CHAPTER ONE

Introduction

Despite the fact that optical fiber is the prevalent transmission medium and the emergence of true optical networking has been anticipated since the early 1990s [1], current networks are far from being termed as optical. The history of optical communication has mostly been about transmission and about ways to provide higher bandwidths while simultaneously reducing the cost per bit transmitted [2].

Telecommunication traffic has been growing at a high and steady rate since the early 1980s. Even though the overly optimistic traffic forecasts of the late 1990s never materialized and the associated investments caused a downturn in the industry, one should not lose sight of the fact that the trend toward office automation, remote access, online transactions, and so on has been steady and will continue to grow [3]. These types of services are supported by two communications technologies that are displacing all others: wireless, which can go everywhere, but with limited capacity, and optical fiber, which, although limited to fixed paths, has almost unlimited capacity [4].

Until the late 1990s, networks using optical fiber were viewed merely as transmission pipes that can carry a huge amount of traffic. With advances in optical technologies, that paradigm is shifting towards optical networks that are capable of providing network flexibility, new services, and operational efficiencies [5–7]. This is the notion behind the intelligent optical network. In addition to the increase in data and wireless traffic volumes, new optical services have become possible due to recent advances in optical technology [3, 8–11].

Widespread deployment of affordable broadband services will depend heavily on the availability of improved optical networks, which already provide the physical infrastructure for much of the world's telecommunications and Internet-related services. Optical technology is also essential to the future development of mobile and wireless communications and cable television networks [12]. In the last couple of years optical backbone equipment development has focused on three

Optical Switching, by Georgios I. Papadimitriou, Chrisoula Papazoglou, and Andreas S. Pomportsis
Copyright © 2007 John Wiley & Sons, Inc.

basic lines: enhanced DWDM, long-haul capabilities, and optical switching [13–20].

Another trend appearing on the horizon is closer interaction between the optical layer and the client layers. With IP routers emerging among the prevailing clients of the optical layer, there has been a great deal of interest recently in trying to obtain a closer interaction between the IP layer and the optical layer from a control and management perspective [2]. Internet protocol traffic, which is dominating the service scene in terms of traffic volume, is largely still of a best-effort nature. Thus, quality-of-service support and traffic engineering functions are supported by a separate layer (ATM); furthermore, a SONET/SDH layer is used to guarantee reliable transport over the DWDM layer, which provides abundant capacity.

This means that although IP traffic is indeed carried over optical fibers, several intermediate technologies are required to provide the required functionalities. These layers are largely unaware of each other, which increases overhead and may lead to duplication of their services [21]. There is a growing industry consensus toward simplifying the network by collapsing these multiple service layers into two layers, a "smart" IP layer and a "smart" optical layer. This paradigm is referred to as "IP over WDM," and the resulting network is referred to as the "Optical Internet" [22–26]. The overall goal is to allow carriers to implement a simpler, lower-cost, more responsive network, which is capable of addressing a wider spectrum of service requirements [25]. The challenge is that the functions supported by the eliminated layers must be transferred to the remaining ones. Obviously, some functionality will have to be migrated to the optical layer, which will no longer be static.

This chapter initially presents a brief overview of the evolution of optical networks followed by a description of current networks and their shortcomings. The benefits of optical networking are then discussed, along with the role of optical switching. The advantages and challenges of optical switching are highlighted and compared with other alternatives. Finally, the optical switching paradigms presented in Chapters 3, 4, and 5 are introduced.

1.1 THE EVOLUTION OF OPTICAL NETWORKS

Interest in optical communications began in the mid-1960s, when early experiments showed that information encoded in light signals could be transmitted over a glass-fiber waveguide. However, optical fiber transmission systems really took off with the advent of the low-loss optical fiber. This silica-based optical fiber has three low-loss windows, which enable transmission of light signals over distances of several tens of kilometers before they needed to be regenerated. A regenerator converts the light signal into an electrical signal and retransmits a fresh copy of the data as a new light signal [2].

The first type of fiber used was multimode fiber, in which light propagates in multiple modes, each traveling over a different path and at a different velocity.

At the end of the fiber, the different modes arrive at slightly different times, resulting in a smearing of the pulse. This phenomenon is referred to as intermodal (or modal) dispersion. In order to cope with intermodal dispersion, regenerators had to be placed every few kilometers to recreate the signal. Regenerators were and continue to be expensive devices; it is therefore desirable to maximize the distance between them.

A significant move in this direction was the introduction of single-mode fiber. This type of fiber has a small core diameter, which forces all the energy in a light signal to travel in the form of a single mode. Single-mode fiber therefore eliminated modal dispersion and enabled dramatic increases in bit rates and distances between regenerators [2]. Since then, continuous improvements in fiber characteristics and transceiver technology have further improved the performance of optical transmission systems in terms of these two aspects (bit rate and regenerator spacing).

The initial use for optical fiber communication, and its prevalent use today, was to provide high-bandwidth, point-to-point pipes. At the ends of these pipes, data are converted from the optical to the electrical domain and all the switching, routing, and intelligent control functions are handled by higher-layer equipment [2]. In such networks, optical fiber is simply used as a transmission medium and no other attempt is made to include functionality in the static optical layer.

The most representative examples of networks based on electronics that utilize optical fiber as a transmission medium are SONET/SDH networks. SONET (Synchronous Optical NETwork) and SDH (Synchronous Digital Hierarchy) are standards developed by the American National Standards Institute (ANSI) and the International Telecommunication Union (ITU), respectively. These two analogous standards describe optical network architectures in terms of the required network elements and the functionality that each node must perform. The goal is to ensure a basic level of interoperability between optical equipment acquired from different vendors. SONET and SDH also define hierarchies of digital data rates along with an efficient multiplexing scheme for combining multiple lower-speed signals into higher ones [27].

The basic network elements for SONET (similar equipment exists for SDH), besides properly spaced regenerators, are

- terminal multiplexers or line terminals,
- add/drop multiplexers,
- digital crossconnect systems.

Terminal multiplexers combine several lower-speed streams into a single, high-speed signal before transmission and separate them at the other end. Both functions are performed electronically and the signal undergoes optical-to-electrical (OE) and electrical-to-optical (EO) conversions. SONET add/drop multiplexers are used to add one or more lower-speed streams to a high-speed stream (add function) or select one or more lower-speed streams out of a high-speed stream (drop function); these functions are performed in transit without terminating the entire traffic. Digital crossconnect systems are able to switch large numbers of individual streams and may also perform add/drop functions.

A number of configurations are possible in a SONET/SDH network, such as point-to-point, ring, and linear networks. The ring is the most common for SONET/SDH networks, both in the access and the backbone parts, because of its high restoration capability in the presence of network faults and its simplicity. In a ring topology, add/drop multiplexers are interconnected with two-fiber links; one is the working fiber and the other is the protection fiber, offering a high degree of availability in case some kind of failure occurs in the working fiber.

