♦ Metro Optical Networking

Yi Chen, Mohammad T. Fatehi, Humberto J. La Roche,
Jacob Z. Larsen, and Bruce L. Nelson

Today's telecommunication networks support an ever-increasing mixture of statistical and deterministic data traffic. Current networks rely on time division multiplexing (TDM) hierarchy, originally invented to be the most efficient multiplexing technique possible for 64-kb/s voice, but this architecture is not particularly suited to statistical data traffic. A router or ATM switch can connect into a dense wavelength division multiplexing (DWDM) transport network by mapping packets or cells directly onto a wavelength without the intervening use of a SONET or SDH TDM. In effect, a TDM can be replaced with an optical DWDM, which can increase bandwidth utilization, facilitate networking, and reduce cost. DWDM can defer or completely eliminate the need for extra fiber, which is especially significant for providers who have a fiber-exhaustion problem, and it can easily coexist with today's SONET/SDH networks or with older fiber-optic terminals (FOTs) operating on asynchronous protocols. DWDM has already revolutionized the telecommunication industry by providing the infrastructure for long-haul optical networks. As the DWDM revolution moves into the metropolitan interoffice (IOF) and access networks, it will rely on the optical add/drop multiplexer (OADM) as its fundamental network element for metro optical networking. In this paper, we examine the market needs for metro optical networking; introduce the concept, network topology, and architectures for metro optical rings; and describe the key requirements that address the needs of service providers. We also discuss the components of the network and review the general requirements that will enable the next step in the evolution—extending the DWDM ring architecture to business access customers. Metro optical networking represents a unique opportunity for service providers to begin deploying a data-centric high-bandwidth services infrastructure.

Introduction

Dense wavelength division multiplexed (DWDM) optical transport technology has revolutionized the telecommunication industry. With the advent of wideband fiber-optical amplifiers, DWDM optical transmission systems are now capable of providing capacities in excess of hundreds of gigabits per second over hundreds of kilometers on a single pair of fibers in long-haul networks. Currently, the DWDM revolution is moving into the interoffice (IOF) and access networks. Consistent with the optical networking vision presented by Alferness et al. in this issue[1] and by Al-Salameh et al. in a previous issue[2] of the *Bell Labs Technical Journal*, the telecommunications infrastructure is evolving toward flexible, all-distance, wavelength-managed optical networking in the near future. To use this enormous capacity to the fullest, we must now focus on managing the traffic at the optical wavelength level. On one hand, we can avoid the bottleneck of electronic processing; on the other hand, we can respond to the fiber exhaust problem. The result is

Panel 1. Abbreviations, Acronyms, and Terms

ADM—add/drop multiplexer
AGC—automatic gain control
ATM—asynchronous transfer mode
BGP4—border gateway protocol, version 4
BLSR—bidirectional line-switched ring
CEPT—Conference of European Postal and Telecommunications Administration
CIO—chief information officer
CO central office
CTO—chief technical officer
CW—continuous wave
DACS—Digital Access and Cross-Connect System
DNI—dual-node interworking
DS1—digital signal level 1, transmission rate of 1.544 Mb/s
DS3—digital signal level 3, transmission rate of 44.736 Mb/s
DSF—dispersion-shifted fiber
DWDM—dense wavelength division multiplexer/multiplexing
E1—2.048 Mb/s rate used by European CEPT carrier to transmit 64-kb/s digital channels for voice or data calls, plus a 64-kb/s signaling channel and a 64-kb/s framing and maintenance channel.
EDF—erbium-doped fiber
ESCON—Enterprise Systems Connectivity
FBG—fiber Bragg grating
FOT—fiber-optic terminal
IOF—interoffice
IP—Internet protocol
IS—intermediate system
ISP—Internet service provider
ITU-T—International Telecommunication Union-Telecommunication Standardization Sector
LSO—local serving office
MCNC—Microelectronics Center of North Carolina
MEMS—microelectromechanical systems
MI—Michelson interferometer
MUMPS—Multi-user MEMS Process
MZI—Mach-Zehnder interferometer
MS—multiplex section
NE—network element
OA—optical amplifier
OADM—optical add/drop multiplexer
OC-*N*—optical carrier digital signal rate of $N \times 53$ Mb/s
OSPF—open shortest path first
OTU—optical translator unit
PDH—plesiochronous digital hierarchy
PON—passive optical network
QoS—quality of service
SDH—synchronous digital hierarchy
SLA—service-level agreement
SN—service node
S/N—signal to noise ratio
SOA—silicon optical amplifier
SONET—synchronous optical network
SPRING—shared protection ring
S-PVC—soft permanent virtual connection
STM-1—synchronous transport module 1, an SDH standard for transmission over OC-3 optical fiber at 155.52 Mb/s
STM-*N*—synchronous transfer module, level *N*
STS-1—synchronous transport signal 1, a SONET standard for transmission over OC-1 optical fiber at 51.84 Mb/s
TDM—time division multiplexing/multiplexer
VT—virtual tributary
WDM—wavelength division multiplexed/multiplexing
WGR—waveguide grating router
XPM—cross-phase modulation

a network that provides seamless connectivity from the long haul to the metropolitan area, to access, and even to the desk top. It also provides connectivity that is bit-rate, format, and protocol transparent for all types of services, including data, video, and telephony.

DWDM technology has already proven its worth in long-haul networking. The ability to effectively "mine" bandwidth from existing fiber strands has saved billions of dollars in the capital expense budgets of providers who would have otherwise had to invest in deploying new fibers. It has also allowed carriers to meet time-to-market needs as they deploy new services. Key to the success of long-haul DWDM has been the optical amplifier (OA), which amplifies multiple signals efficiently, carrying each as a distinct wavelength in the optical fiber.

The metro DWDM network is quite different from a long-haul network. Metro is largely driven by the growth in data traffic, which has created a market opportunity for data networking equipment such as Internet protocol (IP) routers and asynchronous transfer mode (ATM) switches equipped with clear channel

OC-48 or STM-16 interfaces that operate at 2.5 Gb/s. Data networking equipment uses statistical multiplexing rather than time-division multiplexing (TDM), the technique more prevalent in the add/drop multiplexers (ADMs) of today's synchronous optical network (SONET) or synchronous digital hierarchy (SDH).

In TDM, lower-speed tributaries of a fixed bandwidth are aggregated into higher-speed signals, also of a fixed bandwidth. The TDM hierarchy, originally invented to be the most efficient multiplexing technique possible for 64-kb/s voice, is not particularly suited for a network that supports an ever-increasing mixture of statistical data traffic. A statistically multiplexed service can take advantage of data traffic's burstiness by allowing a signal to momentarily fill a larger container, which continues to carry other signals. Customers using a "statistical" service benefit by having a service contract that may allow for bursting at some peak rate (which may be equal to the access pipe bandwidth). This type of service delivers bandwidth on demand. Similarly, service providers benefit by being able to manage capacity flexibly through interface oversubscription, assuming that all customers do not use the peak rate simultaneously.

Metro DWDM can facilitate the deployment of high-speed data networks. A router or ATM switch can connect into a DWDM transport network by mapping packets or cells directly onto a wavelength without the intervening use of a SONET or SDH TDM multiplexer. In effect, the cost of a TDM is replaced with that of an optical DWDM, which increases bandwidth utilization. The equipment cost of transporting an OC-48/STM-16 client signal on DWDM is lower than the cost of transporting a similar signal over an OC-192/STM-64 TDM network.

DWDM can defer or completely eliminate the need to deploy extra fiber. This is especially significant for providers who have a fiber-exhaustion problem, because either they lease *dark* fiber—that is, unused, buried fiber—by the strand or they find the cost of laying new fiber in the ground prohibitive. DWDM can easily coexist with today's SONET/SDH networks or with older fiber-optic terminals (FOTs) operating on asynchronous protocols. The primary motivations for deploying DWDM networks are:

- To promote efficiencies in transport by deploying packet- or cell-based transport directly on a wavelength,
- To eliminate TDM multiplexing equipment at OC-192/STM-16 rates, and
- To reduce the cost of using fiber.

All three factors are likely to have an impact, in varying degrees, on a carrier's deployment decisions.

In the sections that follow, we examine the market need for metro in detail. First, we introduce the concept, network topology, and architectures for metro optical rings, based on the concept of "path in lambda," which has a SONET/SDH counterpart. We then examine the key requirements of the service providers and discuss the components needed to deploy such a network. Finally, we review general requirements that would enable the next step in the evolution—extending the DWDM ring architecture to business access customers.

Metro Optical Networking: The Market

This section describes the market for metro optical networking. In metropolitan networks the market for DWDM is growing in proportion to the bandwidth needs of enterprises. DWDM meets and exceeds enterprise bandwidth needs by relieving fiber exhaustion, by being independent of the various bit rates and protocols used in the network, and by allowing service providers to offer new services and a better grade of service on existing equipment. This section also examines the order in which these criteria will be added to the network and the size of the various markets affected.

DWDM in Metropolitan Networks

Although telephone companies are experiencing strong, consistent growth in the demand for bandwidth, an increase in competition is eroding the profit margins in their core businesses. As a consequence, network service providers must be able to offer services that support their most valuable end users—bandwidth-hungry enterprises.

The paradigm for serving these customers is changing at a rapid pace. Previously, enterprises demanded cheap voice services and some data services, known in the old nomenclature as *value-added services*. Today, however, many enterprises are choosing service providers who offer a wide variety of services predominantly focused on data, along with an end-to-end service quality guarantee.

Enterprises are also becoming more sophisticated. Communications professionals now play an important role in the ranks of an enterprise, as evidenced by the increase in corporate titles such as “chief information officer” and “chief technology officer.” The advent of electronic commerce can only add to the level of sophistication as the communications budget moves from the category of “expenses” to “revenue.” Telephone companies will choose their technologies carefully, as a consequence not only of increasing customer sophistication, but also of increasing competition, an explosive demand for bandwidth, and a demand for proliferation of services (such as ATM, IP, frame relay, and video over ATM and frame relay).

In the long-haul network, DWDM has a clear value proposition because it affords customers a much cheaper way to expand capacity; as a result, it is the technology of choice to satisfy the demand for bandwidth. (Industry-wide revenues from shipments of DWDM systems have already exceeded the revenues from shipments of STM-64/OC-192 systems.) In the metro area most telephone companies have not yet decided whether they want to boost their capacity by using DWDM. In many instances they can lay new fiber or upgrade their existing metropolitan rings with TDM. It is important, however, to realize that these solutions are not viable options for all service providers because:

- The cost of laying new fiber is prohibitively high in many areas, especially in densely populated metropolitan areas, where the anticipated demand is often highest;
- Right-of-way issues are very complicated and often do not favor laying new fiber;
- Several customers have already upgraded to STM-16/OC-48 and cannot go any higher because OC-192 does not offer the granularity desired to carry many data services; and
- The cash flow implications of upgrading with TDM and fiber are unfavorable because all resources have to be committed up front.

As a consequence, an increasing number of service providers are turning to DWDM as a solution to their need for more bandwidth in the metro area.

Value Proposition for Metro DWDM

The metropolitan network, unlike the long-haul network, interconnects with low-speed access rings, high-speed long distance trunk networks, and other interoffice networks. In other words, metropolitan networks are more interconnected, can carry diverse types of traffic, and are more geographically limited than the long-haul network, although they still serve as feeders to the latter. (The various connection types are illustrated later in this paper). This large difference has led many to doubt that DWDM—which so far has predominantly been a backbone technology—would ever play a role in the metropolitan network. Despite this difference, we believe that DWDM has a very strong value proposition for the metro area.

Fiber exhaustion. Fiber exhaustion has occurred most often in the interoffice segment. Deploying DWDM on entire rings rather than in point-to-point applications offers the advantage of creating new “virtual” rings or “virtual” fibers rather than just boosting capacity on a dedicated stretch. As new services and applications emerge, the network can be scaled to the very fast transmission capacities required to service the demand for bandwidth. The ability to scale quickly will give service providers the bandwidth to capture new customers and to start generating revenues sooner than the traditional stepwise capacity upgrade offered by TDM. However, the rationale for deploying DWDM in the metro area goes far beyond relieving the problems of fiber exhaustion.

Bit rate and protocol transparency. The protocol independence of DWDM enables telephone companies to carry native enterprise data traffic such as gigabit Ethernet, ATM, and IP over SONET on discrete channels. The optical layer provides a common transport fabric independent of the types of traffic/service it is carrying. The same advantages are offered in the ability to carry varying bit rates on different channels in the same fiber, which is currently possible in the backbone. This is particularly important in the metropolitan network, where a wide variety of services and bit rates are being generated. The transparency of DWDM and the add/drop capabilities of the metro DWDM systems will allow service providers to add and/or drop entire channels without reformatting (or layering) the original type of traffic. (We describe add/drop capabilities in the metropolitan network later in this paper.) Although these capabilities enable the service provider to save on equipment by eliminating the TDM hierarchy and other "reformatting" hierarchies, it does require optical interfaces on the access equipment.

