Chapter 1

Optical Networking

1.1 Evolution of Optical Network Architectures

The size and complexity of telecommunications networks and the speed of information exchange have increased at an unprecedented rate over the last decades. We live in a new information era, where most people are currently using a number of devices with advanced multimedia applications to obtain and exchange information. The current trends in multimedia communications include voice, video, data and images. These trends are creating a demand for flexible networks with extremely high capacities that can accommodate the expected vast growth in the network traffic volume.

In today's integrated networks, a single communications medium should be able to handle individual sessions with a variety of characteristics, operating in the range of a few megabits to tens of gigabits per second. This will enable it to handle such applications as large-volume data or image transfers (e.g., supercomputer interconnections, supercomputer visualization and high-resolution uncompressed medical images) that have very large bandwidth requirements, as well as applications such as voice or video which require much smaller bandwidth.

The enormous potential of optical fiber to satisfy the demand for these networks has been well established over the last few decades. Optical fiber is highly reliable (Bit Error Rate (BER) in commercially deployed systems is less than 10^{-12}), it can accommodate longer repeater spacings and it has unlimited growth potential. Single mode fiber offers a transmission medium with Tbps bandwidth (enough capacity to deliver a channel of 100 Mbps to hundreds of thousands of users) combined with low loss and low BER. Traditional network architectures, however, that used electrical switches and the optical fiber as a simple substitute for copper wire or other communications media, were limited by an electronic speed bottleneck and could not have been used in telecommunications networks with a growing demand for Gbps applications.

As the next step in the evolution of transport networks, Wavelength Division Multiplexed (WDM) optical networks were proposed [58] which provided concurrency by multiplexing a number of wavelengths for simultaneous transmission within the same medium. Rapid advances in optical fiber communications technology and devices, in terms of performance, reliability and cost over the last few years, were the catalyst in enabling the deployment of optically routed WDM networks as the next generation, high-capacity nationwide broadband networks [289]. This approach can then provide each user with a manageable portion of the enormous aggregate bandwidth.

Path Routing in Mesh Optical Networks Eric Bouillet, Georgios Ellinas,
Jean-François Labourdette, Ramu Ramamurthy

The current information explosion is indeed in large part due to the radical progress in optical communications technology over the last few decades. Dense Wavelength Division Multiplexed (DWDM) mesh networks that route optical connections using optical cross-connects (OXCs) have been proposed as the means to implement the next generation optical networks [308]. Following a wave of timely technological breakthroughs, optical network equipment vendors have developed a variety of optical switching systems capable of exchanging and redirecting several terabits of information per second. The dimensions of the switches range from a few tens to several hundred ports with each single port capable of carrying millions of voice calls, or thousands of video streams. The emergence of new optical technologies is driving down the overall network cost per units of bandwidth, and the trend is accompanied by an explosion of new data service types with various bandwidth characteristics and prescribed Quality of Service (QoS). Optical network architectures as we envision them now not only provide transmission capacities to higher transport levels, such as inter-router connectivity in an Internet Protocol (IP)-centric infrastructure, but also provide the intelligence required for efficient routing and fast failure recovery in core networks [36, 212, 304]. This is possible due to the emergence of optical network elements that have the intelligence required to efficiently manage such networks. Figure 1.1 illustrates the optical network hierarchy, with a core optical network incorporating mesh topologies with optical cross-connects interconnecting WDM metro networks incorporating reconfigurable optical add drop multiplexers (ROADMs), which in turn interconnect various access networks.

The reader should note that there is another alternative architecture in which the IP routers are directly connected to WDM systems (i.e., there is no optical switching). Historically, there have existed two schools of thought concerning the evolution of the core network architecture. The first argued that all of the intelligence should reside within the IP layer, and the optical layer should just be used for transport, while the second argued to move away from a network where all the processing is done at the IP layer to a network where the intelligence is shared between the IP routers and the optical cross-connects. Figure 1.2 illustrates the two different network architecture scenarios. We believe the latter vision is more appropriate for core networks for a variety of reasons. For example, for

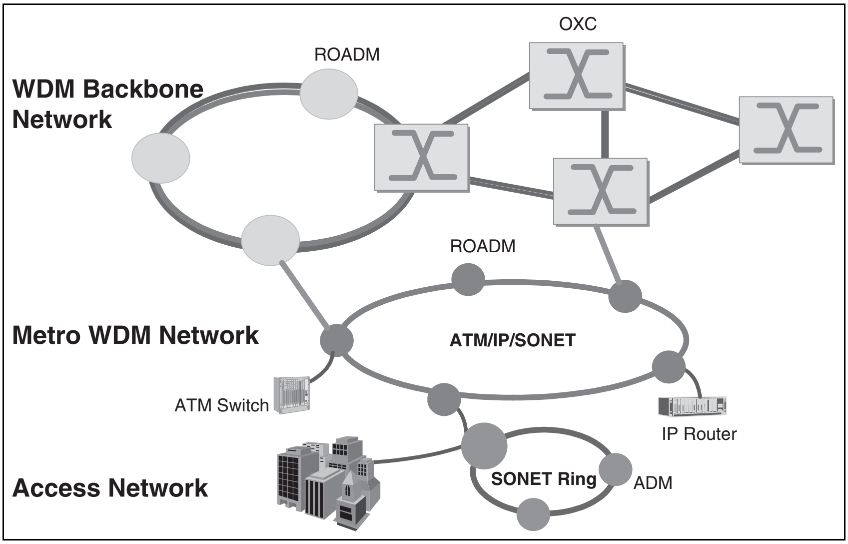

Figure 1.1: Optical network hierarchy.

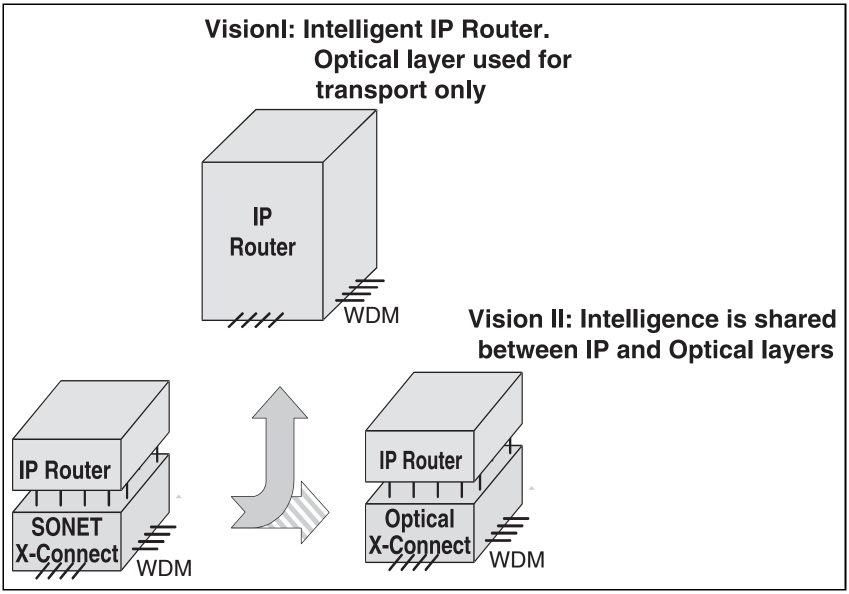

Figure 1.2: Different visions for optical network evolution.

the case of failure recovery that is a main focus of this book, protection/restoration in the optical layer is typically faster, more robust and simpler to plan and upgrade compared to IP/Multiprotocol Label Switching (MPLS)-based recovery. Thus, even though the first vision may indeed be a viable architecture it is not discussed in the remainder of this book, which deals exclusively with the case where intelligent optical switching is present in the network.

Optical networks enable a variety of wavelength services (such as wavelength-on-demand, wavelength brokering, and optical virtual private networks) that open up new opportunities for service providers and their customers alike. In addition to new services, high-speed connections at 2.5 Gbps rates and above are required for the optical core to support trunking between edge service platforms. The dominant traffic carried in today's network is evolving from legacy voice and leased line services to data services, predominantly IP services. Time Division Multiplexing (TDM) aggregation switches, optimized for legacy voice services and leased line services, and acting as edge devices, groom signals at lower bit rates (e.g., 1.728 Mbps and 51.840 Mbps), and feed them into the core, typically, at rates of 2.5 Gbps and 10 Gbps. Equipment operating at 2.5 and 10 Gbps are currently commercially available, and there has been considerable work (both experimentally and in some cases some initial commercial products were developed) on 40 Gbps data rate transmission and switching[1] [47, 69, 70, 150].

Figure 1.3 illustrates the four different node architectures that can comprise a reconfigurable core optical network. The first architecture shows a fixed patch panel. Fixed patch panels located between WDM systems with transponders are currently being replaced by opaque switching nodes (with electrical switch fabrics), as shown in the architecture of Figure 1.3(b), due to their complete lack of flexibility. This is an *opaque network architecture*, as the optical signal now undergoes Optical−Electrical−Optical (OEO) conversion at the switch [30]. The third architecture shows a transparent switch between WDM systems with transponders that would be complemented by an OEO switch for drop traffic. This is again an opaque network architecture, as the optical signal

---

[1] Commercial deployments are also taking place.