Such single-wavelength optical networks utilize only a very small fraction of the available capacity of optical fibers. The most decisive step in the effort to exploit a larger percentage of this capacity was the introduction of the wavelength division multiplexing (WDM) technique. According to this technique, a number of independently modulated wavelength channels (i.e., with potentially different data rates and formats) can propagate in parallel using the same fiber. Rather than further increasing the rate at which data are transmitted on a single wavelength, WDM utilizes multiple channels (at possibly lower rates), each carrying separate data. This approach resembles the deployment of additional optical fibers (but is not as complex and time-consuming) and for this reason, the multiplexed channels are often thought of as virtual fibers. Thus, the aggregate capacity of an optical fiber can be extended by increasing the number of wavelengths (or equivalently by decreasing the spacing between channels) or by increasing the channel data rates. Adequate spacing between channels must be maintained in order to avoid interference. The decrease in channel spacing has led to the development of Dense Wavelength Division Multiplexed (DWDM) optical networks, where the number of channels per fiber reaches the hundreds region.

The most important catalyst fueling the deployment of WDM systems was the development of Erbium Doped Fiber Amplifiers (EDFAs), which occurred in the late 1980s and early 1990s [2]. An EDFA essentially consists of a length of optical fiber, typically a few meters to tens of meters, doped with the rare earth element erbium. By using a pump source, erbium atoms are pumped from their ground state to an excited state at a higher energy level. An incoming signal photon triggers these atoms to come down to their ground state. In the process, each atom emits a photon. Thus, incoming signal photons trigger the emission of additional photons, resulting in optical amplification [2]. Such EDFAs have numerous advantages over electronic regenerators. The most significant advantage of EDFAs is that they are capable of amplifying signals at many wavelengths simultaneously and, of course, without resorting to optical-to-electrical-to-optical (OEO) conversions. The application of WDM and EDFAs dramatically brought down the cost of long-haul transmission systems and increased their capacity. Optical amplifiers replaced arrays of expensive regenerators at a fraction of their cost. Nevertheless, regenerators were not eliminated, because EDFAs provide only a subset of their functionality. If 3R regeneration is considered, which combines signal reamplification, retiming, and reshaping (hence the 3R), then it is obvious that EDFAs can only support the first of the three functions. Furthermore, optical amplifiers are by no means ideal devices, and their gain is not flat over the entire spectrum (i.e., they do not amplify all wavelengths equally).

The advent of WDM and optical amplification gave rise to undesirable phenomena that had not been previously observed in fibers, either because single wavelengths were used or because of the frequent signal regenerations. The first effect that was observed was chromatic dispersion. The speed at which optical signals propagate depends to some extent on the wavelength used. As a result, different spectral components travel with different velocities and pulse smearing is observed. Chromatic dispersion can be compensated by using, for example, specially engineered optical fibers. Unfortunately, this results in other fiber impairments, referred to as nonlinear effects, which involve the interaction of wavelength channels. Four-wave mixing (FWM) is an example of such an effect: In FWM, three light signals at different wavelengths interact in the fiber to create a fourth light signal at a wavelength that may overlap with one of the light signals and interfere with the actual data being transmitted on that wavelength. The relation between fiber nonlinearities and chromatic dispersion can be explained as follows: because of chromatic dispersion, signals at different wavelengths spread in time and go out of phase with one another, which means that the interactions between them are reduced. This trade-off between chromatic dispersion and fiber nonlinearities is taken into account in nonzero, dispersion-shifted, single-mode fibers, which can be used to manage the interaction between these two effects. These fibers are tailored to provide less chromatic dispersion than conventional fibers, and at the same time, reduce nonlinearities [2].

Such techniques have enabled commercial systems to achieve distances of several thousand kilometers between regenerators at high rates [2]. Further increases in terms of data rates and regenerator spacing continue to be pursued. Efforts have focused on the development of special types of fibers with desirable characteristics, special transmission formats and pulse shapes (e.g., solitons), and new approaches to optical amplification such as Raman amplification.

Meanwhile, advances in optical components technology facilitate further reductions in optoelectronic equipment. Optical add/drop multiplexers can be used to insert or extract specified wavelengths from a fiber carrying hundreds of wavelengths. This eliminates the need for the hundreds or thousands of optical transponders typically used to terminate the traffic on a fiber at each node where wavelengths need to be added or dropped. Significant financial benefits arise as a result, because optical transponders account for a large fraction of network cost [9].

All these technological advances relating to the WDM technique also led to the development of the first two types of networks that exploited WDM components to transport and route optical signals without converting them to electrical form, namely broadcast-and-select and wavelength-routed networks.

1.1.1 Broadcast-and-Select Networks

Broadcast-and-select networks are typically based on an optical passive coupler, which interconnects network nodes. The most common topology is the star, but

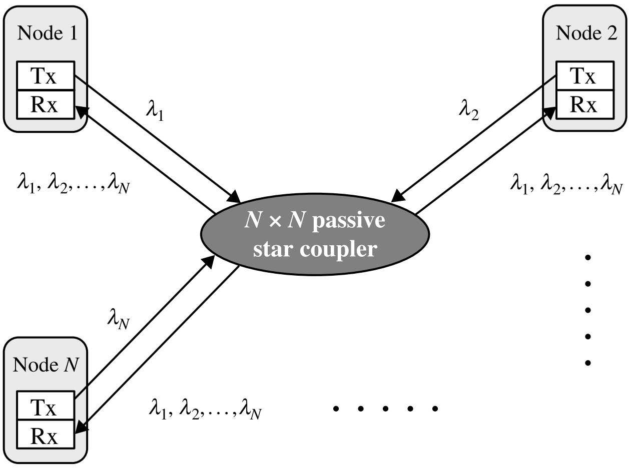

FIGURE 1.1 A broadcast-and-select WDM network with a passive star coupler. Each node is equipped with a fixed transmitter and a tunable receiver.

bus, ring, and tree topologies are also possible. An example of a broadcast-and-select network is illustrated in Figure 1.1, where all node connections are realized via a star coupler and no direct communication between nodes is possible. All nodes are connected with the star coupler via two separate fibers, one for transmission and one for reception. The passive coupler receives all signals transmitted by network nodes and combines them in a multiwavelength signal. The combined signal is subsequently split and broadcast to all N nodes. Each node in a broadcast-and-select network is equipped with a number of transmitters and receivers, which may be fixed or tunable. In Figure 1.1, tunability is assumed to be on the receiving side and nodes transmit on prespecified wavelengths. If one-hop communication is required, the tuning range of the tunable elements (transmitters or receivers) must include all wavelengths used in the system. The choice of a particular implementation is based on technological and financial constraints.

Because all transmissions in a broadcast-and-select network go through the passive coupler, all nodes must transmit using separate wavelengths. Otherwise, a collision will occur, which results in loss of data. This is referred to as the distinct channel assignment constraint and characterizes the broadcast-and-select architecture. A collision may also take place when two or more nodes transmit on distinct wavelengths, but towards the same node equipped with a tunable receiver. Collisions may be avoided by properly designed Medium Access Control protocols.

An attractive feature of broadcast-and-select networks is their inherent multicast capability. Multicasting can be implemented in a very straightforward way by having all nodes in a multicast group tune their receivers on the wavelength used for transmission [51]. On the negative side, the splitting of signals that occurs in the passive coupler limits the scalability of broadcast-and-select networks. Each node receives only a fraction of the signal power (equal to $1/N$), and therefore the number of nodes that can be supported is restricted. For this reason, broadcast-and-select networks are considered practical only for local and metropolitan area networks, where N is relatively small.

