WDM-Based Packet Networks

Eytan Modiano, MIT Lincoln Laboratory

ABSTRACT We discuss design considerations for Wavelength Division Multiplexing (WDM)-based packet networks. In the near term, such networks are likely to consist of WDM links connected using some form of electronic multiplexing. The focus of this article is on the joint design of the electronic and optical layer with the objective of simplifying the network and reducing the protocol stack. To that end, we discuss the benefits of optical flow switching, network reconfiguration, traffic grooming, and optical layer protection. We also discuss the state of all-optical packet networking with particular focus on Local Area Network technology.

The spectacular growth of Internet traffic has led to a dramatic increase in demand for data transmission capacity. In fact, while not too long ago voice traffic was predominant on the backbone networks, requirements for data traffic will soon far surpass those for voice. In anticipation of this added demand, we have seen a tremendous increase in the transmission capacity available in backbone networks. The advancements in the backbone networks can be largely attributed to developments in high-speed transmission and switching systems. Of particular importance has been the emergence of Wavelength Division Multiplexing (WDM) technology, which supports multiple simultaneous channels on a single fiber. WDM transmission systems with transmission capacity exceeding a Tera-bit per second have been demonstrated; and systems supporting hundreds of Giga-bits per second are becoming commercially available. Unfortunately, however, since much of today's infrastructure was developed for supporting voice traffic, it is not ideally suited for meeting the demands of data traffic. As future networks are developed to support the Internet, they must be designed and optimized for that purpose.

Most of today's data networks are supported using the Internet Protocol (IP) and it is widely believed that IP will remain the dominant network layer technology. However, there are a number of issues that keep IP networks from taking advantage of the huge capacity of the underlying transmission systems. For example, in today's internet, IP routers are connected using additional layers of electronic multiplexing (e.g., ATM, SONET) that were designed for supporting voice communication and are a part of the existing network infrastructure. These layers both add complexity and produce unnecessary bandwidth inefficiencies. Furthermore, today's IP routers are difficult to scale to high data rates due to their significant electronic processing requirements. While current developments in IP, such as hardware routing and IP switching, promise improvements for the future, they will still be limited by electronic processing from taking full advantage of the capacity of WDM transmission systems.

This realization has led to a recent wave of research aimed at eliminating the electronic layers that come between IP and WDM. A number of IP router vendors and research institutions have announced future demonstration of "IP-over-WDM." However, it is not clear what putting IP directly over WDM really means. Although the multitude of electronic layer services present in today's networks represent a significant "overkill," they still serve certain functions that are presently not a part of the IP or the WDM layers. For example, IP transmits packets asynchronously, hence at the very least some framing mechanism is needed in order to transport these packets over WDM channels. Also, IP offers a connectionless service and is not designed to provide any service guarantees. While future versions of IP are likely to provide some form of service guarantees, today customers in need of such service guarantees typically use an ATM or Frame Relay connection with guaranteed service, on top of which they set up their IP network. Clearly, in designing an efficient and streamlined IP-over-WDM network it will be necessary to somehow preserve certain functions that are presently provided by the various electronic layers.

The focus of this article is on the design of an IP network that can take full advantage of the optical (WDM) layer. The discussion of WDM-based packet networks can be divided into that of Local Area Networks (LANs) and "Wide" Area Networks (WANs). The primary distinction between the two is that the former consists of a shared broadcast medium (e.g., broadcast star), while the latter consists of multiple fiber links that are connected to each other using switches. One vision for such a network is an all optical network where all of the switching is done in the optical domain. Unfortunately, however, optical packet switching technology is not yet mature. Hence, a near-term vision is to use electronic switches to connect between the different WDM channels. In that context, the goal is to jointly design the electronic and optical layers in order to optimize the network's performance. In the next section of this article we discuss issues in the joint design of the electrical and optical layers in a WDM-based IP network; in the following section we discuss all optical WDM-based packet networks, with particular focus on WDM-based Local Area Networks.