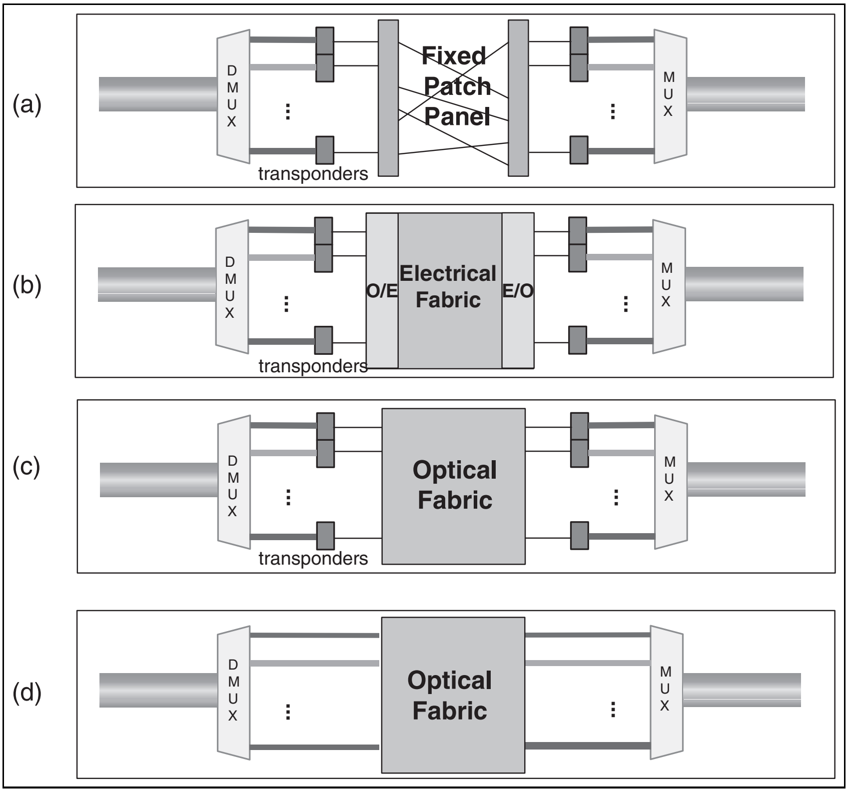

Figure 1.3: Node architectures for a core optical network. (a) Opaque network with fixed patch panel, (b) Opaque network with opaque switch, (c) Opaque network with a transparent switch and (d) Transparent network with a transparent switch. (After [108], Figure 1. Reproduced by permission of © 2003 The International Engineering Consortium.)

undergoes OEO conversion at the WDM transponders. The fourth architecture shows a completely *transparent* network topology, consisting of transparent optical switches and WDM systems that contain no transponders. The transparent switch would be complemented as in Figure 1.3(c) by an OEO switch for drop traffic. In this architecture, the signal stays in the optical domain until it exits the network. Details on the design of each architecture are presented in the sections that follow. There has been extensive research work on the limits of optical transparency, comparisons between transparent and opaque networks, and the benefits and drawbacks of each technique. The reader is referred to [203, 268, 299] for some of the work that has been performed in that area.

1.1.1 Transparent Networks

The transparent node architecture shown in Figure 1.3(d) and elaborated on in Figure 1.4 is a seemingly attractive vision. A signal (wavelength) passing through an office does not undergo opto-electronic (O/E) conversion. Similarly, a client Network Element (NE), such as a router, interfaces with the switch using long-haul optics to interface with the WDM equipment without any O/E conversion. Since a signal from a client NE connected via a specific wavelength must remain on the same wavelength when there is no wavelength conversion, only a small-size switch fabric is needed to

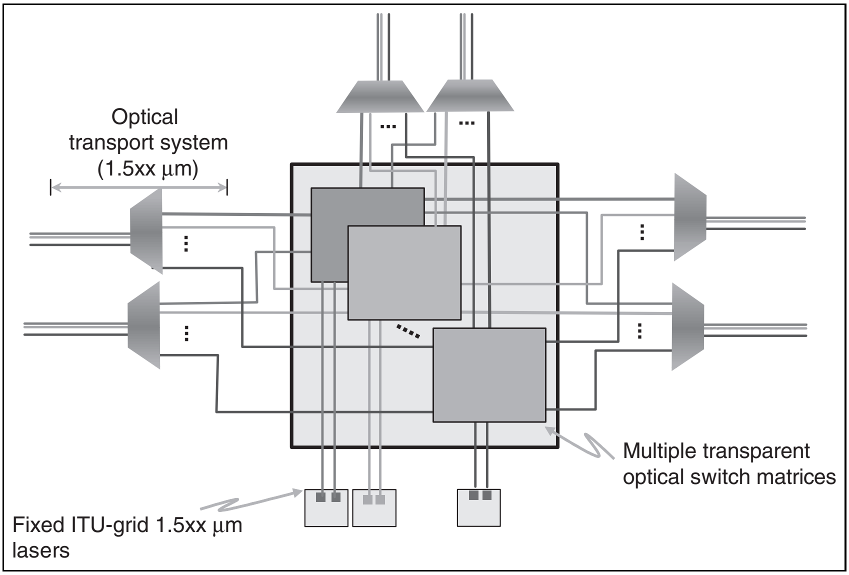

Figure 1.4: Transparent switch architecture in a transparent network. (From [109], Figure 1. Reproduced by permission of © 2004 The Institute of Electrical and Electronics Engineers.)

interconnect the WDMs and NEs in a node. This architecture also implies end-to-end bit rate and data format transparency. Note that another architecture of a transparent switch in a transparent network may include a single large fabric instead of multiple switch matrices of small port counts. However, if one is to provide flexibility, such an architecture design would require the use of tunable lasers at the clients and wavelength conversion.

A transparent network architecture may provide significant footprint and power savings and on the surface suggests cost savings. However, while the transparent network architecture may be a viable option for small-scale networks with pre-determined routes and limited numbers of nodes, it may not be a practical solution for a core mesh optical network for the following reasons:

- This network does not allow wavelength conversion, thus essentially creating a network of *n* (*n* being the number of WDM channels) disjoint layers. Inflexible usage of wavelengths in this network would lead to increased bandwidth and network operational cost, thus negating all savings that may result from the elimination of O/E conversion. In addition, for this technology to be effective and in order to build a flexible network for unrestricted routing and redundant capacity sharing, an all-optical 3R-regeneration function must be available. Such a technology that can be harnessed in a commercial product does not currently exist [220].

- In the absence of wavelength conversion, only client-based dedicated backup path protection (DBPP) can be easily provided [107, 187]. The wavelength continuity constraint on backup paths makes resource sharing very difficult in transparent networks and consequently no shared backup path protection (SBPP)[2] can be easily offered. This in turn means that the capacity requirement for protected services is significantly higher (80–100%) for transparent compared to opaque networks [68].

---

[2] The concepts of sharing, DBPP and SBPP, will be explained in detail in Chapter 2. SBPP is sometimes alternatively termed *backup multiplexed* path protection.

- Physical impairments such as chromatic dispersion, polarization mode dispersion (PMD), fiber nonlinearities, polarization-dependent degradations, WDM filter pass-band narrowing, component crosstalk, amplifier noise, etc. accumulate over the physical path of the signal due to the absence of O/E conversion. The accumulation of these impairments requires engineering of end-to-end systems in fixed configurations [197, 245, 246, 247]. Thus, it may not be possible to build a large network with an acceptable degree of flexibility.

- The design of high-capacity DWDM systems is based on intricate proprietary techniques, eluding any hope of interoperability among multiple vendors in the foreseeable future. Since a signal is launched at the client NE through the all-optical switch directly into the WDM system without O/E conversion, and it is not possible to develop a standard for the interface for a high capacity WDM system, the operators will not have the flexibility to select the client NE vendor and the WDM vendor independently. Consequently, transparent networks by necessity are single vendor (including the client network elements) solutions.

- Finally, in addition to all the limitations discussed above, the challenge of performance-engineering continental-scale transparent reconfigurable wavelength-routed networks remains severe and, in networks that push limits, remains unsolved despite some attempts at formalizing the routing problem [290].

It is apparent that a number of key carrier requirements – dynamic configuration, wavelength conversion, multi-vendor interoperability of transport equipment (WDM), low network-level cost – would be very hard to meet in a transparent network architecture. Therefore, an opaque network solution will remain for now the only practical and cost-effective way of building a dynamic, scalable, and manageable core backbone network. A description of the opaque network architecture is offered in the section that follows.

1.1.2 Opaque Networks

Even though the opaque network solution may be more expensive in terms of equipment costs when the core network capacity increases significantly, the opaque network offers the following key ingredients for a large-scale manageable network:

- No cascading of physical impairments. This eliminates the need to engineer end-to-end systems (only span engineering is required) and allows full flexibility in signal routing.

- Multi-vendor interoperability using standard intra-office interfaces (see Figure 1.5).

- Wavelength conversion enabled. Network capacity can be utilized for service without any restrictions and additional significant cost savings can be offered by *sharing* redundant capacity in a mesh architecture (see Figure 1.5).

- Use of an all-optical switch fabric without any compromise of the control and management functions. Overhead visibility (available through an OEO function that could be complementing the all-optical switch) provides support for the management and control functions that are taken for granted in today's networks.

- The network size and the length of the lightpaths can be large, since regeneration and retiming are present along the physical path of the signal.

- Link-by-link network evolution. Permits link-by-link incorporation of new technology, as the network is partitioned into point-to-point optical links (see Figure 1.6).

