Part I

Introduction

1 Historical overview of optical networks

Optical fiber provides an unprecedented bandwidth potential that is far in excess of any other known transmission medium. A single strand of fiber offers a total bandwidth of 25 000 GHz. To put this potential into perspective, it is worthwhile to note that the total bandwidth of radio on Earth is not more than 25 GHz (Green, 1996). Apart from its enormous bandwidth, optical fiber provides additional advantages such as low attenuation loss (Payne and Stern, 1986). Optical networks aim at exploiting the unique properties of fiber in an efficient and cost-effective manner.

1.1 Optical point-to-point links

The huge bandwidth potential of optical fiber has been long recognized. Optical fiber has been widely deployed to build high-speed optical networks using fiber links to interconnect geographically distributed network nodes. Optical networks have come a long way. In the early 1980s, optical fiber was primarily used to build and study point-to-point transmission systems (Hill, 1990). As shown in Fig. 1.1(a), an optical point-to-point link provides an optical single-hop connection between two nodes without any (electrical) intermediate node in between. Optical point-to-point links may be viewed as the beginning of optical networks. Optical point-to-point links may be used to interconnect two different sites for data transmission and reception. At the transmitting side, the electrical data is converted into an optical signal (EO conversion) and subsequently sent on the optical fiber. At the receiving side, the arriving optical signal is converted back into the electrical domain (OE conversion) for electronic processing and storage. To interconnect more than two network nodes, multiple optical single-hop point-to-point links may be used to form various network topologies (e.g., star and ring networks). Figure 1.1(b) shows how multiple optical point-to-point links can be combined by means of a star coupler to build optical single-hop star networks (Mukherjee, 1992). The star coupler is basically an optical device that combines all incoming optical signals and equally distributes them among all its output ports. In other words, the star coupler is an optical broadcast device where an optical signal arriving at any input port is forwarded to all output ports without undergoing any EO or OE conversion at the star coupler. Similar to optical point-to-point links, optical single-hop star networks make use of EO conversion at the transmitting side and OE conversion at the receiving side. Besides

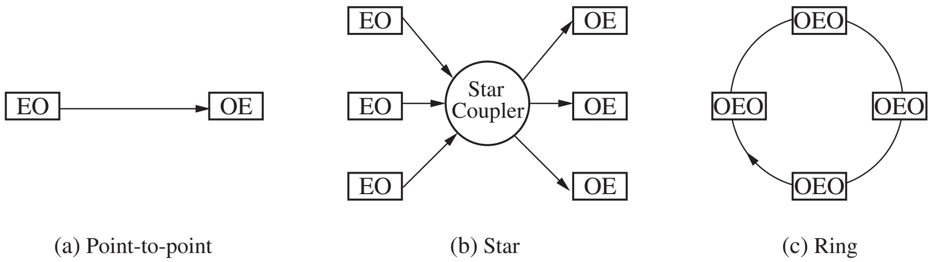

Figure 1.1 Optical single-hop connections: (a) point-to-point, (b) star, and (c) ring configurations.

optical stars, optical ring networks can be realized by interconnecting each pair of adjacent ring nodes with a separate optical single-hop point-to-point fiber link, as depicted in Fig. 1.1(c). In the resultant optical ring network, each node performs OE conversion for incoming signals and EO conversion for outgoing signals. The combined OE and EO conversion is usually referred to as OEO conversion. A good example of an optical ring network with OEO conversion at each node is the fiber distributed data interface (FDDI) standard, which can be found in today's existing optical network infrastructure (Ross, 1986; Jain, 1993).

1.2 SONET/SDH

One of the most important standards for optical point-to-point links is the Synchronous Optical Network (SONET) standard and its closely related synchronous digital hierarchy (SDH) standard. The SONET standardization began during 1985 and the first standard was completed in June 1988 (Ballart and Ching, 1989). The goals of the SONET standard were to specify optical point-to-point transmission signal interfaces that allow interconnection of fiber optics transmission systems of different carriers and manufacturers, easy access to tributary signals, direct optical interfaces on terminals, and to provide new network features. SONET defines standard optical signals, a synchronous frame structure for time division multiplexed (TDM) digital traffic, and network operation procedures.

