

## ◆ Optical Networking

*Daniel Y. Al-Salameh, Mohammad T. Fatehi, William J. Gartner, Stan Lumish, Bruce L. Nelson, and Kamal K. Raychaudhuri*

*Optical networking, as embodied by dense wavelength division multiplexing (DWDM) and optical amplification, has revolutionized long-distance transport and has resulted in capacity expansion, cost reduction, and operations simplification. Extension to more wavelengths, emergence of reconfigurable wavelength add/drop and other optical network elements (NEs), as well as cost reduction of optical amplifiers and laser devices, are pushing the revolution into the local and metropolitan networks. Traffic generated in any format, such as SONET/SDH, asynchronous transfer mode (ATM), and plesiochronous digital hierarchy (PDH), with bit rates ranging from 45 Mb/s to 10 Gb/s, can be economically transported and routed optically without resorting to time division multiplexing (TDM). In this paper, we describe the latest advances in optical NEs and applications in the public and private networks. In addition, we discuss how optical networking creates flexibility in sharing existing fiber inventory between service providers. We present an example of how optical networking is capable of meeting the emerging demand for high-capacity clear-channel end-user services, a demand that TDM systems cannot meet. Finally, we discuss future directions and optical layer restoration as a prime target of network evolution.*

### Introduction

It comes as no surprise that Internet traffic, high-capacity data services, and multimedia services have driven the need for network capacity at unprecedented rates. Service providers report that the demand for network capacity is growing in excess of a 50% compounded annual rate at a time when fiber use rates average 60% or higher.<sup>1</sup> These two trends—high demand coupled with high usage rates—create a capacity crisis for many service providers.

Carriers have three options for relieving the capacity shortage problem:

- *Install new fiber.* This is a perfectly adequate solution for new service providers or for those who have plans to grow their embedded base substantially. However, it is a costly proposition for many service providers, and right-of-way costs may be prohibitive. In addition, customers who have an embedded base of

fiber are seeking ways to exploit fully the inherent capacity of their existing fiber before installing new fiber.

- *Increase the bit rate of the transmission system.* The traditional approach for satisfying increased capacity demand has been to increase the time division multiplexed (TDM) bit rate. Today, the majority of synchronous optical network (SONET) transport deployment is OC-48 (2.5-Gb/s) systems. As 10-Gb/s and higher rate systems become available, inherent fiber and system limitations will constrain the application domains severely, making this an unacceptable long-term strategy.
- *Employ wavelength division multiplexing (WDM).* An alternative approach, WDM allows a customer to combine many wavelengths onto a single fiber. Dense wavelength division multi-

plexing (DWDM) implies closely spaced wavelengths (for example, 0.8-nm spacing).

**Figure 1** shows trends in single-channel (TDM) and multichannel (WDM) systems. As the TDM data show, an asymptote is being approached in the 10-Gb/s to 40-Gb/s range, suggesting a fundamental limit in the ability of TDM-based solutions to deliver capacity consistent with recent demand. The data also show that progress in WDM technology is delivering spectacular growth in system capacity. This suggests that service providers will require some combination of TDM and WDM in their networks to continue growing capacity.

Key figures of merit for DWDM systems include the number of wavelengths (that is, wavelength channels) carried, span distances (that is, the distance between repeaters), and the number of spans allowed before regeneration. For example, an application might require a system to support eight wavelengths over three 120-km spans. Bell Labs researchers have demonstrated a 32-wavelength system, with each wavelength carrying 10 Gb/s over a distance of 640 km. Currently, Lucent Technologies is developing an 80-channel DWDM system as part of the WaveStar™ product line.

Significant deployment of DWDM systems is occurring in long-distance and local markets. As carriers begin deploying more wavelengths in their network, there will be an increased need to manage capacity at the wavelength (that is, optical) layer, avoiding costly conversions between optical and electrical signals. Although the principal customer need being satisfied today is capacity gain, it is our view that an all-optical layer will emerge in the network, where traditional network functions including add/drop multiplexing (ADM), cross-connection, signal restoration, and service deployment will be performed optically. As **Figure 2** shows, this layer will arise due to scale benefits that will exist in managing capacity at the optical level rather than at the electrical level.

Recent advances in optical amplifiers, filters, lasers, and fiber have contributed to the wide-scale deployment of DWDM systems. Continued advances in these areas will be necessary to realize the vision of an all-optical network. This paper describes these

### Panel 1. Abbreviations, Acronyms, and Terms

$\lambda$  A/D—wavelength add/drop  
ADM—add/drop multiplexing  
ATM—asynchronous transfer mode  
CATV—cable television  
DCC—data communications channel  
DWDM—dense wavelength division multiplexing or dense wavelength division multiplexer  
EDFA—erbium-doped fiber amplifier  
LAN—local area network  
NE—network element  
OADM—optical add/drop multiplexer  
OA—optical amplifier  
OCDM—optical code division multiplexing  
OC—optical carrier  
ODU—optical demultiplexer unit  
OFA—optical fiber amplifier  
OLS—Optical Line System  
OMU—optical multiplexer unit  
OSNR—optical signal-to-noise ratio  
OTDM—optical time division multiplexing  
OT—optical translators  
OXC—optical cross-connect  
PDH—plesiochronous digital hierarchy  
PMD—polarization mode dispersion  
POTS—"plain old telephone service"  
S/N—signal-to-noise  
SDH—synchronous digital hierarchy  
SLM—spatial light modulator  
SONET—synchronous optical network  
TDM—time division multiplexing  
WDM—wavelength division multiplexing or wavelength division multiplexer

enabling technologies, as well as their current and required capabilities.

Carriers who adopt the aforementioned advances in optical technology will gain advantages in terms of first cost, operations cost, and flexibility. Clear and compelling cost advantages exist for the deployment of DWDM technology in long-distance applications. Moreover, DWDM increasingly is becoming a viable and attractive option for metropolitan applications. This paper discusses the benefits of optical layer architectures and presents a variety of applications for optical networking.

As wavelengths become the medium of exchange in networks, it will be natural for carriers to demand

more wavelengths per fiber and increased flexibility in managing wavelengths. Satisfying these demands will result in the emergence of a new class of network elements (NEs) that are all optical, including optical cross connects, optical ADMs, and ultra high-capacity optical line systems. This paper looks at some future directions in optical networking, and it concludes by summarizing the future role of the optical layer and the potential roles of its constituent elements.

### Optical Network Elements (NEs)

This section discusses the following key elements that make up an optical network: optical fiber amplifiers (OFAs), wavelength division multiplexers (WDMs), optical ADMs, optical monitoring, optical fiber, and optical sources/detectors.

#### Optical Fiber Amplifiers (OFAs)

OFAs have played a key role in the construction of optical networks. In fact, this invention has driven the progress of all the other optical NEs. The drive to use OFAs has led to practical DWDM devices, extremely narrow line-width transmission lasers, effective channel monitoring schemes, advances in fiber technology, and promising optical cross-connect techniques. The first field trial of an OFA in December 1989 used four wavelength channels on the same fiber, and it was performed by Lucent (then AT&T) in Roaring Creek, Pennsylvania.<sup>2</sup> The OFA offers simultaneous amplification of any number of channels in the 1,550-nm wavelength band. (Current optical transmission systems operate in two wavelength windows: an older one in the 1,310-nm region and a newer, wider one supporting multiwavelength transmission in the 1,550-nm region.) The limits of capacity and distance are set by the gain, output power, noise figure, and gain flatness of the OFA, together with the signal-to-noise (S/N) requirements at the receiver. With the enormous potential in capacity made available by the OFA, the performance of the other optical networking elements has been pushed substantially beyond previously acceptable limits.

Most OFAs currently used to support sixteen wavelengths are the erbium-doped fiber amplifier (EDFA) type. The conventional EDFA amplified region is in the 1,530-nm to 1,565-nm range. The OFA tech-


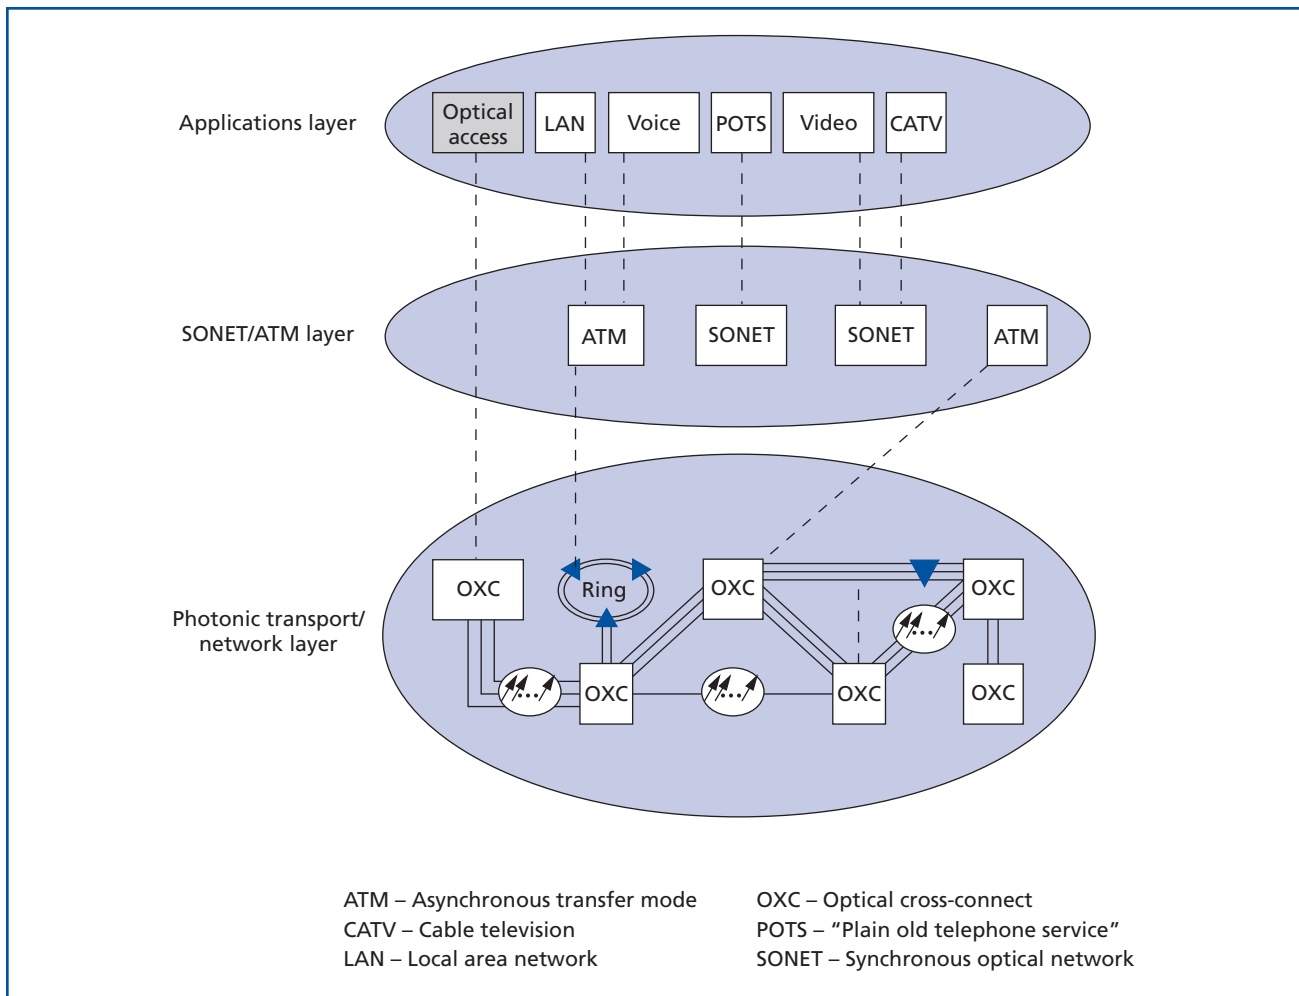


Figure 1. Trends in TDM and WDM.

nology available today is capable of supporting an 80-nm bandwidth with a potential wavelength capacity greater than 100 channels. The next two subsections discuss the key characteristics of OFAs.

