

Optical Networking Technologies: What Worked and What Didn't

Rajiv Ramaswami, Cisco Systems

ABSTRACT

This article takes a broad look at a variety of optical network technologies that have been developed over the past two decades and comments on why some were successful in the real world while others have yet to make it beyond research laboratories or paper publications. We describe each technology briefly and then offer possible explanations for its rapid or slow adoption by the industry.

INTRODUCTION

Fiber optic communications and networking are ubiquitous today in both enterprise data networks and service provider (primarily telephone companies [telcos] and cable operators) networks. Since the early 1980s, billions of dollars have been invested by industry, complemented by academic research, to develop a wide range of technologies, and billions more have been spent on deploying some of these technologies. The purpose of this article is to take a retrospective look at many technologies and architectures explored and examine why some were successful and others were not. For the purposes of this article, "success" is defined as commercial adoption leading to revenue generation. Perhaps more controversially, we also speculate on whether some of the technologies being worked on will be successful or not in the future, recognizing that this is one person's admittedly biased view of the world.

For a technology to be successfully adopted in the commercial marketplace it needs to deliver on several fronts. Generally, it needs to reduce capital and operating costs or enable new revenue streams for its users. In optical communication, capital costs can usually be measured in terms of a simple cost per bit per mile metric. Operating costs are harder to quantify, but some measurable metrics are reductions in power consumption, footprint, and labor costs. On top of this, the technology in question needs to deliver on these fronts better than other alternatives, and be within its window of opportunity from a time to market perspective.

APPLICATIONS

Clearly the two major successes in fiber optic communication have been enterprise data links, and service provider transmission links and networks. Optical fiber is the preferred medium for transmission for data rates larger than a few hundred megabits per second over distances more than tens of meters due to its near-perfect transmission properties, including low attenuation over a multi-terahertz bandwidth window, immunity from interference of most kinds, and requiring no maintenance over a very long life span. To date, no other technology has appeared on the horizon that can compete with these attributes.

Enterprise data links using a variety of protocols (100 Mb/s Ethernet, Gigabit Ethernet, 10 Gigabit Ethernet, Fibre Channel, etc) are widely deployed. Early precursors were the 100 Mb/s fiber distributed data interface (FDDI) and IBM's 200 Mb/s Enterprise Serial Connection (ESCON). The majority of these operate over the widely deployed multimode fiber plant found in enterprises.

Service provider transmission networks operate over single-mode fiber, which enables higher bandwidth transmission over longer distances. While starting in the late 1970s at a humble 45 Mb/s per fiber, today's optical fiber transmission systems can support a couple of hundred wavelengths using wavelength-division multiplexing (WDM), each operating at up to 40 Gb/s, all over a single fiber. The predominant applications have been in long-haul intercity type links as well as trans-oceanic submarine links. The late 1990s saw huge investments in this arena, leading to significant overcapacity, something from which the industry is still recovering. Today, more fiber-based systems are being deployed in metropolitan networks, and there is a huge push worldwide toward bringing fiber closer to individual homes and businesses.


Fiber to the home has been talked about since the mid-1990s but has not happened quite yet. Many factors have impacted this. One was the huge capital investment required to build out the fiber plant. Another was the lack of end-user bandwidth demand. A third was lack of competitive pressures on the telephone companies. A final factor was the effect of telecom regulation

requiring the incumbent local exchange providers (the local telephone monopoly) to *unbundle* their local plant (i.e., allow competing service providers access to the incumbent's plant at pre-determined prices). The availability of appropriate technology (low-cost outside plant network equipment) was perhaps the least of these impediments.

The factors listed above have changed dramatically in the past couple of years. Video (broadcast, high definition, and on demand) has emerged as the main bandwidth driver, with broadband internet access being a second. United States cable companies have deployed a hybrid fiber coax network architecture that is better suited for video delivery than the existing telco network, and are aggressively deploying video services, creating competitive pressures for the telcos. Today's U.S. regulatory environment no longer requires incumbents to unbundle new broadband network deployments. As a result, the incumbent U.S. telco providers have all announced significant plans to deploy fiber close to the home. Finally, a variety of technologies such as passive optical networks (PONs) are commercially available to support this deployment. Elsewhere, notably in Japan, fiber to the home deployments have already ramped up, and many European cities are looking at deployments over the next few years. Perhaps 2005 will go down in history as the year that fiber to the home really started to take off.