1.1.2 Wavelength-Routed Networks

Broadcast-and-select networks are rather static in nature and face scalability issues; therefore their application is limited to the local and metropolitan domains. Wavelength-routed networks are more general and flexible, and have the potential to be deployed in long-haul backbone networks. A fundamental advantage of wavelength-routed networks over broadcast-and-select ones is that they support wavelength reuse in different parts of the network [28]. This cannot be achieved in broadcast-and-select networks because all signals have to traverse the star coupler. This property improves the scalability of wavelength-routed networks, which allows a large number of nodes to be served using relatively few wavelength channels.

Figure 1.2 depicts a wavelength-routed network comprising access nodes and optical switches, which form an all-optical cloud. In such a network data are transported via dedicated lightpaths (i.e., all optical communication paths), which are set up between access nodes. In the simplest case, there is no wavelength conversion capability at intermediate nodes and each lightpath uses a single wavelength on all links. Wavelength reuse is possible even in this simple scenario. In Figure 1.2, the lightpath between nodes 2 and 6 and the one between nodes 3 and 4 use the same wavelength, λ_2. Clearly, the same wavelength can be used in different lightpaths only if they do not share one or more links (i.e., if they do not overlap). For this reason, the lightpath between nodes 2 and 5 must use a different wavelength. Wavelength reuse results in a significant reduction in the number of wavelengths required to support connections between nodes, especially in networks with large node counts.

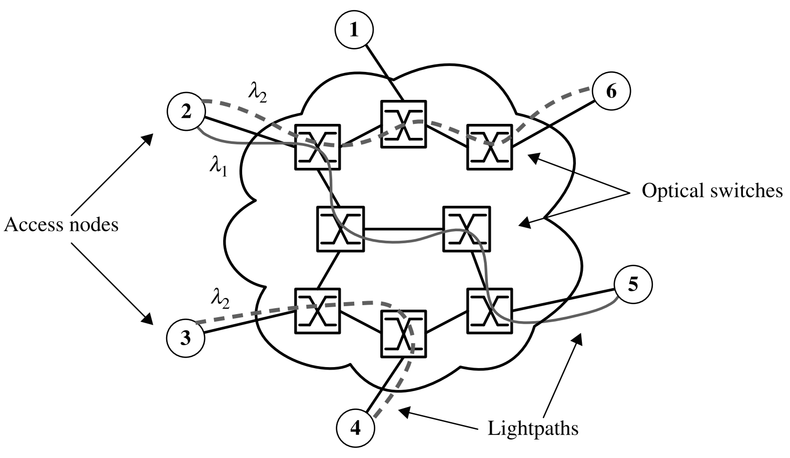

FIGURE 1.2 A wavelength-routed WDM network.

When wavelength conversion is available at intermediate nodes, a lightpath can be set up by assigning a distinct wavelength on each link. This increases the probability that a request for the establishment of a lightpath between two end nodes will be successfully served. The problem of selecting the wavelength(s) and the intermediate nodes (i.e., the path) that will be used by a lightpath is referred to as the routing and wavelength assignment (RWA) problem. The RWA problem can be formally stated as follows: Given a set of lightpaths that need to be established on a particular network topology, determine the routes over which these lightpaths should be established and the wavelengths that should be assigned to them, using the minimum possible number of wavelengths [27–32]. In selecting a wavelength for a lightpath, the two constraints mentioned above must be taken into account. According to the distinct wavelength assignment constraint, all lightpaths sharing a fiber link must be assigned distinct wavelengths to avoid interference. The second constraint only applies when there is no wavelength conversion capability. According to the wavelength continuity constraint, the wavelength assigned to each lightpath must be the same on all links it traverses between the two end nodes.

The establishment of lightpaths may be performed offline or online. In the former (also referred to as static RWA), a set of lightpaths are established in advance and remain fixed for a long period of time. Conversely, when requests for lightpaths are submitted and served online (dynamic RWA [33]), lightpaths are released after a finite amount of time. This type of wavelength routing can be considered the equivalent of circuit switching in optical networks. After the route and the wavelength(s) for a new lightpath have been selected, the necessary resources must be provisioned and nodes must be configured accordingly. This is accomplished via a signaling protocol that informs intermediate nodes about the necessary actions.

Wavelength-routed networks do not face the limitations of broadcast-and-select networks and are a viable alternative for a reconfigurable optical backbone. However, when the volume of traffic that needs to be transported is small compared to a wavelength, this type of routing results in inefficient utilization of resources. Traffic from different ingress nodes or to different egress nodes cannot travel on the same lightpath. Furthermore, wavelength routing is not suitable for transporting bursty traffic streams, that is, traffic with a high intensity and a short duration, because the process of establishing (and tearing down) a lightpath typically takes longer than the transmission of the data to be transported.

1.2 VIEW OF THE CURRENT NETWORK

The current network hierarchy is depicted in Figure 1.3 and includes three levels [34–36]. At the top level, the wide-area backbone networks employ wavelength division multiplexing (WDM) in order to transport vast amounts of data via optical fiber links. The switching of data is performed in the electronic domain following optical-to-electrical conversions. Prior to being retransmitted, data are converted back to the optical form. The second level of the network hierarchy includes metropolitan area (metro) networks that interconnect the backbone networks with the local access networks (third level). Access networks employ

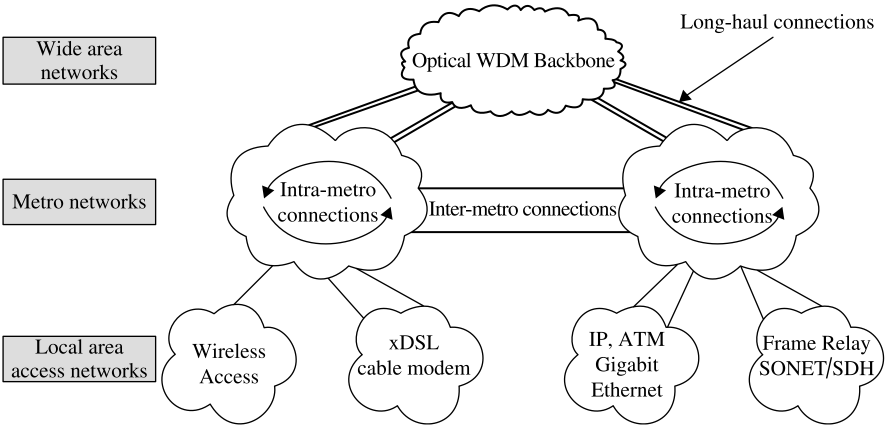

FIGURE 1.3 The current network hierarchy. (© 2004 IEEE.)

advanced LAN technologies (e.g., gigabit Ethernet or broadband access) to provide significant amounts of bandwidth to end users.

In the history of network evolution, three different forces have consistently driven the architecture and evolution of telecommunications networks: traffic growth, development of new services, and advances in technology [3]. These forces are not independent of each other, but each shapes the evolution in a different way. This is depicted in Figure 1.4.

Several broadband applications have recently gained attention. They can generally be classified into two categories [4]:

- medium-size file transfers requiring low latency such as videoconferencing and interactive games; and
- transfer of large files whose latency is not so critical, but for which long transfer times are annoying, such as video on demand or online movies rental.