**OFA main components.** An OFA is a simple device. Its principal constituents are a length of erbium-doped optical fiber, a pump laser, and a WDM component that combine the pump light with the transmission light (see Figure 3). An OFA amplifies signals in the 1,550-nm band, is capable of gains of 30 dB or more, and possesses output power of +17 dBm or more. Such additional components as optical isolators, additional gain stages, tap couplers, and embedded dispersion compensators may be needed to meet critical performance criteria—for example, reflection insensitivity, minimum output power levels, signal monitoring capabilities, and dispersion accommodation. Auxiliary features, such as integrated telemetry and gain profile adjustment,



**Figure 2.**  
**Network layers.**

can be incorporated with additional components.

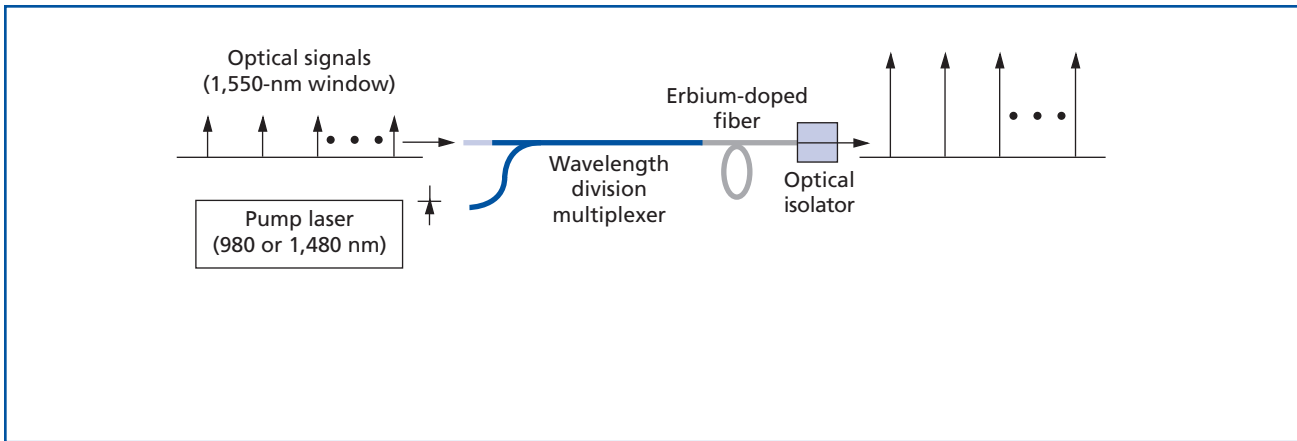
**OFA performance parameters.** The performance requirements for an OFA are determined by the application. There are a few key parameters that determine the performance of the OFA, and they can be optimized depending on the requirements. Some of these parameters are gain, gain flatness, noise level, and output power. In turn, these parameters are influenced by the choice of pump wavelength, number of channels, input optical dynamic range requirement, and overall S/N ratio requirement.

In designing optical networks, careful attention must be given to the optical S/N ratio (OSNR). A simple way to approximate the OSNR for an amplified optical line is to use the equation

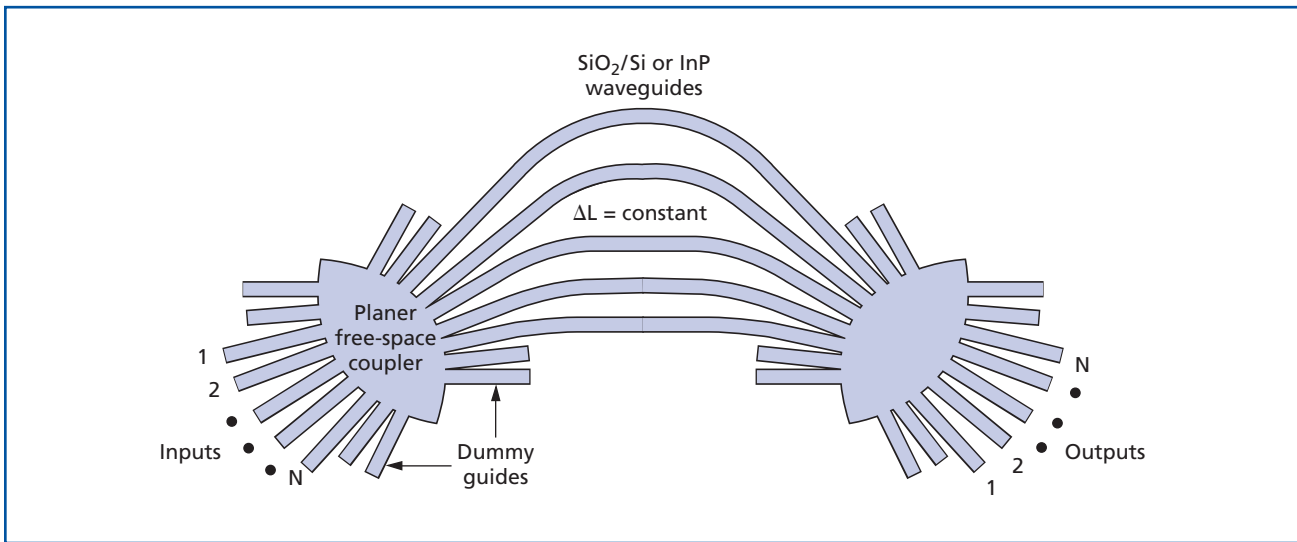
$$OSNR = C - [P - 10\text{Log}(N_\lambda)] - NF - L - 10\text{Log}(N_L),$$

where  $C$  is a constant,  $P$  is the total output power in dBm,  $N_\lambda$  is the number of wavelengths,  $L$  is the span loss,  $NF$  is the noise figure, which is characteristic of the OFA, and  $N_L$  is the total number of spans. This equation allows the performance margin of a system to be estimated, given the receiver sensitivity. Two factors not included in the equation are the impact of unequal gain across the channels and the loss in each span. In estimating OSNR, the designer must allocate sufficient margin to account for additional degradation factors, such as dispersion, multiple path interference, imperfect transmitter-receiver band shapes, nonoptimum decision threshold, finite extinction ratio, aging, and temperature effects.

To support transmission of a multiwavelength signal over long distances without regeneration, it is



**Figure 3.**  
Optical fiber amplifier (OFA).



**Figure 4.**  
An optical waveguide router.

desirable to maximize the number of OFAs that may be concatenated. To achieve this, the key OFA characteristics that a designer must target are *low noise* and *flat gain*.


### Wavelength Division Multiplexers (WDMs)

Wavelength multiplexers and demultiplexers are crucial devices that combine and separate closely spaced optical signals in an optical network. The most commonly used (and lowest cost) wavelength multiplexing device is the  $1 \times N$  combiner. This device, however, exhibits high loss with increasing  $N$ . Recent advances and cost reduction in filters have made it possible to incorporate them in multiplexers, resulting

in lower losses than feasible with a combiner.

Wavelength demultiplexers obviously must be wavelength selective. The ability to separate closely spaced optical signals accurately is challenging and requires innovation and technology improvements. Three popular devices optical network designers use to achieve wavelength selectivity include optical waveguide routers, multilayer interference filters, and optical gratings. The following two subsections briefly discuss optical waveguide routers and multilayer interference filters, which are more commonly used.

**Optical waveguide routers.** A DWDM using a waveguide router is shown in **Figure 4**. This wave-



**Figure 5.**  
**Multilayer interference filters.**

guide device is called a *Dragone router*, named after Corrado Dragone, its inventor in Bell Labs. A passive silicon device—in effect, just a piece of glass—it is an example of an early integrated circuit technology in fiber optics using lithography in silicon.

The principle of operation is optical interference. The device has an array of curved-channel waveguides with a fixed path-length difference between adjacent channels. In addition, two planar free-space slab waveguides are located at both the input and output. Input (multiwavelength) light diffracts into the first slab and enters the waveguide array, where the light signal in each channel is phase-shifted by a different amount with respect to the others, causing interference at the far-end slab. This process results in different wavelengths having interference maxima at different locations. These locations correspond to the output port positions. Thus, the signals are demultiplexed.

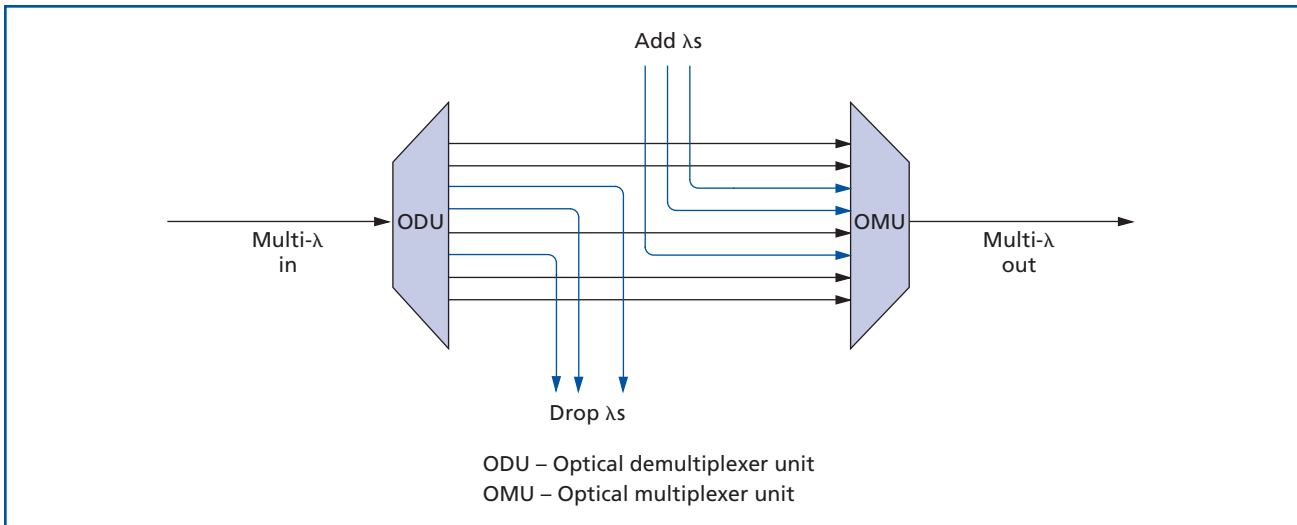
**Multilayer interference filters.** A DWDM using multilayer interference filters (color-coded rectangles) is shown in **Figure 5**. The property of each interference filter is that it will transmit a single wavelength and reflect others. The multilayer composition of the DWDM provides a more rectangular-shaped passband. The multiwavelength light is incident at an angle on to the device. It travels on the light-blue coded filter,

which passes the light-blue coded wavelength through and reflects all others onto the gray-coded filter. The gray-coded filter then passes the black-coded wavelength through and reflects the remainder onto the dark-blue coded filter, and so on. The wavelengths that are passed through the filters are the channels separated by the demultiplexer.

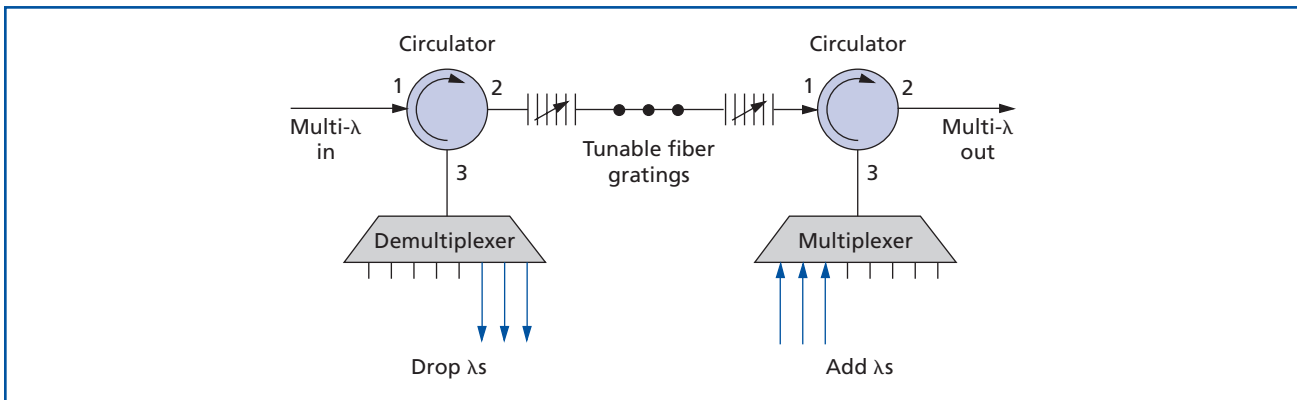
The critical parameters to contend with in the design of wavelength demultiplexers are insertion loss, filter flatness, bandwidth, adjacent and nonadjacent channel isolation, temperature dependence, and polarization dependence.

