Optical Technology Update

Michael Finneran

There's lots of new technology to choose from. But are there any carriers with money to spend?

Nowhere has the Internet bust been felt more acutely than in the optical market. Two years ago, when the Internet engine was running at full power, optical networking seemed like a sure bet. But today, with Global Crossing filing for bankruptcy, carriers slashing capital expenditures and questions arising over carriers' bandwidth swaps, we have what one CEO calls a "nuclear winter" in telecom.

Despite short-term market prospects, however, optical technology continues to march forward. Venture money is still coming in and advances are being made on a variety of fronts. However, many of the developments slated for last year have been deferred until later this year.

Michael Finneran is president of dBrn Associates, Inc., an independent consulting firm in Hewlett Neck, NY, specializing in the design and installation of domestic and international networks. He is the instructor for several of BCR's data communications and optical networking courses. He can be reached at mfinneran@att.net

Optical's Analog Nature

The ultimate goal will be to develop a flexible, transparent optical network—a user would be connected to the network with an optical interface, be assigned a wavelength that in turn is combined with other wavelengths and cross-connected through a network of optical switches. In essence, the customer would be buying a slice of the carrier's optical bandwidth and could transmit any signal, at any bit rate, up to the analog capacity of the channel.

The difficulty and the impetus for many new developments, is the fact that an optical network is analog, with all the attendant problems—like loss, noise and crosstalk. If we separate and recombine wavelengths from different fiber links, we must ensure that all the channels have equal levels of signal, dispersion and noise.

Further, as we increase the transmission rate on each channel, these analog issues become more acute. For the past three years, 10 Gbps has been the maximum rate for a single optical transmitter. New 40-Gbps (OC-768) lasers have opened the door to higher-capacity wave-division multiplexed (WDM) systems, but we might not see a full-scale migration from 10 Gbps to 40 Gbps, according to Paul Haddad, global leader for product solutions at Nortel Networks. He maintains that ranges beyond 1,000 km will likely stay with 10-Gbps terminals for the foreseeable future, and 40 Gbps will be used within 1,000 km.

The primary reason for this is that 40 Gbps is a lot harder to do. First, a 40-Gbps transmitter cannot be directly modulated (i.e., turned on and off to signal 1-bits and 0-bits) without slightly broadening the wavelength band (an effect called a "chirp"). Instead, in 40-Gbps transmitters, the laser is left on continuously, and a "shutter" mechanism or modulator is opened or closed to transmit the signal.

While leaving the laser on solves the chirp problem, the external modulator reduces the transmission power. Further, at 40 Gbps, the pulses of light will be spaced more closely, so pulse spread-