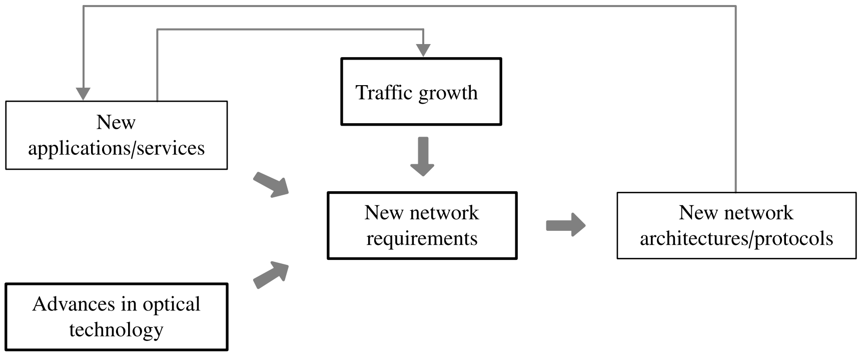

FIGURE 1.4 The drives behind optical network advancement.

Other examples of applications include high-resolution home-video editing, real-time rendering, high-definition interactive television, e-health, and immersive interactive-learning environments. These applications need infrastructures that make a vast amount of storage and computation resources potentially available to a large number of users [37].

Another growing application is remote disk backup at centralized secure servers residing in storage area networks. The potential of this application is constrained today by the excessive time required to communicate large files [4]. Additionally, distributed computing in the form of a computational grid (e.g., for scientific applications) has also gained interest as an application for high-speed, long-haul, optical networks.

Currently, high-speed backbone networks utilize optical fiber almost exclusively as a transmission medium. In this scenario, the optical layer provides circuit-switched, high-bandwidth connections to its users. Whenever necessary, optical signals are converted to electrical and regenerated before further transmission. All routing, forwarding, and switching functions are implemented by electronic routers.

The use of a static optical layer has several disadvantages, all revolving around the way bandwidth provisioning is performed. New connections are established through manual configurations over a timescale of weeks or even months. Provisioning a connection may involve a number of steps. A connection request arrives and the service provider manually looks at his network topology and inventory and decides how to support the connection, including routing it in the network, and determining the appropriate equipment required. Once the connection is turned up, it is then monitored for a while to ensure adequate performance before it is turned over to the user.

This slow provisioning process implies that the establishment of a connection is justified only in cases where its duration is very long. Furthermore, human intervention and manual configuration make the process prone to errors. Fast, automated provisioning within minutes or seconds would open new opportunities for both transport network operators and their clients. Such a shortening of the provisioning time would make more efficient utilization of available resources possible, allowing dynamic allocation of links and nodes, matching rapidly changing needs of client networks [38]. Consider for example two IP routers belonging to separate networks that communicate over an optical network. This situation is depicted in Figure 1.5.

In Figure 1.5, the two routers exchange data over a preconfigured lightpath on wavelength λ_1. If a large volume of bursty traffic has to be transported between them, congestion may be observed. In order to relieve congestion and prevent it from resulting in loss of data, the two routers would have to apply appropriate transport protocol techniques to adapt the flow of traffic to the available bandwidth. A more efficient alternative would be to temporarily allocate additional bandwidth for the communication between the two points of congestion. If dynamic provisioning for short timescales is supported, another lightpath at wavelength λ_2 can be established to aid in the transmission of data. After the traffic decreases, this path can be torn down and assigned to other endpoints.

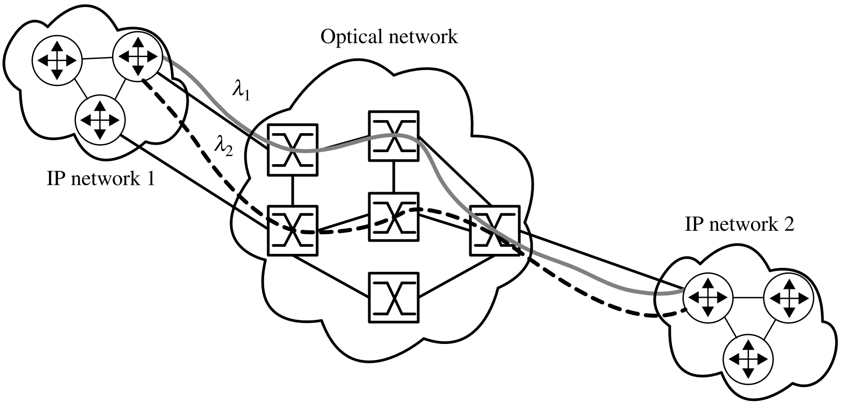

FIGURE 1.5 An example of fast provisioning. (© 2005 IEEE.)

Fast provisioning can be accomplished by migrating all or part of the required functionality to the optical layer, which will no longer be static. Significant benefits may arise from adding intelligence and flexibility to the optical layer. For instance, an agile, smart optical network provides flexible restoration options, which can be utilized by service providers to offer a variety of services tailored to meet the availability requirements of their clients [2]. New protection and restoration schemes for mesh-type networks will improve the reliability performance measures offered to customers. Such measures are especially important if we take into account very-high-bit data rates switched in optical networks [38–41,52].

Meanwhile, significant cost-benefits arise from the removal of unnecessary optoelectronic equipment. These two evolution paths (the addition of intelligence to the optical layer and the removal of unnecessary OEO conversions) lead to the construction of a true optical network.

1.3 OPTICAL NETWORKING

In an all-optical network, the signal stays in the optical domain throughout the source–destination path. This implies that OEO conversions are eliminated. This is a rather strict definition. According to a more realistic definition (more compatible with the current state of technology) an optical network satisfies two conditions [42]:

1. It is composed of long- or ultra-long-reach optical components such that the all-optical reach (i.e., the distance between regeneration sites or other OEO conversions) is larger than a multiple of the typical network internodal distance.
2. Network nodes contain elements that allow optical bypass. Examples of such network elements include optical add–drop multiplexers (OADMs) in degree-two nodes and optical switches in nodes of higher degrees [42]. These elements switch traffic transparently without inspecting it.

The key to constructing an all-optical (or simply optical) network according to the first criterion is the drastic reduction of regeneration sites. Regeneration sites are expensive, require large amounts of space and electrical power, and generate so much heat that active cooling systems may be required [42]. Another consequence of having multiple regeneration sites is slower provisioning times due to the fact that the installation of a circuit requires the installation of regenerators at multiple locations between the two endpoints. Complete elimination of electronic regeneration is currently not possible due to the lack of an optical 3R regeneration technology.

In the absence of all-optical 3R functionality and all-optical wavelength conversion, the all-optical network will continue to remain elusive. Instead, "islands of optical transparency" that are bounded by optoelectronics (OEO) will continue to expand. Long-haul or ultra-long-haul WDM systems, and transparent optical rings, bounded by OEO, will continue to be the norm [18]. These growing islands of transparency, depicted in Figure 1.6, will be composed of WDM links that connect optical crossconnects and optical add/drop multiplexers as they become available. The OEO boundaries between islands will either be defined jurisdictionally or by propagation budget considerations (i.e., due to accumulated attenuation or dispersion) [43].