ELECTRONIC PACKET SWITCHING

One goal of optical packet networks is to eliminate the need for a cumbersome electronic layer. While electronic processing speeds are dramatically increasing, so is the transmission capacity of optical systems. Inevitably, network capacity will always be limited by the electronic bottleneck. Hence, a good network design objective is to bypass or eliminate electronics whenever possible. This can best be achieved by jointly designing the electronic and optical layers so that network functions are performed at the layer best suited to them, and duplication of network functions at multiple layers is eliminated. In this section we discuss a number of issues in the joint design of the electronic and optical layers of a WDM-based packet network.

REDUCING THE PROTOCOL STACK

In the present Internet, IP services are provided using a wide range of electronic multiplexing and switching equipment. A typical network may include as many as three or four different electronic multiplexing and switching layers. For example, as shown in Fig. 1a, internet packets may be carried using Frame Relay where the IP packets are encapsulated in frames that are sometimes mapped into ATM cells, which, in turn, are carried over Synchronous Optical Networking (SONET)

 0163-6804/99/$10.00

transport frames. The multitude of layers produces bandwidth inefficiencies, adds to the latencies of connections, and inhibits providing quality of service assurances. Worse, the layers are largely unaware of each other, causing duplication of network services. Alternatively, one could use the simplified protocol stack shown in Fig. 1b, where the IP traffic is carried directly by a simplified electronic layer. This "new" electronic layer must incorporate the necessary functions that are presently offered by the various electronic layers (e.g., framing and restoration). In addition, this electronic layer can be designed to allow IP to take advantage of new optical layer services. For the remainder of this section we will describe optical layer services and algorithms that can significantly improve the performance of the network and simplify the electronic layer.

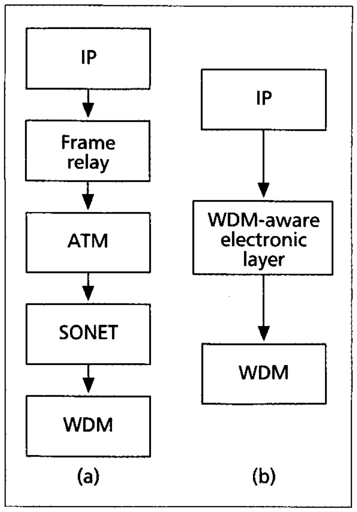

■ **Figure 1.** *a) Typical protocol stack; b) simplified protocol stack.*

Optical Flow Switching

One of the main bottlenecks in the present Internet is routing at the IP layer. Several methods have been proposed to alleviate this bottleneck by switching long duration flows at lower layers. Tag switching uses routing protocols to predefine routes within the network and assign tags to the routes; packets are then switched based on these tags avoiding the need for routing table lookups [1]. IP switching dynamically sets up layer-2 (e.g., ATM) virtual circuits for those connections that are perceived to be long [2].

This concept of lower-layer switching can be extended to switching large volume and/or long duration flows at the optical layer. That is, a light-path can be established for large data transactions such as the transfer of large files or long duration and high bandwidth streams across the network. To achieve ultimate efficiencies, an optical flow switching protocol may need to aggregate flows with similar characteristics in order to switch them together. The simplest form of this technique is to use an entire wavelength at any given time for a single transaction. This concept can give rise to a hybrid multi-layer switching approach in which long duration sessions are switched at the electronic layer, and even longer duration and higher bandwidth flows are switched optically. While it appears that such a multi-layer switching approach can reduce computation loads and processing delays in networks, many issues in the design of such a protocol remain to be resolved. Unlike IP switching where an ATM virtual circuit is set up for a perceived flow, an optical flow switching protocol must be much more judicious in determining when to set up such flows. This is because light-paths in a network are a scarce resource that cannot be arbitrarily assigned. In addition, without the use of wavelength converters, it is necessary to assign the same wavelengths to a lightpath along its entire path, further restricting the number of available wavelengths. Consequently, for optical flow switching, it may be necessary for the application layer to inform the IP layer of the arrival and characteristics of large flows for switching. Also, the size of such a flow must be relatively large when compared to the flow setup time, which will likely exceed a round trip delay through the network. In a sense, optical flow switching provides a step in the direction toward all optical packet switching, the main difference being that optical flow switching is used only for very large transactions, and hence can tolerate the setup time of a flow, while optical packet switching operates on packets of all sizes and hence must be able to switch the packets more rapidly.

