Architectures and Technologies for High-Speed Optical Data Networks

Vincent W. S. Chan, *Fellow, IEEE,* Katherine L. Hall, *Senior Member, IEEE, Member OSA,* Eytan Modiano, *Member, IEEE,* and Kristin A. Rauschenbach, *Member, IEEE*

(Invited Paper)

***Abstract*— Current optical networks are migrating to wavelength division multiplexing (WDM)-based fiber transport between traditional electronic multiplexers/demultiplexers, routers, and switches. Passive optical add–drop WDM networks have emerged but an optical data network that makes full use of the technologies of dynamic optical routing and switching exists only in experimental test-beds. This paper will discuss architecture and technology issues for the design of high performance optical data networks with two classes of technologies, WDM and time division multiplexing (TDM). The WDM network architecture presented will stress WDM aware internet protocol (IP), taking full advantage of optical reconfiguration, optical protection and restoration, traffic grooming to minimize electronics costs, and optical flow-switching for large transactions. Special attention is paid to the access network where innovative approaches to architecture may have a significant cost benefit. In the more distant future, ultrahigh-speed optical TDM networks, operating at single stream data rates of 100 Gb/s, may offer unique advantages over WDM networks. These advantages may include the ability to provide integrated services to high-end users, multiple quality-of-service (QoS) levels, and truly flexible bandwidth-on-demand. We will give an overview of an ultrahigh-speed TDM network architecture and describe recent key technology developments such as high-speed sources, switches, buffers, and rate converters.**

***Index Terms*—Optical data processing, optical fiber communication, optical fiber LAN, optical signal processing, time division multiaccess, time division multiplexing, wavelength division multiplexing.**

I. Introduction

COMMUNICATIONS networks are undergoing rapid developments. In the past, the largest networks were electronic, circuit switched networks providing plain-old telephone service (POTS), and carrying mostly voice traffic. Starting with the ARPANET and the explosive growth of the Internet and the worldwide web, electronic packet switched networks have become ubiquitous. The rapid growth of these packet networks can be attributed to the efficiency of the internet protocol (IP) in servicing bursty traffic or computer users. At the same time, we have seen a revolution in the transport technology that is used to interconnect nodes in a network. Specifically, fiber optic transmission technology has advanced from lower rate multimode fiber links to singlemode amplified fiber links capable of carrying multiple 10 Gb/s channels per fiber. Still, the incredible bandwidth offered by optical transmission technology has only been used as a very high-speed replacement of copper cables. The routing and switching technologies used are currently still electronic based. With the advent of advanced optical devices such as integrated tunable lasers, optical grating routers, and cross-connect switches, a new class of optically routed and switched wavelength division multiplexing (WDM) network is feasible.

One reason that optical data networking is still in its infancy is that optical transport alone has provided an incredible increase in available bandwidth that exceeds any present day demand. It is hard to find applications other than aggregated traffic trunking, that can take advantage of a 10-Gb/s data channel. To date, optical technology has been inserted at the physical layer (bottom layer) of the typical multilayer protocol stack. One reason for this application is that it allows straightforward extension of existing networks into the optical regime, without any modification of the higher layers. However, such an architecture does not maximally utilize the advantages offered by optical networking. We will propose a new architecture with a simplified protocol stack that takes advantage of the unique capabilities of WDM optical routing and switching technologies. We believe the first commercial true optical data networks will employ WDM because the technologies needed to implement these networks are now commercially available.

The nature of WDM technology suggests a network management and control scheme that uses off-band signaling, and network resource scheduling that is either centralized or distributed but coordinated. Tight coordination is required due to the lack of optical buffers that can temporarily store packets before processing and routing. This property is key to optical WDM data network protocol designs and presents interesting and sometimes severe constraints on network synchronization and end-to-end routing.

With the availability of buffering at routers internal to the network, higher network efficiencies in terms of throughputs and delays can be achieved with a time division multiplexing (TDM) scheme where the medium bandwidth is available to the packet in one large channel (>100 Gb/s) instead of fragmented (WDM) channels. The signaling information required for switching packets in this network would be

Manuscript received August 27, 1998; revised September 22, 1998. This work was supported by the Consortium on Wideband All-Optical Networks and the Defense Advanced Research Projects Agency (DARPA).

The authors are with the Lincoln Laboratory, Massachusetts Institute of Technology, Lexington, MA 02420 USA.

Publisher Item Identifier S 0733-8724(98)09304-9.

contained in the packet header and no centralized or highly coordinated scheduler is required. Network management and control could be simpler and more efficient, especially in a bursty data environment with many users attempting to share the same medium. However, the technology needed to implement a network of this type is still in the research stage and some typical networking components may not be feasible at ultrafast rates. At a minimum, optical logic gates and buffering at the channel rate are necessary architectural building blocks. In the second half of this paper, we will describe a novel architecture that takes advantage of the ultrafast components that have been developed to date. A testbed demonstration of this TDM network architecture in the next few years is feasible.

In this paper, we will present what we believe will be the evolution of optical data networks over the next few years. Beginning with WDM networks, we will describe services that may be offered as well as the advantages and tradeoffs that accompany optical networking. In the second major section, we will describe the more futuristic TDM network that may provide enhanced digital services to ultrahigh-speed users. The WDM networking section concentrates mainly on new architectures as many of the components needed to implement such a network are currently commercially available. The TDM section describes both the novel architecture as well as component development, as many of the necessary components are still being designed and demonstrated.

II. WDM Optical Data Networks

In order to design a WDM optical data network for the future that represents a major improvement in performance, substantially lower cost and offerings of new high-end services, it is important to speculate on the type of services that this network may offer and then optimize the network design for these services. Past forecasts of new applications have not been terribly accurate. In fact, the worldwide web application was not in anyone's predictions much before its appearance. Hence, we will not attempt to foretell any specific future applications but rather try to provide an abstraction of the type of services that can be offered in the future. We believe there are four classes of services that can be characterized abstractly and not specifically tied to particular applications.

1) *Transparent Optical Services to Support Conventional Electronic Network Services as Overlays:* For example, the network should be able to provide optical transport for all the usual services in present day networks, such as SONET, ATM, frame relay, and IP services.
2) *Large Point-to-Point Circuit-Switched Trunks on Demand:* OC-48, OC-192, and above to deal with stream traffic but with much less set-up time (of the order of tens of milliseconds) than current networks.
3) *Efficient Very High-Speed "IP" Packet Service Including Multicast:* For bursty, unscheduled large file transfers (100 Mbyte–10 Gbyte) at high access rates (>2 Gb/s).
4) *Analog Services:* narrow- and broad-band analog services with good amplitude, phase, and timing fidelity preserving features.

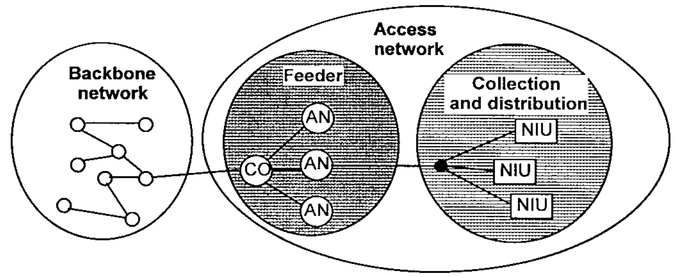

Fig. 1. Partitioning of a WDM optical data network into access and backbone networks (CO = central office, AN = access node, NIU = network interface unit).

In addition, differentiated quality-of-service (QoS) is highly desirable. We will show in the Section II how these services can be addressed by a combination of physical architecture, protocol design, and network management and control.

One way of viewing WDM optical data networks is through a physical and logical partitioning of the network into the access network and the backbone network, as shown in Fig. 1. Thus, local area networks (LAN's) are naturally part of the access network and wide area networks (WAN's) are part of the backbone network. This position emphasizes the strong economic reasons why these two classes of networks should be considered separately and designed and optimized based on very different performance metrics and goals. Though architects usually arrive at radically different designs for the two classes of networks, the two would have to seamlessly interconnect. Long-haul fiber cost (fiber, installation, and pass-through-fees) dominates backbone network total costs. Node equipment costs in that case, though not totally insignificant, are considerably less, especially when optical routing and switching are deployed in the future. Most backbone designs would treat fiber as the precious commodity and try to make maximum use of the capacity. Hence, network topology and scheduling algorithms for optimizing wavelength usage in WDM backbone networks is an area of intense research. On the other hand, in the access network, the tail fiber links have little opportunity to aggregate enough users to require very efficient multiplexing. Connectivity is more important than high throughput, and with the vast bandwidth available to the user at the entrance to the network, fiber bandwidth utilization can be readily and even freely traded for low-cost access equipment and network hardware. WDM optically routed metropolitan area networks (MAN's) should be considered as part of the access network though its design goals would be somewhere between the extreme cases of LAN's and WAN's.

One primary distinction between the two classes of networks is the way passive and active network components are chosen as the architectural building blocks of the networks. For example, to lower access costs, passive optical components without a lot of diagnostics and monitoring are often used for the access network. Broadcast stars, buses and trees, and passively routed optical networks such as the AON (all-optical-network), Level-0 (LAN), and Level-1 (MAN), [11], and the IBM rainbow network are prime examples. In the backbone, higher cost active optical components such as the frequency sensitive switch of the MONET network, [49], and frequency changers can also be used.

The LAN is an important part of the access network. Most work on WDM-based LAN's assume the use of a broadcast architecture and focus on the development of medium access control (MAC) protocols for sharing wavelengths among the different nodes in the network. While the broadcast nature of optical LAN's allows the communication to remain all-optical, it also limits the scalability of the solution. The broadcast architecture, with constraints on the number of available wavelengths, limits the network both in the number of users and in geographic extent and is only practical for networks supporting at most hundreds of users over short distances. In order to extend WDM optical data networks to the backbone and the wide area, some form of electronic packet switching and buffering, or a new all-optical data flow mechanism are necessary. We will discuss some of these concepts in the next few sections.

A. WDM LAN's

In recent years, there has been a wave of research toward the development of WDM-based LAN's [1]–[10]. Most of the proposed protocols and architectures are based on a broadcast network architecture. Some of the protocols are based on random access and consequently result in low throughput due to contention [3], [4]. Other protocols that attempt to minimize contention, through the use of some form of reservation scheme, require that the system be synchronized and slotted and many require multiple transceivers per node [5]–[8]. Despite the added complexity of these systems, most still fail to achieve high levels of utilization due to the use of inefficient scheduling schemes that often fail to deal with receiver contention, or ignore the effects of propagation delays in a high latency, high data rate network. A survey of WDM multiaccess protocols and their properties can be found in [1] and [2].

The purpose of the strawman system described in this paper is to achieve good throughput-delay characteristics, while maintaining simple user terminals. Previous efforts to simplify user terminals involved protocols that use fixed tuned receivers or transmitters [9], [10]. However, those protocols limit the number of users to the number of available wavelengths. Also, protocols using only a single fixed tuned device are often limited to the use of random access protocols that result in low channel utilization.

The system described here is novel in a number of ways. First, it uses an unslotted MAC protocol, yet results in high efficiency even in high latency environments. The choice of an unslotted protocol is driven by a desire, for simplicity, to eliminate the requirement to maintain slotting in the network. This is especially important during cold-start of a terminal. Unfortunately, unslotted MAC protocols such as carrier sense multiple access (CSMA) result in very low utilization in high-latency environments. Alternatively, high-latency protocols such as unslotted Aloha are limited in throughput to less than 18% [3], [5]. Another new attribute of our system is that it uses a centralized master/slave scheduler for each access network

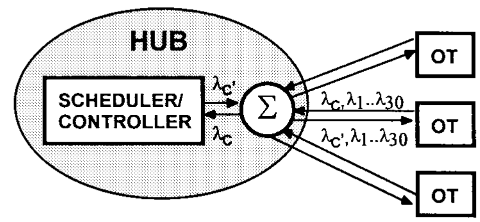

Fig. 2. Scheduler based LAN.

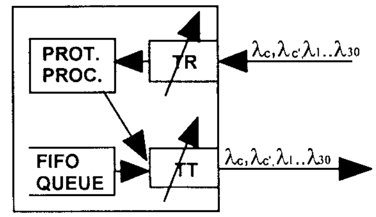

Fig. 3. Optical terminal (OT) with a tunable receiver (TR) and a tunable transmitter (TT).

for efficient resource allocation and network management. To overcome the effects of propagation delays, the scheduler measures the delays between the terminals and the hub and takes that delay into account when scheduling transmissions.

1) LAN Architecture: In the LAN, optical terminals (OT's) are connected via a simple broadcast medium such as a star-coupler located at a hub. As shown in Fig. 2, each OT is connected to the star using two fibers, one in each direction. Transmissions from all OT's on all wavelengths are combined at the star and broadcast to the OT's on the downlink fibers. Each OT is equipped with a single transmitter and receiver, both of which are tunable to all wavelengths, as shown in Fig. 3. This star based architecture has been proposed in the past for use in WDM LAN's [3].

2) MAC Protocol: The proposed protocol is based on a simple master/slave scheduler. All OT's send their requests to the scheduler, which schedules the requests and informs the OT's when and on which wavelength to transmit. Upon receiving their assignments, OT's immediately tune to that wavelength and transmit. Hence, OT's do not need to maintain any synchronization or timing information. There are three major aspects to the protocol. First, the protocol uses ranging to overcome the effects of propagation delays. Second, the protocol uses random access for the control channel and third, the protocol uses a simple scheduling algorithm with first-come–first-serve (FCFS) input queues and a look-ahead window to overcome head-of-line (HOL) blocking. These are described in more detail below.

a) Ranging: The protocol is able to overcome the effects of propagation delays by measuring the round-trip delay of each OT to the hub and using that information to inform the OT's of their turn to transmit in a timely manner. As an illustration, consider the timing sequence depicted in Fig. 4. In order for OT B's transmission to arrive at the hub at time T, the scheduler must send the assignment to OT B at time $T - \tau$, where τ is OT B's round-trip delay to the hub (including tuning delays). In this way the transmissions of different terminals can be scheduled back-to-back, with little dead-time between transmissions.

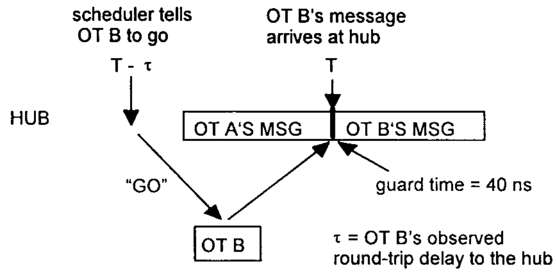

Fig. 4. Use of ranging to overcome propagation delays.

b) Access to the control channel: Reservations are made using a random access protocol in the control channel where terminals send reservation requests periodically and update their requests after waiting a random delay. These reservation messages contain the state of the queues at the requesting terminal. For example, each reservation message can contain the destinations with which the terminal wants to communicate and the duration of the requested transmissions.

Since reservation requests are sent on the control channel at random, it is possible for two or more terminals to send their request during overlapping time intervals. In which case their transmissions would "collide" and not be received by the scheduler. However, since reservation messages containing the state of the queue are sent periodically, all requests will eventually be received by the scheduler. As requests are answered by the scheduler, terminals update their requests to reflect the changes in their request queue. The random access protocol for accessing the control channel is described and analyzed in [18].

c) Scheduling algorithm: In a WDM system with a single transmitter and receiver per node, scheduling is constrained by the number of wavelengths W which limits the number of requests served during a slot to W. It is also constrained by the fact that each node has a single transmitter and a single receiver. Therefore, during a given slot, each node can be scheduled for at most one transmission and one reception. This, in fact, is a very similar problem to that of scheduling transmissions in an input queued switch. In the case of an input queued switch it is known that when a FCFS service discipline is employed, under uniform traffic, throughput is limited to $2 - \sqrt{2} = 0.585$ [13]. This throughput limitation is due to the head-of-line (HOL) blocking effect, where transmissions are prevented because the packet at the head of the queue cannot be scheduled due to a receiver conflict. It is also known that if nodes are allowed to look-ahead into their buffers and transmit a packet other than the one at the head of the queue, the effect of HOL blocking can be significantly reduced [14]. Scheduling algorithms based on bipartite graph matching algorithms have been proposed that achieve full utilization under uniform and nonuniform traffic conditions [15], [16]. However, it is also known that these algorithms are computationally intensive and require O($M^{2.5}$) operations to be implemented, where M is the number of input and output ports on the switch [17].