SONET is based on a digital TDM signal hierarchy where a periodically recurring time frame of 125 μs can carry payload traffic of various rates. Besides payload traffic, the SONET frame structure contains several overhead bytes to perform a wide range of important network operations such as error monitoring, network maintenance, and channel provisioning.

SONET is now globally deployed by a large number of major network operators. Typically, SONET point-to-point links are used in ring configurations to form optical ring networks with OEO conversion at each node, similar to the one depicted in Fig. 1.1(c). In SONET rings there are two main types of OEO nodes: the add-drop multiplexer (ADM) and the digital cross-connect system (DCS). The ADM usually connects to several SONET end devices and aggregates or splits SONET traffic at various speeds. The DCS is a SONET device that adds and drops individual SONET channels at any location. One major difference between an ADM and a DCS is that the DCS can be used to interconnect a larger number of links. The DCS is often used to interconnect SONET rings (Goralski, 1997).

1.3 Multiplexing: TDM, SDM, and WDM

Given the huge bandwidth of optical fiber, it is unlikely that a single client or application will require the entire bandwidth. Instead, traffic of multiple different sources may share the fiber bandwidth by means of multiplexing. Multiplexing is a technique that allows multiple traffic sources to share a common transmission medium. In the context of optical networks, three main multiplexing approaches have been deployed to share the bandwidth of optical fiber: (1) time division multiplexing (TDM), (2) space division multiplexing (SDM), and (3) wavelength division multiplexing (WDM).

- **Time division multiplexing:** We have already seen that SONET is an important example for optical networks that deploy TDM on the underlying point-to-point fiber links. Traditional TDM is a well-understood technique and has been used in many electronic network architectures throughout the more than 50-year history of digital communications (Green, 1996). In the context of high-speed optical networks, however, TDM is under pressure from the so-called "electro-optical" bottleneck. This is due to the fact that the optical TDM signal carries the aggregate traffic of multiple different clients and each TDM network node must be able to operate at the aggregate line rate rather than the subrate that corresponds to the traffic originating from or destined for a given individual node. Clearly, the aggregate line rate cannot scale to arbitrarily high values but is limited by the fastest available electronic transmitting, receiving, and processing technology. As a result, TDM faces severe problems to fully exploit the enormous bandwidth of optical fiber, as further outlined in Section 1.4.
- **Space division multiplexing:** One straightforward approach to avoid the electro-optical bottleneck is SDM, where multiple fibers are used in parallel instead of a single fiber. Each of these parallel fibers may operate at any arbitrary line rate (e.g., electronic peak rate). SDM is well suited for short-distance transmissions but becomes less practical and more costly for increasing distances due to the fact that multiple fibers need to be installed and operated.
- **Wavelength division multiplexing:** WDM appears to be the most promising approach to tap into the vast amount of fiber bandwidth while avoiding the aforementioned shortcomings of TDM and SDM. WDM can be thought of as optical frequency division multiplexing (FDM), where traffic from each client is sent on a different carrier frequency. In optical WDM networks the term wavelength is usually used instead of frequency, but the principle remains the same. As shown in Fig. 1.2, in optical WDM networks each transmitter i sends on a separate wavelength λ_i, where $1 \leq i \leq N$. At the transmitting side, a wavelength multiplexer collects all wavelengths and feeds them onto a common outgoing fiber. At the receiving side, a wavelength demultiplexer separates the wavelengths and forwards each wavelength λ_i to a different

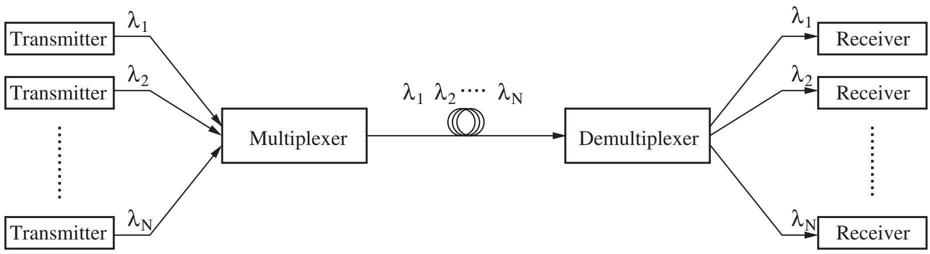

Figure 1.2 Wavelength division multiplexing (WDM).

receiver i. Unlike for TDM, each wavelength channel may operate at any arbitrary line rate well below the aggregate TDM line rate. By using multiple wavelengths the huge bandwidth potential of optical fiber can be exploited. As opposed to SDM, WDM takes full advantage of the bandwidth potential of a single fiber and does not require multiple fibers to be installed and operated in parallel, resulting in significant cost savings. Optical WDM networks have been attracting a great deal of attention by network operators, manufacturers, and research groups around the world, as discussed in Section 1.5.

