WAVELENGTH-DIVISION MULTIPLEXING OPTICAL NETWORKS

EYTAN MODIANO
Massachusetts Institute of Technology
Cambridge, Massachusetts

1. INTRODUCTION

Communication networks were first developed for providing voice telephone service. Early networks were deployed using copper wire as the medium over which traffic was sent in the form of electromagnetic waves. As demand for communication increased, networks began to use optical fiber cables over which information was sent in the form of lightwaves. Thanks to the relatively low attenuation losses of optical fiber, transmitting information over fiber allowed for a significant increase in the transmission capacity of networks. For example, while transmission over copper was limited to only a few tens of thousands of bits per second (kbps), optical fiber transmission enabled data rates exceeding hundreds of millions of bits per second (Mbps). However, the relatively recent development of the Internet has resulted in a tremendous increase in demand for transmission capacity. More recently, developments in optical transmission technology have achieved data rates that exceed many billions of bits per second (Gbps). Even with these enormous data rates, demand still far exceeds the available network capacity. As a result, telecommunication equipment companies are constantly trying to develop new technology that can increase network capacity at reduced costs.

While fiberoptic technology resulted in a significant increase in a network's "bandwidth," or the amount of information that the network could send, the creation of the Internet resulted in an even greater demand for bandwidth. As demand for network capacity increased, service providers exhausted their available transmission capacity. One approach to alleviating fiber exhaust is to deploy additional fiber. This solution, however, is not always economically feasible. As a result, new technologies were developed to increase the transmission capacity of existing fiber.

The simplest approach is to increase the rate of transmission over the fiber (i.e., sending more bits per second). Since 1980, fiber transmission rates have increased from a few Mbps to nearly 100 Gbps. Since most users rarely need such high data rates, a network technology called *synchronous optical networks* (SONET) was developed to allow users to share the capacity of a fiber [1].

SONET is a technology for multiplexing a large number of low-rate circuits onto the high-rate fiber channel. The "basic" transmission rate of SONET is 64 kbps for supporting voice communications. SONET multiplexes large numbers of 64-kbps channels onto higher-rate datastreams. SONET defines a family of supported data rates. These data rates are often referred to as (optical carrier) OC-1, OC-3, OC-12, OC-48, and OC-192; where OC-1 corresponds to a data rate of 51.84 Mbps, or 672 voice circuits, OC-3 is 155.52 Mbps or 2016 voice circuits (3 times OC-1, OC-12 is 12 times OC-1, etc.). SONET uses time-division multiplexing (TDM) for combining traffic from multiple sources onto a common output. TDM multiplexes traffic from different sources by interleaving small "slices" of data from each source. Thus, if traffic from three OC-1 sources is being time division multiplexed onto an OC-3 transmission channel, each source would get access to the channel for a short period of time in a round-robin order, as shown in Fig. 1.

However, because of fundamental limits on optical transmission, the transmission capacity of a fiber cannot be increased indefinitely. Hence, to further increase the capacity of a fiber, a technology called wavelength-division multiplexing (WDM) was developed [1]. Wavelength division multiplexing allows transmissions on the fiber to use different colors of light (each color represents a different wavelength over which light propagates). Whereas in the first optical communications networks, light was transmitted through the fiber using a single wavelength, WDM permits light at multiple, different wavelengths, to be transmitted through a single fiber simultaneously. WDM is analogous to frequency-division multiplexing (FDM), which is often used for transmission over the airwaves.

In WDM systems, incoming optical signals are assigned a specific wavelength and then multiplexed onto the fiber. Moreover, such systems are bit-rate- and protocol-independent, meaning that each incoming signal can be carried in its native format and at a different rate. For example, a WDM system may support the transmission of multiple SONET signals on a single fiber, each operating at transmission rates of 10 Gbps (OC-192). As shown in Fig. 2, WDM systems are designed to operate in the low loss region of optical fiber, around the 1.5-μm band. Typically, wavelengths are assigned in this region with a separation of 25–100 GHz; and systems supporting anywhere from 80 to 160 wavelengths are presently being deployed.

The simplest approach to using WDM is to treat each wavelength as if it were on a separate fiber and continue

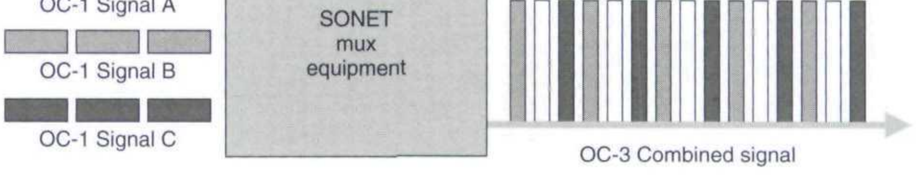

Figure 1. SONET time-division multiplexing.

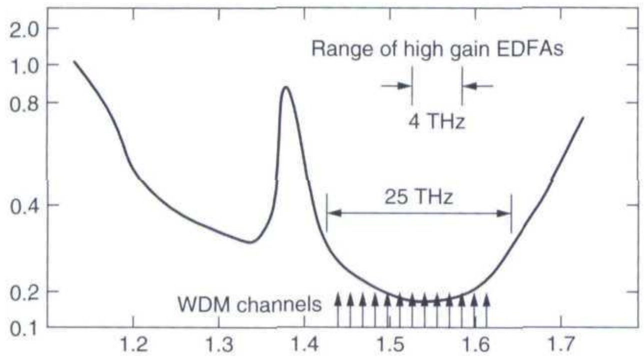

Figure 2. Wavelength-division multiplexing allows the transmission on multiple wavelengths within a single fiber.

to design the network using point-to-point links, as shown in Fig. 3. With this approach the available fiber capacity would in effect be increased by a factor that equals the number of wavelengths, and the network architecture would be largely unchanged. By itself, this approach leads to significant cost savings. First, in many cases existing fiber cannot meet demand, and WDM can help alleviate the fiber exhaust problem. In addition, through the use of erbium-doped fiber amplifiers (EDFAs), as shown in Fig. 3, a number of wavelengths can be simultaneously amplified. This leads to significant cost savings when compared to pre-WDM systems where each fiber would require its own amplification. Since the cost of amplification (and or regeneration) forms a large fraction of the overall network deployment cost, these cost savings through the use of EDFAs play a large role in the commercial success of WDM systems.

