission system has three basic components—transmitter, transmission medium, and receiver—as shown in Fig. 1.3. The transmitter consists of a light source (laser or LED [light-emitting diode]) that can be modulated according to an electrical input signal to produce a beam of light which is transmitted into the optical fiber—the transmission medium. Typically the binary information sequence is converted into a sequence of on/off light pulses which are then transmitted into the optical fiber medium. At the receiver, the on/off light pulses are converted back to an electrical signal by an optical detector. Thus we have a *unidirectional* transmission system (operating only in one direction) which accepts an electrical signal, converts and transmits it by light pulses through the medium, and then reconverts the light pulses to an electrical signal at the receiving end.

Attenuation in fiber leads to loss of signal power as the signal propagates over some distance. Optical fiber had an attenuation loss of 20 dB/km when it was invented in 1970, but within 10 years fibers with a loss of 0.2 dB/km had be-

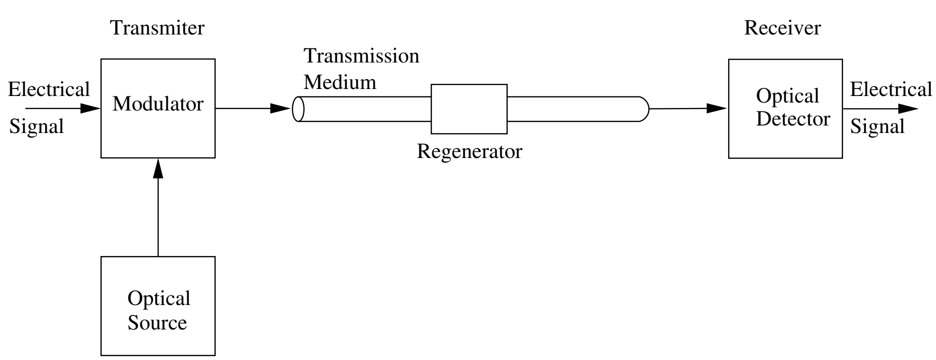

Figure 1.3 Optical transmission system.

come available. The attenuation in decibels is computed as 10 $\log_{10}$(transmitted power/received power). The attenuation of light through fiber depends on the wavelength used. Figure 1.4 shows the attenuation in decibels per (linear) kilometer of fiber. As can be seen from this figure, there are three low-loss windows (bands) centered at 0.85, 1.30, and 1.55 microns. Early optical fiber transmission systems (in the 1970s) operated in the first, that is, 0.85-micron band at bit rates in the tens of megabits per second and used relatively inexpensive LEDs as the light source. Present ones use laser sources and operate in the 1.30- and 1.55-micron bands, achieving rates of gigabits per second. Attenuation is primarily due to impurities (water vapor) in the fiber glass and *Rayleigh scattering* (when the medium is not absolutely uniform, it causes small fluctuations in the refractive index, which in turn causes the light to be scattered and thereby attenuating the propagating signal). To overcome attenuation, electronic regenerators ("refueling stations"), also called repeaters, are used between fiber sections (see Fig. 1.3), which restore a degraded signal for continued transmission.

As the light pulses propagate through the fiber, the pulses spread out, which means the duration of the pulses broadens. This spreading is called *dispersion*, and

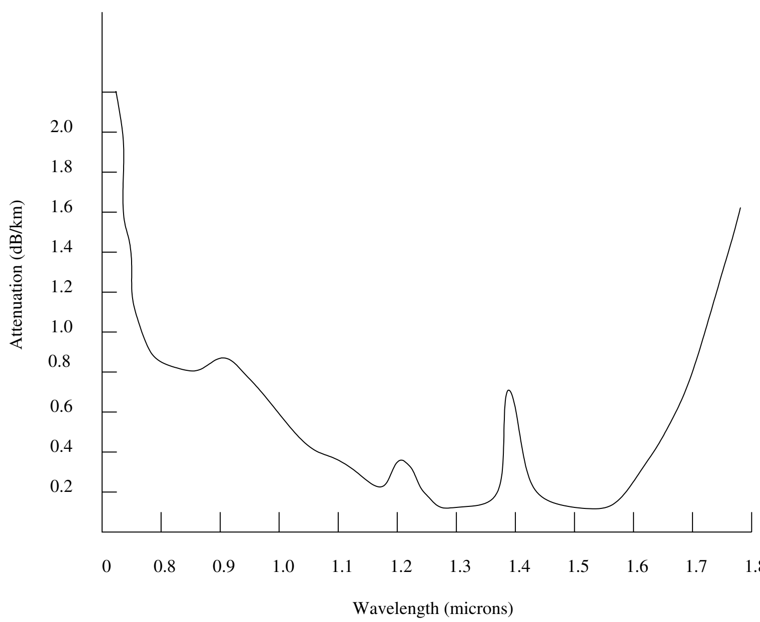

Figure 1.4 Attenuation versus wavelength for optical fiber.

the amount of it is wavelength-dependent. Dispersion limits the minimum time interval between consecutive pulses (because of interference with pulses of light ahead and behind) and hence the bit rate. There are two basic dispersive effects in a fiber. They are *intermodal dispersion* and *chromatic dispersion*. Intermodal dispersion occurs in multimode fibers; in these fibers, as the energy in a pulse travels in different modes, each with a different velocity, the pulse gets smeared after it has traveled some distance along the fiber. Chromatic dispersion is caused by the transmitting laser, which is unable to send all the photons at exactly the same wavelength, so different wavelengths travel at different speeds. A special pulse shape, called a *soliton*, has been discovered that retains its shape as it propagates through the fiber. Thus solitons provide a solution to the dispersion problem. Experiments have shown that solitons can achieve speeds of 80 Gb/s over distances of 10,000 km.

1.3 WAVELENGTH DIVISION MULTIPLEXING

Theoretically, fiber has extremely high bandwidth (about 25 THz [terahertz], i.e., 25 million MHz [megahertz] or 25 $\times 10^{12}$ Hz [Hertz]!) in the 1.55 low-attenuation band, and this is 1,000 times the total bandwidth of radio on the planet Earth [29]. However, only speeds (data rates) of a few gigabits per second are achieved because the rate at which an end user (for example, a workstation) can access the network is limited by electronic speed, which is a few gigabits per second. Hence it is extremely difficult to exploit all of the huge bandwidth of a single fiber using a single high-capacity wavelength channel due to optical-electronic bandwidth mismatch or "electronic bottleneck." The recent breakthroughs (Tb/s) are the result of two major developments: *wavelength division multiplexing* (WDM), which is a method of sending many light beams of different wavelengths simultaneously down the core of an optical fiber; and the *erbium-doped fiber amplifier* (EDFA), which amplifies signals at many different wavelengths simultaneously, regardless of their modulation scheme or speed.

WDM is conceptually similar to frequency division multiplexing (FDM), in which multiple information signals (each corresponding to an end user operating at electronic speed) modulate optical signals at different wavelengths, and the resulting signals are combined and transmitted simultaneously over the same optical fiber as shown in Fig. 1.5. Prisms and diffraction gratings can be used to combine (multiplex) or split (demultiplex) different color (wavelength) signals. FDM is

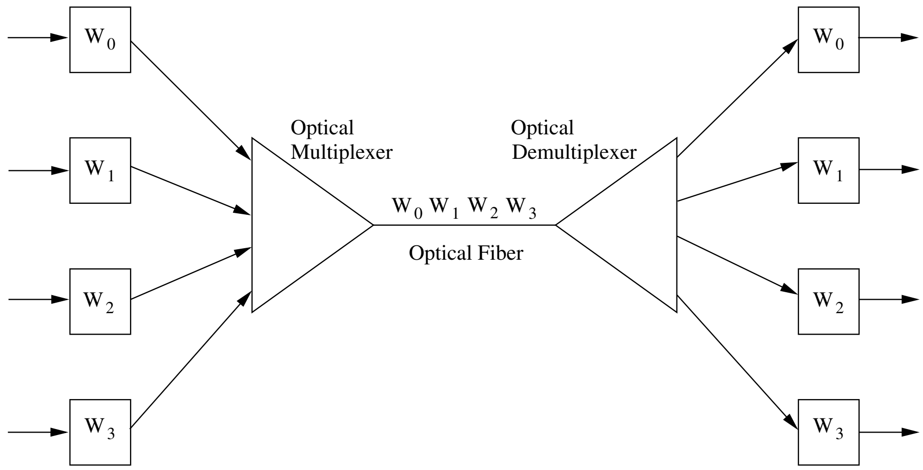

Figure 1.5 Wavelength division multiplexing.

used to carry many radio channels over the air or several simultaneous TV channels over cable. A WDM optical system using a diffraction grating is completely passive, unlike electrical FDM, and thus is highly reliable. Further, a carrier wave of each WDM optical channel is higher than that of an FDM channel by a million times in frequency (THz versus MHz). Within each WDM channel, it is possible to have FDM where the channel bandwidth is subdivided into many radio frequency channels, each at a different frequency. This is called *subcarrier multiplexing*. A wavelength can also be shared among many nodes in the network by electronic *time division multiplexing*. Note that WDM eliminates the electronic bottleneck by dividing the optical transmission spectrum (1.55-micron band) into a number of nonoverlapping wavelength channels, which coexist on a single fiber, with each wavelength supporting a single communication channel operating at peak electronic speed. The attraction of WDM is that a huge increase in available bandwidth can be obtained without the huge investment necessary to deploy additional optical fiber. To transmit 40 Gb/s over 600 km using a traditional system, 16 separate fiber pairs and 224 ($[600/40 - 1]16$) regenerators are required, as regenerators are placed every 40 km. On the other hand, a 16-channel WDM system requires only one fiber pair and 4 ($600/120 - 1$) optical amplifiers, as amplifiers are placed every 120 km. WDM has been used to upgrade the capacity of installed point-to-point transmission systems, typically by adding two, three, or four additional wavelengths. WDM systems that use 16 wavelengths at OC-48 and 32 wavelengths at OC-192 to provide aggregate rates up to 40 and 320 Gb/s, respectively, are available. The dense WDM (DWDM) technique effectively increases the total number of channels in a fiber by using very narrowly spaced (or densely packed) channels. Typical channel spacings range from 0.4 nm (1 nanometer is 10^{-9} meter) to 4 nm (50 GHz to 500 GHz).

As we have seen earlier, when the attenuation problem was solved, dispersion effects became significant. Similarly, when the dispersion problem was solved, nonlinear effects in fiber became dominant. The nonlinear effects—*stimulated Raman scattering*, *stimulated Brillouin scattering*, *self- and cross-phase modulation*, and *four-wave mixing*—may potentially limit the performance (maximum transmission rate) of WDM communication systems. Stimulated Raman scattering is caused by the interaction of the optical signal with silica molecules in the fiber. This interaction can lead to transfer of power from lower-wavelength channels to higher-wavelength channels. Stimulated Brillouin scattering is caused by the interaction between the optical signal and acoustic waves in the fiber. This interaction can make the power from the optical signal be scattered back toward the transmitter. Self- and cross-phase modulation and four-wave mixing are caused because, in optical fiber, the index of refraction depends on the optical intensity of signals propagating through the fiber. Self-phase modulation is caused by variations in the power of an optical signal and results in variations in the phase of the signal. Due to dispersion in fiber, phase shifts get transformed to signal distortion. By contrast, cross-phase modulation is due to a change in intensity of a signal propagating at a different wavelength. Four-wave mixing occurs when two or more optical signals (wavelengths) mix in such a way that they produce new optical frequencies called *sidebands*, which can cause interference if they overlap with frequencies used for data transmission. The above nonlinearities in fiber can be controlled by careful choice of channel power and channel spacing.

The advent of EDFA enabled commercial development of WDM systems by providing a way to amplify all the wavelengths at the same time, regardless of their individual bit rates, modulation scheme, or power levels. Before the invention of EDFAs, the effects of optical loss were compensated every few tens of kilometers by an electronic regenerator, which required that the optical signals be converted to electrical signals and then back again to optical ones. Most important, electronic regenerators work only for the designated bit rate at only one wavelength. The EDFA amplifier contains several meters of silica glass fiber that have been doped with ions of erbium, a rare-earth metal. An optical pump laser is then used to energize the erbium ions, which boost or amplify the optical signals that are passing through (see Fig. 1.6). It is a wonderful coincidence of nature that this amplification band (1.53–1.56 microns), with a gain spectrum of 0.03–0.04 microns, coincides with the 1.55-micron band of optical fibers.

There are three types of signal regenerators available in practice: The 3R regenerator (a regenerator executing *reshaping* and *reclocking* operations; reshaping of the signal reproduces the original pulse shape of each bit, eliminating much of the noise, and reclocking of the signal synchronizes the signal to its original bit timing pattern and bit rate), the 2R regenerator (a regenerator executing only the reshaping operation), and the 1R regenerator (a regenerator, without reshaping and retiming operations, carrying out simple amplification using EDFAs or other amplifiers). A network element which combines an optical receiver, some degree of regeneration, and an optical transmitter is called a *transponder*. If a transponder transmits on a wavelength that is different from that of the received signal, it is also carrying out *wavelength conversion* operation as a by-product of regeneration.