A similar concept is that of Optical Burst Switching (OBS) [3]. With OBS, a control packet is sent on an out-of-band channel to announce an upcoming burst. The control packet is then followed, after a short delay for processing the control message at every node, by a burst of data without waiting for connection establishment. Since a connection is not set up before the burst is transmitted, it is possible that the control packet may fail to reserve resources at some node along its path, in which case, the burst would have to be dropped. In order to reduce the likelihood of dropping a burst, optical delay lines can be used to temporarily store the burst until the resources become available.

In some ways, however, this approach begins to resemble an optical packet-switched network. In an optical packet-switched network the goal is to provide the same functionality as that of an electronic packet network; hence, switches inside the network optically route packets based on packet header information. Presumably, the use of optical packet switching technology would allow the network to operate at much higher speeds. In an optical packet-switched network, all of the functions of the network must be implemented in the optical domain. Some of the challenges include the generation of ultra-fast pulses, synchronization, high-speed switching, optical buffering and packet header processing. Since implementing network functions in the optical domain is complicated, most research efforts toward an all optical packet network focus on implementing a network with very limited functionality. For example, header processing is limited to address recognition; switch buffers, which are implemented using optical delay lines, are typically limited to storing just a few packets; and packet routing is predetermined. All optical-packet networking is still only at the experimental stage and is far from being commercially viable. The reader is referred to [4] for a comprehensive survey of optical packet switching technologies.

Virtual Topology Reconfiguration

In WDM networks, the physical topology is the one seen by the optical layer. It consists of passive or configurable optical nodes interconnected by fiber. The virtual topology, seen by the electronic layer, consists of a set of nodes interconnected by light-paths (wavelengths). In this way, WDM networks provide a way to interconnect electronic switches with high bandwidth bit pipes without dedicating a fiber pair between each pair of switches. However, the configurable nature of WDM also allows the logical topology to be dynamically reconfigured in response to changes in traffic conditions.

WDM networks can reconfigure light-paths, providing the ability to dynamically optimize the network for changing patterns of externally offered traffic. This is achieved by changing the light-path connectivity between electronic switches and routers, thereby reconfiguring the electronic virtual topology. Light-paths can be changed via tuning of the transmitter wavelengths in combination with frequency-selective-switches that can alter the route of a wavelength inside the network. For example, consider four nodes physically connected in a ring. Assume that each node has one port (i.e., single receiver and transmitter), and that the fiber supports two wavelengths,

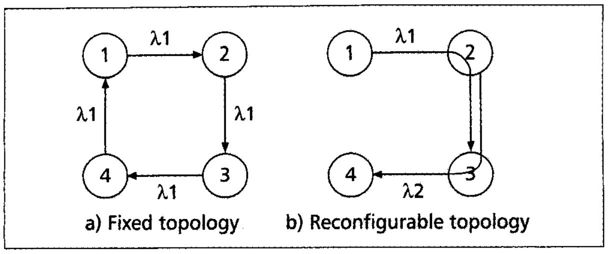

■ **Figure 2.** *Using WDM to reconfigure the electronic topology.*

λ1 and λ2. A connected, fixed logical topology must take the form of a unidirectional ring, as pictured in Fig. 2a. If a full-wavelength call is in progress from node 1 to node 3, and a call request arrives between nodes 2 and 4, then that request must be blocked. In a reconfigurable system, both calls can be supported as shown in Fig. 2b.

Preliminary studies on reconfiguration of a WDM ring show significant promise [5]. The work in [5] assumes that calls take up a full wavelength and cannot be rearranged or rerouted. The assumption about rerouting is made in order to eliminate the possibility of calls being adversely affected by reconfiguration. Shown in Fig. 3 is the gain that can be achieved through reconfiguration. This gain is defined to be the ratio of the load that can be supported by a reconfigurable system to that of a bi-directional, fixed topology system at a given blocking probability (a blocking probability of 0.01 is used in Fig. 3). As can be seen from the figure, the capacity gain due to reconfiguration is most significant when the ratio of wavelengths to ports per node (W/P) is large and the number of ports per node is small. The results from [5] indicate that when the number of wavelengths is much larger than the number of available ports per node, a capacity gain on the order of N can be obtained, where N is the number of nodes in the network. When the number of wavelengths approaches the number of ports per node, the benefits of reconfiguration are significantly diminished with approximately a factor of two gain in capacity.