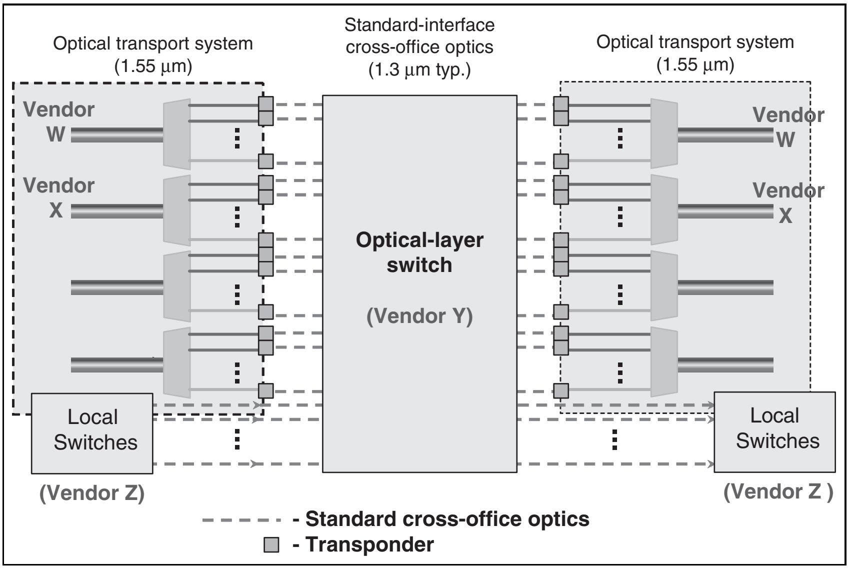

Figure 1.5: Multi-vendor interoperability and wavelength translation as a by-product of an opaque network architecture. (From [109], Figure 2. Reproduced by permission of © 2004 The Institute of Electrical and Electronics Engineers.)

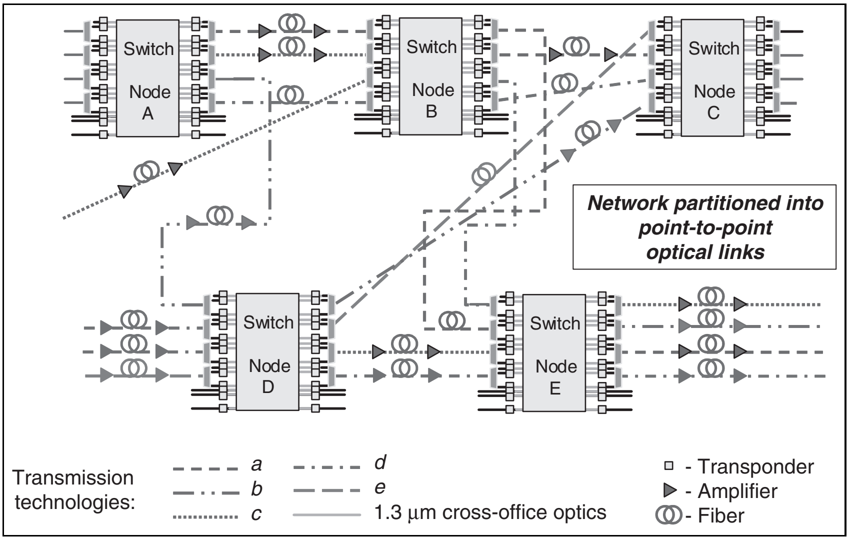

Figure 1.6: Link-by-link network evolution.

Having reasoned that transparent core mesh network architectures are likely to remain unrealistic for quite some time, we now turn our attention to opaque network architectures in which WDM systems utilize transponders. Today's architectures contain, in the most part, opaque switches (with an electronic switch fabric) in an opaque network (with transponders present in the WDM system). This architecture is shown in Figure 1.7. The interfaces to the fabric are opaque interfaces, which means that transceivers are present at all interfaces to the switch, and these transceivers provide an OE (input) and EO (output) conversion of the signal. The presence of the transceivers at the edges of the switch fabric enables the switch to access the Synchronous Optical NETwork/Synchronous Digital Hierarchy (SONET/SDH) overhead bytes for control and management functions. The opaque transceivers provide support for fault detection and isolation, performance monitoring, connection verification, neighbor/topology discovery and signaling, as well as support for implementing the network routing and recovery protocols.

Another design of the opaque switch architecture may be one using Photonic Integrated Circuits (PICs) [222] which can include large numbers of lasers, modulators, and optical multiplexers or optical demultiplexers and photodiodes integrated on the same monolithic Indium Phosphate (InP) chip. This approach allows for low-cost opaque architectures by enabling low-cost OEO conversion of the signal on a semiconductor chip, and is a leap in technology.

The opaque switch approach was, however, faced with a number of challenges when confronted with the traffic growth projections from just a few years ago: it would eventually reach scaling limitations in signal bit rate, switch matrix port count, and network element cost. These were the key motivations behind the attempt to develop large port-count transparent switches to be used in opaque network architectures. For high port-count fabrics, analog gimbal-mirror MEMS (Micro-Electro-Mechanical Systems)-based switches (3D switches) offer the most viable approach [121, 310]. It is important to point out that the opaque switches could still remain in the network architecture in order

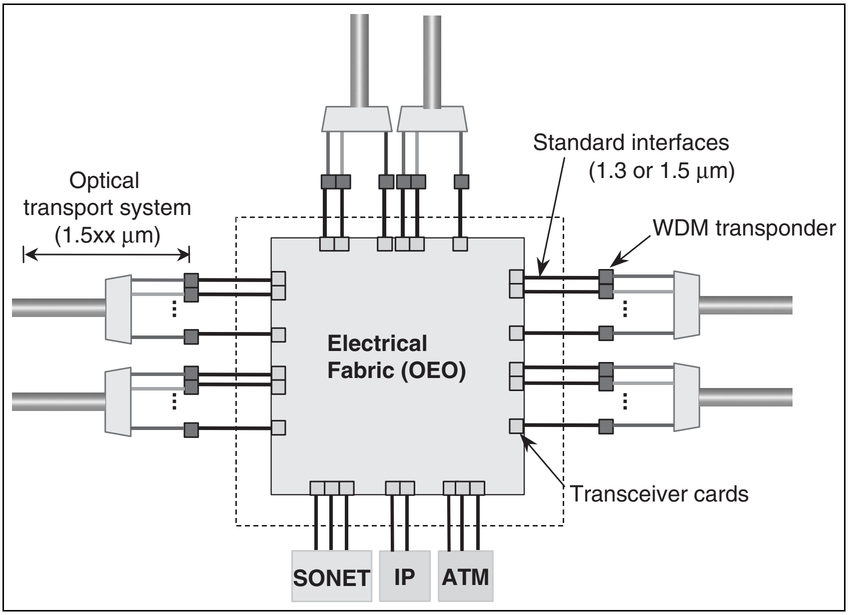

Figure 1.7: Opaque switch architecture. (From [109], Figure 3. Reproduced by permission of © 2004 The Institute of Electrical and Electronics Engineers.)

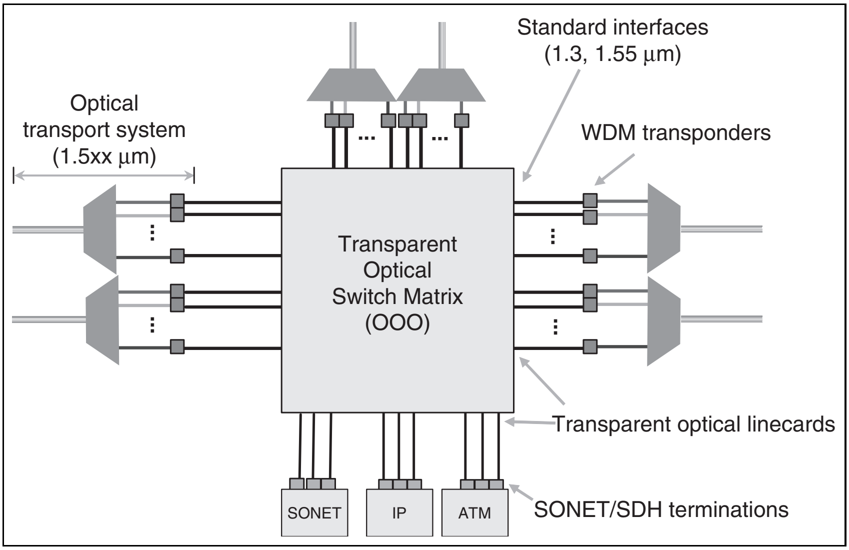

Figure 1.8: Transparent switch architecture. (From [109], Figure 3. Reproduced by permission of © 2004 The Institute of Electrical and Electronics Engineers.)

to provide some key network functions, namely grooming and multiplexing, Service Level Agreement (SLA) verification, and control and management.

Figure 1.8 shows a transparent (OOO) switch architecture. In this architecture, optical signals pass through e.g., a MEMS-based switch fabric, in contrast to the OEO architectures where switching is accomplished using an electrical switch fabric. This switch architecture has transparent interface cards, i.e. no (OEO) transceiver (TR) cards are located at the network ports of the switch fabric that convert the optical into an electrical signal. The switch shown in Figure 1.8 also has no opaque transceiver cards on its add/drop ports. Therefore, it has no direct access to the overhead bytes for control and signaling. The optical switch fabric is bit-rate independent and it accommodates any data rates available (e.g., 2.5, 10 and 40 Gbps).

The advantage of such a switch architecture is that for an $N \times N$ all-optical (OOO) architecture there are N interfaces/ports to the all-optical switch regardless of the type of interfaces. No data-rate-specific interface cards are used, so no replacement is needed when the switch is operating at higher data rates, provided that the optical power budget is sufficient for that rate. This is in contrast to the OEO systems where the number of ports depends on the type of the port. For example, in an OEO system one 10 Gbps interface card will replace four 2.5 Gbps interface cards. From the interface card perspective all ports in the OOO architecture will look one and the same (the same port cards are used for different signal rates and formats). The add/drop-side ports are connected to an OEO switch (or any other client – such as IP/MPLS router or Asynchronous Transfer Mode (ATM) switch) that provides SONET/SDH termination through its opaque ports.