FIGURE 1 Dispersion Effects

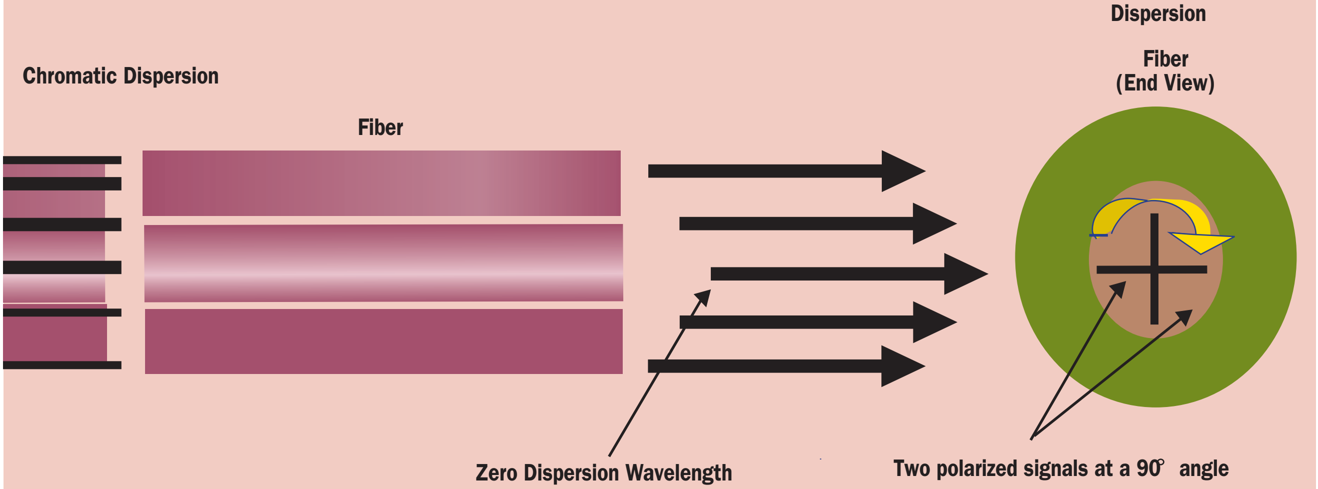

ing or dispersion must be tightly controlled to keep the pulses from overlapping and causing errors.

Refining The Optical Path

One of the major obstacles to 40-Gbps optical networks is the perplexing problem of pulse dispersion; optical pulses tend to spread out as they travel over the fiber. According to Dr. Alan Wilner, CTO of Phaethon Communications, a component manufacturer specializing in dispersion compensation, there are two different causes of pulse dispersion: Chromatic dispersion and polarization mode dispersion (Figure 1).

Chromatic dispersion (CD) results from the fact that different wavelengths travel at slightly different speeds over the fiber; the wavelength that travels fastest is called the zero dispersion point.

Polarization mode dispersion (PMD) results from the fact that the pulse of light is arranged as two polarized signals aligned at a 90-degree angle, and the two signals travel at different speeds.

Both of these problems can vary over time and for different reasons. CD varies with the temperature of the fiber, while the primary cause of PMD is the fact that the fiber core is not perfectly round for the entire length of the fiber.

Fiber manufacturers attacked the PMD problem with a type of fiber called "non-zero dispersion shifted fiber" (NZDSF), which generally has far better core symmetry. However, PMD also varies with stress, temperature and vibration; in addition, fiber amplifiers can introduce PMD.

Furthermore, trying to solve one dispersion problem can worsen the other. The new 40-Gbps transmitters use a form of line coding called return-to-zero, which is designed to reduce PMD, but according to Alan Wilner, this technique results in a wider pulse spectrum, and therefore further exacerbates the CD problem.

Both dispersion issues are more problematic as the bit rate increases. CD is more stable, but it increases faster—i.e., at the square of the bit rate (for example, the CD in a 40-Gbps transmission is 16 times that of a 10-Gbps transmission). In contrast, PMD increases in a more linear fashion, but gyrates more randomly.

According to Wilner, this randomness means that PMD compensation will always require a dynamic compensation technique. On the other hand, CD might be treatable with a fixed-compensation technique, at least at metro area ranges. However, long-haul, 40-Gbps transmission will need tunable CD compensation based on the length and temperature variations the fiber spans will experience.

Fixed chromatic dispersion compensation has traditionally been done with lengths of dispersion compensating fiber (DCF), which allow the slower wavelengths to "catch up." Corning is the leader in this technology. Another technique for dispersion compensation uses specially-designed optical filters, called Bragg Gratings, to reverse the effects of CD. Phaethon builds tunable compensators that mechanically stretch a Bragg grating as one approach to this.

WDM Developments

When I wrote about WDM last year (see *BCR,* May 2001, pp. 24–26), the state of the art for long-haul WDM systems was 160 channels, each running 10 Gbps over a single fiber pair, for a total capacity of 1.6 terabits per second (Tbps). These systems placed 80 channels in the C Band and 80 in the L Band (Table 1); channels were spaced at 50-GHz intervals.