Reconfiguration effectively increases the overall system capacity by allocating bandwidth only where it is needed. However, applying this concept to IP networks poses additional complications that do not exist in a circuit switched system and require further research. For example, how does one determine when to reconfigure the network and how does one prevent routing table instabilities from occurring due to rapid topology changes. Most existing work on topology design for WDM-based packet networks address the optimal virtual topology design problem [6]. In [7], algorithms are developed in order to migrate an existing network topology to another with the minimum number of "branch exchanges" (i.e., exchanging the source and destinations of two links in the network).

Network Restoration

Various failures, such as fiber cuts, line-card and switch failures, and software failures, can occur that disrupt network services. Protection and restoration are two methods networks use to recover from these failures. Protection refers to hardware-based, pre-planned, fast failure recovery. Restoration refers to software-based, dynamic, slower recovery. Protection is generally limited to simple topologies like rings or the interconnection of rings, while restoration can be applied to general mesh networks and is typically more bandwidth efficient.

Failure recovery must be done at the electronic layers in order to recover from line card or electronic switch failures. Electronic recovery mechanisms, e.g., as is done in SONET, can also be used to protect against failures at the optical layer such as a fiber cut or a malfunctioning optical switch. However, in many cases providing recovery at the optical layer can be more robust and efficient. For example, optical layer recovery can protect electronic services that do not have built-in recovery mechanisms or whose recovery mechanisms are slow (e.g., IP). Also, in many cases, optical layer recovery is simpler, more natural and provides enhanced reliability. For instance, consider the case of 32 SONET rings being supported over a WDM ring network with 32 wavelengths. Without optical protection, each of the 32 SONET rings would need to individually recover from a single fiber cut, e.g., by loop-back in a SONET bi-directional ring network. On the other hand, the fiber cut can be optically restored, simultaneously restoring service to many electronic connections. Protection at the optical layer has the added advantage that the failure is transparent to SONET, allowing each SONET ring to individually respond to additional failures such as a line card failure.

Providing protection at the optical layer is of paramount importance if the goal of reducing the protocol stack is to be achieved. In today's networks, SONET is used primarily for its fast recovery capabilities. If the SONET layer is to be eliminated, fast recovery must be implemented at the optical layer. Additional, slower restoration mechanisms may be provided at the electronic (IP) layer to recover from line card failures, etc. Recently, much attention has been paid to optical layer protection and to the interaction between the optical and electronic recovery mechanisms [8]. However, there are problems providing protection at both the optical and electronic layers if the layers are uncoordinated. For instance, protection can be duplicated at both the optical and electronic layers, leading to a 75 percent loss in efficiency (assuming 50 percent efficiency for each layer of protection). In addition, differing time scales may lead to race conditions and topology oscillations.

Traffic Grooming

In addition to its fast recovery mechanisms, SONET also plays an important role as a multiplexer of lower rate streams. Consider the situation where the optical data network is providing a multitude of point-to-point stream connections among routers, switches, and even end users. Unless these "users" require full wavelength connections, sub-wavelength capacity connections need to be allocated. This can be accomplished through the use of SONET multiplexers that can aggregate low rate calls on to a higher rate channel. However, if calls are indiscriminately multiplexed on to wavelengths, then each wavelength entering or leaving a node will need to be convert-