The promise of optical switching was that, unlike integrated electronic switches, an optical switch fabric's complexity is a flat function, independent of the bit rate of the signals it handles (Figure 1.9). Moreover, in the long run, it was projected that few components would be as small, cheap, and low in power consumption as a silicon micro-mirror in the case of MEMS-based switch fabrics. Transparent switches could thus be expected to be cheaper in terms of the switching fabric and interface card cost compared to opaque switches. This would have resulted in significant cost reduction to network

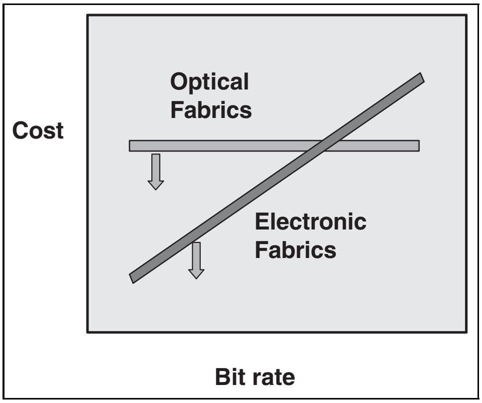

Figure 1.9: Advantages of optical fabrics. (After [108], Figure 5. Reproduced by permission of © 2003 The International Engineering Consortium.)

operators because a large amount of the traffic that passes through an office will be able to bypass the OEO switch (typically approximately 75% through-to-total ratio).

Transparent switches essentially would have helped to relieve the demand for OEO switch ports and reduce the cost of transporting lightpaths. This is accomplished by having all lightpaths pass through the OOO switches (glass through), thus bypassing the OEO switches. Note that this can be a significant portion of the network traffic. For example, if 40 Gbps data rates were used, every time a lightpath passes through an OOO switch, 32 equivalent 2.5 Gbps ports of an OEO switch would be saved (two 40 Gbps ports that correspond to 32 equivalent 2.5 Gbps ports).

Since the OOO switch fabric is bit-rate and data-format independent, the switch matrix can scale more easily than electrical switch fabrics. For these reasons, as bit rates rise, it was thought that optical switch fabrics would eventually prevail. Note, however, that this evolution would only have happened on timescales that were gated by the ability of vendors to meet carrier reliability and operational requirements with all-optical technologies such as lightwave micro-machine (for MEMS-based switch fabric) technology [121]. Even though in early stages of 2.5 Gbps and 10 Gbps development the crossover point shown in Figure 1.9 appeared to be at the 2.5 Gbps and then the 10 Gbps rates, under today's more realistic traffic growth scenario, and given the lack of deployment of 40 Gbps WDM systems and the continued decline in price of OEO components, the crossover point has shifted to the even higher bit rates. Therefore, the need for and the promise of transparent switches appear to have moved beyond the foreseeable future. Provided that the traffic grows and the bit rates increase substantially, there may emerge in the future a potential need for an additional network layer utilizing transparent optical switches. In that case, the main challenge to architectures that use transparent switches will then be to provide the control and management functionalities that are readily available when we have access to the electrical signal and consequently to the SONET/SDH overhead bytes.

Even though the use, in the core, of transparent switches that are cost-effective at very high bit rates is not currently justified, there still exist some niche applications in today's networks that could use a small number of transparent switches. Transparency is mainly limited in metropolitan area networks, utilizing ROADMs, and some ultra long-haul applications in the core, utilizing a small number of wavelength-selective cross-connects/OADMs on high-capacity routes. When OADMs are utilized, selected wavelengths can be added or dropped at a node while the rest of the wavelengths pass through without regeneration [29, 30]. ROADMs further enable any user to access any channel.

ROADMs can be utilized in metropolitan area networks at central offices and customer locations in much the same way that the SONET introduction created a need for large numbers of SONET Add Drop Multiplexers (ADMs). They provide network flexibility and can be used to manage continuous changing traffic patterns and customer service requirements.

Wavelength-selective cross-connects (WSXCs) can also be used in ultra long-haul applications in the core network in a completely transparent manner. Even though these network elements allow for end-to-end bit rate and data format transparency, they face a number of challenges. However, these network elements could be utilized in a few, predetermined and non-reconfigurable high-capacity routes to provide end-to-end transparency between fixed end-nodes. Furthermore, we anticipate that opaque switches will always remain for the embedded service base even after the transparent switches are introduced in the network. These opaque switches will provide the grooming and multiplexing functions, as well as some of the necessary control and management functions, and will scale and decrease in cost with rapid progress in electronics.

While completely transparent core mesh networks have not yet materialized on a large scale, even transparent switches in opaque networks still face technological as well as control and management challenges [109]. Even though most of these issues can be addressed via clever innovation as well as standardization efforts, transparent switches complemented by an opaque function will not be ready for deployment in the network until all the control and management challenges are successfully resolved.

1.1.3 Translucent Networks

There is a third network architecture worth mentioning, the *translucent* network architecture, which is based on optical cross-connects that are a mixture of opaque and transparent cross-connects presented in the previous sections [247, 248]. Figure 1.10 shows an example of such a node. This node is composed of two parts: a transparent and an opaque part. A signal entering this node can pass through (transparently), or can be dropped (or added) and go through a regeneration process [218]. Long lightpaths that cause the optical signal to degrade are the candidates that will go through the regeneration process to recover the original signal. Signal regeneration may occur a number of times before the signal reaches its final destination. Translucent networks allow the connections to stay transparent for as long as their signal quality allows, and then go through a regeneration process. These networks thus have some of the advantages of the transparent networks discussed previously while mitigating some of the drawbacks that appear in networks that are completely transparent (such as allowing

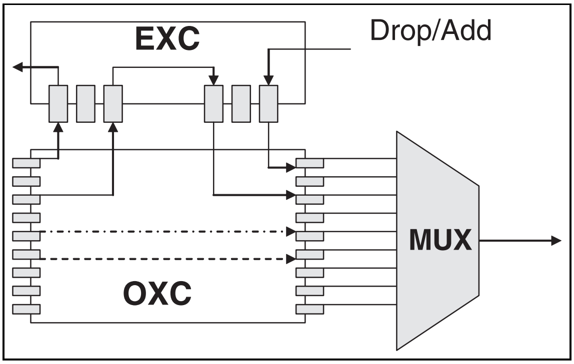

Figure 1.10: Translucent node architecture. (After [218], Figure 1.b. Reproduced by permission of © 2005 DRCN.)

for wavelength conversion and addressing the problem of accumulation of physical impairments on the path).

Potentially not all the nodes in the network will have the regeneration capability. Sparse regeneration can be offered by placing the translucent nodes at strategic locations. Problems such as recognizing and addressing regeneration demands [248], dynamic routing in translucent optical networks [327], and placement of the translucent nodes in order to minimize blocking in these networks [277] have been addressed by the research community.[3]

Studies presented in [217, 218] address the impact of the reach of the WDM systems and the impact of the traffic volume on the cost of these networks. These studies also include cost comparisons between translucent and opaque network architectures. Even though the initial studies presented in [217] showed the translucent architecture to be 50% cheaper than the opaque one, subsequent studies presented in [218] reduced that savings number 10%.[4] The authors in [218] argue that considerable increase in traffic or considerable reduction in the cost of transparent devices will be the factors that can make the translucent network option cost-effective.

1.2 Layered Network Architecture

In this section we review the fundamental parts that constitute a network and its functionality. It goes without saying that many architectures exist or have been suggested and it is not the intention of this book to enumerate them exhaustively (see [39, 155, 243, 270] for further information and useful references on this topic). However, we observe that all the proposed architectures repose on a common denominator. It is this generic model that we present here. The model consists of three superimposed layers. Each layer provides well-defined services to its superjacent layer while concealing implementation details from it.

Figure 1.11 shows a layered architecture with the DWDM network being used as the transport network. The fiber-optic links and optical switching nodes are located in the physical topology. The logical layer represents the view of the network seen by end users, accessing the physical layer through electro-optic interfaces. The access means may be direct (through clear channels) or indirect

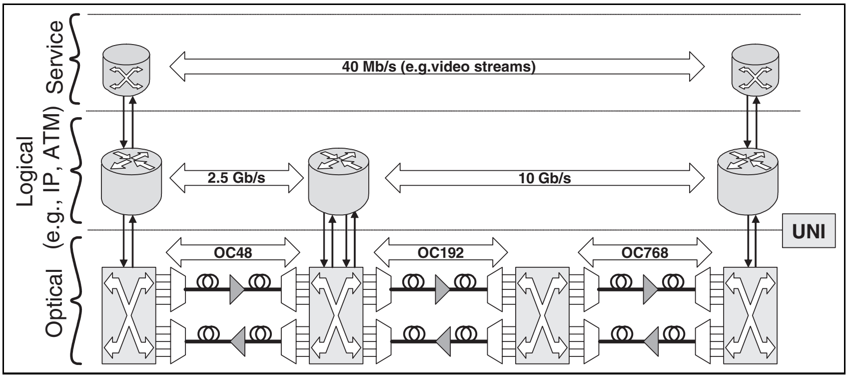

Figure 1.11: Layered network architecture. (From [107], Figure 1. Reproduced by permission of © 2003 The International Society for Optical Engineering.)

