

# Chapter 2

## Optical Network Elements

### 2.1 Introduction

The dramatic shift in the architecture of optical networks that began in the 2000 time frame is chiefly due to the development of advanced optical network elements. These elements are based on the premise that the majority of the traffic that enters a node is being routed *through* the node en route to its final destination as opposed to being *destined for* the node. This transiting traffic can potentially remain in the optical domain as it traverses the node rather than be electronically processed. By deploying technology that enables this so-called *optical bypass*, a significant reduction in the amount of required nodal electronic equipment can be realized.

After briefly discussing some basic optical components in Sect. 2.2, we review the traditional network architecture where all traffic entering a node is electronically processed. The fundamental optical network element in this architecture is the *optical terminal*, which is covered in Sect. 2.3. Optical-terminal-based networks are examined in Sect. 2.4. The economic and operational challenges of these legacy networks motivated the development of optical-bypass technology, which is discussed in Sect. 2.5. The two major network elements that are capable of optical bypass are the *optical add/drop multiplexer* (OADM) and the *multi-degree OADM* (OADM-MD); they are described in Sect. 2.6 and Sect. 2.7, respectively. These two elements are more generically referred to in the industry as *reconfigurable OADMs*, or *ROADMs*; in large part, that terminology is adopted here.

There are three principal ROADM design architectures: *broadcast-and-select*, *route-and-select*, and *wavelength-selective*, all of which are covered in Sect. 2.8. The chief attributes that affect the flexibility, cost, and efficiency of ROADMs are covered in Sect. 2.9, including the *colorless*, *directionless*, *contentionless*, and *gridless* properties. A variety of designs are presented to illustrate several possible ROADM operational alternatives.

ROADMs are one type of optical switch. A more complete taxonomy of optical switches is covered in Sect. 2.10. Hierarchical, or multigranular, optical switches, which may be desirable for scalability purposes, are presented in Sect. 2.11.

In a backbone network, bypass-capable network elements must be complemented by extended *optical reach*, which is the distance an optical signal can travel

before it degrades to a level that necessitates it be “cleaned up,” or regenerated. The interplay of optical reach and optical-bypass-enabled elements is presented in Sect. 2.12.

Integration of elements or components within a node is a more recent development, motivated by the desire to eliminate individual components, reduce cost, and improve reliability. There is a range of possible integration levels as illustrated by the discussions of Sect. 2.13 (integrated transceivers), Sect. 2.14 (integrated packet-optical platforms), and Sect. 2.15 (photonic integrated circuits, PICs).

Throughout this chapter, it is implicitly assumed that there is one fiber pair per link; e.g., a degree-two node has two incoming and two outgoing fibers. Due to the large capacity of current transmission systems, single-fiber-pair deployments are common. However, the last section of the chapter addresses multi-fiber-pair scenarios. (The related topic of fiber capacity is covered in Chap. 9.)

Throughout this chapter, the focus is on the functionality of the network elements, as opposed to the underlying technology.

## 2.2 Basic Optical Components

Some of the optical components that come into play throughout this chapter are discussed here. (Several of these components are illustrated in the various optical-terminal architectures shown in Fig. 2.3.) One very simple component is the *wavelength-independent optical splitter*, which is typically referred to as a *passive splitter*. A splitter has one input port and  $N$  output ports, where the input optical signal is sent to all of the output ports. Note that if the input is a wavelength-division multiplexing (WDM) signal, then each output signal is also WDM. In many splitter implementations, the input power level is split equally across the  $N$  output ports, such that each port receives  $1/N$  of the original signal power level. This corresponds to a nominal input-to-output optical loss of  $10 \cdot \log_{10} N$ , in units of decibels (dB). Roughly speaking, for every doubling of  $N$ , the optical loss increases by another 3 dB. It is also possible to design optical splitters where the power is split nonuniformly across the output ports so that some ports suffer lower loss than others.

The inverse device is called a *passive optical coupler* or *combiner*. This has  $N$  input ports and one output port, such that all of the inputs are combined into a single output signal. The input signals are usually at different optical frequencies to avoid interference when they are combined. The nominal input-to-output loss of the coupler is the same as that of the splitter.

Another important component is the  $1 \times N$  *demultiplexer*, which has one input port and  $N$  output ports. In the most common implementation, a WDM signal on the input line is demultiplexed into its constituent wavelengths, with a separate wavelength sent to each output port. The inverse device is an  $N \times 1$  *multiplexer*, with  $N$  input ports and one output port, where the wavelengths on the input ports are combined to form a WDM signal.

