Optical Communication Networks

EE654

Lectuer - 6

Spring 2016

WDM Network Elements

In this chapter we will explore the architectural aspects of the network elements that are part of this network. The architecture of such a network is shown in Figure 7.1. The network consists of *optical line terminals* (OLTs), *optical add/drop multiplexers* (OADMs), and *optical crossconnects* (OXCs) interconnected via fiber links. Not shown in the figure are optical line amplifiers, which are deployed along the fiber link at periodic locations to amplify the light signal. In addition, the OLTs, OADMs, and OXCs may themselves incorporate optical amplifiers to make up for losses. As of this writing, OLTs are widely deployed, and OADMs are deployed to a lesser extent. OXCs are just beginning to be deployed.

The architecture supports a variety of topologies, including ring and mesh topologies. OLTs multiplex multiple wavelengths into a single fiber and also demultiplex a composite WDM signal into individual wavelengths. OLTs are used at either end of a point-to-point link. OADMs are used at locations where some fraction of the wavelengths need to be terminated locally and others need to be routed to other

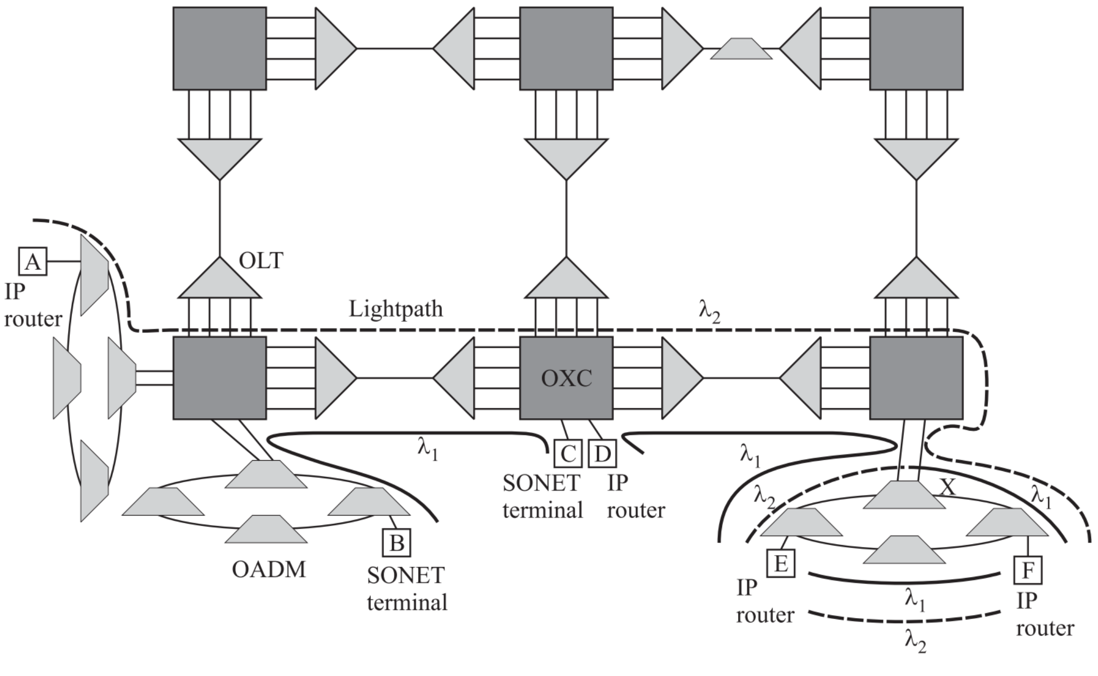

Figure 7.1 A wavelength-routing mesh network showing optical line terminals (OLTs), optical add/drop multiplexers (OADMs), and optical crossconnects (OXCs). The network provides lightpaths to its users, such as SONET boxes and IP routers. A lightpath is carried on a wavelength between its source and destination but may get converted from one wavelength to another along the way.

destinations. They are typically deployed in linear or ring topologies. OXCs perform a similar function but on a much larger scale in terms of number of ports and wavelengths involved, and are deployed in mesh topologies or in order to interconnect multiple rings. We will study these network elements in detail later in this chapter. The users (or clients) of this network are connected to the OLTs, OADMs, or OXCs. The network supports a variety of client types, such as IP routers, Ethernet switches, and SONET terminals and ADMs.

We next describe several noteworthy features of this architecture:

Wavelength reuse. Observe from Figure 7.1 that multiple lightpaths in the network can use the same wavelength, as long as they do not overlap on any link. This spatial reuse capability allows the network to support a large number of lightpaths using a limited number of wavelengths.