Improved service offerings. Using DWDM in the metro network enables optical restoration, which can be done more cost effectively than in the electrical domain. Given that the restoration in the optical domain is independent of the service and the bit rate, traffic not inherently protected by its own scheme—such as plesiochronous digital hierarchy (PDH)—can also be protected in a metro DWDM environment. Service providers can then offer a better service grade on both existing and new services. In addition, a service provider will be able to resell bandwidth to the enterprise and to other service providers, creating a new revenue opportunity.

Metro Market Development

The development of the metro DWDM market will be driven by three factors, ranked in their order of importance:

1. Fiber exhaustion in the service provider's metropolitan network,
2. The enhanced networking capabilities of DWDM, and
3. The service provider's ability to add new services.

The first two items save equipment costs, and the third enhances revenues (that is, protects profit margins). Figure 1 illustrates this evolution.

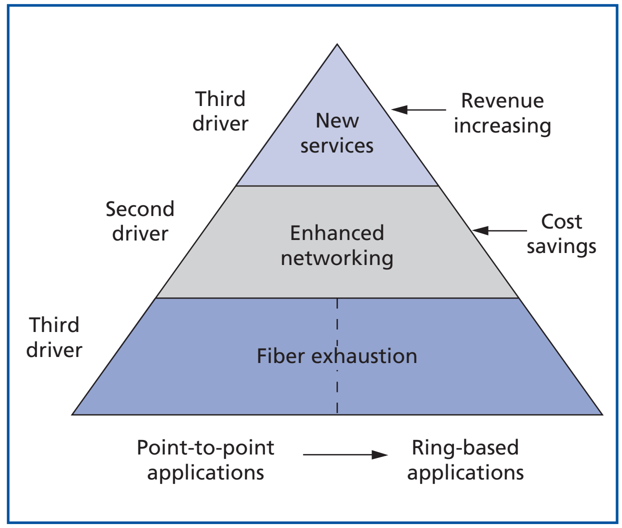

Figure 1.
The DWDM evolution.

The chronology of these market drivers strongly suggests that, initially, the market will demand point-to-point applications to resolve network congestion in areas of heavy traffic. As we mentioned earlier in this paper, DWDM is not the only technology that can relieve fiber exhaustion. The size of this market will be capped to areas where the economics of using DWDM are more favorable than the economics of using TDM or laying new fiber. Given that this "fire-fighting" model is not sustainable, a ring-based market will emerge.

Currently, a number of service providers in the United States, Europe, and Asia are participating in trials of point-to-point systems. However, a subtle difference exists in the rate of adoption we can expect from the different markets. The difference is predominantly driven by two factors—the probability of choosing the alternatives to metro DWDM as a solution, and the degree of previous exposure to DWDM. Not surprisingly, the first to adopt this technology will be customers who have used DWDM before (typically in long-haul networks) and those for whom the cost of using the alternatives will be high.

In general terms this market evolution suggests that the global market for metro DWDM will emerge first in the United States, where the value of DWDM has been proven in the long-haul network and where the expected increase in demand for bandwidth is highest. It is highly probable that the European market will follow next, delayed only by two factors. First, the fiber plant—on average—is newer in Europe, and thus the percentage using fiber is lower than it is in the United States. Second, the demand for bandwidth is slightly less pervasive in Europe, because the Internet has not made as many inroads into European households or enterprises.

Metro DWDM Market Size

Ovum, a global market analysis firm, recently published its 1998 forecast for the market size of metro DWDM systems, the results of which are shown in

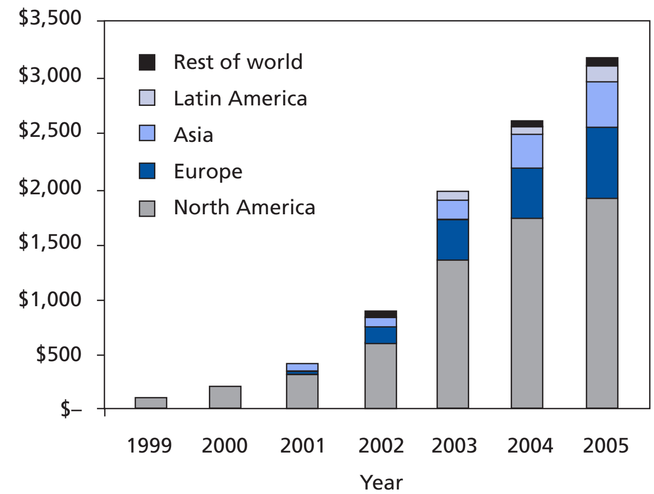

Figure 2.
The global metro market size.

Figure 2. Their market forecast seems to confirm the three claims about market opportunity made earlier in this paper, namely, that the market opportunity is large, has a strong momentum, and will develop gradually from one geographic region to the next.

Metro Optical Ring Architectures

Metro optical networks consolidate the transport of all types of services. The common infrastructure used in these networks enables the ring architecture to consistently support multiple forms of traffic. This section describes optical ring architectures.

Transport Consolidation

Figure 3 shows an explicit application of transport consolidation. In Figure 3a we see the variety of point to-point demands that may occur between different types of network elements (indicated as arrows between network elements), and in Figure 3b we see how DWDM can provide the connectivity required. If traffic between the two SONET/SDH ADMs is already protected by a TDM (SONET or SDH) scheme, it should be transported unprotected at the optical layer. On the other hand, traffic between IP routers may require protection beyond the painfully slow router table update implemented with protocols such as border gateway protocol, version 4 (BGP4) or open shortest path first (OSPF). In that case, the optical layer protection is switched on as an option. Traffic between ATM switches may consist of soft permanent virtual connections (S-PVCs) and may therefore be restorable at the ATM layer. For these as well, it may be possible to turn off the optical layer protection.

We proposed a "path-in-lambda" ring architecture,[3] which is being developed by Lucent and other equipment vendors to replace the "path-in-line" SONET/SDH ring architecture. Using this architecture along with today's technology could significantly reduce costs. The architecture is also well suited to meet the needs of tomorrow's metro IOF optical networking. It offers bit rate, protocol, and format-independent transport with full connectivity between a number of local serving offices (LSOs), or central offices (COs), in a metropolitan area. Metro optical access rings connect one or two LSOs to a number of business or residential customers, providing a full line of IP, data, telephony, and video services on dedicated wavelengths.

Path-in-Line Architecture: SONET/SDH Rings

Figure 4 shows a typical metropolitan area network architecture that provides local connectivity for various types of services. Today, these networks can be built using SDH/SONET rings based on "path-in-line" architecture. Path-in-line architecture (SONET/SDH rings) is structured around a high-capacity backbone

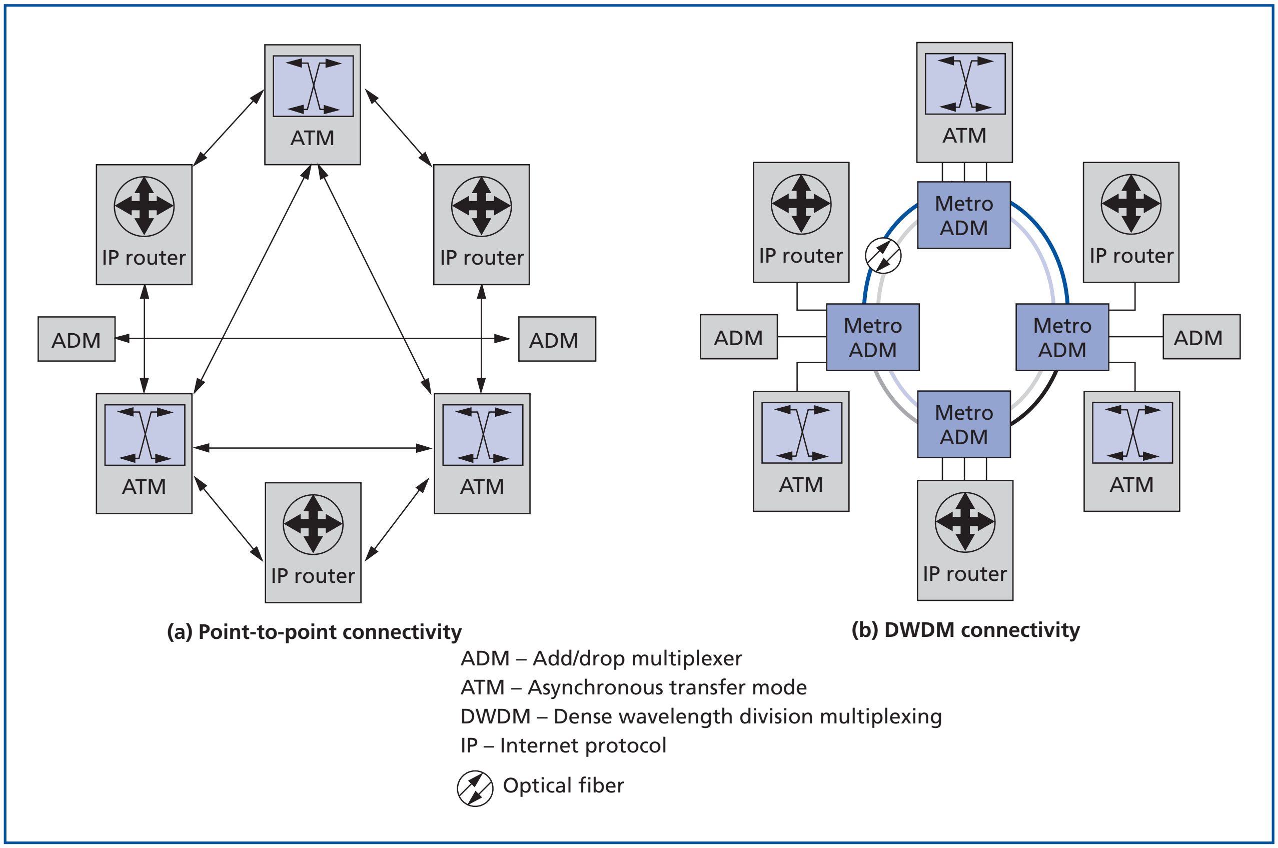

Figure 3.
Consolidating transport in an optical ring.

ring, or *IOF ring*. This ring provides high-speed connections (currently at OC-48/STM-16) between the main hub, called the *service node* (SN), and a number of COs or LSOs, denoted by LSO-1, LSO-2, …, LSO-*m* in Figure 4. In a typical metropolitan area network, *m* can be as few as 4 or as many as 15 offices. Multiplexed onto the backbone ring is the traffic of several lower-capacity distribution rings at each LSO. The distribution rings are path switched at the virtual tributary (VT) or STS-1 level, which is 1+1 protected. The lower-speed distribution rings are carried over a single rail, or *time slot*, in the interoffice ring. Multiplexing the distribution rings onto a backbone ring conserves fiber pairs.

The IOF ring is a two-fiber or four-fiber bidirectional line-switched ring (BLSR)[4] or multiplex section/shared protection ring (MS/SPRING)[5] that provides service and protection paths based on dual homing to the service node or among LSOs. Time slot reuse enables the IOF ring to support efficient mesh connectivity. The network elements (NEs) at the LSOs are SDH/SONET add/drop multiplexers (ADMs), which drop and add (that is, peel off) a set of designated SDH/SONET or PDH tributaries (OC-3/12, STM-1/4, DS1/DS3, E1/E3) and interface them to the traffic on the access rings. In addition to the ADMs, the IOF ring is equipped with regenerators at locations where the loss budget so dictates.

In access, or distribution, rings the hub located at the LSO collects traffic from multiple customers on the access side. Requirements for the total number of wavelengths are not high, because one would expect low-capacity traffic from a small number (2-6) of business customers to be collected by the access rings and routed to the IOF. The access rings are not expected to support wavelength reuse, since mesh-type connectivity is unusual in access networks.