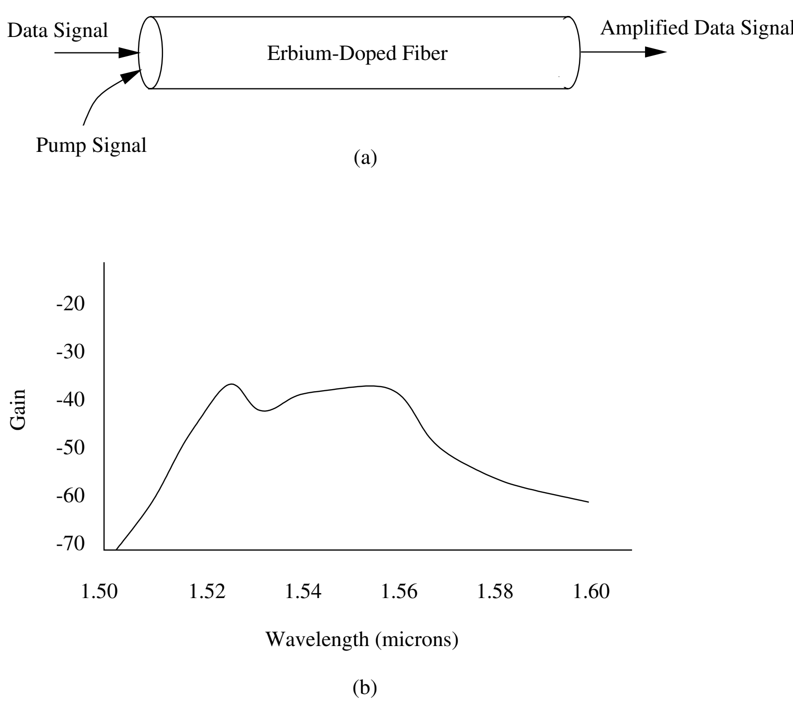

Figure 1.6 (a) EDFA. (b) EDFA gain profile.

1.4 WDM OPTICAL NETWORKING EVOLUTION

1.4.1 WDM Point-to-Point Link

WDM point-to-point links, an example of which is shown in Fig. 1.7, are being deployed by several telecommunication companies due to the increasing demands on communication bandwidth. The capacity of the fiber link $A \rightarrow B$ is now increased by a factor of 2, the number of wavelength channels used. These links are more cost-effective, when the demand exceeds the capacity in existing fibers, compared to installing new fiber. WDM mux/demux (multiplexer/demultiplexer) in point-to-point links with 64 channels is currently available [112].

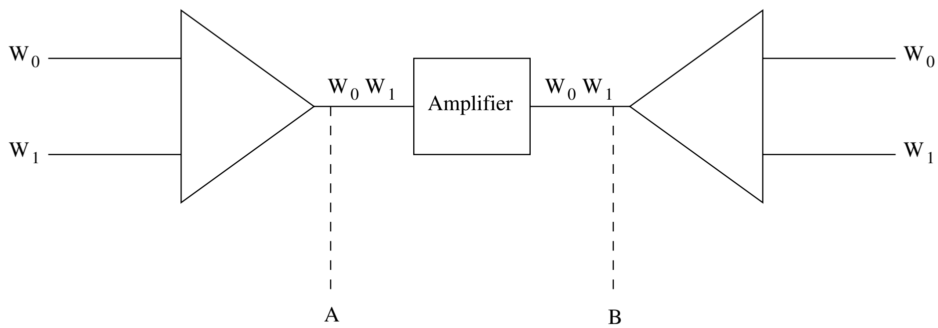

Figure 1.7 WDM point-to-point link.

1.4.2 Wavelength Add/Drop Multiplexer

While WDM point-to-point links provide very large capacity between two widely spaced (300 to 600 km) (end) points (or terminals), in many networks it is necessary to drop some traffic at intermediate points along the route between the end points. Using a wavelength add/drop multiplexer (WADM), which can be "inserted" on a fiber link as shown in Fig. 1.8, one can add/drop necessary traffic (wavelengths) at the WADM location. A WADM can be realized using a demultiplexer, 2×2 switches (one switch per wavelength), and a multiplexer. If a 2×2 switch (S_1 in the figure) is in "bar" state, then the signal on the corresponding wavelength passes through the WADM. If the switch (S_0 in the figure) is in "cross" state, then the signal on the corresponding wavelength is "dropped" locally, and another signal can be "added" on to the same wavelength at this WADM location.

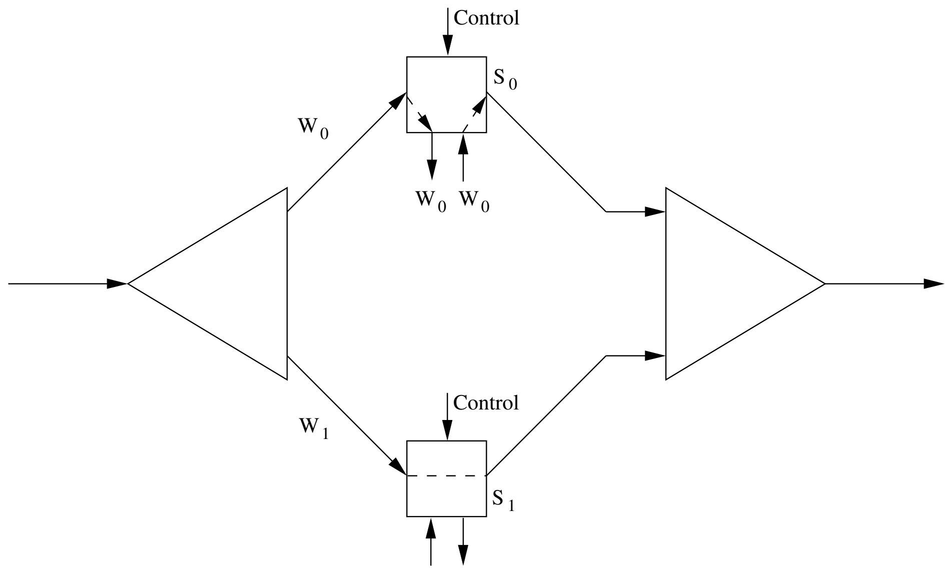

Figure 1.8 Wavelength add/drop multiplexer.

1.4.3 Wavelength Crossconnect

In order to build a flexible multipoint WDM optical network, apart from WADMs, we need another optical network element called a wavelength crossconnect. Functionally, WADM and wavelength cross-connects are quite similar, differing mainly in the number of input fibers that need to be handled [3]. The function of each element is to provide, under network control, the ability to connect (switch) any input wavelength channel from an input fiber (port) to any one of the output fibers (ports) in optical form or to drop a channel. The wavelength crossconnect is also referred to as a *wavelength selective crossconnect* (WXC) or *wavelength routing switch*. Figure 1.9 shows a 2×2 wavelength crossconnect which can be realized by demultiplexers, optical switches, and multiplexers. Note that a WXC may also allow addition and dropping of wavelengths.

1.5 ENABLING TECHNOLOGIES FOR WDM OPTICAL NETWORKS

WDM optical networking is enabled by a range of technologies. At the foundation is the extremely high-bandwidth (25-THz), low-attenuation-loss (0.2-dB/km

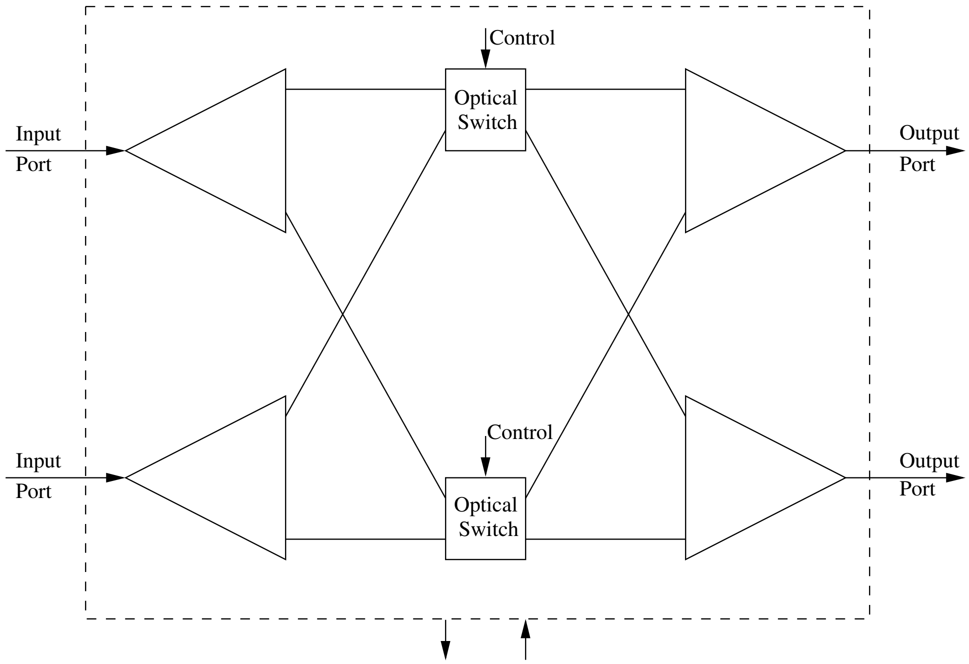

Figure 1.9 Wavelength crossconnect.

in the 1.55-micron low-attenuation band) single-mode optical fiber allowing long-distance transmission. A new type of fiber, called *AllWave* fiber, does not have the "waterpeak window" at 1.39 microns, which previously prevented the use of that part of the spectral region of the fiber, and hence it provides a more usable optical spectrum. EDFAs provide optical amplification of all the wavelengths at the same time to compensate for power loss in optical signal transmission. Conventional or C-band EDFAs amplify signals in the range 1.530–1.562 microns. Long-wavelength- or L-band EDFAs (which use longer lengths of erbium-doped fiber) allow amplification of signals in the wavelength range 1.570–1.610 microns and will be the next generation of EDFAs [50]. A combined (C+L)-band EDFA can provide 10 THz of bandwidth. While EDFAs work over 1.530–1.610 microns wavelength range, Raman amplifiers can amplify signals from 1.270 to 1.670 microns to further increase the capacity of optical fibers. Most optical transmission systems use semiconductor lasers as their light sources. The most commonly used light sources are distributed Bragg reflector and distributed feedback lasers. The transmitters used in WDM networks often require the capability to tune to different wavelengths. The well-established approaches for realizing tunable optical sources include mechanically tuned lasers, acousto-optically and electro-optically tuned lasers, and injection current tuned lasers. Two recent developments include array sources and switched sources. These tunable lasers differ mainly in two characteristics: tuning range (this refers to the range of wavelengths over which the laser may be operated) and tuning time (this specifies the time required for the laser to tune from one wavelength to another). Semiconductor lasers with a tunable range of 0.04 microns are becoming commercially available. Single-channel transmission speeds greater than 10 Gb/s are enabled by Mach-Zehnder laser modulators with modulation bandwidths in excess of 10 GHz. Tunable filters, which allow splitting and combining of the available wavelength band into several individual wavelength channels, are another key technology. A variety of tunable filters such as fiber Fabry-Perot and fiber Bragg grating filters are commercially available. The switching times of electromechanical optical switches are now typically between a few and tens of milliseconds. In the very near future it will be possible to realize high-performance, low-cost optical components (such as optical switches, tunable lasers, and variable optical attenuators—variable optical attenuators are used inside optical amplifier, add/drop multiplexers, and crossconnects to control the optical power) using MEMS (Micro-ElectroMechanical Systems) technologies [119]. MEMS devices are miniature structures fabricated on silicon substrates in a similar manner to silicon integrated circuits. However, unlike electronic circuits, these are mechanical devices. Optical MEMS switches built using several micromirrors on silicon have already been demonstrated. Commercial PIN (p-type, intrinsic, n-type) photodetectors and avalanche photodetectors provide bit rates of 10 Gb/s and higher at the receiver. EDFA-based optical preamplifers, which raise the power level at the input of the receiver, enable high sensitivity of these optical receivers.

Today's widely installed WDM optical networks are *opaque*, that is, a signal path (connection) between any two end nodes or users in these networks is not totally optical. This means the path involves optical–electronic–optical conversion operations at intermediate nodes and these conversion operations affect the network speeds or bit rates. WDM optical networks are migrating from just point-to-point WDM links to *all-optical* networks, where more and more switching and routing functions are carried out in optical domain. In all-optical networks each connection (*lightpath*) is totally optical except at the end nodes.

Note that the designers of WDM networks must be aware of the properties and limitations of optical fiber and components (without making unrealistic assumptions about these) in order to realize a practical network and its associated protocols/algorithms which exploit the full potential of WDM. For network operators, who like to deploy equipment from multiple vendors that operate together in a single network, interoperability among the equipment is very important. WDM network standards are being developed to achieve WDM multivendor interoperability. Standards (which focus on interfaces that specify how equipment is physically interconnected and what procedures are used to operate across different equipment) allow network operators to have choice of buying equipment from different, competing vendors, rather than being committed to buying equipment from a single vendor. Achieving optical interoperability is not an easy task because it requires standardizing parameters such as wavelength, optical power, signal-to-noise ratio, pilot tone (for keeping track of the origin of each optical signal that is monitored), and supervisory channel (for carrying control information). There are several standards organizations and industry consortia working in this area (standards), including the International Telecommunications Union–Telecommunication Standardization Sector (ITU-T), Internet Engineering Task Force (IETF), Optical Domain Service Initiative (ODSI), and Optical Internetworking Forum (OIF). Appendix A gives a Web resources list, linking to home pages of several WDM optical component/system vendors.

1.6 WDM OPTICAL NETWORK ARCHITECTURES

There are three classes of WDM optical network architectures: broadcast-and-select networks, wavelength routed networks, and linear lightwave networks. We now explain each of these networks in detail.