Over the past decade, many companies developed free-space optical (FSO) systems as an alternative to wired or wireless options to provide broadband connectivity to businesses. These point-to-point systems operate by transmitting optical signals over the air at speeds up to 1 Gb/s over a distance of several kilometers. These systems, however, are highly susceptible to rain and fog, which can significantly increase signal attenuation and reduce transmission distances to a few hundred meters. Today they occupy a small niche in the overall access market, compared to other wired (digital subscriber line [DSL], cable, fiber) or wireless alternatives.

One last application is optical interconnects for shorter distances, such as between racks or shelves in a multishelf system or between multiple line cards within a system. Some large multishelf systems such as digital crossconnects and terabit routers do use optical interconnects, but it has remained a niche application. The technologies deployed are not that different from standard data communications links for the most part. In some cases multiple transmitters and receivers are integrated within a package along with ribbon fiber cables to provide large-bandwidth parallel optical interconnects. Optical interconnects have not been successful in displacing electrical interconnects for card-to-card interconnects across a backplane, primarily because electrical serializer-deserializer technologies have continued to scale well here, extending to multiple gigabits per second per trace today. With these types of bit rates, electrical backplanes can provide hundreds of gigabits per second to each line card slot in a system, sufficient for most applications.



■ **Figure 1.** Form factors for different types of optical transceiver modules. Generally, the modules keep getting smaller and the bit rates higher.

ETHERNET, FIBRE CHANNEL, AND SONET/SDH

Today's enterprise networks are built primarily using Ethernet. The annual market for Ethernet equipment is close to \$15 billion. Ethernet works over any transmission medium, including copper, wireless, and optical fiber links. Ethernet adoption has been driven by the explosive need for data networking in enterprises, its first-mover advantage and simplicity compared to some of the other network options such as token ring, FDDI, and asynchronous transfer mode (ATM), and continued innovation in silicon application-specific integrated circuits (ASICs) driving higher and higher bandwidths and densities. Many high-end enterprise networks also use Fibre Channel for their storage area networks.

A variety of optical interfaces exist for 100 Mb/s, Gigabit, and 10 Gigabit Ethernet. The Ethernet and Fiber Channel market helped drive the miniaturization of optical transceiver modules, including the transition to user-pluggable modules. Today's Ethernet and Fiber Channel switches use a wide variety of pluggable optical transceiver modules at various rates, as shown in Fig. 1. These pluggable modules provide several benefits: they reduce up front expenditure as multiport line cards can be deployed first and the modules inserted later as needed. Each port can be customized to a given bit rate and reach by inserting the appropriate optical module, providing flexibility to the end user as well as reducing R&D costs for network equipment providers. These modules are now making their way into the service provider market as well.

Today's service provider networks are built primarily using synchronous optical network/synchronous digital hierarchy (SONET/SDH). The annual market for SONET/SDH equipment is over \$5 billion. SONET/SDH is optimized for multiplexing circuit-switched low-speed streams (e.g., DS1, DS3), and provides excellent operations, administration, maintenance, and provisioning capabilities important in service provider networks. But there are significant changes in the offing. Most applications and services are moving toward Ethernet. Ethernet-based business services are growing, and most video is now digitally transmitted over Ethernet. As a result, the service provider network infrastructure is in the process of making a transition from SONET/SDH to Ethernet.

WDM NETWORKS

An OADM is an element that allows some wavelengths to be dropped and added while allowing the remaining wavelengths to pass through optically, without an OEO conversion. The alternative is to demultiplex all the wavelengths, convert them to the electrical domain, and do the local add/drop electrically.

WDM has enjoyed tremendous success in the marketplace, with today's annual WDM equipment market sitting at around \$2.5 billion. In WDM signals at multiple wavelengths are combined and transmitted on a single fiber, and demultiplexed at the other end. Each wavelength can carry data at tens of gigabits per second (today). WDM, combined with optical amplification, enabled dramatic improvements in the unregenerated bandwidth distance product, providing huge reductions in the cost per bit per mile of transmission. In addition, WDM systems provide orders of magnitude increases in the bandwidth that can be delivered over an existing fiber plant, and enable a significant time-to-market improvement in service delivery, compared to the alternative of laying additional fiber to provide more bandwidth. WDM systems are deployed widely in long haul service provider networks, and are increasingly being deployed in metro service provider networks and for enterprise data center connectivity applications.

