1 The Big Picture

Since the beginning of the 21st century there has been a burgeoning demand for communications services. From the ubiquitous mobile phone, providing voice, images, messaging, and more, to the Internet and the World Wide Web, offering bandwidth-hungry applications such as interactive games, music, and video file sharing, the public's appetite for information continues to grow at an ever-increasing pace. Underneath all of this, essentially unseen by the users, is the optical fiber-based global communications infrastructure – the foundation of the information superhighway. That infrastructure contains the multiwavelength optical networks that are the theme of this book.

Our purpose is to present a general framework for understanding, analyzing, and designing these networks. It is applicable to current network architectures as they have evolved since the mid-1990s, but more importantly it is a planning and design tool for the future. Our approach is to use a generic methodology that will retain its relevance as networks, applications, and technology continue to evolve.

1.1 Why Optical Networks?

Since the fabrication of the first low-loss optical fiber by Corning Glass in 1970, a vision of a ubiquitous and universal all-optical communication network has intrigued researchers, service providers, and the general public. Beginning in the last decades of the 20th century enormous quantities of optical fiber were deployed throughout the world. Initially, fiber was used in point-to-point transmission links as a direct substitute for copper, with the fibers terminating on electronic equipment. Glass fiber was and is the ideal medium because of its many superior properties: extraordinary bandwidth, low loss, low cost, light weight and compactness, strength and flexibility, immunity to noise and electromagnetic interference, security and privacy (it is difficult to tap them), and corrosion resistance. Although all of these qualities make the fiber a technological marvel, fibers do not become networks until they are interconnected in a properly structured architecture. For our purposes, an *optical network* is a telecommunications network with transmission links that are optical fibers, and with an architecture that is designed to exploit the unique features of fibers. (Most of the communication systems in use today, including many specialized networks such as cable TV and mobile phone systems, have optical fiber in them somewhere; however, this does *not* make them optical networks.) As we shall see, suitable architectures for high-performance lightwave networks involve complex combinations of both optical and electronic devices. Thus, as used here, the term *optical* or *lightwave network* does not necessarily imply a *purely* optical network, but it does imply something more than a set of fibers interconnecting electronic switches.

As optical and photonic technology has advanced, applications to point-to-point transmission have preceded advances in networking. For example, it was clear in the early years of optical fiber transmission that by introducing wavelength division multiplexing (WDM) on existing fibers the capacity of a fiber link could be increased manyfold at minimum cost. However, it was only since the early 2000s that the optical switching technology necessary to convert isolated fiber transmission links to optical networks matured sufficiently to permit the commercial deployment of these networks. In the mid-1990s, the optical network (as opposed to optical fiber transmission alone) was still a "blue sky" concept. New optical and photonic devices were being developed and incorporated into experimental networks. But full-fledged multiwavelength networks integrating optical transmission, switching, and user access were still in the research and development stage. At that time the *technology push* for networking was out in front, but demand for the seemingly unlimited capacity of these networks was essentially nonexistent. As this is being written, the promise of optical networking is finally being fulfilled. The *demand pull* for these networks has materialized. As low-cost broadband services are made available to the general public, demand for Internet-based applications continues to increase. Equipment manufacturers, communications carriers, and service providers have joined in moving optical networking from feasibility studies to commercial viability in both cost and performance. The focus in the networking community has now shifted to organization, control, manageability, survivability, standardization, and cost-effectiveness, a trend that reflects the maturing of the optical technology as well as the recognition that the optical network is the only way of supporting current and future demand. These networks have played a critical role in reducing communications costs, promoting competition among carriers and service providers, and thereby increasing the demand for new services.