A significant portion of the network cost lies in the equipment used to convert signals from the electrical to the optical domain. The OEO devices are the main consumers of electrical power in the network and account for a big fraction of the

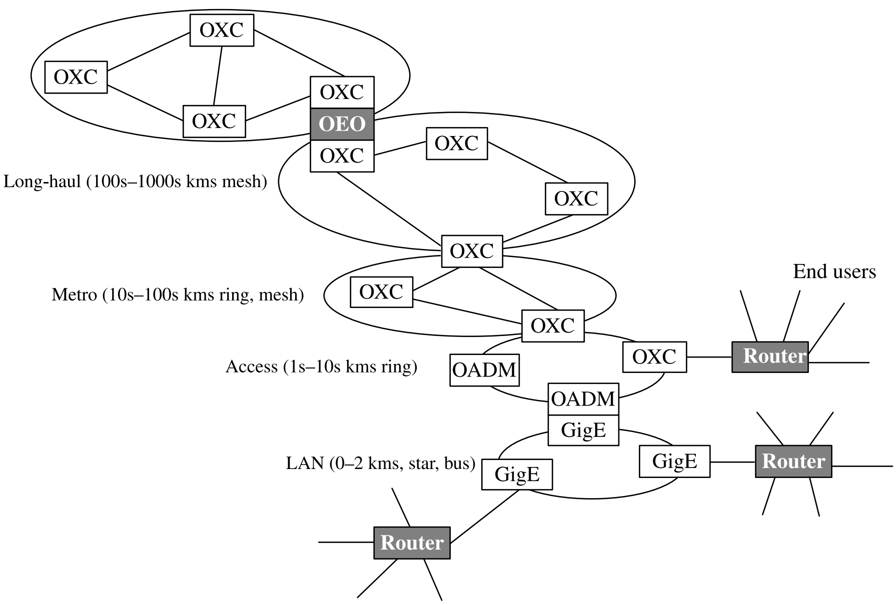

FIGURE 1.6 Expanding islands of all-optical transparency connect to each other through OEO interfaces and to users through IP routers. (© 2001 IEEE.)

overall footprint of the equipment. Therefore, the elimination of unnecessary optoelectronic conversions from a signal path in core optical mesh networks promises significant cost, footprint, and power savings. Because some optoelectronic components will unavoidably be used in the network, another approach is to reduce the cost of components used in OEO conversions. This can also prove advantageous during the transitional period of replacement. Cost reduction in OEO has led to the exploration of photonic integrated circuit technology, based on the expectation that meaningfully large-scale integration will deliver sustainable network cost reductions through volume manufacturing efficiencies, greater functional integration, and increased device density [42].

Regarding the second criterion for an all-optical network, optical layer equipment can deal with significantly higher capacities than higher-layer electronic equipment. Additionally, the operation of optical equipment is independent of the data rate and format, that is, fully transparent with respect to these two factors. Therefore, it is better to handle transit traffic at a node in the optical layer, rather than switching it through higher-layer equipment. The notion of transparency is very important in optical networks. The OEO devices are referred to as opaque because the signal goes through OEO conversions and independence from the traffic characteristics is lost. Higher-layer OEO devices should ideally be used only at the edge of the network to aggregate traffic into the optical network, rather than inside the network to handle traffic that is passing through intermediate nodes [2]. This implies that the switching of traffic should be performed in the optical domain. Therefore, optical switching is an important step in the effort towards optical networking.

1.4 SWITCHING IN OPTICAL NETWORKS

Switching is defined as the process of directing input traffic to the proper output interface as instructed by a forwarding process. Switched networks are capable of directing traffic to its destination without relying on fixed prespecified connections. Reconfigurability is a very important concern in core optical networks [44, 45]. A reconfigurable switching node can be constructed with varying degrees of transparency. Three different architectures are possible. In the first architecture, optical signals are switched by an electrical switch fabric after undergoing optical to electrical conversion and are subsequently converted back to optical for transmission. This is a case of an opaque network that utilizes an opaque switch. The opaque switch fabric may be replaced by a transparent optical switch fabric; in this second case the overall network remains opaque, because OEO conversions and associated dependencies are not avoided. A third option would be a transparent network with a transparent switch fabric that eliminates all optoelectronic conversions. This architecture is currently not feasible and will require substantial technological advances (breakthroughs to be exact) before it becomes viable.

Network designs span a continuum between fully transparent and fully opaque. Generally, hybrid networks (opaque networks with transparent switches and some regeneration sites) retain some of the advantages of full transparency but use opaque solutions, where necessary, for practical implementation. In many cases, the hybrid solutions relax the constraints on the design of the transmission lines of the network. A fully connected transparent network may require very-high-performance transmission lines, as every node of the network has to communicate with every other node without regeneration. The drawback is that hybrid solutions need some OEO regeneration, which raises the cost [42].

The following requirements for a core mesh network can be identified and should be taken into account when comparing network architectures [46]:

- *Cost minimization*. The selected network architecture should be the one that results in the lowest overall network cost for a given performance level. This means that cost comparisons should not be based on individual components but on all network elements that are necessary to achieve the desired performance objectives. For example, the cost of an optical switch with a high port count may be smaller compared to an electronic switch. The use of the optical switch without wavelength conversion capability, however, may result in a less efficient bandwidth utilization compared to the opaque electronic alternative.
- *Interoperability between components from different vendors*. A network operator should not be constrained to buying all equipment from a single vendor.
- *Elimination or minimization of manual configuration*. This means that functions such as network topology discovery, fault management, and performance management should be as automated as possible.
- *Scalability* in terms of increases in the number of channels and/or transmission rate. Such upgrades should not result in significant increases in implementation cost and footprint.

1.4.1 Optical Switching

Optical switches are devices that are capable of directing input traffic to the appropriate output interface without resorting to OEO conversions, at least in the data path. Optical switches feature a number of advantages: transparency to bit rate and data format, scalability to large port counts, cost factors, and functional simplicity.

1.4.1.1 Transparency to Bit Rate and Data Format (Protocol Independence). Optical switches are completely agnostic to the characteristics of the input traffic because they switch light and not data. This means that an optical switch fabric is bit-rate-independent and can accommodate any supported data format. Interface cards are neither data rate nor data format specific and therefore no replacements are needed when the switch is configured to operate at higher data rates (or different formats), provided the optical power budget is sufficient for that rate. Another practical implication is that the application of a new protocol does not require changes in the switching equipment and therefore can be performed quickly and easily. Because transparent architectures utilize transparent interface cards, they have no direct access to the payload, the header, or any overhead bytes for control and signaling. This means that for an $N \times N$ transparent architecture there are N interfaces/ports to the switch fabric, regardless of the type of interfaces. This is in contrast to opaque architectures where changes in data rate or format result in changes in the node architecture. Unlike electronic switches, the complexity of an optical switch fabric is a flat function independent of the bit rate of the signals it handles [46]. As bit rates rise, there is a crossover point at which the cost of a transparent switch becomes smaller than an opaque one. This crossover point is estimated based on the selected technologies and the overall costs of the two architectures.