TABLE 1 ITU Wavelength Bands

Label	Wavelengths
Unlabeled	820-900 nm
O Band (Original)	1260-1360 nm
E Band (Extended)	1360-1460 nm
S Band (Short)	1460-1530 nm
C Band (Conventional)	1530-1565 nm
L Band (Long)	1565-1625 nm
U Band (Ultra-long)	1625-1675 nm

Vendors have been pushing beyond these limits, and there are two basic strategies: Increase the number of bits carried on a channel or increase the number of channels. In turn, there are two ways to increase the number of channels:

1. Expand the usable frequencies (aka the bandwidth) to make room for more channels.
2. Pack channels closer together..

A primary limit on the first option, expanding optical bandwidth, is the design of the erbium-doped fiber amplifiers (EDFAs) that are commonly deployed on fiber routes. EDFAs only boost light in the C and L Bands, and a different EDFA is required in each band. To use both bands, the C Band and L Band wavelengths are physically separated, fed into separate amplifiers and then recombined. To make that band separation simpler, a 5-nm separation is left between the two bands—that unused bandwidth represents about 12 10-Gbps channels, so you lose 120 Gbps due to EDFA design constraints.

The second option—packing channels closer together—is primarily feasible in long-haul systems. Current long-haul WDM spaces channels at 50 GHz. Ciena pioneered the use of 25- and 12.5-GHz spacing, but packing channels that densely requires precise lasers and channel separators—which in turn affects the transmission range, amplifier spacing and, ultimately, cost.

In contrast to long haul, metro systems sacrifice efficient use of optical transmission bandwidth in the interest of lower cost. Metro systems employ coarse WDM (CWDM), which spaces channels at 100 GHz or 200 GHz, so while they get fewer channels (typically 32 or 64), the devices are far cheaper to build.

WDM vendors continue to pack more channels into less space

The major development in tunable lasers is the ability to put the whole device on one chip

There's another drawback to the tight channel spacing of long-haul systems: The smaller-bandwidth channels (i.e. 50 GHz) cannot support 40-Gbps transmissions. To see how this plays out in real-life equipment design, consider the systems manufactured by Ceyba Corp., which mix 10- and 40-Gbps channels in the same band. The device supports 160 10-Gbps channels which use 50-GHz spacing; however, only 80 40-Gbps channels are supported, and 100-GHz spacing is required to support the higher bit-rate. Nevertheless, the capacity is impressive—you still wind up with 3.2 Tbps per link.

Xtera Communications is going in a slightly different direction by using Raman amplifiers. Unlike an EDFA, Raman amplifiers use a powerful laser source to boost the signal power in standard optical fiber, and a single amplifier can boost both the C and L bands. This obviates the need for the 5-nm separation band required in EDFAs, so more capacity is available. According to Chip Pratt, Xtera's product portfolio manager, the system can deliver 240 50-GHz channels for a total link capacity of 2.4 Tbps.

The last entry in the area of amplifiers is the semiconductor optical amplifier (SOA). Unlike EDFAs or Raman amplifiers that use pump lasers, SOAs are built on a single chip. Their small size allows SOAs to be integrated onto multifunction optical chips, where they can serve as miniature switches, and potentially even be incorporated into designs for wavelength conversion—which would be a major breakthrough, as discussed below.

The next 12 months could be the coming out party for SOAs as they find their way into a variety of low-cost optical devices, but there is one big downside: Noise. Where EDFAs typically degrade the signal/noise ratio by 3dB, SOAs are in the 6–9 dB range.

Flexibility Factor: Tunable Lasers

Beyond raw capacity, the other area for development in optical networks is flexibility. Since the goal of a transparent optical network is to separate and recombine wavelengths, what happens if the particular wavelength we need on a given route is already assigned to another transmission? That problem is called wavelength contention.

The ideal solution would be a device that could directly translate light from one wavelength to another, while not diminishing power or introducing dispersion or other unwanted effects. Unfortunately, wavelength converters do not exist yet.