In addition to the technology push and demand pull, a number of other recent developments are contributing to the expansion and effectiveness of optical networks. One is the accelerating removal of the bottleneck in the "last mile" – the distribution network that is the bridge between the high-speed fiber core network and the end users. Until the last decade of the 20th century this distribution network – composed of twisted pairs of copper wires connecting each residential subscriber to the local telephone Central Office – was specifically engineered to a limited bandwidth of 3000 Hz. As a result the user bit rates were restricted to a tiny trickle. This low-speed access link separated the various high-speed communications and computing devices located on the premises of the end users (e.g., PCs, TV displays, and music/image/video storage equipment in the home) from the high-capacity network serving this equipment. Considering that the processors in today's PCs operate at speeds six orders of magnitude faster than a low-speed access link, and the optical fibers in the network have bandwidths nine orders of magnitude wider than the bandwidth of the access link, it is obvious that access was – and is – a severe problem. As long as the last-mile bottleneck is present, the information superhighway is still a dirt road; more accurately, it is a set of isolated multilane highways with cow paths for entrance and exit ramps. The introduction of broadband access to residential customers by the telephone carriers and the cable operators is a step toward eliminating those cow paths. However, digital subscriber line (DSL) and cable modems are half-measures at best. Direct access to the fiber network by the end user [i.e., fiber to the home (FTTH) or business user] is the ultimate way of removing the bottleneck so that the network remains effective as demand for bandwidth grows. Although FTTH was deployed many years ago in a few demonstration projects, it did not take hold for several reasons, including cost and the absence of services of interest to the customers. Today that has changed because of the proliferation of broadband Internet services. Deployment of glass is now moving from the network core through fiber access networks to the end users. This will undoubtedly stimulate interest in new broadband services that take advantage of high-speed access and in turn produce demand for more bandwidth. At this writing, most of the world's installed fiber capacity is underutilized – arguably due to the last-mile bottleneck. That should change rapidly as progress in the removal of the bottleneck results in a quantum jump in network traffic, making high-performance optical networks indispensable.

Higher level issues such as deregulation, new ideas for improving the economics of networking, and standardization of control and management techniques in multivendor networks are also contributing to the growing effectiveness of optical networks. Deregulation, which began in the United States in 1984 with the dismantling of AT&T, has brought with it a new level of competition, with long-haul carriers, local carriers, Internet service providers (ISPs), and cable operators poaching on each other's domains and using optical fiber capacity to do so. Bandwidth trading has been introduced as a way of improving the utilization of fibers and thereby optimizing profits. A carrier with idle capacity sells it to another carrier with excess demand. This type of exchange requires sophisticated control and management tools for network reconfiguration. More generally, any large network requires complex control and management systems and intelligent network elements for performance monitoring, network reconfiguration, and fault recovery. The systems, protocols, and equipment for performing these functions in traditional telephone and data networks were built over many years by the public carriers and equipment manufacturers. The new optical networks require similar tools, and this is especially important in multivendor environments. These are now making their appearance in the form of a proposed control plane for optical networks and protocols for systems management in these networks. As more sophisticated control and management functions are incorporated into optical networks the network operators are in a better position to offer high-quality service to their customers, improving the operator's revenue stream and customer loyalty.

Above all, the lessons of the past show us that tomorrow's networks must be flexible and versatile enough to adapt to a continuing barrage of new and as-yet-unknown services. It is interesting to note that when optical networks were still in an embryonic form, the typical uses envisaged for them were high-tech applications such as high-resolution medical image archiving and remote supercomputer visualization – basically usages generated by a minuscule, elite segment of the population. Today these applications represent but a tiny part of the global network traffic, submerged in a torrent generated by the common man, who has only recently gained access to the enormous opportunities our worldwide communication system has to offer. The networks we conceive today must be "futureproof" so as to be ready for the next unforeseen developments.

1.2 Objectives of an Optical Network Architecture

Today's and tomorrow's optical networks must provide the capacity, connectivity, and intelligence necessary to link together a global community of information providers and consumers. A well-designed network performs this function efficiently and reliably. To facilitate a systematic study of networks that achieve this goal, it is useful to formulate a generic model in the form of a multiwavelength network architecture (MWNA). As background for the MWNA we briefly review the current network structures and the services they support.

Until the end of the second millennium, the world of networking consisted of two separate spheres: the traditional telephone networks mainly devoted to providing voice services (operated in a circuit-switched mode) and data networks (operated in a packet-switched mode) for communication between computers. Each type of network was specially engineered and optimized for its own type of service. Circuit switching was the preferred approach to voice transmission, because the voice signal was transmitted as a continuous stream, whereas packet switching was invented to carry data traffic because data signals were bursty in nature, making the circuit-switching approach very inefficient. Because the voice networks operated by the public carriers contained virtually all of the world's installed communication capacity, the early data networks were constructed as overlays on these networks, running on lines leased from the public carriers – mainly AT&T in the United States and the government administrations in Europe. The traffic flow in the early data networks was minute compared to voice traffic – essentially confined to businesses, universities, and research laboratories. For this reason the main players in data networking were originally government, research, business, and educational organizations and data-processing equipment manufacturers.