---

[3] Apart from the indicative references given here the reader is encouraged to investigate the large body of work that exists on translucent networks, sparse wavelength assignment, etc.

[4] Savings depend in a large part on the architectures of the nodes in the network and on the system reach.

through the intermediary of electronic (e.g., SONET, ATM) switching equipment. The service layer demonstrates the large number of applications that can be provided in such networks. As shown in Figure 1.11, from bottom to top the layers are (1) optical layer, (2) logical (electrical) layer and (3) service/application layer.

1.2.1 Optical Layer

The optical layer offers and manages the capacity required to transport traffic between clients in the logical layer. The optical layer includes wavelength transmission equipment (DWDM), wavelength switching or cross-connect equipment (also called optical switches) handling 2.5 and 10 Gbps wavelengths, and wavelength grooming equipment, handling subrate circuits (in multiples of STS-1) into 2.5 Gbps and/or 10 Gbps wavelengths.

Figure 1.12 depicts an example of a logical network (two IP/MPLS routers) linked to an optical network (four optical switches). Optical switch ports are either: (1) add/drop-ports, interfacing the optical layer to the client's logical layer, or (2) network ports, interconnecting optical switches. Using our graph representation, nodes are optical switches, and links are bundles of bidirectional optical channels between pairs of optical switches. An optical channel is a wavelength that connects the network ports of adjacent optical switches. A link in the logical layer is realized by way of optical channels in tandem forming a lightpath (circuit) between the end-nodes of that link.

The optical layer faces the same challenges, and conceptually even borrows solutions from the logical layer. For instance, it relies on Generalized MPLS (GMPLS) [26, 182, 210] also formerly known as MPLambdaS (an extension of MPLS) to encompass all types of architectures, including wavelength-oriented traffic engineering and management. It also relies on Neighbor Discovery Protocol (NDP)[5]/Link Management Protocol (LMP) [193] and Open Shortest Path First (OSPF) protocol [219] together with Link State Advertisements (LSAs) exchanges, to create and publicize the network's topological views. Differences that set apart the optical layer from its logical counterpart

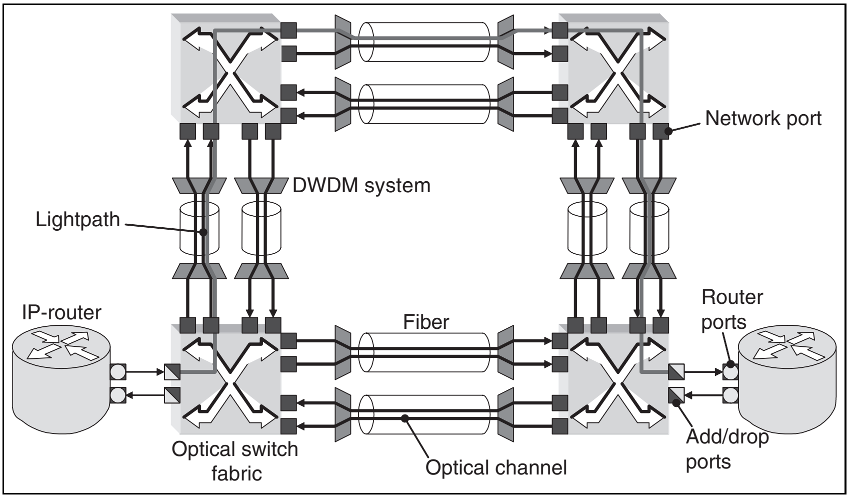

Figure 1.12: The optical layer. (From [107], Figure 2. Reproduced by permission of © 2003 The International Society for Optical Engineering.)

---

[5] The Hello Protocol is also used.

are among others: (1) routing in the optical layer is exclusively circuit oriented, (2) circuit set-up and tear-down is done at a much slower timescale and (3) the bandwidth granularity of the logical layer is much lower than the granularity of the optical layer.

In the *overlay approach* assumed throughout this book the layers work individually, with the client logical layer leasing resources from the optical layer. The User Network Interface (UNI) harmonizes communication of control messages between the two domains [12]. In addition, since an optical carrier will normally acquire network components from several vendors, a suite of protocols is being developed in the Internet Engineering Task Force (IETF) to allow for the seamless interaction between the various network components. As part of that suite, the Link Management Protocol, for example, is used to maintain control channel connectivity, verify component link connectivity and isolate link, fiber or channel failures within the network [12].

1.2.2 Logical Layer

Also known as electrical or digital layer, the logical layer aggregates services into large *transmission pipes* and assures their proper routing from Point of Presence (PoP) to PoP with prescribed QoS [172, 173]. Using a graph representation, a logical node corresponds, for example, to an IP/MPLS core router, an ATM backbone switch or a digital cross-connect (DCS), and a logical link connects the ports of two adjacent nodes. The logical layer may consist of several interconnected subnetworks, either for scalability reasons, as it is easier to manage several smaller networks than a large network (hierarchical decomposition), or because the subnetworks belong to several independent carriers or employ different technologies (e.g., IP/MPLS versus ATM). In either case, boundaries and proper network interfaces within the logical layer delimit the subnetworks and their respective domains of operation.

The logical layer fulfils several roles: (1) it maintains a consistent topological view of its layer, (2) it manages the address space, (3) it routes streams on request, and (4) it polices the traffic to ensure a fair share of capacity among data streams and to guarantee each individual's QoS. The first part, also called topology discovery, can be achieved, for example, by way of an NDP in conjunction with the OSPF protocol [219]. NDP operates in a distributed manner through in-band signaling to construct local port-to-port connectivity databases at each node. OSPF completes the topology discovery by assembling and globally disseminating pieces of information collected by NDP, plus additional information such as link states, to the logical plane [38]. Logical nodes have only a few tens of ports, and with the exception of very small networks, a full connectivity featuring one link between every pair of node is not probable. Instead, services may have to be routed in the logical layer through one or more transit nodes to the desired destination using, for example, Constrained-based Routing Label Distribution Protocol/Resource Reservation Protocol (CR-LDP/RSVP) explicit routing and bandwidth reservation protocols [27]. The computation of a logical path must satisfy a set of constraints, such as round-trip delays and spare bandwidth, defined in the service layer in accordance to prescribed QoS [173]. Note that the failure of a logical link or logical node can be detected, for example, by NDP, and advertised by OSPF. That is, the layer has the primitives to detect a failure and resume interrupted services.

1.2.3 Service/Application Layer

In the service layer, clients such as edge or service routers or Multi-Service Provisioning Platforms (MSPPs) located in a provider's POP represent users and the data communication among them. Using a graph representation, a node corresponds to a client who emits and receives data, and a link represents a service or a two-way data stream between clients. Link attributes in this layer correspond to minimum QoS requirements, which transpose into bit rates, jitter, and bit-propagation or round-trip delay constraints. SLAs for instance are negotiated and crafted in this layer.

Section 1.3 that follows applies to the optical layer as defined in this section. It deals with the wavelength switching and wavelength grooming functions, their distribution across equipment and layers, their interplay, and their impact in terms of transport efficiency and transport failure recovery performance.

1.3 Multi-Tier Optical Layer

An opaque core optical switch, as described in Section 1.1.2, converts optical signals into the electrical domain at the ingress port, switches the electrical signals through an electrical switch matrix, and then converts signals into the optical domain at the egress port. The switch fabric of the OEO switch is typically strictly nonblocking, allowing any interconnection pattern among the ports of the switch (e.g., a Clos switch fabric [76]). This is not always the case though, as some of the switch fabrics for an OEO switch can also be rearrangeably nonblocking (e.g., a Benes switch fabric [35]) or wide-sense nonblocking. The reader is referred to [154] for a comprehensive review of the switch interconnection fabrics.

An OEO switch performs aggregation and grooming functions. As an aggregation device, the OEO switch takes multiple bit-streams and maps them onto wavelengths. As a grooming device, the OEO switch can interchange time slots between different bit-streams. If the OEO switch can switch at the STS-N (Synchronous Transport Signal[6] level N) granularity, we call it an STS-N switch, and if it can switch at the STS-M granularity, we call it an STS-M switch. For example, an STS-M switch can aggregate STS-N traffic onto OC-N wavelengths ($M < N$, N a multiple of M), and can switch STS-M frames between OC-N wavelengths. However, the STS-M switch cannot switch STS-K ($K < M$, K a multiple of M) frames between different STS-N frames.

Figure 1.13 illustrates the aggregation and switching functions of an OEO switch. Switch A aggregates frames from four Input/Output (I/O) interfaces and multiplexes them onto a wavelength

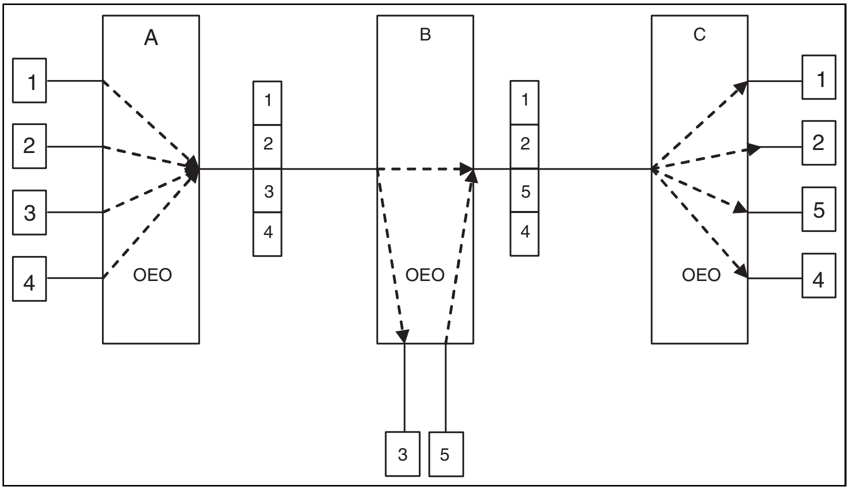

Figure 1.13: Grooming (aggregation and switching) performed by an OEO switch.