1.6.1 Broadcast-and-Select Networks

A broadcast-and-select network consists of a passive star coupler connecting the nodes in the network as shown in Fig. 1.10(a). Each node is equipped with one or more fixed-tuned or tunable optical transmitters and one or more fixed-tuned or tunable optical receivers. Different nodes transmit messages on different wavelengths simultaneously. The star coupler combines all these messages and then broadcasts the combined message to all the nodes. A node selects a desired wavelength to receive the desired message by tuning its receiver to that wavelength. Note that the star coupler offers an optical equivalent to radio systems: each transmitter broadcasts its signal or message on a different wavelength and the receivers are

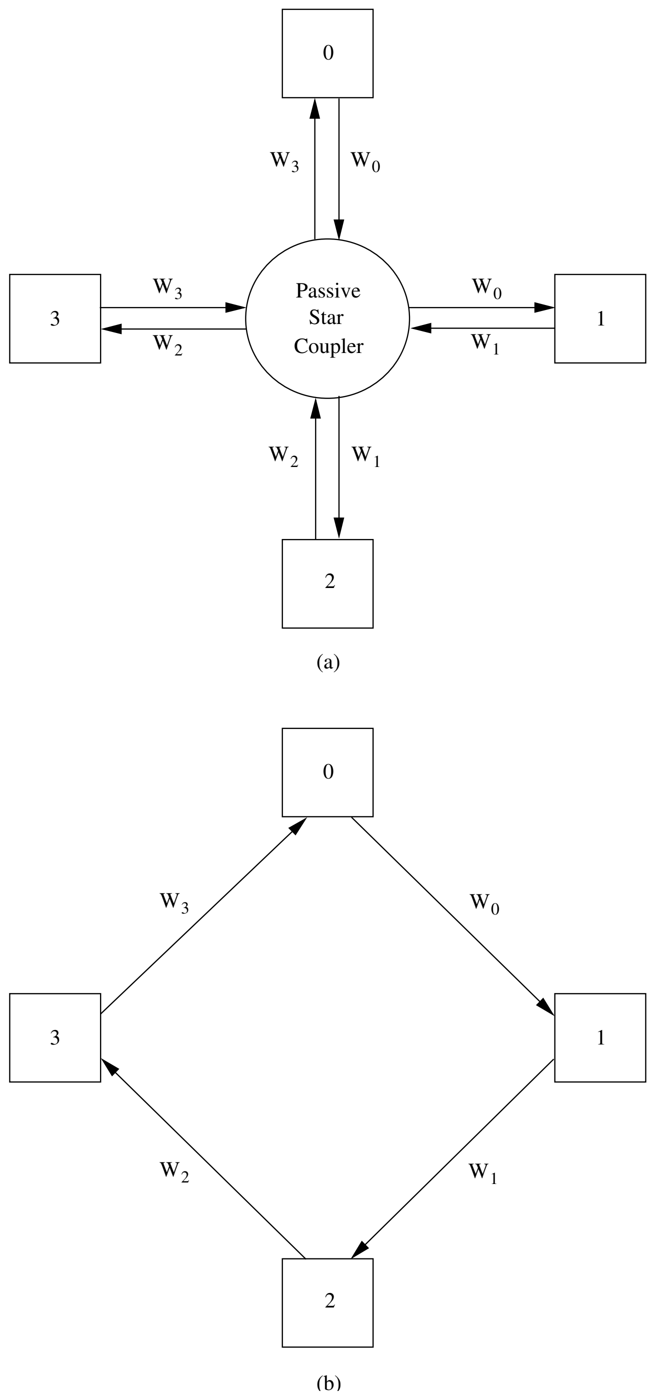

Figure 1.10 (a) Broadcast-and-select network. (b) Logical topology.

tuned to receive the desired signal. An $N \times N$ star coupler can be realized using a multistage interconnection network which has $\log_2 N$ stages of 2×2 couplers with $N/2$ couplers per stage (assuming N is a power of 2) or directly in integrated optics form with a common coupling region. Integrated optics refers to integration of optical components with fiber interconnections onto a single optical substrate, similar to the way in which electrical components (such as resistors, capacitors, and inductors) are combined in an electronic integrated circuit.

In single-hop broadcast-and-select networks, a message, once transmitted as light, reaches its final destination directly, without being converted to electronic form in between. In order to support packet switching in these networks, we need to have optical transmitters and receivers that can tune rapidly. This is because, in a packet-switched network, a node must be able to transmit (receive) successive packets to (from) different nodes on different wavelengths. The main networking challenge in these networks is the coordination of transmissions between various nodes. In the absence of coordination or efficient medium access control (MAC) protocol, collisions occur when two or more nodes transmit on the same wavelength at the same time. Also, destination conflicts occur if two or more nodes transmit on different wavelengths to the same destination when the destination has only one tunable optical receiver. Moreover, the destination must know when to tune to the appropriate wavelength to receive a packet. Several MAC protocols have been proposed to prevent such collisions/conflicts for single-hop broadcast-and-select networks, assuming the availability of rapidly tunable transmitters and/or receivers [111]. (See Problem 1.10.)

To support packet switching efficiently in broadcast-and-select networks, a multihop approach, which avoids rapid tuning altogether, can be used. Each node has a small number of fixed-tuned optical transmitters and fixed-tuned optical receivers. Each transmitter is at a different wavelength. We can represent the network as a graph, wherein a node corresponds to a network node and an edge corresponds to a transmitter–receiver pair on the same wavelength. Thus we obtain a *virtual* or *logical topology* over the physical broadcast topology. Figure 1.10(a) shows a four-node broadcast-and-select network. Each node transmits at one fixed wavelength and receives on one fixed wavelength. For example, node 0 can transmit directly to node 1 using wavelength w_0, but not to node 2. To transmit to node 2, node 0 sends a packet to node 1 on wavelength w_0, which receives it, converts it to electronic form, and retransmits it on wavelength w_1. The packet then reaches node 2. The virtual topology of the network in Fig. 1.10(a) is shown in Fig. 1.10(b). In these networks, a packet may have to go through more than one hop before reaching its destination. This leads to an increase in propagation delay in addition to queueing delays at intermediate nodes, and wastage of network capacity.

The advantage of broadcast-and-select networks is in their simplicity and natural *multicasting* capability (ability to transmit a message to multiple destinations). However, they have severe limitations. (1) They require a large number of wavelengths, typically at least as many as there are nodes in the network, because there is no *wavelength reuse* in the network. Thus the networks are not scalable beyond the number of supported wavelengths. (2) They cannot span long distances since the transmitted power is split among various nodes and each node receives only a small fraction of the transmitted power, which becomes smaller as the number of nodes increases. For these reasons, the main application for broadcast-and-select is high-speed local area networks (LANs) and metropolitan area networks (MANs).

Broadcast-and-Select Network Demonstrators

A number of demonstrators based on the broadcast-and-select approach have been built [39]. Bellcore's LAMBDANET (late 1980s and early 1990s), which employed 18 wavelengths separated by 2 nm, was one of the first demonstrators. Transmission at 1.5 Gb/s per wavelength over 57 km was also demonstrated. Nippon Telegraph and Telephone Corporation's (NTT's) test bed in the early 1990s used 100 wavelengths, spaced 10 GHz apart, each carrying data at 622 Mb/s. International Business Machines Corporation's (IBM's) RAINBOW-I and RAINBOW-II test beds were designed to support 32 wavelengths separated by 1 nm, in a star configuration. RAINBOW-I interconnected computers; after a connection was set up, transmission at 300 Mb/s on each wavelength was possible. In RAINBOW-II, an extension of RAINBOW-I, the transmission rate per wavelength was upgraded to 1 Gb/s, and the control protocol processing was improved. STARNET-I and II WDM computer networks, developed at Stanford University, resemble RAINBOW. STARNET-I had two data rates at each node, both transmitted on a unique wavelength using a single laser. These rates were a fast 2.5 Gb/s for establishing circuit-switched connections with other nodes as in RAINBOW and 100 Mb/s for establishing a logical ring network. In STARTNET-II, the high data rate was modified to 1.25 Gb/s. The European Research and Development for Advanced Communications in Europe (RACE) program developed a star network and used it for video applications in a BBC television studio test bed. Each node in the network, called a local routing center (LRC), interconnected up to 16 (electrical signal) sources and destinations operating at 155 Mb/s each. Up to 16 such LRCs were then interconnected using a star coupler. LIGHTNING is another test bed that is a tree hierarchical WDM network used for interconnecting processors which form the leaves of the tree. Another example is the supercomputer supernet, developed jointly by the University of California at Los Angeles, Jet Propulsion Laboratory, and the Aerospace Corporation, for connecting supercomputers separated geographically by several kilometers.

1.6.2 Wavelength Routed Networks

Wavelength routed WDM networks have the potential to avoid the three problems—lack of wavelength reuse, power splitting loss, and scalability to wide area networks (WANs)—of broadcast-and-select networks. A wavelength routed network consists of WXCs (routing nodes) interconnected by point-to-point fiber links in an arbitrary topology. Each end node (end user) is connected to a WXC via a fiber link. The combination of end node and its corresponding WXC is referred to as a (network) node. Each node is equipped with a set of transmitters and receivers, for sending data into the network and receiving data from the network, respectively, both of which may be wavelength-tunable.

In a wavelength routed network, a message is sent from one node to another node using a wavelength continuous route called a *lightpath*, without requiring any optical–electronic–optical conversion and buffering at the intermediate nodes. This process is known as *wavelength routing*. Note that the intermediate nodes route the lightpath in the optical domain using their WXCs. The end nodes of the lightpath access the lightpath using transmitters/receivers that are tuned to the wavelength on which the lightpath operates. A lightpath is an *all-optical communication path* between two nodes, established by allocating the same wavelength throughout the route of the transmitted data. Thus it is a high-bandwidth pipe, carrying data up to several gigabits per second, and is uniquely identified by a physical path and a wavelength. The requirement that the same wavelength must be used on all the links along the selected route is known as the *wavelength continuity constraint*. Two lightpaths cannot be assigned the same wavelength on any fiber. This requirement is known as *distinct wavelength assignment constraint*. However, two lightpaths can use the same wavelength if they use disjoint sets of links. This property is known as *wavelength reuse*.

Example 1: *Consider a wavelength routed network with five nodes and two wavelengths per fiber as shown in Fig. 1.11. Assume that lightpaths are to be established*

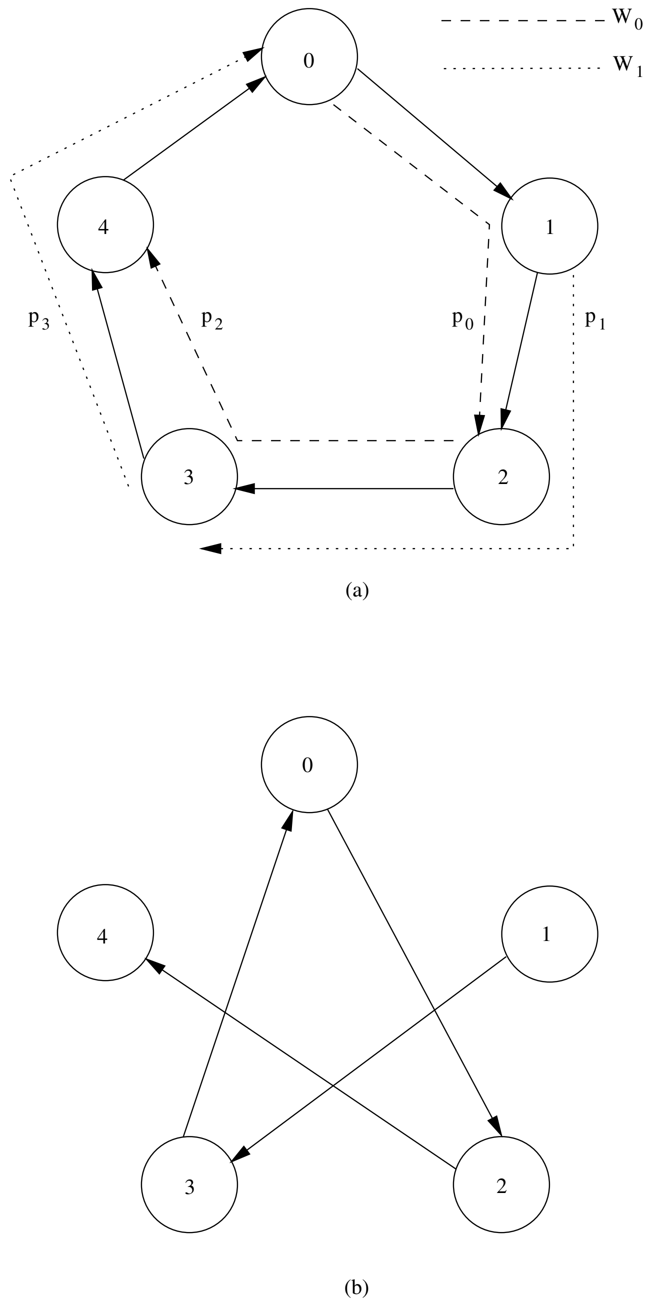

Figure 1.11 (a) Wavelength routed network. (b) Logical topology.