As optical fiber became the dominant transmission medium, various standards for exploitation of fiber were developed, including the synchronous optical network (SONET) standard in the United States and a similar synchronous digital hierarchy (SDH) standard in Europe. The SONET/SDH transmission, multiplexing and switching equipment, adapted primarily to circuit-switched applications, was soon augmented by asynchronous transfer mode (ATM) switches and Internet Protocol (IP) routers (cell-switched and packet-switched, respectively) to handle a wide variety of data and multimedia services. By the late 1990s the traditional separation of voice and data networks changed significantly. In a very short time we moved from a voice-centric world to a data-centric world, and, more importantly, the techniques of carrying data (packet switching) were extended to an infinite variety of services having no resemblance to those in the traditional computer world. Internet/Web services, running the gamut from interactive computer games through telemedicine to peer-to-peer file sharing, now use IP for transmitting anything from computer data to video (Internet Protocol TV; IPTV) to old-fashioned voice (Voice over IP; VoIP).[1]

This brings us to the characteristics and requirements of the services supported by the optical networks discussed here. These are extremely diverse in terms of connectivity, bandwidth, performance, survivability, cost, and a host of other features.

Consider the common Internet services offered to the general public (e.g., e-mail and search engines). They serve a vast globally distributed user community. In terms of connectivity, these types of services push networking to its ultimate limits; any end user wants rapid connectivity to anyone or anything in the network. However, in terms of performance, they are undemanding – they can tolerate errors, delays, and occasional downtimes due to congestion, programming bugs, and equipment failure. Total costs may be high, but they are spread over an enormous user base resulting in a very low cost per user.

In contrast, consider a different type of application, the virtual private network (VPN). This is a subnet carved out of a larger network by a telecommunications carrier and put at the disposal of a single enterprise, which typically controls and manages it. Consequently it has a much smaller user group with more intense utilization per user, far fewer active connections, and tighter control of network performance, including security and reliability. Customer costs per user will be higher, but this is offset by higher performance and more responsiveness to the needs of the customer.

Another example is telemedicine, which requires high-quality communication (e.g., high fidelity medical image encoding and transmission, and rapid response) and where cost is secondary. Different requirements apply to public safety services (e.g., police, fire, and disaster relief), which depend on a high degree of survivability, fault recovery, and availability[2] in the face of equipment and line failures, natural disasters, or malicious attacks. Transmission quality is secondary. Similar requirements hold to a lesser degree for financial services (e.g., banks and brokerage houses). In public safety and financial service applications, cost is not the primordial issue.

To ensure satisfactory service, large users of network services (e.g., enterprises operating VPNs) enter into service-level agreements (SLAs) with the service providers. For example, the SLA might specify a level of availability, network delay, packet loss, and other features. These represent promises from the provider to the user, and as such they must be backed up by suitable controls within the underlying network to achieve the performance stated in the SLA. These controls are enforced within a large network by identifying *differentiated services*, that is, traffic flows that are singled out to be provided with a predictable quality of service (QoS) (e.g., limits on packet loss and delay). Traffic routed through a large network can be tagged to recognize its class of service (CoS), thereby facilitating the satisfaction of service requirements through mechanisms such as priority packet queueing, bandwidth allocation, and service recovery priority.

[1] The increased interest in VoIP, because of its low cost and growing ubiquity is, to paraphrase Shakespeare, the most unkindest cut of all from the computer community to the traditional telephone carriers.

[2] Availability is the percentage of time that a network is operational. For example, "five 9s" (99.999%) availability, which is a goal for public carriers, implies 5.25 minutes of downtime per year.