As described so far, from a network perspective, WDM systems are not vastly different from any other optical transmission systems. However, in addition to pure transmission technology, a number of WDM network elements have been developed that make it possible to design networks that actually route and switch traffic in the optical domain [2]. While these optical networking functions are rather limited, they have the potential of significantly enhancing network performance. In Section 2, we describe some of the basic WDM network elements. Subsequently, in Section 3 we describe architectures for future WDM optical local-area networks (LANs), and in Section 4 we describe architectural issues in the design of all-optical WDM wide-area networks (WANs). Since all-optical networks are not likely to emerge in the near future, we devote the last section of this article to discussing future WDM-based networks that use a combination of optical and electronic processing.

2. WDM NETWORK ELEMENTS

A number of optical network elements have been developed that allow for simple optical processing of signals. The simplest such device is a "broadcast star," shown in Fig. 4. In a broadcast star, each node is connected using one input and one output fiber. The fibers are then coupled or "fused" together so that a signal coming in from any input fiber will propagate on all output fibers. In this way, all nodes can communicate with each other. Because of its simplicity and broadcast property, a star has often been proposed as a suitable technology for optical LANs. In Section 3 of this article we describe WDM-based LAN architectures utilizing a broadcast star.

A somewhat more sophisticated device is a *wavelength router*, a passive optical device that combines signals from a number of input fibers and "routes" those signals in a static manner, based on the wavelength on which they propagate, to the output fibers. The operation of a wavelength router with four input fibers is shown in Fig. 5. Notice that the signal propagating on wavelength 1 of the first input fiber is routed to the first output fiber, while the signal propagating on wavelength 2 is routed to the second output fiber, and so on. Similarly, the signal propagating on wavelength 1 of the second fiber is routed directly to the second output fiber, while the signal on wavelength 2 is routed to the third output fiber, and so forth. A wavelength router is different from a broadcast star in that it separates the wavelengths onto different output fibers. Notice that the signals propagating on the first input fiber are routed onto the four output fibers so that the first wavelength is routed to the first output, the second to the second output, and so on. In this way, a wavelength router is commonly used in optical networks as a WDM demultiplexer. WDM demultiplexers can be used to form optical add/drop multiplexers (OADMs), that are able to drop or add any number of wavelengths at a node. OADMs play an important role in a network and can be designed in a number of ways (not necessarily using a

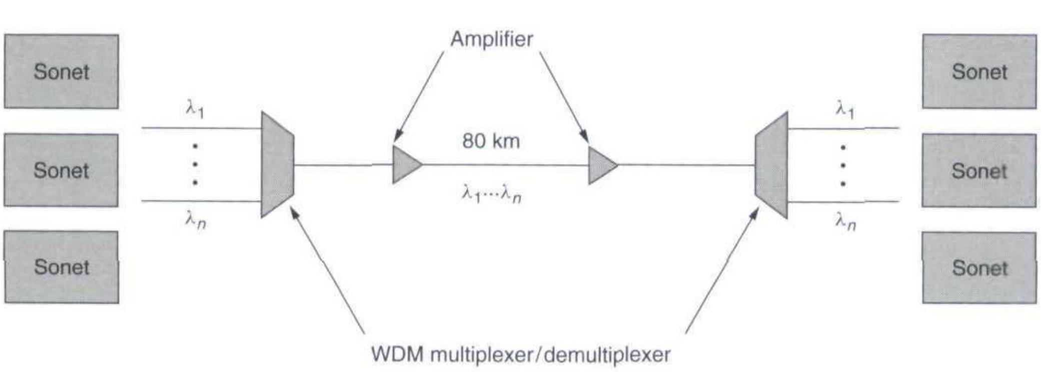

Figure 3. Typical SONET/WDM deployment.

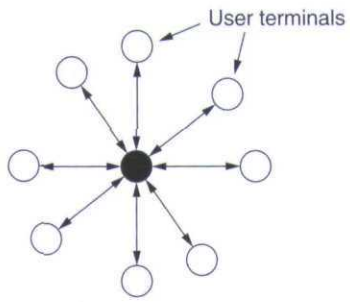

Figure 4. Broadcast star.

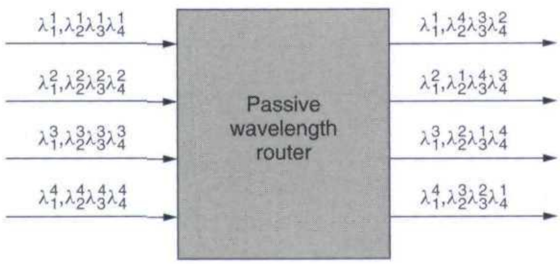

Figure 5. Wavelength router.

wavelength router); the operation of an OADM is shown in Fig. 6.

WDM demultiplexers can be used in conjunction with optical switches to form an even more sophisticated optical switching device known as a frequency-selective switch (FSS). Unlike a wavelength router that routes wavelength from input fibers onto output fibers in a static manner, a FSS is a configurable device that can take any wavelength from any input fiber and switch it onto any output fiber. In this way, a FSS allows for some flexibility in the operation of the network. The basic operation of a FSS is shown in Fig. 7. The signals traveling on each input fiber are demultiplexed into the different wavelengths. Each wavelength is then connected to a switching element that can switch any of the input fibers onto any of the output fibers. The outputs of the switch elements are then connected to WDM multiplexers that combine the signals onto the output fibers.