To simplify the implementation of the scheduler we propose the following simple algorithm for scheduling traffic. The algorithm is based on input queues and is made efficient through the use of a "look-ahead" window that allows the scheduler to look-ahead into each input queue and schedule requests that are not necessarily at the head of their queue. A look-ahead capability of k, allows the scheduler to look as far as the kth request in the queue. The algorithm is implemented on a slot-by-slot basis to form a schedule for the given slot. The algorithm works by maintaining N request queues, each containing the transmission requests from one of the N nodes in the network. The algorithm visits every node in random order and starting with the first request in the queue it searches for a request that can be scheduled. That is, it searches for a request for a transmission to a receiver that has not been assigned yet. The algorithm searches the queue until depth k has been reached. If a request has been found, a wavelength is assigned to it. This process is continued until either all of the request queues have been visited or all W wavelengths have been assigned. During the next slot, the algorithm starts anew with the first request in each queue.

Table I shows the maximum achievable throughput under uniform traffic, with 30 data wavelengths. When the number of nodes is equal to the number of channels and no look-ahead is employed (i.e., $k = 1$), HOL blocking limits throughput to 59% as predicted in [13]. However, a look-ahead window of just four packets can increase throughput to over 80%. As the number of nodes exceeds the number of channels, the effect of HOL blocking is drastically reduced. This is due to two factors. First, the probability that multiple nodes have a packet at the HOL to the same destination is reduced due to the increase in the number of destinations, and second, with fewer channels than nodes the algorithm has many more requests from which to choose a schedule of W transmissions. As can be seen from the table, the combination of more nodes than channels and a look-ahead window of four or five packets virtually eliminates the effects of HOL blocking on throughput, under uniform traffic. Scheduling multicast traffic in a WDM broadcast-and-select system is even more of a challenge because multicast messages have multiple intended receivers and trying to schedule transmissions in order to avoid receiver conflicts can be very inefficient. Simple and efficient multicast scheduling algorithms, based on random scheduling, are presented in [12].

In order to analyze the average queuing delay in this system we assume that packets arrive to each of the N nodes according to a Poisson random process of rate λ. We also assume that all packets are of the same length and take one slot to transmit and that the scheduler uses the slotted scheduling algorithm described in the previous section and that all transmissions are scheduled to occur at the beginning of a time slot.

Shown in Fig. 5 is the simulated delay for a system with 100 nodes and 30 wavelengths. Notice that with these values the arrival rate of new packets to a user cannot exceed 0.3 due to the channel constraint. Also notice from the figure that a look-ahead of just two packets can significantly help in reducing delays. However, a larger look-ahead window does not reduce delay any further because for these values of N and W, a look-ahead of just two packets essentially eliminates the HOL blocking effect.

TABLE I
THE MAXIMUM ACHIEVABLE THROUGHPUT FOR A NETWORK WITH 30 WAVELENGTHS, N NODES, AND A LOOK-AHEAD WINDOW k

N	k=1	k=2	k=3	k=4	k=5	k=6	k=7
30	0.59	0.71	0.77	0.81	0.83	0.85	0.86
35	0.69	0.83	0.90	0.94	0.96	0.98	0.99
40	0.79	0.95	0.99	0.99	0.99	0.99	0.99
45	0.89	0.99	0.99	0.99	0.99	0.99	0.99
50	0.96	0.99	0.99	0.99	0.99	0.99	0.99
60	0.99	0.99	0.99	0.99	0.99	0.99	0.99

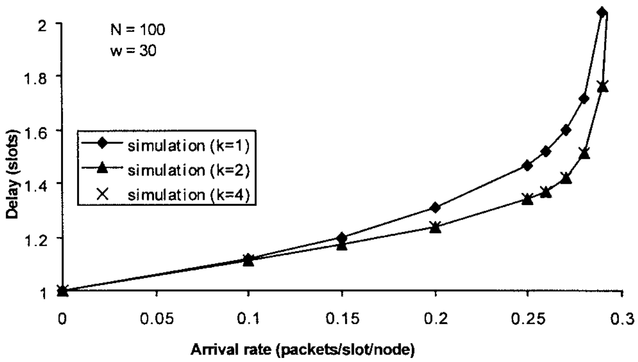

Fig. 5. Delay versus load for a system with 100 nodes and 30 wavelengths and a look-ahead capability (k).

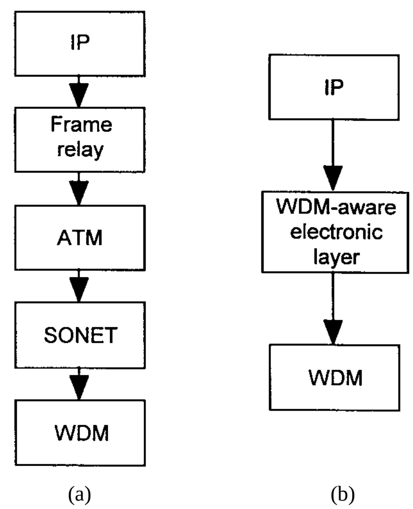

Fig. 6. (a) Typical protocol stack and (b) simplified protocol stack.

The algorithms described can be readily applied to passively routed optical networks such as the Level-1 MAN of the AON, [11], where the only dynamically tunable optical elements are lasers and receivers at the user terminals. Applications to networks with dynamically reconfigurable elements such as the frequency-sensitive-switch (FSS's) of MONET, [49], will require additional timing measurements to these elements and command and telemetry to reconfigure these elements for "just in time" service.

B. WDM Data Networks Beyond LAN—WDM Aware IP

The MAC protocol described in the previous section can be used to efficiently statistically multiplex packets from bursty data users in splitting losses and the broadcast LAN. The protocol limits this solution to relatively short distances and the number of available wavelengths also limits scalability to larger user populations. In order to efficiently extend WDM to WAN's, mechanisms are needed for statistically multiplexing bursty data in wide areas that cannot be spanned with a broadcast architecture. The approach we will describe in this section consists of putting more intelligence at the IP layer to recognize the resources available at the WDM layer that can be put to bear in the efficient delivery of data.

1) IP over WDM: 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. 6, 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) transport frames. This 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 and, in some scenarios, uncoordinated network management and control algorithms at the different layers acting against each other creating oscillations and further degrading network performance. Alternatively, one could use the simplified protocol stack shown in Fig. 6(b), where the IP traffic is carried directly by a simplified electronic layer. Such an arrangement would not only reduce the overhead associated with the different layers but would also allow the electronic layer to be "WDM-aware" and take advantage of network services offered at the optical layer. 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 design of the electronic layer.

2) Flow-Switching/Multilayer-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 [44], [46], [48], [56]. Tag switching uses routing protocols to predefine routes within the network; tags are then used to quickly assign flows to these routes. IP switching dynamically sets up layer-2 (e.g., ATM) virtual circuits for those connections that are perceived to be long.

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 multilayer switching approach where long duration sessions are switched at the ATM layer, and even longer duration and higher bandwidth flows are switched optically. While it appears that such a multilayer switching approach can reduce computation loads and processing delays in networks, many issues in the design of such a protocol remain to be resolved. While the detailed protocol is yet to be developed, there are some properties this protocol is likely to possess. For example, it is highly desirable for the application layer and the transport layer to inform the IP layer of the arrival and characteristics of large flows for switching. This can be done via an augmentation of the transport control protocol (TCP) layer. Predictions by the IP layer without higher layer inputs as is sometimes done today can be inaccurate and lead to inefficiencies. Since optical flows require several round-trips for setup and are likely to be assigned an entire wavelength for each transaction, the size of the flows should be at least of the order of the product of the physical propagation delay and the highest data rate of each wavelength for high throughput operations. Any finer grain assignment of network resources for flows such as the scheduled time division multiplexed WDM "B" service of the AON, [11], would require substantially more complexity and probably is hard to justify, especially if flow switching at the ATM layer is also available. Also, the point of electronic to optical conversion will most likely occur at the user interface to the network such as the network interface card of a high-end workstation. Thus, the WDM optical data network will almost appear to provide an all-optical service to the user except for an intimate dialogue and timed handshake between the network and the user terminal.

3) Virtual Topology Reconfiguration: In WDM networks, the physical topology is the one seen by the optical layer. It consists of passive or dynamically configurable optical nodes interconnected by fiber. The virtual topology, seen by the electronic layer, consists of a set of nodes interconnected by light-paths and in some cases time-shared light-paths. 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. The design of static network topologies has been studied extensively in the past [19], [51]. 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, subject to availability of wavelengths and node equipment. 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 wavelength sensitive passive routing elements such as a grating router or reconfigurations via commands of internal network element such as cross-bar switches and frequency-sensitive-switches. For example, consider a WDM ring network with one transceiver per node. Shown in Fig. 7 are two of the many ways in which the ring can be configured: on the left, nodes are connected in sequence using a single wavelength that is dropped at every node; on the right, nodes are connected in a different order using three wavelengths. Notice that the connectivity on the left does not allow both calls 1 and 2 to be admitted simultaneously, while the one on the right allows both.

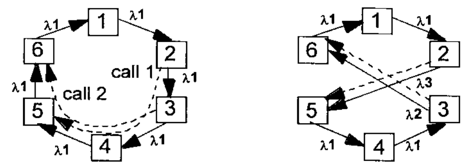

Fig. 7. Using WDM to reconfigure the electronic topology.

Reconfiguration effectively increases the overall system capacity by allocating bandwidth only where it is needed. However, there is an overhead associated with reconfiguration in that signaling messages need to be sent between switches to coordinate topology changes. In addition, reconfiguration may impact existing connections and these connections may need to be rerouted in the virtual topology. However, it should be possible to design a reconfiguration strategy that does not require the rearranging of existing calls. Such strategies may provide fewer benefits but would be simpler to implement and eliminate many coordination and network management problems.

Preliminary studies on reconfiguration of a WDM ring show significant promise [21]. Fig. 8 shows the call blocking probability versus call arrival rate for a WDM ring with 20 nodes and a bandwidth granularity of one call per wavelength. As can be seen from the figure, a reconfigurable WDM topology can support six times the traffic of a fixed WDM topology for the same blocking probability. Similar results were found for other topologies.

4) Traffic Grooming: 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, subwavelength capacity connections need to be allocated. This can be accomplished through the use of electronic multiplexing equipment that can aggregate low rate calls on to a higher rate channel (e.g., SONET multiplexers). However, if calls are indiscriminately multiplexed on to wavelengths, then each wavelength entering or leaving a node will need to be converted 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 [40], [41].

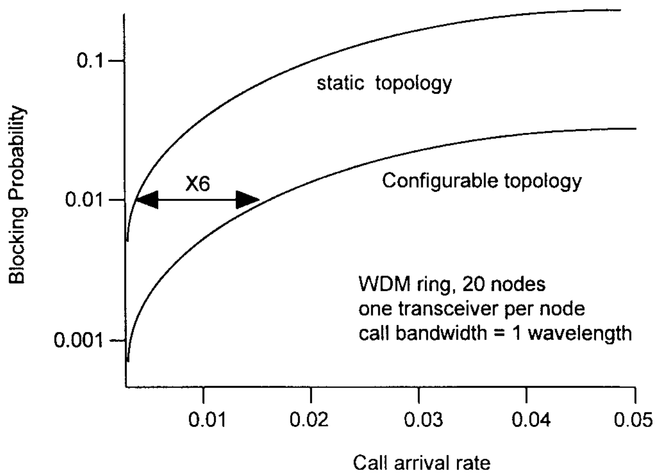

Fig. 8. Performance of reconfiguration in a WDM ring network.

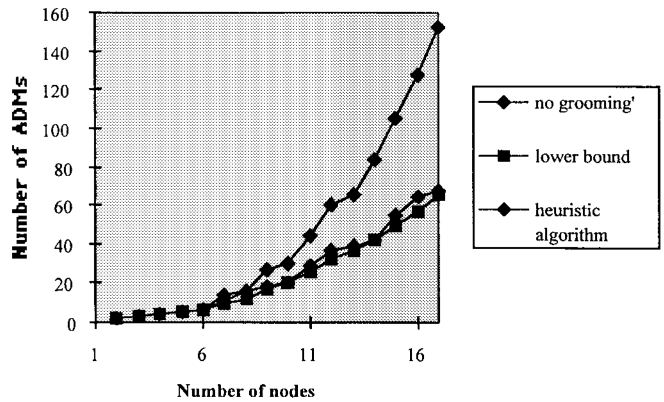

Fig. 9. Performance of grooming in a WDM/SONET ring network.

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 (ADM's) 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 ADM's 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. Optimal traffic grooming has been shown to be an NP-complete problem [40]. Specifically, traffic grooming studies on SONET ring networks with uniform all-to-all traffic show that grooming can result in a significant reduction in the number of required ADM's [40], [54]. Fig. 9 shows the number of ADM's required in a unidirectional SONET ring network using an algorithm developed in [40]. The algorithm provides a significant reduction in SONET ADM's, by as much as a factor of 2.5. Also, in many optimization algorithms studied, the solution that minimizes the number of ADM's also minimizes the number of wavelengths. Hence, a good grooming algorithm has the potential of minimizing network cost both in terms of efficient use of fiber and reducing use of electronics.

5) Protection and Restoration: Various failures, such as fiber cuts, line-card and switch failures, and software failures, can occur that disrupt network services [58]. Protection and restoration are two methods networks used to recover from these failures. Protection refers to hardware-based, preplanned, 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. Restoration works on general mesh networks and is typically more bandwidth efficient. Recently, fast protection mechanisms at the optical layer have been proposed for general mesh networks [22]–[24], and for ring networks [29], [37].

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. Therefore, at first sight, it might appear that optical layer recovery is not needed. There are, however, three significant advantages to providing protection and restoration optically. First, optical layer recovery can protect electronic services that do not have built-in recovery mechanisms or whose recovery mechanisms are slow (e.g., IP). Second, 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 bidirectional ring network. On the other hand, the fiber cut can be optically restored with a simple 2×2 optical switch, thereby 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. If protection were only performed electronically, there would be no guarantee that the SONET ring would be resilient to both a fiber cut and a line card failure. Finally, optical recovery allows the construction of arbitrary virtual topologies resilient to fiber failures. For instance, consider IP routers connected in a virtual ring topology over a WDM ring. Traffic may dictate that the IP routers be connected such that more than one IP link travels over the same fiber and hence a single fiber cut will disrupt many IP links. Since rings are in general only resilient to a single link failure, a single fiber cut of the optical ring may disconnect the electronic ring. Optical layer restoration solves this problem by restoring the fiber directly. For a simple example of optical loop-back protection, consider the two fiber bidirectional ring shown in Fig. 10. On each ring half of the wavelengths are used for working traffic and the other half are reserved for protection against a cut in the fiber on the other ring. In the event of a fiber cut, the wavelengths from the cut fiber can be switched onto the uncut fiber, using a two-by-two switch at the node before the fiber cut. These wavelengths can then be looped back to bypass the cut fiber and rejoin their original ring using another switch at the node immediately following the fiber cut.

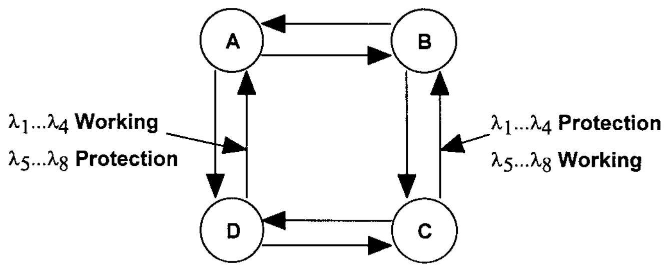

Fig. 10. Protection in WDM ring networks.

There are problems providing restoration at both the optical and electronic layers if the layers are uncoordinated. For instance, restoration can be duplicated at both the optical and electronic layers leading to a 75% loss in efficiency (assuming 50% efficiency for each layer of protection). In addition, differing time scales may lead to race conditions and topology oscillations. In the case of a SONET network, optical protection must somehow be completed before SONET starts its protection process or provisions made at the SONET layer for SONET protection process interrupts when the WDM protection process is going on.