Many wavelength bands are available for transmission in single-mode fiber. Many early systems used the 1.3 μ band over standard single-mode fiber, where the fiber loss is 0.5 dB/km and there is no chromatic dispersion. Today's WDM systems operate in the 1.55 μ band, primarily because practical optical amplification is available only in this band. This band also offers a lower attenuation (0.2 dB/km) but we do have to deal with chromatic dispersion (using dispersion compensation). Several bands exist within the 1.55 μ band: C (1530–1565 nm), L (1565–1625 nm), and S (1460–1530 nm).

The Erbium doped fiber amplifier (EDFA) was a key enabler for WDM, dramatically altering its economics and increasing its rate of adoption. A single EDFA amplifies multiple wavelengths over the entire C or L band (EDFAs unfortunately do not work in the other bands). In a typical transmission system, EDFAs are inserted every 80–120 km to amplify the signal, and today's systems can transmit over thousands of kilometers before regeneration (optical-to-electrical-to-optical) conversion is needed. Compare this to having to insert a regenerator every 80 km for each wavelength — this was the status quo prior to the development of EDFAs. EDFAs provided huge cost and reliability improvement. Also, unlike regenerators, EDFAs are transparent to the bit rate and protocols employed by the signals, providing flexibility for mixing and matching different traffic types and allowing future network upgrades. One practical example is that many of the wavelengths lit up in systems today are 10 Gb/s SONET/SDH, whereas we are seeing more 10 Gb/s Ethernet traffic today, which can be carried over the same systems by changing the electronics at the endpoints rather than throughout the network.

WDM can be categorized as coarse or dense. Coarse WDM systems typically provide 4–16 wavelengths spaced tens of nanometers apart. Because the channels are relatively far apart, coarse WDM can use cheaper transmitters and filters than dense WDM. However, the wavelengths are spaced so far apart that they do not

fit within the EDFA bandwidth. Also, today at 1 Gb/s, there is a significant cost differential between coarse and dense WDM transmitters, but the difference is not as significant at 10 Gb/s. Therefore, coarse WDM is good for low-cost, low-bit-rate, short-distance, unamplified applications. Coarse WDM typically gets used in campus applications and potentially also in the access part of a service provider network; these systems typically operate at about 1 Gb/s/wavelength. Dense WDM (DWDM) can provide hundreds of wavelengths spaced less than a nanometer apart within the available amplifier bandwidth. Dense WDM is used in metro core and long haul networks as well as in high-end enterprise data center interconnects.

There has been a lot of research into enabling these different transmission bands and also trying to pack wavelengths tightly together in dense WDM systems. Most deployed WDM systems use the C band and wavelength spacings of 50 or 100 GHz. This provides up to 80 wavelengths at 10 Gb/s each or 40 wavelengths at 40 Gb/s each, usually sufficient for most applications.

Equipment vendors have developed so-called hyper WDM systems with 25 GHz wavelength spacing, as well as systems that operate in the L and S bands. These have enjoyed limited commercial success for a variety of reasons. The simpler 50/100 GHz spaced C-band systems provide sufficient capacity for most practical applications at lower first installed costs. L band EDFAs are available but cost more than C band EDFAs because they do not operate as efficiently. And EDFAs are not available in the S band, so more expensive Raman amplifiers will need to be used in this band. Moreover, by the time a provider comes close to exhausting the C band capacity of the simpler systems, newer-generation systems are typically available that offer even lower cost per bit. Finally, because of the overinvestment during the late 1990s, there is currently an abundance of long haul fiber available to light up additional capacity as needed.

L band systems have enjoyed some limited deployment, primarily in Japan over dispersion shifted fiber. In Japan there is a large installed base of this type of fiber, which has a chromatic dispersion zero in the C band (and a small amount of dispersion in the L band) Unfortunately, not having any chromatic dispersion is bad for WDM systems as fiber nonlinearities induce severe degradation (these effects are significantly reduced when dispersion is present). So the C band cannot be used for WDM in dispersion shifted fiber, prompting the Japanese to deploy L band systems.

The majority of WDM deployment has occurred in the form of point-to-point links with amplifiers in between as needed. The WDM *lightpaths* (an end-to-end WDM channel) are static. Once set up, they remain in place, essentially forever.

The network topology is becoming more of an all-optical network, incorporating network elements such as optical add/drops and optical crossconnects, as shown in Fig. 2. Today, this is primarily being implemented in the form of simple ring topologies, particularly in metro networks, but there have been a small amount of

all-optical mesh networks also deployed in long haul networks. There is also a desire to make the lightpaths more dynamic, a topic discussed later.