one for each of the node pairs $< 0, 2 >$, $< 1, 3 >$, $< 2, 4 >$, $< 3, 0 >$, *and* $< 4, 1 >$. *Further assume that every node is equipped with one transmitter and one receiver. For the given set of node pairs, a node is a source for one lightpath and destination for one lightpath. There exists only one physical path between any node pair. Every fiber link in the network would carry physical paths corresponding to two lightpaths, if lightpaths were successfully established for all the given node pairs. Since two wavelengths are available on any fiber link, it should be possible to route all the lightpaths. However, due to the wavelength continuity constraint it is not possible to establish lightpaths for all five node pairs. The figure shows a possible way of routing four lightpaths* p_0, p_1, p_2, *and* p_3, *where* p_i *is the lightpath emanating from node* i. *Since* p_0 *uses wavelength* w_0, p_1 *can use only* w_1, *as* p_0 *and* p_1 *share a link. Lightpath* p_2 *can use only* w_0, *as* p_1 *and* p_2 *share a link. Lightpath* p_3 *can use only* w_1, *as* p_2 *and* p_3 *share a link. As a consequence, wavelength* w_0 *is free on link* $4 \rightarrow 0$ *and* w_1 *is free on link* $0 \rightarrow 1$. *Therefore, a lightpath cannot be established from node 4 to node 1 even though bandwidth (wavelength) is available on links* $4 \rightarrow 0$ *and* $0 \rightarrow 1$, *a transmitter is available at node 4, and a receiver is available at node 1. As we will see later, this bandwidth loss caused by the wavelength continuity constraint can be overcome by using a wavelength converter. However, observe that both the wavelengths* w_0 *and* w_1 *are reused two times, thus helping increase the number of lightpaths established while employing a limited number of wavelengths.* □

Wavelength reuse (an important feature that refers to simultaneous transmission of messages on the same wavelength over fiber-link-disjoint lightpaths) in wavelength routed networks makes them more scalable than broadcast-and-select networks. Another important characteristic which enables wavelength routed networks to span long distances is that the transmitted power invested in the lightpath is not split to irrelevant destinations. Given a WDM network, the problem of routing and assigning wavelengths to lightpaths is of paramount importance in these networks, and clever algorithms are needed in order to ensure this function (routing and wavelength assignment) is performed using a minimum number of wavelengths. The number of available wavelengths in a fiber link plays a major role, in these networks, which currently varies between 4 and 32, but is expected to increase (with announcements of over 100 wavelengths already made).

Packet switching in wavelength routed networks can be supported by using either a single-hop or a multi-hop approach, in a way similar to broadcast-and-select networks. In the multi-hop approach, a virtual topology (a set of lightpaths or *optical layer*) is imposed over the physical topology (which is not broadcast here) by setting the WXCs in the nodes. Over this virtual topology, a packet from one node may have to be routed through some intermediate nodes before reaching its final destination. At each intermediate node, the packet is converted to electronic form and retransmitted on another wavelength. A virtual topology, formed by lightpaths p_0 through p_3, corresponding to the physical network shown in Fig. 1.11(a), is given in Fig. 1.11(b). A packet from node 3 (source) is routed through intermediate node 0 undergoing optical–electronic–optical conversion at this node before reaching node 2 (destination).

Existing Internet backbone networks consist of high-capacity IP (Internet Protocol—developed for providing connectionless transfer of packets across an internetwork) routers interconnected by point-to-point fiber links. Traffic is transported between routers through high-speed gigabit links. These links are realized by SONET or ATM-over-SONET technology. The backbone routers use IP-over-SONET or IP-over-ATM-over-SONET technology to route IP traffic in the backbone network. Most of the SONET-based backbone transport networks provide data interface at the rate of OC-3 and OC-12. The traffic demand is growing at a faster rate and a point has been reached where data interfaces at the rate of OC-48 and more are required. Upgrading the existing SONET transport infrastructures to handle these high-capacity interface rates is not desirable, as it is impractical to go for upgrading every time the interface rate increases. Also, such upgrading is not economical. A viable and cost-effective solution is to use WDM technology in backbone transport networks. In such (for example, IP-over-WDM) networks, network nodes are interconnected by WDM fiber links (where each link is capable of carrying multiple signals simultaneously, each on a different wavelength), and the nodes employ WXCs and electronic processing elements. Figure 1.12 shows a typical WDM backbone network. The electronic processing element can be an IP router, ATM switch, or a SONET system.

Any two IP routers in this network can be connected together by a lightpath. Two nodes that are not connected directly by a lightpath communicate using multi-hop approach, i.e., by using electronic packet switching at the intermediate nodes. This electronic packet switching can be provided by IP routers, ATM switches, or SONET equipment, leading to an IP-over-WDM or an ATM-over-WDM, or a SONET-over-WDM network, respectively.

A WDM-based transport network can be decomposed broadly into three layers, *a physical media layer, an optical layer, and a client layer*, as shown in Fig. 1.13. Application of WDM technology has introduced the optical layer between the lower

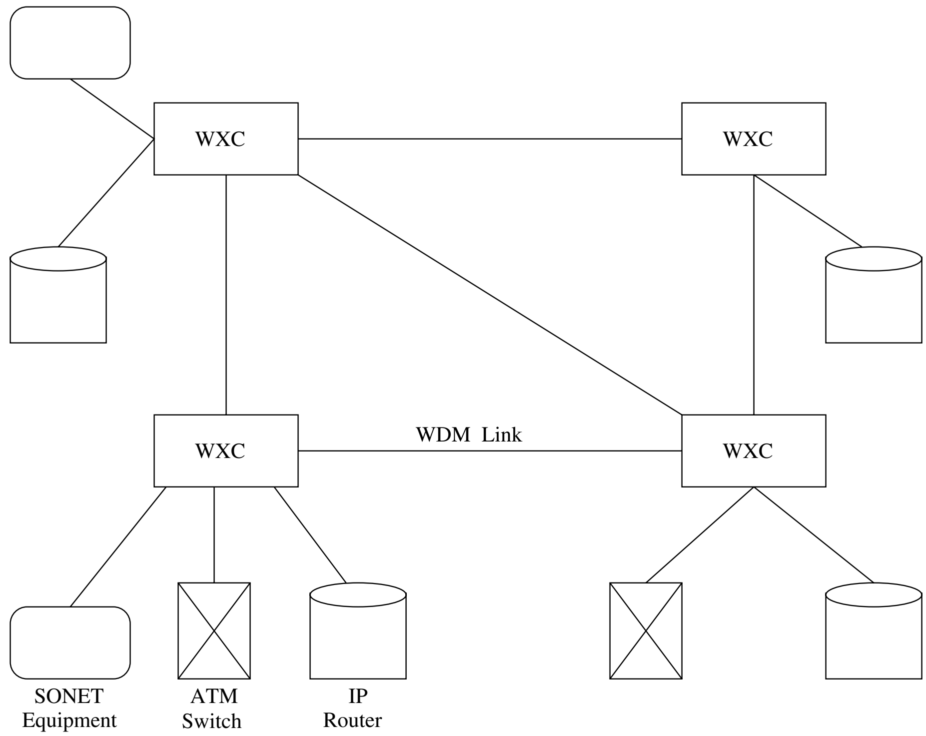

Figure 1.12 A WDM backbone network.

physical media layer and upper client layer. A set of lightpaths constitutes the optical layer (virtual topology). The optical layer provides client-independent or protocol-transparent circuit-switched service to a variety of clients that constitute the client layer. This is possible because the lightpaths can carry messages at a variety of bit rates and protocols. Thus the optical layer can support a variety of clients concurrently. For example, some lightpaths could carry SONET data, whereas others could carry IP packets/datagrams or ATM cells. A network with an optical layer can be configured such that in the event of failures, lightpaths can be rerouted over alternate paths automatically. This provides a high degree of reliability in the network. According to International Telecommunications Union–Telecommunication Standardization Sector (ITU-T) Recommendation G.872, an optical layer can be further decomposed into three sublayers: *an optical channel layer, an optical multiplex section layer*, and *an optical transmission section layer*. The functionality of the optical channel layer is to provide end-to-end networking of

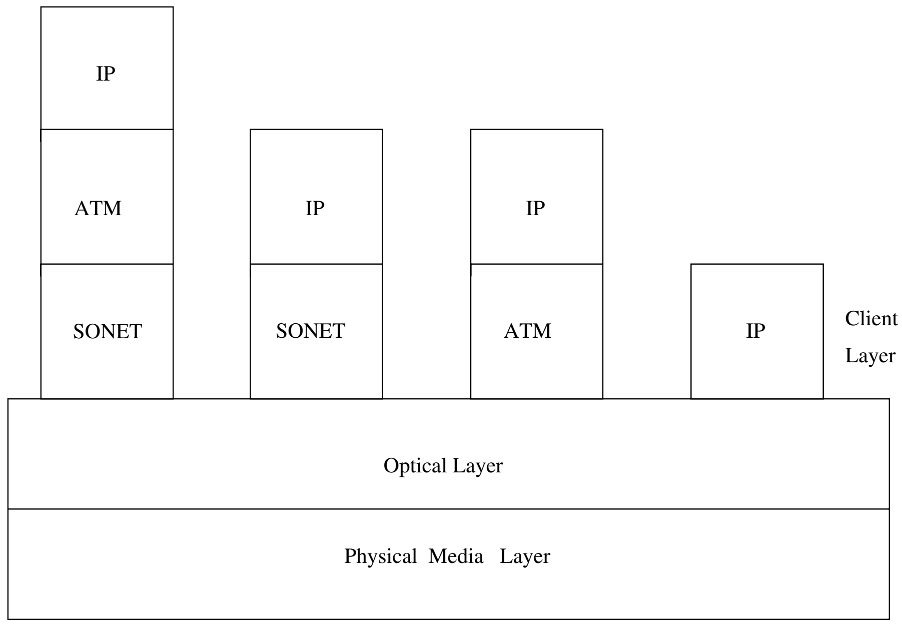

Figure 1.13 Possible layers in a WDM optical transport network.

optical channels (lightpaths) for transparently conveying the client data. The optical multiplex section layer concerns networking of aggregate multiwavelength optical signals. The optical transmission section layer concerns the transmission of optical signals on different kinds of optical media such as single-mode and multimode transmission.

These attractive features—wavelength reuse, protocol transparency, and reliability—make wavelength routed networks suitable for WANs. This book concentrates on wavelength routed (wide area) optical networks, mainly addressing the key issues—design, reconfiguration, and failure recovery in the optical layer, and architectures and technologies which will enable the realization of wavelength routed future optical Internet networks to transport high-speed IP traffic. Designing an optical layer to meet the traffic demand is an important problem in order to use the wavelength and fiber resources efficiently and to provide quality service to the users. Reconfiguring the optical layer is necessitated by the changing traffic demand. Since a huge amount of traffic is carried by the optical layer, rapid service recovery in case of network component failures is critically important.

Wavelength Routed Network Demonstrators

One of the most comprehensive field deployment trials is the Defense Advanced Research Projects Agency (DARPA)-sponsored Multiwavelength Optical Network (MONET) program. It is a wavelength routed test bed that uses eight wavelengths spaced 200 GHz apart, with a data transmission rate of 2.5 Gb/s per wavelength. The field trial consists of a local ring network with several WADMs in it and a WXC to interconnect the local ring network to a long-distance network. The program has demonstrated 10-Gb/s data transmission over 2,000 km. The All-Optical Network (AON) test bed, deployed in Boston between MIT's Lincoln Laboratory and Digital Equipment Corporation, uses 20 wavelengths spaced by 50 GHz. Each node in the test bed has a tunable transmitter (distributed Bragg grating laser with 10 ns [1 ns is 10^{-9} second] tuning time) and a tunable filter. NTT's test bed consists of a unidirectional ring with a central hub node and many access nodes. It uses six wavelengths spaced by 100 GHz with a distance of 40 km between nodes at 622 Mb/s, with a hub node and two access nodes. The ring employs a failure protection mechanism. Another ring architecture was demonstrated by the Optical Networks Technology Consortium (ONTC). The ONTC test bed consists of two unidirectional rings, with two access nodes per ring, connected by a 2×2 WXC. It has four wavelengths spaced by 4 nm, with a data transmission rate of 155 Mb/s over a total distance of 150 km. MultiWavelength Transport Network (MWTN) was one of the first wavelength routed test beds, developed by the European RACE program, to demonstrate successfully the concept of wavelength routing in optical networks. The test bed uses four wavelengths spaced 4 nm apart, with each wavelength carrying SDH data at 622 Mb/s or 2.5 Gb/s in a Stockholm field trial.

1.6.3 Linear Lightwave Networks

A usable portion of the optical spectrum (for example, the 1.55-micron band) can be partitioned into a number of either wavelengths or wavebands as shown in Fig. 1.14 [162]. Observe that in Fig. 1.14(b) each waveband is further subdivided into a number of wavelengths. Note that sufficient spacing or guard bands have to be placed between any two wavelengths to allow for imprecision and drift in laser transmitter tuning and to make it possible to separate adjacent signals at the receivers. Wavelength routed networks use wavelength (one-level) partitioning and in these networks several wavelengths are multiplexed on a fiber link. Linear lightwave networks, on the other hand, use waveband (two-level) partitioning, and in these networks several wavebands are multiplexed on a fiber and several wavelengths are

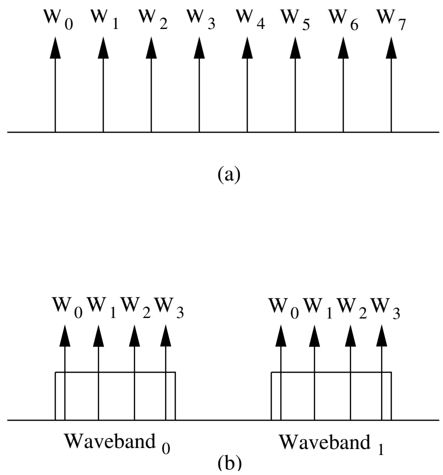

Figure 1.14 (a) Wavelength partitioning. (b) Waveband partitioning.

multiplexed on each waveband. In a wavelength routed network, routing nodes demultiplex, switch, and multiplex wavelengths, whereas in a linear lightwave network, routing nodes demultiplex, switch, and multiplex wavebands, but not wavelengths within a waveband. Thus the hardware requirements at the nodes, by grouping a set of wavelengths into a waveband, in linear lightwave networks get simplified because the number of optical switches required in a node is equal to the number of wavebands, not the number of wavelengths. Since a linear lightwave network as a whole does not distinguish between wavelengths within a waveband, individual wavelengths within a waveband are separated from each other at the end nodes (optical receivers).