1.4.1.2 Scalability to Large Port Counts. When optical switches are employed, capacity upgrades can be readily applied because they only require the replacement of OEO equipment at the edges of the network and no changes to intermediate nodes are necessary. In contrast, electronic routers switch data using the individual channels within a WDM link and this implies that hundreds (or even thousands) of switch interfaces must be used to terminate a single link with a large number of channels. Moreover, there can be a significant loss of statistical multiplexing efficiency when parallel channels are used simply as a collection of independent links, rather than as a shared resource [47]. As traffic increases, parallel systems have to be deployed, which raise the system volume and cost [42].

1.4.1.3 Cost Factors. The cost of an optical switch is a flat function of the port count. Furthermore, the large first cost and service lifetime cost of communication software and its supporting hardware are present only in the end nodes in an all-optical network and not in intermediate nodes [43].

1.4.1.4 Functional Simplicity. The lack of bit-level processing can be considered a disadvantage of optical switches as it complicates the routing and forwarding operations as well as network control and management functions. However, this "cut-through" approach can also be considered advantageous from certain aspects. Each signal path no longer needs to climb its way through several layers of software or firmware and back down again at each intermediate node, accumulating in the process software path length delays and exposure to the many possible failure modes intrinsic in very-high-speed electronics and all software [43].

1.4.2 Opaque Switching

An opaque network solution may be more expensive in terms of equipment costs when the core network capacity increases significantly, but it meets the following key requirements for core mesh networks [46]:

- No cascading of physical impairments. Optical signals can be regenerated whenever it is considered necessary and physical impairments are eliminated before they even become an issue; thus, signal quality concerns are not taken into account in routing decisions. This means that the lengths of paths are no longer a restricting factor for lightpaths.

- Multivendor interoperability is guaranteed.
- Wavelength conversion capability is present throughout the network.
- Access to subwavelength (sub-lambda) services is easy, because the signal is converted at each node.
- The network can be thought of as consisting of point-to-point WDM links. Therefore, each link can evolve independently and incorporate new technology, regardless of the status of other links. On the contrary, in an all-optical network all decisions for upgrades are made on an end-to-end basis [46].

In terms of network control and management operations, access to the optical signal and overhead bytes (if any) provides the following advantages [46]:

- It allows the unified handling of control and data planes because all signals are converted to electrical form.
- It allows an opaque switch to run automated neighbor and topology discovery protocols. The ability of the network autonomously to create and maintain its resource databases is the fundamental building block for an efficient, flexible, and manageable network [46].
- It allows an opaque switch to perform fault detection and performance monitoring. Lightpath-based restoration or switch fabric protection switching can be triggered by a detected failure condition. Fault isolation in such architecture relies on the alarms generated by the interface cards after a failure is detected.
- It enables the switch to provide service assurances such as performing connection verification and control to avoid misconnections.
- It allows an opaque switch to generate a keep-alive signal at every idletransceiver on the switch's network side to prevent alarms in other equipment connected to the switch. This feature is required in switch architectures that implement shared mesh restoration, because the channels have been provisioned but do not carry any signal until a failure event occurs [46].

Opaque switches are faced with a number of challenges relating mainly to traffic growth. Scaling issues arise when the signal bit rate and/or the switch matrix port count increase and lead to increased network cost. It must be noted, however, that some opaque switches may be preserved in a network that utilizes transparent switches in order to provide some key network functions: grooming and multiplexing, service level agreement verification, and control and management [46].

1.4.3 Challenges for Optical Switching

A network architecture where signals remain entirely in the optical domain guarantees end-to-end bit rate and data format transparency. The following factors, however, challenge the applicability of such an architecture in core mesh networks [46]:

- Network performance heavily depends on the presence of wavelength conversion capability. If wavelength conversion is not possible, a network with W WDM channels can be thought of as a network with W disjoint layers, because traffic cannot be transferred between layers [46]. Such inflexible use of wavelengths in the network leads to increased bandwidth and network operational cost, and thus negates savings resulting from the elimination of OE conversions.
- Physical impairments such as chromatic dispersion, polarization mode dispersion, fiber nonlinearities, and amplifier noise accumulate over the physical path of the signal due to the absence of OE conversions [46]. Furthermore, signal-to-noise ratio concerns may arise, especially for data paths with high hop counts. In general, fiber impairments need to be taken into account in end-to-end routing decisions. If signals are to remain in the optical domain throughout a backbone network that spans for example Europe, an all-optical 3R regeneration function (combining retiming, reshaping, and reamplification) must be available. Although promising technologies in this direction have been demonstrated, no all-optical 3R products have reached the market [46].
- Transparent-based restoration is a lower-cost approach than opaque-based restoration, as it saves a number of opaque interface ports. However, it may result in slower restoration (possibly on the order of seconds), and it requires new out-of-band signaling channels between transparent switches, and between transparent switches and client equipment [46].
- The lack of wavelength conversion capability hampers the application of shared mesh protection schemes in all-optical networks. This is due to the fact that in the case of a network fault, a signal that needs to be transferred to an alternative backup path must remain in the same wavelength because there is no conversion capability. If the desired wavelength is not free in the shared backup path, data loss is inevitable. This means that either more backup paths are needed to achieve performance objectives comparable to shared mesh protection in opaque networks or dedicated protection schemes need to be applied [46].
- Operators in an all-optical networks do not have the flexibility to select the client network elements and WDM vendors independently. This is because the interface optics at the client network element launch the signals through the all-optical switch directly into the WDM system without OE conversion.

Opaque architectures with optical switches also face significant challenges regarding the support of control and management functionalities that are readily available when there is access to the electrical signal and to any overhead bytes [46]. Network control and management features are collectively very difficult to achieve in a transparent switch. Important control and management functions include the following:

- *Automatic port/neighbor and topology discovery*. This is an important requirement for a flexible and manageable network. This function can be implemented by a transparent switch either by employing a specialized link management protocol or by using a small number of opaque interface cards providing electrical signal generation functions. These cards conduct neighbor/port discovery in an automated process by establishing a connection between the opaque interface cards of two adjacent switches [46].
- *Keep-alive signal generation*. This feature is useful in shared mesh protection schemes where a number of channels are provisioned in advance to serve as backup for primary paths. Until these channels are used for data transmission (i.e., until a failure occurs), these channels require the presence of a keep-alive (unequipped) signal. The lack of such a signal results in (1) alarms generated at the WDM systems that have knowledge of provisioned channels but detect no light on those channels, (2) lack of monitoring of the restoration channels to ensure availability when/if a failure occurs, and (3) increased restoration time if a failure occurs, because of the additional time required to turn on the WDM lasers and perform power adjustments and equalization. This function is handled in transparent switches by using a limited number of lasers at the drop side of the optical switches and having each signal propagate along several idle channels, looping back and forth between the switches [46].

1.5 OPTICAL SWITCHING PARADIGMS

Chapters 3, 4, and 5 present the three emerging optical switching paradigms, namely optical packet switching (OPS), generalized multiprotocol label switching (GMPLS), and optical burst switching (OBS).