While a FSS adds significant functionality to the operation of the network, it requires the ability to dynamically switch optical signals from one input fiber onto another. Such switching can be accomplished optically using a number of techniques that are beyond the scope of this article. Optical switching technology, while rapidly progressing, is still relatively immature and costly. Alternatively, the signals can be switched in the electronic domain by first converting the signals traveling on each wavelength to electronics, switching the signals electronically, and retransmitting the signal onto the output fiber. Such optoelectronic switching is often less costly, and results in faster switching times, than does optical switching. However, in order to perform the optoelectronic conversion, the signal format (e.g., modulation technique, bit rate) must be known. This eliminates the signal "transparency" that is so desirable in optical networks.

A FSS adds flexibility to network operations by allowing wavelengths to be dynamically switched among the different fibers. However, notice that wavelength cannot be switched arbitrarily. For example, it is not possible for a FSS to route the signal propagating on a given wavelength from more than one input fiber onto the same output fiber. This limitation can be overcome by a *wavelength converter*, which can convert a signal from one wavelength onto another. Wavelength converters come in a number of variations. The simplest is a *fixed-wavelength converter*, which can convert a given wavelength onto another in a predetermined static fashion. More flexible converters can convert a given wavelength onto one of a number of wavelengths dynamically; such converters are known as *limited-wavelength converters*. The most flexible wavelength converters can convert any wavelength onto any other. Wavelength conversion can be accomplished in either the optical or electronic domain. Optical wavelength conversion is a rather immature technology primarily implemented in experimental laboratories; while electronic wavelength conversion suffers from the need for optoelectronic conversion and the consequent loss of transparency. Hence, while desirable, wavelength conversion in optical networks is still very limited.

Of course, in order to use WDM technology, one must be able to transmit and receive the signal on the different wavelengths. Transmission is accomplished using lasers that operate at a given wavelength, while reception is accomplished using WDM filters and light detectors. Typically lasers and filters are designed to operate at a single, fixed, frequency; such devices are commonly referred to as *fixed tuned devices*. Fixed tuned WDM transmitters and receivers again limit the capability and flexibility of an optical network because a signal that is transmitted on a given wavelength must travel throughout the network, and be received on that wavelength. Hence, without wavelength conversion, a node that has a transmitter that operates on a given frequency can

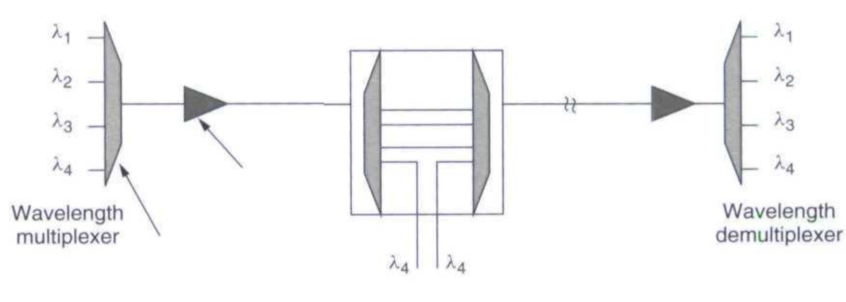

Figure 6. Optical add/drop multiplexer.

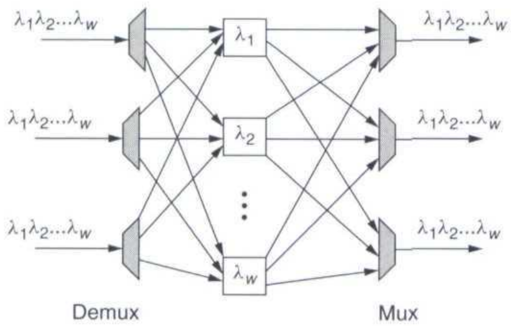

Figure 7. A frequency-selective switch.

communicate only with nodes that are equipped with a receiver for that frequency. Tunable transmitters and receivers are hence very desirable for optical networks; however, much like wavelength conversion, that technology is still at its inception. Tunable transmitters and receivers are often characterized according to the speed with which they can tune to different wavelengths. Slow tuning lasers, which can tune on the order of a few milliseconds, are now becoming commercially available, while fast lasers that can tune in microseconds are emerging.

3. WDM LOCAL AREA NETWORKS

Typically, local-area networks (LANs) span short distances, ranging from a few meters to a few thousands of meters. Because of the relatively close proximity of nodes, LANs are typically designed using a shared transmission medium. In this section we discuss WDM-based LANs, where users share a number of wavelengths, each operating at moderate rates (e.g., 40 wavelengths at 2.5 Gbps each).

Typically WDM-based LANs assume the use of a broadcast star architecture [3]. An optical star coupler is used to connect all the nodes. Each node is attached to the star using a pair of fibers: one for transmission and the other for reception. The star coupler is a passive device that simply connects all the incoming and outgoing fibers so that any transmission, on any wavelength, on an incoming fiber is broadcast on all outgoing fibers. In order for nodes to communicate, they must tune their transmitters and receivers to the appropriate wavelength.

A WDM LAN based on a broadcast star architecture can provide a transmission capacity that can easily exceed 100 Gbps. Perhaps the greatest reason preventing such systems from emerging is the cost of WDM transceivers. In order for a WDM LAN to allow flexible bandwidth sharing, both transmitter and receiver must be rapidly tunable over the available wavelengths. Transceiver tuning times that are smaller than the packet transmission times are desirable if efficient use of the bandwidth is to be obtained. With packets that are just a few thousands of bits in length, this calls for tuning times on the order of microseconds or faster. Present technology for fast tuning lasers is largely at the experimental stage; and while such lasers are slowly becoming commercially available, they are very expensive. Similarly, fast tuning receivers are also complex and expensive.