Two constraints—wavelength continuity and distinct wavelength assignment—on optical connections applicable to wavelength routed networks also apply to linear lightwave networks. Further, there are two routing constraints unique to linear lightwave networks: *inseparability*, that is, channels belonging to the same waveband when combined on a single fiber cannot be separated within the network; and *distinct source combining*, that is, on any fiber, only signals from distinct sources are allowed to be combined.

Inseparability

Figure 1.15 illustrates the inseparability constraint. Here nodes 0 through 5 are end nodes, while nodes A through H are routing nodes. The figure also shows two

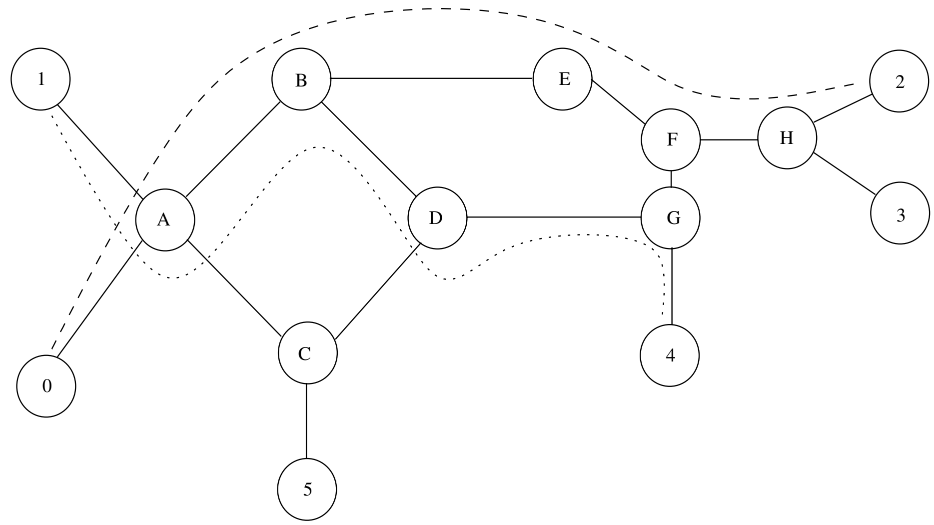

Figure 1.15 Linear lightwave network.

connections, one between nodes 0 and 2, and the other between nodes 1 and 4. The notation $< x, y >$ is used to denote a connection from node x to node y. $< 0, 2 >$ passes through nodes A, B, E, F, and H, while $< 1, 4 >$ passes through nodes A, B, D, and G. Since the two connections share the fiber $A \rightarrow B$, they have to use different wavelengths (because of the distinct wavelength assignment constraint). Suppose wavelength w_0 is used for $< 0, 2 >$ and wavelength w_1 for $< 1, 4 >$. The two signals, though on different wavelengths, may be in the same waveband. The two signals (that is, the power from both sources, nodes 0 and 1) are combined at node A. At node B, however, the two signals cannot be separated since they belong to the same waveband. Therefore, to route the two connections to their destinations, the combined signal power is split equally between the output ports of node B, one leading to node E and the other to routing node D. Only the end nodes (destinations), 2 and 4, filter out w_0 and w_1, respectively, rejecting the other wavelength. Note that some unintended destinations appear—node 2 in the case of connection $< 1, 4 >$, and node 4 in the case of connection $< 0, 2 >$. These unintended destinations are called *fortuitous destinations*; they tend to waste fiber resources and power, and are therefore to be dispensed with, if possible. For example, in this case the fortuitous destinations could have been avoided by rerouting connection $< 1, 4 >$ on the path 1–A–C–D–G–4.

The inseparability constraint can lead to a *color clash*. Suppose connections $< 0, 2 >$ and $< 1, 4 >$ routed on paths 0–A–B–E–F–H–2 and 1–A–C-D–G–4, respectively, use the same wavelength w_0. Now a third connection $< 5, 3 >$, routed on the path 5–C–D–G–F–H–3 using the same waveband is assigned wavelength w_1, which is different from w_0 because connections $< 1, 4 >$ and $< 5, 3 >$ share the fiber $C \rightarrow D$. Note that, due to inseparability, all three signals (generated at sources 0, 1, and 5) are carried on fiber $F \rightarrow H$. Because signals generated at sources 0 and 1 are now on the same fiber $F \rightarrow H$, using the same wavelength w_0, they produce a color clash.

Distinct Source Combining

Distinct source combining disallows an optical signal from splitting at a node, taking multiple paths in the network, and then recombining with itself. We now illustrate how a correct but poor routing algorithm can violate the distinct source combining constraint. Consider the network shown in Fig. 1.15. Suppose that when two connections 0–A–B–E–F–H–2 and 1–A–B–D–G–4, on wavelengths w_0 and w_1, respectively, are in progress, a new connection $< 5, 3 >$ is routed along the path 5–C–D–G–F–H–3 using a wavelength w_2. All three connections are assumed to be in the same waveband. Due to inseparability, signal generated at source 5 carries with it (fortuitously) portions of signals generated at sources 0 and 1 after combining with them on fiber $D \rightarrow G$. This causes the signal generated at source 0 (which split at node B) to recombine with itself on fiber $F \rightarrow H$, violating the distinct source combining constraint. Thus the new connection $< 5, 3 >$ must not be routed along the path 5–C–D–G–F–H–3. However, this problem could have easily been avoided had the choice of routes been more prudent. Routing connection $< 1, 4 >$ via A, C, D, and G, or routing connection $< 5, 3 >$ via nodes C, D, B, E, F, and H would have made it possible for all three connections to be routed successfully.

Setting up connections with the above routing constraints in a linear lightwave network is significantly more complicated when wavebands contain more than one wavelength. Apart from this, the combining and splitting losses present at the nodes is preventing linear lightwave networks from becoming practical. Since at each node the power at an output port is a linear combination of the powers at the input ports, these networks are called linear lightwave networks.

1.7 ISSUES IN WAVELENGTH ROUTED NETWORKS

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

1.7.1 Routing and Wavelength Assignment

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

Static Versus Dynamic Traffic Demand

The connection requests (traffic demand) can be either static or dynamic. In case of a static traffic demand, connection requests are known *a priori*. The traffic demand may be specified in terms of source–destination pairs. These pairs are chosen based on an estimation of long-term traffic requirements between the node pairs. The objective is to assign routes and wavelengths to all the demands so as to minimize the number of wavelengths used. The dual problem is to assign routes and wavelengths so as to maximize the number of demands satisfied, for a fixed number of wavelengths. The above problems are categorized under the *static lightpath establishment* (SLE) problem. The SLE problem has been shown to be NP-complete (that is, it is computationally intractable or, in other words, the only known algorithms that find an optimal solution require exponential time in the worst case) [33]. Therefore, polynomial-time algorithms which produce solutions close to the optimal one are preferred.

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

Example 2: *Consider a network with four nodes and two wavelengths w_0 and w_1 per fiber, as shown in Fig. 1.16. It is required to establish lightpaths between node pairs $< 0, 1 >$, $< 2, 3 >$, $< 1, 3 >$, and $< 0, 2 >$. Assuming that the requests arrive in the same order as given here, a dynamic RWA algorithm establishes lightpaths p_0, p_1, and p_3 to the first three requests as shown in Fig. 1.16(a). This algorithm uses the first free wavelength for the chosen route. Lightpaths p_0 and p_1 use wavelength w_0 and p_2 uses w_1. As a result, no lightpath can be established between node 0 and node 2, as the route between these two nodes is not wavelength-continuous. A static RWA is able to establish lightpaths q_0, q_1, q_2, and q_3 for all the four node pairs as shown in Fig. 1.16(b). This algorithm considers the connections in nonincreasing order of hop length, and the first free wavelength is assigned to a connection. Here, node pairs are considered in the order $< 1, 3 >$, $< 0, 2 >$, $< 0, 1 >$, $< 2, 3 >$.* □

Centralized Versus Distributed Control

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

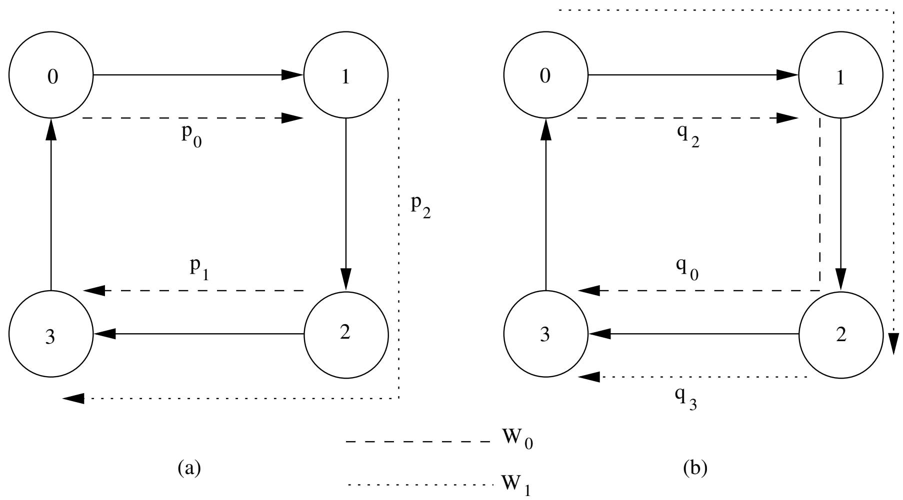

Figure 1.16 Performance of dynamic and static RWA algorithms. (a) Lightpaths established by a dynamic RWA algorithm. (b) Lightpaths established by a static RWA algorithm.

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

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

Route and Wavelength Selection Methods

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

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

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

1.7.2 Wavelength-Convertible Networks

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

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

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

Example 3: *Consider the network and the set of lightpaths in Example 1. With no node having conversion capability, lightpath p_4 cannot be established from node*

4 to node 1. If node 0 is equipped with a converter which can convert wavelength w_0 *to* w_1*, then* p_4 *can be established. This lightpath uses* w_0 *on link* $4 \rightarrow 0$ *and* w_1 *on link* $0 \rightarrow 1$. □

1.7.3 Multifiber Networks

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

Example 4: *Again consider the network and the set of lightpaths in Example 1. We have already observed that lightpath* p_4 *cannot be established from node 4 to node 1 because of the wavelength continuity constraint. However, if there is an additional fiber between node 4 and node 0, then* p_4 *can be established. This lightpath uses* w_1 *on the additional fiber on link* $4 \rightarrow 0$ *and* w_1 *on link* $0 \rightarrow 1$*. Thus multifiber wavelength routed networks employ space division multiplexing in addition to wavelength division multiplexing.* □

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

A multifiber wavelength-selective network with F fibers per link and W wavelengths per fiber is equivalent to a single-fiber network with FW wavelengths with the nodes having conversion degree F. The equivalence is depicted in Fig. 1.17. A node with an incoming and an outgoing link in a multifiber network is shown in Fig. 1.17(a). The link is comprised of two fibers each carrying two wavelengths w_0 and w_1. This node does not have conversion capability. A wavelength on an