### Optical Add/Drop Multiplexers (OADMs)

At present, optical add/drop multiplexing is based on the following simple but costly principle, which **Figure 6** illustrates. An incoming multiwavelength line is demultiplexed by an optical demultiplexer unit (ODU). Desired wavelengths are dropped to local terminals and the remainder are passed through to an optical multiplexer unit (OMU). The OMU combines this through traffic with wavelengths added from the local terminals and generates the new output signal, as shown in the figure. This simple design requires two jumpers (one for each direction) per wavelength, connecting the OMU and ODU back to back. For 32-wavelength systems, up to 64 fiber jumpers are



**Figure 6.**  
*Optical add/drop multiplexer.*



**Figure 7.**  
*Programmable optical add/drop multiplexer.*

required, resulting in a significant intraoffice wavelength management challenge. The main performance limitation of this architecture is that each through-wavelength is unnecessarily filtered and combined at each add/drop site, resulting in optical signal narrowing and distortion leading to performance degradation.

In general, an OADM should have the ability to:

- Add/drop any wavelength in any order.
- Pass through wavelengths undisturbed by the add/drop operation. This is necessary to maximize the number of add/drop sites that may be connected in cascade without the need for regeneration.
- Be locally and remotely controllable.
- Exhibit low optical loss for through-

wavelengths. This feature allows for designing optical networks with fewer amplifiers, leading to improved OSNR.

To satisfy the above needs, teams from the Lucent product development and Bell Labs research organizations conceived several designs. Currently, they are evaluating one of the most promising designs for commercial availability—the *programmable optical add/drop* shown in **Figure 7**. The first generation of this device will support sixteen wavelengths and maintenance channels.

As the figure shows, multiwavelength light entering the left-hand optical circulator at port 1 and exiting at port 2 travels onward into a series of tunable fiber gratings. The circulator works according to the follow-

ing rule: light entering port 1 (2, 3) exits port 2 (3, 1). Each fiber grating is transparent to all component wavelengths of the multiwavelength light in the untuned state, but if tuned to a given wavelength, the grating will reflect that wavelength back in the incident direction. (Common methods of tuning are to use thermally induced or piezoelectrically induced strain.) Thus, in Figure 7, three of the gratings are tuned to the blue-coded drop wavelengths, and these wavelengths are reflected back into port 2 of the left circulator, whereupon they are directed to port 3 and then into the demultiplexer. As Figure 7 shows, the wavelengths are separated and make their appearance at the individual ports. The wavelengths not reflected by the tunable gratings (the through-wavelengths) enter port 1 of the right-hand circulator and exit port 2, joining the multiwavelength signal output.

Similarly, on the add side, the three wavelengths are directed into the multiplexer and combined. Next, they enter the right-hand circulator at port 3. The light is directed to port 1, enters the tunable grating, and then the wavelengths are reflected into port 1 of the right-hand circulator and directed to port 2 to join the multiwavelength signal output. There are no other added wavelengths, so nothing is transmitted toward the left through the series of gratings.

### Optical Monitoring

Monitoring the health of optical networks is vital when large traffic streams are involved. At present, monitoring is accomplished by providing a system that performs off-line measurements on the condition of the optical signals at various points in the network. The basic functions of optical monitoring are:

- *Power level measurements of individual wavelengths.* This function causes the triggering of an alarm condition if there is a reduction in the power level of an optical channel below a pre-specified threshold.
- *Wavelength identification.* This function is critical to add/drop operation in which confirmation of a wavelength being added or dropped is required.

Optical monitoring is in a relatively nascent stage, and advanced methods are currently being investigated. The subject is further discussed in the "Network Maintenance" subsection later in this paper.

### Optical Fiber

Today, the following two types of fiber are commonly used:

- *Single mode*, which is a very thin fiber (on the order of a 10- $\mu\text{m}$  core diameter) that supports the propagation of only the lowest order bound mode at the wavelength of interest. This type of fiber is used for long-distance and high bit-rate applications.
- *Multimode*, which has a larger core (on the order of a 50- $\mu\text{m}$  to 60- $\mu\text{m}$  diameter) and thus allows the propagation of more than one propagation mode at the wavelength of interest. In addition to the chromatic dispersion described below, this type of fiber leads to modal dispersion, which significantly limits the bandwidth-distance product. This type of fiber is used for low-speed and intraoffice communication.

Several types of single-mode fiber are commonly used, including:

- *Nondispersion-shifted fiber*, considered the standard type that supports wavelengths in both the 1,310-nm region with no dispersion and in the 1,550-nm region with chromatic dispersion of 18 ps/nm-km;
- *Dispersion-shifted fiber*, which supports wavelengths in the 1,550-nm region with no dispersion; and
- *Nonzero dispersion-shifted fiber*, such as TrueWave<sup>®</sup> fiber, with around  $\pm 2$  ps/nm-km dispersion in the 1,550-nm region.

Regardless of the type of fiber, the following factors are the main causes of performance impairment:

- *Attenuation*, which limits span distance. The typical loss allowed in optical network design is 0.25 dB/km.
- *Dispersion*, which also limits span distance and is bit-rate dependent. Other factors, such as chromatic dispersion and polarization mode dispersion (PMD), limit deployment of systems at 10 Gb/s and greater transmission speeds.
- *Nonlinearities*, which cause noise and distortion.
- *Connectorization and splicing*, which introduce additional loss.



## Optical Sources and Detectors

An optical source is a laser or a light-emitting diode whose sole purpose is to convert an electrical signal stream into pulses of light for transmission over an optical fiber. Lasers suitable for multiwavelength operation are required to conform to a set of standard wavelengths, the so-called *compliant wavelengths*. The complete set of these wavelengths have yet to be accepted universally by the standards bodies. Laser output is characterized by the wavelength, line width, power, and chirp (wavelength variation with changing amplitude).

The majority of optical sources used in optical networking employ distributed feedback single-mode lasers, which can be characterized by type as follows:

- *Directly modulated*, in which, as the name implies, the laser output is modulated (turned on or off) by modulating the laser current directly. These sources are generally simple and efficient. However, the bit rate is limited to 2.5 Gb/s, and there is a tendency to generate a broader laser line width as the modulation current changes. Dispersion limits the transmission distance on standard single-mode fiber to about 150 km.
- *Externally modulated*, in which the source maintains a steady output that passes through an external modulator functioning like a shutter, which in turn is controlled by the signal. These sources are generally more complex, exhibit high loss, and are more costly. They are used for bit rates greater than 2.5 Gb/s or for extremely long distances. A commonly used external modulator is the Mach-Zehnder interferometer. This device, which is fabricated on a lithium-niobate substrate, receives the laser output and splits it into two parallel rails, recombining the rails to produce the output. An electric field that varies with the impressed signal (other modulation methods exist) modulates one of the rails. If there is no field, the combined output is the same as the input that was split into the two rails because the rails are identical in length and carry the same signal. If a field exists and is of such a value that it

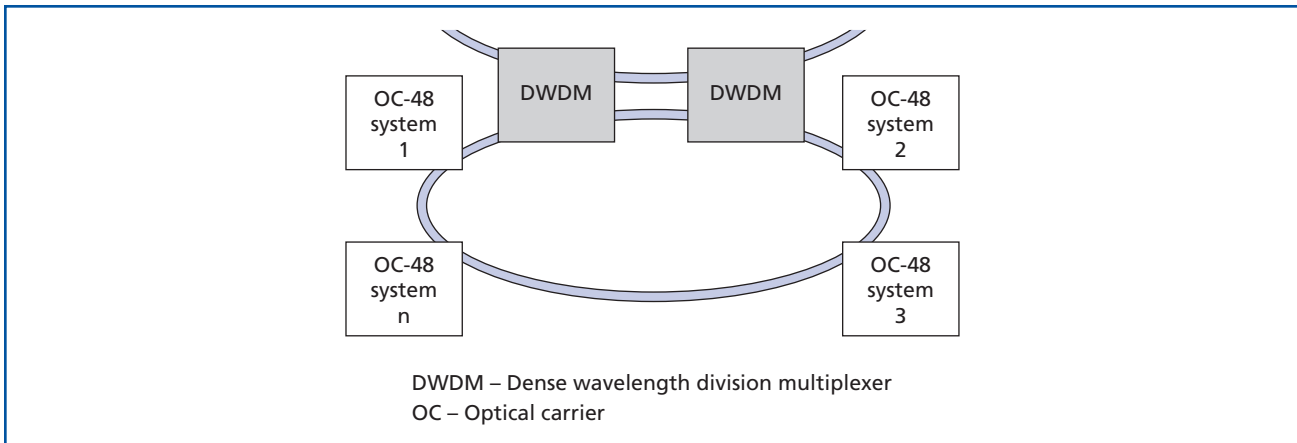
inverts the phase of the signal on the modulated rail, destructive interference will occur on recombination and the output will be suppressed. Because the laser current is steady, external modulation eliminates chirp.

An optical detector is the device used to receive the optical signal from the fiber and to convert it into the electrical signal we originally used to generate it. Receivers generally are of the wideband type and receive wavelengths in both the 1,310-nm and 1,550-nm windows, unless they are equipped with filters to reject undesirable wavelengths generated by random white noise and spontaneous emission of amplifier noise along the signal path. Optical detectors used in optical networking are generally the same as those used in conventional single-wavelength optical transmission. The factor that influences the optical detector of the optical network is the difference in the noise characteristic. If OFAs are used in the optical network, the noise at the receiver is asymmetric—that is, there is more noise in the *ones* than in the *zeroes*. This being the case, the decision point in the optical detector should be adjusted appropriately, which is the ideal solution.

## Why Optical Networking?

The most straightforward advantage introduced by optical networking, or DWDM in particular, is the potential of essentially unlimited bandwidth on a fiber pair—invaluable in situations of fiber exhaust. DWDM is also insensitive to bit rate or format, which means that the aggregated bit rate of the tributaries over the line does not get locked into a particular value as in a conventional SONET system (for instance, OC-48). In fact, the signal does not even have to adhere to SONET/SDH specifications.

DWDM allows the efficient use and rearrangement of embedded optical capacity as demand patterns change, contributing to transport networking flexibility as would be the case, for example, with a carrier deploying a ring network. On one hand, the carrier faces the high first cost of OC-192, as well as the fact that existing fiber facilities may not support OC-192 bandwidth, suggesting OC-48 deployment for initial service requirements. On the other hand, shortage of fiber facilities may inhibit the multiple



**Figure 8.**  
Span capacity relief with a DWDM in the OC-48 ring.

OC-48 rings necessitated by growth.

The DWDM alternative addresses this dilemma on all fronts. It solves fiber shortage problems on a per-span basis and facilitates first-cost and growth economic savings where OC-48 rings offer the best solution (see Figure 8). In the example, the span between nodes 1 and 2 is faced with a fiber shortage due to additional capacity being required on that span, perhaps because a pair of nodes from an adjacent ring need to use that span. The solution is to install a DWDM system in that span, which is equivalent to adding seven or fifteen (or more) fiber pairs (depending on the capacity of the DWDM system).

DWDM allows the elimination of high-capacity terminals in low-capacity sites, and it allows a multiplicity of OC-48s in a single span (for example, 16 would correspond to a 40-Gb/s capacity). Exploitation of DWDM minimizes capital expenditure by matching capacity expansion to actual demand (see Figure 9). In the figure, SONET terminals are added as capacity on a route is needed. Individual SONET terminals can be upgraded individually in capacity (OC-3, OC-12, OC-48, OC-192) as each route/ring capacity requirement grows.