It is reasonable to expect that as the commercial market for these devices develops, their cost will decrease and they will become more widely available. However, in the near future, if WDM-based LANs are to become a reality they must limit the use of tunable components. WDM-based LANs are usually classified according to the number and tunability of the transmitters and receivers [4]. For example, a system utilizing one tunable transmitter and one tunable receiver is referred to as a TT-TR system. Similarly, a fixed tuned system would be referred to as FT-FR. Obviously, a FT-FR system can only use one wavelength if full connectivity among the nodes is desired. In order to provide full connectivity over multiple wavelengths, it is necessary that either the receivers or the transmitters be tunable. Systems employing either a tunable transmitter and a fixed tuned receiver (TT-FR) or a fixed transmitter and a tunable receiver (FT-TR) have been proposed in the past for the purpose of reducing the network costs.

Particularly attractive is the use of a fixed tuned receiver, because with a fixed tuned receiver all communication to a node is done on a fixed wavelength. Hence, this eliminates the need for any coordination before the transmission takes place. Of course, having a fixed tuned receiver means that nodes will have to be assigned to wavelengths in some fashion. For example, in an N node–W wavelengths network, N/W nodes can be assigned to receive on each wavelength. This, of course, creates a number of complications. First, when nodes are assigned to wavelengths in such a fixed manner, it is possible that certain wavelengths will be carrying a larger load than others and so, while some wavelengths may be lightly loaded, others may be overly saturated. In addition, such a network is complicated to administer because whenever adding a new node, care must be taken to determine on which wavelength it must be added, and a transceiver card tuned to that wavelength must be used.

In order to obtain the full benefit of the WDM bandwidth, a WDM-based LAN must have a TT-TR architecture. With this architecture, some form of transmission coordination is necessary for three reasons: (1) if two nodes transmit on the same wavelength simultaneously, their transmissions will interfere with each other (collide) and so some mechanism must be employed to prevent such collisions; (2) if two or more nodes transmit to the same node at the same time (albeit on different wavelengths), and if that node has only a single receiver, it will be able to receive only one of the transmissions; and (3) for a node to receive a transmission on a wavelength, it must know in advance of the upcoming transmission so that it can tune its receiver to the appropriate wavelength.

Most proposed WDM LANs use a separate control channel for the purpose of pretransmission coordination. Often, these systems use an additional fixed tuned transceiver for the control channel. Alternatively, the control and data channels can share a transceiver, as shown in Fig. 8.

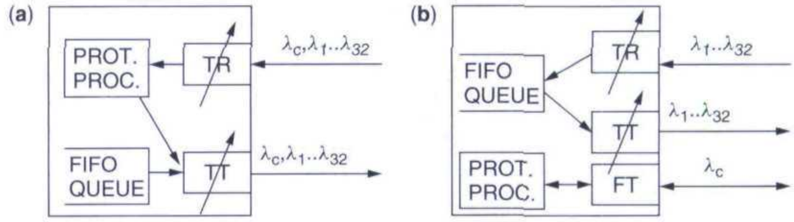

Figure 8. User terminals: (**a**) single turnable transceiver; (**b**) two-transceiver configuration.

In order for one node to send a packet to another, it must first choose a wavelength on which to transmit, and then inform the receiving node, on the control channel, of that upcoming transmission. A number of medium access control (MAC) protocols have been proposed to accomplish this exchange [5]. These protocols are more complicated than single-channel MAC protocols because they must arbitrate among a number of shared resources: the data channels, the control channel, and the receivers.

Early MAC protocols for WDM broadcast networks attempted to use ALOHA for sharing the channels [6]. With ALOHA, nodes transmit on a channel without attempting to coordinate their transmissions with any of the other nodes. If no other node transmits at the same time, the transmission is successful; however, if two nodes transmit simultaneously, their transmissions "collide" and both nodes must retransmit their packets. To reduce the likelihood of repeated collisions, nodes wait a random delay before attempting retransmission. When the load on the network is light, the likelihood of such a collision is low; however, with increased load such collisions occur more often, limiting network throughput. Single-channel versions of ALOHA have a maximum throughput of approximately 18%. A slotted version of ALOHA, where nodes are synchronized and transmit on slot boundaries, can achieve a throughput of 36%.

In a WDM system using a control channel, a MAC protocol must be used both for the control and the data channels. Early MAC protocols attempted to use a variation of ALOHA on both the control and the data channels [7]. In order for a transmission to be successful, the following sequence of events must take place: (1) the transmission on the control channel must be successful (i.e., no control channel collision), (2) the receiving node must not be receiving any other transmission at the same time (i.e., no receiver collision), and (3) the transmission on the chosen data channel must also be successful. In a system that uses ALOHA for both the control and data channels, it is clear that throughput will be very limited. It has been shown that systems using slotted ALOHA for both the data and control channels achieve a maximum utilization of less than 10% [7].

In view of the discussion above, a number of MAC protocols that attempt to increase utilization by coordinating and scheduling the transmissions more carefully have been proposed [4]. For example, the protocol described by Modiano and Barry [8] uses a simple master/slave scheduler as shown in Fig. 9. All nodes send their requests to the scheduler on a dedicated control wavelength, λ_C. The scheduler, located at the hub, schedules the requests and informs the nodes on a separate wavelength, $\lambda_{C'}$, of their turn to transmit.

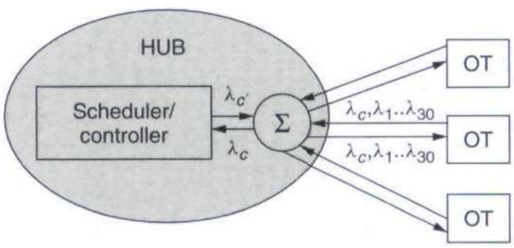

Figure 9. A WDM-based LAN using a broadcast star and a scheduler.