---

[6] STS-1 (Synchronous Transport Signal level 1) is the basic building block of SONET. STS-N signals are created by concatenating multiple STS-1 signals or via a combination of other concatenated signals STS-M with $M < N$ and N a multiple of M. Before transmission, the STS-N signal is converted to an OC-N (Optical Carrier level N) signal. Refer to [4, 90] for additional details on the SONET protocol.

channel. Switch B drops frame 3, and adds frame 4 and multiplexes them onto a wavelength channel to switch C. Switch C demultiplexes the frames onto the I/O interfaces. The granularity of the switches must be at least the frame-rate.

Historically, as the network has evolved so has the granularity of grooming. In the early days of the transport networks, the grooming granularity at the core was 64 kbps (DS0 – Digital Signal 0 is the lowest level of the Plesiochronous Digital Hierarchy (PDH) system [90]). Over time as networks grew and traffic volume increased, the grooming granularity also increased to 1.5 Mbps (DS1) and then to $\approx$ 45 Mbps (DS3/STS-1) to keep the networks scalable and manageable and to improve network performance and cost. As Digital Cross-Connect Systems (DCSs) have been introduced into digital core transport networks over the past 25 years, the rate of the core transmission speed has traditionally been about 20 to 40 times the rate of the *core switching* (cross-connect) rate. DS0 (64 kbps) signals were switched when core transport systems were on the order of DS1 signals (1.5 Mbps); similarly, DS1 signals were switched within DS3 ($\approx$45 Mbps) signals. Most recently, DS3 signals were switched with DCSs when the core transmission speeds were on the order of 1.5 Gbps to 3 Gbps ($\approx$STS-48) [124].

The right granularity for grooming at the core has been a question that has been continuously investigated by engineers. There are advantages and drawbacks for either of the following approaches: switching with fine (e.g., STS-1) granularity (e.g., offering enhanced flexibility to manage all services in the network) or switching with coarse (e.g., STS-N) granularity (e.g., offering increased manageability and scalability and keeping the network complexity under control). Clearly, the drivers behind such a decision will be the expected growth in traffic volume, the traffic composition (emerging applications), the need for scalable and manageable networks, the performance of the network in terms of service provisioning and recovery, the advances in enabling transport network technology, and finally the total network cost [124]. These are factors that need to be considered very carefully before a decision on the grooming granularity is taken.

Apart from the granularity of the network switching elements, another issue of crucial importance is the architecture of the core mesh optical networks in terms of layering in the optical domain. There are two possible architectures for the backbone network: a *flat* (one-tier) architecture or a *layered* or *hierarchical* (multi-tier) architecture.[7] Historically, large networks have always been organized in multi-level hierarchies. It has been a network provider's dream to accommodate all services at all rates with a single box that is scalable, manageable, and low-cost. However, practical considerations such as hardware and software scalability and manageability have led mostly to hierarchical network architectures, taking advantage of the optimization of each layer independently. *All-purpose boxes* may be well suited for enterprise and some metro applications, but typically not for core applications that require specialized, carrier grade products. In a hierarchical architecture, scalability and manageability are achieved by multiplexing traffic flows into larger streams as they traverse from the edge to the core of the network and demultiplexing them as they traverse from the core to the edge. In other words, traffic flows are groomed at a coarser granularity at the network core than at the edges of the network [194, 216, 334].

The sections that follow will define each of these network architectures and will analyze them, trying to identify under what conditions each of them should be used.

1.3.1 One-Tier Network Architecture

The discussion in this section (and in the next section on two-tier network architectures) uses STS-1 and STS-48 as the explicit notations of granularity in order to more easily explain the grooming and multi-tier concepts to the reader. The general argument could be applied for networks with any STS-M and STS-K switches (e.g., $M < K$).

Figure 1.14 shows a sketch of the one-tier network architecture. Ubiquitous STS-1 switches in different offices are connected through physical links. DWDM systems carry the physical connections

---

[7] This layering applies to the optical layer described in Section 1.2.1.

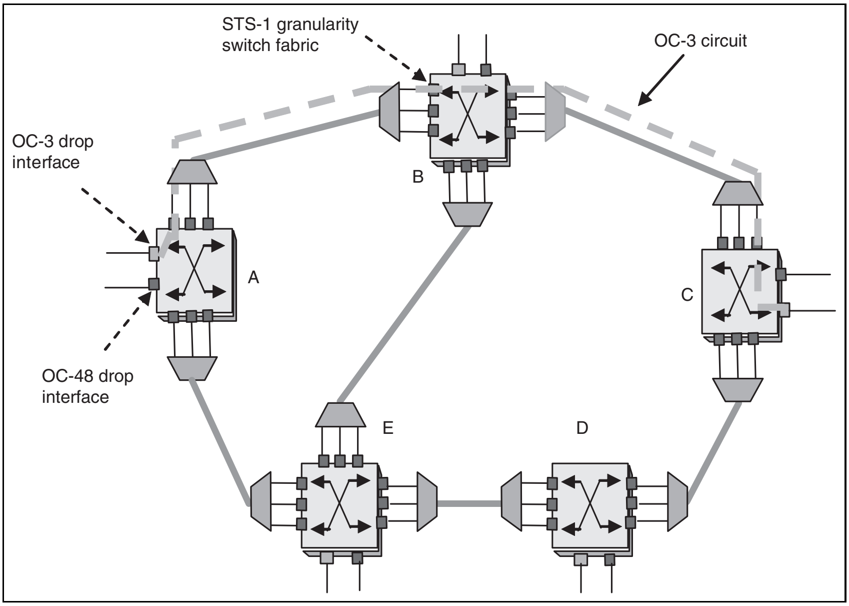

Figure 1.14: One-tier network architecture. (From [189], Figure 3. Reproduced by permission of © 2003 The Optical Society of America.)

from one node to another. In the one-tier (flat) architecture, the core optical switch can switch at the STS-1 granularity. The term *one-tier* comes from the fact that for the STS-1 traffic, the network looks *flat*, i.e., STS-1 traffic in principle is switched end-to-end on the shortest path across the network. However, it can be argued that for rates below STS-1 (e.g., DS3) we need a second tier of (subrate) switches. Nevertheless, we will call it a one-tier architecture, since our focus is the core optical network where traffic rates at or above STS-1 dominate. The STS-1 switch handles all the STS-N services (whose rates are multiples of the STS-1 rate) from the client equipment (such as IP routers, ATM switches, Frame Relay (FR) switches, MSPPs). The STS-1 switch also terminates wavelengths (e.g., OC-48/OC-192) from the DWDM equipment as illustrated in Figure 1.14. In general, an STS-1 switch can switch STS-1 frames from the ingress bit-streams onto egress bit-streams, and allows wavelengths to be managed in increments of STS-1.

An STS-1 switch fabric may be implemented using both space-division and time-division multiplexing. The STS-1 granularity of the switch fabric allows the switch to terminate STS-N traffic with interfaces that can multiplex/demultiplex STS-N traffic onto the switch fabric. A consequence of the STS-1 granularity of the switch is that multiplexing several STS-1 streams from different inputs to a single output could involve the configuration of a large number of cross-point configurations (routing across the switch fabric, and setting appropriate time slot switches). The availability of dense time and space division cross-point chips allows such a switch fabric to scale theoretically to 2000 × 2000 I/O ports.

Figure 1.14 illustrates an example of provisioning an OC-3 circuit from backbone node A to backbone node C. A route is first found across the optical network, with available bandwidth on each link of the route. Then each switch can be configured to set up the OC-3 circuit.

Note that an OC-N circuit between regional PoPs that has to traverse the core optical network is provisioned with multiple legs. The first leg is an OC-N circuit that traverses the first regional network onto the core optical switch at the backbone node. The second leg is an OC-N circuit between

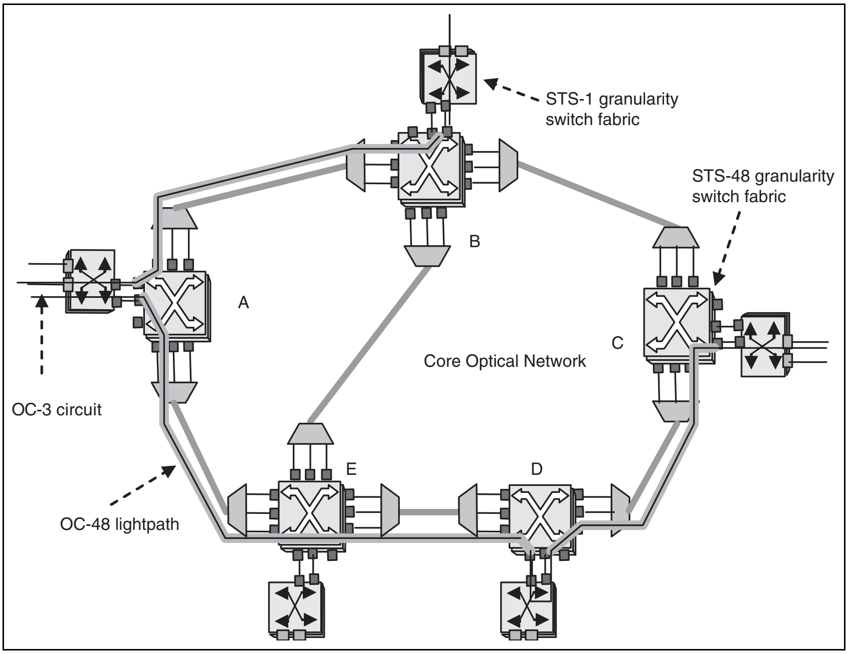

Figure 1.15: Two-tier network architecture. (From [189], Figure 4. Reproduced by permission of © 2003 The Optical Society of America.)

two backbone core switches. The third leg is an OC-N circuit from the backbone node core switch, onto the destination regional PoP. Such provisioning in multiple legs is due to the practical constraint that the regional and core networks are controlled by different subnetwork management systems, and recovery of circuits can be performed by different mechanisms on each individual leg of the circuit.