The current state of the art in wavelength conversion is an O-E-O transponder that receives the optical signal at one wavelength, converts it into an electronic signal and then drives a tuned laser that produces the desired (changed) wavelength. But these tuned lasers cost thousands of dollars, and a separate component is needed for every channel on the WDM system. Further, the carrier would have to stock spare parts for each of those transponders.

To address those difficulties, the manufacturers are turning to lasers that can produce a range of wavelengths. There are a number of challenges in producing such tunable lasers including cost, output power, tuning range and stability.

Tunable lasers have been available for laboratory systems for years, but they've been large, bulky devices unsuitable for an optical transmission systems. The major development in tunable lasers is the ability to build the entire device onto a single silicon chip.

One way this has been done is to change the way semiconductor lasers are designed. In traditional semiconductor lasers, the resonating cavity—the reflecting chamber where the optical power is built up—is constructed horizontally and the light emits from the side of the chip in an elliptical cone. In the alternate design, called a vertical cavity surface emitting laser (VCSEL), the resonating cavity operates vertically and the light emits from the surface of the chip in a round beam that couples much more efficiently (i.e., with less loss) to the round core of the fiber.

VCSELs are widely used for 850-nm transmitters, but generating longer wavelengths has been problematic, because they require a longer resonating cavity. That, in turn, means depositing more layers of material on the chip.

However, one vendor, Bandwidth9, has developed a VCSEL that can tune to wavelengths in the 1500 nm range to support either C or L band applications. In their design (Figure 2, left side), a mirror is placed on top of the device, its position controlled with a device called a micro-electromechanical systems (MEMS) actuator. Changing the position of the mirror changes the size of the resonating cavity, which allows it to tune to a particular wavelength.

According to Charles Duvall, senior director of applications for Bandwidth9, one of the major advantages of VCSELs is that a large number of them (1,000 or more) can be built on a single wafer, and they can be tested early in the manufacturing process before a lot of money has been committed. All the VCSELs on the wafer are tested and if the yield appears to be poor, the equipment manufacturer throws the whole thing out and starts over.

The major tradeoff with VCSELs is output power. Bandwidth9's device is directly modulated, so it eliminates a power-sapping external modulator, but it can generate only 1 mW. In laboratory tests, VCSELs have reached 5 mW, but that is still far short of the 20 mW power required for long-haul transmitters. As a result, Bandwidth9 is focusing its efforts on metro applications.

Meanwhile, iolon Corp. is tackling the problem of long-haul tunable lasers, manufacturing 10-Gbps lasers for the C and L Bands as well as a 40-Gbps C Band tunable laser (Figure 2, right side). The 10-Gbps devices can provide 100 50-GHz channels or 200 25-GHz channels in each band, according to John Clark, iolon's CEO. Their device uses an external tuning mechanism where the different wavelengths are separated using diffraction grating, and a MEMS-controlled mirror is positioned to select the desired wavelength.

New Approaches In Optical Switching

As noted above, the inherent difficulty in delivering a transparent optical network is dealing with the analog issues of signal and noise levels when separating and combining wavelength channels from different fiber routes. We are now starting to see products that address those issues. At Supercomm this June, Innovance Networks plans to introduce a new type of optical switching system or "Agile Photonic Solution."

According to James Frodsham, chief operating officer, the key element in Innovance's approach is a management system that assigns channels and selects the wavelengths to be used. A single channel might pass through multiple switches, thereby introducing the possibility of wavelength contention. In the Innovance switch, the management system tracks what wavelengths are available on each link, and assigns a wavelength end-to-end. The management system then sets the tunable lasers at each end of that channel to the correct wavelength.

To address the problem of signal equalization, the management system will know the link length and loss parameters for each span, and automatically tune the amplification as a channel is provisioned. When a channel is assigned, the management system will determine the level requirement for the channel on each link it will traverse, and set the gain accordingly.

Given the large number of possible channel combinations, there are limits on the end-to-end channel length. Innovance will be able to configure channels with maximum end-to-end range on the order of 2,500 to 3,000 km; the company claims that 80 percent of channel links in North America are 1,000 km or less.