It is worthwhile to note that in existing and emerging optical networks all three multiplexing techniques are used together to realize high-performance network and node architectures. By capitalizing on the respective strengths of TDM, SDM, and WDM and gaining a better understanding of their duality relationships, novel space–time–wavelength switching node structures may be found that enable future performance-enhanced optical networks (Kobayashi and Kaminow, 1996).

1.4 Optical TDM networks

Progress on the development of very short optical pulse technology enables the realization of optical time division multiplexing (OTDM) networks. OTDM networks aim at operating at an aggregate line rate of 100 Gb/s and above. At such high data rates, the transmission properties of optical fiber come into play and need to be taken care of properly. In particular, dispersion has a major impact on the achievable bandwidth-distance product of OTDM networks. Simply put, dispersion makes different parts of the optical signal travel at different speeds along a fiber link. As a result, parts of the optical signal arrive at the receiving side at different time instances, resulting in the so-called intersymbol interference (ISI), where the optical power of a given received bit interferes with that of adjacent bits. As a consequence, the optical power level of adjacent bits is changed and may lead to wrong decisions at the threshold detector of the receiver and transmission errors. This effect is exacerbated for increasing data rates and fiber lengths, translating into a decreasing bandwidth-distance product. Therefore, OTDM networks appear better suited for short-range networks where the impact of dispersion is kept small. For long-distance networks, dispersion effects can be avoided by the use of the so-called “soliton” propagation. With the soliton propagation, dispersion effects are canceled out by nonlinear effects of optical fiber, resulting in a significantly improved bandwidth-distance product (Green, 1996).

OTDM networks have been receiving considerable attention due to the progress of optical short-pulse technology. Apart from the aforementioned transmission issues in very-high-speed OTDM networks, other important topics have been addressed. Among others, research efforts have been focusing on the design of OTDM network and node architectures and advanced components (e.g., ultra-short-pulse fiber laser, soliton compression source, and optical short-pulse storage loop; Barry et al., 1996).

OTDM networks suffer from two major disadvantages: (1) due to the underlying TDM operation, nodes need to be synchronized in order to start transmission in their assigned time slot and thus avoid collisions on the channel; more important, (2) OTDM networks do not provide *transparency*. Synchronization is a fundamental requirement of OTDM networks and becomes more challenging for increasing data rates of 100 Gb/s and above. As for the missing transparency, note that OTDM network clients are required to comply with the underlying TDM frame structure. As a result, the TDM frame structure dictates the transmission and reception of client traffic and thereby destroys the transparency against arbitrary client protocols in that clients need to match their traffic and protocols to the OTDM framing format. To build optical networks that are transparent against different protocols, the optical signal must be able to remain in the optical domain until it arrives at the destination. Clearly, this can be achieved by avoiding OEO conversions at intermediate nodes. In doing so, data stays in the optical domain and is optically switched all the way from the source to the destination node, enabling end nodes to communicate with each other using their own protocol. By using optical switching components that are electronically controlled, transparent OTDM networks are getting closer to feasibility and deployment (Seo et al., 1996).

Transparent OTDM networks are an interesting type of optical network but are still in their infancy. Alternatively, optical WDM networks are a promising solution to realize transparent optical networks. In optical WDM networks, each wavelength channel may be operated separately without requiring network-wide synchronization, thus providing a transparent channel not only against protocol but also against data rate and modulation format, as opposed to OTDM networks. Compared to OTDM networks, optical WDM networks are widely considered more mature and are discussed at length in the following section.