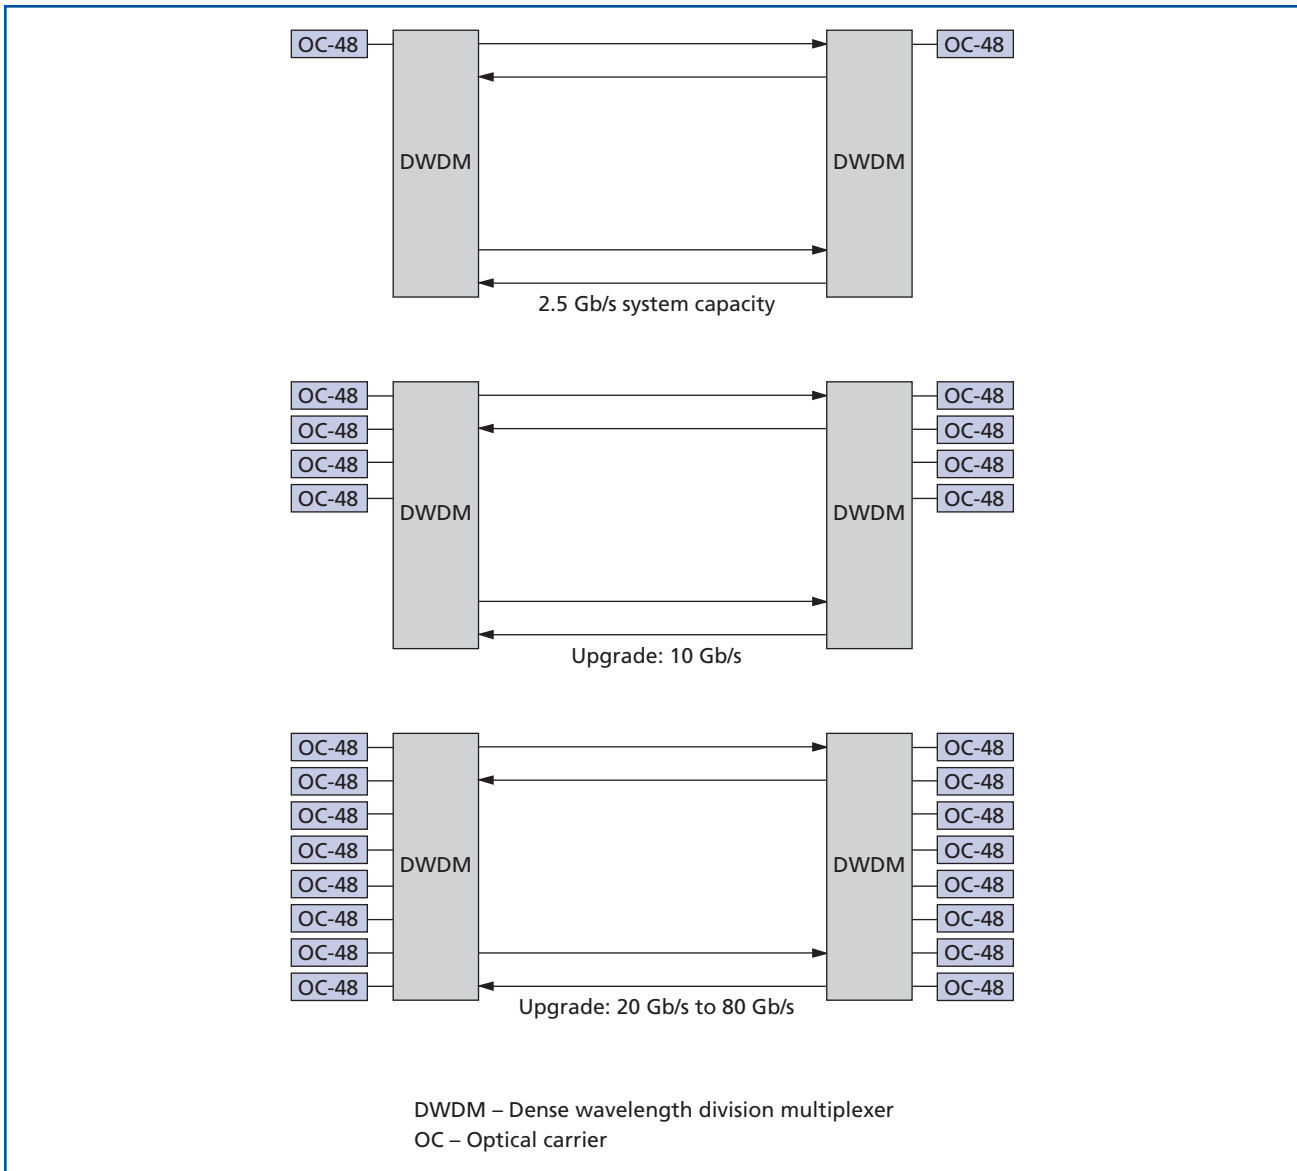
Optical networking also allows for simpler operations. Today's transport infrastructure consists predominantly of SONET/SDH signals. Such a signal carries with it an overhead component known as the data communications channel (DCC), which provides an in-band communications path between optically interconnected SONET/SDH transport NEs. Unfortunately,

there are limits to the number of NEs over which DCC connectivity can be maintained. First, there is an upper bound imposed by the NEs themselves, primarily due to processing constraints. Moreover, because DCCs have a finite bandwidth, a DCC serving large numbers of interconnected nodes can experience performance problems during periods of heavy message activity—for example, during “alarm storms.”

Thus, if a carrier were to multiplex a number of OC-3 rings onto an OC-48 ring, each with a substantial number of nodes, the above-mentioned problems would apply. By contrast, if the OC-3 rings were carried in DWDM channels in an optical ring network, they would retain their optical characters throughout, DCC and all. The NEs on each OC-3 ring in the DWDM system would communicate via their individual OC-3 DCCs, and the node limitation and bottleneck issues would not arise. Additionally, explicit segregation of the DCC is important in some cases for independent reasons, such as the need for maintenance segmentation (see the “Wavelength leasing” subsection later in this paper).

## Applications

In this section, we discuss the major applications of optical networking. These include long distance, the metropolitan area, and lastly, two special applications—high-speed parallel data transport and wavelength leasing.

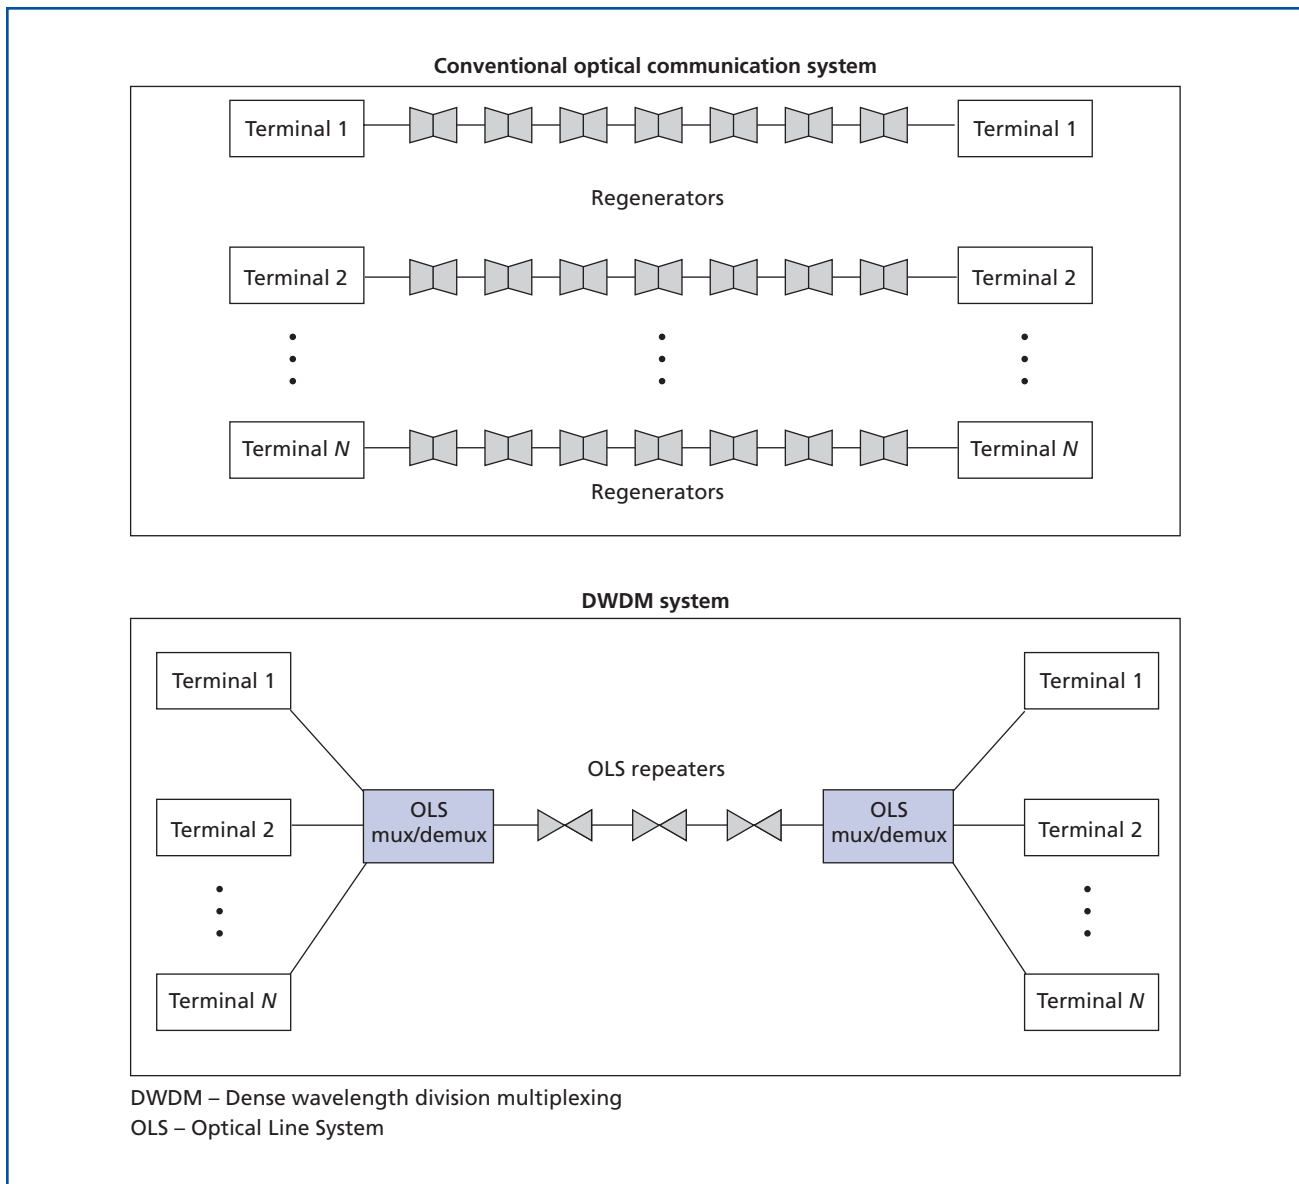


**Figure 9.**  
**Upgrade with DWDM.**

### Long Distance

The advent of the optical fiber amplifier has caused a revolution in long distance. **Figure 10** shows a long-haul application for the Lucent DWDM Optical Line System (OLS) product and stresses the benefit that the optical amplifier offers. Today's conventional long-haul systems in the network use multiple channels in numerous fibers with regenerators spaced every 40 km. This conventional regenerator spacing can be expanded to 120 km between optical fiber amplifiers. Then, because optical amplifiers are capable

of amplifying multiple channels simultaneously, N 2.5-Gb/s lines using eight fibers can be replaced by N 2.5-Gb/s DWDM channels using one fiber. Figure 10 clearly indicates that the equipment count is reduced by a large factor, which translates into substantial savings. One optical amplifier replaces twenty-four repeaters for an eight-channel system (that is, N=8). These amplifiers, unlike the regenerators, are format and bit-rate independent. They also support existing terminal equipment and make fiber available for further capacity expansion.