1.3.2 Two-Tier Network Architecture

This section focuses on the special case of two tiers for the multi-tier network architectures. In the two-tier network architecture, shown in Figure 1.15, the core optical switches switch at the STS-48 granularity (perform *core grooming* at STS-48 rates), and connected to a core optical switch are one or more switches that can switch at the STS-1 (or lower) granularity (perform *edge grooming* at STS-1 rates). The STS-48 switches terminate OC-48 and OC-192 services, and the STS-1 switches terminate STS-N services below STS-48. The STS-48 switches also groom STS-48 frames onto STS-192 frames. OC-48 or OC-192 services between backbone nodes are set up as lightpaths by finding a route in the core network, and configuring the STS-48 switches along the route. We term this the two-tier architecture because there is a core STS-48 switching tier, and there is a second tier that switches STS-1 traffic. It can be argued that a third tier is necessary for switching at granularities lower than STS-1, but we will nevertheless consider it two-tier because STS-1 and above granularities dominate the backbone traffic [88, 229].

We could also have a network architecture where the node architecture is not uniform. For example, all nodes could still have an STS-1 level switch, which is essential for handling sub-STS-48 traffic. However, some nodes, either carrying a large amount of add/drop traffic or occupying strategic switching locations, called hub nodes, could also have an STS-48 level switch. This architecture is considered a *mixed two-tier architecture*, because not all nodes in the architecture are of the same kind: some nodes are nonhub nodes, while other nodes are hub nodes. Nodes could now be connected not only through STS-48 switches, but also through neighboring STS-1 switches.

Figure 1.15 illustrates three OC-48 lightpaths (A–D, A–B and D–C), and two OC-3 circuits (A–B, and A–C). The OC-3 circuit A–B rides on the direct lightpath A–B. The OC-3 circuit A–C rides on lightpath A–D, *hairpins* into the STS-1 switch at node D, and rides onto lightpath D–C to its final destination. In this example, the STS-48 switch terminates all OC-48 and OC-192 services on the drop side.

In general, subrate services between backbone nodes are set up as follows: A set of OC-48 or OC-192 lightpaths between the core switches serves as an *overlay* topology for purposes of routing subrate services. For example, initially, the overlay topology may be identical to the physical topology, with a direct lightpath between all neighbors. If there is a direct lightpath between a pair of backbone nodes, and there is enough capacity on that lightpath, a subrate service between the node-pair can use that lightpath. In this case, the STS-1 switch serves the role of a multiplexer/demultiplexer device (as opposed to a switch). For this reason, subrate services that take the direct lightpath between two backbone nodes may be terminated on the STS-1 multiplexer/demultiplexer. If there is no direct lightpath with enough capacity, then the subrate service has to traverse multiple lightpath hops, and at each intermediate node *hairpin* into the STS-1 switch (get regroomed) and get switched onto the lightpath on the next hop. Intermediate grooming is a natural aspect of routing on an overlay topology [231, 298, 335, 336]. Figure 1.15 illustrates this. If the total subrate traffic between a pair of backbone nodes exceeds a threshold, then a direct lightpath may be set up between the pair of nodes terminating on drop ports that are connected to the STS-1 multiplexer, and services are terminated on the STS-1 multiplexer. For those services that cannot be routed on a direct lightpath, they can be terminated at the STS-1 switch, and routed over multiple lightpaths between the STS-1 switches with hairpinning.

The size of the STS-1 multiplexer at a node depends on the total subrate traffic demand, but the size of the STS-1 switch depends only on the number of backbone nodes. Hairpinning is a natural inefficiency of routing on an overlay topology. Hairpinning does not occur in the one-tier architecture because the core switch can switch at the STS-1 granularity. However, the amount of traffic that hairpins is bounded because, in principle, as soon as hairpinned traffic exceeds a threshold, a direct lightpath can be set up between a pair of nodes. In practice, this means that service routes have to be reoptimized periodically which is not an easy task. It is desirable to have subrate services routed on direct lightpaths, while at the same time ensure that network capacity is utilized efficiently by having lightpaths *well-packed.*

Note that if there are more than one STS-1 switch/multiplexer boxes connected to the STS-48 switch, then the network is blocking for STS-N services. A blocking second tier may be acceptable if, for example, each second-tier STS-1 switch/multiplexer terminates services for different customers, or if the STS-N services are expected to be static, and not dynamically changing. If the network needs to be nonblocking for STS-N services, then there has to be a single STS-1 nonblocking switch that terminates all STS-N services at the second tier. The size of the second-tier switch depends on the sub-OC-48 component of the traffic demand, and plays an important role in the two-tier architecture.

Also, to set up subrate services in the two-tier architecture, there needs to be coordination between the STS-48 switches and the STS-1 switches on the management plane. In contrast, in the one-tier architecture, all services can be provisioned using the management system that controls the STS-1 switches. An OC-N circuit between regional PoPs that has to traverse the core optical network is provisioned with multiple legs, just as in the one-tier architecture, with coordination among multiple management systems.

1.3.3 Network Scalability

Although switch fabrics in one-tier architectures that are designed to switch data with finer granularity can theoretically scale to large port numbers, the control and management of the core network (e.g.,

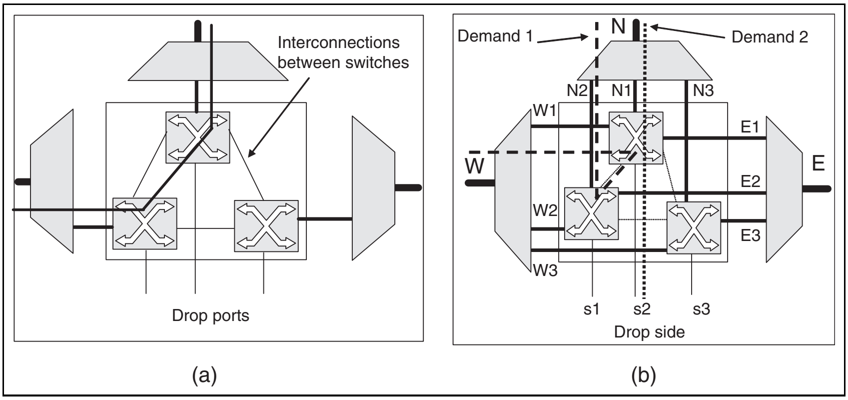

Figure 1.16: Multiple interconnected switches at a single site: (a) Channels from each WDM link are connected to a single switch, (b) Channels from each WDM link are connected to multiple switches.

large port-state databases and a large amount of performance monitoring information), as well as fast mesh recovery at that granularity, will present scalability problems. For example, fast SBPP will be difficult in a one-tier architecture even with bundling, because of the need to perform multiple cross-connects at the end-nodes of the failed link (simulation studies in [16] have shown this). Fast recovery could be achieved utilizing DBPP at finer granularity, however, this incurs the capacity penalty of DBPP. On the other hand, a two-tier architecture using higher-capacity switches is easier to control and manage, and it will be more scalable when recovery is performed at higher granularities.

In addition, as the traffic grows beyond the capacity of the switch over time (something which is more likely to happen with switches of finer granularity), multiple switches will need to be interconnected to yield a larger switch. Otherwise, the network will block some of the connections, or the network capacity will be used inefficiently by routing connections on longer paths. Figure 1.16 shows two possible architectures for multiple interconnected switches at a single site: in Figure 1.16(a) channels from each WDM link are connected to a single switch, whereas in Figure 1.16(b) channels from each WDM link are connected to multiple switches. Both configurations waste ports for interconnecting multiple switches together, as traffic passing through the node can potentially pass through multiple switch ports (interconnect penalty). As the percentage of the traffic that passes through the node is a large portion of the total network traffic,[8] this waste in terms of extra ports used will be significant [254]. In addition, the interconnect penalty is dependent on the traffic forecast accuracy, and it increases when the traffic forecast exhibits considerable uncertainty [189].

Studies show that for arbitrary but uniform set of connection requests, when the switches are interconnected in a mesh topology and 70% of the traffic is pass-through traffic passing through a single switch at the site, approximately 30% of the ports need to be designated as interconnect ports [254]. Furthermore, the design of the interconnection network and the algorithms on how to incrementally add switches and how to route connections in the interconnection network are of crucial importance to the usage of network capacity. Blocking conditions in the interconnection network (the interconnection network may be a blocking network, even though the switches that comprise this network are themselves nonblocking) will lead to stranded or inefficient usage of the network capacity, or a network may not recover from a failure condition, even though there might be enough protection capacity in the network.

---

[8] Typically 70–75% of the traffic that reaches a node will be pass-through traffic. The rest of the traffic will be dropped at the node.