Wavelength conversion. Lightpaths may undergo *wavelength conversion* along their route. Figure 7.1 shows one such lightpath that uses wavelength λ_2 on link EX, gets converted to λ_1 at node X, and uses that wavelength on link XF. Wavelength conversion can improve the utilization of wavelengths inside the network. We will study this aspect in Section 7.4.1 and in Chapter 10. Wavelength conversion is also needed at the boundaries of the network to adapt signals from outside the network into a suitable wavelength for use inside the network.

Transparency. Transparency refers to the fact that the lightpaths can carry data at a variety of bit rates, protocols, and so forth and can, in effect, be made protocol insensitive. This enables the optical layer to support a variety of higher layers *concurrently*. For example, Figure 7.1 shows lightpaths between pairs of SONET terminals, as well as between pairs of IP routers. These lightpaths could carry data at different bit rates and protocols.

Circuit switching. The lightpaths provided by the optical layer can be set up and taken down upon demand. These are analogous to setting up and taking down circuits in circuit-switched networks, except that the rate at which the setup and take-down actions occur is likely to be much slower than, say, the rate for telephone networks with voice circuits. In fact, today these lightpaths, once set up, remain in the network for months to years. With the advent of new services and capabilities offered by today's network equipment, we are likely to see a situation where this process is more dynamic, both in terms of arrivals of lightpath requests and durations of lightpaths.

Note that packet switching is *not* provided within the optical layer. The technology for optical packet switching is still fairly immature; see Chapter 12 for details. It is left to the higher layer, for example, IP or Ethernet, to perform any packet-switching functions needed.

Survivability. The network can be configured such that, in the event of failures, lightpaths can be rerouted over alternative paths automatically. This provides a high degree of resilience in the network. We will study this aspect further in Chapter 9.

Lightpath topology. The *lightpath topology* is the graph consisting of the network nodes, with an edge between two nodes if there is a lightpath between them. The lightpath topology thus refers to the topology seen by the higher layers using the

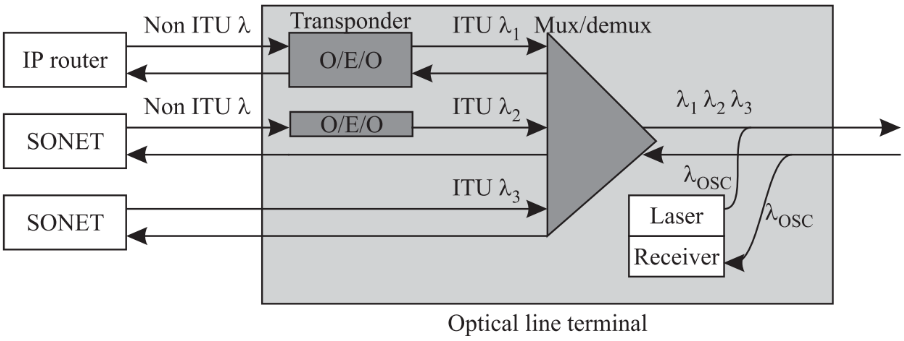

Figure 7.2 Block diagram of an optical line terminal. The OLT has wavelength multiplexers and demultiplexers and adaptation devices called transponders. The transponders convert the incoming signal from the client to a signal suitable for transmission over the WDM link and an incoming signal from the WDM link to a suitable signal toward the client. Transponders are not needed if the client equipment can directly send and receive signals compatible with the WDM link. The OLT also terminates a separate optical supervisory channel (OSC) used on the fiber link.

6.1 Optical Line Terminals

OLTs are relatively simple network elements from an architectural perspective. They are used at either end of a point-to-point link to multiplex and demultiplex wavelengths. Figure 7.2 shows the three functional elements inside an OLT: *transponders, wavelength multiplexers,* and optionally, *optical amplifiers* (not shown in the figure). A transponder adapts the signal coming in from a client of the optical network into a signal suitable for use inside the optical network. Similarly, in the reverse direction, it adapts the signal from the optical network into a signal suitable for the client. The interface between the client and the transponder may vary depending on the client, bit rate, and distance and/or loss between the client and the transponder. The most common interface is the SONET/SDH short-reach (SR) interface described in Section 6.1.5. There are also cheaper very-short-reach (VSR) interfaces at bit rates of 10 Gb/s and higher.