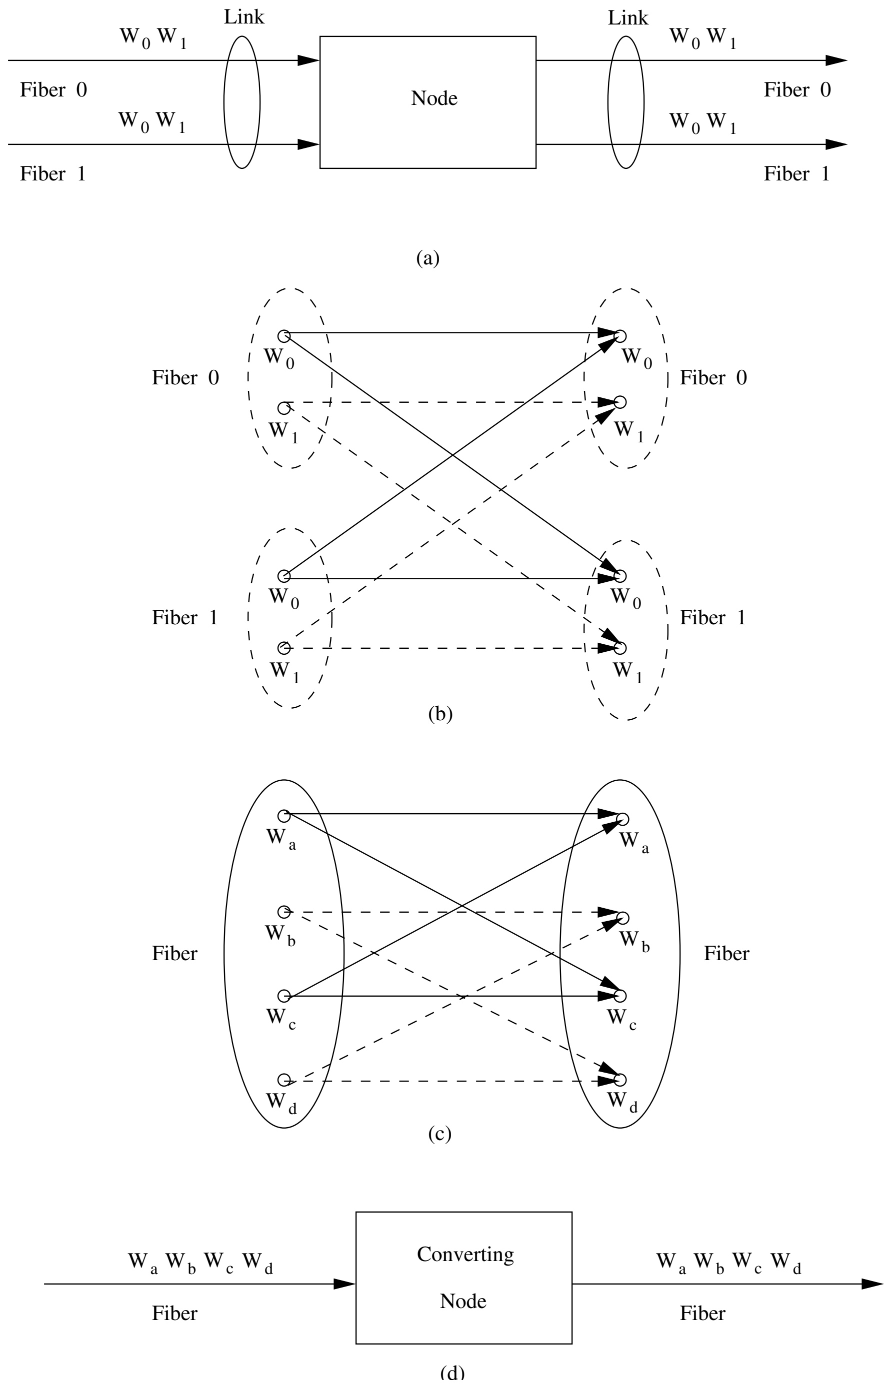

Figure 1.17 Equivalence between a multifiber wavelength-selective network and a wavelength-convertible network. (a) A node with two-fiber links. (b) A mapping of wavelengths from incoming fibers to outgoing fibers. (c) Wavelengths with new labels. (d) An equivalent wavelength-converting node with wavelength conversion degree = 2.

incoming fiber can be mapped to the same wavelength on any of the two outgoing fibers as shown in Fig. 1.17(b). Since each fiber carries two wavelengths, there are four wavelengths on a link. The wavelengths are relabeled as w_a, w_b, w_c, and w_d as shown in Fig. 1.17(c). Here, wavelengths w_0 and w_1 on fiber 0 are relabeled as w_a and w_b, respectively, and wavelengths w_0 and w_1 on fiber 1 are relabeled as w_c and w_d, respectively. If a single fiber is to carry the above four wavelengths, then the node is capable of shifting a wavelength to one of the two wavelengths. Figure 1.17(d) shows such a node with conversion degree 2, whose capability is the same as that of the node shown in Fig. 1.17(a).

1.7.4 Wavelength Rerouting

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

Example 5: *Consider the dynamic RWA problem in Example 2. Here, a lightpath from node 0 to node 2 cannot be established (see Fig. 1.16[a]), as there is no wavelength-continuous route available. If lightpath p_0 can be migrated from wavelength w_0 to wavelength w_1, then a lightpath can be established from node 0 to node 2 on wavelength w_0.* □

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

1.7.5 Virtual Topology Design

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

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

1.7.6 Virtual Topology Reconfiguration

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

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

Example 6: *Consider a network with four nodes labeled 0, 1, 2, and 3. The traffic demand between node pairs is specified in matrix* T_1 *as shown in Fig. 1.18(a). The value specified at row* i *and column* j *is the traffic flow between node* i *and node* j*. It is assumed that the demand between node pairs is bidirectional. The virtual topology* V_1 *designed to meet the demand in* T_1 *is shown in Fig. 1.18(b). The traffic is assumed to be routed on the virtual topology using the shortest paths. Accordingly, node pairs* $< 0,1 >$, $< 0,2 >$, $< 0,3 >$, $< 1,2 >$, $< 1,3 >$, *and* $< 2,3 >$ *use the paths* 0–1, 0–3–2, 0–3, 1–3–2, 1–3, *and* 2–3, *respectively. The load on each of the lightpaths is shown over the corresponding edge in Fig. 1.18(b). The network congestion is measured to be 7, which corresponds to the maximum load carried by the lightpath between node 2 and node 3. This lightpath carries the traffic between node pairs* $< 0,2 >$, $< 1,2 >$, *and* $< 2,3 >$.

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

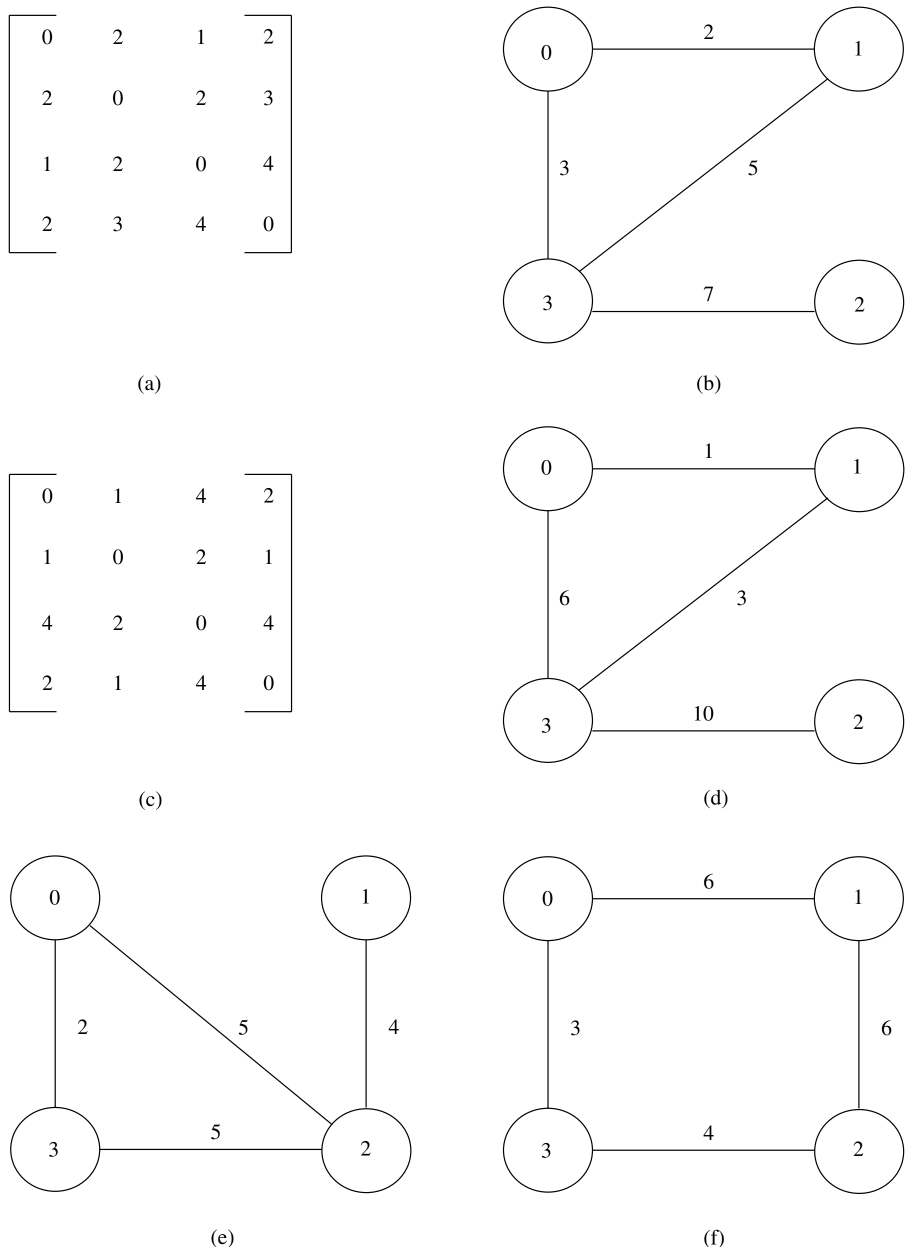

Figure 1.18 Illustration of virtual topology and its reconfiguration. (a) Traffic demand matrix T_1. (b) Virtual topology V_1 for T_1. (c) Traffic demand matrix T_2. (d) Virtual topology V_1 for T_2. (e) Virtual topology V_2 for T_2. (f) Virtual topology V_3 for T_2.

demands, then the network congestion increases to 10. Figure 1.18(d) shows the lightpath loads for the new traffic matrix T_2. *If the virtual topology is allowed to change to a new topology* V_2 *as shown in Fig. 1.18(e), then the network congestion is drastically reduced to 5. Here, node pairs* $< 0, 1 >$, $< 0, 2 >$, $< 0, 3 >$, $< 1, 2 >$, $< 1, 3 >$, *and* $< 2, 3 >$ *use the shortest paths* 0–2–1, 0–2, 0–3, 1–2, 1–2–3, *and* 2–3, *respectively to route their traffic. Comparing* V_1 *and* V_2, *there are two lightpaths deleted and two lightpaths added, and therefore, the number of changes is 4. Alternatively, by using virtual topology* V_3 *(shown in Fig. 1.18[f]), traffic* T_2 *can be routed. Topology* V_3 *is better than* V_2 *from the perspective of number of changes, as* V_3 *requires removal and addition of only one lightpath. However, the network congestion in* V_3 *is 6, which is slightly more that in* V_2. *In* V_3, *node pairs* $< 0, 1 >$, $< 0, 2 >$, $< 0, 3 >$, $< 1, 2 >$, $< 1, 3 >$, *and* $< 2, 3 >$ *use the shortest paths* 0–1, 0–1–2, 0–3, 1–2, 1–0–3, *and* 2–3, *respectively.* □

1.7.7 Survivable Networks

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

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

1.7.8 Optical Multicast Routing

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

Example 7: *Consider the network shown in Fig. 1.19(a). Here, node 0 is source of a multicast session and nodes 2 through 6 are destinations. If unicast paths (lightpaths) are used for multicasting, then node 0 needs to establish five lightpaths, each to a different destination as shown in Fig. 1.19(b). The resources required are thus five wavelengths and 10 wavelength channels. Also, node 0 needs to transmit the same data five times, each time to a different destination. In Fig. 1.19(c), it is assumed that all nodes in the network have both DaC and splitting capabilities. Here, a multicast tree (a treelike structure that contains the source and all its destinations) is constructed for routing multicast traffic. Node 2 uses its splitting capability to transmit the incoming optical signal on link* $0 \rightarrow 2$ *to the links* $2 \rightarrow 3$, $2 \rightarrow 4$, *and*

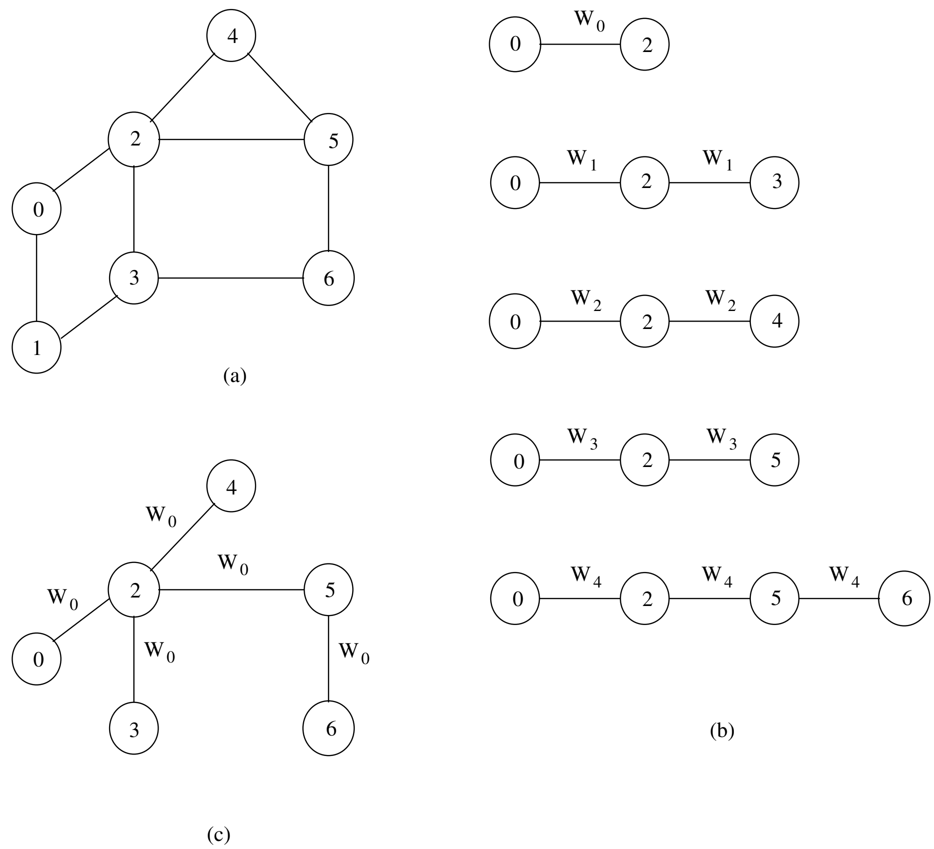

Figure 1.19 Benefit of constructing a multicast tree for multicast routing than using unicast paths. (a) Example network. (b) Multicast routing using unicast paths. (c) Multicast routing using split and tap capable nodes.