C. LAN/MAN Access Network Physical Architecture

As we have alluded to before, WDM optical data networks can be partitioned into the access network and the backbone network. In the backbone, fiber is the precious commodity, especially when optical routing and switching will substantially lower node costs in the future. The network design problem there will concentrate on high utilization of the fiber assets and areas such as scheduling and protection and restoration algorithms using minimal overheads are important. It is reasonable to assume that by the time traffic appears at the backbone it would be of an aggregated form so that the backbone network would usually be providing entire light-path services to the access networks it interconnects. Finer grain services such as time division multiplexing of individual wavelengths would be much too complex for coordination over a wide area. In the access network, the situation is almost the inverse. Bandwidth efficiency is not paramount and should be freely traded for lower access cost. This trade is the key to providing high-end services in the future at substantially lower costs than what is achievable today.

The access network can be further subdivided into the *feeder network* (can be viewed as the MAN) and the *distribution network* (can be viewed as the LAN) as shown in Fig. 11(a). Generally, bandwidths of fiber plant closer to the end users are less precious than bandwidths of fiber closer to the backbone when some aggregation can and will have occurred. Thus, the designs of the distribution network and the feeder network can be very different. Some of the critical design issues are highlighted in the next two sections.

1) Passive Distribution Network: The distribution network is the interface network to the users. The number of users supported by a single distribution network should be the largest number without having to have active amplification in this part of the network. Depending on the design, this number

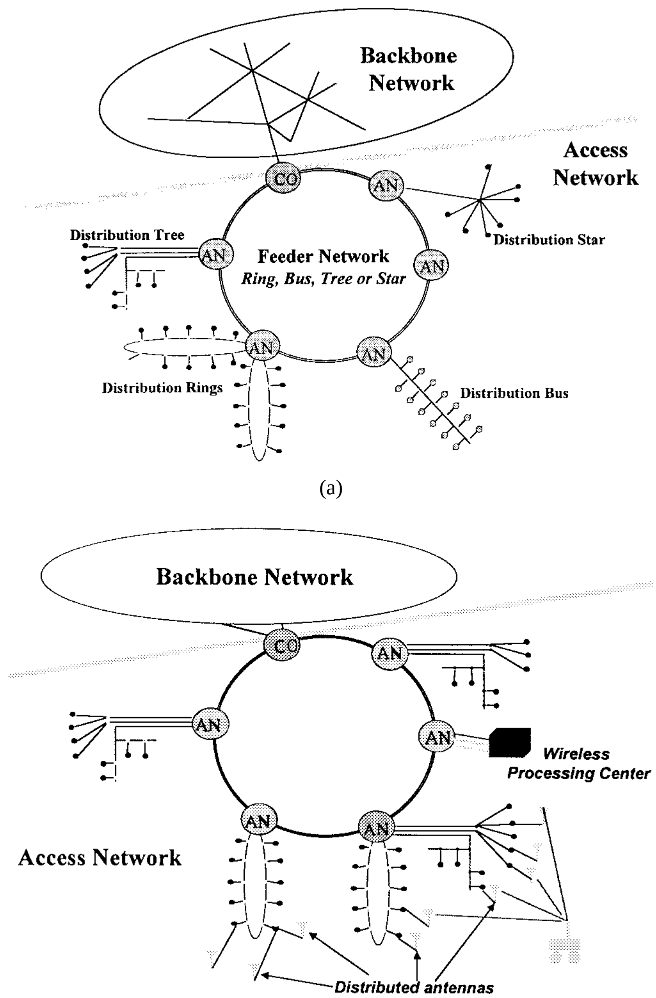

(a)

(b)

Fig. 11. (a) Partitioning of the access network into the feeder network and the distribution network (CO = central office, AN = access node) and (b) overlay high-performance wireless network on WDM optical access network

can range from tens to hundreds. A user can be a high-end workstation, a high-speed router, an ATM switch, or a gateway to an electronic LAN. Since the fiber to the end user is not a highly shared medium, bandwidth should be freely traded for a low-cost distribution network. Passive optical components such as couplers of various types should be used for this network and expensive routing elements should be avoided. The network topology can be a ring, star, bus, or tree. Since a passive ring, even if it is unidirectional, must be terminated or it will feedback on itself, it is really a bus, topologically. For unidirectional fiber operations, stars, buses, and trees can have duals for up-stream and down-stream operations, with local loop-backs at the access node. These passive topologies have very different attributes that should be considered carefully. A star splits power evenly but it needs more fiber than the other topologies and losses may become a problem if the physical span of the distribution network is large. The bus is awkward in the way it has to meander and touch every user. Power distribution is also tricky with a bus. To support the maximum number of users so that each user can get the minimum power but no more, the bus needs nonuniform taps.

This is very difficult logistically, particularly when a new user is added in midrun. Also, the last user of the bus experiences the maximum amount of fiber loss having the longest run (length proportional to the number of users) to the access node. This would ultimately limit the number of users supported by the network, often long before other resource limitations, such as the number of wavelengths, become an issue. The tree is perhaps the best topology for this network where the fiber length to the access node is only proportional to the log N, where N is the number of users. Since the most disadvantaged user suffers much less loss, this topology is more scalable to larger number of users. There are only a few levels of the tree structure so the number of different power splitters can also be small.

2) Optically Routed Feeder Network: The simplest form of the feeder network can also be a broadcast network, but since it is collecting traffic from the distribution network, some form of passive and/or active routing and switching will increase fiber utilization. In any case, the feeder will need to package traffic efficiently before sending it onwards to the backbone at its gateway (labeled CO in Fig. 11). Thus, a star topology is not favorable. Ring, bus, and tree each have their own merits. A ring topology that provides symmetric passive add–drop of wavelengths at each access node probably presents the most efficient use of fiber if the traffic pattern is symmetric and this topology is easy to protect with bidirectional dual rings. In the case where most of the traffic is going outward into the backbone (as is often the case in the present Internet), a tree structure may be more efficient. There is an access-node-to-central-office near-far problem for a passive ring, but long fiber runs incurring substantive losses can be off-set by optical amplifiers. At the gateway (CO) to the backbone network, full function frequency-sensitive-switches will aggregate wavelengths for the efficient use of long-haul fibers. A full spectrum of electronic routers and switches and SONET add–drop multiplexers does not have to be supported at every access node. They can be populated strategically at a few nodes of the feeder network, and point-to-point connections from each distribution network to the node with the desired service is a way to trade wavelength usage for fewer expensive electronic entities.

D. Overlay Services on WDM Networks

One advantage of a transparent all-optical WDM network is that the network will support any kind of signaling format, even analog signaling. This flexibility allows the overlay of other special application networks on top of the WDM infrastructure. As an illustrative example, consider the wireless network depicted in Fig. 11(b). Wireless antennas are proliferated at the entrance to the WDM optical access network. The radio frequency (RF) signals directly modulate WDM optical carriers and subcarriers instead of being digitized at the antennas. These signals are then optically routed to a centralized wireless processing center within a single access network service area. Thus, the processing center can take advantage of the many channels of RF inputs and use standard array antenna techniques to form simultaneous customized spatial antenna beams that can individually track users while simultaneously reject interfering users. This concept allows, in principle, as many antenna elements as there are to form user beams and the traditional notion of wireless systems broken into cells with hand over algorithms no longer has to be held. Such architecture will allow the realization of much higher data rates due to the greatly enhanced antenna gain, interference rejection, and the possibility of extensive frequency reuse.

III. Ultrafast TDM Optical Data Networks

In the WDM optical data network described above, network management and control are provided by off-band signaling and a centralized scheduler that grants requests for access to the network and tells sources and destinations which wavelengths to use for transmission and reception of messages. Bit interleaved TDM systems are similar to WDM systems because they divide the fiber bandwidth into a large number of lower rate channels, each operating at electronic rates. Again, centralized network management is required to regulate access to the network resources and in general, no user has access to the entire bandwidth of the network at any given time. Therefore, from an architectural point-of-view, bit interleaved TDM systems are isomorphic to WDM systems. However, implementing bit interleaved TDM systems is quite different from implementing WDM systems. For example, most components necessary for implementing a WDM point-to-point link are commercially available, while components necessary for TDM links are still in the research stage. Still, research on bit-interleaved TDM systems continues because there are advantages to using short optical pulses to represent digital data in optical networks. One advantage is that optical pulse streams may be all-optically regenerated, expanding the loss margins in transmission systems. In addition, utilizing high-speed optical logic, TDM systems may provide enhanced digital services such as packet routing and optical data stream encryption [60].

A different approach to shared-media optical networking, commonly referred to as slotted TDM, statistically multiplexed TDM, or packet switching, is fundamentally different from WDM and bit interleaved TDM. In a slotted TDM network, one ultrafast channel is shared by access nodes capable of bursting at the full media rate ($>$100 Gb/s). The signaling required for switching is contained in the header of the packet launched. There is no reservation necessary and no centralized scheduler. The total capacity of such a single channel local or metropolitan area network may be the same as for a WDM network or bit interleaved TDM network operating with a large number of lower rate channels. However, slotted TDM networks provide potential improvements in terms of user access time, delay, and throughput, depending on the user rates and traffic statistics.

Currently, existing lower-rate shared media networks such as ethernet at 10 Mb/s, token ring at 10 Mb/s, or FDDI at 100 Mb/s, utilize multiplexing schemes resulting in a single data stream rather than in many lower rate channels. In these networks, the media access nodes process the single stream data at the full media rate electronically. However, extensive data processing at 100 Gb/s rates is currently not feasible. Therefore, much of the data processing in a high-speed optical network will be performed at the slot rate, 10–100 Mslot/s, with some rudimentary processing, such as address recognition, performed by optical logic at the bus rate. As the high-speed optical processing components mature, we believe that slotted TDM networks will be the obvious choice to provide packet service in networks where users are capable of bursting at very high data rates. However, even if the users of the optical network are operating at lower data rates, the media access control (MAC) protocol is simpler in a slotted TDM system and may be able to more efficiently provide bandwidth-on-demand (BOD) services to a set of lower-rate users. In the following sections of this paper, we will describe our work on architecture and technology development for an ultrafast, slotted TDM multiaccess network.

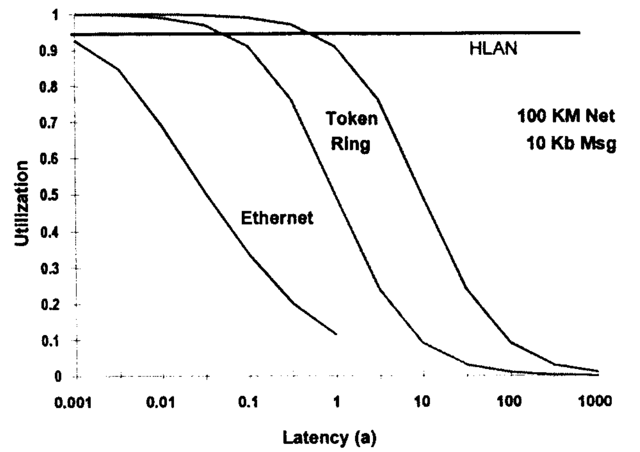

Fig. 12. Network utilization as a function of latency for ethernet, token ring, FDDI, and HLAN protocols. This calculation assumes a 10-Kb message size and a 100-km network.

A. HLAN Architecture

The optical TDM network we envision will run at a bus rate of at least 100 Gb/s but will serve a heterogeneous population of users such as high-end single users, high speed video servers, terabyte media banks and networks of supercomputers, operating at speeds from 1 to 100 Gb/s. The most important features of the network are to provide a backbone to interconnect high-speed networks, to transfer quickly very large data blocks, to switch large aggregations of traffic, and to provide flexible, lower-rate access to users. Although we are designing a 100 Gb/s network, the architecture will be scalable to much higher rates.

There are two characteristics of these 100 Gb/s slotted TDM networks that significantly impact the architectural design. First, they operate with long propagation delays, i.e., there are many data packets in flight at one time in the network, a high-latency environment. The well-known protocols governing today's electronic packet networks, where a typical latency "a" may range from $a = 0.01$–10, were never designed to operate in high-latency networks. Here, "a" is the packet propagation time divided by the packet transmission time. For example, in ethernet-based networks, "a" must remain much less than one in order to maintain reasonably high throughputs. In extending the operating rate of standard (10 Mb/s) ethernet to 100 Mb/s and Gigabit speeds, certain restrictions are necessary. For 100 Mb/s ethernet, the physical length of the data bus is limited to spans less than 100 m. Alternatively, a switch may be inserted in the bus that gives each transmitter its own line, so that collisions are impossible. For Gigabit ethernet, to assure collisions could be detected a data bus of 10 m is required. This length is not practical so the Gigabit ethernet bus is kept at 100 m but he minimum packet length is extended. Alternatively, a switch may be inserted in the bus as in the 100 Mb/s ethernet case. Clearly, these extension techniques will not scale gracefully to 100 Gb/s.

The optical TDM network we envision will have thousands of packets of data in flight in the network at any given moment ($a > 1000$). To study how standard protocols would perform in this high latency environment, we performed simulations to predict the network utilization as a function of the latency in the network for a 1-Mb/s ethernet protocol, a 10 Mb/s token ring protocol and a 15-Mb/s FDDI protocol. The results, shown in Fig. 12, indicate that existing protocols will not scale to ultrahigh speed optical TDM networks.

Second, there is a limited amount of processing that can be performed at the 100 Gb/s rate. These two characteristics, high latency and limited processing, make it difficult to provide simultaneously guaranteed bandwidth (GBW) and bandwidth-on-demand (BOD) services. Also, it is difficult to insure both efficient and "fair" bandwidth sharing in the presence of moderately loaded or overloaded network traffic conditions.

The algorithms designed to manage the network must be simple enough to be executable with a minimal amount of processing available at the optical bit rate. There have been some articles addressing these problems in the literature [61], [62], but they do not appear to satisfy simultaneously all of the above criteria. We have developed a frame-based slotted architecture called HLAN (helical local area network) [63], [64] which satisfies the above criteria. In early work, the HLAN architecture was implemented on a helical ring physical structure, but most recently has been explained using a folded bus implementation.

HLAN is a frame-based architecture, which is implemented on a folded unidirectional bus as shown in Fig. 13. A head-end generates frames of empty slots and puts them on the bus. Guaranteed bandwidth traffic is transmitted on the GBW segment, bandwidth on demand traffic is transmitted on the BOD segment and data is received on the RCV segment. Note that only one bus is required and that all receivers are downstream of the transmitters. GBW services are provided to users who request them by the head-end. The head-end allocates reserved slots to the node on the GBW segment. Users can access the HLAN at their guaranteed rate using only a counter, a flip-flop, a few gates, and slot marker detection logic [63], [64]. Using existing logic, HLAN slot rates of 10–100 Mslots/s can be supported.

Fair and efficient BOD service is provided via the BOD segment. The head-end creates credits and distributes them at

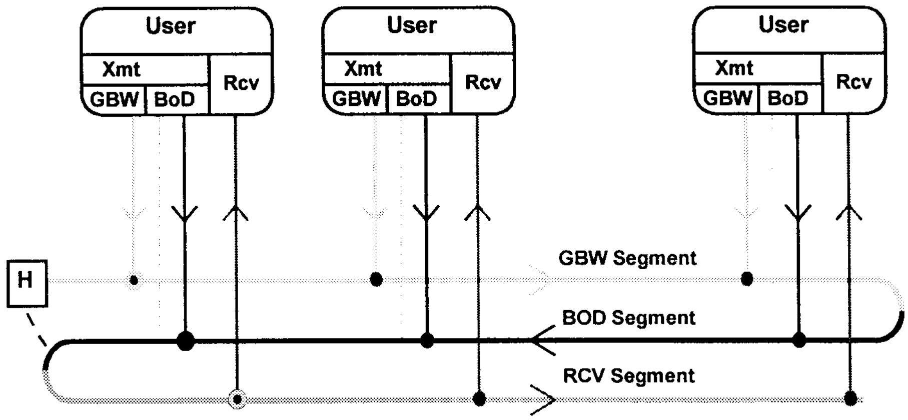

Fig. 13. Helical local area network topology.

a given rate. A user who has data to send and who has a credit can burst data onto the network when the first empty slot comes by. Users with data to send but no credits must wait for the head-end to distribute more credits. Users who receive credits but have no data to send may accumulate only a limited number of reserve credits. To prevent lockouts from high-traffic upstream nodes, the head-end monitors free slots at the end of the BOD segment. If no free slots are observed, the credit allocation rate is reduced, decreasing the bandwidth available to the individual nodes. If, on the other hand, multiple free slots are observed, the credit allocation rate is increased, returning more bandwidth to the individual nodes. The system is "fair" because all users will get equal throughput when they have data to send and the network is at high loading.