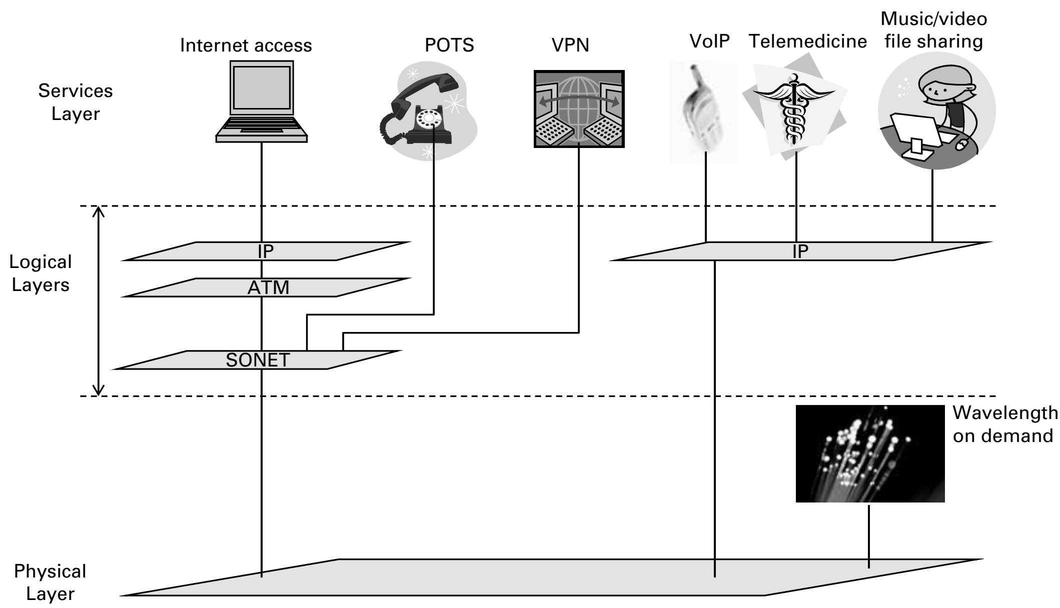

Figure 1.1 Multilayered network.

The various functions executed by the network operator, such as load balancing and QoS-based traffic handling, are known as *traffic engineering*.

Considering the wide diversity of these service requirements (and we have only mentioned a small sample), it is a challenge to support all of them on a single network. Yet this is not only possible, but it is generally the most efficient way of doing the job. As we will show, the building blocks are now available to assemble multiwavelength optical networks that can sustain large user populations with diverse service requirements of the type just described. This means that, ideally, the MWNA must be structured to offer a special set of features adapted to each service it supports. To see how this is achieved it is convenient to think of the network in terms of its constituent layers, with client–server relations between the neighboring layers. An illustration is the multilayered view shown in Figure 1.1. The architecture is composed of an underlying optical infrastructure – the *physical layer* – which provides basic communication services to a number of independent *logical networks* (LNs) residing in the *logical layer*. Each LN organizes the raw capacity offered by the physical layer, adapting it to the needs of the clients it serves, shown in the *services layer* of the figure. For example, the SONET network shown in Figure 1.1 uses optical wavelength channels provided by the physical layer, transmits optical signals on them, and carries multiplexed communication channels on those signals. The SONET channels can be tailored to support a wide variety of services; two services shown in the figure are plain old telephone service (POTS) and a VPN. In our example, the SONET layer also supports an ATM layer that in turn supports a client IP layer providing Internet access services to end users. Another independent IP network shown in Figure 1.1 is supported directly by the physical layer, providing a telemedicine service, VoIP, and a music/video file-sharing service. In addition, the physical layer provides purely optical connections directly to end users via demand-assigned wavelengths (also known as *clear channels*), thereby bypassing the logical layer altogether.

Thus, the logical layer shown in the figure contains several LNs; some are stacked in a client–server relationship, and others are independent of each other, offering specialized features to the service layer. Stacked logical layers; e.g., IP over ATM over SONET over WDM, have both advantages and disadvantages. For example, different services (e.g., POTS and VPNs) require channels running at different bit rates. The SONET layer supports these different speeds and in addition provides a *grooming* function, packing the diverse channels onto a common optical wavelength, using time division multiplexing. This “fills up” the wavelength channel for efficient utilization. However, stacked layers mean additional equipment, which is costly, introduces delays and potential points of failure, and is difficult to manage. Therefore, it is desirable to reduce superfluous layers wherever possible. For example, the IP equipment manufacturers propose to provide *IP over WDM*, short-circuiting the commonly used configurations involving stacked intermediate layers.[3]

Another view of an optical network is the physical picture of Figure 1.2, showing the network elements in the layers of Figure 1.1. Here the physical layer is portrayed for simplicity as a *transparent purely optical* core.[4]

The “glue” in the physical layer that holds the transparent optical network together fits roughly into two basic classes: the optical network nodes (ONNs), which connect the fibers within the network, and the network access stations (NASs), which interface end-user systems and other nonoptical equipment to the network. Shown as rectangles in Figure 1.2, the NASs (or *stations* for short) provide the terminating points (sources and destinations) for the optical signal paths within the physical layer. The communication paths continue in electrical form outside the purely optical part of the network, either terminating at *end systems* (for example, PCs, telephones, and servers) or traversing electronic multiplexing and switching equipment (e.g., ATM switches, IP routers, or SONET digital cross-connect systems [DCSs]), shown as hexagons in Figure 1.2. The ONNs (or *nodes* for short), shown as circles in Figure 1.2, provide the switching and routing functions that control the optical signal paths (also called *lightpaths*), configuring them to create desired source-destination connections. The stations and nodes contain the optoelectronic and photonic components of the network: lasers, detectors, couplers, filters, optical switches, amplifiers, and so on. These components work together with the fibers to produce the required optical signal connectivity. Although the underlying optoelectronic and photonic technologies have matured considerably since the mid-1990s, they are not as well developed as their electronic counterparts. Thus, electronics (in the logical layer) is currently an equal partner with photonics (in the physical layer).