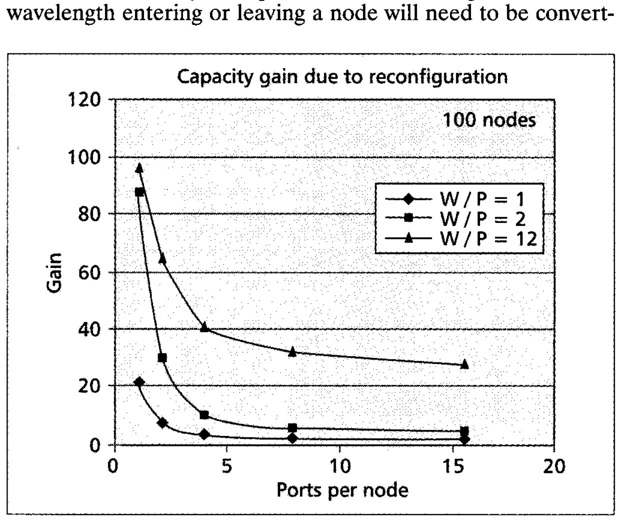

■ **Figure 3.** *Performance of reconfiguration in a WDM ring network.*

ed to electronics to make drop/forwarding decisions. Alternatively, if calls are *groomed* with foresight onto wavelengths, then the number of wavelengths that need to be processed at each node can be significantly reduced.

For example, when a SONET ring network is used to provide point-to-point OC-3 circuits between pairs of nodes, SONET Add/Drop Multiplexers (ADMs) are used to combine up to 16 OC-3 circuits into a single OC-48 that is carried on a wavelength. If a wavelength carries traffic that originates or terminates at a particular node, then that wavelength must be dropped at that node and terminated in a SONET ADM. In order to reduce the number of ADMs used, it is better to groom traffic such that all of the traffic to and from a node is carried on the minimum number of wavelengths (and not dispersed among the different OC-48's). Traffic grooming algorithms can be designed to minimize electronic costs while simultaneously making efficient use of wavelengths. Traffic grooming on SONET ring networks with uniform all-to-all traffic show a significant reduction in the number of required ADMs [9, 10]. Figure 4 shows the significant savings in the number of ADMs needed in a SONET ring network using an algorithm developed in [9]. The results from [9] assume the use of a unidirectional ring. Similar results, however, were obtained in [10] for a bi-directional ring.

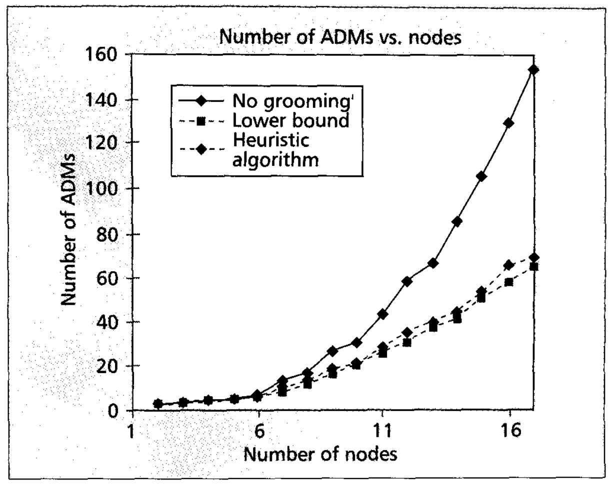

■ **Figure 4.** *Performance of grooming in a WDM/SONET ring network.*

WDM-Based Local Area Networks

In the previous section we focused on the joint design of the electronic and optical layer of a WDM-based packet network, the intrinsic assumption being that an electronic switching and multiplexing layer is necessary and cannot be altogether eliminated. That assumption is certainly true in today's Wide Area Networks (WANs) because in the wide area, packet switching and multiplexing are needed to provide both flexible connectivity and efficient use of the scarce bandwidth. Unfortunately, optical packet switching and buffering technology is in no way a mature technology [4]. Hence, at least in the near term, WANs will have to rely on electronic switching to provide connectivity.

The future of all optical packet networks is much brighter in Local Area Networks (LANs), because LANs span short distances, ranging from a few meters to a few thousand meters, and are an inherently shared medium. Therefore, in the LAN there is no need for any switching or buffering inside the network. This makes it possible to design optical LANs that entirely eliminate the electronic layer and transmit the packets directly over the optical light-paths. 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 Gb/s each).