There are two types of access rings—access loops and access arcs—both shown in Figure 4. An *access loop* emanates from and terminates within a single LSO,

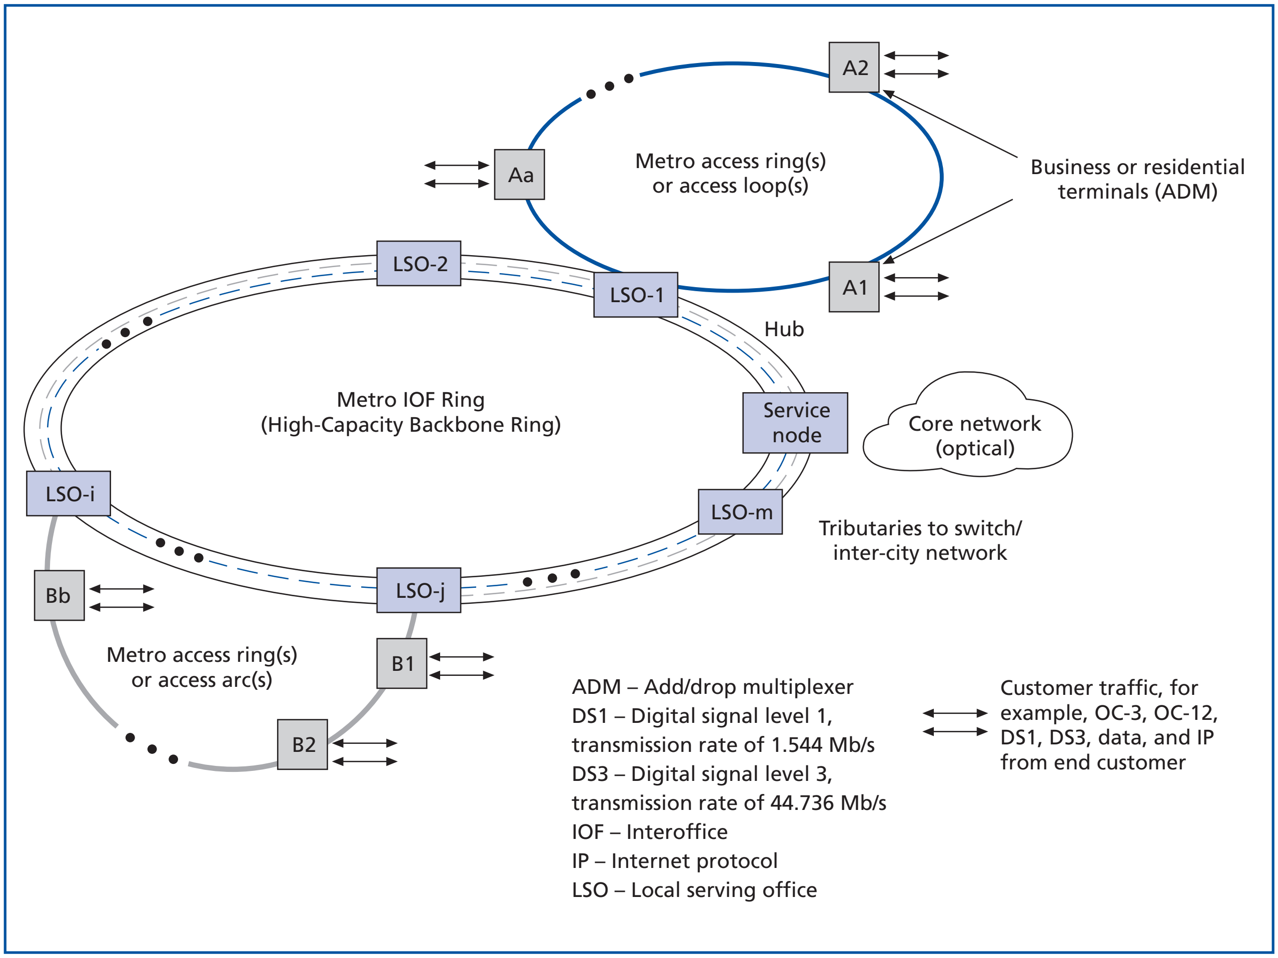

Figure 4.
A metro area optical network topology.

whereas an *access arc* emanates from one LSO and terminates in another. The access loop provides path diversity (protected) connectivity between the LSO and other locations. An access arc may use a service node as one terminating node and another LSO as the other terminating node.

SONET/SDH rings in metro IOF applications have a proven track record. They effectively support mesh traffic patterns and provide both line and network element protection on a per-tributary basis. However, SONET/SDH rings do not improve fiber efficiency, nor are they particularly well suited for data traffic.

DWDM optical rings will address such deficiencies, as well as possessing and improving on the advantages of SONET/SDH rings. For example, the ability of SONET/SDH to be flexible and reconfigurable can become an attribute of metro IOF rings. The sample application described below, the path-in-lambda architecture, illustrates how the optical layer can acquire some attributes of SONET/SDH.

Path-in-Lambda Architecture

DWDM optical networking replaces the SONET rings in Figure 4. Since the initial use of optical networking in metropolitan applications is expected to be in metro IOF rings, we describe it first. The characteristics and requirements for introducing DWDM optical networking in access rings are discussed in "Metro Access Rings."

Figure 4, described earlier, also illustrates how path-in-lambda architecture can be used in metro DWDM optical rings. This architecture replaces the OC-192/STM-64 backbone with a multi-wavelength optical channel (lambda) that can transport as many as 40 wavelengths, each capable of carrying OC-3/STM-1, OC-12/STM-4, or OC-48/STM-16 signals. Central offices (LSOs) on the backbone ring are required to drop multiple signals at various rates, formats, and protocols. This rate, format, and protocol transparency makes the IOF optical ring suitable for networking the new services in the metro area. In addition, the added multiplexing dimension (wavelength) allows a significant capacity upgrade without the need to increase the rate to OC-192/STM-64 and beyond. As a result, the DWDM ring solves the problem of fiber exhaust in the metro area.

At each LSO on the IOF backbone ring, optical translator units (OTUs) convert (translate) the customer traffic from the access rings, or arcs, to a set of designated wavelengths and then wavelength multiplex them with the rest of the traffic on the ring. For example, in Figure 4, the wavelength λ_i, depicted by the blue circle, is used to interconnect the traffic from the access loop to the service node. Similarly, the wavelength λ_j, depicted by the gray circle, is used to interconnect the traffic from the access arc interfaced at LSO-i and LSO-j (that is, customer sites B1, B2, ..., Bb).

Optical Transport Network Considerations

The emerging picture of metro optical networking portrays the optical layer as the universal transport layer for all services, consistent with the vision expressed at the beginning of this paper. Wavelength reuse is important, because it efficiently supports the mesh-based connectivity typical in interoffice networks. While this is a trivial statement from the perspective of DWDM capabilities, transport consolidation across a DWDM backbone entails management capabilities that are far from trivial. For example, the problem of provisioning a path across a ring in the most cost-effective (and capacity-saving) manner is not a trivial one. Consider the simple situation shown in Figure 5. A point-to-point demand for traffic to go from A to B can be routed directly, as in Figure 5a, or it can be routed the long way around the ring, as in Figure 5b. At first blush, one could invent a routing rule that one should always proceed with the shortest path on a first-come, first-served basis (in fact, SONET provisioning systems have often been designed to do this). The route chosen, however, always depends on the other traffic present on the ring. For example, the wavelengths on span from A to B may be filled with other traffic, making it impossible to add more traffic to the short path. In a ring, the long path may have hidden cost consequences as well, since it may be necessary to regenerate or optically amplify signals affected by optical path distance impairments.

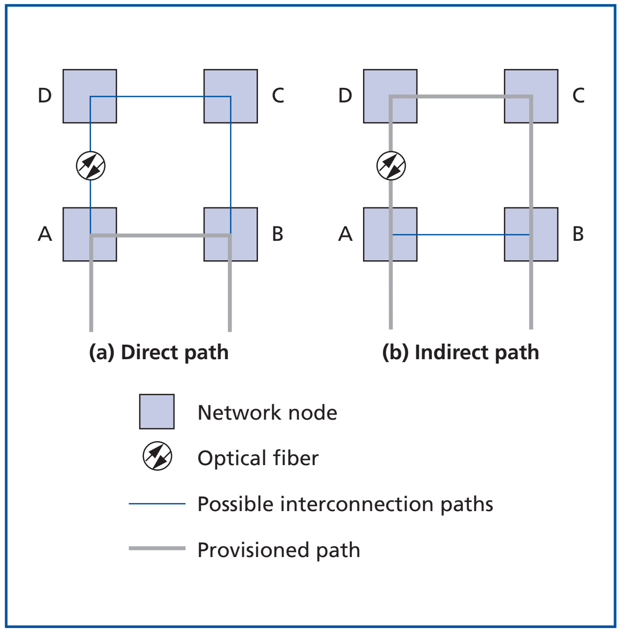

Figure 5.
Routing a point-to-point demand.

The problem of routing traffic around an optical ring in a way that minimizes cost is important to service providers, who must find ways to maximize revenues. The most obvious solution to limiting routing costs is for element managers to find ways to provision rings intelligently by reusing the wavelengths within the ring.

Strategies for Fast Restoration of DWDM Traffic

Another consideration in the quest to maximize revenues is the self-healing network strategy. Optical networking will, in time, support powerful self-healing capabilities consistent with those that SONET/SDH offers today. As we evolve toward all-optical networking, the key attributes of SONET/SDH needed in a totally encompassing optical infrastructure are:

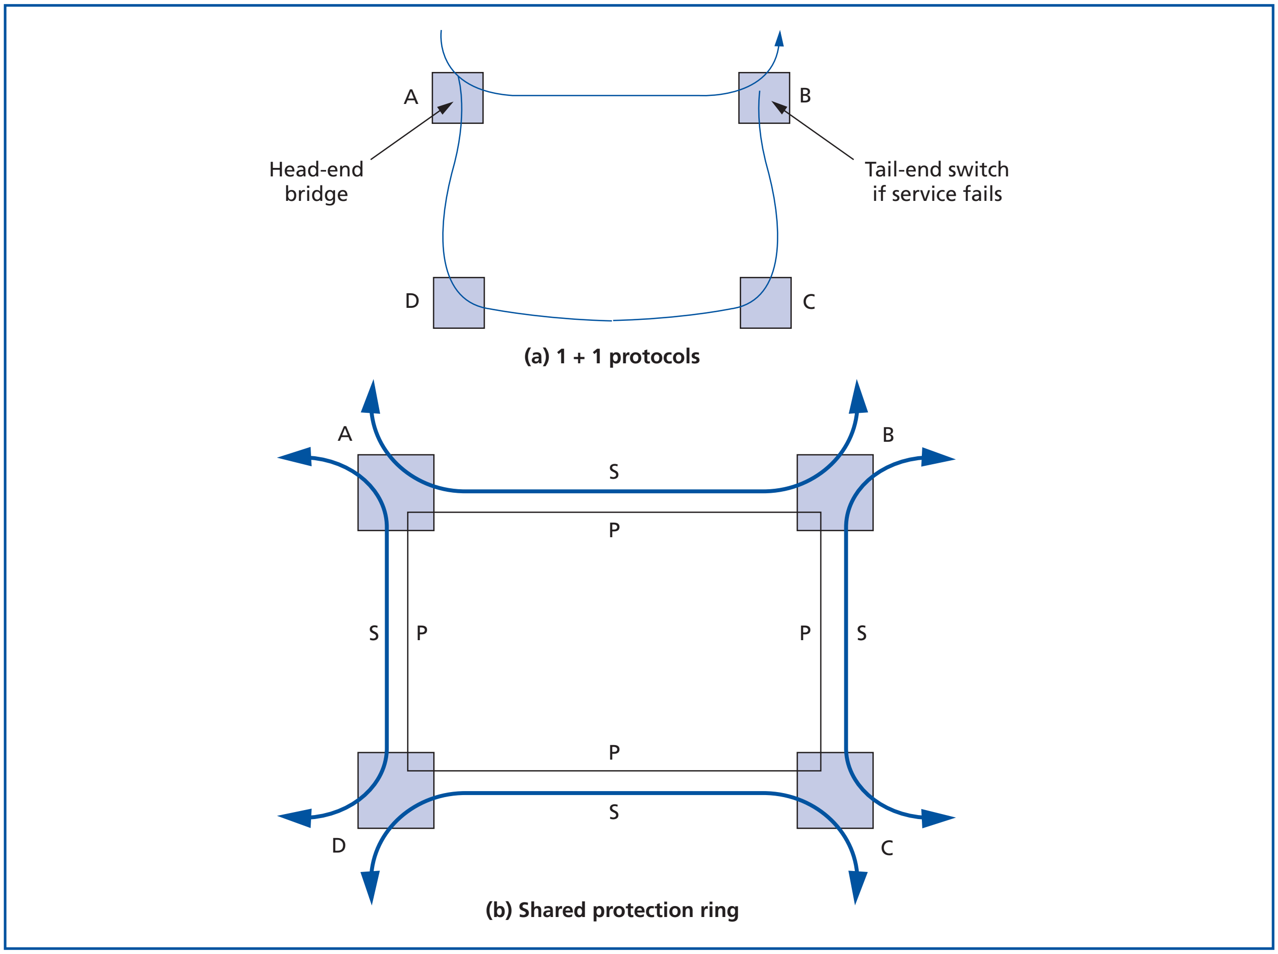

Figure 6.
Optical layer ring restoration.