[3] In many cases, superfluous LN stacking results from a reluctance of carriers to write off a large investment in legacy systems, which would require a complete revision of existing control and management structures.

[4] As will be seen, the physical layer in optical networks often includes electronic components in the form of signal regenerators or electronic switch fabrics, so it is not always purely optical nor is it completely transparent. We have more to say about purely optical signal paths and the meaning of transparency in Sections 1.3 and 1.4.

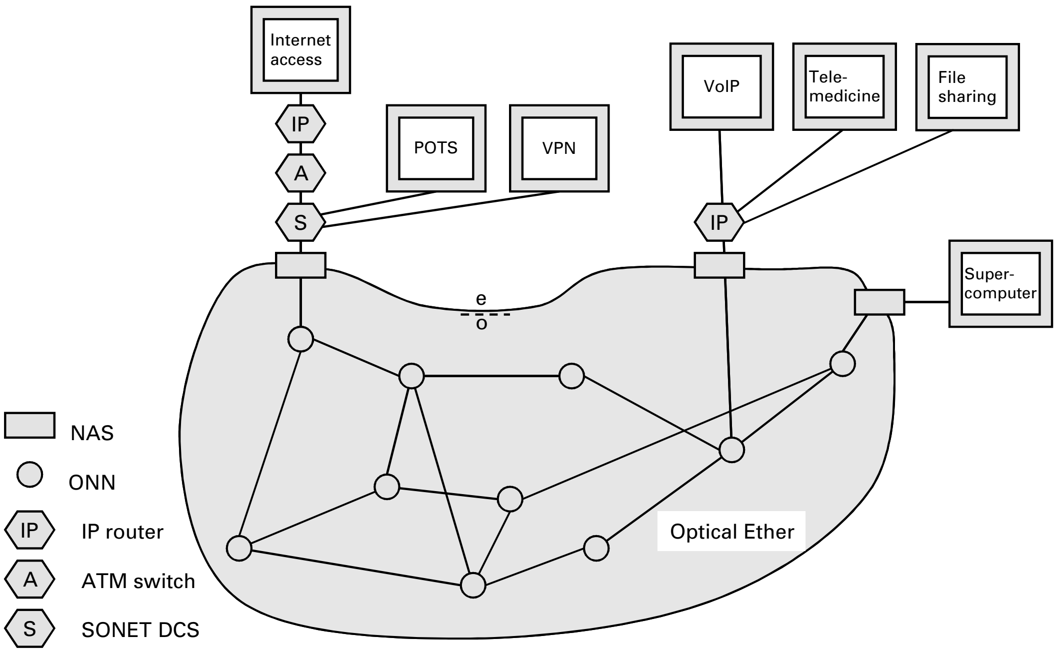

Figure 1.2 Physical picture of the network.

The line between the optical and electronic parts of the network has become fuzzy as technology has advanced, but at this point in our discussion we retain the simplified view that the physical layer is transparent and optical, whereas the logical layers are electronic and "opaque."

The electronic switching node plays the same role in the logical layer as ONNs play in the physical layer. Our generic term for an electronic switching node (each hexagon in Figure 1.2) is a *logical switching node* (LSN). The LSNs sort, multiplex, switch, and route signals in the various LNs. In this way they create *virtual connections* among the entities they serve. These entities may be service provider equipment or user equipment (end systems), as in the case where the LSNs are IP routers connecting ISP servers to customer PCs, or they may be higher layer switching nodes in a client network, as in the case of SONET DCSs serving ATM switches serving IP routers.

Although the focus of this book is the physical layer and its optical components, the logical layers are an integral part of the overall network architectures we discuss here. Therefore, our MWNA *includes* the logical layers and their electronic components. An understanding of the design and operation of multiwavelength optical networks requires an awareness of the close coupling between the physical layer and the logical layers it serves.