1.5 Optical WDM networks

Optical WDM networks do not necessarily have to be transparent. Strictly speaking, optical WDM networks are networks that deploy optical WDM fiber links where each fiber link carries multiple wavelength channels rather than only a single one. Like any other optical network, optical WDM networks may consist of one or more simple point-to-point WDM links with OEO conversion at each network node, similar to the

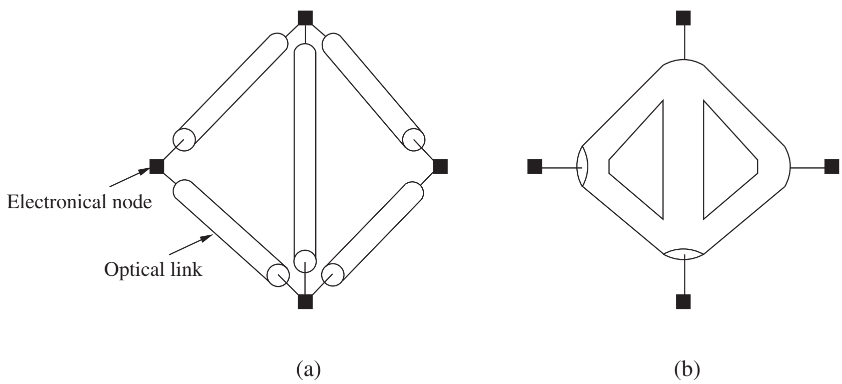

Figure 1.3 Optical WDM networks: (a) opaque and (b) transparent network architectures.

point-to-point link in Fig. 1.1(a) and ring network in Fig. 1.1(c). Optical WDM networks like that depicted in Fig. 1.1(c) are multihop networks where traffic traverses multiple intermediate nodes between any pair of source and destination nodes. Due to the fact that OEO conversion takes place at intermediate nodes, source and destination nodes are prevented from choosing their own protocol, line rate, and modulation format but have to follow the transmission requirements imposed by intermediate nodes. Thus, optical multihop networks with OEO conversion at intermediate nodes are unable to provide transparency to end nodes. In contrast, optical single-hop star networks similar to that shown in Fig. 1.1(b) inherently provide transparency for any pair of source and destination nodes. To see this, recall that the central star coupler is an optical device which does not perform any OEO conversion and leaves all in-transit traffic in the optical domain. As a result, end nodes are free to communicate with each other using their own agreed-upon protocol, data rate, and modulation format and are not hindered by any transmission requirements of intermediate nodes. The inherent transparency together with the simplicity of optical single-hop networks have led to a family of optical WDM networks known as *broadcast-and-select* networks (Mukherjee, 1992). Broadcast-and-select networks are WDM networks that are based on a central star coupler. Each transmitter is able to send on one or more different wavelengths. The star coupler broadcasts all incoming wavelengths to every receiver. Each receiver deploys an optical filter that is either fixed-tuned to a specific wavelength or tunable across multiple wavelengths. In either case, the optical filter selects a single wavelength and the destination is thus able to retrieve data sent on the selected wavelength.

Optical single-hop star WDM networks received considerable attention both from academia and industry. They are suitable for local area networks (LANs) and metropolitan area networks (MANs) where the number of nodes and distances are rather small. To build networks that are scalable in terms of nodes and coverage, optical WDM networks must be allowed to have any arbitrary topology (e.g., mesh topology). These networks can be categorized into two generations of optical WDM networks: (1) *opaque* and (2) *transparent* optical network architectures (Green, 1993). As shown in Fig. 1.3(a), in opaque optical WDM networks all wavelength channels are OEO converted at each network node, whereas in transparent optical WDM networks, as depicted in Fig. 1.3(b), intermediate nodes can be optically bypassed by dropping only a subset of wavelength channels into the electronical domain while leaving the remaining wavelength channels in the optical domain. Consequently, data sent on optically bypassing wavelengths can stay in the optical domain all the way between source and destination nodes, enabling transparent optical WDM networks. (For the sake of completeness, we note that there also exist so-called *translucent* optical networks. Translucent optical networks may be viewed as a combination of transparent and opaque optical networks where some network nodes provide optical bypassing capability while the remaining nodes perform OEO conversion of all wavelength channels. That is, translucent optical networks comprise both transparent and opaque network nodes.) Optical WDM networks with optical bypassing capability at intermediate nodes are widely referred to as *all-optical networks* (AONs) since the end-to-end path between source and destination is purely optical without undergoing any OEO conversion at intermediate nodes. AONs can be applied at any network hierarchy level. Unlike optical star networks, AONs are well suited for building not only optical WDM LANs and MANs but also optical WDM wide area networks (WANs). Due to their wide applicability, AONs have been attracting a great amount of attention by research groups and network operators worldwide.