Demultiplexers and multiplexers may be built, for example, using *arrayed waveguide grating* (AWG) technology [Okam98, RaSS09, DoOk06]; such a device is

simply referred to as an “AWG,” or a *wavelength grating router* (WGR). For large  $N$ , the loss through an AWG is on the order of 4–6 dB. AWGs are generally  $M \times N$  devices, where individual wavelengths on the  $M$  input ports can be directed only to one specific output port. In the typical AWG  $1 \times N$  demultiplexer implementation, the number of output ports and the number of wavelengths in the input WDM signal are the same, such that exactly one wavelength is sent to each output port. Similarly, with an AWG  $N \times 1$  multiplexer, where the number of input ports typically equals the number of wavelengths, each input port is capable of directing only one particular wavelength to the output port.


Throughout this chapter, various types of switches are mentioned; it is advantageous to introduce some switch terminology here. There is a broad class of switches known as *optical switches*. Contrary to what the name implies, these switches do not necessarily perform the switching function in the optical domain. Rather, the term “optical switch” is used to indicate a switch where the ports operate on the granularity of a wavelength or a group of wavelengths, as opposed to on finer granularity substrate signals.

*Wavelength-selective* is a term used to classify devices that are capable of treating each wavelength differently. For example, a  $1 \times N$  *wavelength-selective switch* (WSS) can direct any wavelength on the one input port to any of the  $N$  output ports [MMMT03; Maro05; StWa10], thereby serving as a demultiplexer. An  $N \times 1$  WSS performs a multiplexing function. More generally, an  $M \times N$  WSS can direct any wavelength from any of the  $M$  input ports to any of the  $N$  output ports [FoRN12]. Note that a WSS is capable of directing multiple wavelengths to an output port. However, it is typically *not* possible to multicast a given wavelength from one input port to multiple output ports, nor is it typically possible for multiple input ports to direct the same wavelength to one output port (although in principle, a WSS could support both of these functions, depending on the technology). WSSs play a prominent role in many of the architectures discussed in this chapter.

Micro-electro-mechanical-system (MEMS) technology [WuSF06] is often used to fabricate switches with an optical switch fabric. (The switch fabric is the “guts” of the switch, where the interconnection between the input and output ports is established.) This technology essentially uses tiny movable mirrors to direct light from input ports to output ports. Note that an individual MEMS switching element is not wavelength selective; it simply switches whatever light is on the input port without picking out a particular wavelength. However, when combined with multiplexers and demultiplexers that couple the individual wavelengths of a WDM signal to the ports of the MEMS switch, the combination is wavelength-selective, capable of directing any input wavelength to any output port.

## 2.3 Optical Terminal

In traditional optical network architectures, optical terminals are deployed at the endpoints of each fiber link. Figure 2.1 illustrates a single optical terminal equipped with several WDM transponders. An optical terminal is typically depicted in fig-



**Fig. 2.1** A representation of an optical terminal equipped with wavelength-division multiplexing (*WDM*) transponders


ures as a trapezoid to capture its multiplexing/demultiplexing functionality. In most architectures, there are individual wavelengths on the client side of the terminal and a WDM signal on the network side. Unfortunately, a trapezoid is often used to specifically represent a  $1 \times N$  AWG. While an optical terminal can be built using AWG technology, there are other options as well, some of which are discussed in Sect. 2.3.1. *Throughout this book, the trapezoid is used to represent a general optical terminal, or any device performing a multiplexing/demultiplexing function, not necessarily one based on a specific technology.*

In Fig. 2.1, two Internet Protocol (IP) routers are shown on the client side of the optical terminal. Tracing the flow from left to right in the figure, both IP routers transmit a 1,310-nm signal that is received by a WDM transponder. The transponder converts the signal to a WDM-compatible optical frequency, typically in the 1,500-nm range of the spectrum. The optical terminal multiplexes the signals from all of the transponders onto a single network fiber. In general, the transponders plugged into an optical terminal generate different optical frequencies; otherwise, the signals would interfere with each other after being multiplexed together by the terminal. Note that the 1,310-nm signal is sometimes referred to as *gray optics*, to emphasize that the client signals are nominally at the same frequency, in contrast to the different frequencies (or *colored optics*) comprising the WDM signal.