Optical Packet Switching

Besides the difficulty of equalizing the signal, noise and dispersion levels, the other challenge in optical packet switching is storing and processing light, with the ultimate goal of applying routing techniques to the wavelengths.

There is a joke in the optical business that says all problems with optics can be solved with electronics—and so far, that's proving to be the case with optical routing. A case in point is AcceLight Networks, which will introduce its Photonic Burst Switch at NFOEC this September.

According to Mark Milinkovich, director-strategic marketing, AcceLight's product will be the first optical switching system capable of routing; unlike optical cross connects, its switching fabric is optical rather than electronic. But that doesn't mean there are no O-E-O conversions in the AcceLight switch—in fact, it's really an O-E-O-E-O switch (Figure 3).

Here's how it works: Individual wavelength channels come into the optical interface and go through an O-E-O conversion into a 1310-nm wavelength, which is the common denominator on which the optical switching matrix is built. Upon exiting the matrix, each channel is run through another O-E-O conversion to produce the

The ultimate goal is to apply routing techniques to wavelengths

FIGURE 2 Tunable Lasers

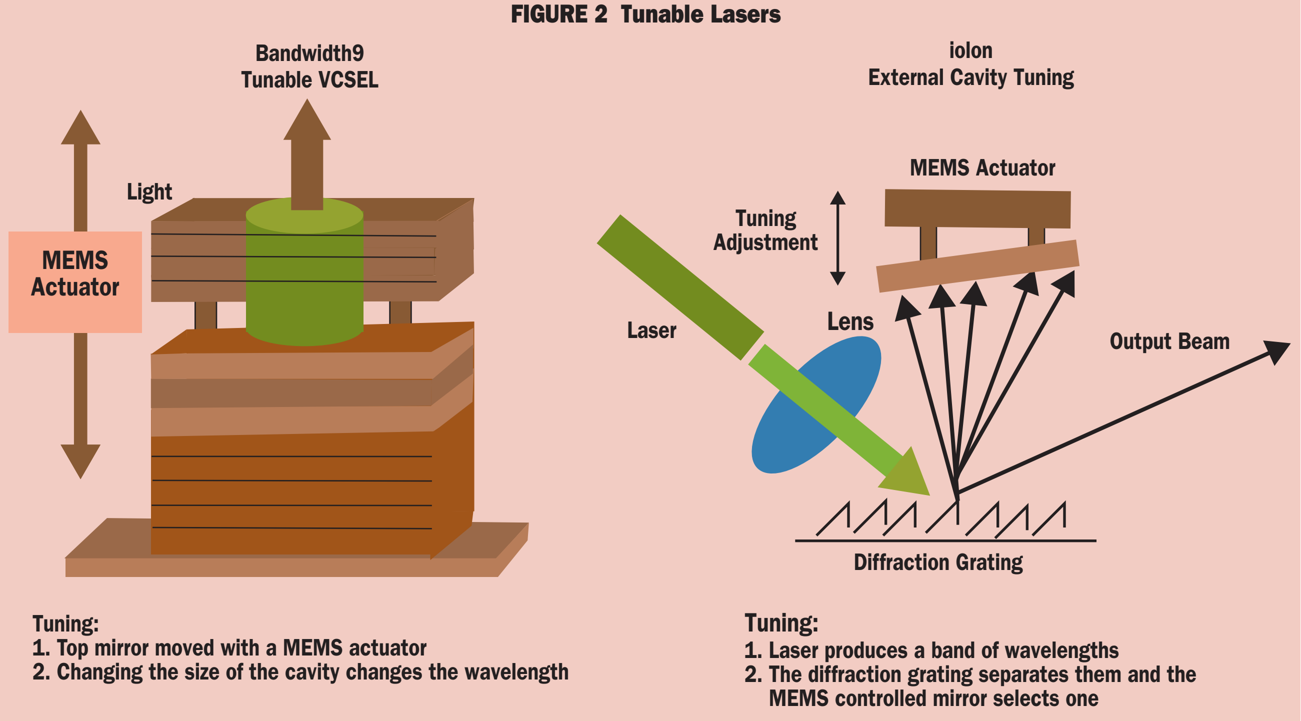

Vendors' control planes don't interoperate—a big drawback to the carriers

wavelengths needed for the next optical link in the end-to-end connection.