OPTICAL ADD/DROP AND RECONFIGURABLE OPTICAL ADD/DROP MULTIPLEXERS

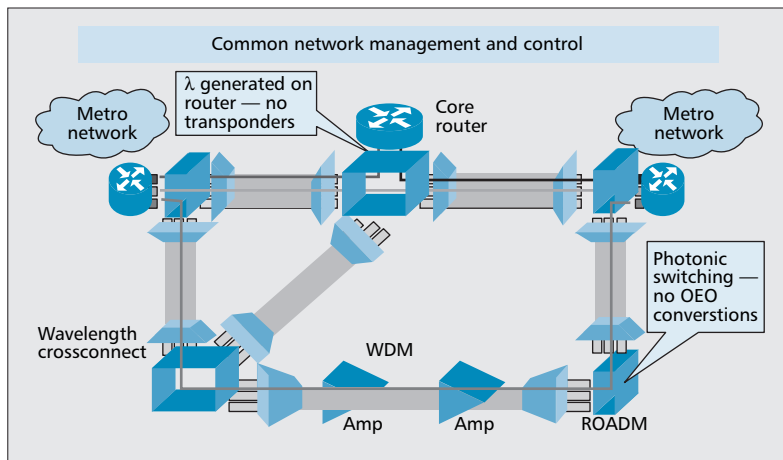
An optical add/drop multiplexer (OADM) is an element that allows some wavelengths to be dropped and added while allowing the remaining wavelengths to pass through optically, without optical-electrical-optical (OEO) conversion. The alternative is to demultiplex all the wavelengths, convert them to the electrical domain, and do the local add/drop electrically. As we saw earlier, OEOs tend to be the most expensive parts of WDM systems, so reducing their number reduces the network cost. For example, if at a site we need to drop two out of, say, 40 wavelengths passing through, a two-channel OADM is significantly cheaper than having to demultiplex all 40 wavelengths and use 40 receivers backed up by 40 transmitters.

While OADMs do provide this cost benefit, their application in real networks has been limited by the operational complexity associated with using them. Early OADMs were all fixed configuration devices, in that you had to determine the set of wavelengths to be dropped a priori and deploy the appropriate filters to support that. Any changes to this plan could result in a significant redesign of the entire network and cause service outages, making this an operational nightmare for service providers. In addition, these OADMs did not provide power equalization among the different wavelengths or adequate monitoring capabilities, and did not come with good network planning tools to support their deployment.

Today we have a new generation of OADMs, called reconfigurable OADMs (ROADMs), that address all the issues described previously. ROADMs allow any wavelength to be dropped and added without impacting other wavelengths, provide power balancing and monitoring, and come with good planning tools. Traditionally, ROADMs have cost more than fixed OADMs; however, network equipment vendors have (for now) priced ROADMs to be comparable to their fixed counterparts, stimulating a rapid deployment ramp over the past year. ROADMs are rapidly becoming ubiquitous in both metro and long haul deployment, and over 2000 ROADM nodes have been deployed as of this writing.

OPTICAL CROSSCONNECTS

ROADMs are perfect network elements for ring and linear network topologies, but cannot handle mesh nodes. While most all-optical networks use ring and linear topologies, long haul networks tend to have mesh topologies, and some providers are expressing a desire to build out metro mesh networks as well. Mesh nodes require a wavelength crossconnect (WXC), as



■ **Figure 2.** An all-optical WDM network with integrated wavelength generation on attached routers, using ROADMs and optical wavelength crossconnects.

shown in Fig. 2. A WXC performs the same function as an ROADM in a ring node, switching a wavelength from one input port to another output port independent of the other wavelengths. In this configuration a WXC integrates wavelength demultiplexers and multiplexers along with some switching. A typical device may have 4–8 input ports and output ports, with each port capable of handling 32–40 wavelengths. These types of WXC are just emerging today. They have been deployed in a few long haul network nodes but could find broader application in both metro and long haul networks over the next few years, and will be tightly integrated with WDM systems, just like ROADMs.

During the late 1990s, large investments were made in developing large-scale standalone optical crossconnects, with hundreds to thousands of ports. At that time large service providers had plans to deploy multiterabit nodes and anticipated needing these large crossconnects to switch the capacity in these nodes. As reality set in with the collapse of the optical bubble, it became clear that there would be no near-term real applications for such large-scale crossconnects. Today people are trying to identify other niche applications for these crossconnects, such as automatic patch panels for interconnecting fibers in large installations.