The algorithm is also efficient because few slots that could be used are left empty. Fig. 14 shows a plot of the average wait the first and last user nodes encounters in accessing the network as a function of network load. The solid lines show the average wait when the greedy algorithm is used and the dashed line shows the performance when the HLAN protocol is used. This simple HLAN algorithm performs like the "greedy" algorithm under lightly loaded network conditions and transitions to fixed TDMA as the load increases [63], [64]. Stability and speed of convergence of the HLAN credit allocation rate to its correct value is controlled through the use of an integration time constant which filters the "free slot" feedback information at the head-end.

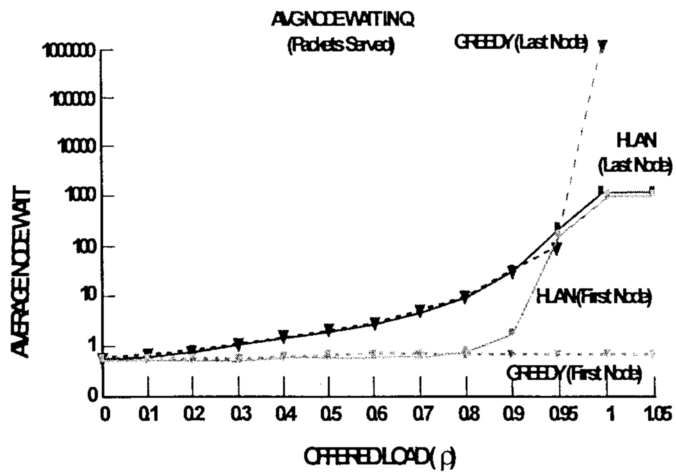

Fig. 14. Average wait for the first and last user to access the network as a function of network load using the greedy algorithm (dashed line) and the HLAN algorithm (solid line) to regulate access.

B. Ultrafast Optical Components

A high-speed single-channel network has many desirable features, as outlined above. The components needed to implement such a network are currently in the research stage. In some instances, there exist no ultrahigh-speed counterparts to vital electronic network components. For example, ultrafast optical random access memory has not been invented. Therefore, protocols that are developed to manage these high-speed TDM networks must reflect the limited set of technological building blocks that are available. While advancements in components for WDM networks are still being made, WDM networks can be assembled today with commercially available technology. The same cannot be said for TDM networks. As a result, TDM network development is tied to both technology and architectural development. Therefore, we will highlight the state-of-the-art of TDM technology and use this section to point out some of the hardware limitations and capabilities that determine the applicability of newly developed TDM access protocols.

Fig. 15 shows a block diagram of some of the hardware components needed to implement this high-speed, yet simple HLAN architecture. The block diagram shows the components in the receiver node, but the components in the transmitter node are nearly identical. The slow protocol logic unit is used to perform such functions as start-up, failure recovery, bypass control in the case of component failure, and interfacing to the fast protocol logic unit. The slow protocol logic unit operates at a rate much lower than the slot rate. The fast protocol logic unit operates at the slot rate and regulates access to the media. This unit is interfaced to the ultrafast optical components shown in the gray boxes, at the interface between the optical bus and the node electronics. These ultrafast optical components consist of all-optical switches and buffers capable of operating at the 100 Gb/s bus rate. Currently, the only logic gates operating at 100 Gb/s are all-optical logic gates [65], [66]. These gates may be used to perform pattern matching,

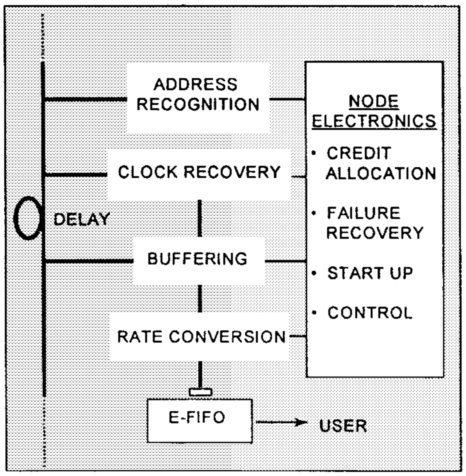

Fig. 15. Block diagram of hardware components needed in a user receiver node. Components shown in the gray boxes must operate at the bus data rate and may be implemented all-optically.

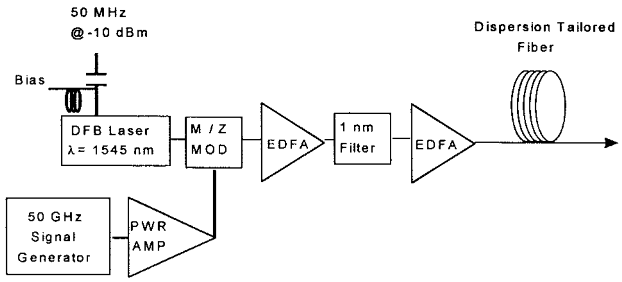

Fig. 16. Block diagram of a 100-GHz soliton compression source. The dark lines symbolize electronic connections and the narrow lines, optical connections.

or address recognition, clock recovery, rate conversion, and possibly buffering. In the transmitter node, the optical buffers will hold data packets while they wait for access to the network and in the receiver node they will present multiple copies of the incoming data packets to a demultiplexer or rate converter. After rate conversion, the data will be detected and further processing will be performed electronically. In the following section, we will describe our latest results on developing the technological building blocks necessary for demonstrating a 100 Gb/s slotted TDM network. These technologies include high-speed sources and transmission, all-optical switching, buffering, and rate conversion. After a brief description of access node operation we will discuss other important high-speed network services that may be enabled by this technology.

1) High-Speed Short Pulse Generation and Transmission:

a) Sources: These high-speed optical TDM networks will utilize return-to-zero (RZ) coding of the data. Therefore, optical clock sources, capable of generating short optical pulses are essential components in these systems. We require sources for optical modulation and demultiplexing, that may operate at electronic rates, 10–40 Gb/s, as well as sources for synchronization to buffers and optical switches that operate at the 100 Gb/s bus rate. In a previous publication on these types of networks, we described the state-of-the-art for harmonically modelocked erbium doped fiber ring lasers [64]. These lasers were capable of generating short pulses at high repetition rates, but were susceptible to pulse-to-pulse energy fluctuations because of the long upper state lifetime of the erbium and the multiple pulses in the laser. Since that time, a number of more stable erbium doped fiber lasers capable of generating pulses with 1–2 ps pulse durations have been demonstrated. These sources may rely on pulse shortening mechanisms within the laser cavity, such as soliton compression or nonlinear polarization rotation [67], [68]. Other sources capable of supplying short pulse clock streams have relied on compression techniques, external to the laser cavity [69]–[72].

Laser modelocking techniques that require gain or loss modulation at the optical clock rate do not scale gracefully to pulse train generation above 40 GHz [73]. At the higher repetition rates, 100 GHz and above, clock streams have been generated by passively multiplexing multiple lower rate streams. However, these clock streams suffer from timing and amplitude fluctuations resulting from imperfections in the passive multiplexers [66]. One alternative to passive multiplexing is soliton compression. Soliton compression sources have generated very high repetition rate streams of very short pulses with high peak powers [74]–[76]. In addition, these sources are electronically tunable and made be synchronized using standard electrooptic phase locked loop techniques. Recently, we extended a soliton compression technique utilizing a Mach–Zehnder amplitude modulator to 100 GHz [65].

100 GHz clock pulses may be generated by the soliton compression source shown in Fig. 16 [65], [75]. This source consists of a DFB laser, a high-speed electrooptic modulator, an erbium-doped fiber amplifier (EDFA), and a length of dispersion tailored fiber (DTF) whose dispersion profile is a step-wise approximation to dispersion decreasing fiber. The DFB laser is driven by a CW bias current as well as a small dither signal at 50 MHz. The dither signal broadens the laser linewidth and increases the stimulated brillouin scattering (SBS) power threshold in the DTF. The output from the DFB laser is coupled through a high-speed Mach–Zehnder amplitude modulator that is DC-biased at the transmission null. RF modulation, at 50 GHz is also applied to the modulator. RF modulation of the continuous wave (CW) optical input provides a periodically modulated intensity at twice the applied frequency, yielding two phase synchronous optical carriers, offset by 100 GHz, at the modulator output. These two tones beat at their difference frequency, are amplified in a high-power EDFA, and compress into a 100 GHz soliton pulse stream. An autocorrelation of the 100 GHz, 1.6 ps pulse stream, is shown in Fig. 17.

High repetition rate optical clocks have also been demonstrated utilizing frequency multiplication [77] and rational harmonic modelocking techniques [73]. In conventional harmonically modelocked lasers, the frequency of the RF drive applied to the gain or loss modulator is equal to the laser output pulse rate and is a harmonic of the cavity fundamental frequency. To achieve frequency multiplication, the modulator in the laser oscillator may be overdriven so that the laser output pulse rate is two or three times the RF drive frequency. Another frequency multiplication scheme utilizes a modulator biased at its transmission null to double the laser output pulse rate as described above. In the rational harmonic modelocking

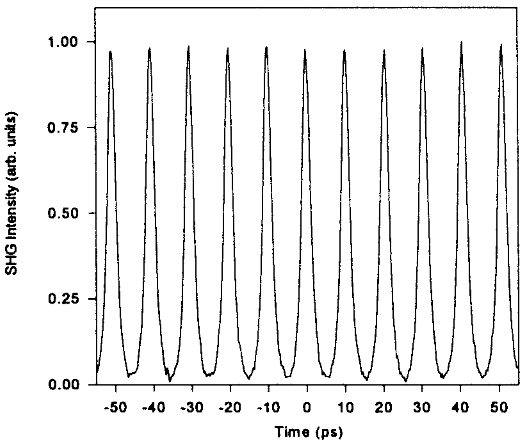

Fig. 17. Autocorrelation of the 100 GHz stream of 1.6 ps pulses generated by the soliton compression source shown in Fig. 16.

approach, the ratio of the RF drive frequency to the cavity fundamental is a noninteger rational number. In this case, the laser output pulse rate is the least common multiple of the RF rate and the cavity fundamental [77]. In this rational harmonic technique, pulses circulating in the optical oscillator are shifted with respect to the RF gain or loss modulation by a fraction of that RF drive period on each round trip. After a number of round trips, equal to the ratio of the data rate to the modulation rate, a high-rate effective modulator response is realized.

b) Transmission: An important issue affecting the success of any future testbed demonstration is the propagation of 100 Gb/s data packets over 100 km long lengths of fiber. Recently, 640 Gb/s single-stream soliton transmission was demonstrated over a 60 km link [78]. In that experiment, the transmission distance was limited by polarization mode dispersion (PMD). In another transmission scheme, multihundred Gb/s TDM streams have been propagated over near-zero dispersion fiber links with low optical powers [79], [80]. Using low optical powers minimizes nonlinear effects in the fiber and low dispersion minimizes pulse spreading. However, in a multiaccess network there will be large gain and loss variations along the fiber bus and it may be difficult to restrict the entire fiber span to zero-dispersion fiber. Still, transmission of 100 Gb/s data streams over long distances at large amplifier spacings may be possible if dispersion managed fiber links and all-optical regenerators are used. Dispersion managed soliton transmission is a hot research topic and many exciting theoretical results have recently been reported [81]–[85]. However experimental studies of the limitations of this new transmission scheme for 100 Gb/s data streams have not been reported. Another method that may be used to extend the propagation distance for ultrafast data streams is optical regeneration. A regenerator consists of a clock recovery unit and an optical switch. The clock recovery unit synchronizes a local clock with the incoming data stream. The incoming data stream and the synchronized clock become the logical inputs to an optical AND gate and the noisy input data signal is transferred to the local clock signal. The regenerator output is an optical data stream that has been retimed and reshaped. If the regenerators are placed closely enough, transmission distances may become limitless. All-optical regenerators have been demonstrated [86]–[88], and will be discussed further in the section on optical switching applications.

2) User Node Data Processing:

a) Address recognition: In electronic networks, address recognition is a straightforward process and is usually done by loading the received address into a register and comparing it with a stored bit pattern. For very-fast photonic networks, storing an address in a register and accessing it randomly is problematic. Attractive solutions for address recognition in very-fast photonic networks include source signal coding and orthogonal signal coding. Source signal coding includes special flag codes in the address stream that are prohibited in the packet data stream. Source signal coding can simplify system design but usually involves special hardware to encode the data and uses extra signaling bandwidth. In the optical domain such hardware may be impractical. Orthogonal signal coding has the potential to take advantage of the special properties of optical transmission and conserve data bandwidth. Wavelength and time are two dimensions that one can use for orthogonal signal coding. Signaling in a separate wavelength is attractive because it is easy and inexpensive in the optical domain but dispersion may delay the signaling channel differently for different receivers in a distributed network. Using the time dimension for signaling, such as using special marker codes or self-synchronizing codes, is attractive because the signal coding travels on the same path as the data and remains invariant in a distributed network. In initial experiments, self-synchronizing codes have been implemented using uniquely spaced pulse pairs [89] and different amplitude marker pulses [90] to signify the beginning of the address stream. The determination of this absolute phase of the data packet is called stream capture. After stream capture, ultrahigh-speed address recognition can be performed using all-optical logic gates.

All-optical logic gates and switches have been investigated primarily for wavelength conversion of high-speed data signals in WDM systems and for demultiplexing in bit-interleaved TDM systems. For TDM applications, all-optical AND and NOT functions have been demonstrated using four-wave-mixing in waveguides [91], [92], cross-phase/gain modulation in nonlinear optical loop mirrors (NOLM's) [93]–[95], semiconductor NOLM's [96], [97], Mach–Zehnder devices [98], [99], and ultrafast nonlinear interferometers [100]–[102].

Optical switches fundamentally rely on material nonlinearities. While many optical switching experiments have utilized gain and absorption nonlinearities in optical waveguides, much more popular have been demonstrations utilizing interferometric switches containing materials whose refractive index is intensity-dependent [103]. Interferometric switches are popular because they are the most versatile, capable of demonstrating AND, NOT, and XOR functionality using a single switch configuration. As an example, consider the simple Mach–Zehnder interferometer shown in Fig. 18. To operate as an optical switch, at least one arm in the interferometer must contain a material whose refractive index, n, is intensity dependent, $n = n_0 + n_2 I_c$. Here, n_0 is the linear refractive index, n_2 is the nonlinear refractive index coefficient and I_c is the intensity of light. Then, assuming the beamsplitters in the

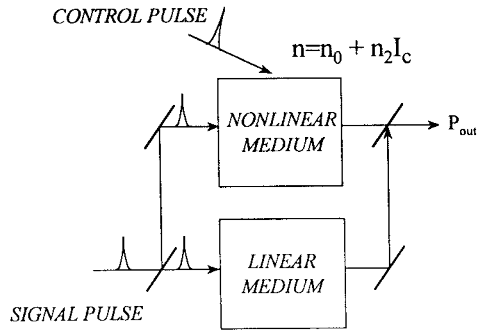

Fig. 18. Nonlinear Mach–Zehnder interferometer.