On receiving their assignments, nodes immediately tune to their assigned wavelength and transmit. Hence nodes do not need to maintain any synchronization or timing information. By measuring the amount of time that nodes take to respond to the assignments, the scheduler is able to obtain an estimate of each node's round-trip delay to the hub. This estimate is used by the scheduler to overcome the effects of propagation delays. The system uses simple scheduling algorithms that can be implemented in real time. Unicast traffic is scheduled using first-come–first-serve input queues and a window selection policy to eliminate head-of-line blocking, and multicast traffic is scheduled using a random algorithm [9].

We should point out, however, that despite their appeal, WDM-based LANs still face significant economic challenges. This is because the cost of WDM transceivers (especially tunable) is far greater than the typical cost of today's LAN interfaces. Since tunable WDM transceivers are just beginning to emerge in the marketplace, it is difficult to provide an accurate cost estimate for these devices, but it is certainly in the thousands of dollars. While a 100-Gbps LAN is very attractive, few would be willing to pay thousands of dollars for such LAN interfaces. Hence, in the near term, it is reasonable to expect WDM LANs to be used only in experimental settings or in networks requiring very high performance. However, as the cost of transceivers declines, it is not unlikely that this technology will become commercially viable.

4. WDM WIDE-AREA NETWORKS

In the WDM broadcast LAN architecture described above, nodes communicate by tuning their transmitter and receiver to common wavelengths. In a wide-area network (WAN), where a broadcast architecture is not scalable, traffic must be switched and routed at various communication nodes throughout the network. In electronic networks, this switching is accomplished using either circuit switching or packet switching techniques.

With packet switching, each network node must process each packet's header to determine the destination of the packet and make suitable routing decisions; with circuit switching, circuits are set up in advance of the communication and routing and switching decisions are predetermined for the duration of the call. Hence, with circuit switching there is no need for nodes to process the incoming data.

Optical packet switching involves a number of rather complex functions, such as header recognition, packet synchronization, and optical buffering. These technologies are rather crude and largely experimental at present [10]. Hence most efforts at optical networking have focused on circuit-switched networks. With WDM, much of the effort has been on the design of wavelength-routed networks, where connections between end nodes in the network utilize a full wavelength. There are a number of challenges in the design of an optical WDM network including the choice of a network architecture, performing the functions of routing and switching wavelengths, as well as assigning wavelengths to the various connections.

Since optical network elements are relatively expensive and of limited capabilities, the choice of a network architecture is particularly critical. Early efforts at designing all-optical WDM networks have been focused on a hierarchical architecture where different network elements are employed at different levels of the hierarchy. For example, as shown in Fig. 10, an optical star may be used in the local areas of the network and frequency-selective switches, optical amplifiers, wavelength converters, and other components may be used in the backbone of the network.

An early prototype of an all-optical-network is the all-optical network (AON) testbed developed by scientists at MIT, AT&T, and Digital Equipment Corporation (DEC) [2]. The AON testbed used the hierarchical architecture shown in Fig. 11. The lowest level in the

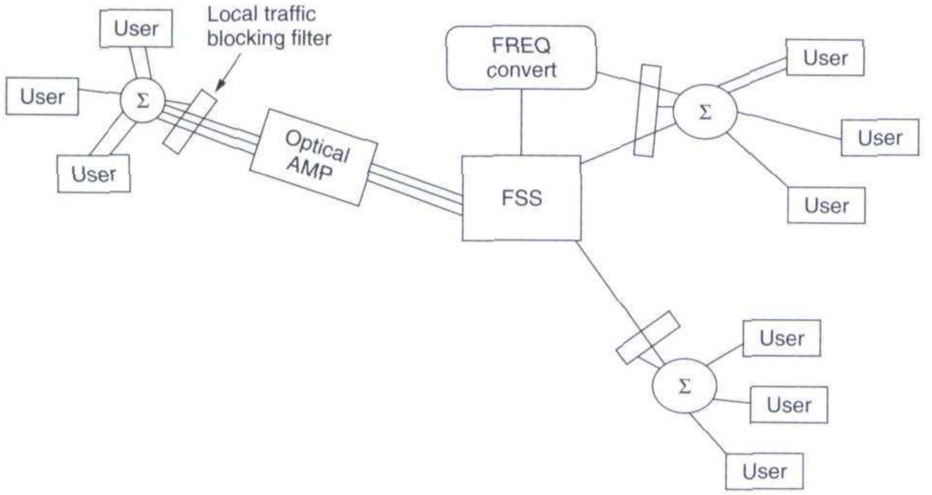

Figure 10. A partitioned all-optical WDM network.