**Figure 10.**  
*Conventional optical communications system versus simplified DWDM system.*

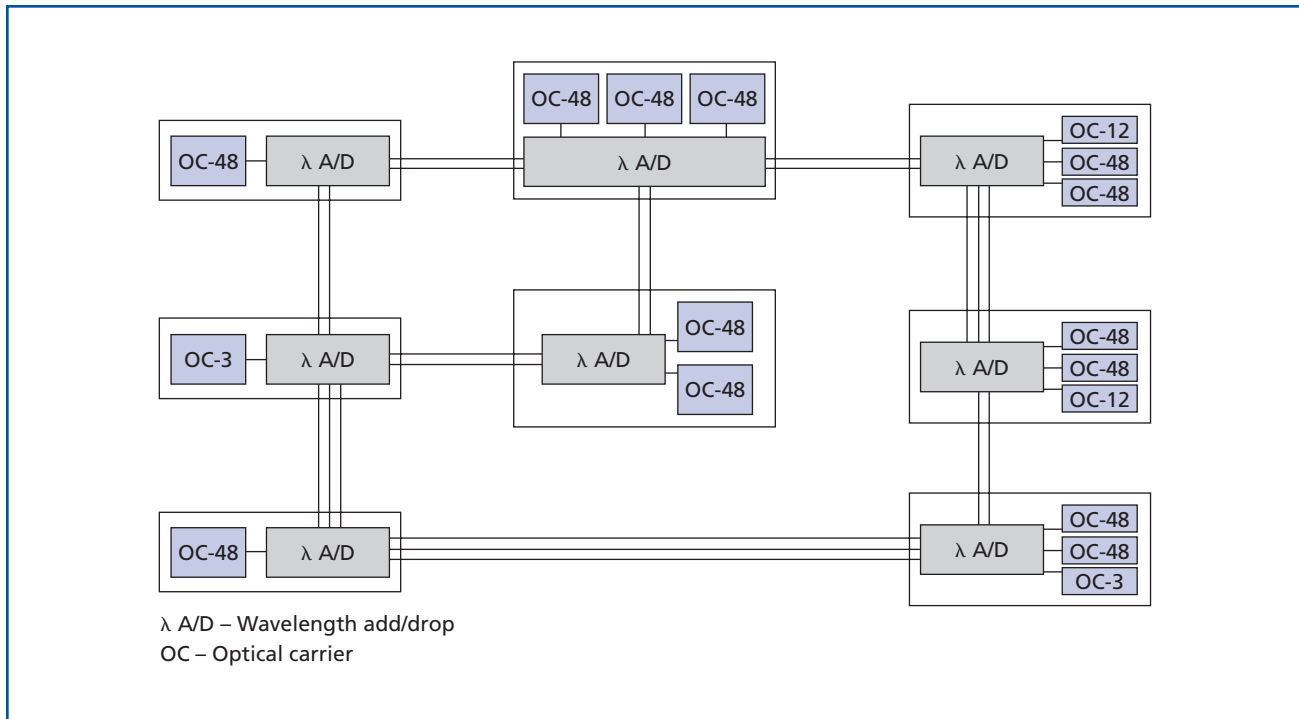
Finally, upgrading the two systems is vastly different. The TDM system requires new terminals, including higher speed transmitters and receivers. The WDM system requires the addition of another terminal at another wavelength. The amplifiers remain as they are, whereas the electronic regenerators all need to be changed.

The economics of the WDM solution are compelling. The WDM system avoids both new equipment system costs (end terminals and fiber) and saves ongoing maintenance costs on repeater sites.

For the long-distance network, moving beyond optical multiplexing into programmable wavelength ADM and full optical cross-connection promises to be the most effective way to provision, rearrange, and restore interoffice bandwidth. In the future, the optical layer will become the facility network.

**Metropolitan Area Network Architecture**

Optical networking is expanding beyond its initial application in the long-distance network to the metropolitan area. The key benefits of optical networking—



**Figure 11.**  
**Metropolitan interoffice backbone network.**

essentially unlimited bandwidth on fiber (new and existing), bit rate and format transparency, and efficient bandwidth use and management—directly affect the following metropolitan area network needs:

- Evolving the embedded interoffice transport network to provide higher capacity and support new services while not requiring the installation of new fiber,
- Providing cost-effective high-bandwidth transport for new end-customer services and for new network capabilities,
- Introducing alternate metropolitan area networks,
  - Private networks, and
  - New local service providers.

Next, we describe the application of optical networking to several specific situations.

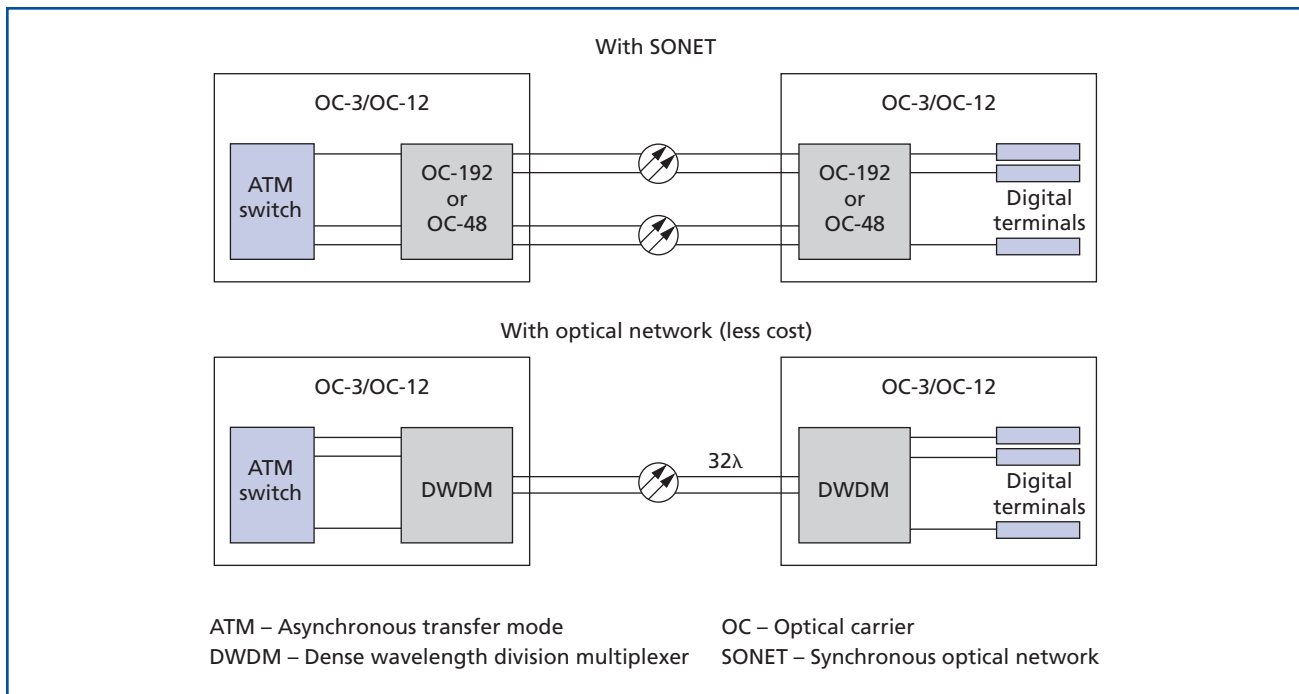
**Metropolitan interoffice network.** In addition to long-haul applications, DWDM may be used to solve fiber exhaust problems in metropolitan areas. **Figure 11** shows a local exchange carrier interoffice network design based on overlapping OC-48 rings riding on multiwavelength optical line systems. The rings

interconnect at key offices and have the following design requirements:

- No new fiber,
- Maximum flexibility in continued growth of the network,
- Support of new high-capacity end-user services (OC-3 and OC-12), and
- Lowest cost.

This design meets applicable requirements, and our studies indicate that it is 30% lower in cost than an OC-192 backbone network.

Typically, metropolitan applications require short terminal-to-terminal spans because there are more add/drop sites in smaller geographic domains in these areas. Because of the shorter spans, optical amplifiers are not required at mid span but only at the end terminals. In addition, the nature of a metropolitan application demands a mix of transport signals, such as OC-48, OC-12 and OC-3. Therefore, the metropolitan OLS will support interfaces to different transport signals. Any of the wavelengths can be upgraded from OC-3 to OC-48 and ultimately to OC-192 capacity



**Figure 12.**  
*Interactive multimedia.*

without affecting the bandwidth use on the other optical channels. This situation is in marked contrast to a TDM backbone.

**Direct optical transport of ATM.** High-speed multimedia services are beginning to appear on the network. **Figure 12** shows a core network ATM switch that routes traffic from a service provider's server to appropriate end offices for ultimate distribution to end users. The ATM switch and the end-office distribution systems both have OC-3 and OC-12 optical interfaces. Two network options are available: SONET multiplexing of the OC-3/12s onto a higher speed interoffice SONET transport system, versus optical multiplexing of the already-optical OC-3/12s onto a multiwavelength optical interoffice line system. For the example modeled, the optical multiplexing option was 50% less expensive.

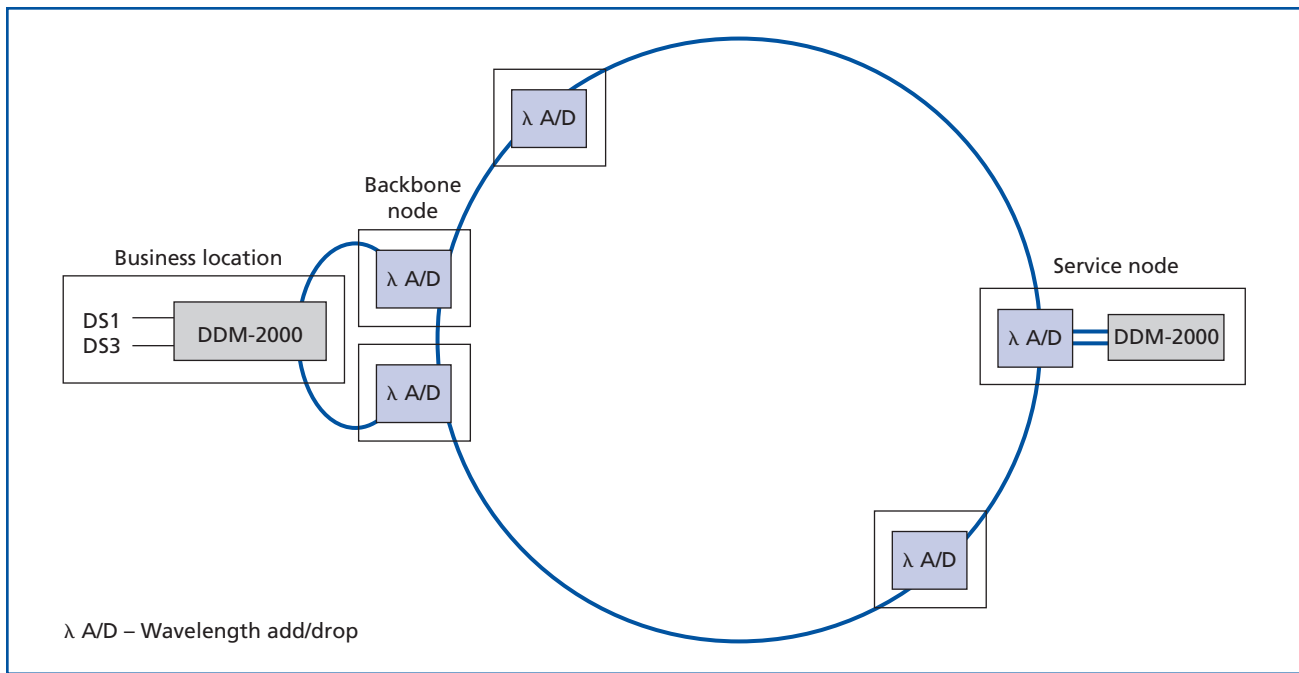
**Path-in-lambda architecture.** SONET rings have been used by private networks and competitive local service providers to provide high-capacity backbone networks. Two problems arise when installing a SONET backbone ring. The first is achieving good fill on backbone STS-1 rails. Many of the sites served by the backbone ring may require significantly less than a full

STS-1 drop. There is a need to aggregate DS1 traffic from several backbone node sites onto a single STS-1. The second problem is offering a high-reliability service option to business customers served by the backbone ring. Specifically, some customers want protection from a possible failure of a node on the backbone.