- *Fast restoration of traffic*. Currently, SONET/SDH systems can restore service around link failures in about 50 ms. Because service-level agreements (SLAs) between service providers and end customers of these systems often contain stringent survivability clauses, any optical network will be required to conform to these.
- *Interface protection*. An optical system should be able to reliably interface with other network elements. Today, these interfaces are protected by using the 1+1, 1:1, 1:*N* interface protection, described in the specifications GR-253[4] and ITU-T G.841.[5] In optical systems where the key issue is to terminate the protection boundary (as defined in GR-253), it is necessary to introduce an optoelectronic conversion to process the K1 and K2 bytes, part of the SONET (and SDH) line overhead. This conversion entails the additional cost of the detection equipment and expenses associated with decision-making and performing the rearrangement and switching.

Optical layer ring restoration can be based either on simple 1+1 protocols or on a shared protection ring scheme, both shown in **Figure 6**. In the restoration strategy, shown in Figure 6a, traffic is split into two identical copies (head-end bridging) at the head end, or traffic origination point; at the receive end, a selector picks the better of the two signals (tail-end switching). The two copies of the circuit are routed in separate directions on the ring. Protection switching may also be performed on each tributary.

In a shared protection ring scheme, shown in Figure 6b, traffic is protected on a per-wavelength basis by a shared protection channel. Each of the separate services—A-B, B-C, C-D, and D-A—shares the same protection channel (P). While very similar to the MS/SPRING and BLSR schemes described in ITU-T G.841[5] and Bellcore GR-1230,[6] respectively, an optical shared protection ring has no tributaries. In fact, this type of protection ring is defined by a set of 2+2 counter-rotating wavelengths and protection logic. Two of these can be reused for service traffic on different spans, and the third and fourth are reserved as shared protection channels for all paths on the ring. The protection logic consists of a set of algorithms communicated over an optical supervisory channel that enables the ring to self-heal in the event of a failure. Optical shared protection rings and 1+1 ring protection are options for implementing an optical layer survivability strategy for metro optical networking.

The second issue we will discuss in connection with survivability is interface protection. This type of scheme is implemented to reliably interface optical rings to other network elements. The key consideration of this scheme is the location of the protection boundary. For interfaces, the protection boundary terminates at the edges of the ring node. Optical ring elements, in particular, must have additional electronic processing internal to ring elements to accommodate traffic that is added or dropped locally. Currently, strategies for implementing this capability rely on OTUs, which are required for wavelength conversion (see "Metro IOF Ring Node Architecture"). Clearly, shared protection requires coordination among the nodes involved. This, in turn, may slow down the protection switching process, unless a "predetermined" or "precomputed" protection strategy is put in place and followed.

In addition to the previous considerations, we must discuss the issue of multi-layer survivability, part of any introduction to the metro IOF networks. Often, protected traffic will be introduced into the optical network because, for example, it is part of a SONET/SDH ring. Data traffic emanating from ATM switches may be protected by ATM layer restoration procedures (so-called S-PVCs). In addition, nonpremium IP traffic may be reliably, though slowly, restored by using network layer procedures, which rely on router table updates. For such services it is desirable to be able to turn off the optical layer restoration on a per-channel basis to prevent any potential race condition among various protection switching layers.

Data Applications

A wide variety of data applications can be deployed with metro optical networking. Consider, for example, today's typical Internet service provider (ISP) network. Metro router networks are implemented as fully connected meshes, with paths from each router to every other router implemented over the ATM layer. Often, two ATM switches are introduced consistent with a survivable architecture (if one switch fails, the other one intervenes). These designs help support efficient networking, with bandwidth management provided by ATM quality of service (QoS). Typically, routing protocols such as integrated intermediate system to intermediate system (IS-IS) operate over these links.

Figure 7 shows two methods of implementing meshes. An alternative to implementing the mesh at the ATM layer, shown in Figure 7a, is to implement it directly at the optical layer using DWDM on a wavelength add/drop ring. A reconfigurable optical ring, shown in Figure 7b, provides a direct packet over SONET path between optical add/drop multiplexers (OADMs). IP-level bandwidth management (such as that found in differentiated services or priority queuing) is assumed. DWDM has, in effect, replaced the ATM switch(es) supporting router paths.

Another application of metro DWDM is interconnecting enterprises. IBM has introduced an optical interconnect for its mainframe computers known as Enterprise Systems Connectivity, or ESCON. ESCON allows an IBM mainframe to access remotely located resources (such as high-speed storage devices) with fiber-optic technology. The ESCON logical frame definition, however, is not a standard definition used in a transport network infrastructure, so supporting transport over existing TDM networks would require complex adaptation equipment. The clear-channel transport capability of DWDM networks can help here.

So-called broadband OTUs act as the interface between bit-rate and format-independent traffic and the DWDM optical ring. These OTUs support a broad-

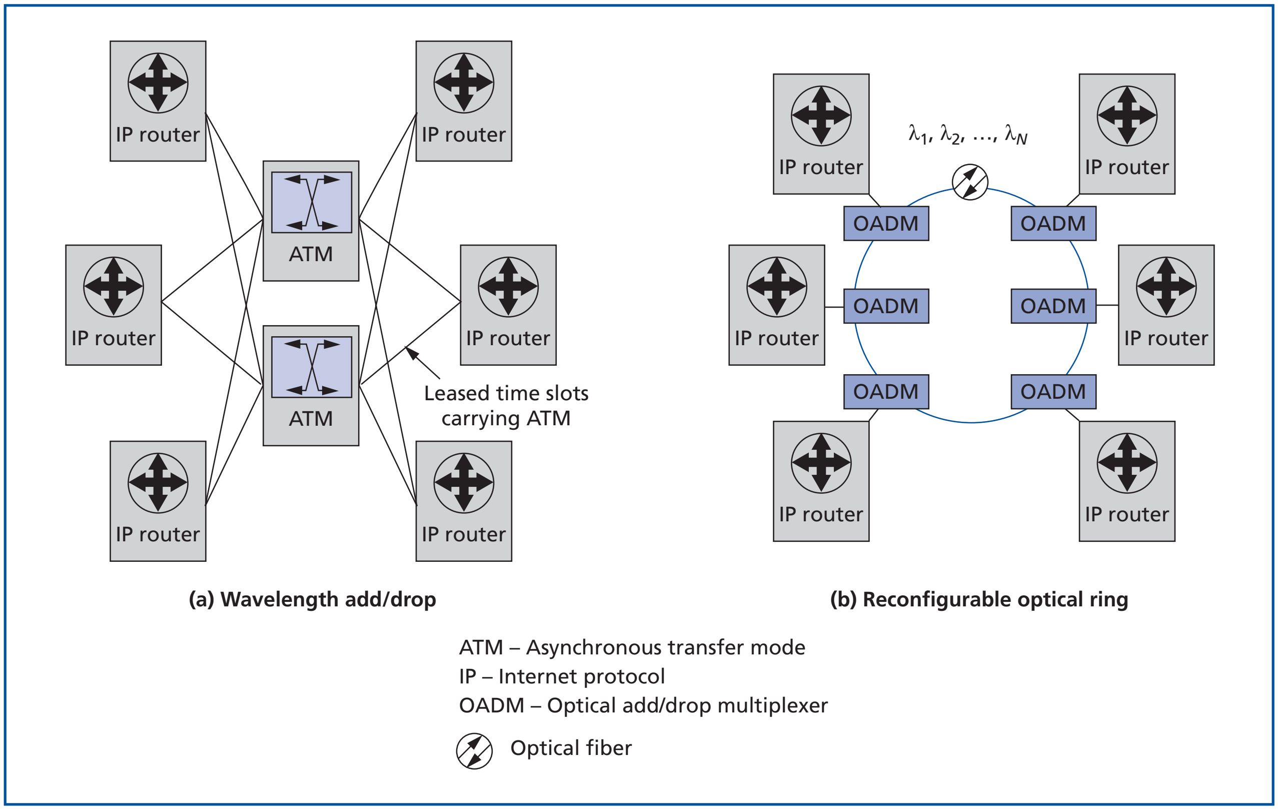

Figure 7.
Two methods of implementing mesh.

band phased-locked loop circuitry, which can latch onto signals in a specified range of bit rates.

Emerging Markets Case: Asia/Pacific Region

We have chosen to use the introduction and evolution of DWDM in metropolitan networks within the Asia/Pacific region as an example of an application tailored to the global emerging markets.

Figure 8 shows a typical layered network structure in which the existing SDH metro network has two layers:

- Layer 1, the metro access (622-Mb/s) network, and
- Layer 2, the metro backbone (2.5-Gb/s) layer.

In Layer 1 the access network connects all distant suburban offices, mostly with SONET/STM-4 rings. The suburban areas are usually divided into zones according to their local government structure. Figure 8 illustrates this by showing three zones. These zones, or rings, collect traffic and send it to central/tandem offices in their corresponding zones. The Layer 2 metro backbone network connects all central/tandem offices. Interzone traffic is carried at this layer. The metro backbone typically consists of a wideband Digital Access and Cross-Connect System (DACS), in this case a DACS 4/1-based mesh network.

Under the paradigm of traditional network buildout, more access rings would be added at Layer 1 and more central/tandem offices would be created at Layer 2 as the demand for bandwidth grows. Such architecture will be complicated and difficult to manage. In addition, since the DACS 4/1-based network performs network resource management at the DS1/E1 level, it is difficult to incorporate advanced services such as data and video, which use bigger transport pipes than those in DS1/E1. A superior alternative would be to create a third layer in the metro network. This layer—the *metro express layer*—would manage resources at OC-3/STM-1 or higher rates. Metro DWDM is the natural choice for this layer.