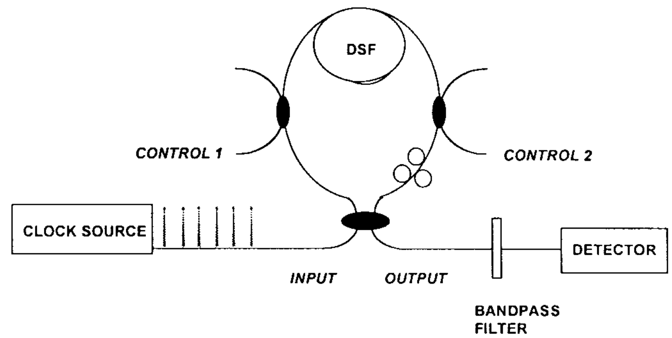

Fig. 19. NOLM configured for XOR operation.

interferometer have a 50/50 splitting ratio, the interferometer output signal intensity is related to the input signal intensity by $I_{\rm out} = I_{\rm in} \cos((\phi_l + \Delta\phi_{nl})/2)$. Here, ϕ_l is the initial phase (length) difference between the two arms of the interferometer and $\Delta\phi_{nl} = (2\pi/\lambda)n_2 L I_c$ is the phase difference induced by the presence of a control beam with intensity I_c. In this expression, L is the interaction length of the signal and control pulses and λ is the wavelength of the light. The control pulses may be distinguished from the signal pulses by wavelength, polarization, or the direction of propagation. From this simple picture it is clear that to obtain all-optical switching, that is changing the interferometer transmission from a maximum to a minimum, or vice versa, it is necessary to induce a phase shift of π in the nonlinear material. The peak intensity necessary to induce this π phase shift is given by $I_c = (\lambda/2n_2L)$. In optical fiber, n_2 is approximately -2.8×10^{-16} cm^2/W. A rule of thumb is that an optical control pulse with a peak power of 1 W, will induce a π phase shift in a fiber interaction length of 1 km.

In the most common interferometric switch demonstrated to date, a 50/50 fiber coupler and a long length of optical fiber are used to construct a Sagnac interferometer. This type of interferometric switch has been referred to as a nonlinear optical loop mirror. Fiber has been a common switching material because the refractive index nonlinearity, owing to the Kerr effect, is essentially instantaneous and has no corresponding transmission nonlinearity. Therefore, terabit per second switching may be possible in these fiber gates. The disadvantage of using fiber as the nonlinear material is that n_2 is relatively small, and long interaction lengths are necessary to achieve reasonable switching energies. Long fiber lengths in optical switches are undesirable because of pulse distortion effects and the large physical size of the device. Still, many important network components have been demonstrated using fiber based NOLM's.

For example, all-optical address recognition has been demonstrated using a NOLM configured as an XOR gate. While previous experiments utilized AND gates based on four-wave-mixing in semiconductor optical amplifiers (SOA's) to perform address recognition, the number of potential addresses was limited to a fixed set of "keywords" because some address patterns could be matched by patterns other than their logical compliment [104]. Using all-optical XOR gates to perform address recognition removes the limitation on the available number of keywords. Also, utilizing XOR gates in the address recognition circuit may facilitate the recognition of more than one address per hardwired node. For example, a single circuit might be able to recognize a local address as well as a broadcast address.

We have demonstrated all-optical pattern generation and matching at 10 Gb/s using a NOLM configured as an XOR gate as shown in Fig. 19 [105]. Clock pulses are coupled to the input port of the NOLM, are split at the fiber coupler into two counterpropagating components that traverse the 2-km length of fiber, and are interfered back at the coupler. In the absence of a control pulse, the recombined clock pulses exit the NOLM through the input port. In other words, for zero control input to the NOLM, there is zero signal (clock) output through the output port. A control pulse, coupled into either of the two control ports shown, will copropagate with one of the clock components and nonlinearly shift its phase via cross-phase modulation. The control peak power is chosen to induce a phase shift of π on the copropagating clock component pulse, thereby switching the NOLM from the reflecting state to the transmitting state. At the output port of the NOLM, an optical bandpass filter distinguishes the switched out signal (clock) pulses from the control pulses. If a control pulse is simultaneously present at each control port, both counterpropagating clock components are nonlinearly phase shifted by cross phase modulation and the interferometer remains in the reflecting state. Therefore, the NOLM acts as an all-optical XOR, with the two control streams as the logical inputs to the gate. Fig. 20(a) and (b) shows the output from the NOLM when a controlling data pattern is present at control port 1 only, and control port 2 only. Fig. 20(c) shows the output from the NOLM when both control signals are simultaneously present at the switch. Note the excellent contrast ratio (25 : 1) at the output of the all-optical XOR gate.

b) Clock recovery: Besides self-synchronizing codes, the most likely form of clock recovery for a slotted OTDM access node will be phase synchronization using all-optical or electrooptic phase locked loops. Picosecond accuracy synchronization of 40 GHz clock streams has been demonstrated using electrooptic phase locked loops [94]. Typical lockup times have been on the order of microseconds, the equivalent of 10–100 slot times, but improvements in the lock-up time are expected using specialized electronic circuits. Much slower lock-up times have been demonstrated for all-optical clock (typically laser injection locking) recovery

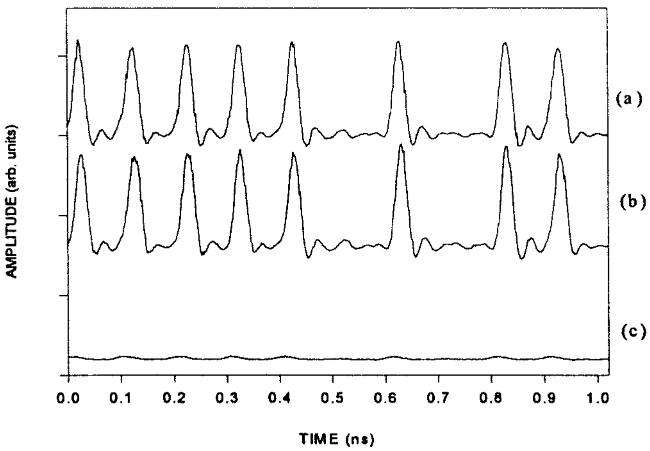

Fig. 20. Output from the XOR NOLM when (a) data is coupled to control port 1 only, (b) data is coupled to control port 2 only, and (c) data is coupled simultaneously to control ports 1 and 2.

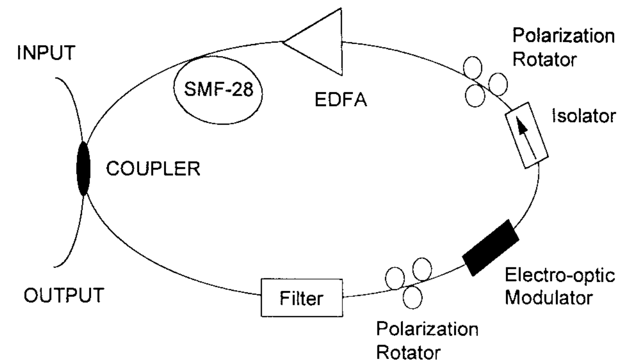

Fig. 21. Block diagram of a compensating recirculating loop buffer.

techniques. A weakness of this phase locking technique is that it synchronizes only the relative phase of two optical signals. That is, electrooptic and all-optical synchronization schemes based on injection locking optical sources and phase-locked loops cannot be used for stream capture.

c) Buffering: Buffering of ultrafast optical data packets will be an important function in the user access nodes. Since bandwidth-on-demand represents an important network service and because all users may not generate or accept data at the ultrafast bus rate, optical buffers operating at the network data rate are key components. To date, a variety of optical storage elements have been demonstrated [106]. Most promising for slotted TDM applications are the regenerative buffers employing optical logic gates [107]–[110] and the compensating fiber loop buffers [111]–[113]. These two types of buffers are attractive from a network standpoint because they can be designed to store a single packet at a single wavelength for hundreds of circulations. Because of the cyclical nature of the memory, reading of the information is restricted to multiples of the round trip loop time (packet length). Therefore, truly random access memories do not yet exist—one practical consequence of this is that they cannot be used to synchronize signals as might be required in a wide area packet network. Also, current implementations of optical memory are not practical for conventional electronic network applications because they are complex and costly. Conventional network architectures using link-by-link flow control require about a round trip propagation delay worth of memory (10^3–10^4 buffers at 100 Gb/s) per node and therefore must be modified for OTDM applications.

A typical configuration for a compensating fiber loop buffer is shown in Fig. 21. The loop consists of an erbium doped fiber amplifier (EDFA), a length of standard single mode fiber (SMF-28), a 10% output coupler, a polarization-sensitive electrooptic modulator, polarization rotators, and an isolator. The EDFA provides gain to signals circulating in the loop and the isolator forces unidirectional operation. The modulator provides timing stability because pulses experience lower loss if they arrive at the modulator at the transmission peak. The polarization rotators, the optical fiber, and the polarizing waveguide in the modulator combine to create an artificial fast saturable absorption effect that provides bistable operation in the loop. That is, the loop is simultaneously stable for the pulses (ONE's) and for the ZERO's. Much of the recent work on optical buffers has centered on techniques to extend the attainable data rates. In active loop buffers using electrooptic (EO) modulators to provide the data timing stability, the data rate (determined by the RF driving frequency) can be multiplied by factors of two or three by overdriving the (EO) modulator [76], [77]. Utilizing semiconductor optical amplifiers (SOA's) as the modulator in the cavity is another way to extend the buffer operation to high data rates [113]. Also, an asynchronous phase-modulated optical fiber ring buffer has been demonstrated in which the amplitude modulator in the loop was replaced by a phase modulator and a filter. In this type of memory, the loop round trip time is slightly detuned from the frequency driving the phase modulator. One advantage of this scheme in the receiver node is that incoming data packets can be injected into the loop without precise phase synchronization to the phase modulator [114].

A block diagram of a regenerative buffer, in which the data is reshaped and retimed on every circulation, is shown in Fig. 22(a). These types of buffers are very attractive because their regenerative properties minimize the effects of accumulated timing jitter. In proof-of-concept demonstrations, these regenerative buffers are typically configured as circulating shift registers with inverters, as shown in Fig. 22(b), because the experimental set-up is less complex. However, fiber-based optical switches are not suitable for these single-packet buffer designs because of the long lengths of fiber needed to achieve switching. Instead, these regenerative buffers will utilize compact interferometric switches where the nonlinear material is a semiconductor waveguide.

Typically, relevant semiconductor nonlinearities are orders of magnitude larger than in fiber and required interaction lengths are on the order of millimeters, rather than kilometers. This comparatively small interaction length in conjunction with monolithic integration techniques can realize potentially huge gains in size reduction, manufacturability, and reliability. Semiconductor materials are not without their drawbacks, however. In addition to subpicosecond components of the semiconductor nonlinearity due to two photon absorption, carrier heating, and virtual electronic processes [103], custom-

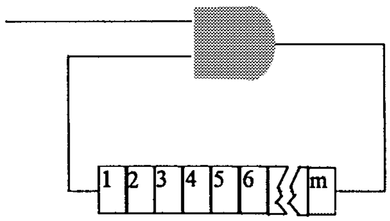

(a)

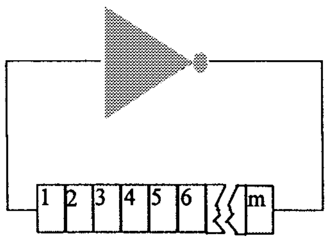

(b)

Fig. 22. The logical equivalents of (a) a regenerative buffer and (b) an inverting circulating shift register.

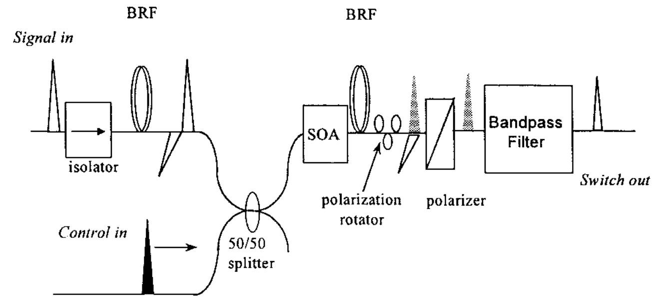

Fig. 23. Block diagram of the ultrafast nonlinear interferometer or UNI. BRF is birefringent fiber and SOA is semiconductor optical amplifier.

arily there are free carrier (recombination) components with relaxation times of tens of picoseconds to nanoseconds. The free carrier recombination time is the primary effect limiting the speed of SOA-based all-optical switches. At high carrier densities, the carrier lifetime may be reduced to a few tens of picoseconds by nonradiative recombination mechanisms such as Auger recombination [115], and by the presence of optical holding beams [116]. Still, switches based on these free carrier effects will be limited to operating rates of approximately 100 Gb/s in the best case. This limitation can potentially be overcome using semiconductor materials with primarily ultrafast nonlinearities, such as passive semiconductor waveguides. Alternatively, it can be circumvented by using innovative optical switch geometries, such as balanced interferometric switches.

For example, the problem associated with utilizing semiconductor nonlinearities in a two arm interferometer optical switch geometry is related to components of the nonlinear response that relax on time scales longer than the bit period. If the control stream induces carrier density changes in the semiconductor with recombination times of up to a few nanoseconds, the phase modulation imposed on one of the signal components is pattern dependent leading, ultimately, to fluctuating output signal pulse intensities. Stated another way, the value of a particular output bit no longer depends exclusively on its associated input signal and control bits, but rather on the exact pattern of control bits over a time approximately equal to the carrier recombination time, i.e., a form of nonlinear intersymbol interference (ISI).

A device geometry that circumvents the effects of long lived refractive index nonlinearities is the single-arm interferometer (SAI) [65], [100]–[102], originally developed in the context of femtosecond pump-probe studies of various nonlinear processes in semiconductor waveguides. Fig. 23 is a block diagram of one particular SAI implementation known as the ultrafast nonlinear interferometer or UNI. An input signal pulse enters the device and is split, via a polarization sensitive delay element into orthogonal polarizations separated approximately by a pulsewidth. These two signal components travel through a nonlinear medium in which a control pulse, coupled in via a 50/50 splitter, is temporally coincident with one of the signal pulse components. The two signal components are retimed to overlap in a second polarization sensitive delay and are subsequently interfered using a polarizer set at 45° with respect to the orthogonal signal polarizations. The control pulse is filtered out at the output of the device.

One advantage of this geometry is interferometric stability since all signals travel along the same path and are thus exposed to identical optical length variations, if any. The more important characteristic of the SAI from the perspective of optical switching, however, is immunity to long lived refractive index changes in the nonlinear medium. More specifically, since both signal components traverse the nonlinear medium, they each accrue phase shifts that are due to effects that persist for time scales longer than twice the temporal pulsewidth. However, the signal component that overlaps the control experiences, in addition to the slowly recovering phase changes, the ultrafast changes of the refractive index nonlinearity induced by the control pulse. Thus, ultrafast differential phase modulation between orthogonal signal components can be achieved with the SAI. Other balanced interferometric implementations include the Mach–Zehnder geometry with identical nonlinear media placed in each arm such that the long-lived nonlinear refractive index responses are canceled at the output [98]. Also, the NOLM may be modified to include an SOA, placed off center from the loop, as the nonlinear element [96], [97].

We have used the UNI to demonstrate a 40 Gb/s circulating shift register (CSR) with an inverter [110]. In this case, the UNI was configured to operate with a counterpropagating control stream as shown in Fig. 24. This experimental setup has the logical equivalent shown in Fig. 22(b). Initially, there are no control pulses in the feedback loop and the signal (clock) pulses are transmitted by the UNI. These output pulses are amplified and coupled to the control port of the UNI. When these control pulses arrive at the SOA, they induce gain and refractive index nonlinearities in the SOA and change the state of the UNI from transmitting the signal (clock) pulses to extinguishing the signal (clock) pulses. Data within the inverting CSR is monitored at the output of the 90/10 coupler as shown in Fig. 24. The temporal output of the inverting CSR consists of alternating blocks of "zeros" and "ones" where the lengths of the data blocks correspond to the transit time

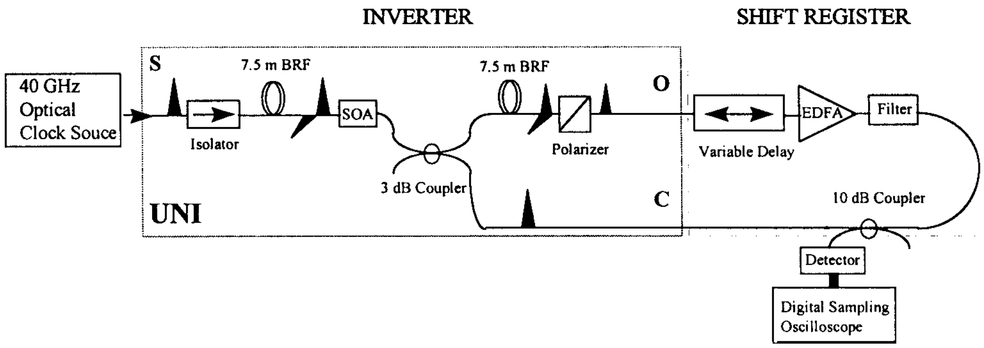

Fig. 24. Experimental set-up of the circulating shift register with an inverter. The UNI is configured as an inverter and the feedback loop constitutes the feedback shift register.