TUNABLE LASERS

Tunable lasers address two important problems in WDM networks. They eliminate the operational cost associated with having to manufacture and stock multiple part numbers to address different wavelengths by component suppliers, equipment makers, and the ultimate service provider or end-user customer; it is also quite difficult to predict the demand for each wavelength in the network as traffic growth can be unpredictable. Tunable lasers also allow connections to be provisioned dynamically on demand without manual intervention, when coupled with ROADMs and WXC. 10 Gb/s tunable transmitters are available today at a modest premium over fixed, and are a recent success story, driven primarily by the first application. At lower




Figure 3. A 3-wavelength network example to illustrate the need for wavelength conversion. The network in a) has no wavelength conversion. Although a free wavelength is available on AB and BC, we cannot set up a lightpath from A to C. In b), if wavelength conversion is possible at node B, the new lightpath can be set up using different wavelengths on links AB and BC.

speeds, the premium for tunable compared to fixed is still too high to justify their deployment. Much of this market is served by the pluggable transceivers described earlier.

BACK TO OEOs?

WDM systems to date have focused on increasing the unregenerated bandwidth distance product. This in turn reduces the overall cost per bit per mile. One important assumption built into this model is that the cost of OEO conversions is relatively very high, and that it is therefore beneficial to maintain the signal in the optical domain as much as possible. However, if we can reduce the cost of OEOs dramatically, this assumption is no longer true. In this case we may well go back to the days of deploying more OEOs rather than trying to extend the transmission reach. OEOs also provide good signal monitoring capabilities and reset the optical impairments. However, they generally also have poorer reliability compared to all-optical devices, take up more floor space and electrical power, increasing the operating costs, and generally continue to be protocol- and bit-rate-specific devices.

There are two approaches toward realizing low-cost OEOs. The first is to continue to miniaturize optical transceiver modules and exploit the volumes found in the enterprise market. For example, 1 Gb/s, 2.5 Gb/s, and 10 Gb/s DWDM pluggable modules are available today. Another is optical integration, where multiple elements, such as lasers, modulators, semiconductor amplifiers, receivers, multiplexers, and demultiplexers, are integrated together.

WAVELENGTH BLOCKING AND WAVELENGTH CONVERSION

In the all-optical network shown in Fig. 2, a lightpath from one node to another is established by setting up a wavelength between the two nodes. One of the constraints imposed by the network is that the same wavelength must be used along all the intermediate links in this lightpath. Depending on the other existing lightpaths in the network, it may not be possible to find the same wavelength free on every link in the network, as shown in Fig. 3. Instead, if it were possible to convert wavelengths along the path, we would be able to set up lightpaths that could not be set up otherwise.

Much has been said and written about this wavelength blocking problem and the benefit of having wavelength conversion in the network. Most of the theoretical work focuses on determining the blocking probability of a connection in networks where the lightpaths are assumed to be dynamically established and taken down, exploring a variety of wavelength assignment algorithms and locations and types of wavelength converters. Unfortunately, this is not the way optical networks operate today. Lightpaths are mostly static, and if a demand is made, the lightpath must be set up and cannot be blocked, and additional capacity is added as needed to support the demand.

Wavelength converters can reduce the number of wavelengths required to support a given traffic demand in the network. In practical terms, this means we can get better utilization out of a WDM system before we run out of wavelengths in the network and have to light up additional fibers. If the network is lightly loaded, for instance, say 4 out of 40 wavelengths are in use, then there is no need to add wavelength converters as we can simply light up a fifth wavelength to set up a new connection. As the network gets heavily loaded, it becomes increasingly difficult to find a common free wavelength on all the links in the desired path. At this time, a cost and time-to-market trade-off needs to be made in determining whether to add wavelength converters to continue to support additional traffic on the network or light up a new fiber and bring on more wavelengths.

Wavelength converters are rarely needed in today's networks because:

- Many of them have relatively low utilization levels.
- Many all-optical networks are simple topologies, such as rings, where some up front planning and knowledge of traffic demand can be used to plan the wavelength allocation for connections.
- Many connections tend to be protected in rings and take up a wavelength all the way around the ring, making wavelength converters moot.

In the rare case where wavelength converters are needed, they are manually added to the network at the desired locations. (Note that wavelength converters are used at the edge of the network to adapt incoming non-WDM signals into WDM signals — a different and very practical need for these devices.)

Today wavelength converters are essentially OEOs that convert signals back into the electrical domain and retransmit them on a different wavelength. All-optical wavelength converters have been explored in research laboratories, but do not have the cost or performance of OEOs.