Additional analysis of the one-tier versus the hierarchical architecture was performed in [189]. In the hierarchical architecture the network is scaled by organizing it into layers and these layers are optimized to switch and groom at different rates (STS-1 switching was used for the flat architecture and STS-1 and STS-48 switching for the hierarchical architecture). Studies have shown that there is a crossover point beyond which the layered architecture becomes more cost-effective as the total traffic grows and as the traffic mix evolves towards higher rates [189]. This study assumed uniform traffic pattern (with inaccurate traffic forecast), 30% interconnection ports for the STS-1 switches in the flat network architecture (with nonblocking interconnection), switch sizes in the two-tier network architecture that are not exceeded, and SBPP protection in both cases. Simulations showed that when the proportion of the OC-48 traffic becomes bigger than 50% in the two-tier architecture, that architecture becomes cheaper than the flat topology. In addition, the two-tier becomes cheaper than the flat architecture when the traffic scales beyond the capacity of the STS-1 switch. Similar crossover points were also detected for all unprotected traffic as well.

Furthermore, Capital Expenditure (CapEx) simulations of the flat and hierarchical architectures have shown that the two-tier architecture operating at the STS-48 level in one tier and at the STS-1 level in other, exhibits overall network capital savings on the order of 24−36% compared to the one-tier architecture operating at the STS-1 level [124, 229].

While clearly a number of aspects come into play for a quantitative and fair comparison between flat and hierarchical network architectures, it appears that a layered network exhibits economic efficiency compared to its flat counterpart, as well as improved scalability and network performance.

1.4 The Current State of Optical Networks

The first historical testbeds that were created utilizing optical networking included the Optical Network Technology Consortium (ONTC) [63], Multiwavelength Optical Networking (MONET) [307, 308] and the European Multiwavelength Transport Network (MTWN) [170, 230]. Currently, optical networking has been introduced in both the metro [160] and the long-haul arenas.

Initial testbed experiments [19], and the introduction of network elements such as Reconfigurable Add Drop Multiplexers (ROADMs) [331] have shown the applicability (and cost-effectiveness [267]) of optical networking in the metro space. The typical metro architecture consists of a number of interconnected rings (in a hierarchical fashion) but some mesh network topologies have also appeared. Several testbeds for metro WDM network have been deployed and are described in detail in [226, 301, 303, 333]. For additional architecture and simulation work on designing WDM metro networks the reader is referred to [21, 22, 207, 237, 269, 317], etc.

Long-haul and ultra-long-haul networks initially utilized optical fibers solely as the transmission medium (point-to-point links) and a number of experiments dealt with the enabling technologies and the problem of expanding the reach of these links [223, 302, 311]. With the addition of intelligent and reconfigurable optical cross-connects discussed in Section 1.1, these networks evolved to mesh-based architectures providing enhanced capabilities such as point-and-click provisioning and failure recovery. As analyzed extensively in Section 1.1, these networks can be opaque, transparent or translucent, utilizing various cross-connect switch architectures. Several experiments (and simulations) reported in [239, 266, 283, 309] discuss the design of applicable cross-connects and the viability of these architectures for long-haul and ultra-long-haul systems.

The use of optical networking for long-haul networks has left the laboratory and was successfully commercially implemented in at least three real networks, namely the Dynegy Global Communications nationwide mesh network utilizing Tellium's optical cross-connects (utilizing an opaque design) [66], the AT&T nationwide network utilizing Ciena's optical cross-connects (in an opaque design as well but different than Tellium's) [87, 259], and the Broadwing Communications Services deployment of Corvis' transparent cross-connect. The first two deployments utilized the intelligence of the optical cross-connects to address such issues as point-and-click provisioning, fast failure recovery, and traffic grooming.

The network deployments mentioned above also addressed control and management functionalities of the network, mostly in a proprietary manner. However, there have been significant advances in IETF in developing a control plane for optical networks. An extension of MPLS, namely GMPLS [210], is used for the control plane protocols in optical networks. The three main control functions that are addressed are neighbor discovery, signaling and routing. The reader is referred to [38] for an extensive presentation of MPLS, GMPLS and their traffic engineering extensions. As pointed out in Section 1.1, there are a number of issues that arise during the implementation of control and management functions when completely transparent switch architectures are utilized. A complete analysis of these issues and possible ways to solve these problems in a completely transparent architecture is presented in [108].

Some of the issues that are essential in the successful implementation of mesh optical networks, namely provisioning and failure recovery of connections (lightpaths) and dimensioning of the network, are exactly what the rest of the book is about. The main goal of this book is to present efficient techniques on path routing and single failure recovery[9] in opaque optical networks with arbitrary mesh topologies. These problems are emerging problems in optical networking that have for the large part been dealt with in the research community. This book provides a solid foundation for these problems that can be used by researchers to further their understanding in these issues, as well as by network architects and engineers for the design and implementation of real mesh optical network deployments. The section that follows briefly describes the content of each chapter and the interdependencies between the various chapters in the book.

1.5 Organization of the Book

The advancement of today's optical networks reveals three main trends: arbitrary mesh network topologies, large capacity sessions and the need for reliable services. These trends have motivated us to explore techniques that can provide lightpath provisioning and failure recovery in a fast and efficient manner in mesh optical networks.

This book presents an in-depth treatment of a specific class of optical networks, namely path-protection-oriented mesh optical networks, and focuses on the routing, failure recovery, dimensioning and performance analysis of such networks. Readers who are generally interested in survivability principles are referred to [319] for an extensive analysis of ring-based survivability (with a focus on SONET networks), to [139] for a meticulous account of other mesh-based approaches to survivability, including the *p*-cycles technique ([139] addresses survivability techniques for optical, ATM, SONET and MPLS networks), and to [305] for a detailed description mainly of the MPLS layer recovery mechanisms.

Furthermore, readers who are interested in optical transmission at the physical layer, or the technologies required for the deployment of intelligent optical networks, are referred to books such as [13, 233], and readers who are interested in optical networking in general are referred to books such as [221, 257, 289].

The rest of the book is organized as follows:

Chapter 2 is devoted mostly to background material on survivability techniques. It motivates the reader with a discussion on failures and the need for survivable networks and presents a survey of the existing optical network fault recovery techniques proposed in the literature. This is by no means a complete listing of all the recovery approaches, but they cover the main techniques that have been proposed over the years. It distinguishes between what are termed protection and restoration

---

[9] The focus of the book is on single failure recovery. Double failures are only considered in Chapter 12, in the context of service availability.

failure recovery techniques and discusses ring, link, path and segment-based survivability approaches. Discussions on multi-layer recovery and integrated protection/restoration approaches in IP-over-WDM networks even though somewhat out of the scope for this book, were added in order to show the direction where survivability research is progressing and to motivate the reader to investigate these areas in more detail.

Chapter 3 is tied directly to Chapter 2 and explores further the classification of fault recovery approaches, focusing on path-based protection techniques for mesh optical networks. The general notion of network components sharing a failure risk is introduced, and Shared Risk Groups (for links, nodes and equipment) are defined and analyzed. Chapter 3 also introduces routing approaches for survivable connections and examines briefly the cases of distributed and centralized routing without discussing the implementation details.

Chapter 4 continues from where Chapter 3 leaves off and describes in detail routing and recovery for the specific case addressed in the remainder of the book, that of failure independent preplanned path-protection for mesh optical networks. The chapter starts with a framework for routing path-protected connections in a mesh network, and then discusses protected connections via the Dedicated Backup Path Protection (DBPP) and Shared Backup Path protection (SBPP) techniques as well as other types of connections such as preemptible, unprotected, etc.

Chapter 5 analyzes the complexity of such routing problems, essentially the complexity of routing working and backup paths in mesh networks, and Chapter 6 introduces, discusses and presents results for various routing algorithms (mostly a variety of heuristic approaches).

Chapter 7 investigates an enhanced algorithm cost model to control trade-offs in provisioning SBPP lightpaths, and Chapter 8 describes three approaches for limiting the number of lightpaths protected by a shared channel for SBPP services in optical mesh networks.

Chapter 9 presents an extension to the computation of SBPP paths using statistical techniques and Chapter 10 investigates lightpath reoptimization and shows how reoptimization offers the network operator the ability to better adapt to the dynamics of the network (demand churn and network changes) that causes the routing to become suboptimal.

Finally, Chapters 11 and 12 address two very timely subjects at this time of writing, namely dimensioning and availability of mesh optical networks. Specifically, Chapter 11 describes analytical approaches to dimension mesh optical networks for backup path protection, and presents techniques that can be used to quickly estimate the network size and failure recovery performance with limited inputs. Chapter 12 ends the book with the modeling and analysis of the service availability mainly of the DBPP and SBPP services, which is a critical tool for the establishment of service level agreements for these services.

The book is organized and demarcated in such a way that readers who may want to just focus on some specific topics can do so without having to read the entire book. For example, readers who are interested generally in survivability can read Chapter 2 and readers who want to gain an insight into path-based protection approaches for mesh optical networks can read Chapters 3 and 4 as well. Readers who want to read through general information on routing algorithms for working and backup paths in mesh networks are referred to Chapters 5 and 6 and readers who are interested in further details on these routing approaches (enhanced cost metrics, limited sharing, routing using probabilistic methods and lightpath reoptimization) are encouraged to read Chapters 7–10. The dimensioning chapter (Chapter 11) and the availability chapter (Chapter 12) can be treated as stand-alone chapters (requiring only some limited background information from the previous chapters) for readers who are interested only in these subjects.