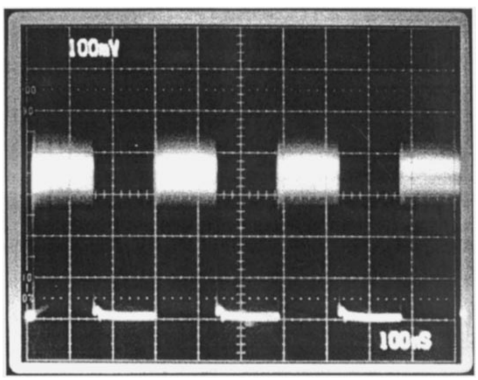

Fig. 25. Temporal output from the inverting CSR. The data are detected and displayed on a 1-GHz 3-dB bandwidth oscilloscope.

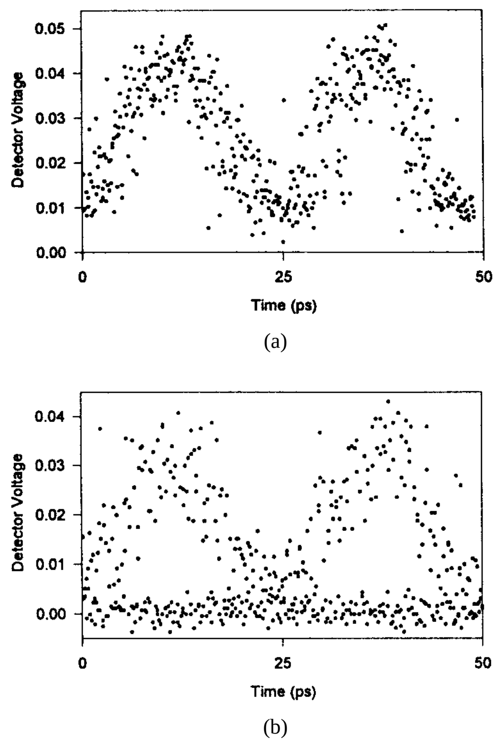

Fig. 26. (a) Oscilloscope trace of the 40-GHz clock stream used to drive the inverting CSR and (b) the eye diagram of the data stream output from the inverting CSR. The data are displayed on a 45-GHz 3-dB bandwidth oscilloscope.

through the feedback loop. In our case, the feedback loop length is approximately 28 m (in fiber) corresponding to a transit time of 142 ns.

Fig. 25 shows the output of the shift register on a microsecond time scale. The data are detected by a high-speed photodiode and displayed on a 1 GHz analog oscilloscope. Individual pulses in the data stream are not resolved. These data clearly show the alternating blocks or packets of “zeros” and “ones.” The CSR was very stable and these alternating packets could be maintained for hours at a time, corresponding to tens of billions of circulations, with no temperature or length stabilization applied to the feedback loop. Fig. 26(a) shows digital sampling oscilloscope traces of the 40 GHz input clock stream on a 50 ps time scale. Note that the jitter of the displayed pulses is primarily due to timing jitter in the oscilloscope time base (6 ps) and not to timing jitter on the pulses themselves. The timing jitter of the soliton compression source is less than 0.5 ps [110]. Fig. 26(b) shows an eye-diagram of the data output from the inverting CSR. Again, the jitter in the data is due to limitations in the detection electronics.

d) Rate converters: Data stored in the recirculating buffers will be rate converted down to lower rate signals that the user node electronics can process. One method for performing rate conversion is to demultiplex every Kth bit of an N-bit long data packet such that N and K are relatively prime. The time required to rate convert all of the data bits within a packet using this scheme is NKT_b where T_b is the bit period of the original data stream. Unfortunately, such schemes may scramble the data bits and, consequently, require further processing in electronics in order to faithfully reconstruct the original packet. An alternative method for rate conversion is based on time dilation via optical sampling. This method has the advantage that the rate-converted stream requires no reordering of bits. The essence of this technique is that the aliasing brought about by sampling a high speed optical stream at a rate lower than the Nyquist rate gives rise to a time dilated version of the packet which can then be detected and displayed on a low speed analog oscilloscope. The sampling can be achieved via optical cross-correlation in a nonlinear crystal or an optical switch, for example. The

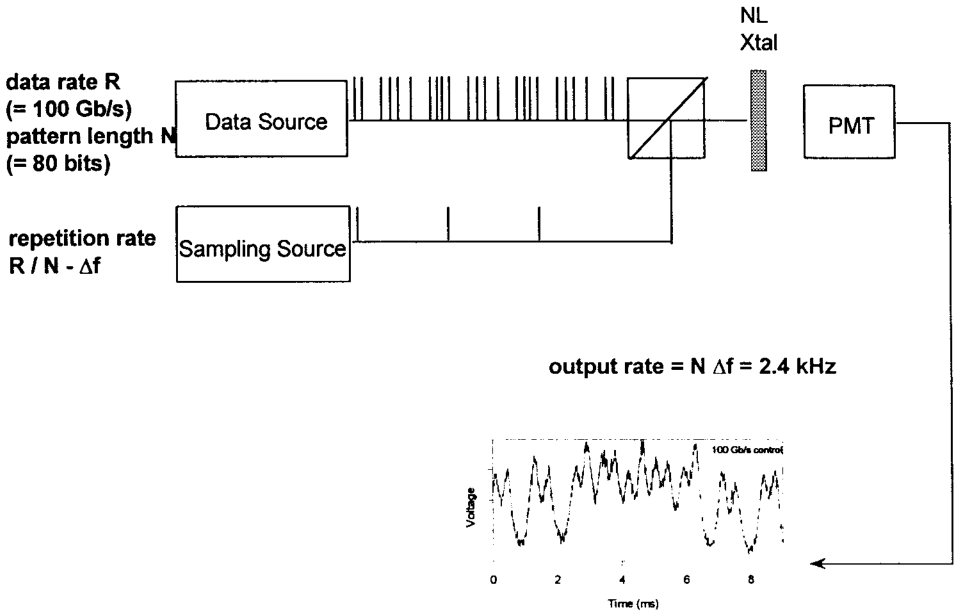

Fig. 27. Block diagram of the rate converter setup.

cross-correlated data are detected and low-pass filtered so that the cross-correlation signal envelope and not the individual sampling pulse envelopes are measured. It is important to note that both rate conversion techniques, demultiplexing of all the data bits within a packet, and time dilation, require a repeating input optical data pattern. Therefore, in a real system, these schemes must be employed in conjunction with the recirculating or regenerative loop buffers.

We demonstrate the operation of a time dilating rate converter using a passive multiplexer to assemble a 100 Gb/s optical data stream and a low-speed photomultiplier tube and oscilloscope to display the data stream. Fig. 27 shows a block diagram of the experimental set-up. The 100 Gb/s data stream is generated by passively multiplexing a lower rate (12.5 Gb/s) 10-bit data stream. A short, repetitive, data pattern is chosen to facilitate the measurement, via rate-conversion, of the high-speed data stream. The 12.5 Gb/s data stream is split and recombined three times to generate an 80-bit, 100 Gb/s data stream. The sampling source is a gain-switched DFB laser, driven at 1.249999970 GHz, offset 30 Hz from the 1.25 GHz pattern repetition rate. The output pulses are linearly compressed in a short length of dispersion compensating fiber and then further compressed to 3-ps utilizing soliton effects in a 10-km spool of standard fiber. Cross-correlation signals between the switched out data pulses and the sampling pulses are generated via type-II phase matching in a KTP crystal, are detected by a photomultiplier tube, and are displayed on an oscilloscope [117]. The inset in Fig. 27 shows the control data stream as it appears at the output from the rate converter. Notice that even at 100 Gb/s, the rate-converted data pulses are clearly resolved. Also note that the rate converted pulse spacing is approximately 0.4 ms, consistent with the 30 Hz offset between the sampling source repetition rate and the 80-bit, 100 GHz data stream [118]. The nonuniformity in pulse amplitude in the multiplexed stream is caused by variations in the splitting ratios of the couplers in the passive multiplexer.

C. Network Node Operation

With the technologies described above, we believe a high-speed slotted TDM network demonstration will be feasible in the next few years. To highlight how these technologies will be incorporated into the network node; we will give a brief description of the transmitter and receiver nodes for such an ultrafast, multiaccess network.

Fig. 13 shows a number of passive taps between the optical bus and the user transmitter and receiver nodes. In the receiver node, the data packets enter via a passive tap, at the slot rate from 10 to 100 Mslots/s. The first function that must be performed is address/header recognition. In the transmitter nodes, header recognition includes determining if there is a credit marker present and if the slot is empty or full. In the receiver node, the incoming address must be compared with a local address to determine whether or not the packet requires further processing. Address recognition must be performed on-the-fly, at 100 Gb/s. Ultrahigh-speed optical switches, configured as AND gates and XOR gates can be used for this purpose. Synchronization may be achieved by using self-synchronizing codes or by using all-optical switches as phase comparators in electrooptic phase locked loops.

If the packet address matches the local address, the data packet must be stored while it is rate converted down to a rate that the node electronics can process. In the transmitter node, lower rate optical data will be assembled (up-converted) in fiber loop buffers and stored while the user waits for access to the network. In the receiver node, after the data have been loaded successfully into the buffer, the data will be rate converted and detected for electronic processing.

D. All-Optical Switching

Recently, world-record switching speed in a semiconductor based UNI was demonstrated using the components described in the previous sections. Bitwise logic at 100 Gb/s was reported

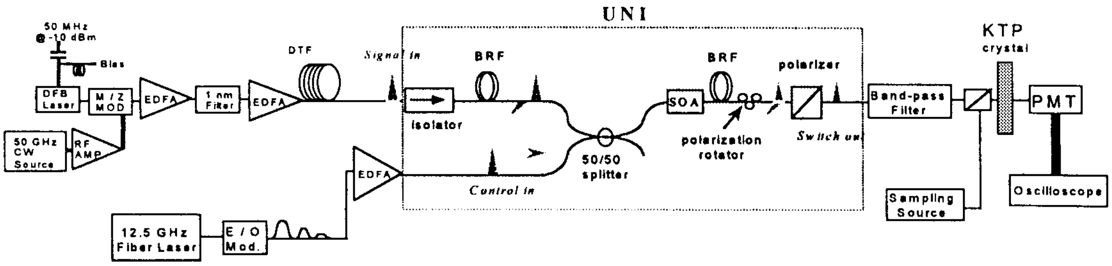

Fig. 28. Block diagram of the experimental setup used to demonstrate 100 Gb/s bitwise logic in semiconductor based all-optical switch.

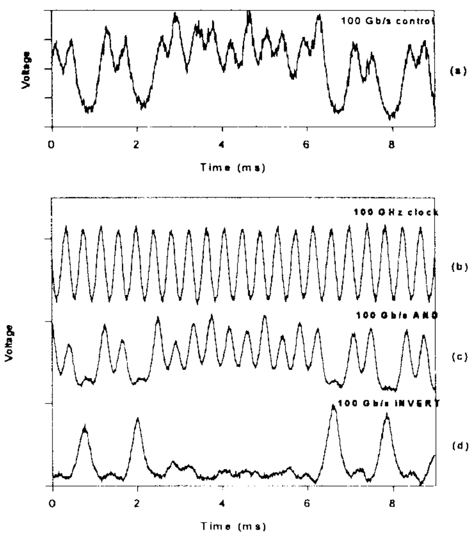

Fig. 29. A portion of the rate converted (a) input control data, (b) input clock stream, (c) switched-out AND signal, and (d) switched-out NOT signal demonstrating 100 Gb/s bitwise logic.

for the ultrafast nonlinear interferometer (UNI) configured as AND and NOT gate [65]. Here we distinguish "bitwise logic" from demultiplexing because every bit is switched. In demultiplexing experiments, bits are switched at the demultiplexed output rate, so many bits travel through the switch but are not acted on by the control stream. Such experiments may be less sensitive to SOA dynamics and average control power effects. The experimental configuration is shown in Fig. 28 along with the experimental results in Fig. 29. The previous record speed for bitwise logic was 40 Gb/s [102]. The 100 Gb/s AND operation was also reported for a fiber-based NOLM [66]. With these record all-optical switching results, new applications may be envisioned. For example, an ultrafast optical encrypter may be used to perform an XOR operation between a data stream and a cryptographic key stream. However, owing to limitations in all-optical switch performance, electronic cryptographic circuits cannot be duplicated using optical hardware. For example, creating a linear feedback shift register with optical logic, as proposed in [108], would permit the use of very few taps, making the register vulnerable to a sparse matrix (Berlekamp–Massey) attack (see, for instance, [119]).

One way to increase the rate at which pseudo-noise (PN) sequences are generated is by passively multiplexing lower rate optical data streams that have been encrypted using standard electronic schemes. The drawback of a passive multiplexing approach is that the probability of breaking any one of the lower rate sequences is increased owing to the fact that M sequences are observed in a given time t, where M is the number of multiplexed sequences. This problem can be partially circumvented by multiplexing sequences with relatively prime periods. In this case, the period of the multiplexed sequence grows exponentially in M. However, requiring that the periods be relatively prime restricts the choice of PN sequences, forcing some of them to be relatively short, and hence, vulnerable.

An alternative way to increase the rate for PN sequences is to use slower, electronic PN generators to dynamically reconfigure a limited number of optical taps on a high-speed optical shift register [120], [121]. Since there is a difference in rate between the electronics and the optics, the reconfiguration of the optics is done every d optical bits, where d, an integer, is the ratio of the optical bit rate to the electronic bit rate. This scheme generates high rate optical sequences with extremely high periods, on the order of the product of the electronic controller period, P, and 2^L, where L is the latency (circulation time), in bits, of the optical feedback shift register. In a real implementation, the latency of the feedback shift register might consist solely of the latency of the concatenated optical switches and may be augmented by additional lengths of fiber. At 100 Gb/s, the latency required to achieve high periods is on the order of a million bits. Moreover, it can be proved that for linear feedback the number of taps is unimportant as long as it greater than or equal to d. Therefore, we can derive the same benefits from having ten taps as if we had an arbitrarily large number of taps. This approach is different from reconfiguration or control that runs at the same speed as the logic that it controls. The main thrust behind using optical feedback shift registers controlled by electronics is to rely on the complexity afforded by the electronics for the long-range security of the sequence while relying upon the high-speed optics to provide the ultrafast short-term randomness. We believe this scheme could be demonstrated with all-optical gate counts as low as 20. An example of an implementation of this encryption scheme is shown in Fig. 30.