Figure 8 shows an example in which metro

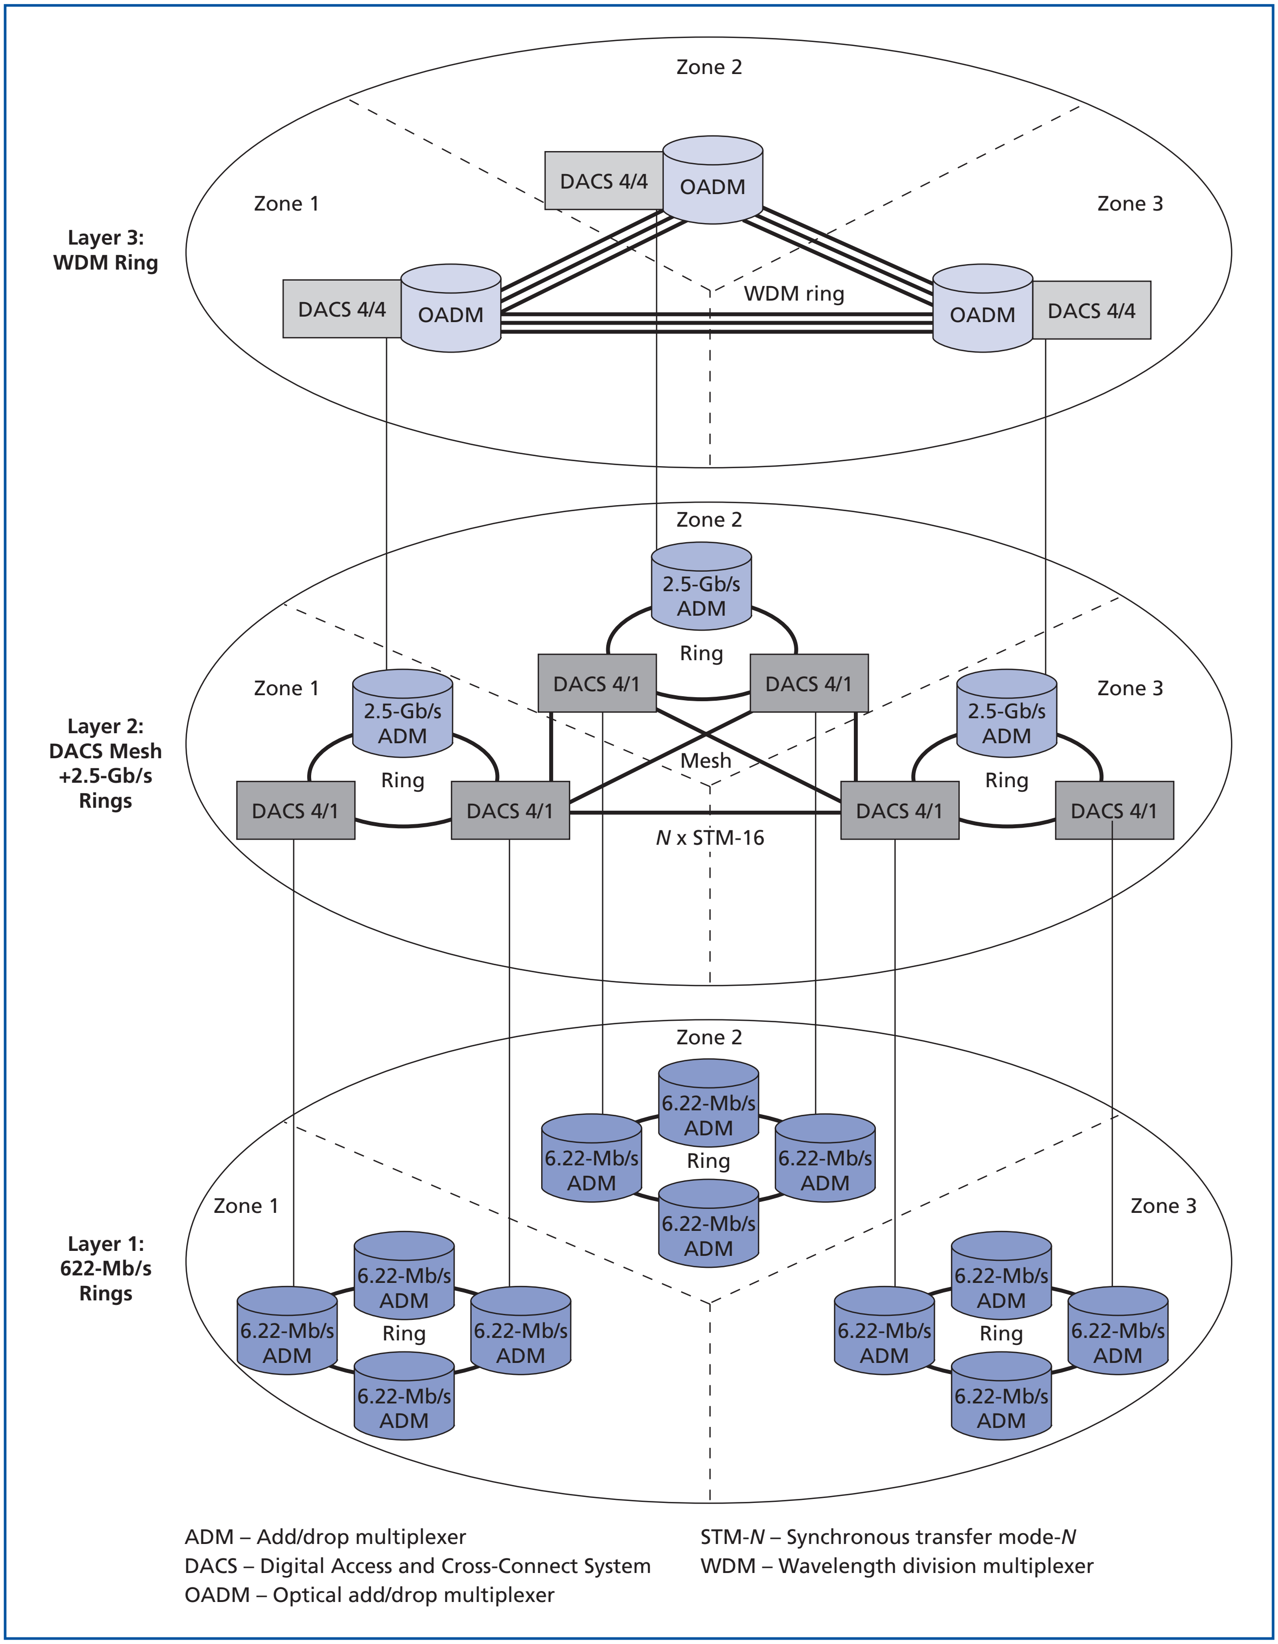

Figure 8.
A typical layered metro network architecture with DWDM.

DWDM is introduced in Layer 3. This DWDM ring will serve as the express highway for inter-zone traffic. Long distance and international gateways could also be connected to this DWDM ring. Broadband DACS (DACS 4/4) will be deployed in those central/tandem offices that form the express ring. Network resources are managed at OC-3/STM-1 in Layer 3, while network resources are managed at DS1/E1 in Layer 2. The biggest advantage of this layered architecture is that it will allow the metro network to expand to include traditional voice services and new broadband services such as high-speed data and video services.

As the network evolves, central/tandem offices at Layer 3 can be equipped with optical cross connects, or cross connects that contain an ATM and an IP fabric for direct broadband traffic processing. DWDM will be used at Layer 2 as well. The DACS 4/1 at Layer 2 will evolve into DACS 4/4/1, and the access rings at Layer 1 will be upgraded to OC-48/STM-16. Further evolution of the network will incorporate DWDM at all layers. Eventually, true all-optical networking can be performed across the entire metropolitan network. In other words, the deployment of DWDM will migrate from the initial express ring all the way to the access rings.

Benefits Summary

The metro IOF over SONET/SDH solution has significant benefits:

- It needs fewer fiber lines, which is especially important where fiber exhaustion is an issue. It can also provide additional wavelength capacity (in the form of dark wavelength channels) for revenue-generating leases in the segments with fiber exhaustion.
- It can use embedded fiber to increase the total capacity. If new fiber can be deployed, we can use the new generation broadband low-loss fiber that includes a 1.4-μm transmission window (for example, Lucent's new AllWave™ fiber).
- It permits seamless upgrades from OC-3 to OC-12 to OC-48, and beyond, using the existing equipment.
- It simplifies timing architecture by eliminating the need to time synchronize tributary channels.
- It provides clear channel service capability.
- It can provide bit rate, format, and protocol transparency, making it capable of accommodating all types of existing and emerging services.

These potential benefits explain the interest in metropolitan inter-office optical rings.

We can infer a set of features and requirements that will make metro optical networking successful. This type of networking needs to provide full mesh-type connectivity among as many as 16 nodes, implying that the ring must support wavelength reuse. A large number of wavelengths (about 40) are needed to provide multiple tributary service to each node. As an option, traffic should be protected. (If a wavelength is part of a SONET/SDH ring, it may not require any optical layer protection.) Since the traffic pattern may change in time, metro IOF must be fully flexible and able to be provisioned from a remote site. It is possible to introduce OADMs that differ in size, add/drop capacity, and cost structure into a metro optical ring, but this ring must remain competitive in cost with the SONET/SDH ring alternative.

Metro IOF Ring Node Architecture

To meet the network requirements described earlier, the central office, whether at an LSO or at the service node, must provide the following key functions:

- It must selectively drop the required wavelengths, λ_1, λ_2, …, λ_k from, and selectively add the local tributaries to, the pass-through wavelength bundle.
- It must translate the noncompliant traffic (new traffic with noncompliant wavelengths) to compliant wavelengths, λ_1, λ_2, …, λ_m before adding such traffic to the backbone ring.
- It must provide a power budget sufficient to handle both express and local traffic in order to meet the link losses in the loop.
- It must meet the specified optical signal characteristics such as maximum signal-to-noise ratio (S/N) and power uniformity.

Figure 9 depicts a functional node architecture and its major components. Next we focus on the architecture, implementation, and technologies suitable for a metropolitan central office using optical networking.

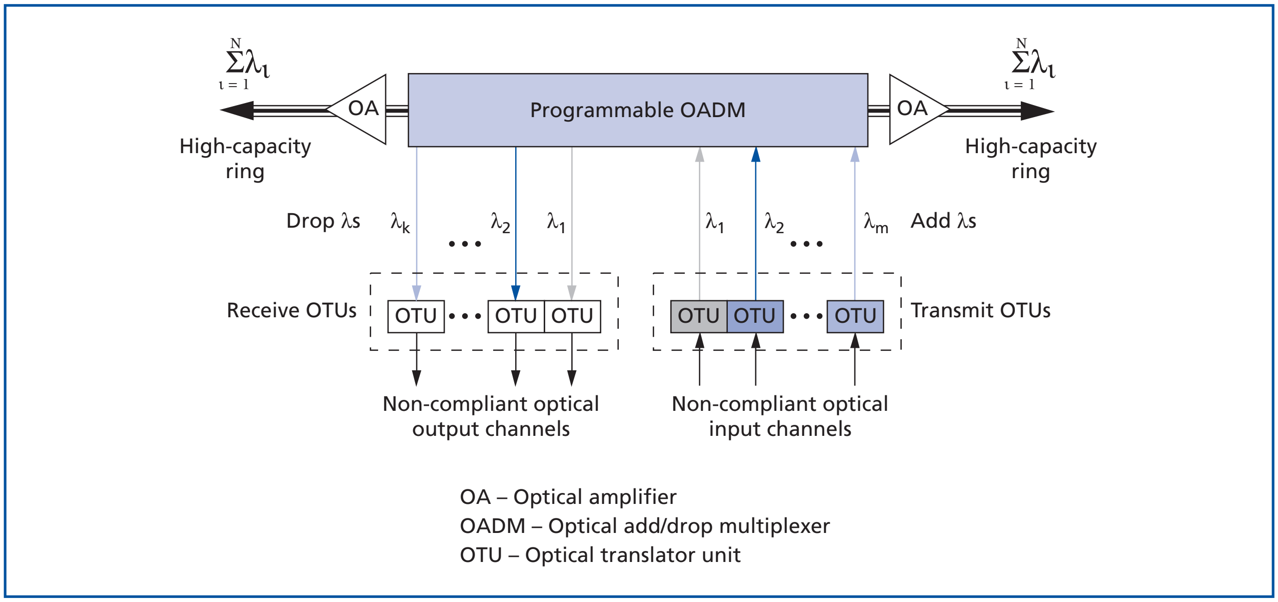

Figure 9.
Reference model for a metro optical networking node.

Building Blocks of Metro IOF Optical Networking

The main building block in metro optical networking, shown in Figure 9, is the programmable OADM. This element permits selective adding and dropping of a number of optical wavelength channels (from 1 to the maximum number of wavelengths, *N*) of any rate, format, or protocol at the central office (LSO or SN). In addition to the OADM, the other enablers of this architecture are wideband OAs and OTUs.

OTUs are used in the DWDM wavelength set to convert a client's noncompliant wavelength to a compliant one. In addition, OTUs perform the regeneration function at points in the network where the optical signal quality has deteriorated. Another important function provided by the OTUs is *overhead processing*, in which the digital overhead is extracted from the signal and processed. The OTU can readily perform overhead processing because it has access to electrical signals.

At the present time, optical networks use optoelectronic OTUs. All-optical OTUs are very promising and may replace optoelectronic OTUs in the next generation of optical networks. All-optical OTUs, however, cannot provide overhead processing at the present time, because they do not allow access to the electrical bit stream. Mikkelsen et al.[7] of Bell Labs demonstrated that the cascadability performance of bit-rate transparent all-optical interferometric wavelength translators is superior to that of optoelectronic translators.

Figure 10 shows several types of OTUs. The operation of the optoelectronic OTU is straightforward, as shown in Figure 10a. At the front end, a high-speed photodetector converts the received optical signal to an electrical one and amplifies it to the designated level. In general, the detectors are broadband, making them identical for all wavelengths (at their input sites).

The detected signal is then reshaped, regenerated, and (optionally) retimed by the electronic circuitry. Clearly, this electronic processing limits the operation of the OTU to specific bit rates and formats. Bit-rate and format-transparent OTUs (within a limited range of rates, such as 45 to 700 Mb/s) have been developed by eliminating the retiming circuitry from the unit. As a regenerator, an OTU performs clock recovery for retiming, which cascades the signal through many network elements in the optical network. The electronic signal then modulates a laser with the desired output wavelength, λ_{out}, using either an internal modulator or an external modulator (such as a Mach-Zehnder modulator). As a result, the output optical signal is a "cleaned" (regenerated) replica of the input signal, but at the desired wavelength.