In the reverse direction, a WDM signal is carried by the network fiber into the optical terminal, where it is demultiplexed into its constituent frequencies. Each transponder receives a signal on a particular optical frequency and converts it to a 1,310-nm client-compatible signal.

Recall from Sect. 1.6 that each fiber line shown in Fig. 2.1 actually represents two fibers, corresponding to the two directions of traffic. Also, recall from Fig. 1.7b that the WDM transponder encompasses both a client-side receiver/network-side transmitter in one direction and a network-side receiver/client-side transmitter in the other direction. Similarly, the optical terminal is composed of both a multiplexer and a demultiplexer. Note that it is possible for the network-side signal transmitted by a transponder to be at a different optical frequency than the network-side signal received by the transponder; however, in most scenarios these frequencies are the same.

**Fig. 2.2** An example of an optical-terminal shelf, with the transponder slots fully populated

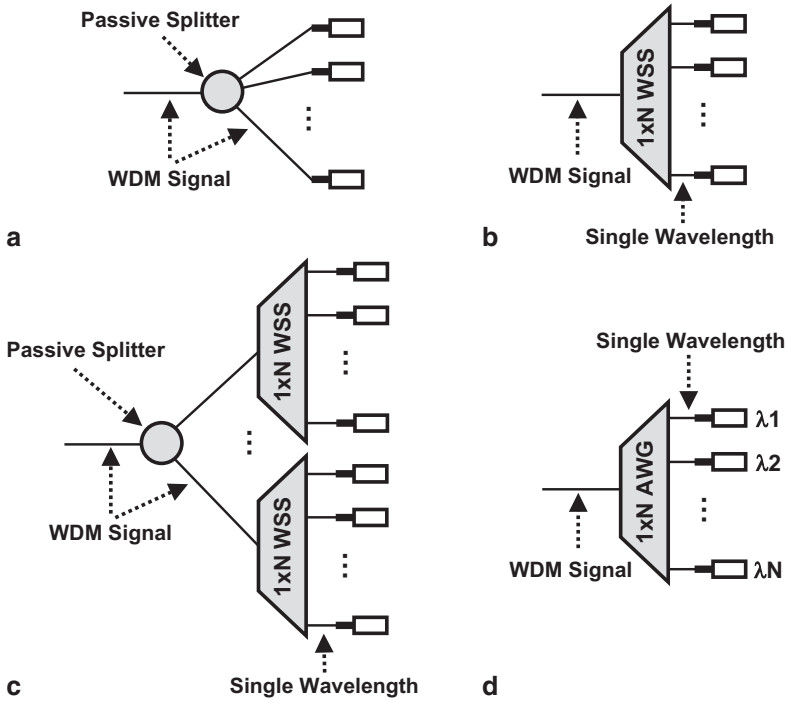


### 2.3.1 Colorless Optical Terminal (Slot Flexibility)

An optical terminal is deployed with equipment shelves in which the transponders are inserted. A prototypical optical-terminal shelf is shown in Fig. 2.2. As depicted, the shelf is fully populated with WDM transponders (i.e., the vertically oriented circuit boards inserted in the slots at the center of the shelf). One figure of merit of an optical terminal is the density of the transponders on a shelf, where higher density is preferred. For example, if a shelf holds up to sixteen 10-Gb/s transponders, the density is 160 Gb/s per shelf. The density is typically determined by properties such as the physical size or power requirements of the transponders (there are industry-wide accepted maxima for the heat dissipation in a fully populated shelf [Telc05a]).

Another desirable feature of an optical terminal is a “pay-as-you-grow” architecture. A fully populated optical terminal may require multiple equipment racks, with multiple shelves per rack, to accommodate all of the transponders. However, ideally the optical terminal can be deployed initially with just a single shelf, and then grow in size as more transponders need to be installed at the site.

The flexibility of the individual slots in the transponder shelves is another important attribute. In the most flexible optical-terminal architecture, any slot can accommodate a transponder of any frequency. Such an architecture is referred to as *colorless*. Clearly, the colorless property simplifies network operations, as a technician can plug a transponder into any available slot. This architecture also maximizes the benefits of tunable transponders, as it allows a transponder to tune to a different frequency without needing to be manually moved to a different slot. Additionally, a colorless optical terminal is typically pay-as-you-grow; i.e., the number of slots deployed needs to be only as large as the number of transponders at the node (subject to the shelf granularity).