Another application for high-speed all-optical switches and ultrafast optical pulse sources is an electrooptic analog-to-digital (A/D) converter. An important issue in very fast A/D conversion is the timing jitter of the sampling pulses. Using optical pulses, such as those generated by mode-locked lasers or soliton compression sources, may offer sampling pulse timing stability much better than can be obtained by electronic [123]. Feasibility demonstrations have used such optical sampling

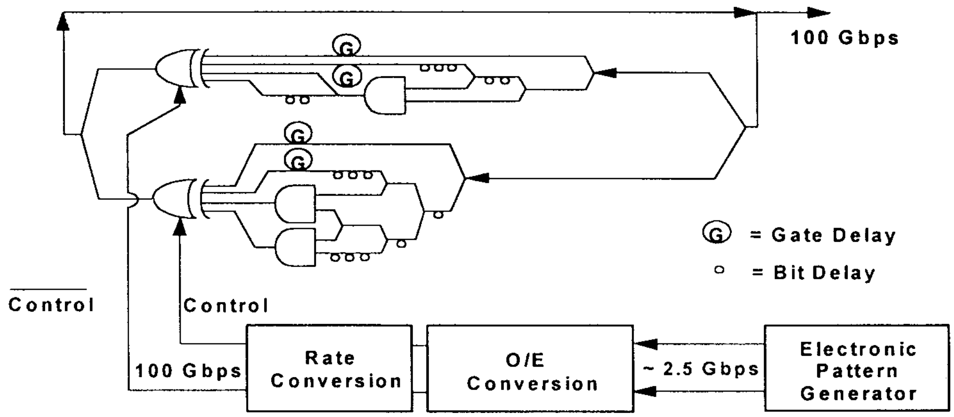

Fig. 30. Example of an all-optical encryption circuit.

pulses in electrooptic modulators, with the RF signal applied to the electrical modulation port. The modulation amplitude was sensed and digitized by subsequent optical detection of the two quadrature signals, followed by digital signal processing. Other potential detection and quantization techniques might use more direct phase (proportional to the electrical modulation) sensing techniques, such as interferometric measurement of the modulator optical output. In principle, such A/D conversion could be accomplished as fast as the phase-sensitive wavefront can be measured, which could be an order of magnitude faster than pure electronic methods. The major challenge would seem to be providing sufficient dynamic range (or bits of resolution).

In addition to the all-optical signal processors described above, ultrahigh-speed optical techniques may offer advantages over traditional electrical techniques in the generation of high bit rate, high-fractional bandwidth RF signals. For example, low timing jitter 100 GHz optical soliton clock streams and all-optical switches may be used to generate phase shift keyed (PSK) RF signals. In this application, optical switches can be used to switch between two phase-shifted optical clock streams or optical switches may be used to rotate the polarization of bits in an optical stream. Polarization rotation can be converted to phase modulation in a polarization sensitive delay. In a proof of concept demonstration, we generated a 2 Gb/s (PSK) signal on a 20 GHz carrier [123]. This generation technique could scale to data rates exceeding 10 Gb/s and fractional bandwidths well above 25% using existing sources and switches.

IV. Summary

In summary, we have described WDM and TDM optical data networks. Due to commercial availability of necessary components, and the ability to meet near-term bandwidth requirements, we believe that WDM optical networks will be deployed in the near future. By smart design of the protocol stack and the physical layer architecture we believe orders of magnitude increase in data rate with a corresponding drastic reduction in cost will be realized. Specifically, if the IP layer is modified to take into account directly the WDM layer underneath and eliminate the ATM and SONET layers, the performance of the resulting WDM optical data network will far exceed that of current networks.

In the more distant future, as high-speed optical processing components mature, we believe TDM networks will be the obvious choice to provide packet service in networks capable of bursting at very high data rates. The algorithms needed to regulate access to these ultrafast optical networks will differ from their electronic counterparts in that they will operate in a high latency environment and they will have to provide both fair and efficient bandwidth utilization using a minimal amount of processing at the shared media rate. Components necessary to implement such algorithms include high-speed optical sources, switches, buffers, and rate converters. State-of-the-art laboratory demonstrations suggest that these components can operate at rates exceeding 100 Gb/s, but commercial development is necessary for widespread deployment of these types of networks.

Acknowledgment

The authors are grateful for their colleagues who have made significant contributions to this project. They thank B. S. Robinson, J. D. Moores, S. G. Finn, A. L. Chiu, S. R. Chinn, E. A. Swanson, M. Medard, A. H. Chan, N. S. Patel, and R. A. Barry from the Massachusetts Institute of Technology (M.I.T.), Lincoln Laboratory, Lexington. On the M.I.T., Cambridge, campus they thank, Prof. R. S. Gallager, Prof. E. P. Ippen, and Prof. H. A. Haus, as well as graduate students D. J. Jones, C. R. Doerr, and W. S. Wong.

References

[1] B. Mukherjee, "WDM-based local lightwave networks, Part I: Single-hop systems," *IEEE Network,* May 1992.
[2] G. N. M. Sudhakar, M. Kavehrad, and N. D. Georganas, "Access protocols for passive optical star networks," *Comput. Networks ISDN Syst.,* pp. 913–930, 1994.
[3] N. Mehravari, "Performance and protocol improvements for very high speed optical fiber local area networks using a passive star topology," *J. Lightwave Technol.,* Apr. 1990.
[4] M. S. Chen, N. R. Dono, and R. Ramaswami, "A new media access protocol for packet switched wavelength division multiaccess metropolitan network," *IEEE J. Select. Areas Commun.,* Aug. 1990.
[5] H. B. Jeon and C. K. Un, "Contention based reservation protocols in multiwavelength protocols with passive star topology," in *Proc. ICC'92* Chicago, IL, June 1992.
[6] I. Chlamtac and A. Ganz, "Channel allocation protocols in frequency-time controlled high-speed networks," *IEEE Trans. Commun.,* Apr. 1988.

[7] F. Jia and B. Mukherejee, "The receiver collision avoidance (RCA) protocol for single hop WDM lightwave networks," in *Proc. ICC'92,* Chicago, IL, June 1992.

[8] I. M. I. Habib, M. Kavehrad, and C.-E. W. Sundberg, "Protocols for very high-speed optical fiber local area networks using a passive star topology," *J. Lightwave Technol.,* Dec. 1987.

[9] G. N. M. Sudhakar, N. D. Georganas, and M. Kavehrad, "A multichannel optical star LAN and its application as a broadband switch," in *Proc. ICC'92,* Chicago, IL, June 1992.

[10] P. Dowd, "Random access protocols for high speed interprocess communications based on a passive optical star topology," *J. Lightwave Technol.,* vol. 9, June 1991.

[11] I. P. Kaminow *et al.,* "A wideband all-optical WDM network," *IEEE J. Select. Areas Commun.,* June 1996.

[12] E. Modiano, "Unscheduled multicasts in WDM broadcast-and-select networks," in *Proc. INFOCOM '98,* San Francisco, CA, Mar. 1998.

[13] M. J. Karol, M. G. Hluchyj, and S. P. Morgan, "Input versus output queueing in a space-division packet switch," *IEEE Trans. Commun.,* Dec. 1987.

[14] M. G. Hluchyj and M. J. Karol, "Queueing in high-performance packet switching," *IEEE J. Select. Areas Commun.,* Dec. 1988.

[15] K. M. Sivalingam and J. Wang, "Media access protocols for WDM networks with on-line scheduling," *J. Lightwave Technol.,* June 1996.

[16] N. McKeown, V. Anantharam, and J. Walrand, "Achieving 100% throughput in an input queued switch," in *Proc. INFOCOM'96,* San Francisco, CA, Apr. 1996.

[17] J. E. Hopcroft and R. M. Karp, "An $n^{5/2}$ algorithm for finding maximal matching in bipartite graphs," *Soc. Industr. Appl. Math. J. Comput.,* Feb. 1973.

[18] E. Modiano, R. Barry, and E. Swanson, "A novel architecture and medium access control protocol for WDM networks," *OSA Trends in Optics and Photon. (TOPS),* vol. 20, 1998.

[19] J. Bannister, L. Fratta, and M. Gerla, "Topological design of WDM networks," in *Proc. INFOCOM'90.*

[20] L. S. Brakmo and L. Peterson, "TCP vegas: End-to-end congestion avoidance on a global internet," *IEEE J. Select. Areas Commun.,* Oct. 1995.

[21] B. Schein and E. Modiano, "Increasing traffic capacity in WDM ring networks via topology reconfiguration," *OFC'99,* to be published.

[22] G. Ellinas and T. E. Stern, "Automatic protection switching for link failures in optical networks with bidirectional links," in *Proc. GLOBECOM'96.*

[23] S. G. Finn, M. Medard, and R. A. Barry, "A novel approach to automatic protection switching using trees," in *Proc. ICC'97.*

[24] ——, "A new algorithm for bidirectional link self healing for arbitrary redundant networks," in *Proc. OFC'98.*

[25] D. W. Faulkner *et al.,* "Optical networks for local loop applications," *J. Lightwave Technol.,* Nov. 1989.

[26] N. J. Frigo *et al.,* "A wavelength division multiplexed passive optical network with cost shared components," *IEEE Photon. Technol. Lett.,* June 1994.

[27] N. J. Frigo, "Local access optical networks," *IEEE Network,* Dec. 1996.

[28] B. Glance *et al.,* "Novel optically restorable WDM ring network," in *Proc. OFC'96.*

[29] O. Gerstel and R. Ramaswami, "Multiwavelength ring architectures and protection schemes," in *Photonic Networks, Advances in Optical Communications.* New York: Springer-Verlag, 1997, pp. 42–51.

[30] P. P. Iannone, N. J. Frigo, and T. E. Darcie, "WDM passive optical network architecture with bidirectional optical spectral slicing," in *Proc. OFC'95.*

[31] V. Jacobson, R. Braden, and D. Dorman, "TCP extensions for high performance," Internet Engineering Task Force RFC-1323, May 1992.

[32] I. P. Kaminow *et al.,* "A wideband all-optical WDM network," *IEEE J. Select. Areas Commun.,* June 1996.

[33] T. V. Lakshman and U. Madhow, "The performance of TCP/IP for networks with high bandwidth-delay products and random loss," *IEEE/ACM Trans. Networking,* 1997.

[34] Y. M. Lin and D. R. Spears, "Passive optical subscriber loops with multiaccess," *J. Lightwave Technol.,* Nov. 1989.

[35] X. Lu *et al.,* "Minifiber node hybrid fiber/coax networks for two way broadband access," in *Proc. OFC'96.*

[36] A. Luvison, "The architecture of full service access network," in *Proc. OFC'98,* San Jose, CA, Mar. 1998.

[37] J. Manchester and P. Bonenfant, "Fiber optic network survivability: SONET/optical protection layer interworking," in *Proc. Nat. Fiber Optic Eng. Conf.,* 1996.

[38] E. Modiano, R. Barry, and E. Swanson, "A novel architecture and medium access control (MAC) protocol for WDM networks," in *Proc. OFC'98,* San Jose, CA, Feb. 1998.

[39] ——, "A novel architecture and medium access control protocol for WDM networks," *TOP's,* vol. 20, (Optical Networks and their Applications), Optical Society of America, 1998.

[40] E. Modiano and A. Chiu, "Traffic grooming algorithms for minimizing electronic costs in unidirectional SONET/WDM ring networks," in *Proc. CISS'98,* Princeton, NJ, Mar. 1998.

[41] A. L. Chiu and E. H. Modiano, "Reducing electronic multiplexing costs in unidirectional SONET/WDM ring networks via efficient traffic grooming," in *Proc. IEEE GLOBECOM'98,* Sydney, Australia, Nov. 1998.

[42] E. Modiano, "Unscheduled multicasts in WDM broadcast-and-select networks," in *Proc. IEEE INFOCOM'98,* San Francisco, CA, Mar. 1998.

[43] B. Mukherjee, "WDM-based local lightwave networks, Part I," *IEEE Network,* May 1992.

[44] K. Nagami *et al.,* "Toshiba's flow attribute notification protocol spec.," in *IETF RFC 2129,* Apr. 1997.

[45] P. Narvaez and K.-Y. Siu, "An acknowledgment bucket scheme to regulate TCP flow over ATM," in *Proc. IEEE GLOBECOM,* Nov. 1997.

[46] P. Newman, T. Lyon, and G. Minshall, "Flow labeled IP: A connection approach to ATM," in *Proc. IEEE INFOCOM,* 1996.

[47] W. Puh and G. Boyer, "Broadband access: Comparing alternatives," *IEEE Commun. Mag.,* Aug. 1995.

[48] Y. Rekhter *et al.,* "Cisco systems' tag switching architecture overview," in *IETF RFC 2105,* Feb. 1997.

[49] R. E. Wagner, R. C. Alferness, A. A. M. Saleh, and M. S. Goodman, "MONET: Multiwavelength optical networking," *J. Lightwave Technol.,* June 1996.

[50] M. N. Ransom and D. R. Spears, "Applications of public giga-bit networks," *IEEE Network,* Mar. 1992.

[51] R. Ramaswami and K. Sivarajan, "Design of logical topologies for wavelength routed optical networks," *IEEE J. Select. Areas Commun.,* June 1996.

[52] ——, *Optical Networks.* San Francisco, CA: Morgan Kaufmann, 1998.

[53] S. Subramanian and R. Barry, "Wavelength assignment in fixed routing WDM networks," in *Proc. ICC '97,* June 1997.

[54] J. M. Simmons, E. L. Goldstein, and A. A. M. Saleh, "On the value of wavelength-add/drop in WDM rings with uniform traffic," in *Proc. OFC'98,* San Jose, CA, Feb. 1998.

[55] G. N. M. Sudhakar, M. Kavehrad, and N. D. Georganas, "Access protocols for passive optical star networks," *Comput. Networks ISDN Syst.,* pp. 913–930, 1994.

[56] A. Viswanathan *et al.,* "ARIS: Aggregate route-based IP switching," IETF Internet Draft, Mar. 1997.

[57] S. S. Wagner and H. L. Lemberg, "Technology and system issues for a WDM-based fiber look architecture," *J. Lightwave Technol.,* Nov. 1989.

[58] T.-H. Wu, *Fiber Network Service Survivability.* Norwood, MA: Artech House, 1992.

[59] T.-H. Wu and W. I. Way, "A novel passive protected SONET bidirectional self-healing ring architecture," *J. Lightwave Technol.,* vol. 10, Sept. 1992.

[60] S. G. Finn and R. A. Barry, "Optical services in future broadband networks," *IEEE Network Mag.,* p. 7, Nov./Dec. 1996.

[61] C. Yeh, M. Lin, M. Gerla, and P. Rodrigues, "RATO-net: A random-access protocol for unidirectional ultrahigh-speed optical fiber networks," *J. Lightwave Technol.,* vol. 8, Jan. 1990.

[62] G. C. Watson and S. Tohme, "S++—A new MAC protocol for Gb/s local area networks," *IEEE J. Select. Areas Commun.,* vol. 11, May 1993.

[63] S. G. Finn, "HLAN-an architecture for optical multiaccess networks," in *Tech. Dig. LEOS Summer Top. Mtg.,* 1995, vol. 45.

[64] R. A. Barry, V. W. S. Chan, K. L. Hall, E. S. Kintzer, J. D. Moores, K. A. Rauschenbach, E. A. Swanson, L. E. Adams, C. R. Doerr, S. G. Finn, H. A. Haus, E. P. Ippen, W. S. Wong, and M. Haner, "All-optical network consortium-ultrafast TDM networks," *J. Lightwave Technol./IEEE J. Select. Areas Commun.,* vol. 14, no. 5, p. 999, 1996.

[65] K. L. Hall and K. A. Rauschenbach, "100 Gb/s bitwise logic," *Opt. Lett.,* vol. 23, no. 16, p. 1271–1273, 1998.

[66] H. Takara, I. Shake, K. Uchiyama, O. Kamatami, S. Kawanishi, and K. Sato, "Ultrahigh-speed optical TDM signal generator utilizing all-optical modulation and optical clock multiplication," in *Tech. Dig. Optic. Fiber Commun. Conf.,* 1998, vol. 2, PD16-1-PD16-4.

[67] T. F. Carruthers and I. N. Duling, "10-GHz, 1.3-ps erbium fiber laser employing soliton pulse shortening," *Opt. Lett.,* vol. 21, no. 23, p. 1927, 1996.

[68] M. Nakazawa, E. Yoshida, and K. Tamura, "10 GHz, 2 ps regeneratively and harmonically FM mode-locked erbium fiber ring laser," *Electron. Lett.,* vol. 32, no. 14, p. 1285, 1996.

