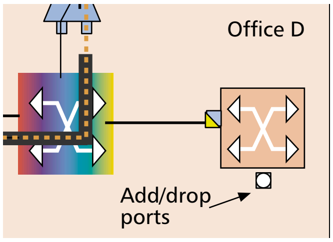

NETWORK CONTROL AND MANAGEMENT CHALLENGES IN OPAQUE NETWORKS UTILIZING TRANSPARENT OPTICAL SWITCHES

G. Ellinas, City College, J.-F. Labourdette, Tellium Inc.,
J. A. Walker, JayWalker Technical Consulting, S. Chaudhuri, Tellium Inc.,
L. Lin, University of Washington, E. Goldstein, University of Washington, K. Bala, Tellium Inc.

ABSTRACT

There is a potential for significant cost, footprint, and power savings by eliminating unnecessary opto-electronic conversions on a signal path in a core optical mesh network. This article addresses and clarifies some fundamental issues surrounding all-optical networking and all-optical switching, and analyzes the trade-offs between transparent and opaque networking. It investigates a number of networking and interface compatibility issues that arise for an opaque network with transparent switches, and presents a number of ways to address these issues.

INTRODUCTION

There is a potential for significant cost, footprint, and power savings by eliminating unnecessary opto-electronic (OE) conversions on a signal path in a core optical mesh network. Current networks have seen the deployment of wavelength-division multiplexing (WDM) technology, followed more recently by the deployment of an optical transport layer where optical crossconnects (OXCs) are connected using WDM links. Both currently deployed WDM systems and OXCs use electronics in the signal path, thereby creating an opaque network. It is very compelling to imagine an optical transport layer where signals remain in the optical domain from the time they enter the network until they leave the network, thereby creating a transparent network [1]. This article considers both opaque networks — OEO conversions occur in the signal path at either the WDM systems if transponders (that incorporate a transmitter and receiver) are present or the switches if transceivers are present — and transparent networks (no OEO conversions in the signal path). It addresses and clarifies some fundamental issues surrounding all-optical networking and switching, and analyzes the trade-offs between transparent and opaque networking, and between transparent and opaque switching in opaque networks. To carry out our assessment of opaque and transparent networks, we make the following basic assumptions on the requirements for core mesh networks:

- Network operators require a lowest-cost network, not just lowest-cost network elements. For example, even though optical may be cheaper than electrical network elements, a network without wavelength conversion and tunable wavelength access in the optical domain could lead to higher network cost due to inefficient capacity usage than a network with wavelength conversion in the electrical domain.
- A network operator must not be constrained to buy the entire network from a single vendor.
- In order to build a dynamic, scalable, and manageable backbone network it is essential that manual configuration be eliminated as much as possible. This requires automatic port/neighbor and network topology discovery, and other networking functions such as service assurance (e.g., access point performance monitoring for service level agreement [SLA] verification), interworking with other network equipment (e.g., unequipped signal), fault management, and performance management regardless of the switching technology.
- An optical switching system must be easily scalable with low cost and a small footprint as the network grows to many hundreds of wavelength channels per fiber and to a speed of 40 Gb/s.

Based on these requirements, we identify the challenges faced by completely transparent core mesh networks. The results of this exercise tend to indicate that core mesh networks will remain opaque for some time. This article focuses on opaque networks with opaque or transparent switches, and explores the potential opportunity for cost reduction and scalability by introducing transparent switches in opaque networks. It also addresses several challenges in network architecture that must be addressed before the potential benefits of transparent switches in opaque networks can be achieved.

The article is organized as follows. The following section describes the possible choices for network architectures. It discusses the challenges associated with transparent network architectures, and addresses the advantages and drawbacks of opaque network architectures. We then investigate the control and management issues that arise in opaque network architectures with opaque or transparent switches, and the role of transparency in the near future. The final section offers some concluding remarks.

NETWORK ARCHITECTURES

Increased traffic volume due to the introduction of new broadband services is driving carriers to deployment of an optical transport layer based on WDM [2]. The network infrastructure of existing core networks is currently undergoing a transformation from rings using synchronous optical network (SONET) add/drop multiplexers (ADMs) to mesh topologies using OXCs [3]. Even though the applications driving large-scale deployment of transparent optical switches are not currently in place (niche applications in today's networks only use a very small number of transparent switches), and the traffic demand does not currently justify the use of transparent switches that are cost effective at very high bit rates, it is possible that at some point in the future transparent switches may be deployed in the network.