**Fig. 2.3** Four optical-terminal architectures, the first three of which are colorless. Only the receive sides of the architectures are shown. **a** The passive splitter architecture has high loss and the transponders must be capable of selecting a particular frequency from the wavelength-division multiplexing (*WDM*) signal. **b** This wavelength-selective switch (*WSS*) architecture limits the number of transponders that can be accommodated to  $N$ . **c** A *WSS* tree architecture increases the number of supported transponders, but at an increased loss. **d** The architecture based on the arrayed waveguide grating (*AWG*) is not colorless; a transponder of a given frequency must be inserted in one particular slot

Four optical-terminal architectures are shown in Fig. 2.3, the first three of which are colorless, while the fourth is not. Only the receive sides of the architectures are shown; the transmit sides are similar.

Figure 2.3a depicts a colorless optical terminal based on a passive splitter in the receive direction, and a passive coupler in the transmit direction (not shown). The received *WDM* signal is passively split, rather than demultiplexed, and directed to each of the transponders. The transponder receiver (on the network side) is equipped with an optical filter (or other frequency-selective technology) to select the desired optical frequency from the *WDM* signal. For maximum transponder flexibility, this filter should be tunable. In the reverse direction, the signals from the transponders are passively coupled together into a *WDM* signal; again, for maximum flexibility, the transponders should be equipped with tunable lasers. Because passive splitters and couplers can result in significant optical loss if the number of supported transponders is large, this architecture often requires optical amplifiers to boost the signal level. Additionally, if the outputs of a large number of transmitters are directly

combined in the passive coupler (i.e., without any filtering to clean up the signals), there may be issues resulting from adding all of the spontaneous emission and other noises of the various lasers; the problem is exacerbated with tunable transmitters.

Another colorless optical-terminal architecture, shown in Fig. 2.3b, is based on WSSs. In the receive direction, a  $1 \times N$  WSS demultiplexes the signal; it is capable of directing any wavelength from the input WDM signal to any of the  $N$  transponders. A second  $N \times 1$  WSS (not shown) is used in the reverse direction to multiplex the signals from the transponders. In this architecture, the transponder receiver does not need to have an optical filter because the wavelength selection is carried out in the WSS (i.e., the transponder is capable of receiving whatever optical frequency is directed to its slot). One drawback of the WSS approach is the relatively high cost compared to the other architectures, although the cost difference is shrinking as WSS technology matures. Another drawback is the limited size of the WSS, which limits the number of transponders that can be supported by the terminal. Commercially available WSSs in the 2015 time frame have a maximum size on the order of  $1 \times 20$ , although they continue to increase in size. If more than  $N$  transponders need to be installed in the optical terminal, then a “tree” composed of a passive splitter/coupler and multiple WSSs can be deployed, as shown for the receive direction in Fig. 2.3c ([WFJA10]; also see Exercise 2.6).

In contrast to these colorless optical-terminal architectures, there are also *fixed* optical terminals where each slot can accommodate a transponder of only one particular frequency. This type of optical terminal is often implemented using AWG technology, as shown in Fig. 2.3d. A  $1 \times N$  AWG demultiplexes the WDM signal in the receive direction; a second  $N \times 1$  AWG (not shown) multiplexes the signals from the transponders in the transmit direction. The transponder receiver does not need to have an optical filter. Using current commercially available technology, an AWG can accommodate many more transponders than a WSS (i.e., much larger  $N$ ). Though relatively cost effective and of low loss, this fixed architecture can lead to inefficient shelf packing and ultimately higher cost in networks where the choice of optical frequencies is very important; i.e., a new shelf may need to be added to accommodate a desired frequency even though the shelves that are already deployed have available slots. The fixed architecture also negates the automated configurability afforded by tunable transponders.

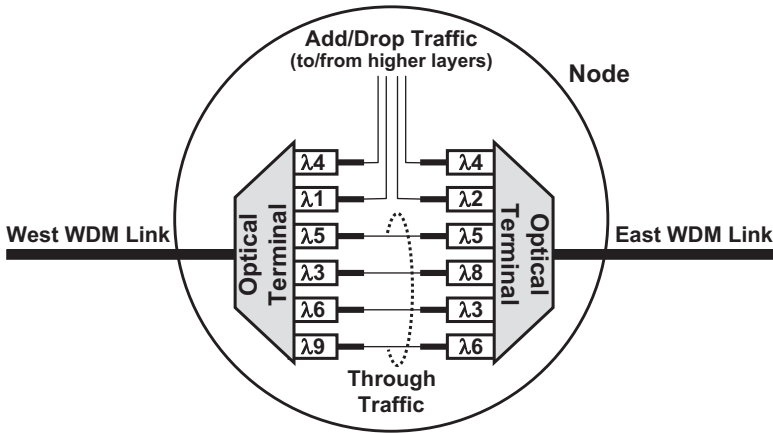
In an intermediary optical-terminal architecture, the WDM spectrum is partitioned into groups, and a particular slot can accommodate transponders only from one group [ChLH06]. This type of terminal can be architected with lower loss and/or cost than a fully colorless design, but has limited configurability.