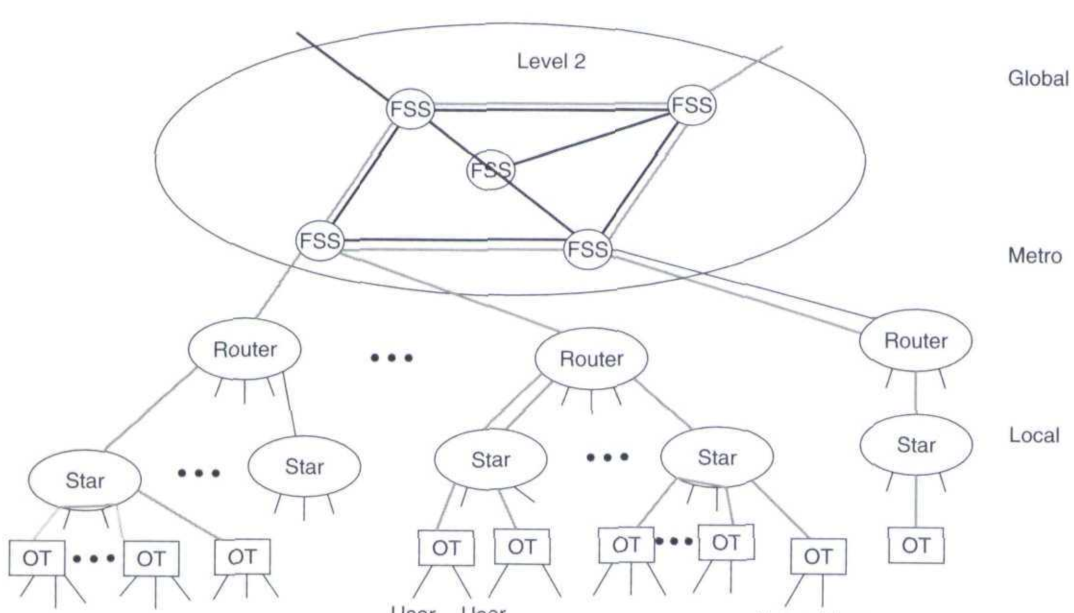

Figure 11. The AON architecture.

hierarchy was the LAN employing a broadcast star. A number of wavelengths were allocated for use within the local area, and those were separated from the other levels of the hierarchy using a wavelength blocking filter. Separating the local wavelengths from the rest of the network allows for those wavelengths to be reused at different local areas. The next level of the hierarchy used a wavelength router for connecting (in a static manner) different local areas. Finally, a FSS was used for providing connectivity in the wide area. A prototype of the architecture consisting of the lowest two levels was deployed in the Boston area connecting between facilities at MIT, MIT Lincoln Laboratory, and DEC.

The AON testbed supported two primary types of services: (1) a circuit-switched wavelength service that can establish wavelength connectivity between different nodes and (2) a circuit-switched time-slotted service by which a fraction of a wavelength can be assigned to a connection. The time-slotted service allows the flexibility of provisioning at the subwavelength level. However, implementing such a service requires very precise synchronization between the nodes so that the different circuits can be aligned on time-slot boundaries. The AON testbed demonstrated a 20-wavelength network, separated by 50 GHz and transmitting at rates of up to 10 Gbps per wavelength. AON also employed tunable transceivers. The transmitter was implemented using a DBR (distributed Bragg reflector) laser that can tune between wavelengths in 10 ns.

The early wavelength routing networks raised a number of architectural questions for all-optical networks. Perhaps the one that received the most attention is that of dealing with wavelength conflicts. Without using a wavelength converter, the same wavelength must be available for use on all links between the source and the destination of the call. A wavelength conflict may occur when each link on the route may have some free wavelengths, but the same wavelength is not available on all of the links. This situation can be dealt with through the use of wavelength converters that can switch between the wavelengths. However, because of the high cost of wavelength conversion, a number of studies quantifying the benefits of wavelength conversion in a network have been published [11,12]. Others have considered the possibility of placing the wavelength converters only at some key nodes in the networks [13]. However, the detailed results of these studies are beyond the scope of this article.

Another promising approach for dealing with wavelength conflicts is the use of a good wavelength assignment algorithm that attempts to reduce the likelihood of a wavelength conflict occurring. The wavelength assignment algorithm is responsible for selecting a suitable wavelength among the many possible choices for establishing the call. For example, the three calls illustrated in Fig. 12 can be established using three wavelengths ($\lambda_1\lambda_2\lambda_3$) as shown on the left or just two wavelengths as shown on the right. By choosing the assignment on the right, λ_3 remains free for use by future potential calls. A number of wavelength assignment schemes have been proposed [14,15], and the subject remains an active area of research.

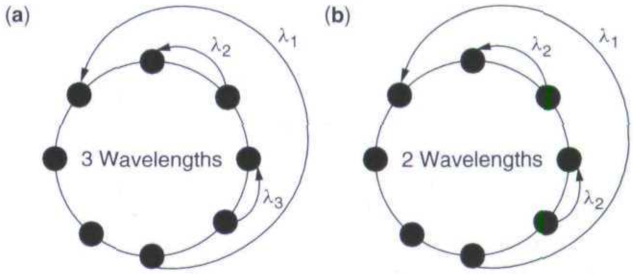

Figure 12. Two possible wavelength assignments for three calls on a ring: (**a**) bad and (**b**) good assignments.

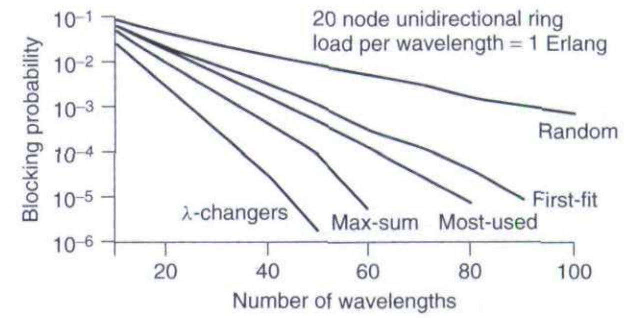

Figure 13. Performance of wavelength assignment algorithms in a ring network.

Figure 13 compares the performance of some proposed wavelength assignment algorithms. The simplest algorithm is to randomly select a wavelength from among the available wavelengths along the path. Clearly such an algorithm would be very inefficient, and, as can be seen from the figure, the random algorithm results in the highest blocking probability. A first-fit heuristic assigns the first available (i.e., lowest index number) wavelength that can accommodate the call. The most frequently used heuristic assigns the wavelength that is used on the most number of fibers in the network and lastly, the max-sum algorithm assigns the wavelength that maximizes the number of paths that can be supported in the network after the wavelength has been assigned [16]. All of these algorithms attempt to pack the wavelengths as much as possible, leaving free wavelengths open for future calls. Also shown in Fig. 12 is the blocking probability that results when wavelength changers are used. This represents an upper bound on the performance of any wavelength assignment algorithm. The significance of this illustration is that a good wavelength assignment algorithm can result in a blocking probability that is nearly as low as if wavelength changers are employed. Hence, a significant reduction in network costs can be obtained by using a good wavelength assignment algorithm.