Optical packet switching can be thought of as the equivalent of electronic packet switching. In OPS networks, optical packets are composed of a header and payload, which are tightly coupled in time. At each node, the header undergoes processing (electronic processing for the time being and the near future) and routing information is extracted. Based on this information, the switching node reconfigures its (optical) switching fabric and directs the packet to the appropriate output port transparently and without OEO conversions. Because the packet size can be almost arbitrarily small, optical packet switching offers bandwidth allocation with fine granularity and hence guarantees a high degree of statistical multiplexing. Furthermore, the dynamic manner in which network resources are allocated to packets (on demand) offers excellent scalability.

For these reasons, optical packet switching is often viewed as the ultimate goal in optical networks evolution. The application of optical packet switching in practical systems, however, is hindered by the two fundamental limitations of optics, namely the lack of bit-level processing and buffering capabilities. Challenges surrounding

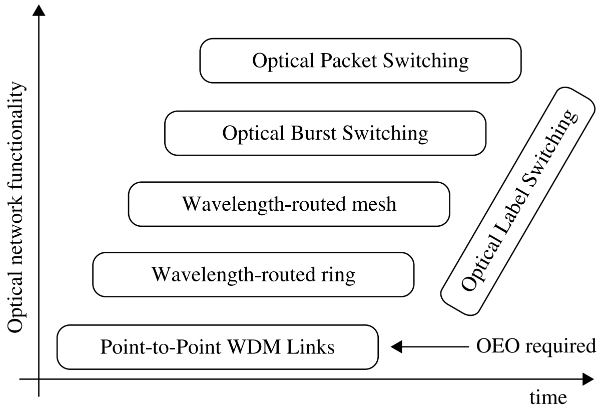

FIGURE 1.7 The gradual adoption of optical networking paradigms.

optical packet switching include the synchronization/alignment of headers and payloads, the processing of optical headers and the resolution of contentions between packets. As can be seen in Figure 1.7 [48, 49], the applicability of OPS is placed well into the future compared to the other techniques.

Generalized multiprotocol label switching, as its name suggests, is a generalization of the multiprotocol label switching technique, which emerged as a way to simplify IP-based forwarding and to incorporate support for service differentiation. In GMPLS, bandwidth is provisioned on demand by setting up label switched paths with the use of signaling and resource reservation protocols specifically designed for optical networks. In addition, GMPLS facilitates the aggregation of lower-order paths into higher-order ones for more efficient bandwidth utilization.

Optical burst switching is an approach that combines the advantages of packet switching and wavelength routing. At the edges of an OBS network, packets are assembled to form bursts, which will traverse the optical core transparently and be disassembled into the original packets on the other side. Shortly before the transmission of each burst, a control packet is sent to attempt to reserve the resources necessary for the burst on all nodes in its path. The burst follows its control packet without waiting for an acknowledgement that the required resources have been reserved. Optical burst switching facilitates the statistical sharing of resources, but with an intermediate granularity (a burst) compared to wavelength routing and optical packet switching. As a result, the technical issues surrounding the application of OBS that need to be resolved are fewer than those for optical packet switching. At the same time, OBS guarantees a more efficient utilization of resources compared to wavelength routing, in which the holding times of lightpaths need to be much longer than their setup times [50]. On the other side, because bursts are transported over a bufferless core and without prereservation of bandwidth, there are no guarantees for successful delivery, and losses are relatively high.

REFERENCES

1. C. A. Brackett, "Dense Wavelength Division Multiplexing Networks: Principles and Applications," *IEEE Journal on Selected Areas in Communications*, 8, 948–964 (1990).
2. R. Ramaswami, "Optical Fiber Communication: From Transmission to Networking," *IEEE Communications Magazine*, 40, 138–147 (2002).
3. M. El-Sayed and J. Jaffe, "A View of Telecommunications Network Evolution," *IEEE Communications Magazine*, 40, 74–81 (2002).
4. P. E. Green, Jr., "Fiber to the Home: The Next Big Broadband Thing," *IEEE Communications Magazine*, 42, 100–106 (2004).
5. R. Ramaswami and K. N. Sivarajan, *Optical Networks, A Practical Perspective*. Morgan Kaufmann Publishers, Inc., San Francisco, USA, 1998.
6. T. E. Stern and K. Bala, *Multiwavelength Optical Networks, A Layered Approach*. Addison Wesley Longman, Inc., Boston, USA, 1999.
7. P. E. Green, Jr., *Fiber Optic Networks*. Prentice-Hall, Inc., New Jersey, USA, 1993.
8. P. Gravey, S. Gosselin, C. Guillemont, D. Chiaroni, N. Le Sauze, A. Jourdan, E. Dotaro, D. Barth, P. Bertome, C. Laforest, S. Vial, T. Atmaca, G. Hebuterne, H. El Biaze, R. Laalaoua, E. Gangloff, and I. Kotuliak, "Multiservice Optical Network: Main Concepts and First Achievements of the ROM Program," *IEEE/OSA Journal of Lightwave Technology*, 19, 23–31 (2001).
9. S. J. B. Yoo, F. Xue, Y. Bansal, J. Taylor, Z. Pan, J. Cao, M. Jeon, T. Nady, G. Goncher, K. Boyer, K. Okamoto, S. Kamei, and V. Akella, "High-Performance Optical-Label Switching Packet Routers and Smart Edge Routers for the Next-Generation Internet," *IEEE Journal on Selected Areas in Communications*, 21, 1041–1051 (2003).
10. B. St. Arnaud, J. Wu, and B. Kalal, "Customer-Controlled and -Managed Optical Networks," *IEEE/OSA Journal of Lightwave Technology*, 21, 2804–2810 (2003).
11. M. Veeraraghavan, R. Karri, T. Moors, M. Karol, and R. Grobler, "Architectures and Protocols that Enable New Applications on Optical Networks," *IEEE Communications Magazine*, 39, 118–127 (2001).
12. A. Houghton, "Supporting the Rollout of Broadband in Europe: Optical Network Research in the IST Program," *IEEE Communications Magazine*, 41, 58–64 (2003).
13. D. Colle, P. Demeester, P. Lagasse, P. Falcao, and P. Arijs, "Application, Design, and Evolution of DWDM in Pan-European Transport Networks," *IEEE Communications Magazine*, 41, 48–50 (2003).
14. G. I. Papadimitriou, C. Papazoglou, and A. S. Pomportsis, "Optical Switching: Switch Fabrics, Techniques, and Architectures," *IEEE/OSA Journal of Lightwave Technology*, 21, 384–405 (2003).
15. S. A. Hamilton, B. S. Robinson, T. E. Murphy, S. J. Savage, and E. P. Ippen, "100 Gb/s Optical Time-Division Multiplexed Networks," *IEEE/OSA Journal of Lightwave Technology*, 20, 2086–2100 (2002).
16. A. Rodriguez-Moral, P. Bonenfant, S. Baroni, and R. Wu, "Optical Data Networking: Protocols, Technologies, and Architectures for Next Generation Optical Transport Networks and Optical Internetworks," *IEEE/OSA Journal of Lightwave Technology*, 18, 1855–1870 (2000).