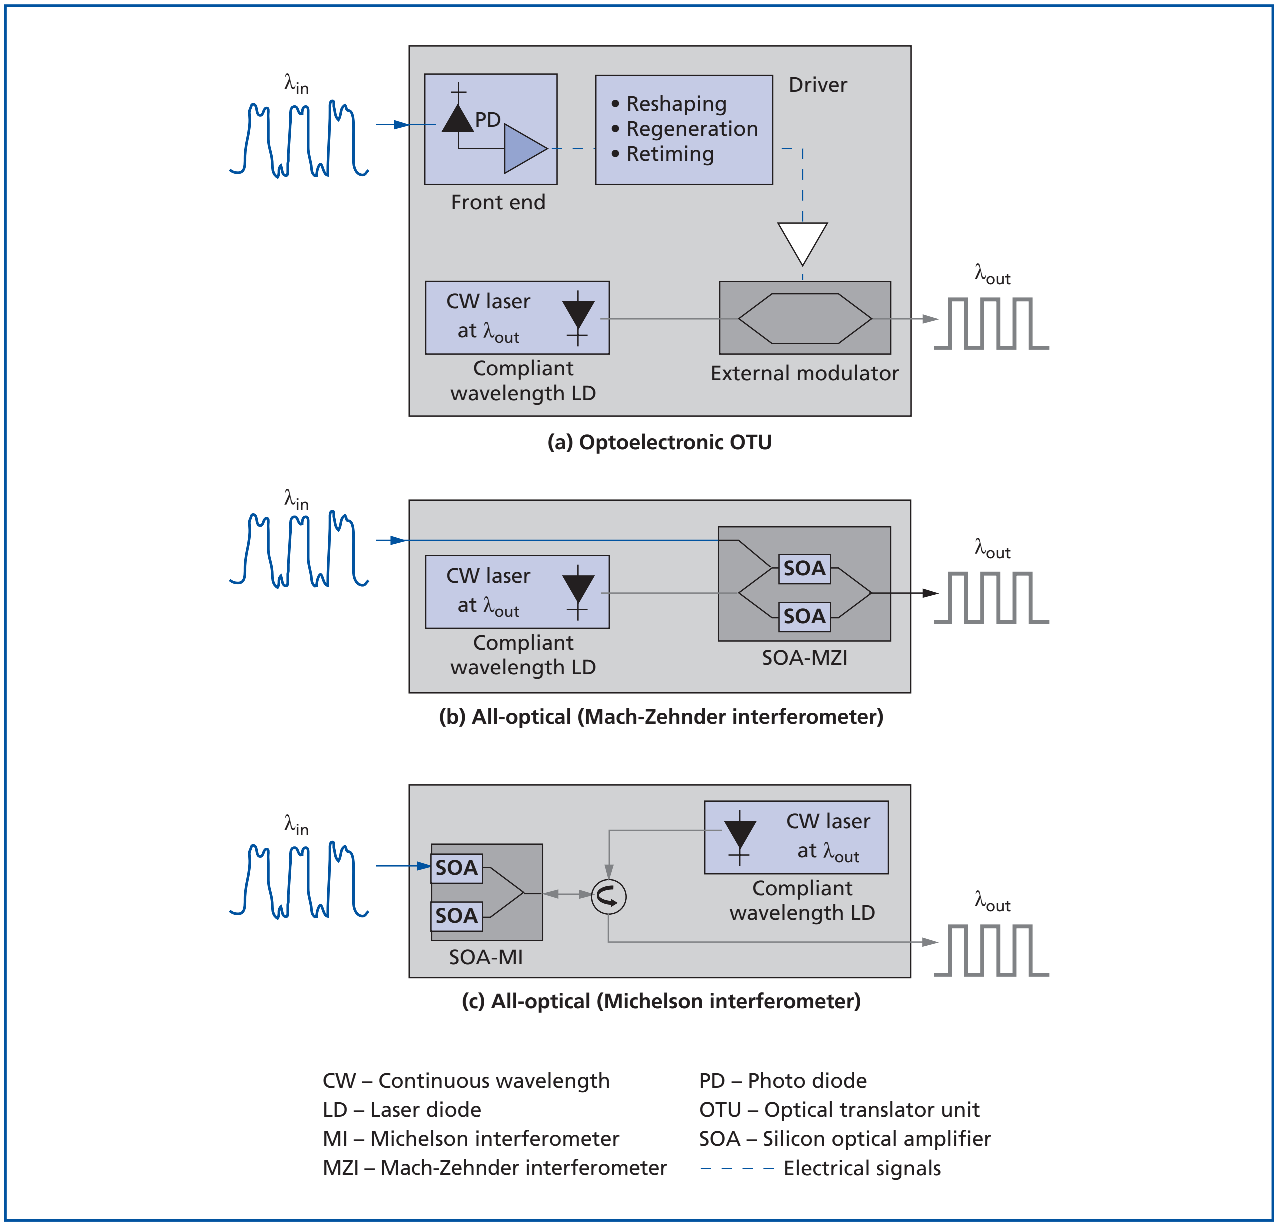

Figure 10.
Optical translator units.

The most promising bit-rate-transparent all-optical translators are based on cross-phase modulation (XPM) in silicon optical amplifiers (SOAs), integrated into either a Mach-Zehnder interferometer[8] (MZI) or a Michelson interferometer (MI).[9] Figure 10b shows a schematic of polarization-insensitive SOA-MZI, and Figure 10c shows a schematic of SOA-MI wavelength translators. In these interferometric XPM devices, the unmodulated, or continuous, wave (CW) light from a laser operating at the target wavelength is split into two paths containing SOAs. A relative phase shift is induced by the input optical signal entering one of the SOAs, which saturates the gain. When the light is coherently recombined, constructive or destructive interference will occur, depending on the phase difference between the two paths. The unperturbed state of the interferometer can be set up for constructive or destructive interference, causing an input signal injected into one path to either decrease (invert) or increase (non-invert), respectively, the output converted signal.[10] As a result of gain saturation, optical waveform shaping is inherent, as shown by the waveforms in Figure 10.

Hybrid integration of SOAs and couplers or monolithic integration can create structures analogous to Mach-Zehnder and Michelson interferometers. Early experiments used hybrid devices consisting of discrete SOAs and fiber couplers; however, the need to maintain relative path lengths to within a fraction of a wavelength makes these devices sensitive to environmental conditions such as temperature and vibration. Monolithic integration offers considerable advantages in stability and compactness.

Optical Amplifiers

The advent of wideband fiber OAs has played a major role in enabling the practical implementation of metro IOF optical networks. Since most wavelength add/drop elements are passive and inherently lossy, the OA must compensate for outside plant link losses and also provide additional power gain to meet the losses associated with these devices. In addition to providing the gain for the express (pass-through) wavelengths, in some architectures the OAs must amplify the "add" and "drop" channels as well. The OA is a costly component and must be used sparingly in the metro network. Any OA used in the metro ring must also possess the following key characteristics:

- *Noise figure*. The OA must maintain a low noise level, because some channels may pass through a cascade of all the amplifiers in the loop, and all the noise and impairments accumulate.
- *Bandwidth*. The amplifier bandwidth must be wide enough to cover the entire set of wavelength channels (N) in the DWDM spectrum.
- *Gain flatness*. In addition to providing sufficient gain, the OA must have gain uniformity for all channels, from one end of the spectrum to the other.
- *Per-channel automatic gain control (AGC)*. Even if the incoming multi-wavelength optical signals are not uniform in power, either the OA (using its internal components) or other external devices associated with the OADM—such as attenuator arrays—should equalize the outgoing wavelength channels in the bundle.
- *High-speed AGC*. One of the cumbersome scenarios resulting from a reconfiguration or a link/component failure is the sudden loss or appearance of a number of (a subset of the total number, N) channels. The amplifier must quickly respond to such total power changes and re-equalize the per-channel gains, or the perturbation will propagate through the entire network and affect all the other channels.

In this issue of the *Bell Labs Technical Journal*, Sun et al.[11] provide a complete discussion of wideband fiber-optical amplifier technologies and their role in DWDM optical networks.

OADMs

Flexible (or reconfigurable) OADMs are the most critical enablers of metropolitan optical networking. The term "flexible OADM," when applied to a metro IOF optical ring, implies that the system has the ability to provide a connection between any two nodes for all wavelength channels, up to the limit of the available channels, under remote software control. Only with careful technology selection and OADM architectural design could the requirements for metro IOF be met in a cost-effective manner.

For completeness, we describe a number of technological and architectural options for implementing metro OADMs. In this issue of the *Journal*, Giles[12] presents a complete discussion and classification of various technologies, philosophies, and techniques for implementing OADMs.

The most straightforward configuration for implementing a fixed (preset and not flexible) OADM is by using back-to-back multiplexing, as shown in **Figure 11**. In this configuration, all the N wavelength channels in the multi-wavelength optical signal are demultiplexed into their wavelength tributaries. The tributary signals are then remultiplexed to form the outgoing multi-wavelength signal. Various multiplexer/demultiplexer devices are commercially available for this purpose. By using these two devices, one can access the desired tributaries and add/drop any wavelength, as depicted in Figure 11. In a fixed OADM, optical jumpers connect the drop (add) wave-

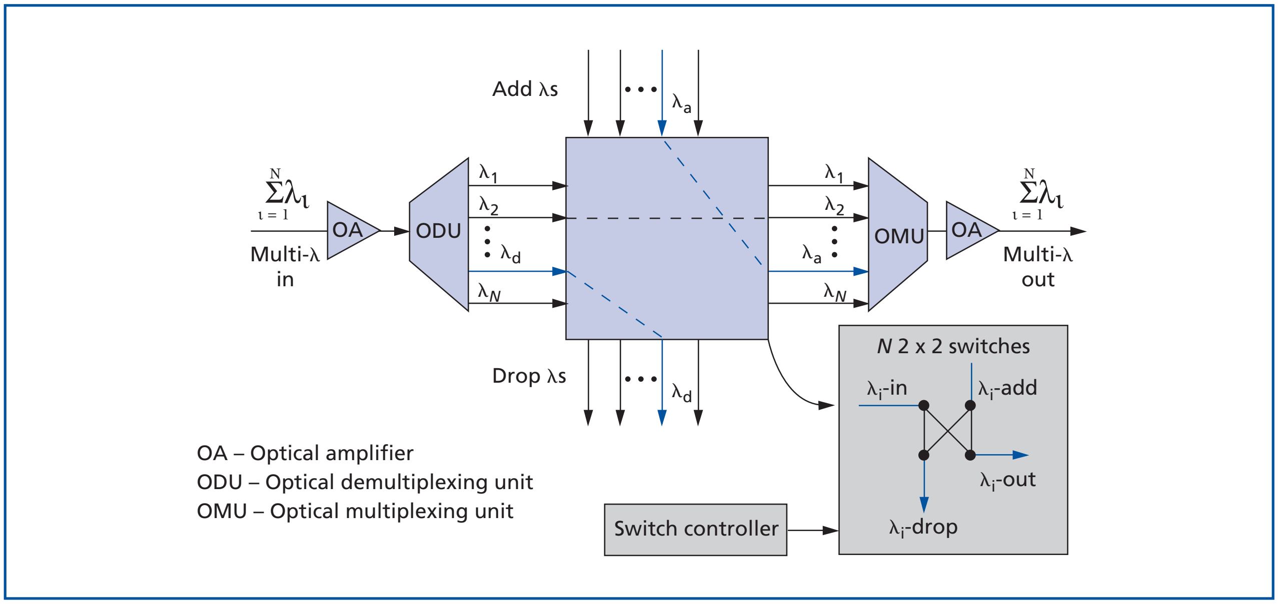

Figure 11.
OADM architecture based on back-to-back multiplexing.

lengths from the demultiplexer (multiplexer) to the local receive (transmit) OTUs. This OADM structure can be converted to a flexible OADM by employing *N* fiber-optical (mechanical, thermal, or otherwise) switches to selectively pass through or add/drop each wavelength, as shown by the box in the lower part of Figure 11.

In addition to multiplexing/demultiplexing losses that need to be compensated for by external optical amplifiers shown in this architecture, other impairments may deteriorate the signal if the express channels are demultiplexed and recombined at every LSO. The design must particularly consider impairments such as filter bandwidth narrowing and multi-path interference. A more efficient method for implementing OADM entails accessing only the add/drop channels without significantly affecting the express wavelengths (except for added loss). This requires tunable/switchable filters that only demultiplex the local (add/drop) channels. A fixed OADM based on a multi-layer interference filter[13] can demultiplex only the add/drop channels without demultiplexing the pass-through wavelengths.

Figure 12 depicts loss factors in OADM. In 1995 Giles and Mizrahi[14] reported a low-loss flexible (programmable) OADM. This structure, depicted in Figure 12a, uses a number of tunable (or switchable) fiber Bragg gratings (FBGs) sandwiched between two optical circulators. The FBGs reflect the add/drop wavelengths and pass the express wavelengths with little impairment. The reflected (drop) wavelengths are dropped on the port indicated in Figure 12a and must be separated (demultiplexed) to feed the receive OTUs. The "add" wavelengths are combined (multiplexed) and inserted into the "add" port. Several variations of this basic architecture have been proposed for OADM and for optical cross-connect applications[15] that use series or parallel FBGs.