The optical networking solution is to build an optical add/drop backbone ring and use the wavelength channels as virtual fibers over which to run separate SONET path-switched rings. The path-switched rings provide end-to-end protection switching and aggregation of DS1s across several nodes into a single STS-1, and they support dual homing from business locations (see **Figure 13**).

Each wavelength in the multiwavelength optical ring used in this path-in-lambda architecture is capable of carrying OC-3, OC-12, or higher bit-rates, adding versatility to the solution. The basic building block is a programmable wavelength ADM, which is essentially a passive device permitting any of the sixteen wavelengths on the ring (regardless of the bit rate) to be added or dropped.

On the transmit end of each traffic direction, optical amplifiers (not shown in the figure) are used to



**Figure 13.**  
*Path-in-lambda backbone ring with dual-homed customer access.*

amplify all wavelengths simultaneously, if needed. Such use provides the power budget to circumvent the losses associated with both the fibers and the programmable ADMs along the entire backbone ring. With the path-in lambda architecture, timing synchronization is not an issue because the various wavelength components do not have to be synchronized with respect to each other, resulting in additional cost savings.

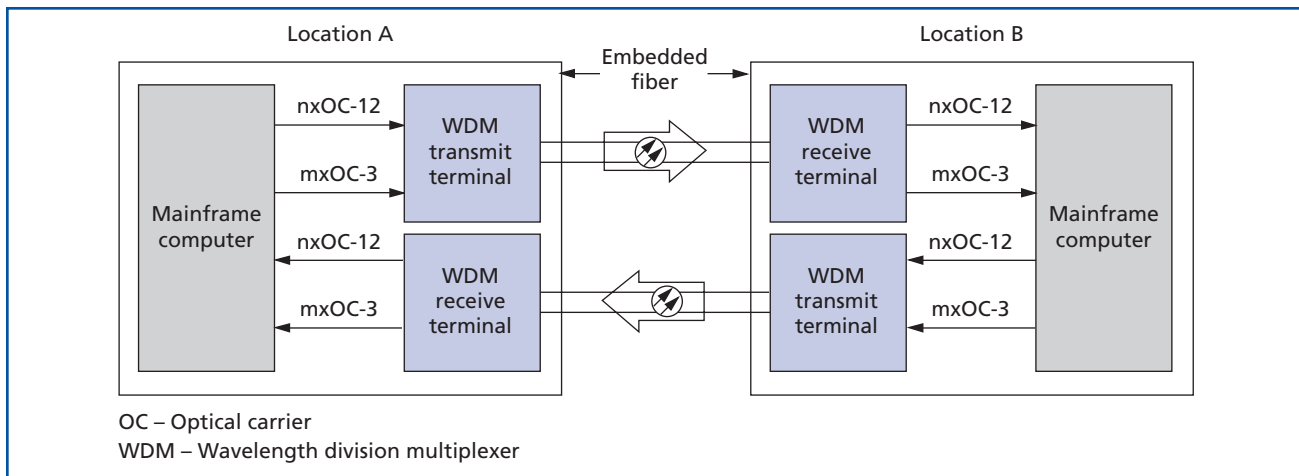
Overall, the path-in-lambda architecture provides the same capabilities as a comparable architecture that uses a TDM backbone, with a significant advantage in bandwidth capacity and flexibility on each of its optical channels. In addition, network equipment costs are lower and fewer fibers are needed with the path-in-lambda architecture.

### Special Applications

We now discuss two special applications of optical networking: long-reach high-speed parallel interfacing and wavelength leasing.

**DWDM for long-reach high-speed parallel interfacing.** Certain computer applications require that computer centers be interconnected with multiple high-speed channels that have capacity and availability requirements, as well as interlink delay restrictions

that cannot be met with TDM transport systems. **Figure 14** shows the configuration for such a computer application. Two mainframe computers need to connect over a long distance, and the high rate of information being exchanged between the two computers requires multiple parallel OC-N connections, as shown in the figure. The bits being transmitted at any instant on these lines have a tight skew-delay requirement—that is, they need to arrive at the far end within a very short deterministic (predetermined) time interval of each other. If the OC-N signals were time-multiplexed to a higher SONET rate, then the randomness introduced by pointer processing in the SONET terminal would introduce sufficiently high nondeterministic delay between channels, making it impossible to meet this skew-delay constraint. However, if the OC-Ns are wave-division multiplexed onto the same fiber, the signals are constrained to traverse the same physical path and, therefore, they have identical transit and processing delays (this ignores the minor skew introduced due to spectral dispersion among the channels). Thus, the skew-delay requirement is met. In general, DWDM optical transport benefits all delay-sensitive applications.



**Figure 14.**  
*Long-reach high-speed parallel interfacing application.*

**Wavelength leasing.** Network customers (both retail and wholesale) are beginning to demand high-capacity network transport that affords high reliability and security, as well as maintenance segmentation from the provider's network. **Figure 15** shows an interoffice optical network with several wavelengths being used for the provider's core network SONET/SDH ring applications. A spare wavelength (leased  $\lambda$ ) is used to provide clear-channel transport to a customer. There are distinct maintenance boundaries, but the customer can use the full set of SONET/SDH capabilities to provision and maintain private network equipment. The customer's bandwidth requirements are cleanly separated from the provider's core network needs.

### Future Directions

In the future, the optical layer will be the transport of choice, supporting all the current and prospective applications (see Figure 2). As noted previously, bandwidth demands will continue to increase, forcing network operators to manage the network on a per-wavelength basis, as they are doing today with DS3. This simple observation will have major implications, influencing future research and development with respect to new devices, and evolving current and future NEs.

In this section, we discuss some future directions relating to the following areas: optical network restoration, optical cross-connect evolution, loss-less

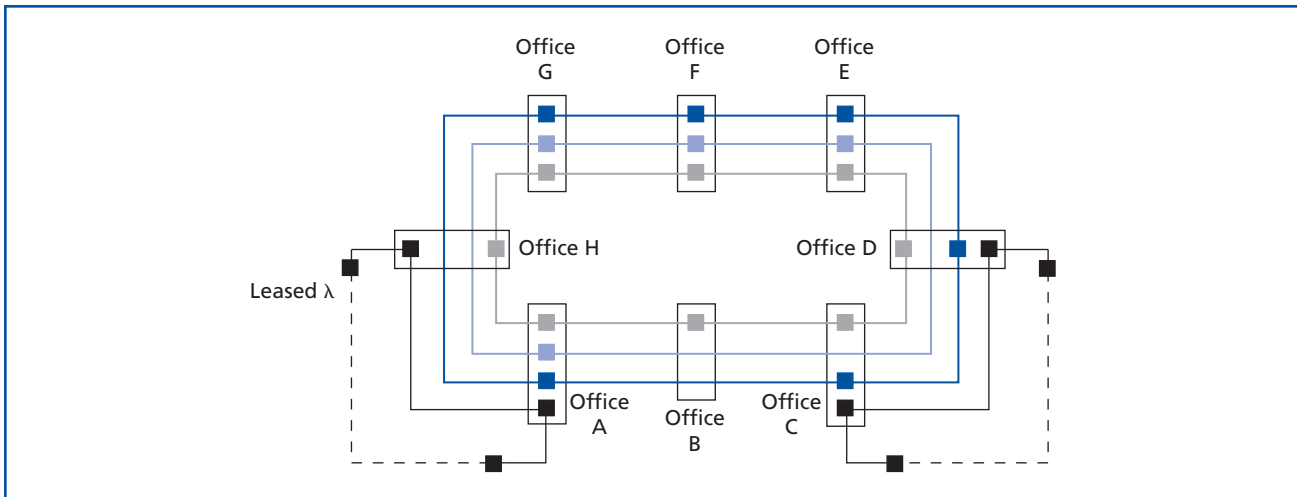
in-service upgradable optical ADMs, network maintenance, and wavelength capacity expansion.

### Optical Network Restoration

In optical networks, failure of a single fiber loaded with 32 wavelengths with each wavelength carrying an OC-48 signal will result in loss of service capacity equivalent to more than one million telephone connections. If the number of wavelengths is greater, the vulnerability is proportionately larger. If the line fill is high, it then becomes cost effective to perform restoration at the optical line level rather than at the constituent level, with optical restoration assuming great importance. Conceptually, network restoration via optical NEs is no different than restoration using the elements' electronic counterparts. It can be done either centrally under the command of a network management system or autonomously in a distributed fashion by the NEs themselves.

As an example, we discuss restoration via optical cross-connects (OXC). **Figure 16** illustrates a subnetwork consisting of the nine OXCs A through I. Each OXC has the capability of terminating a large number of wavelengths, and each one is connected to its neighbor with fiber carrying multiple wavelengths and one or more maintenance channels. These wavelengths may be cross-connected to other lines, added/dropped to TDM terminals, or earmarked as standby wavelengths for restoration. In the example Figure 16 illustrates, the black-coded wavelength is





**Figure 15.**  
**Wavelength leasing.**

terminated at F and H, where access to its electrical signals is required. It then passes through G, forming the wavelength path F-G-H. The blue-coded wavelength terminates at F and D and passes through G and C, forming the wavelength path F-G-C-D.

If a fiber cut takes place between G and F, the two wavelengths (coded black and blue) will fail. In a proposed distributed restoration scenario, terminating OXCs F and H on realization of the failure will initiate the restoration of the black-coded wavelength path F-G-H, while OXCs F and D will be responsible for initiating the restoration of the blue-coded wavelength path F-G-C-D. The failure realization and restoration process will be conducted via communication between the OXCs over the maintenance channels. The distributed restoration approach requires that the OXCs in the network maintain knowledge of the current state of the network (such factors as the cross-connections carrying service and spare capacity, for example). In addition, the OXCs must have knowledge of the necessary algorithms required to determine and implement the restoration paths, and they must reinstate the normal paths on repair of the failure. Restoration paths can be determined dynamically from knowledge of the existing and available cross-connections at the time of failure. They can also be determined in advance and maintained in the OXCs as *prestored maps* of cross-connections.

In a centralized scheme, however, the OXCs

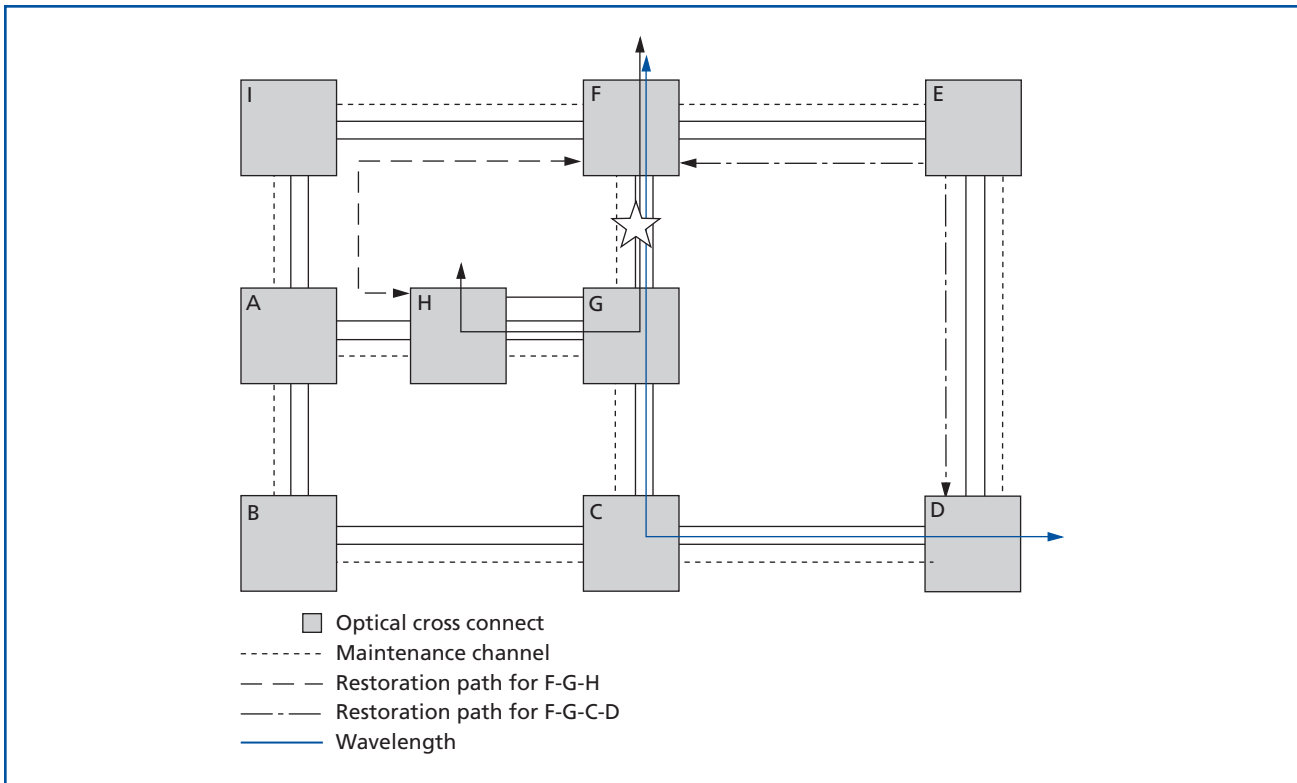
would report wavelength failures to a central management system (or hierarchy), which then would issue commands to the appropriate OXCs to reconfigure themselves to the restoration paths. In such a case, the intelligence (network state data and algorithms for restoration) would be maintained by the management system and not by the OXCs, and both the dynamic and prestored approaches would be feasible.