Typically WDM-based LANs assume the use of a broadcast star architecture (Fig. 6). An optical star coupler is used to connect all of the nodes. Each node is attached to the star using a pair of fiber, one for transmission and the other for reception. The star coupler is a passive device that simply connects all of 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 Gb/s. 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 the 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 micro-seconds 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. Discussions of WDM-based LANs usually classify them based on the number and tunability of the transmitters and receivers. 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 fixed tuned receiver (TT-FR) or a fixed transmitter and 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, thus eliminating 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 and 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. Second, such a network is complicated to administer because whenever a new node is added, 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

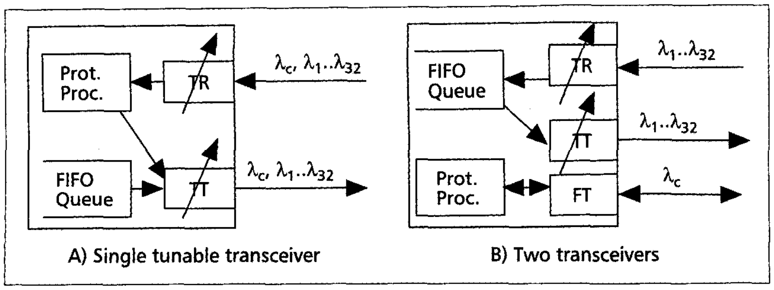

■ **Figure 5.** *User terminals.*

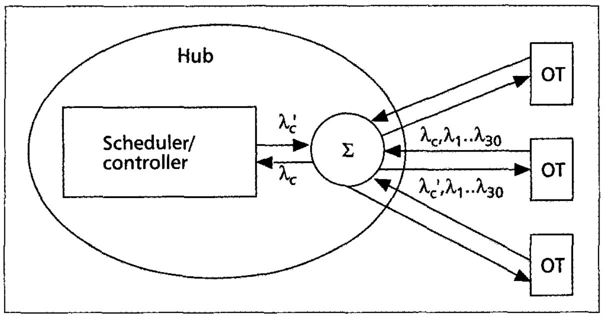

■ **Figure 6.** *A WDM-based LAN using a broadcast star and a scheduler.*

necessary for three reasons. First, if two nodes transmit on the same wavelength simultaneously, their transmissions will interfere with each other (collide), so some mechanism must be employed to prevent such collisions. Second, if two or more nodes transmit to the same node at the same time (albeit on different wavelengths), if that node only has a single receiver, it will only be able to receive one of the transmissions. Finally, in order 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 pre-transmission 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. 5.

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. These protocols are more complicated than single channel MAC protocols because they must arbitrate among a number of shared resources: data channels, the control channel, and the receivers.

Early MAC protocols for WDM broadcast networks attempted to use Aloha for sharing the channels. 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. In order 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 percent. A slotted version of Aloha, where nodes are synchronized and transmit on slot boundaries, can achieve a throughput of 36 percent.

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 [11]. In order for a transmission to be successful, the following sequence of events must take place. First, the transmission on the control channel must be successful (i.e., no control channel collision); second, the receiving node must not be receiving any other transmission at the same time (i.e., no receiver collision); and last, 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 channel achieve a maximum utilization of less than 10 percent.

In view of the above, a number of MAC protocols that attempt to increase utilization by coordinating and scheduling the transmissions more carefully have been proposed [11]. For example, the protocol in [12] uses a simple master/slave scheduler as shown in Fig. 6. All nodes send their requests to the scheduler on a dedicated control wavelength, λ_C. The scheduler, located at the star, schedules the requests and informs the nodes on a separate wavelength, $\lambda_{C'}$, of their turn to transmit. Upon 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; multicast traffic is scheduled using a random algorithm [13]. Simulations show that the total system delay is relatively low even at high traffic loads [12].

We should point out, however, that despite their appeal, WDM-based LANs still face significant economic challenges because the cost of WDM transceivers (especially tunable transceivers) 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 Gb/s 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.

Conclusion

Since their inception, the premise of optical networks has been to eliminate the electronic bottleneck. Yet 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 and the fact that all-optical packet networking is still at the experimental stage. 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.