The adaptation includes several functions, which we will explore in detail in Section 8.6.3. The signal may need to be converted into a wavelength that is suited for use inside the optical network. The wavelengths generated by the transponder typically conform to standards set by the International Telecommunications Union (ITU) in the 1.55 μm wavelength window, as indicated in the figure, while the incoming signal may be a 1.3 μm signal. The transponder may add additional overhead for purposes of network management. It may also add forward error correction (FEC), particularly for signals at 10 Gb/s and higher rates. The transponder typically also monitors the bit error rate of the signal at the ingress and egress points in the network. For these reasons, the adaptation is typically done through an optical-to-electrical-to-optical (O/E/O) conversion. Down the road, we may see some of the all-optical wavelength-converting technologies of Section 3.8 being used in transponders; these are still in research laboratories.

The signal coming out of a transponder is multiplexed with other signals at different wavelengths using a wavelength multiplexer onto a fiber. Any of the multiplexing technologies described in Chapter 3, such as arrayed waveguide gratings, dielectric thin-film filters, or fiber Bragg gratings, can be used for this purpose. In addition, an optical amplifier may be used to boost the signal power if needed. In the other direction, the WDM signal is amplified again, if needed, before it is sent through a demultiplexer that extracts the individual wavelengths. These wavelengths are again terminated in a transponder (if present) or directly in the client equipment.

Finally, the OLT also terminates an *optical supervisory channel* (OSC). The OSC is carried on a separate wavelength, different from the wavelengths carrying the actual traffic. It is used to monitor the performance of amplifiers along the link as well as for a variety of other management functions that we will study in Chapter 8.

6.2 Optical Line Amplifiers

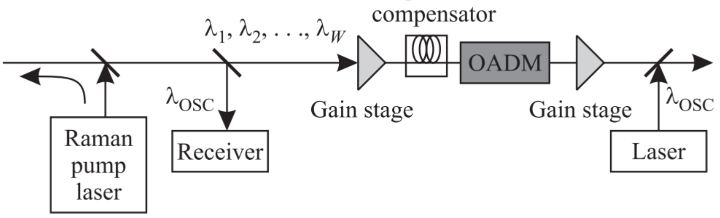

Figure 7.3 Block diagram of a typical optical line amplifier. Only one direction is shown. The amplifier uses multiple erbium gain stages and optionally includes dispersion compensators and OADMs between the gain stages. A Raman pump may be used to provide additional Raman gain over the fiber span. The OSC is filtered at the input and terminated, and added back at the output.

Optical line amplifiers are deployed in the middle of the optical fiber link at periodic intervals, typically 80–120 km. Figure 7.3 shows a block diagram of a fairly standard optical line amplifier. The basic element is an erbium-doped fiber gain block, which we studied in Chapter 3. Typical amplifiers use two or more gain blocks in cascade, with so-called midstage access. This feature allows some lossy elements to be placed between the two amplifier stages without significantly impacting the overall noise figure of the amplifier (see Problem 4.5 in Chapter 4). These elements include dispersion compensators to compensate for the chromatic dispersion accumulated along the link, and also the OADMs, which we will discuss next. The amplifiers also include automatic gain control (see Chapter 5) and built-in performance monitoring of the signal, a topic we will discuss in Chapter 8.

There are also Raman amplifiers, where a high-power pump laser is used at each amplifier site to pump the fiber in the direction opposite to the signal. The optical supervisory channel is filtered at the input and terminated, and added back at the output. In a system using C- and L-bands, the bands are separated at the input to the amplifier and separate EDFAs are used for each band.

6.3 Optical Add/Drop Multiplexers

Optical add/drop multiplexers (OADMs) provide a cost-effective means for handling passthrough traffic in both metro and long-haul networks. OADMs may be used at amplifier sites in long-haul networks but can also be used as stand-alone network elements, particularly in metro networks. To understand the benefits of OADMs, consider a network between three nodes, say, A, B, and C, shown in Figure 7.4, with IP routers located at nodes A, B, and C. This network supports traffic between A and B, B and C, and A and C. Based on the network topology, traffic between A and C passes through node B. For simplicity, we will assume full-duplex links and full-duplex connections. This is the case for most networks today. Thus the network in Figure 7.4 actually consists of a pair of fibers carrying traffic in opposite directions.