Figure 12b shows a loss-compensated OADM composed of appropriate lengths, l_1 and l_2, of optical rare earth (erbium)-doped fiber (EDF), sandwiched between two optical circulators. The FBGs are incorporated (or spliced) between the rare earth-doped fiber segments. At least one grating is used for each wavelength (of the set of wavelengths in the DWDM aggregate) that needs to be added or dropped. By pumping the EDF segments (either through the "add" port or through additional wavelength-selective couplers) by using a laser source operating at an appropriate wavelength (such as 1480 nm or 980 nm), the EDFs act as multiwave optical amplifiers, compensating simultaneously for losses associated with the "express," "drop," and "add" wavelengths. The relative gains in the EDF segments are design parameters. This architecture

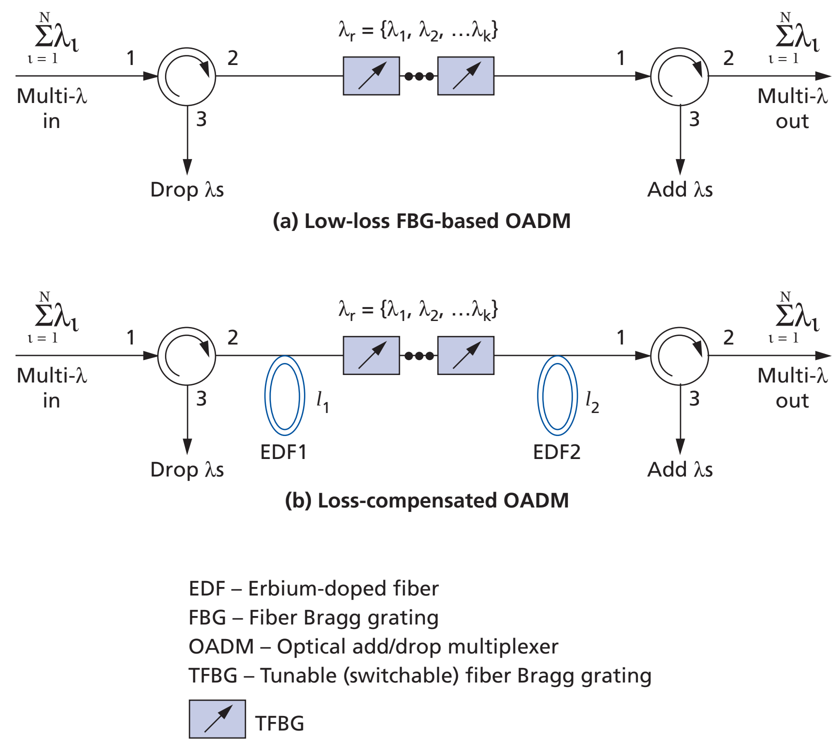

Figure 12.
Loss factors in OADMs.

could eliminate the need for external fiber amplifiers surrounding the OADM. This OADM structure is lossless, (re)configurable, and flexible, and it can be concatenated (with minimal optical impairments for "through" wavelengths). The internal distributed amplifier compensates for the losses experienced in the through (express), add, and drop wavelengths. Since all three paths experience gain, passive trees can replace $1 \times N$ passive multiplexers and demultiplexers in the "add" or "drop" legs. This OADM architecture could lead to an architecture that is expandable, upgradable while in service, and less costly at startup.

Figure 13 depicts the simplest structure for an OADM. It uses one (or more) passive coupler(s) that drop a fraction of the total (aggregate) signal at the splitting port or insert a portion of the "add" signals at the combiner port. This low-cost device has several drawbacks, the main one being that it does not permit wavelength reuse. As a result it does not satisfy the requirements stated for metro IOF optical rings. Its simplicity and low cost, however, may make it suitable

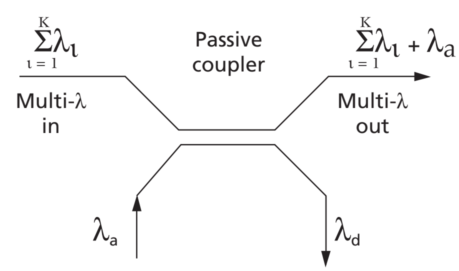

Figure 13.
A passive-coupler-based OADM.

for metro access rings, described in the final section of this article. Clearly, the dropped wavelengths must be separated by adding optical circuitry.

Doerr of Bell Labs[16] proposes a novel OADM structure—shown in **Figure 14**—which can lead to a low-cost integrated device and may find applications in metro access as well as in metro IOF. It can function as a wavelength ADM or a two-line wavelength cross connect. This OADM consists of two waveguide grating routers (WGRs) that are interleave-chirped and

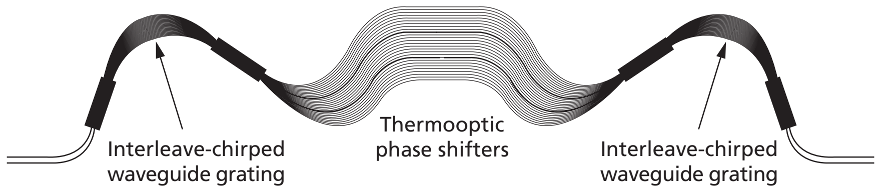

Figure 14.
A waveguide layout of an integrated OADM in silica.

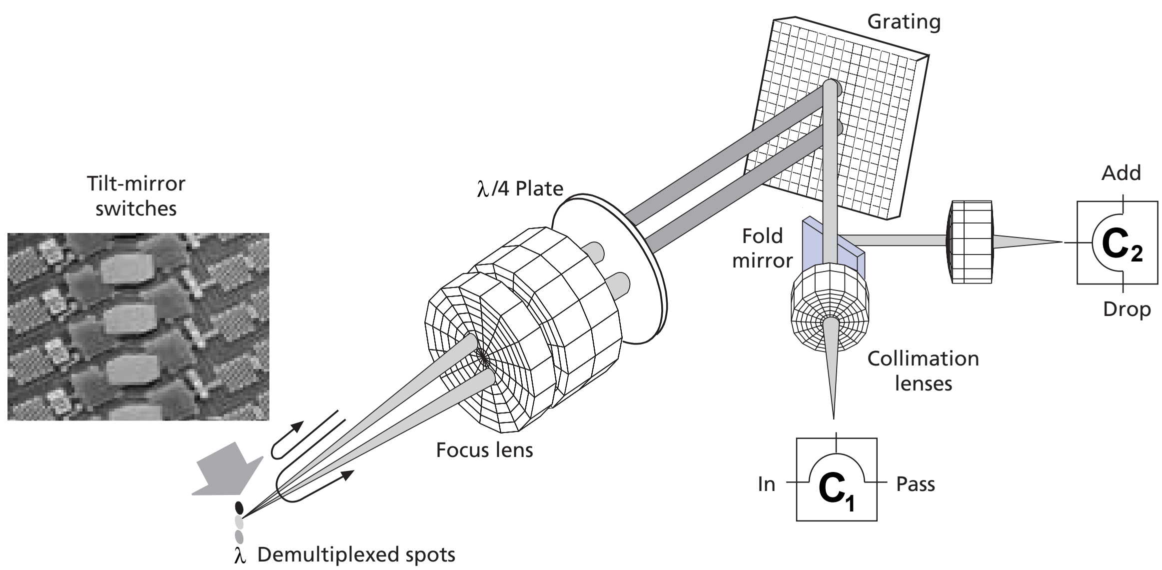

Figure 15.
An OADM based on spatial optics and micromechanical tilting mirrors.

connected by a waveguide lens consisting of equal-length waveguides with phase shifters.[16]

The entire OADM can be fabricated as an integrated device. Figure 14 shows a waveguide layout of one way to make an integrated add/drop multiplexer in silica. Each channel has three thermooptic phase shifters whose settings determine whether that channel is in a bar or a cross state. The main advantages of such a design are low loss (Doerr et al.[17] have demonstrated 6.6-7.6 insertion loss for a 16-channel version) and compactness (at least four such 16-channel devices can fit on a 5-inch wafer).

All the OADM structures described above are based on guided wave optics, where fibers or fiber-optic devices guide the multi-wavelength signals as well as the wavelength tributaries. It is also possible to design OADM systems based on "spatial optics," which incorporate bulk optical components and use space as their optical interconnect. The spatial interconnect adds an additional dimension for manipulating signals, which can be beneficial.

For example, Ford[18] et al. of Bell Labs report a promising OADM, which is the spatial counterpart of the back-to-back demultiplexing structure shown in Figure 11. Free-space imaging performs the wavelength demultiplexing, remultiplexing, and switching in this structure through a planar diffraction grating, as shown in **Figure 15**. After passing through an input circulator, the input signal from the fiber is collimated, then diffracted by a 600-line/mm grating and projected onto a micromirror array. The micromechanical mirrors (shown on the left of Figure 15) tilt in response to applied electrostatic forces. Depending on the angle of the mirror, one of two scenarios occurs. The light is either retro-reflected into the original input port, or it is tilted by about 9 degrees and carried into a second output port, where the desired wavelengths are added/dropped by the second circulator. A quarter-wave plate between the lens and grating rotates the reflected light polarization to compensate for any grating polarization dependence. This device is packaged in a $5 \times 9 \times 19$-cm aluminum housing. The total fiber-to-fiber insertion loss is 5 dB for the passed signals and 8 dB for the dropped or added signals, with an 0.2-dB polarization dependence. The 16-mirror switch device was fabricated by the Microelectronics Center of North Carolina (MCNC)[19] through the Multi-User MEMS Process (MUMPS) foundry. The switching contrast was better than 30 dB for all input and output states, with a 20-µs switching time.

Metro Access Rings

The metro access ring can be extremely useful for enhancing the services offered in the local loop. Consider an application where a service provider offers a fiber-based access service. Such broadband capabilities are expected to be increasingly popular as the bandwidth requirements of enterprise and business customers increase. Today, a bit-rate transparent "buy-a-lambda" service can be sold for profit to an ISP.

A service provider can use the metro access ring to support multiple customers. The DWDM advantage is that a customer, or group of customers, can have dedicated wavelengths consistent with both new and old services. **Figure 16** shows a core network with a metro IOF system deployed. At two of its nodes, it interfaces with an access ring. The nodes can provide a form of optical dual-node interworking (DNI) to ensure that the access ring can survive nodal failures (dual homing).

The access ring shown in Figure 16 supports separate wavelengths to:

- TDM equipment,
- ATM switches, and
- A next-generation "data-centric" access multiplexer, which acts as a statistical vehicle offering services through IP or ATM interfaces.

A metro access ring can be deployed using a "right-sizing" philosophy. It is possible to define a small shelf, a "module," or a plug-in to an existing network element that has as its key function the ability to extract a specified channel from a set of multiple channels defined on a fiber using DWDM. This right-sizing approach makes it possible to introduce a family of optical ring terminals, each optimized for a different capacity of wavelength add/drop. In Figure 16, the hub nodes have add/drop requirements that are very different from those of the nodes at remote locations. Optical add/drop multiplexers at the access nodes can be designed around specific sizes and can be optimized by selecting the appropriate technology.

Metro Access Ring Characteristics

The requirements for the DWDM metro access optical ring—as opposed to those of the metro IOF—have been relaxed to permit low-cost implementation. The DWDM metro access optical ring has the following key attributes:

- It provides only hub-type connectivity among a small number of nodes (that is, from two to six nodes).
- A limited number of wavelength add/drops (1-8) suffice at the customer location.
- Wavelength reuse is not required.
- The traffic must be optionally 1+1 protected.