OPTICAL PROTECTION

Resilience is an important part of network design, as networks performing mission-critical functions are expected to be up at least 99.999 percent of the time (that is a down time of less than 5 min/year). Protection switching is an important part of enabling this resiliency. The goal of protection switching is to detect failures and reroute traffic around these failures as

quickly as possible, typically ranging from within tens of milliseconds to several seconds.

Protection schemes can be implemented at any of the network layers. A variety of protection schemes have been standardized for SONET/SDH, all of which provide restoration in less than 50 ms after a failure has been detected. IP networks typically employ IP rerouting, which can take seconds to complete, but newer restoration schemes using multiprotocol label switching (MPLS) can restore traffic in tens of milliseconds. Ethernet switches typically use the Spanning Tree Protocol to converge routes after a failure, and this can take seconds. Resilient packet rings (RPR) providing 50 ms restoration times have been standardized for IP and Ethernet networks.

At the optical WDM layer, a variety of protection schemes have been developed, starting from simple point-to-point dedicated protection schemes to more complex shared protection schemes including ring and mesh protection mechanisms. Protection at the optical WDM layer does provide some benefits, including the possibility of detecting failures quicker, restoring multiple higher-layer connections quickly, and being able to share the optical bandwidth more efficiently among multiple higher-layer connections (for shared protection schemes). However, higher-layer restoration can potentially be more service-aware and manage the optical bandwidth better. A higher-layer scheme could use the entire bandwidth during normal operation and do a graceful failover of traffic by dropping low-priority traffic in case of failures.

To date, simple dedicated protection schemes have been deployed in some applications, particularly when no other form of rapid protection is available in the layers above the optical layer. This is driven by the complexity of shared protection ring and mesh schemes, a lack of standardization at the optical WDM layer, and the availability of excellent protection mechanisms at the higher layers. This will likely continue to be the case.

UNIFIED CONTROL PLANE

Traditionally, optical networks have been managed in a centralized fashion using network management systems to provision services, provide fault and performance monitoring, and assist in network maintenance functions. Since the late 1990s a lot of work has gone into trying to establish a distributed unified control plane (UCP) to do these functions. The UCP employs generalized multiprotocol label switching protocol (GMPLS). The main drivers behind this effort are to:

- Enable service providers to offer new dynamic bandwidth services
- Reduce network operations costs
- Promote multivendor interoperability

However, each of these areas has its challenges. Service providers are trying to constantly innovate on services, but it does not appear that the dynamic bandwidth aspect of it is very compelling for very-high-capacity connections. For instance, it is not clear that a service provider can share this bandwidth among multiple users, because requests for these types of connections cannot be denied.

Another related application is the ability of a device attached to the optical network, such as a router, to request additional bandwidth from the network or re-vector bandwidth used by it based on, say, load conditions. Today, for the most part, links between routers are static and created at the time the network is planned rather than on the fly.

Even for static services, it is possible that UCP may enable service providers to turn up services faster than relying on a management system, improving their response time, providing higher revenue, and giving them a competitive advantage. However, these items are hard to quantify.

A similar argument can be made with respect to operations cost. UCP could reduce the labor required to provision and maintain connections. A solid business case proving this point could help trigger deployment.

Many multivendor interoperability demonstrations have taken place, but we are still a long way from completely defining the standards to enable full interoperability. For example, UCP is particularly challenging for optical networks in which it needs to keep track of optical impairments as part of the routing decision. Without compelling applications and business cases, it is hard to justify the amount of effort required to make this happen.

The first real applications for UCP may come into play as IP and WDM networks get more tightly integrated. Equipment vendors are increasingly integrating WDM interfaces directly on routers and having them tied to the rest of the WDM layer in a single flat network. Many new networks are being planned in this fashion. In these types of networks, elements of UCP are essential to provide topology auto discovery across the router and the other WDM elements. UCP can also be used to signal from the router to the WDM elements to alert them of failure conditions detected by the router that the WDM elements may not be able to detect themselves. For instance, if an optical link degrades, the WDM elements may be unable to detect this degradation, but the router, having visibility into the bitstream, can. The router can then signal to the WDM layer to switch over to another protection path. Finally, even though we have a flat IP+WDM network, the IP and WDM equipment may be managed by separate management systems due to their legacies, and UCP can help bridge between these two systems.