Wavelength assignment was the first fundamental architectural problem in the design of all-optical networks. It has received much attention in the literature, and remains an active area of research. Beyond wavelength assignment, other important areas of research include the use of wavelength conversion (e.g., where wavelength converters should be deployed), the use of optical switching, and mechanisms for providing protection from failures. While a meaningful discussion of these topics is beyond the scope of this article, Ref. 1 provides a recent overview on most of these topics.

5. JOINT OPTICAL AND ELECTRONIC NETWORKS

Since all-optical networks are not likely to become a reality in the near future, the current trend in networking is to design networks that use a combination of optical and electronic techniques. A simple example would be a SONET-over-WDM network where the nodes in a SONET ring are connected via wavelengths rather than point-to-point fiber links. As we explained in the introduction, this use of WDM transmission is beneficial because it both reduces network cost and increases network capacity due to the large number of wavelengths. In fact, using WDM at the optical layer also introduces an additional flexibility in the design of the network.

Consider, for example, the networks in Fig. 14, where the optical topology consists of optical nodes [e.g., optical switches or ADMs (add/drop multiplexes)] that are connected via fiber and the electronic topology consists of electronic nodes (e.g., SONET multiplexers) that are connected using electronic links. Without WDM, the electronic topology shown in the figure cannot possibly be realized on the optical topology because the optical topology does not have a fiber link between nodes 1 and 3. However, with WDM an electronic link can be established between nodes 1 and 3 using a wavelength that is routed through node 2. The optical switch (or ADM) at node 2 can be configured to pass that wavelength through to node 3, creating a virtual link between nodes 1 and 3. This approach allows for various electronic topologies to be realized on optical topologies that do not necessarily have the same structure. Electronic nodes can be connected via wavelengths that are routed on the optical topology.

This approach can be used to realize a variety of electronic networks, such as ATM, IP, or SONET [17]. The connectivity of the electronic nodes determines the required wavelength connection that must be established. In other words, each link in the electronic topology requires a wavelength connection (also referred to as *lightpath*) between the optical nodes. In order to realize a particular electronic topology, the corresponding set of wavelength connections must be realized on the optical topology. This very practical problem leads to another version of the routing and wavelength assignment (RWA) problem known as the *batch RWA*. Given a set of lightpaths that must be established, a RWA must be found such that each lightpath must use the same wavelengths along its route from the source to the destination (assuming no wavelength conversion) and no two lightpaths can use the same wavelength on a given link. This problem is closely related to the well-known NP-complete graph coloring problem, and in fact Chlamtac et al. [17] showed that the static RWA problem is indeed NP-complete by suitable transformation from graph coloring.

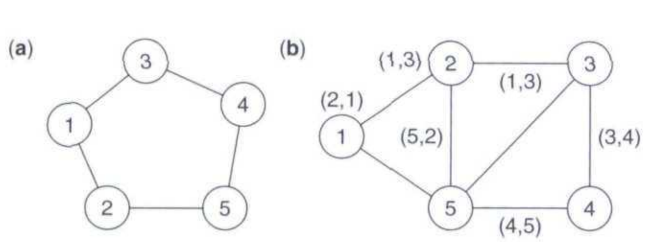

Figure 14. The electronic (**a**) and optical (**b**) topologies of a network.

WDM allows the electronic (logical) topology to be different from the physical topology over which it is implemented. This ability created the interesting opportunity for "logical topology design"; that is, given the traffic demand between the different nodes in the network, what is the best logical topology for supporting that demand. For example, suppose that one is to implement a ring logical topology as in Fig. 14. If a large amount of traffic is being carried between nodes 1 and 4, and virtually no traffic between 1 and 3, it makes more sense to connect the ring in the order 1–4–3 rather than 1–3–4. In this way, the length of the path that the traffic must traverse is reduced, and consequently, the load on the links is also reduced. Designing logical topologies for WDM networks has typically been formulated as an Integer Programming problem, solutions to which are obtained using a variety of search heuristics [18,19]. Furthermore, with configurable WDM nodes (e.g., wavelength switches), it is even possible to reconfigure the logical topology in response to changes in traffic conditions [20].

6. FUTURE DIRECTIONS

Since their inception, the premise of optical networks has been to eliminate the electronic bottleneck. Yet, so far, all optical networks have failed to emerge as a viable alternative to electronic networks. This can be attributed to the tremendous increase in the processing speeds and capacity of electronic switches and routers. Nonetheless, the enormous capacity and configurability of WDM can be used to reduce the cost and complexity of the electronic network, paving the way to even faster networks.

While in the near term optical networking will be used only as a physical layer beneath the electronic network, researchers are aggressively pursuing all-optical packet-switched networks. Such networks may first emerge in the local area, where a broadcast architecture can be used without the need for optical packet switching. As all-optical packet switching technology matures, all-optical networks may become viable. However, as of this writing, it is very unclear whether truly all-optical networks will ever become a reality.

BIOGRAPHY

Eytan Modiano received his B.S. degree in electrical engineering and computer science from the University of Connecticut at Storrs in 1986 and his M.S. and Ph.D. degrees, both in electrical engineering, from the University of Maryland, College Park, Maryland, in 1989 and 1992, respectively. He was a Naval Research Laboratory fellow between 1987 and 1992 and a National Research Council postdoctoral fellow during 1992–1993, while he was conducting research on security and performance issues in distributed network protocols.