The networks we examine will generally be designed to serve large, heterogeneous, geographically dispersed user populations. Given this fact, and the various service requirements discussed, we can infer a list of general design and operating objectives:

- Connectivity
  - Support a very large number of end systems
  - Support a very large number of concurrent connections, including multiple connections per station and per end system
  - Support multicast connections efficiently
- Performance
  - High aggregate network throughput (hundreds of terabits per second)
  - High fiber transmission capacity (terabits per second)
  - High user bit rate (gigabits per second)
  - Small end-to-end delay
  - Low error rate (digital)/high signal-to-noise ratio (analog)
  - Low processing load in nodes and stations
  - Adaptability to changing and unbalanced loads
- Technology: cost-effective design and utilization
  - Access stations: limited number of optical transceivers per station, limited complexity of optical transceivers, simple tuning techniques
  - Optical network nodes: high throughput, minimal signal impairment, and low complexity
  - Logical switching nodes: efficient channel grooming, simple packet-routing procedures, and controlled traffic load
  - Network: Limited number and length of cables and fibers, efficient use (and reuse) of the optical spectrum, controlled signal impairment in the physical layer, minimization of logical layer complexity.
- Structural features
  - Scalability
  - Modularity
  - Survivability (fault tolerance)
- Control and management
  - Efficient, rapid, automated connection provisioning, and reconfiguration
  - Built-in intelligence in the network elements for monitoring and control
  - Efficient and rapid automatic fault identification and recovery
  - An integrated network management system to monitor and coordinate all network layers

As we look at existing and proposed network architectures, it is important to keep these goals in mind.

1.3 Optics versus Electronics: The Case for Transparent Multiwavelength Networks

There are certain functions that come naturally to each technology. Referring to the somewhat idealized view of a network in Figure 1.2 based on the assumption of a purely optical physical layer, there is a clean separation between optical/photonic technology, on the one hand, and electronic technology, on the other. The NASs represent the optoelectronic interface (denoted by the boundary labeled e/o) between the electronic domain (the equipment outside the purely optical portion of the network) and the optical domain, sometimes called the *optical ether*. This interface is the point of demarcation between the physical layer and the logical layers. In a typical purely optical physical layer the optical signal paths are as transparent as a piece of glass.

The stations provide the basic functions of getting the light into the fibers (with lasers) and getting it out (with photodetectors). When the signals are in optical form, photonic technology is well suited to certain simple signal-routing and switching functions within the nodes. With static photonic devices, it is fairly easy to perform functions such as optical power combining, splitting, filtering, and wavelength multiplexing, demultiplexing, and routing. By adding suitable control, the static devices can be controlled dynamically (switched) at slow to fast speeds (milliseconds in the case of mechanical or thermal control and microseconds or nanoseconds in the case of electronic control).

The enormous usable bandwidth of a single fiber (tens of terahertz) is at the same time a great asset and a great challenge. It is technologically impossible to exploit all of that bandwidth using a single high-capacity channel. Thus, to make efficient use of the fiber it is essential to *channelize* its bandwidth. This is most easily accomplished by superimposing many concurrent signals on a single fiber, each on a different wavelength; that is, by using WDM. Thus this book focuses on *multiwavelength* or WDM network architectures. The relative ease of signal manipulation in the wavelength (or optical frequency) domain, as opposed to the time domain, suggests that current optical technology is particularly suited to multiwavelength techniques. In WDM networks each optical transmitter (receiver) is tuned to transmit (receive) on a specific wavelength, and many signals operating on distinct wavelengths share each fiber – possibly more than 100 in dense WDM (DWDM) transmission systems.

It should be observed that all photonic routing and switching functions within the optical domain in these networks are *linear* operations. Thus, at the optical level the network typically consists of only linear devices, either fixed or controllable. It is the property of linearity that makes multiwavelength networking simple and cost-effective. To distinguish these linear networks from other types of optical networks, we refer to them frequently as *transparent* optical networks. Typical *nonlinear* operations performed in networks include signal detection, regeneration, reading, and modifying the information in the signal, buffering, and logic functions (e.g., packet routing based on header information). Although many nonlinear functions can be performed in the optical domain with present-day technology the current state of the art for these nonlinear devices is not nearly as advanced as it is for linear components. For these reasons, we frequently use the terms *transparent optical network* and *purely optical network* interchangeably in this book.[5] Nonlinearities make the signal path *opaque* rather than transparent. Some of the advantages of keeping nonlinear operations out of the signal path are (1) the

[5] As the state of the art progresses photonic technology is becoming a viable alternative for many nonlinear signal processing operations, so that the linkage between "transparent" (i.e., linear) and "purely optical" is becoming tenuous.