Suppose the traffic requirement is as follows: one wavelength between A and B, one wavelength between B and C, and three wavelengths between A and C. Now suppose we deploy point-to-point WDM systems to support this traffic demand. The resulting solution is shown in Figure 7.4(a). Two point-to-point systems are deployed, one between A and B and the other between B and C. As we saw earlier in Section 7.1, each point-to-point system uses an OLT at each end of the link. The OLT includes multiplexers, demultiplexers, and transponders. These transponders constitute a significant portion of the system cost.

Consider what is needed at node B. Node B has two OLTs. Each OLT terminates four wavelengths and therefore requires four transponders. However, only one out of those four wavelengths is destined for node B. The remaining transponders are used to support the passthrough traffic between A and C. These transponders are hooked back to back to provide this function. Therefore, six out of the eight transponders at node B are used to handle passthrough traffic—a very expensive proposition.

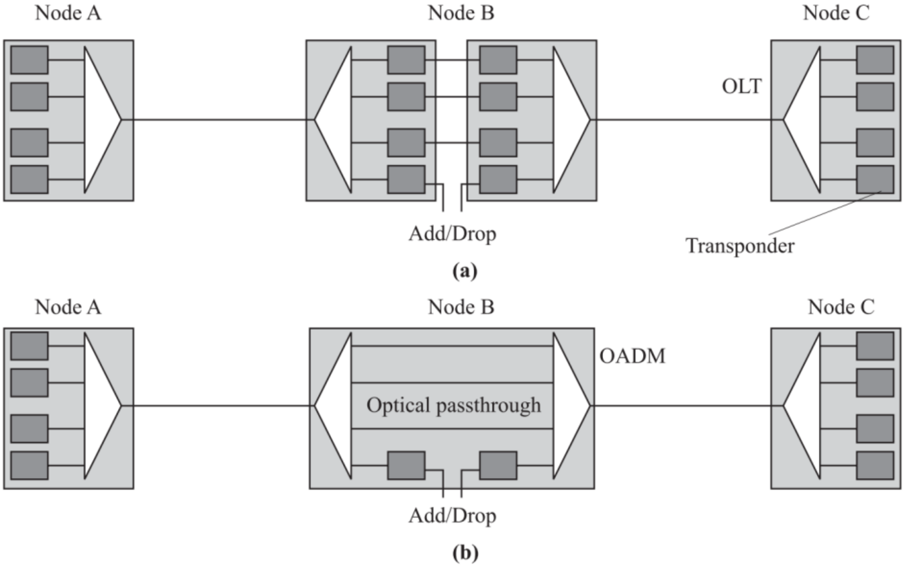

Figure 7.4 A three-node linear network example to illustrate the role of optical add/drop multiplexers. Three wavelengths are needed between nodes A and C, and one wavelength each between nodes A and B and between nodes B and C. (a) A solution using point-to-point WDM systems. (b) A solution using an optical add/drop multiplexer at node B.

Consider what is needed at node B. Node B has two OLTs. Each OLT terminates four wavelengths and therefore requires four transponders. However, only one out of those four wavelengths is destined for node B. The remaining transponders are used to support the passthrough traffic between A and C. These transponders are hooked back to back to provide this function. Therefore, six out of the eight transponders at node B are used to handle passthrough traffic—a very expensive proposition.

Consider the OADM solution shown in Figure 7.4(b). Instead of deploying point-to-point WDM systems, we now deploy a wavelength-routing network. The network uses an OLT at nodes A and C and an OADM at node B. The OADM drops one of the four wavelengths, which is then terminated in transponders. The remaining three wavelengths are passed through in the optical domain using relatively simple filtering techniques, without being terminated in transponders. The net effect is that only two transponders are needed at node B, instead of the eight transponders required for the solution shown in Figure 7.4(a). This represents a significant cost reduction. We will explore this subject of cost savings in detail in Section 10.1.

In typical carrier networks, the fraction of traffic that is to be passed through a node without requiring termination can be quite large at many of the network nodes. Thus OADMs perform a crucial function of passing through this traffic in a cost-effective manner.

6.4 Optical Crossconnects

OADMs are useful network elements to handle simple network topologies, such as the linear topology shown in Figure 7.4 or ring topologies, and a relatively modest number of wavelengths. An additional network element is required to handle more complex mesh topologies and large numbers of wavelengths, particularly at hub locations handling a large amount of traffic. This element is the optical crossconnect (OXC). We will see that though the term *optical* is used, an OXC could internally use either a pure optical or an electrical switch fabric. An OXC is also the key network element enabling reconfigurable optical networks, where lightpaths can be set up and taken down as needed, without having to be statically provisioned.