17. E. Le Rouzic and S. Gosselin, "160-Gb/s Optical Networking: A Prospective Techno-Economical Analysis," *IEEE/OSA Journal of Lightwave Technology*, 23, 3024–3033 (2005).
18. P. Bonenfant, "Optical Data Networking: What Bubble?," *IEEE Communications Magazine*, 41, 46–47 (2003).
19. G. I. Papadimitriou1, M. S. Obaidat, and A. S. Pomportsis, "Advances in Optical Networking," *Wiley International Journal of Communication Systems*, 15, 101–113 (2001).
20. M. Tatipamula, F. Le Faucheur, T. Otani, and H. Esaki, "Implementation of IPv6 Services over a GMPLS-Based IP/Optical Network," *IEEE Communications Magazine*, 43, 114–122 (2005).
21. S. Dixit and Y. Ye, "Streamlining the Internet-Fiber Connection," *IEEE Spectrum*, 52–57 (2001).
22. N. Ghani, S. Dixit, and T. S. Wang, "On IP-WDM Integration: A Retrospective," *IEEE Communications Magazine*, 41, 42–45 (2003).
23. J. Y. Wei, "Advances in the Management and Control of Optical Internet," *IEEE Journal on Selected Areas in Communications*, 20, 768–785 (2002).
24. A. Kaheel, T. Khattab, A. Mohamed, and H. Alnuweiri, "Quality-of-Service Mechanisms in IP-over-WDM Networks," *IEEE Communications Magazine*, 40, 38–43 (2002).
25. D. Benjamin, R. Trudel, S. Shew, and E. Kus, "Optical Services over the Intelligent Optical Network," *IEEE Communications Magazine*, 39, 73–78 (2001).
26. N. Ghani, S. Dixit, and T. Wang, "On IP-over-WDM Integration," *IEEE Communications Magazine*, 38, 72–84 (2000).
27. G. I. Papadimitriou, P. Tsimoulas, M. S. Obaidat, and A. S. Pomportsis, *Multiwavelength Optical Networks*. John Wiley & Sons, West Sussex, England, 2003.
28. M. Scheutzow, M. Maier, M. Reisslein, and A. Wolisz, "Wavelength Reuse for Efficient Packet-Switched Transport in an AWG Based Metro WDM Network," *IEEE/OSA Journal of Lightwave Technology*, 21, 1435–1455 (2003).
29. Y. Xin, G. N. Rouskas, and H. G. Perros, "On the Physical and Logical Topology Design of Large-Scale Optical Networks," *IEEE/OSA Journal of Lightwave Technology*, 21, 904–915 (2003).
30. C. Nuzman, J. Leuthold, R. Ryf, S. Chandrasekhar, C. R. Giles, and D. T. Neilson, "Design and Implementation of Wavelength-Flexible Network Nodes," *IEEE/OSA Journal of Lightwave Technology*, 21, 648–663 (2003).
31. E. Iannone and R. Sabella, "Optical Path Technologies: A Comparison among Different Cross-Connect Architectures," *IEEE/OSA Journal of Lightwave Technology*, 14, 2184–2196 (1996).
32. J. Strand, A. Chiu, and R. Tkach, "Issues for Routing in the Optical Layer," *IEEE Communications Magazine*, 39, 81–87 (2001).
33. H. Zang, J. P. Jue, L. Sahasrabuddhe, R. Ramamurthy, and B. Mukherjee, "Dynamic Lightpath Establishment in Wavelength-Routed WDM Networks," *IEEE Communication Magazine*, 39, 100–108 (2001).
34. C. Papazoglou, G. I. Papadimitriou, and A. S. Pomportsis, "Design Alternatives for Optical-Packet-Interconnection Network Architectures," *OSA Journal of Optical Networking*, 3, 810–825 (2004).

35. C. Guan and V. W. S. Chan, "Topology Design of OXC-Switched WDM Networks," *IEEE Journal on Selected Areas in Communications*, 23, 1670–1686 (2005).
36. M. Herzog, M. Maier, and M. Reisslein, "Metropolitan Area Packet-Switched WDM Networks: A Survey on Ring Systems," *IEEE Communications Surveys & Tutorials*, 6, 2–20 (2004).
37. D. Simeonidou, R. Nejabati, G. Zervas, D. Klonidis, A. Tzanakaki, and M. J. O'Mahony, "Dynamic Optical-Network Architectures and Technologies for Existing and Emerging Grid Services," *IEEE/OSA Journal of Lightwave Technology*, 23, 3347–3357 (2005).
38. A. Jajszczyk, "Automatically Switched Optical Networks: Benefits and Requirements," *IEEE Communications Magazine*, 43, S10–S15 (2005).
39. C. V. Saradhi, M. Gurusamy, and L. Zhou, "Differentiated QoS for Survivable WDM Optical Networks," *IEEE Communications Magazine*, 42, S8–S14 (2004).
40. S. Ramamurthy, L. Sahasrabuddhe, and B. Mukherjee, "Survivable WDM Mesh Networks," *IEEE/OSA Journal of Lightwave Technology*, 21, 870–883 (2003).
41. O. Gerstel and R. Ramaswami, "Optical Layer Survivability: A Post-Bubble Perspective," *IEEE Communications Magazine*, 41, 51–53 (2003).
42. J. Livas, "Optical Transmission Evolution: From Digital to Analog to ? Network Trade-offs Between Optical Transparency and Reduced Regeneration Cost," *IEEE/OSA Journal of Lightwave Technology*, 23, 219–224 (2005).
43. P. Green, "Progress in Optical Networking," *IEEE Communications Magazine*, 39, 54–61 (2001).
44. S. Sengupta, V. Kumar, and D. Saha, "Switched Optical Backbone for Cost-Effective Scalable Core IP Networks," *IEEE Communications Magazine*, 41, 60–70 (2003).
45. D. Chiaroni, "Packet Switching Matrix: A Key Element for the Backbone and the Metro," *IEEE Journal on Selected Areas in Communications*, 21, 1018–1025 (2003).
46. G. Ellinas, J. A. Walker, S. Chaudhuri, L. Lin, E. Goldstein, and K. Bala, "Network Control and Management Challenges in Opaque Networks Utilizing Transparent Optical Switches," *IEEE Communications Magazine*, 42, S16–S24 (2004).
47. M. Listanti, V. Eramo, and R. Sabella, "Architectural and Technological Issues for Future Optical Internet Networks," *IEEE Communications Magazine*, 38, 82–92 (2000).
48. K. Dolzer, C. Gauger, J. Spath, and S. Bodamer, "Evaluation of Reservation Mechanisms for Optical Burst Switching," *AEU International Journal of Electronic Communications*, 55, 18–26 (2001).
49. T. Battestilli and H. Perros, "An Introduction to Optical Burst Switching," *IEEE Communications Magazine*, 41, S10–S15 (2003).
50. C. Qiao and M. Yoo, "Choices, Features and Issues in Optical Burst Switching," *Optical Networks Magazine*, 1, 36–44 (2000).
51. J. He, S. H. G. Chan, and D. H. K. Tsang, "Multicasting in WDM Networks," *IEEE Communications Surveys & Tutorials*, 4, 2–20 (2002).
52. V. W. S. Chan, K. L. Hall, E. Modiano, and K. A. Rauschenbach, "Architectures and Technologies for High-Speed Optical Data Networks," *IEEE/OSA Journal of Lightwave Technology*, 16, 2146–2168 (1998).