At present, distributed restoration appears to be the better choice for meeting providers' currently targeted network restoration times, which are on the order of 250 ms or less.

### Optical Cross-connect Evolution

Today's state-of-the-art optical cross-connect architecture, developed by Lucent and demonstrated at SUPERCOMM '97,<sup>3</sup> is shown in **Figure 17**. This wavelength-selective cross-connect architecture is transparent to signal bit rate, format, modulation scheme, and other signal characteristics. In addition, Lucent is developing a prototype of a large-size opaque fiber cross-connect that allows SONET/SDH formatted signals to be cross connected. The Lucent opaque cross-connect architecture is designed around an open architecture and is intended primarily as a restoration vehicle in multivendor optical networks.

For an M-line N-wavelength system, the architecture uses N independent layers of  $M \times M$  lithium niobate digital switch fabrics. Each incoming signal (labeled 1 through M in Figure 17) is demultiplexed to



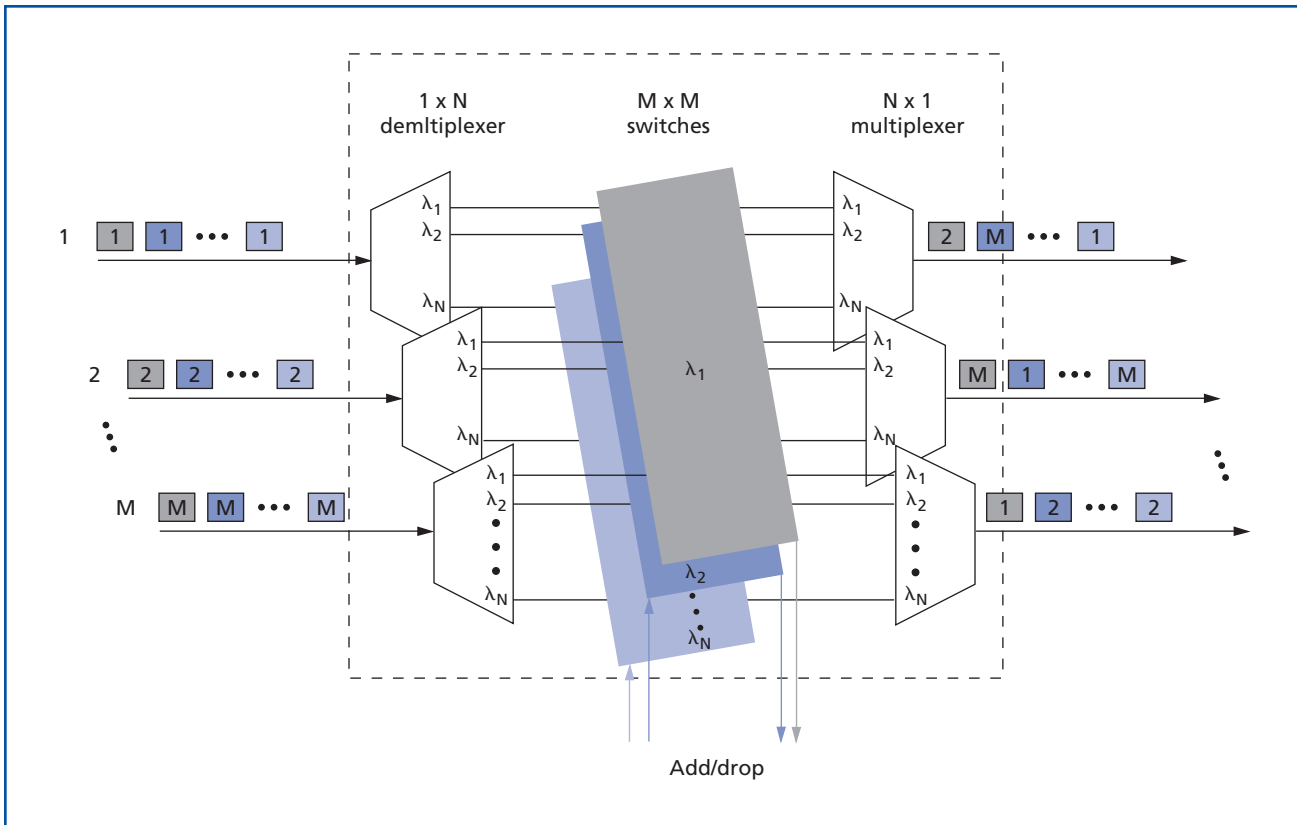
**Figure 16.**  
**Optical network restoration example.**

its  $N$  component wavelengths (distinguished by the colors and indicated by  $\lambda_i$ ), and each wavelength is routed to a corresponding switch fabric layer. The signals out of the switch fabric layers are re-multiplexed for routing (after amplification to the specified power level) to the output DWDM lines. At the input and output lines, a fraction of the DWDM signal is tapped off and monitored for integrity using a wavemeter for fault isolation purposes (not shown in the figure). In this architecture, local add/drop is implemented by eliminating the multiplexer/demultiplexer for a subset of the input/output ports (shown at the bottom of the figure).

A true  $(M \times N)$  cross-connect would require wavelength interchangers in addition to space switches. This feature is needed to permit wavelength reuse in a large optical network. The architecture that Figure 17 depicts can be converted to a wavelength interchanging OXC by incorporating wavelength changers in front or back of the switch fabrics, depending on the type of wavelength changers. Today, electro-optical wavelength changers, also known as optical translators

(OTs), are commercially available. All-optical wavelength changers based on a variety of principles, such as fiber amplifier cross-phase modulation in which one incoming wavelength induces the modulation of a second wavelength, will be available in the future.

Optical cross-connect technology is very new, and it is expected to mature considerably over the next several years. In near-term optical networks, the present architecture (with larger  $M$  and  $N$ ), perhaps using a different switch technology, may be applicable. In the next few years, other architectures not requiring demultiplexing are possible.<sup>4,5</sup> Later, radically different technologies may be developed to produce extremely compact, large capacity, and functionally versatile cross-connects. Promising technologies among them exploit the parallelism of light in space and provide a very efficient interconnect means for optical switching and, indeed, optical computing.<sup>6</sup> These kinds of two-dimensional optical interconnect technologies, based on both imaging and Fourier transform lens systems, can operate on entire arrays of optical signals. With



**Figure 17.**  
**Wavelength-selective optical cross-connect architecture.**

suitable high-speed spatial light modulators (SLM), one could expect to see very large size (for example,  $512 \times 512$ ) spatial cross-connects. These devices would be extremely small in physical size and could consist of lenses, prisms, mirrors and other passive optical components, in addition to an SLM powered by an array of laser diodes.

### Loss-Less In-Service Upgradable Optical Add/Drop Multiplexer (OADM)

The programmable OADM discussed earlier in the “Optical Add/Drop Multiplexers” subsection and depicted in Figure 7 has losses that require compensation by external OFAs. A loss-compensated OADM is realizable if one replaces the conventional fiber segment(s) interconnecting the two circulators by erbium-doped fiber segment(s). By optically pumping these segments, the fiber acts as a multiwavelength amplifier compensating simultaneously for losses associated with the through, drop, and add channels. The

pump light may be coupled in via an additional wavelength-selective coupler (or couplers) at one or both ends of the erbium-doped fiber segment(s). Alternatively, the pump light may be coupled in from the add side of the OADM.

By having gain on the add and drop channels, one could use passive splitters followed by tunable filters in the drop leg and combiners in the add leg of the OADM instead of in the demultiplexers and multiplexers as suggested by Figure 7. This alternative has the significant added advantage of allowing the OADMs to be upgradable in service.

### Network Maintenance

Maintenance of optical networks, as well as fault detection and isolation in the networks, are subjects of growing importance as the optical layer matures and as optical add/drop and cross-connect technologies are refined and become more cost effective. In this subsection, we briefly discuss specific proposals for monitor-

ing, fault detection and isolation, and end-to-end signal tracking.

**Monitoring, fault detection, and isolation.** In principle, monitoring optical networks is no different from monitoring electronic networks. However, because of the technological differences, the implementation problem has to be addressed anew. Various methods are being considered in addressing this challenge.

Bell Labs has suggested the use (at each add/drop site) of an optical monitoring unit that can tap off a fraction of the optical signals entering and leaving the add/drop systems, as well as the injection of a channel-equivalent signal into the outgoing fiber (as required)—all under automated control. The procedures would allow fault isolation and channel continuity verification by means of reports from the network monitoring units to the controller, which would then perform a comparison to determine the point of failure or discontinuity. In addition, the procedures would allow optical performance monitoring whereby wavelength, peak power, and OSNR for every channel would be received by the controller. Thus, degradations could be inferred and alarmed. This method would allow both in-service and out-of-service testing.

**End-to-end signal tracking and monitoring.** In an optical network, such as the one shown in the low-level layer of Figure 2, a signal (of any bit rate and format) riding on one of the compliant wavelengths may originate at any point, A, in the network and leave the network at any other point, B. Along the route from A to B, that wavelength may pass through several optical NEs, such as OXCs and OADMs, and temporarily coexist with other wavelengths conveying other information on the same fiber. Further, that signal may even change in wavelength several times by wavelength interchanging NEs. For these reasons, it may be desirable that this signal channel in the network be monitored and tracked from entry to exit.

One way of doing this is to tag each optical tributary at the entrance to the optical network and track it all along the route until it exits the network. At every NE—including the optical amplifier repeaters—the tag is monitored and, if necessary, it is refreshed, regenerated, or replaced. The tag follows each optical signal

tributary end to end and carries all the information regarding that signal, its origin, destination, and each NE through which it passes. **Figure 18** depicts this philosophy, which forms a virtual layer within the physical layer.