The above attributes translate into these initial key requirements for an access OADM:

- It must operate in a customer premises (non-central office) environment, which requires:
  - High reliability;
  - Low, simplified maintenance;
  - A single module (shelf) OADM or a stand-alone box;
  - A minimum number of active components;
  - Wide tolerance of variations in environmental conditions; and
  - Low power consumption.
- The add/drop capacity must be kept low to minimize cost—that is, 1-~8 wavelength

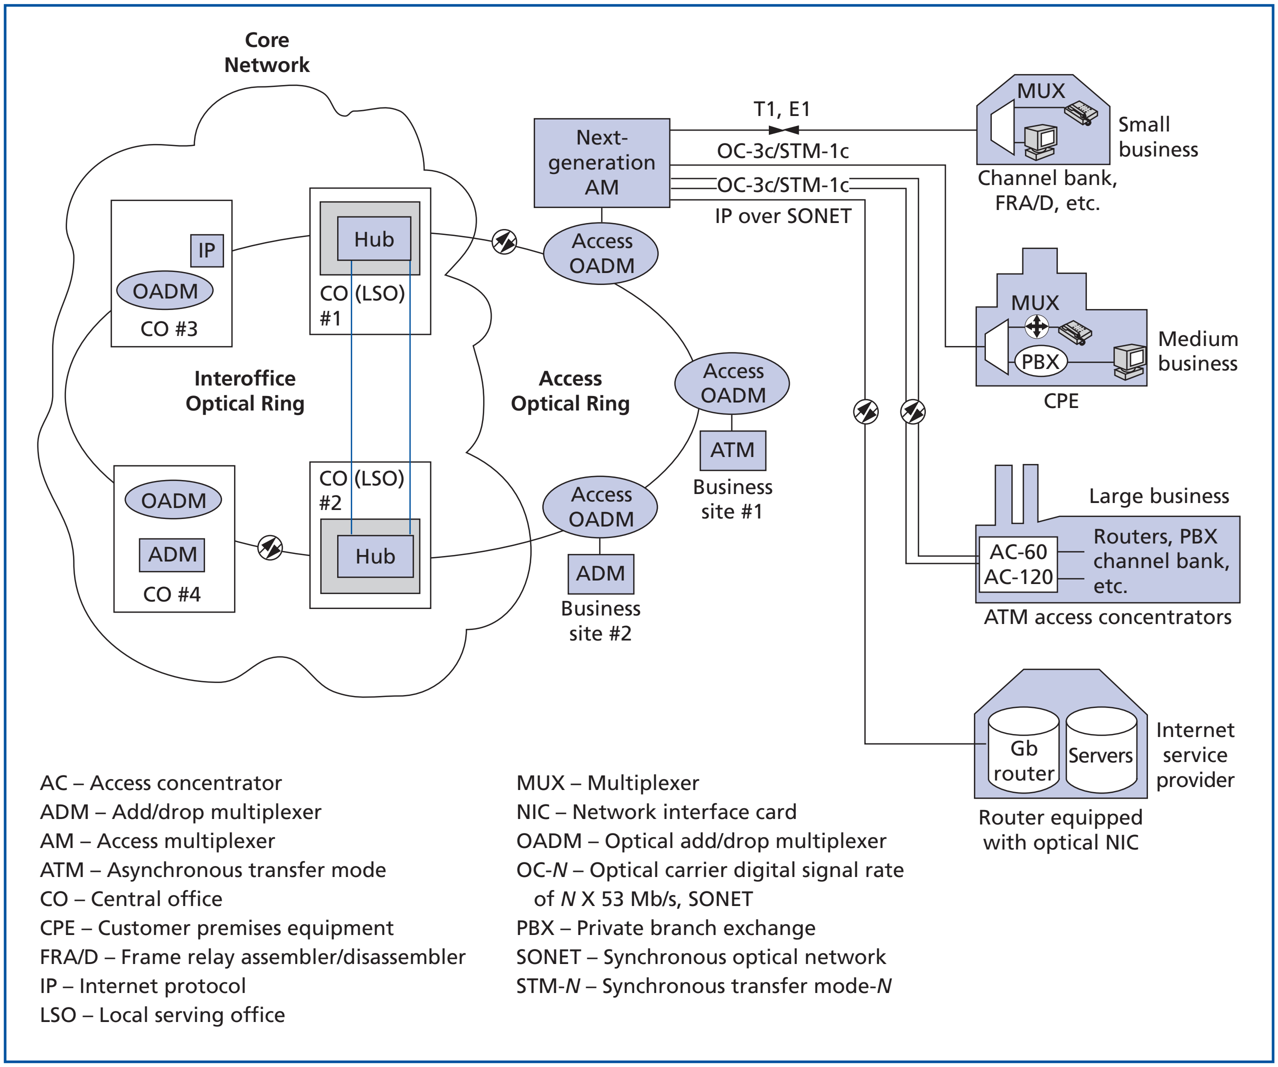

Figure 16.
Metro access rings and interoffice networking.

channels—or as many as can be dropped in a single shelf. The exact number is an architectural decision.

- It can have "limited" flexibility. Plug-in modules can provision or rearrange the assigned "add" and "drop" wavelengths.
- It supports shorter span lengths than the metro IOF.
- It supports all types of services (bit rates, formats, and protocols).
- It must be compatible with the metro IOF product.
- It supports mainly "hubbed" traffic between the LSO(s) and the customer premises equipment.

The last item implies that there is no requirement for efficient support of meshed traffic patterns (between customer locations). When this factor is combined with the fact that the number of wavelengths needed to serve all the customers in a typical access ring is small, there is no advantage to wavelength reuse. Simplification of the OADM design can lead to significant cost saving. For example, a simple passive coupler followed by appropriate wavelength-selective filters can serve as an OADM. The key requirement for a metro access OADM is low cost. Specifically, it needs to be competitive on a per-customer-location basis with the cost of installing TDM electronics. Currently, Lucent Technologies is conducting research to move DWDM optical networking to the access loop.

Summary

In this paper we described three key drivers for metro optical networking—how to achieve transport efficiencies, eliminate high-speed TDM, and reduce the cost of fiber utilization. The fundamental vehicle for metro optical networking is the OADM, introduced from an applications perspective. We reviewed the DWDM applications in local connectivity (introducing the "path-in-lambda" ring topology) and surveyed self-healing schemes for optical rings, including 1+1 and optical shared protection. The data applications we discussed included router connectivity (and ATM switch elimination) and the access application. We introduced the concept of "right-sizing" an OADM terminal and defined it in this context. In a review of the optical elements used to build OADM network elements, we described OTUs, or transponders, optical amplifiers, and add/drop multiplexer primitives, all of which can be used to build an OADM. Metro optical networking represents a unique opportunity for a service provider to begin deploying a data-centric high-bandwidth services infrastructure.

Acknowledgments

Defining metro optical networking at Lucent Technologies has been a largely cooperative effort involving multiple organizations, from whose work we have borrowed substantially to write this paper. We thank C. R. Doerr, J. E. Ford, and the many indirect contributors to this paper, too numerous to be acknowledged here.

References

1. R. Alferness, P. A. Bonenfant, C. Newton, K. Sparks, and E. Varma, "A Practical Vision for Optical Transport Networking," *Bell Labs Tech. J.*, Vol. 4, No. 1, Jan.–Mar. 1999, pp. 3–18.
2. D. Y. Al-Salameh, M. T. Fatehi, W. J. Gartner, S. Lumish, B. L. Nelson, and K. K. Raychaudhuri, "Optical Networking," *Bell Labs Tech. J.*, Vol. 3, No. 1, Jan.-Mar. 1998, pp. 39–61.
3. T. E. Adams, B. G. Cortez, M. T. Fatehi, K. K. Raychaudhuri, B. L. Nelson, and D. Y. Al-Salameh, "Metropolitan Area Optical Networking," *Proc. of the Nat'l. Fiber-Optic Eng. Conf. (NFOEC) '98*, Orlando, Fla., Sept. 13-17, 1998, Vol. II, pp. 301–309.
4. Telcordia, "Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria," Bellcore GR-253, Dec. 1993. http://telecom-info.bellcore.com/site-cgi/ido/index.html
5. International Telecommunication Union-Telecommunication Standardization Sector, "Types and Characteristics of SDH Network Protection Architectures," ITU-T G.841, July 1995. http://www.itu.int/publications/telecom.htm
6. Telcordia, "SONET Bidirectional Line Switched Ring Equipment Generic Criteria," Bellcore GR-1230, Dec. 1996. http://telecom-info.bellcore.com/site-cgi/ido/index.html
7. B. Mikkelsen, G. Raybon, T. N. Nielsen, U. Koren, B. I. Miller, and K. Dreyer " Opto-electronic and all-optical wavelength translators and their cascadability," *Proc. of the OFC/IOOC '99*, invited paper FJ1, San Diego, Feb. 1999, pp. 146–148.
8. B. Mikkelsen et al., "Use of Semiconductor Optical Amplifiers for High-Speed Optical Signal Processing," *Dig. of OAA '98*, Paper WA1, Victoria, Canada, July 1997, pp. 190–193.
9. L. H. Spiekman, J. M. Wiesenfeld, U. Koren, B. I. Miller, and M. D. Chien, "All-Optical Mach-Zehnder Wavelength Converter with Monolithically Integrated Preamplifiers," *Photon. Technol. Lett.*, Vol. 10, No. 8, 1998, pp. 1115–1117.
10. D. Nesset, T. Kelly, and D. Marcenac, " All-Optical Wavelength Conversion Using SOA Nonlinearities," *IEEE Commun. Mag*. Vol. 36, No. 12, Dec. 1998, pp.56–61.
11. Y. Sun, A. Srivastava, J. Zhou, and J. Sulhoff, "Optical Fiber Amplifiers for WDM Optical Networks," *Bell Labs Tech. J.*, Vol. 4, No. 1, Jan.–Mar. 1999, pp. 187–206.
12. C. R. Giles and M. Spector, "The Wavelength Add/Drop Multiplexer for Lightwave Communication Networks," *Bell Labs Tech. J.*, Vol. 4, No. 1, Jan.–Mar. 1999, pp. 207-229.
13. N. A. Jackman, "Multilayer Optical Filters with Random Errors," *Bell Labs. Tech. J.*, Vol. 3, No. 1, Jan.-Mar. 1998, pp. 112-121.
14. C. R. Giles and V. Mizrahi, "Low-Loss Add/Drop Multiplexers for WDM Lightwave Networks," *Proc. of the10th Int'l Conf. on Integrated Optics and Optical Fibre Commun. (IOOC '95)*, Hong Kong, June 26–30, 1995, Paper ThC2-1, pp. 66–67.

15. Y. K. Chen and C. C. Lee, "Fiber Bragg Grating-Based Large Nonblocking Multiwavelength Cross-Connects," *J. Lightwave Tech.*, Vol. 16, No. 10, Oct. 1998, pp. 1746-1756.
16. C. R. Doerr, "Proposed WDM Cross Connect Using a Planar Arrangement of Waveguide Grating Routers and Phase Shifters," *IEEE Photon. Technol. Lett.*, Vol. 10, Apr. 1998, pp. 528-530.
17. C. R. Doerr, L. W. Stulz, J. Gates, M. Cappuzzo, E. Laskowski, L. Gomez, A. Paunescu, A. White, and C. Narayanan, "Arrayed Waveguide Lens Wavelength Add-Drop in Silica," *IEEE Photon. Technol. Lett.*, Vol. 11, No. 5, May 1999.
18. J. E. Ford, V. A. Aksyuk, D. J. Bishop, and J. A. Walker, "Wavelength Add/Drop Switching Using Tilting Micromirrors," *J. Lightwave Technol.*, Vol. 17, No. 5, May 1999.
19. The Microelectronics Center of North Carolina, http://mems.mcnc.org/

(Manuscript approved March 1999)

YI CHEN, a technical manager in the Optical Networking Group in Holmdel, New Jersey, is in charge of product marketing and development in the Asia Pacific region and China. She is responsible for systems engineering and network planning, focusing on the China market and DWDM products. Ms. Chen received a B.S. in physics from Peking University in China, after which she earned an M.S. in electrical engineering and a Ph.D. in applied physics, both from the University of Michigan in Ann Arbor.

MOHAMMAD T. FATEHI is a distinguished member of technical staff in the Product Evolution Planning Department, part of Lucent's Optical Networking Group in Holmdel, New Jersey. He received a B.E. degree from The American University of Beirut in Lebanon, after which he earned M.S. and Ph.D. degrees from the Ohio State University in Columbus, all in electrical engineering. Since joining Bell Labs, Dr. Fatehi has focused on the vision of optical networking and on optical network architectures and technologies. Currently, he is leading the Metro Optical Networking Product Evolution Planning team. Dr. Fatehi holds more than a dozen patents, and he has authored numerous technical papers on related areas of research. Formerly a professor at Sharif University of Technology in Tehran, Iran, and at the Ohio State University, Dr. Fatehi is also a senior member of the IEEE.

HUMBERTO J. LA ROCHE, a manager in the Product Management Department of Lucent's Optical Networking Group in Holmdel, New Jersey, received an undergraduate degree in physics from the Universidad Simon Bolivar in Caracas, Venezuela, and a Ph.D. in physics from the University of Texas at Austin. Dr. La Roche is responsible for product and feature planning for optical systems.

JACOB Z. LARSEN, a business planning manager in Lucent's Optical Networking Group in Holmdel, New Jersey, is responsible for the market assessment and investment planning of Lucent's next-generation optical products. Mr. Larsen earned a B.A. in international business from The Copenhagen Business School in Denmark and an M.B.A. from The Amos Tuck School of Business at Dartmouth College in Hanover, New Hampshire.

BRUCE L. NELSON is director of Product Evolution Planning and Portfolio Management, part of Lucent's Optical Networking Group in Holmdel, New Jersey. He received a B.A. from the University of Tennessee in Knoxville, and M.A. and Ph.D. degrees from the University of Texas in Austin, all in physics. ◆