While this seems a rather round-about way of doing optical switching, it allows any incoming wavelength to be converted into any outgoing wavelength, and the signal is regenerated in the process, so we start out again with a brand new pulse of light for the next leg of its journey.

Furthermore, handling the packets in electronic form allows AcceLight's switch to apply routing intelligence to the process. When the incoming signal is converted into electronic form, the routing decision is made by a traditional electronic computer. The packet is buffered and a connection request is forwarded to a scheduling system that controls the optical switching matrix. The scheduling system sets up the path through the switch matrix, the data packet is converted into light and the optical burst is fired.

AcceLight's approach combines the functions of an optical cross connect with those of a router, according to Milinkovich. Like an O-E-O cross connect, the AcceLight device will be able to switch individual STS-n channels from OC-48 or OC-192 ports. However, the system also will include full routing capabilities including MPLS and plans for GMPLS support (more on GMPLS below).

Optical Control Plane

The last facet of flexibility is the control system used to assign channels through a network of optical switches. While every vendor has such a system, none will work with other vendors' products—a major drawback from the carriers' perspective.

There are three major standards initiatives for the optical control plane, according to Joe Berthold, VP of network architecture for Ciena. These are:

- The Optical UNI from the Optical Internetworking Forum.
- Generalized Multi-Protocol Label Switching (GMPLS) from the IETF.
- The Automated Switched Optical Network (ASON) from the ITU.

The Optical UNI allows for neighbor discovery and channel assignment at the STS level. The control interface can be an out-of-band Ethernet interface or an in-band command carried in the SONET overhead.

Neighbor discovery, or the ability for a device to automatically learn the availability of capacity in adjacent devices, is powerful though potentially problematic. While the carriers would like this capability, they also want to keep their topology secret from other carriers and from customers.