[69] S. V. Chernikov, M. J. Guy, J. R. Taylor, D. G. Moodie, and R. Kashyap, "Duration-tunable 0.2–2.0-ps 10-GHz source of transform

limited optical pulses based on an electroabsorption modulator," *Opt. Lett.*, vol. 20, no. 23, p. 2399, 1995.
[70] R. Ludwig, S. Diez, A. Ehrhardt, L. Kuller, W. Pieper, and H. G. Weber, "A tunable femtosecond modelocked semiconductor laser for applications in OTDM systems," *IEICE Trans. Electron.*, vol. E81-C, no. 2, p. 140, 1998.
[71] E. A. Swanson and S. R. Chinn, "23-GHz and 123-GHz soliton pulse generation using two CW lasers and standard single-mode fiber," *IEEE Photon. Technol. Lett.*, vol. 6, p. 796, July 1994.
[72] R. Kashyap, S. V. Chernikov, and J. R. Taylor, "Fiber laser and beat frequency sources based on fiber gratings for microwave and ultrafast processing," *Int. J. Optoelectron.*, vol. 11, no. 2, p. 87, 1997.
[73] J. D. Moores, W. S. Wong, and K. L. Hall, "50 Gb/s optical pulse storage ring using novel rational-harmonic modulation," *Opt. Lett.*, vol. 20, no. 24, p. 2547, 1995.
[74] E. A. Swanson, S. R. Chinn, K. L. Hall, K. A. Rauschenbach, R. S. Bondurant, and J. W. Miller, "100-GHz soliton pulse train generation using soliton compression of two phase side bands from a single DFB laser," *IEEE Photon. Technol. Lett.*, vol. 6, p. 1194, Oct. 1994.
[75] E. A. Swanson and S. R. Chinn, "40-GHz pulse train generation using soliton compression of a Mach–Zehnder modulator output," *IEEE Photon. Technol. Lett.*, vol. 7, p. 114, Jan. 1995.
[76] F. Futami and K. Kikuchi, "Generation of 113-GHz, 1.8-ps optical pulse trains by Fourier synthesis of four-wave-mixing signals obtained from semiconductor optical amplifiers," *Opt. Lett.*, vol. 22, no. 24, p. 1873, 1997.
[77] N. Onodera, A. J. Lowry, L. Zhai, Z. Ahmed, and R. S. Tucker, "Frequency multiplication in actively modelocked semiconductor lasers," *Appl. Phys. Lett.*, vol. 62, no. 12, p. 1329, 1993.
[78] M. Nakazawa, E. Yoshida, T. Yamamoto, E. Yamada, and A. Sahara, "TDM single channel 640 Gb/s transmission experiment over 600 km using a 400 fs pulse train and a walk-off-free, dispersion flattened nonlinear optical loop mirror," in *Tech. Dig. Optic. Fiber Commun. Conf.*, 1998, vol. 2, PD14-1–PD14-4.
[79] S. Kawanishi, H. Takara, T. Morioka, O. Kamatani, and M. Saruwatari, "200 Gb/s, 100 km time division multiplexed optical transmission using supercontinuum pulses with prescaled PLL timing extraction and all-optical demultiplexing," *Electron. Lett.*, vol. 31, no. 10, p. 816, 1995.
[80] S. Kawanishi, H. Takara, K. Uchiyama, I. Shake, O. Kamatami, and H. Takahashi, "1.4 Tb/s (200 Gb/s $\times$ 7 ch) 50 km optical transmission experiment," *Electron. Lett.*, vol. 33, no. 20, p. 1716, 1997.
[81] S. Kumar, M. Wald, F. Lederer, and A. Hasegawa, "Soliton interaction in strongly dispersion-managed optical fibers," *Opt. Lett.*, vol. 23, no. 13, p. 1019, 1998.
[82] Y. Chen and H. A. Haus, "Dispersion-managed solitons with net positive dispersion," *Opt. Lett.*, vol. 23, no. 12, p. 1013, 1998.
[83] V. S. Grigoryan and C. R. Menyuk, "Dispersion-managed solitons at normal average dispersion," *Opt. Lett.*, vol. 23, no. 8, p. 609, 1998.
[84] J. F. L. Devaney, W. Forysiak, A. M. Niculae, and N. J. Doran, "Soliton collisions in dispersion-managed wavelength-division-multiplexed systems," *Opt. Lett.*, vol. 22, no. 22, p. 1695, 1997.
[85] J. N. Kutz and S. G. Evangelides, "Dispersion-managed breathers with average normal dispersion," *Opt. Lett.*, vol. 23, no. 9, p. 685, 1998.
[86] M. Jinno and M. Abe, "All-optical regenerator based on nonlinear fiber Sagnac interferometer," *Electron. Lett.*, vol. 28, no. 14, p. 1350, 1992.
[87] J. K. Lucek and K. Smith, "All-optical signal regenerator," *Opt. Lett.*, vol. 18, no. 15, p. 226, 1993.
[88] W. Pieper, K. Weich, R. Ludwig, E. Patzak, and H. G. Weber, "All-optical polarization and wavelength independent 3R signal regenerator," *Electron. Lett.*, vol. 32, no. 14, p. 1316, 1996.
[89] M. Shabeer, J. K. Lucek, K. Smith, D. Cotter, and D. C. Rogers, "Self-synchronization scheme for high-speed photonic networks," *Electron. Lett.*, vol. 31, no. 17, p. 1476, 1995.
[90] Y. H. Kao, Y. Liang, T. J. Xia, K. H. Ahn, J. W. Lou, G. A. Nowak, O. Boyraz, and M. N. Islam, "Self-synchronization of 100 Gb/s TDM packets using a semiconductor laser amplifier and an intensity discriminator," in *Tech. Dig. Conf. Lasers Electro-Optics*, 1998, p. 503.
[91] D. Cotter, J. K. Lucek, M. Shabeer, K. Smith, D. C. Rogers, D. Nesset, and P. Gunning, *Electron. Lett.*, vol. 31, p. 1475, 1995.
[92] S. Kawanishi and O. Kamatani, "All-optical time division multiplexing using four wave mixing," *Electron. Lett.*, vol. 30, no. 20, p. 1697, 1994.
[93] N. J. Doran and D. Wood, "Nonlinear optical loop mirror," *Opt. Lett.*, vol. 13, no. 1, p. 56, 1988.
[94] K. L. Hall, K. A. Rauschenbach, E. A. Swanson, S. R. Chinn, and G. Raybon, "Picosecond-accuracy all-optical bit phase comparison using a nonlinear optical loop mirror," *IEEE Photon. Technol. Lett.*, vol. 7, p. 935, 1995.
[95] K. A. Rauschenbach, K. L. Hall, J. C. Livas, and G. Raybon, "All-optical pulsewidth and wavelength conversion at 10 Gb/s using a nonlinear optical loop mirror," *IEEE Photon. Technol. Lett.*, vol. 6, p. 1130, 1994.
[96] M. Eiselt, W. Pieper, and H. G. Weber, "All-optical high speed demultiplexing with a semiconductor laser amplifier in a loop mirror configuration," *Electron. Lett.*, vol. 29, no. 13, p. 1167, 1993.
[97] J. P. Sokoloff, P. R. Prucnal, I. Glesk, and M. Kane, "A terahertz optical asymmetric demultiplexer (TOAD)," *IEEE Photon. Technol. Lett.*, vol. 5, p. 787, 1993.
[98] K. I. Kang, T. G. Chang, I. Glesk, and P. R. Prucnal, "Demonstration of ultrafast, all-optical low control energy, single wavelength polarization independent cascadable and integratable switch," *Appl. Phys. Lett.*, vol. 67, p. 605, 1995.
[99] S. Nakamura and K. Tajima, "Analysis of subpicosecond full-switching with a symmetric Mach–Zehnder all-optical switch," *Japan J. Appl. Phys. 2*, vol. 35, no. 11, p. L1426, 1996.
[100] N. S. Patel, K. L. Hall, and K. A. Rauschenbach, "40 Gb/s cascadable all-optical logic with an ultrafast nonlinear interferometer," *Opt. Lett.*, vol. 21, no. 18, p. 1466, 1996.
[101] N. S. Patel, K. A. Rauschenbach, and K. L. Hall, "40 Gb/s demultiplexing using an ultrafast nonlinear interferometer," *IEEE Photon. Technol. Lett.*, vol. 8, p. 1695, Dec. 1996.
[102] N. S. Patel, K. L. Hall, and K. A. Rauschenbach, "Interferometric all-optical switches for ultrafast signal processing," *Appl. Opt.*, vol. 37, no. 14, pp. 2831–2842, 1998.
[103] K. L. Hall, G. Lenz, A. M. Darwish, and E. P. Ippen, "Subpicosecond gain and index nonlinearities in InGaAsP diode lasers," *Opt. Commun.*, vol. 111, pp. 589–612, 1994.
[104] D. Cotter and S. C. Cotter, "Algorithm for binary word recognition suited to ultrafast nonlinear optics," *Electron. Lett.*, vol. 29, no. 11, p. 945, 1993.
[105] K. L. Hall and K. A. Rauschenbach, "All-optical bit pattern generation and matching," *Electron. Lett.*, vol. 32, no. 13, p. 1214, 1996.
[106] K. L. Hall, "All-optical buffers for high-speed slotted TDM networks," in *Proc. IEEE/LEOS Summer Topic. Mtg. Tech. Dig., Advanced Semiconductor Lasers and Applications*, 1997, pp. 15–16.
[107] H. Avamopoulos and N. Whitaker, "Addressable fiber-loop buffer," *Opt. Lett.*, vol. 18, no. 1, p. 22, 1993.
[108] N. A. Whitaker, M. C. Gabriel, H. Avrampoulos, and A. Huang, "All-optical, all fiber circulating shift register with an inverter," *Opt. Lett.*, vol. 16, no. 24, p. 1999, 1991.
[109] A. J. Poustie, R. J. Manning, and K. J. Blow, "All-optical circulating shift register using a semiconductor optical amplifier in a fiber loop mirror," *Electron. Lett.*, vol. 32, no. 13, p. 1215, 1996.
[110] K. L. Hall, J. P. Donnelly, S. H. Groves, C. I. Fennelly, R. J. Bailey, and A. Napoleone, "40 Gbit/s all-optical circulating shift register with an inverter," *Opt. Lett.*, vol. 22, no. 19, pp. 1479–1481, 1997.
[111] C. R. Doerr, W. S. Wong, H. A. Haus, and E. P. Ippen, "Additive pulse mode-locking/limiting storage ring," *Opt. Lett.*, vol. 19, no. 21, p. 1747, 1994.
[112] J. D. Moores, K. L. Hall, S. M. LePage, K. A. Rauschenbach, W. S. Wong, H. A. Haus, and E. P. Ippen, "20-GHz optical storage loop/laser using amplitude modulation, filtering and fast saturable absorption," *IEEE Photon. Technol. Lett.*, vol. 7, p. 1096, Sept. 1995.
[113] K. L Hall, J. D. Moores, K. A. Rauschenbach, W. S. Wong, E. P. Ippen, and H. A. Haus, "All-optical storage of a 1.25 kbit packet at 10 Gb/s," *IEEE Photon. Technol. Lett.*, vol. 7, p. 1093, Sept. 1995.
[114] D. J. Jones, K. L. Hall, H. A. Haus, and E. P. Ippen, "Asynchronous phase-modulated optical fiber-ring buffer," *Opt. Lett.*, vol. 23, no. 3, pp. 177–179, 1998.
[115] R. J. Manning, D. A. O. Davies, D. Cotter, and J. K. Lucek, "Enhanced recovery rates in semiconductor laser amplifiers using optical pumping," *Electron. Lett.*, vol. 30, no. 10, pp. 787–788, 1994.
[116] J. M. Wiesenfeld, "Gain dynamics and associated nonlinearities in semiconductor optical amplifiers," *Int. J. High Speed Electron.*, vol. 7, no. 1, pp. 179–222, 1996.
[117] H. Takara, S. Kawanishi, T. Morioka, K. Mori, and M. Saruwatari, "100 Gb/s optical waveform measurement with 0.6 ps resolution optical sampling using subpicosecond supercontinuum pulses," *Electron. Lett.*, vol. 30, no. 14, pp. 1152–1153, 1994.
[118] N. S. Patel, K. L. Hall, and K. A. Rauschenbach, "Optical rate conversion for high-speed TDM networks," *IEEE Photon. Technol. Lett.*, vol. 9, pp. 1277–1279, 1997.
[119] H. C. A. van Tilborg, *An Introduction to Cryptography.* New York: Kluwer Academic, 1988.
[120] A. H. Chan and M. Medard, "Reconfigurable feedback shift registers" in *Proc. ISIT Tech. Dig. 1997*, p. 178.

[121] M. Medard, A. H. Chan, J. D. Moores, K. L. Hall, K. A. Rauschenbach, and S. Parikh, "Ultrafast cryptography using optical logic in reconfigurable feedback shift registers," in *Multimedia Networks: Security, Displays, Terminals, and Gateways, Proceedings of the SPIE,* 1998, vol. 3228, p. 342.
[122] J. C. Twichell and R. Helky, "Solid state technical report," M.I.T. Lincoln Laboratory, Lexington, MA, Tech. Rep. 3, 1996.
[123] K. L. Hall, S. R. Chinn, D. M. Boroson, and K. A. Rauschenbach, "Wideband Rf signal generation using all-optical switching," in *Proc. Conf. Lasers Electro-Optics Tech. Dig. 1997,* 1997, p. 296.

Vincent W. S. Chan (S'69–M'88–SM'92–F'94) received the B.S., M.S., and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology (M.I.T.), Cambridge, in 1971, 1972, and 1974, respectively, in the area of communications.

From 1974 to 1977, he was an Assistant Professor with the School of Electrical Engineering at Cornell University, Ithaca, NY. In 1977, he joined the M.I.T. Lincoln Laboratory, Lexington, MA, as a staff member of the Satellite Communication System Engineering Group. In January 1981, he became the Assistant Leader of the Communication Technology Group starting a research and development program on optical space communications. In July 1983, he formed and became Leader of the Optical Communication Technology Group and Manager of the LITE (Laser Intersatellite Transmission Experiment) Program. He is currently the Head of the Communications and Information Technology Division and serving as the principal investigator of a new Next-Generation Internet consortium formed among AT&T, Cabletron, M.I.T., and JDS Fitel Company. On January 1, 1999, he will become the Director of the Laboratory for Information and Decision Systems on the M.I.T. Campus and serve on the faculty as Professor of Aeronautics and Astronautics and Electrical Engineering and Computer Science. His research interests are in optical communications, wireless communications, space communications and networks.

Katherine L. Hall (S'91–M'92–SM'96) was born in Burlington, VT, on June 15, 1962. She received the B.A. degree with honors in physics from Wellesley College, Wellesley, MA, in 1984 and the M.S. and Ph.D. degrees in electrical engineering for the Massachusetts Institute of Technology (M.I.T.), Cambridge, in 1990 and 1993, respectively.

From 1984 to 1987, she worked in the Lightwave Systems Research Group, Crawford Hill Laboratory, AT&T Bell Laboratories, Holmdel, NJ. In 1993, she joined the Optical Communications Technology Group at M.I.T. Lincoln Laboratory, Lexington, MA, and is now a senior staff member in the Advanced Networks Group. Her interests include high-speed optical communication systems and networks, subpicosecond optical techniques and nonlinear processes in semiconductor materials.

Dr. Hall is a member of Sigma Xi, the Laser and Electro-Optic Society (LEOS), and the Optical Society of America (OSA).

Eytan Modiano (S'90–M'93) received the B.S. degree in electrical engineering and computer science from the University of Connecticut, Storrs, in 1986 and the M.S. and Ph.D. degrees in electrical engineering from the University of Maryland, College Park, 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. He joined the Communications Division of the M.I.T. Lincoln Laboratory, Lexington, MA, in 1993, where he has been working on communication protocols for satellite, wireless, and optical networks. He has published over 20 papers on various aspects of data networks including multiple access, queueing systems, and distributed protocols. Since 1994, he has also been an Adjunct Professor at the College of Computer Science, Northeastern University, Boston, MA, where he teaches graduate level courses on data networks.

Kristin A. Rauschenbach (M'98) received the B.S. degree in electrical engineering from the University of Minnesota, Minneapolis, in 1984 and the M.S., E.E., and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology (M.I.T.), Cambridge, in 1986, 1987, and 1989, respectively.

In 1989, she joined Arthur D. Little, Inc., in Cambridge, where she worked on a variety of technical product development and business consulting projects. In 1991, she rejoined the M.I.T. Lincoln Laboratory, Lexington, MA, as a staff member in the Optical Communications Group. She is currently an Associate Division Head of the Communications and Information Technology Division at the M.I.T. Lincoln Laboratory. Her technical interests are focused in the areas of networking, optical switching, lasers, and integrated optics.