WDM LOCAL AREA NETWORKS

In the late 1980s and early 1990s it was thought that one of the first applications of WDM would be in enterprise local and metropolitan area networks to deliver hundreds of megabits to high-end computers. There was a fair bit of interest in broadcast-and-select architectures, where a station would transmit at a given wavelength that would be broadcast using passive couplers to all the other stations. A receiver would tune to the appropriate wavelength to receive the signal. Some early prototypes were built that provided on-demand (within milliseconds) connectivity between stations at a few hundred megabits per second. However, they remained prototypes for

Higher-layer restoration can potentially be more service-aware and manage the optical bandwidth better. A higher-layer scheme could use the entire bandwidth during normal operation and do a graceful failover of traffic by dropping low-priority traffic in case of failures.

Continued advances in application-specific integrated circuit technology has enabled electronic switches and routers to scale well into multi terabit-per-second capacities while supporting interfaces at up to 40 Gb/s today, making the prospect for optical packet switching rather dim.

two good reasons: high cost and the inability to provide packet switching. The notion was that packet switching would be done by rapidly tuning from one wavelength to another in sub-microsecond time frames, coordinated by a media access protocol.

However, even today we are extremely challenged to accomplish stable sub-microsecond switching between wavelengths and get to practical cost points compared to other technologies, primarily Ethernet. This did not prevent researchers from exploring a huge variety of media access protocol variants for these types of networks, none of which, unfortunately, helped solve the basic issues of this technology. WDM broadcast-and-select networks will probably never make it out into the real world.

As history has shown, Ethernet has been the hands-down winner here, constantly extending its capabilities with respect to capacities and physical media interfaces while being extremely cost effective for enterprises to deploy.

OPTICAL PACKET SWITCHING

As we have seen so far, optical networks have offered primarily static connections. For the most part, packets are transmitted over these connections, with the packets being switched by electronic Ethernet switches or IP routers. A lot of work has gone into trying to build optical packet switches. The rationale here is that as bandwidths continue to increase, there could come a point in time where optical packet switching could be more practical and economical than electronic packet switching.

However, despite many years of research, there continue to be major impediments to making optical packet switching a reality. Large optical switches that can switch in microseconds do not exist, and the smaller ones that can suffer from high loss, polarization dependence, and are expensive to fabricate. Optical random access memory does not exist, and the only buffering available is via fiber delay lines that simply delay the packet for a fixed duration depending on the length of fiber used, and are therefore lossy. Also, all-optical header processing techniques are still primitive. Therefore, packet headers must continue to be processed electronically. Unfortunately, the sophistication of header processing required is increasing, because routers and switches today are increasingly looking at higher-layer tags in the packet to handle security and application-layer functions. Finally, continued advances in ASIC technology has enabled electronic switches and routers to scale well into multi-terabit-per-second capacities while supporting interfaces at up to 40 Gb/s today, making the prospect for optical packet switching rather dim.

OPTICAL BURST SWITCHING

Optical burst switching (OBS) is a technique that falls between packet switching and circuit switching. The idea is to transmit data in units of bursts, which can be thought of as rather long packets with durations of, say, milliseconds to even seconds. OBS is perhaps easier to implement than optical packet switches because networks can be designed without optical buffers.

Bandwidth can be reserved ahead of time and with proper synchronization, the burst can be transmitted through the network without much intermediate buffering. Alternatively, bursts could be transmitted without reserving the bandwidth, in which case they can be dropped if there is contention. However, OBS is significantly more complex to implement than static or circuit-switched optical networks.

An appropriate question to ask is: what practical network/application problem does OBS solve? Given that electronic packet switching will continue to exist for the foreseeable future, what additional benefit does OBS provide as a layer under this, compared to a static or circuit-switched optical layer? One argument may be that OBS can improve the bandwidth utilization of the optical layer compared to a static optical layer. However, we could argue that we can also achieve similar efficiencies by appropriately characterizing the "burstiness" of the packet traffic offered to the optical layer and then dimensioning the right amount of bandwidth to support it, rather than incurring an expensive additional packet multiplexing layer underneath it. Another argument is that OBS promises lower cost as the signal is kept in the optical domain. However, optical circuit switching in conjunction with electrical packet processing at the ends of the circuit, provides the same — or better — cost savings.

Finally, given that electronics keeps pushing toward higher data rates (currently 40 Gb/s), optical packet and burst switching will probably have to prove themselves at even higher transmission rates, say, 100 Gb/s. At these rates, transmission impairments are challenging even for a static optical layer, but become even more difficult to deal with when the network itself is dynamic, where we need to switch between paths that have significantly different impairments and delays. All in all, it is hard to imagine OBS playing a role in real networks. Unfortunately, most of the research on OBS focuses on proposing various protocols and analyzing their performance, rather than dealing with the practical aspects of whether the technology makes sense and can be made to work in practical networks.