## 2.4 Optical-Electrical-Optical (O-E-O) Architecture

### 2.4.1 O-E-O Architecture at Nodes of Degree-Two

The traditional, non-configurable, optical-terminal-based architecture for a node of degree-two is shown in Fig. 2.4. There are two network links incident on the node,





**Fig. 2.4** O-E-O architecture at a degree-two node (without automated reconfigurability). Nodal traffic is characterized as either add/drop traffic or through traffic. All traffic entering and exiting the node is processed by a transponder. Note that the through traffic can undergo wavelength conversion, as indicated by interconnected transponders of different wavelengths (e.g., the bottom pair of interconnected transponders converts the signal from  $\lambda_6$  to  $\lambda_9$  in the East-to-West direction)

where it is common to refer to the links as the “East” and “West” links (there is not necessarily a correspondence to the actual geography of the node). As shown in the figure, the node is equipped with two optical terminals arranged in a “back-to-back” configuration. The architecture shown does not support automated reconfigurability. Connectivity is provided via a manual patch panel, i.e., a panel where equipment within an office is connected via fiber cables to one side (typically in the back), and where short patch cables are used on the other side (typically in the front) to manually interconnect the equipment as desired. Providing automated reconfigurability is discussed in the next section in the context of higher-degree nodes.

Tracing the path from right to left, the WDM signal enters the East optical terminal from the East link. This WDM signal is demultiplexed into its constituent wavelengths, each of which is sent to a WDM transponder that converts it to a 1,310 nm optical signal. (Recall that 1,310 nm is the typical wavelength of the client-side optical signal.) At this point, it is important to distinguish two types of traffic with respect to the node. For one type of traffic, the node serves as the exit point from the optical layer. This traffic “drops”<sup>1</sup> from the optical layer and is sent to a higher layer (the higher layers, e.g., IP, are the clients of the optical layer). The other type of traffic is transiting the node en route to its final destination. After this transiting traffic has been converted to a 1,310 nm optical signal by its associated transponder, it is sent to a second transponder located on the West optical terminal. This transponder converts it back into a WDM-compatible signal, which is then multiplexed by the West optical terminal and sent out on the West link. There are also transponders on

<sup>1</sup> While “drop” often has a negative connotation in telecommunication networks (e.g., dropped packets, dropped calls), its usage here simply means a signal is exiting from the optical layer.



the West terminal for traffic that is being “added” to the optical layer, from higher layers, that needs to be routed on the West link.

In the left to right direction of the figure, the operation is similar. Some of the traffic from the West link drops from the optical layer and some is sent out on the East link. Additionally, there are transponders on the East terminal for traffic that is added to the optical layer at this node that needs to be routed on the East link.

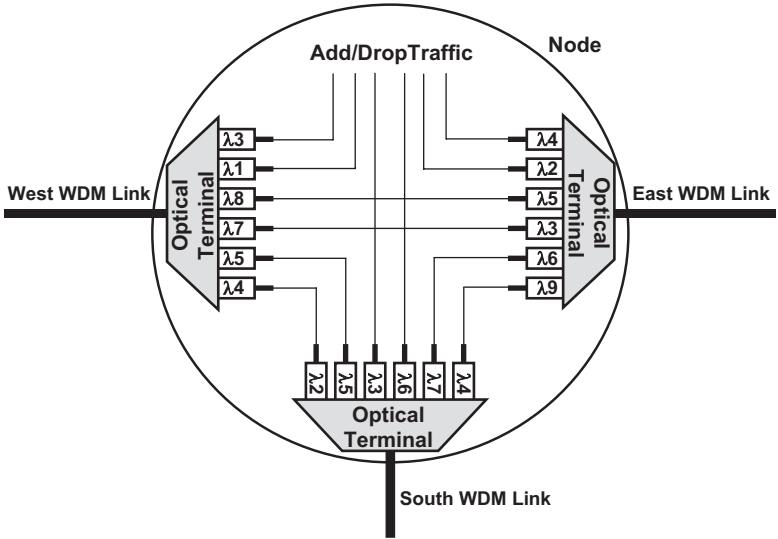
The traffic that is being added to or dropped from the optical layer at this node is termed *add/drop* traffic; the traffic that is transiting the node is called *through* traffic. Regardless of the traffic type, note that all of the traffic entering and exiting the node is processed by a WDM transponder. In the course of converting between a WDM-compatible optical signal and a client optical signal, the transponder processes the signal in the electrical domain. Thus, all traffic enters the node in the optical domain, is converted to the electrical domain, and is returned to the optical domain. This architecture, where all traffic undergoes optical-electrical-optical (O-E-O) conversion, is referred to as the *O-E-O architecture*.