FIGURE 3 AcceLight's Photonic Burst Switch

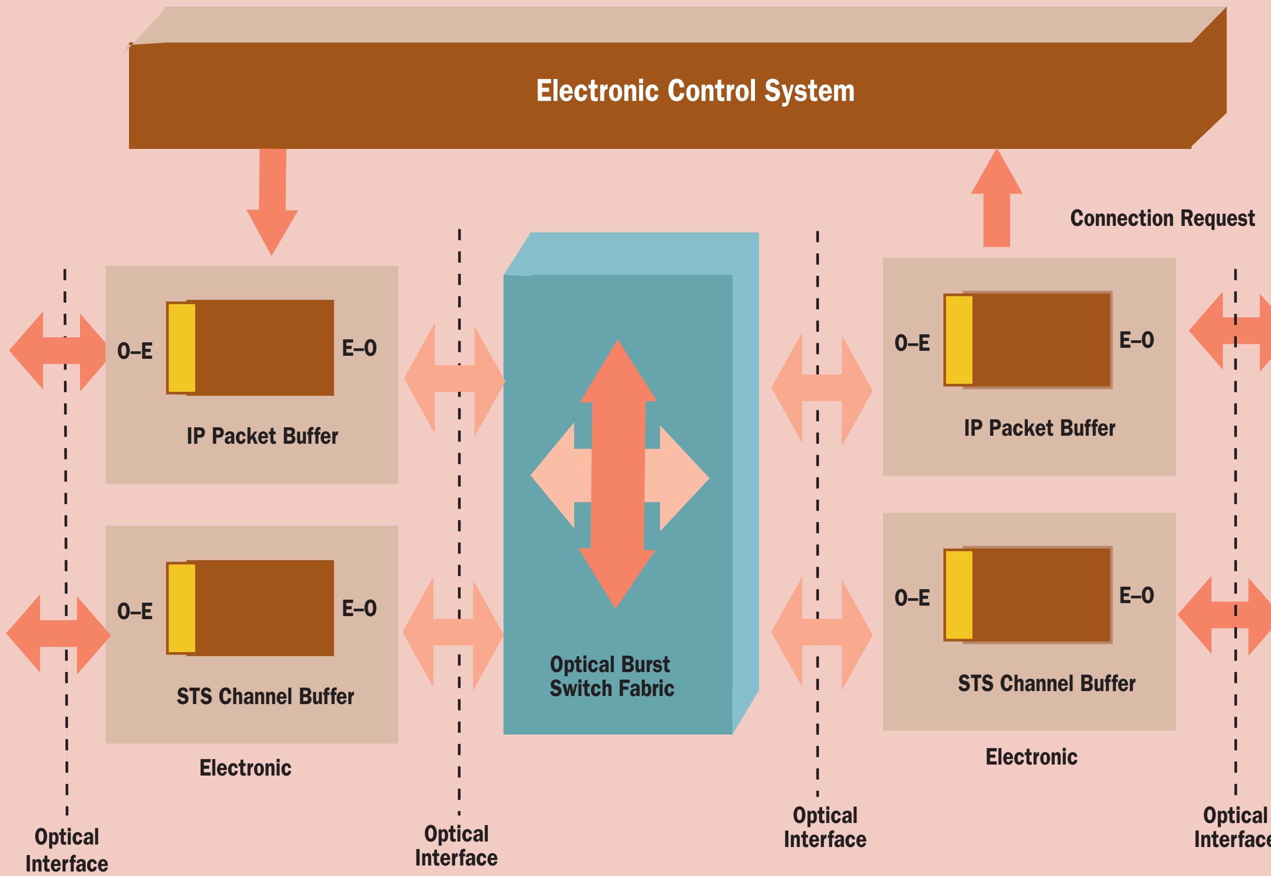

Optical UNI provides that security, and multivendor compatibility was demonstrated last year.

The emerging GMPLS standard would provide a finer level of detail and would bring the optical control plane into synch with the other MPLS protocols, according to Joe Berthold. GMPLS is actually a family of protocols, including a link-management protocol for neighbor discovery, signaling protocols for establishing links and restoration paths, and routing protocols for selecting paths through the network.

Most of these functions will use existing routing protocols modified for the requirements of circuit rather than packet networks. For example, the signaling protocols are based RSVP and CR-LDP, while the routing protocols use IS-IS, OSPF and BGP. Most vendors see GMPLS capabilities coming on line in the next 12 to 24 months.

Finally, ASON is at an early stage but will likely complement Optical UNI and GMPLS.

Conclusion

The optical networking industry is ready to take a significant leap forward—the vendors just need to find a carrier with money to spend. The speed of recovery will also vary according to which parts of the optical network are being served.

For example, the picture for WDM equipment is unclear at best; these systems originally were deployed to provide additional capacity on routes where the carriers had run out of fiber pairs. But the current estimate is that less than 5 percent of existing long-haul fiber is lit. Carriers are much more likely to focus on grooming and aggregation at the edge, rather than capacity in the core.

The carriers will have to digest these new economics, decide what services they wish to sell and then buy the equipment they need to make it happen. The outlook for optical networking equipment is a market delayed but not defeated□

Companies Mentioned In This Article

AcceLight Networks (www.accelight.com)
Bandwidth9 (www.bandwidth9.com)
Ceyba Corp. (www.ceyba.com)
Ciena (www.ciena.com)
Global Crossing (www. global crossing.com
iolon Corp (www.iolon.com)
Innovance Networks (www.innovance.com)
Nortel Networks (www.nortelnetworks.com)
Phaethon Communications (www.phaethoncommunications.com)
Xtera Communications (www.xtera.com)

Technology advances will continue to focus on the edge, rather than the core