$2 \rightarrow 5$. *Node 5 taps a part of the signal and forwards the remaining signal onto the link* $5 \rightarrow 6$. *The resources required in this case are one wavelength and five wavelength channels. Also, node 0 needs to transmit data only once.* □

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

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

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

1.7.9 Network Control and Management

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

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

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

1.7.10 Transmission Impairment

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

Optimizing Amplifier Placements

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

1.7.11 Ring Networks and Traffic Grooming

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

SONET rings are a popular architecture used in many existing networks. A SONET ring is built from SONET add drop multiplexers (SADMs). An SADM is capable of multiplexing and demultiplexing a number of low-speed SONET connections to and from a high-speed connection onto an optical ring using time division multiplexing (TDM) technique. For example, four OC-3 connections at 155.52 Mb/s can be multiplexed onto an OC-12 ring at 622.08 Mb/s. Similarly, 16 OC-3 connections at 155.52 Mb/s can be multiplexed onto an OC-48 ring at 2.48832 Gb/s and four OC-12 connections at 622.08 Mb/s can be multiplexed onto an OC-48 ring at 2.48832 Gb/s. By applying WDM technology, several rings each on a different wavelength can be formed on a physical ring network. Similar to SONET ring networks, WDM ring networks can also be classified as two-fiber unidirectional ring (UR-2), two-fiber bidirectional ring (BR-2), and four-fiber bidirectional ring (BR-4) networks. In a UR-2 network, there are two unidirectional rings in opposite directions. One ring carries traffic during normal operation and the other ring is used for service protection against failures. In a BR-2 network, half of the wavelengths in one ring and half of the wavelengths in the other ring are used for routing normal traffic. All the remaining wavelengths are used for protection purposes. In a BR-4 network, one pair of rings in opposite directions is used for carrying normal traffic while the other pair of rings is used for protection purpose. A multifiber ring network uses a number of fibers in multiples of two or four. In a multifiber ring network, WXCs can be used in place of WADMs to provide flexibility in switching a wavelength to any of the fibers. Another useful architecture for backbone networks is an interconnected ring network, wherein several rings are interconnected. The nodes in a ring can employ WADMs except for those nodes connecting different rings. Such bridging nodes can employ WXCs.

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

Traffic Grooming

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

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

Example 8: *Consider a unidirectional ring network with four nodes labeled 0 through 3. Assume that each node has two units of traffic (that is, two SONET OC-x connections) to every other node. The multiplexing factor is assumed to be 4. Therefore, a wavelength can carry four OC-x connections. Routing all the OC-x connections on the unidirectional ring requires that each fiber link carry 12 OC-x connections. This means that a minimum of three wavelengths are required to carry the traffic. We give below two different ways of grooming the traffic onto wavelengths. Accordingly, we get two different WDM-SONET ring networks. While both networks use only three wavelengths, they require different numbers of SADMs.*

Figure 1.20(a) shows a WDM-SONET ring obtained by way of aggregating traffic as depicted in Fig. 1.20(b), where an entry in row i and column j indicates the wavelength onto which the traffic from node i to node j is multiplexed. Figure 1.20(c) shows the traffic carried on each of three wavelengths on different fiber links. The entry corresponding to fiber link $x \rightarrow y$ and wavelength w_k is the set of node pairs

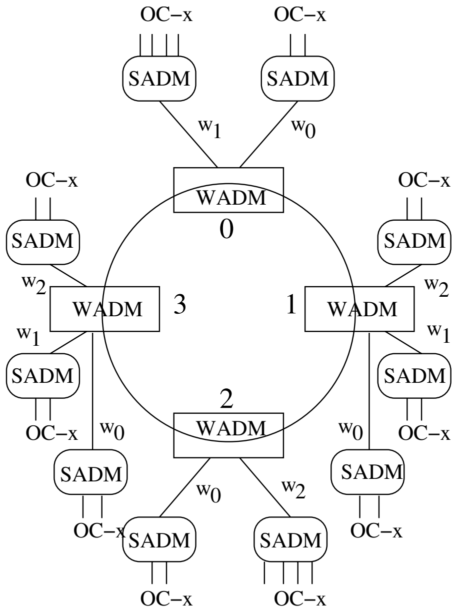

(a)

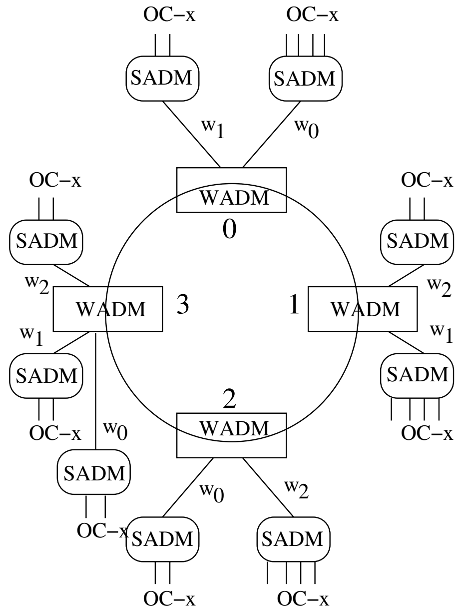

(d)

	0	1	2	3
0		w_1	w_0	w_1
1	w_1		w_2	w_0
2	w_0	w_2		w_2
3	w_1	w_0	w_2	

(b)

	0	1	2	3
0		w_1	w_0	w_0
1	w_1		w_2	w_1
2	w_0	w_2		w_2
3	w_0	w_1	w_2	

(e)

	w_0	w_1	w_2
0–>1	<0,2> <3,1>	<0,3> <0,1>	<2,1> <3,2>
1–>2	<0,2> <1,3>	<0,3> <1,0>	<3,2> <1,2>
2–>3	<1,3> <2,0>	<0,3> <1,0>	<2,1> <2,3>
3–>0	<3,1> <2,0>	<1,0> <3,0>	<2,1> <3,2>

(c)

	w_0	w_1	w_2
0–>1	<0,2> <0,3>	<3,1> <0,1>	<2,1> <3,2>
1–>2	<0,2> <0,3>	<1,3> <1,0>	<3,2> <1,2>
2–>3	<2,0> <0,3>	<1,3> <1,0>	<2,1> <2,3>
3–>0	<2,0> <3,0>	<3,1> <1,0>	<2,1> <3,2>

(f)

Figure 1.20 A WDM-SONET ring network. (a) WDM-SONET ring (design 1). (b) Traffic aggregation (design 1). (c) Traffic carried by wavelengths on different fibers (design 1). (d) WDM-SONET ring (design 2). (e) Traffic aggregation (design 2). (f) Traffic carried by wavelengths on different fibers (design 2).

$< i, j >$ *whose traffic is carried by fiber link* $x \rightarrow y$ *in the SONET ring corresponding to* w_k. *Since each wavelength on a fiber carries only four OC-x connections, the traffic grooming is valid. From Fig. 1.20(b), it is clear that node 0 requires two SADMs corresponding to* w_0 *and* w_1, *node 1 requires three SADMs corresponding to* w_0, w_1, *and* w_2, *node 2 requires two SADMs corresponding to* w_0 *and* w_2, *and node 3 requires three SADMs corresponding to* w_0, w_1, *and* w_2. *Therefore, the ring requires 10 SADMs.*

An alternative design is shown in Fig. 1.20(d), wherein the WDM-SONET ring requires only nine SADMs. The traffic grooming is valid because, for any wavelength on a fiber, the number of OC-x connections carried does not exceed the multiplexing factor. The details of wavelength assignment and the traffic carried on fiber links in different wavelength rings are depicted in Fig. 1.20(e) and Fig. 1.20(f), respectively.

□

In typical telecommunication backbone networks, multiple rings are interconnected together to provide large geographical coverage. In [181], traffic grooming in interconnected multi-WDM rings is studied addressing two issues: how to interconnect WDM ring networks, and how to groom the traffic in interconnected rings. In [171] and [199], traffic grooming in mesh topology WDM networks is examined. Although mesh topology WDM networks will be of greater importance in the future—at least in the near future—ring topology WDM networks are viable because SONET/SDH architectures are ring-oriented.

1.7.12 Virtual Private Networks over WDM Optical Networks

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

Example 9: *Consider a wide area WDM backbone network with five wavelength routed nodes (A, B, C, D, E) as shown in Fig. 1.21. Each node is equipped with*

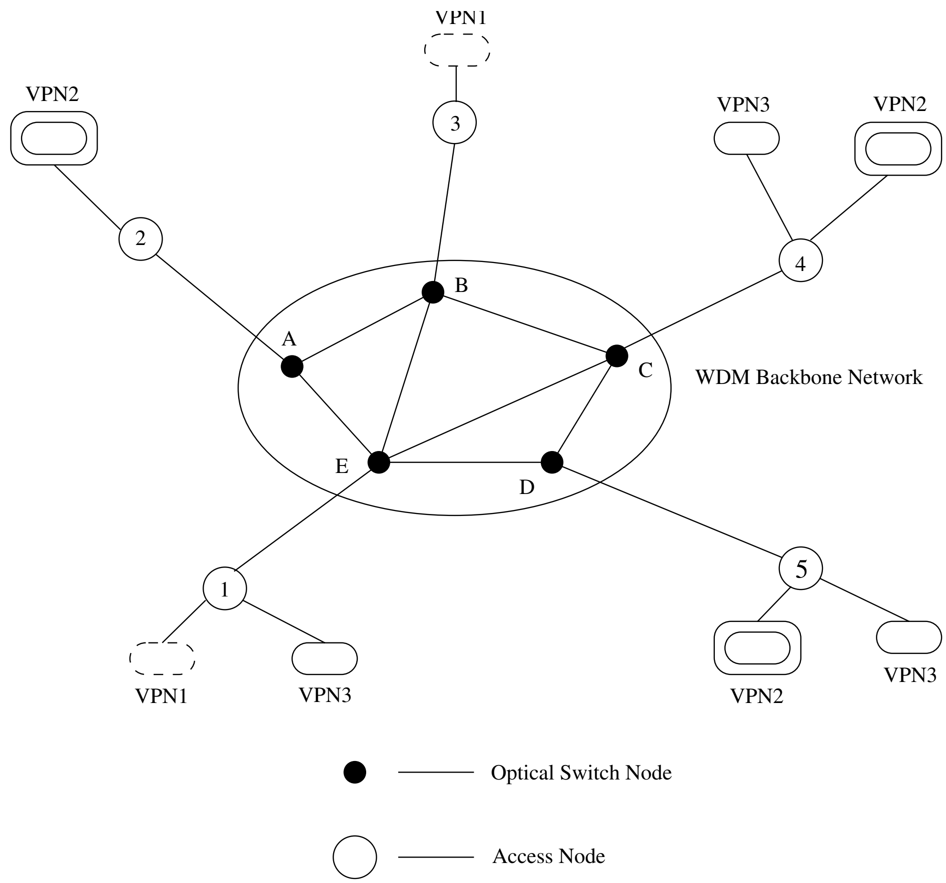

Figure 1.21 A virtual private network.

some transmitters and receivers. The access nodes at the border of the backbone provide the interface (electronic–optical, optical–electronic conversion and electronic buffering when contention arises for a lightpath) to the end users, which may be branch offices (or regional networks) that feed into the WDM backbone network. Here there are three VPNs: VPN1 is specified by a topology that consists of two access nodes (1, 3), VPN2 is specified by a topology consisting of three access nodes (2, 4, 5), and VPN3 is specified by a topology consisting of three access nodes (1, 4, 5).

□

Given the traffic demand from the three VPNs, the problem of VPN design over WDM optical networks involves establishing lightpaths between (a) nodes $< 1, 3 >$ for VPN1, (b) the pairs of nodes $< 2, 4 >, < 4, 5 >,$ and $< 2, 5 >$ for VPN2, and (c) the pairs of nodes $< 1, 4 >, < 4, 5 >,$ and $< 1, 5 >$ subject to physical constraints (limited transmitters, receivers, wavelength converters, etc.). Note that this problem is similar to the virtual topology design problem mentioned earlier; instead of one we need to design a set of logical topologies. In [137] the problem of VPN design over WDM networks is formulated as a mixed-integer linear problem. In [131] a framework for supporting VPNs with different QoS requirements over WDM optical networks is presented. Depending on the QoS requirement of a VPN, one or more of the following three lightpaths is used to carry the VPN's traffic: (1) a dedicated lightpath (which is used by exactly one VPN), (2) a shared lightpath (which is shared by multiple VPNs), and (3) a multi-hop lightpath (a series of lightpaths with optical–electronic and electronic–optical conversion at the junction between two lightpaths). Three types of traffic are defined: Type1, carried over dedicated lightpaths as far as possible; Type2, carried over shared lightpaths; and Type3, carried over multi-hop lightpaths. A VPN's traffic specifies information on each of the three different types. The solution (heuristic) algorithm tries to accommodate diverse QoS (delay) requirements of different VPNs (by trying to establish dedicated lightpaths for all Type1 traffic, shared lightpaths for Type2 traffic, and multi-hop lightpaths for Type3 traffic), maximizing the total traffic carried on the backbone network.