In order to achieve these cost savings, it will be useful to jointly design the electronic and optical layers of the networks so that they can take advantage of the capabilities of one another. One important objective in this joint design is to reduce the multitude of layers that are used in today's networks, while preserving their functionality. For example, in order to eliminate the need for SONET equipment, it will be necessary for the optical layer to provide fast recovery mechanisms; traffic grooming algorithms can be used to reduce the need for electronic multiplexers, and optical flow switching techniques may help reduce the size and complexity of IP routers.

While in the near term optical networking will only be used as a physical layer beneath the electronic network, researchers are aggressively pursuing all-optical packet-switched networks. Such networks are most likely to emerge in the local area, where a broadcast architecture can be used without the need for optical packet switching. In recent years, a number of protocols have been proposed for WDM-based LANs using a simple broadcast star architecture. Such WDM-based LANs can achieve throughputs exceeding 100 Gb/s without the need for any high-speed electronics in the network.

REFERENCES

[1] Y. Rekhter *et al.*, "Cisco Systems' Tag Switching Architecture Overview," IETF RFC 2105, Feb. 1997.

[2] P. Newman, T. Lyon, and G. Minshall, "Flow Labeled IP: A Connection Approach to ATM." *IEEE INFOCOM*, 1996.

[3] M. Yoo and C. Qiao, "Just-enough-time (JET): A high speed protocol for bursty traffic in optical networks," IEEE/LEOS Summer Topical meeting, Aug. 1997.

[4] K. Hall, *et al.*, "Architectures and Technologies for Optical Data Networks," *J. Lightwave Tech.*, vol. 16, no. 12, Dec. 1998.

[5] B. Schein and E. Modiano, "Increasing traffic capacity in WDM ring networks via topology reconfiguration," *Proc. CISS '99*, Mar. 1999.

[6] J. P. Labourdette and A. Acampora, "Logically rearrangeble multihop lightwave networks," *IEEE Trans. Commun.*, Aug. 1991.

[7] J. P. Labourdette, G. W. Hart and A. Acampora, "Branch-Exchange sequences for reconfiguration of lightwave networks," *IEEE Trans. Commun.*, Oct. 1994.

[8] G. Ellinas and T.E. Stern, "Automatic Protection Switching for Link Failures in Optical Networks with Bi-directional Links," *Proc. GLOBECOM '96*.

[9] A. L. Chiu and E. H. Modiano, "Reducing Electronic Multiplexing Costs in Unidirectional SONET/WDM Ring Networks via Efficient Traffic Grooming," *Proc. GLOBECOM '98*, Sydney, Australia, Nov. 1998.

[10] J. Simmons, E. Goldstein, and A. Saleh, "On The Value of Wavelength-Add/Drop in WDM Rings With Uniform Traffic," *Proc. OFC '98*, San Jose, CA, Feb. 1998.

[11] B. Mukherjee, "WDM-Based Local Lightwave Networks Part I," *IEEE Network*, May 1992.

[12] E. Modiano and R. Barry, "Design and analysis of an asynchronous WDM local area network using a master/slave scheduler," *Proc. INFOCOM '99*, New York, NY, Mar. 1999.

[13] E. Modiano, "Unscheduled Multicasts in WDM Broadcast-and-Select Networks," *Proc. INFOCOM '98*, San Francisco, CA, Mar. 1998.

BIOGRAPHY

EYTAN MODIANO (modiano@ll.mit.edu) received a 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, MD, in 1989 and 1992, respectively. He was a Naval Research Laboratory Fellow between 1987 and 1992 and a National Research Council Post Doctoral Fellow from 1992 to 1993 while he was conducting research on security and performance issues in distributed network protocols. In 1993 he joined the Communications Division of MIT Lincoln Laboratory, where he has been working on communication protocols for satellite, wireless, and optical networks. He has published more than 30 papers on various aspects of data networks. Presently he is the lead architect for MIT Lincoln Laboratory's Next Generation Internet (NGI) project. The scope of the project is to develop a WDM-based access network architecture for supporting broadband internet services. Since 1994 he has also been an adjunct professor in the College of Computer Science at Northeastern University, where he teaches graduate level courses on Data Networks.