OPTICAL CODE DIVISION MULTIPLE ACCESS

Code-division multiple access (CDMA) is an undisputed success for cellular and military applications. Like time- or frequency-division multiplexing, CDMA allows multiple users to share a common spectrum. Each user uniquely encodes his/her signal in either the time or frequency domain by spreading his/her transmitted spectrum so that interference from other users is minimized. The receiver needs to know the unique encoding sequence to be able to decode the signal. With optical CDMA, a signal at, say, a 10 Gb/s data rate would need to be encoded into an effective spectrum of, say, 100 GHz to provide sufficient immunity from interference from other users. At these bandwidths, the encoding and decoding must be done optically, and existing optical technologies are unable to give a big enough code space to isolate overlapping users. Transmission impairments at the encoded rate are much more limiting than at the

native data rate. Overall, optical CDMA appears to be an unlikely candidate to achieve commercial success.

SUMMARY

As we have seen, technologies that have been successfully adopted in the marketplace provide clear quantifiable benefits around which industry can develop business plans and fund deployments. Networking technologies include SONET/SDH, optical Ethernet, and WDM point-to-point links. All-optical WDM networks using ROADMs and tunable lasers appear to be on the road toward widespread deployment and could evolve to all-optical mesh networks using wavelength crossconnects. Fiber to the home (or close to it) in a variety of forms, including direct point-to-point fiber as well as PONs, is poised to become the next major success story for optical fiber communications.

Technologies that have yet to fulfill their promise generally fall into two categories. The first category consists of technologies that attempted to solve a real world problem and provided tangible value, but were upstaged by other technologies that could solve the problem better. Networking technologies in this category include large-scale optical crossconnects, all-optical wavelength converters, and WDM local area networks.

The second category is a bit more problematic in that it consists of fields where substantial efforts have been expended to make a technology work practically, when it is not even clear what problem the technology really addresses in the first place. This makes it difficult to create business cases, which slows or stalls indefinitely commercial funding and deployment. Some of these technologies are still early in their development cycles (e.g., the unified control plane) and could find successful adoption down the road. Others, such as optical packet switching, optical burst switching, and optical CDMA, have been worked on for many years and as of this

writing are unlikely to make it out of the laboratories.

FURTHER READING

There is a vast body of literature on various aspects of optical fiber communication. References [1–3] provide good overall starting points. Reference [4] covers recent developments in fiber to the home technologies and deployments. Reference [5] provides an overview of the various protection techniques available in the different network layers. See [6, 7] for details on the unified control plane. The theoretical aspects of OBS are covered in [8].

ACKNOWLEDGMENTS

I would like to thank the reviewers, Paul Green, Loukas Paraschis, Jim Theodoras, and Ori Gershtel for their constructive comments on this article.

REFERENCES

- [1] R. Ramaswami and K. N. Sivarajan, *Optical Networks — A Practical Perspective*, 2nd ed., Morgan Kaufmann, 2001.
- [2] J. Hecht, *City of Light: The Story of Fiber Optics*, Oxford Univ. Press, 1999.
- [3] I. P. Kaminow and T. Li, *Optical Fiber Telecommunications IV*, Elsevier Science, 2002.
- [4] P. E. Green, *Fiber to the Home*, Wiley, 2005.
- [5] J.-P. Vasseur, M. Pickavet, and P. Demeester, *Network Recovery: Protection and Restoration of Optical, SONET-SDH, IP, and MPLS*, Morgan Kaufmann, 2004.
- [6] G. Bernstein, B. Rajagopalan, and D. Saha, *Optical Network Control: Architecture, Protocols, and Standards*, Addison-Wesley, 2003.
- [7] A. Farrel and I. Bryskin, *GMPLS: Architecture and Applications*, Morgan Kaufman, 2005.
- [8] J. P. Jue and V. M. Vokkarane, *Optical Burst Switched Networks*, Springer, 2005.

BIOGRAPHY

RAJIV RAMASWAMI (rajivr@cisco.com) is a vice president and general manager at Cisco Systems, where he is responsible for Cisco's mid-range Ethernet switching and optical modules business units. Previously he led Cisco's optical business unit, and had a long career in optical networking at IBM Research, Tellabs, Xros, and Nortel Networks. He holds a Ph.D. in electrical engineering from the University of California at Berkeley.

Fiber to the home (or close to it) in a variety of forms, including direct point-to-point fiber as well as passive optical networks, is poised to become the next major success story for optical fiber communications.