1.7.13 Access Networks

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

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

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

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

In the past, cable television networks have been used as downstream networks for delivering analog television. Newer cable television systems are HFC networks in which transmission of information from the cable television company head end to remote (fiber) nodes is over fiber optic links and then over a coaxial cable network to subscribers' homes. The current 300- to 450-MHz coax cables will be replaced by 750-MHz coax cables, thereby upgrading the capacity from 50 to 75 6-MHz channels to 125 6-MHz channels (in North America channels are 6 MHz wide). Seventy-five of the 125 channels are for transmitting analog television and the new 50 channels using sophisticated modulation techniques provide a total 2 Gb/s (50×40 Mb/s per channel) of new bandwidth. However, this requires the cable operators to replace all the existing cables with 750-MHz coax and unidirectional analog amplifiers with bidirectional split-band amplifiers for allowing information to flow in both up/down streams. As in the case of ADSL, the upstream/downstream transmission rates are asymmetric here for two-way services (such as high-speed Internet access via cable modem, cable telephony, and pay-per-view video). Note that in DSL the final segment is a point-to-point local loop using twisted pairs, whereas in HFC the final segment is a shared coaxial cable. While DSL is fully switched, HFC uses a shared medium without switching.

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

PON technology has several attractive features. A PON allows for a longer distance (over 20 km) between a central office and customer premises compared to that (about 5.5 km) in DSL. A PON provides higher bandwidth due to deeper fiber penetration (FTTH; however, FTTC is the most economical deployment today).

PONs use passive optical devices instead of active devices (such as multiplexers and demultiplexers) in the splitting locations, thus relieving network operators from maintaining them and providing power to them. Moreover, the fiber infrastructure being transparent to bit rates and modulation formats, PONs allow easy upgrades to higher bit rates or additional wavelengths without changing the infrastructure itself.

We now describe two types of PON architectures: WDM PONs and TDM PONs (see Fig. 1.22). A WDM PON uses a wavelength routing device (for example, arrayed waveguide grating, AWG) at the passive optical splitter to provide a single, dedicated wavelength to every ONU. Different multiplexing techniques can be used to implement bidirectional (up/down stream) transmission. When a single fiber is dedicated to every customer, the downstream traffic can be transmitted employing wavelengths belonging to the 1.55-micron optical window, while for upstream traffic the wavelengths in the 1.3-micron window can be used. When two fibers are dedi-

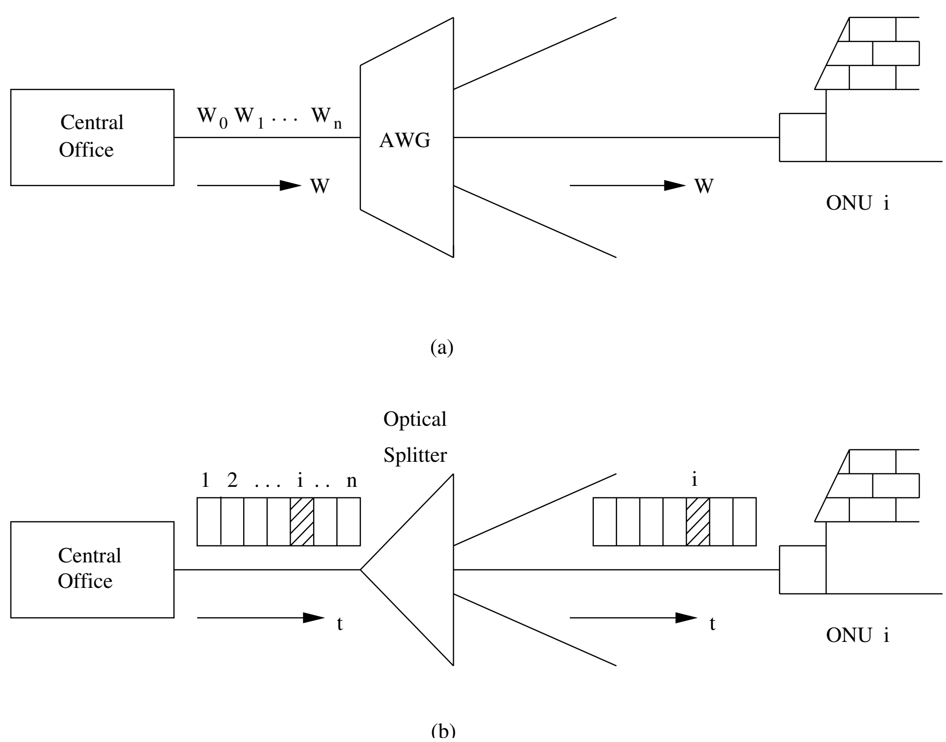

Figure 1.22 FTTH PON architectures. (a) WDM PON. (b) TDM PON.

cated to every customer, each can carry the traffic in one direction (space division bidirectionality). Unfortunately, implementing WDM PONs is a very costly affair. A *time-division multiplexed* (TDM) PON uses an optical power splitter instead of wavelength routing device as in WDM PON, and can be implemented with low-cost components. Since many customers share the fiber, an important issue is how different customers should access the fiber bandwidth. A very simple way is for all customers to operate in the same wavelength band and share the fiber bandwidth in time slots. A TDM PON uses TDM for downstream traffic and *time-division multiple access* (TDMA) for upstream traffic. A large international group of network operators and equipment vendors—the Full Services Access Network (FSAN) Consortium—has been working to standardize PON access systems [3]. By standardizing fiber access systems that will be purchased by a large number of network operators, the FSAN hopes that PON access systems will become more affordable.

1.8 NEXT-GENERATION OPTICAL INTERNET NETWORKS

WDM-based optical networks are becoming the right choice for the next-generation Internet networks to transport high-speed IP traffic. In the first phase, lightpath-based circuit switching WDM networks are deployed as a means of carrying IP traffic. SONET and ATM networks have been widely deployed in the transport networks. SONET systems have several attractive features such as high-speed transmission and network survivability. ATM networks have several attractive features such as flexible bandwidth allocation and QoS support. Therefore, ATM and/or SONET layers can be used in between the IP layer and the WDM optical layer for transporting IP packets. A major drawback of this multilayer approach is that it incurs increased control and management overhead.

WDM technology is evolving from circuit switching technology to burst switching and packet switching technologies. The granularity of the basic switching entity is large in circuit switching networks, medium in burst switching networks, and small in packet switching networks. While circuit switching technology is mature, the other technologies are not. In a circuit switching network, a wavelength channel on a link is used by a circuit (lightpath) for a long time, until it is torn down. In a burst (packet) switching network, a wavelength channel on a link is reserved only for the duration of the burst (packet). The bandwidth utilization in burst and packet switching networks is higher when compared to that in circuit switching networks. This is because, the former networks use *statistical multiplexing* while the latter does not.

In a burst switching network, the basic switching entity is a burst. A number of IP packets which are destined for the same egress router are assembled into a burst at the ingress router. The major challenges in burst switching networks include the design of cost-effective and fast switches, burst switching protocols, and wavelength channel scheduling. In a packet switching network, the basic switching entity is a packet. The major challenges in packet switching networks include the design of cost-effective and fast switches, packet synchronization, and contention resolution. Since optical processing is technologically and economically infeasible, the packet/burst header is processed electronically while the payload is switched optically. Since optical random access memory (RAM) is not available, a packet or burst cannot be buffered optically for a long time. A possible way is to use *fiber delay lines* (FDLs) to buffer (delay) a packet or burst for a short time. As we will see in Chapter 9, a *multiprotocol label switching* (MPLS)[4] framework has several advantages, such as traffic engineering, explicit path routing, fast packet forwarding, and network survivability. Due to the above advantages, future Internet networks employing circuit/burst/packet switching are likely to use the MPLS approach.

1.9 BOOK OVERVIEW

This chapter provided the "big picture" of the WDM optical networks. As mentioned earlier, this book focuses on wavelength routed (wide area) optical networks. In this chapter we introduced several issues concerning the design of these networks. The rest of the book addresses some of these issues in depth and is organized into eight chapters followed by a bibliography and appendices.

Chapter 2 deals with the problem of routing and wavelength assignment (RWA). It first presents different methods for route selection and wavelength selection with illustrations and then describes several important RWA algorithms in detail with suitable examples. It also presents distributed control protocols for wavelength routing and several solutions for the connection fairness problem which arises as longer-hop connections are more likely to be blocked (rejected) compared to shorter-hop connections.

Chapter 3 deals with wavelength-convertible networks. It first highlights the benefit of employing wavelength converters and then explains different wavelength-convertible switch architectures. It then presents a routing algorithm in convertible networks and some analytical models which predict the performance of these networks. It also addresses converter placement and converter allocation problems.

[4]See Appendix D.

Chapter 4 deals with wavelength rerouting algorithms to improve wavelength channel utilization. It first describes several lightpath migration operations and then presents several algorithms for wavelength rerouting in single-fiber and multi-fiber networks with and without wavelength conversion. It also analyzes the computational complexity of these algorithms.

Chapter 5 deals with the problem of virtual topology design, which is an important and challenging one. It first discusses the limitations on choosing the best possible virtual topology and then explains the different virtual topology design subproblems with necessary illustrations. It also presents a mixed-integer linear programming formulation of the virtual topology design problem, and several heuristic solutions for the problem and its variants with suitable examples. It finally introduces the problem of designing multifiber networks and also presents a solution for this problem.

Chapter 6 deals with the problem of virtual topology reconfiguration. It first presents several schemes to reconfigure the existing virtual topology in response to a change in the traffic. It then discusses the exact solution methods and also several heuristic methods for reconfiguring a virtual topology. It also presents an approach to reconfigure a virtual topology when a network component fails.

Chapter 7 deals with the problem of network provisioning and survivability. It presents a classification of lightpath restoration methods which are employed to reroute affected traffic upon a network component failure. It also presents an integer linear programming formulation of the survivable network design problem, and it discusses several algorithms for designing survivable networks. It finally explains survivability mechanisms in WDM ring networks.

Chapter 8 deals with optical multicast routing. It first describes architecture of multicast-capable switches which help in carrying out multicast routing in an efficient way. It then presents several algorithms for multicast routing in networks with full and sparse splitting capabilities.

Chapter 9 deals with three different WDM technologies: circuit switching, burst switching, and packet switching, that are likely to be adopted in the next-generation Internet networks. It presents different protocol stack options and packet encapsulation and framing techniques that can be used for transporting IP traffic over the WDM (optical) layer. It then presents an architecture and protocols for optical burst switching networks. It also describes a burst switching protocol to support differentiated services. It then explains two switch architectures and contention resolution techniques for optical packet switching networks. It finally explains how the traditional IP/MPLS framework can be extended to WDM-based optical networks.

1.10 PROBLEMS

1. What is the bandwidth of an OC-12 connection?

2. Show that 0.17 micron of spectrum at a wavelength of 1.3 microns corresponds to 30 THz of bandwidth.

3. The refractive indices of core and cladding are 1.47 and 1.44, respectively. Compute the critical angle.

4. Explain how the effects of (a) attenuation, (b) dispersion, and (c) nonlinearities in a fiber can be reduced.

5. What do you understand by "electronic bottleneck"? Suggest a way to eliminate it.

6. What are the differences between FDM and WDM?

7. Compare and contrast electronic regenerators and optical amplifiers (EDFAs).

8. List the different WDM optical components/systems available in the market by visiting the Web sites given in Appendix A.

9. Show a design of an 8×8 star coupler which has three stages of 2×2 couplers with four couplers per stage. If each node transmits with a power P, what is the power received by another node? Show that the overall splitting loss is 10 log 8 dB.

10. Consider the following MAC protocol for single-hop broadcast-and-select networks: "It assumes that each node transmits on a unique fixed wavelength, different from the other nodes, and uses a tunable filter to select one among all the different wavelengths for reception. In addition, there is a control channel shared by all the nodes, in which each node is assigned a time slot for transmission. A node transmits the packet header in its assigned time slot on the control channel and follows this by transmitting the actual packet on its data channel. After listening to the packet header on the control channel, the intended receiver tunes to the appropriate channel to receive the packet." Can collisions occur here? Suggest a way to deal with destination conflicts.

11. Show a design of an $N \times N$ WXC. How many simultaneous connections can it support through itself?

12. Consider the physical and logical topologies shown in Fig. 1.11. How is a packet from node 1 to node 4 routed?

13. Explain why the two routing constraints, inseparability and distinct source combining, do not apply to wavelength routed networks.

14. What factors limit throughput in the following networks: (a) broadcast-and-select network, (b) wavelength routed network, and (c) linear lightwave network.

15. Let an N-node WDM-SONET ring use W number of SADMs at each node, each for one wavelength. If D is the number of SADMs required when using WADMs and traffic grooming to support the given traffic pattern, what is the percentage saving on the number of SADMs? Is it always possible to minimize both the number of wavelengths (W) and the total number of SADMs (D) at the same time?

16. Consider a combined WDM-SONET ring with four nodes and links connected as $0 \rightarrow 1 \rightarrow 2 \rightarrow 3 \rightarrow 0$. It is required to route one OC-48 SONET connection from each node to every other node in the ring. The wavelengths operate at a speed of OC-48. (a) Determine the lower bound on the number of wavelengths and SONET ADMs required. (b) Give a solution for grooming and wavelength assignment so that a minimum number of wavelengths and SONET ADMs is used. (c) How many wavelengths and SONET ADMs are required if point-to-point WDM links are used?

17. What are the differences between a WDM PON and a TDM PON?