### 2.4.2 *O-E-O Architecture at Nodes of Degree-Three or Higher*

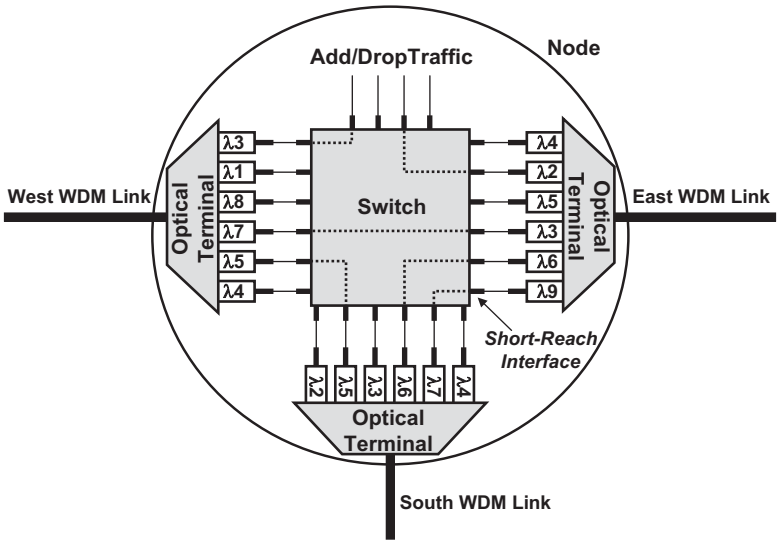
The O-E-O architecture readily extends to a node of degree greater than two. In general, a degree- $N$  node will have  $N$  optical terminals. Figure 2.5 depicts a degree-three node equipped with three optical terminals, with the third link referred to as the “South” link. The particular architecture shown does not support automated reconfigurability.

As with the degree-two node, all of the traffic entering a node, whether add/drop or through traffic, is processed by a transponder. The additional wrinkle with higher-degree nodes is that the through traffic has multiple possible path directions. For example, in the figure, traffic entering from the East could be directed to the West or to the South; the path is set by interconnecting a transponder on the East optical terminal to a transponder on the West or the South optical terminal, respectively. In many real-world implementations, the transponders are interconnected using a manual patch panel. Modifying the through path of a connection requires that a technician manually rearrange the patch panel, a process that is not conducive to rapid reconfiguration and is subject to operator error.

The reconfiguration process can be automated through the addition of an optical switch, as shown in Fig. 2.6. (The traffic patterns shown in Fig. 2.5 and Fig. 2.6 are not the same.) Each transponder at the node feeds into a switch port, and the switch is configured as needed to interconnect two transponders to create a through path. Additionally, the add/drop signals are fed into ports on the switch, so that they can be directed to/from transponders on any of the optical terminals. Furthermore, the switch allows any transponder to be flexibly used for either add/drop or through traffic, depending on how the switch is configured. Note that the degree-two architecture of Fig. 2.4 could benefit from a switch as well with respect to these latter



**Fig. 2.5** O-E-O architecture at a degree-three node (without automated reconfigurability). There are three possible directions through the node. The path of a transiting connection is set by inter-connecting a pair of transponders on the associated optical terminals



**Fig. 2.6** A switch is used to automate node reconfigurability. The particular switch shown has an electronic switch fabric and is equipped with short-reach interfaces on all of its ports

<http://www.springer.com/978-3-319-05226-7>

Optical Network Design and Planning

Simmons, J.M.

2014, XXV, 516 p. 203 illus., 5 illus. in color., Hardcover

ISBN: 978-3-319-05226-7