Between 1993 and 1999 he was with the Communications Division at MIT Lincoln Laboratory where he designed communication protocols for satellite, wireless, and optical networks and was the project leader for MIT Lincoln Laboratory's Next Generation Internet (NGI) project. Since 1999, he has been a member of the faculty of the Aeronautics and Astronautics Department and the Laboratory for Information and Decision Systems (LIDS) at MIT, where he conducts research on communication networks and protocols with emphasis on satellite and hybrid networks, and high speed optical networks.

BIBLIOGRAPHY

1. R. Ramaswami and K. Sivarajan, *Optical Networks*, Morgan Kaufmann, San Francisco, CA, 1998.
2. I. P. Kaminow, A wideband all-optical WDM network, *IEEE JSAC/JLT* **14**(5): 780–799 (June 1996).
3. E. Modiano, WDM-based packet networks, *IEEE Commun. Mag.* **37**(3): 130–135 (March 2000).
4. B. Mukherjee, WDM-based local lightwave networks part I: Single-hop systems, *IEEE Network* **6**(3): 12–27 (May 1992).
5. G. N. M. Sudhakar, M. Kavehrad, and N. D. Georganas, Access Protocols for passive optical star networks, *Comput. Networks ISDN Syst.* 913–930 (1994).
6. D. Bertsekas and R. Gallager, *Data Networks*, Prentice-Hall, Englewood Cliffs, NJ, 1992.
7. N. Mehravari, Performance and protocol improvements for very high-speed optical local area networks using a star topology, *IEEE/OSA J. Lightwave Technol.* **8**(4): 520–530 (April 1990).
8. E. Modiano and R. Barry, A novel medium access control protocol for WDM-based LAN's and access networks using a master/slave scheduler, *IEEE J. Lightwave Technol.* **18**(4): 2–12 (April 2000).
9. E. Modiano, Random algorithms for scheduling multicast traffic in WDM broadcast-and-select networks, *IEEE/ACM Trans. Network.* **7**(3): 425–434 (June 1999).
10. V. Chan, K. Hall, E. Modiano, and K. Rauchenbach, Architectures and technologies for optical data networks, *J. Lightwave Technol.* **16**(12): 2146–2186 (December 1998).
11. R. Barry and P. Humblet, Models of blocking probability in all-optical networks with and without wavelength changers, *IEEE J. Select. Areas Commun.* **14**(5): 858–867 (June 1996).
12. S. Subramanian, M. Azizoglu, and A. Somani, Connectivity and sparse wavelength conversion in wavelength-routing networks, *IEEE INFOCOM* 148–155 (1996).
13. R. Ramaswami and G. Sasaki, Multiwavelength optical networks with limited wavelength conversion, *IEEE INFOCOM* 490–499 (1997).
14. R. Ramaswami and K. N. Sivarajan, Routing and wavelength assignment in all-optical networks, *IEEE/ACM Trans. Network.* 489–500 (Oct. 1995).
15. S. Subramanian and R. Barry, Wavelength assignment in fixed routing wdm networks, ICC '97, Montreal, June 1997.
16. I. Chlamtac, A. Ganz, and G. Karmi, Lightpath communications: An approach to high-bandwidth optical WANs, *IEEE Trans. Commun.* **40**(7): 1171–1182 (July 1992).
17. J. Bannister, L. Fratta, and M. Gerla, Topological design of WDM networks, *Proc. Infocom '90*, Los Alamos, CA, 1990.
18. R. Ramaswami and K. Sivarajan, Design of logical topologies for wavelength routed optical networks, *IEEE J. Select. Areas Commun.* **14**(5): 840–851 (June 1996).
19. J. P. Labourdette and A. Acampora, Logically rearrangeble multihop lightwave networks, *IEEE Trans. Commun.* **39**(8): 1223–1230 (Aug. 1991).
20. A. Narula-Tam and E. Modiano, Dynamic load balancing in WDM packet networks with and without wavelength constraints, *IEEE J. Select. Areas Commun.* **18**(10): 1972–1979 (Oct. 2000).

WAVELETS: A MULTISCALE ANALYSIS TOOL

HAMID KRIM
ECE Department, North Carolina State University
Centennial Campus
Raleigh, North Carolina

One's first and most natural reflex when presented with an unfamiliar object is to carefully look it over, and hold it up to the light to inspect its different facets in the hope of recognizing it, or of at least relating any of its aspects to a more familiar and well-known entity. This almost innate strategy pervades all science and engineering disciplines.

Physical phenomena (e.g., earth vibrations) are monitored and observed by way of measurement in the form of temporal and/or spatial data sequences. Analyzing such data is tantamount to extracting information useful to further understand the underlying process (e.g., frequency and amplitude of vibrations may be an indicator for an imminent earthquake). Visual or manual analysis of typically massive amounts of acquired data (e.g., in remote sensing) are impractical, causing one to resort to adapted mathematical tools and analysis techniques to better cope with potential intricacies and complex structure of the data. Among these tools figure a variety of functional transforms (e.g., Fourier transform) that in many cases may facilitate and simplify an analytical track of a problem, and frequently (and just as importantly) provide an alternative view of, and a better insight into, the problem. (This, in some sense, is analogous to exploring and inspecting data under a "different light.") An illustration of such a "simplification" is shown in Fig. 1, where a rather intricate signal $x(t)$ shown in the leftmost figure may be displayed or viewed in a different space as two elementary tones. In Fig. 2, a real bird chirp is similarly displayed as a fairly rich signal which, when considered in an appropriate space, is reduced and "summarized" to a few "atoms" in the time–frequency (TF) representation. Transformed signals may formally be viewed as *convenient* representations in a different domain that is itself described by a set of vectors/functions $\{\phi_i(t)\}_{i=\{1,2,\ldots,N\}}$. A contribution of a signal $x(t)$ along a direction "$\phi_i(t)$" (its projection) is given by the following inner product

$$C_i(x) = < x(t), \phi_i(t) > = \int_{-\infty}^{\infty} x(t)\phi_i(t)\,dt \qquad (1)$$