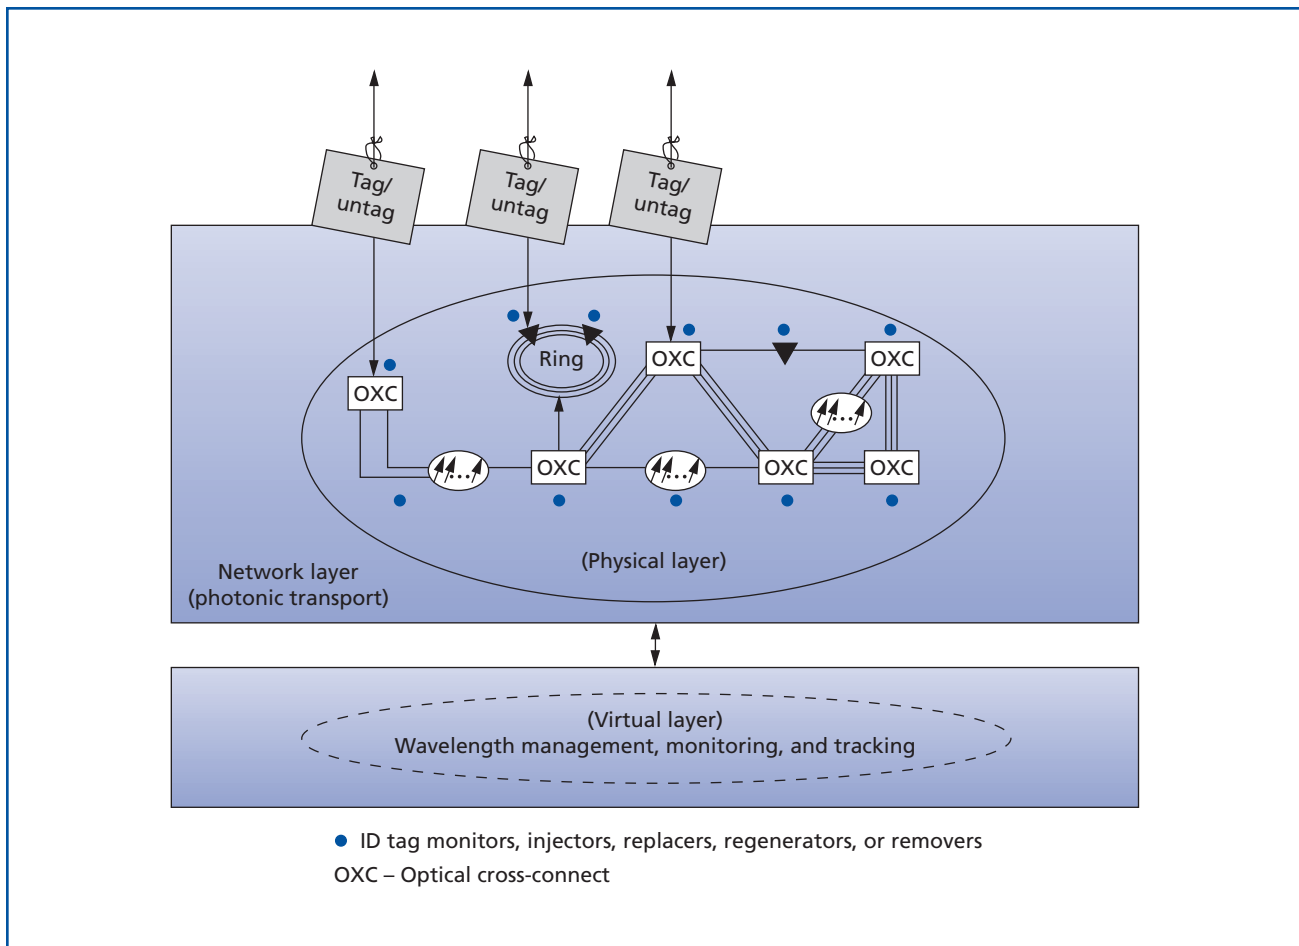
One proposed implementation<sup>7</sup> is to tag each tributary channel with a unique low-frequency overhead tone. This tag is inserted into each single-wavelength channel at the first encountered optical NE and can be (optionally) removed from each single-wavelength channel at the latest encountered optical NE. The tones may be further modulated with low-speed analog or digital information, such as the characteristic codes (an identification tag), or other information at each NE where the signal is demultiplexed. When added to a channel at any input point, identification tags may follow the signal, even through wavelength changers.

### Wavelength Capacity Expansion

Clearly, the larger the wavelength capacity of an optical network the better. Thus, it is important to strive to increase as much as possible the number of wavelengths that can be supported in a DWDM system using in-line OFAs. Currently, 8-wavelength and 16-wavelength DWDM systems have been widely deployed. In addition, an 80-channel system is planned, and Lucent has shown through experimentation that more than one hundred channels are feasible.

Channels in an optical network are limited primarily by the spectral bandwidth and shape of the fiber's optical amplifier. The spectral shape of optical demultiplexing filters also has a channel-limiting effect. Using in-fiber Bragg gratings and other techniques to flatten the amplifier gain, the entire bandwidth of the fiber's optical amplifier can be used. Employing co-dopants—for example, tellurides and fluorides—may cause wider bandwidth optical amplifiers to find their way into optical networks. Finally, using multiple fiber amplifiers—each responsible for a subset of the entire set of wavelengths, which are separated by a guard band—can further increase the number of wavelength channels to perhaps as many as 128.

Even though the available bandwidth in the current erbium-doped optical amplifiers is on the order of several THz, the number of channels would still be limited by the required channel spacing in the discrete



**Figure 18.**  
**End-to-end signal tracking.**

DWDM systems. Today's systems are based on 100-GHz and 200-GHz wavelength spacings. Closer spacing requires improvements in laser (source) and demultiplexing filter stability. DWDM channels with a 50-GHz spacing are possible with current technology.

In the future, to tap into the essentially unlimited capacity, the spectral spacing between the current DWDM-compliant wavelengths must be used to carry additional information. With an analogy to radio, optical spread-spectrum transmission may be among the candidates. With optical logic and storage elements,<sup>8</sup> it is not unreasonable to expect that the employment of optical time division multiplexing (OTDM) and optical code division multiplexing (OCDM) will use spectral real estate more effectively.

### Summary

The world has an insatiable appetite for bandwidth, and optical networking provides the means to deliver and manage almost unlimited capacity. Applied first in the long-distance network, DWDM has increased the capacity of existing fiber routes by a factor of eight in the past and sixteen today. In the future, fiber routes could experience increases in capacity by a factor of 80 and perhaps up to 128.

The cost of deploying a DWDM network is significantly lower than the cost of installing new fiber routes or route reinforcement. And, unlike higher capacity TDM systems, optical multiplexing is not challenged by embedded fiber characteristics. In addition, the longer distance (by a factor of 3) and multichannel (8, 16, 32, and possibly 128) capability of a single opti-

cal amplifier significantly reduces the cost of installing, upgrading, and maintaining an "army" of regenerators. Because of the fit of optical networking with existing network topologies, it is natural to expect long-distance networks to evolve to an optical network having wavelength-based facility provisioning, management, and restoration using optical cross-connects.

Low-loss programmable optical add/drop terminals having both a large channel capacity and bandwidth transparency on the individual optical channels are creating a new cost-effective backbone alternative to the metropolitan area network. Lower initial costs, fiber reuse, bandwidth growth capabilities, and maintenance segmentation solve the problems of the growing metropolitan transport network.

Finally, the format independence, bandwidth flexibility, low cost, security, performance, and reliability of an optical transport layer introduce new service opportunities for emerging high-capacity end-to-end customer applications that cannot be met by the public facility network today.

#### \*Trademark

Fastar is a registered trademark of AT&T.

#### References

1. *Fiber Optics News*, Vol. 15, May 1995.
2. D. A. Fishman, J. A. Nagel, T. W. Cline, R. E. Tench, T. C. Pleiss, T. Miller, D. G. Coult, M. A. Milbrodt, P. D. Yeates, A. Chraplyvy, R. Tkach, A. B. Piccirilli, J. R. Simpson, and C. M. Miller, "A High-Capacity Noncoherent FSK Lightwave Field Experiment Using Erbium-Doped Fiber Optical Amplifiers," *IEEE Photonics Tech. Lett.*, Vol. 2, No. 9, Sept. 1990, pp. 662–664.
3. "Lucent Technologies Announces Industry's First All-Optical Cross Connect," Lucent Technologies Press Release, SUPERCOMM '97, New Orleans, La., June 2-5, 1997.
4. W. Warzanskyi, F. Heismann, and R. C. Alferness, "Polarization-Independent Electro-Optically Tunable Narrow-Band Wavelength Filter," *Appl. Phys. Letts.*, Vol. 53, No. 1, July 4, 1988, pp. 13–15.
5. D. A. Smith, J. E. Baran, K. W. Cheung, and J. J. Johnson, "Polarization-Independent Acoustically Tunable Optical Filter," *Appl. Phys. Letts.*, Vol. 56, No. 3, Jan. 15, 1990, pp. 209–211.
6. M. T. Fatehi, K. C. Wasmundt, and S. A. Collins, Jr., "Optical Flip-Flops and Sequential Logic Circuits Using a Liquid Crystal

Light Valve," *Appl. Opt.*, Vol. 23, No. 13, July 1, 1984, pp. 2163–2171.

7. F. L. Heisman, M. T. Fatehi, S. K. Korotky, and J. J. Veselka, "Signal Tracking and Performance Monitoring in Multiwavelength Optical Networks," *Proc. ECOC '96*, Vol. 3, Oslo, Nor., 15-19 Sept. 1996, pp. 47–50.
8. M. T. Fatehi and C. R. Giles, "Erbium-Doped Fiber Amplifiers with Wavelength-Selective Optical Feedback," *IEEE Photonics Tech. Lett.*, Vol. 8, No. 8, Aug. 1996, pp. 1012–1014.

(Manuscript approved February 1998)

**DANIEL Y. AL-SALAMEH** is a distinguished member of technical staff in the Lightwave Hardware Department at Bell Labs in Holmdel, New Jersey. He received a B.S. degree from the University of Jordan, as well as M.S. and Ph.D. degrees from the Polytechnic Institute of New York, all in physics. Currently, Dr. Al-Salameh designs, specifies, and tests passive optical components and devices for DWDM systems, such as multiplexers, demultiplexers, filters, and routers. Previously, he worked on DWDM network design, optical protection switching, and metropolitan optical networks. He received an Eagle Award from the AT&T core network division for his contributions to the design and implementation of the Fastar\* restoration network. Dr. Al-Salameh is a patent holder and a member of the American Physical Society.



**MOHAMMAD T. FATEHI** is a distinguished member of technical staff in the Interexchange Carrier Business Development Department at Lucent Technologies' Optical Networking Group in Holmdel, New Jersey. He received a B.E. degree from The American University of Beirut in Lebanon, as well as M.S. and Ph.D. degrees from the Ohio State University in Columbus, all in electrical engineering. Since joining Lucent Technologies, Dr. Fatehi has focused on the vision of optical networking and on optical network architectures and elements. He holds numerous patents for his work in these fields, and he has authored many technical papers on related areas of research. Dr. Fatehi is a former professor at Tehran University in Iran and also at the Ohio State University. He is a senior member of the IEEE.



*WILLIAM J. GARTNER is product realization vice president of the Optical Networking Group at Lucent Technologies in Holmdel, New Jersey. His organization is responsible for the overall development of large-capacity SONET lightwave products and optical networking products. Mr. Gartner has a B.S. degree from the Georgia Institute of Technology in Atlanta, and an M.S. from Cornell University in Ithaca, New York, both in electrical engineering. In addition, he has an executive masters degree from the University of Pennsylvania in Philadelphia. Mr. Gartner, a patent holder, is a Moore Fellow in the management of technology at the University of Pennsylvania and a member of the IEEE.*



*planning and product feature definition for SONET and optical networks. His principal Bell Labs contributions are in the areas of self-healing ring and mesh network architectures, and network management and control. Dr. Raychaudhuri, a senior member of the IEEE, has written numerous technical papers detailing his work. ♦*

*STAN LUMISH is a Bell Labs technical manager in the Lightwave Hardware Development Department at Lucent Technologies' Optical Networking Group in Holmdel, New Jersey. He received B.E., M.S., and Ph.D. degrees in electrical engineering from the State University of New York at Stony Brook. Dr. Lumish, a Bell Labs Fellow, is responsible for the development of optical amplifiers, DWDM devices, and wavelength add/drop components for applications in global terrestrial lightwave systems. His department supports the 16-channel system and is now developing Lucent's next-generation terrestrial lightwave system.*



*BRUCE L. NELSON is director of the Interexchange Carrier Business Development Department at Lucent Technologies' Optical Networking Group in Holmdel, New Jersey. He received a B.A. degree from the University of Tennessee in Knoxville, as well as M.A. and Ph.D. degrees from the University of Texas in Austin, all in physics. Dr. Nelson's department offers development for, and marketing of, SONET and optical networking solutions to interexchange carriers.*



*KAMAL K. RAYCHAUDHURI is a distinguished member of technical staff in the Interexchange Carrier Business Development Department at Lucent Technologies' Optical Networking Group in Holmdel, New Jersey. He received an M.S. degree in physics from the University of Calcutta in India. In addition, he has an M.S. in computer and information science and a Ph.D. in physics, both from the University of Pennsylvania in Philadelphia. Dr. Raychaudhuri is working on network evolution*