Consider a large carrier central office hub. This might be an office in a large city for local service providers or a large node in a long-haul service provider's network. Such an office might terminate several fiber links, each carrying a large number of wavelengths. A number of these wavelengths might not need to be terminated in that location but rather passed through to another node. The OXC shown in Figure 7.10 performs this function. OXCs work alongside SONET/SDH network elements as well as IP routers, and WDM terminals and add/drop multiplexers as shown in Figure 7.10. Typically, some OXC ports are connected to WDM equipment and other OXC ports to terminating devices such as SONET/SDH ADMs, IP routers, or ATM switches. Thus, the OXC provides cost-effective passthrough for express traffic not terminating at the hub as well as collects traffic from attached equipment into the network. Some people think of an OXC as a crossconnect switch together with the surrounding OLTs. However, our definition of OXC does not include the surrounding OLTs because carriers view crossconnects and OLTs as separate products and often buy OXCs and OLTs from different vendors.

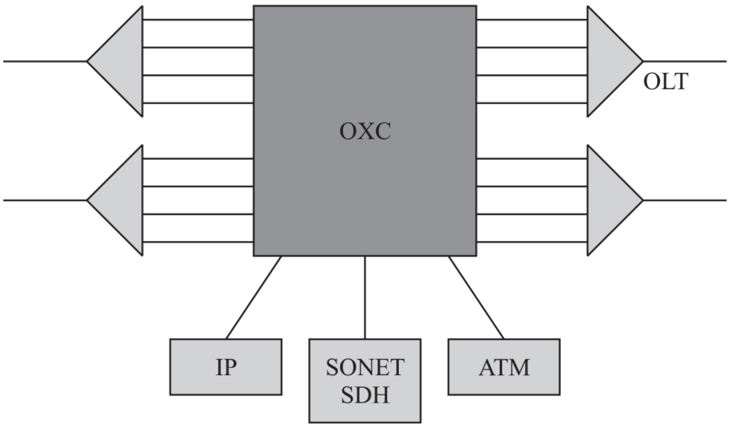

Figure 7.10 Using an OXC in the network. The OXC sits between the client equipment of the optical layer and the optical layer OLTs.

An OXC provides several key functions in a large network:

Service provisioning. An OXC can be used to provision lightpaths in a large network in an automated manner, without having to resort to performing manual patch panel connections. This capability becomes important when we deal with large numbers of wavelengths in a node or with a large number of nodes in the network. It also becomes important when the lightpaths in the network need to be reconfigured to respond to traffic changes. The manual operation of sending a person to each office to implement a patch panel connection is expensive and error prone. Remotely configurable OXCs take care of this function.

Protection. Protecting lightpaths against fiber cuts and equipment failures in the network is emerging as one of the most important functions expected from a crossconnect. The crossconnect is an intelligent network element that can detect failures in the network and rapidly reroute lightpaths around the failure. Crossconnects enable true mesh networks to be deployed. These networks can provide particularly efficient use of network bandwidth, compared to the SONET/SDH rings we discussed in Chapter 6. We discuss this topic in detail in Chapter 9.

Bit rate transparency. The ability to switch signals with arbitrary bit rates and frame formats is a desirable attribute of OXCs.

Performance monitoring, test access, and fault localization. OXCs provide visibility to the performance parameters of a signal at intermediate nodes. They usually allow test equipment to be hooked up to a dedicated test port where the signals passing through the OXC can be monitored in a nonintrusive manner. Nonintrusive test access requires *bridging* of the input signal. In bridging, the input signal is split into two parts. One part is sent to the core, and the other part is made available at the test access port. OXCs also provide loopback capabilities. This allows a lightpath to be looped back at intermediate nodes for diagnostic purposes.

Wavelength conversion. In addition to switching a signal from one port to another port, OXCs may also incorporate wavelength conversion capabilities.

Multiplexing and grooming. OXCs typically handle input and output signals at optical line rates. However, they can incorporate multiplexing and grooming capabilities to switch traffic internally at much finer granularities, such as STS-1 (51 Mb/s). Note that this time division multiplexing has to be done in the electrical domain and is really SONET/SDH multiplexing, but incorporated into the OXC, rather than in a separate SONET/SDH box.

An OXC can be functionally divided into a switch core and a port complex. The switch core houses the switch that performs the actual crossconnect function. The port complex houses port cards that are used as interfaces to communicate with other equipment. The port interfaces may or may not include optical-to-electrical (O/E) or optical-to-electrical-to-optical (O/E/O) converters.