1.5.1 All-optical networks

AONs are usually optical circuit-switched (OCS) networks, where circuits are switched by (intermediate) nodes at the granularity of a wavelength channel. Accordingly, OCS AONs are also called wavelength-routing networks. In wavelength-routing OCS networks, optical circuits are equivalent to wavelength channels. As mentioned earlier, AONs provide end-to-end optical paths by deploying all-optical node structures which allow the optical signal to stay partly in the optical domain. Such all-optical nodes are also called OOO nodes to emphasize the fact that they do not perform OEO conversion of all wavelength channels and in-transit traffic stays in the optical domain.

To understand the rationale behind the design of AONs, it is instructive to look at the similarities between AONs and SONET/SDH networks, which were discussed in Section 1.2. Note that both AONs and SONET/SDH networks are circuit-switched systems. The multiplexing, processing, and switching of TDM time slots in SONET/SDH networks are quite analogous to the multiplexing, processing, and switching of WDM wavelength channels in AONs. More precisely, in SONET/SDH networks lower-speed channels are multiplexed via byte interleaving to generate a higher-speed signal, where a SONET/SDH TDM signal can carry a mix of different traffic types and data rates. Furthermore, in SONET/SDH, ADMs and DCSs enable the manipulation and access to individual channels. Analogous functions can be found in AONs. As a matter of fact, the OOO node architectures used in AONs may be considered optical replica of the ADM and DCS node architectures of SONET/SDH, where electrical components are replaced with their optical counterparts. Accordingly, the resultant optical AON node architectures are called optical add-drop multiplexer (OADM) and optical cross-connect

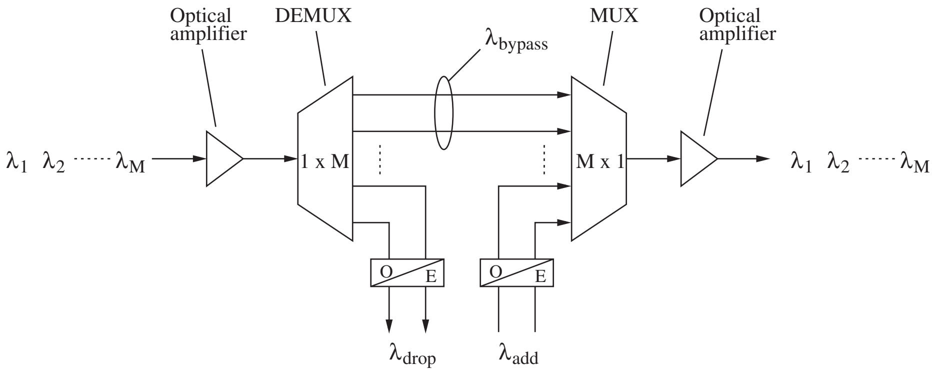

Figure 1.4 Optical add-drop multiplexer (OADM) with a single fiber link carrying M wavelengths.

(OXC), which are also known as wavelength ADM (WADM) and wavelength-selective cross-connect (WSXC), respectively (Maeda, 1998).

Figure 1.4 shows the basic schematic of an OADM with a single input/output fiber link that carries M different wavelength channels. At the input fiber the incoming optical signal comprising a total of M wavelengths $\lambda_1, \lambda_2, \ldots, \lambda_M$ is preamplified by means of an optical amplifier. A good choice for an optical amplifier is the so-called Erbium doped fiber amplifier (EDFA). A single EDFA is able to amplify multiple WDM wavelength channels simultaneously. After optical preamplification the WDM wavelength comb signal is partitioned into its M separate wavelengths by using a $1 \times M$ wavelength demultiplexer (DEMUX). In general, some bypass wavelengths λ_{bypass} remain in the optical domain and are thus able to optically bypass the local node. The remaining wavelengths λ_{drop} are dropped by means of OE conversion for electronic processing and/or storing at the local node. In doing so, the dropped wavelengths become available. The local node may use each of these freed wavelengths to insert local traffic on the available added wavelengths λ_{add}. Note that the dropped wavelengths λ_{drop} and added wavelengths λ_{add} operate at the same optical frequency but carry different traffic (locally dropped and added traffic, respectively). Subsequently, all M wavelengths are combined onto a common outgoing fiber by using an $M \times 1$ wavelength multiplexer (MUX). The composite optical WDM comb signal may be amplified by using another optical amplifier at the output fiber (e.g., EDFA).

The generic structure of an OXC with N input/output fiber links, each carrying M different wavelength channels, is shown in Fig. 1.5. An OXC is an $N \times N \times M$ component with N input fibers, N output fibers, and M wavelength channels on each fiber. A demultiplexer is attached to each input fiber (and optionally also an optical amplifier, similar to the previously discussed OADM). Each output from a demultiplexer goes into a separate wavelength layer. Each wavelength layer has a space division switch that directs each wavelength channel to a selected multiplexer. Each multiplexer collects light from M space division switches and multiplexes the wavelengths onto an output fiber. OXCs improve the flexibility and survivability of networks. They provide restoration and can reconfigure the network to accomodate traffic load changes and to compensate

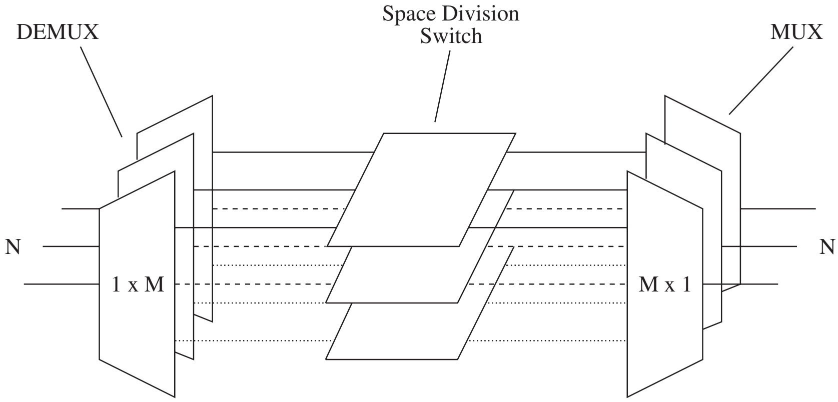

Figure 1.5 Optical cross-connect (OXC) with N fiber links, each carrying M wavelengths.

for link and/or node failures. An AON that deploys OXCs and OADMs is commonly referred to as an *optical transport network* (OTN). OTNs are able to provide substantial cost savings due to their flexibility, optical bypass capability, reconfigurability, and restoration (Sengupta et al., 2003).

AONs were examined by various research groups. Two major design goals of AONs are *scalability* and *modularity* (Brackett et al., 1993). Scalability is defined as the property that more nodes may always be added to the network, thereby permitting service to be offered to an arbitrarily large service domain. Modularity is defined as the property that only one more node needs to be added at a time. Besides scalability and modularity, AONs are intended to support a very large degree of wavelength reuse. Wavelength reuse allows wavelengths to be used many times in different locations throughout the network such that signals sent on the same wavelengths never interfere with each other. With wavelength reuse, bandwidth resources are used highly efficiently, resulting in an increased network capacity and decreased network costs. Toward the realization of scalable and modular AONs, significant progress has been made in the area of device technology, for example, acousto-optic tunable filters (AOTFs), multiwavelength lasers, multiwavelength receiver arrays, and other components (Brackett et al., 1993; Chidgey, 1994).

AONs are expected to support a number of different services and applications. For instance, provided services may comprise point-to-point as well as point-to-multipoint optical high-speed circuits. These services may be used to support applications such as voice, data, video, uncompressed high-definition TV (HDTV), medical imaging, and the interconnection of supercomputers (Alexander et al., 1993). AONs hold great promise to support all these different applications in a cost-effective fashion due to their transparency. To build large transparent AONs one must take the impact of physical transmission impairments on transparency into account, for example, signal-to-noise ratio (SNR), fiber nonlinearities, and crosstalk. For instance, the SNR poses limitations on the number of network nodes, and fiber nonlinearities constrain the number of used wavelengths and distances. As a result, transparency can be achieved only to a certain