Four different node architectures can constitute a reconfigurable core optical network. Fixed patch panels located between WDM systems with transponders are currently being replaced by opaque switching nodes (with electrical switch fabrics), due to the complete lack of flexibility of patch panels (manual intervention is required to change the connections). This second architecture is an opaque network architecture, as the optical signal undergoes OEO conversions at the WDM transponders and the switch fabric [4]. In the third architecture the opaque switching node is replaced by an optical (transparent) switch. This architecture is again an opaque network architecture as the transparent switch resides between WDM systems with transponders. The fourth architecture is a completely transparent network topology, consisting of transparent optical switches and WDM systems that contain no transponders.

TRANSPARENT NETWORK ARCHITECTURE

The transparent network shown in Fig. 1 is a seemingly attractive vision. Since a signal from a client network element (NE), such as a router, connected via a specific wavelength must remain on the same wavelength when there is no wavelength conversion, only a small size switch fabric is needed to interconnect the WDMs and NEs in a node. This architecture also implies end-to-end bit rate and data format transparency. Another architecture of a transparent switch in a transparent network may include a single large fabric instead of multiple switch matrices of small port counts. If one is to provide flexibility, such an architecture design would require the use of tunable lasers at the clients and wavelength conversion.

This network architecture may provide significant footprint and power savings, and on the surface suggests cost savings. However, while the transparent network architecture may be a viable option for small-scale networks with predetermined routes and limited numbers of nodes, it is not a practical solution for a core mesh network for the following reasons:

- This network does not allow wavelength conversion,[1] thus essentially creating a network of n (n being the number of WDM channels) disjoint layers. Inflexible usage of wavelengths in this network would lead to increased bandwidth and network operational cost, thus negating all savings that may result from elimination of OE conversion. In addition, for this technology to be effective and to build a flexible network for unrestricted routing and restoration capacity sharing, an all-optical 3R-regeneration[2] function must be available. Such a technology that can be harnessed in a commercial product does not currently exist [5].

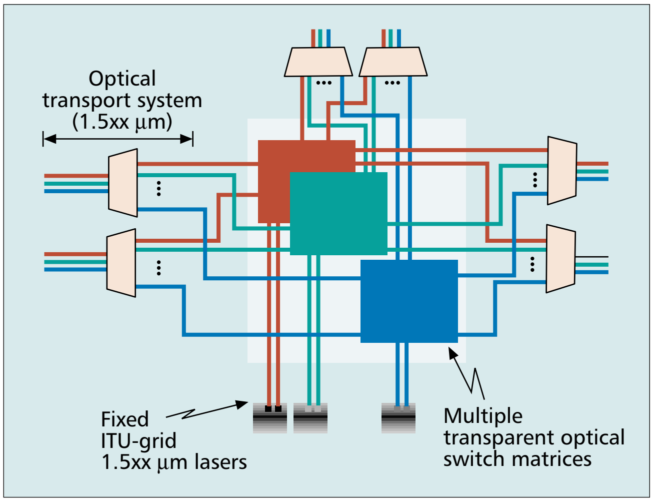

FIGURE 1. Transparent switch architecture in a transparent network.

- In the absence of wavelength conversion, the wavelength continuity constraint on backup paths makes dynamic resource sharing almost impossible in transparent networks; consequently, no dynamic shared mesh restoration can easily be offered. Studies have shown that with shared mesh restoration, opaque networks can accommodate 95 percent of a demand set before they experience blocking. On the other hand, with dedicated protection (which is the only type of protection that can be offered easily in all-optical networks), transparent networks accommodate only between 45 and 50 percent of the same demand set before they experience blocking [6]. This in turn means that the capacity requirement for protected services is significantly higher (80–100 percent) for transparent networks than for opaque ones [6].
- Physical impairments such as chromatic dispersion, polarization mode dispersion, fiber nonlinearities, polarization-dependent degradations, WDM filter passband narrowing, component crosstalk, and amplifier noise accumulate over the physical path of the signal due to the absence of OE conversion. The accumulation of these impairments requires engineering of end-to-end systems in fixed configurations [7]. It is thus not possible to build a large network with an acceptable degree of flexibility.
- The design of high-capacity dense WDM (DWDM) systems is based on intricate proprietary techniques, eluding any hope of interoperability among multiple vendors in the foreseeable future. Also, operators do not have the flexibility to select the client NE and WDM vendors independently. This is because the interface optics at the client NE launch the signals through the all-optical switch directly into the WDM system without OE conversion. Consequently, transparent networks by necessity are single-vendor (including the client NEs) solutions.
- Polarization mode dispersion (PMD) and signal-to-noise ratio (SNR) constraints are considered in determining the

[1] *Our assumption here is that there will be no commercially viable wavelength conversion technology in the optical domain available in the next several years.*

[2] *3R regeneration function implies retiming, reshaping, and reamplification of the signal.*

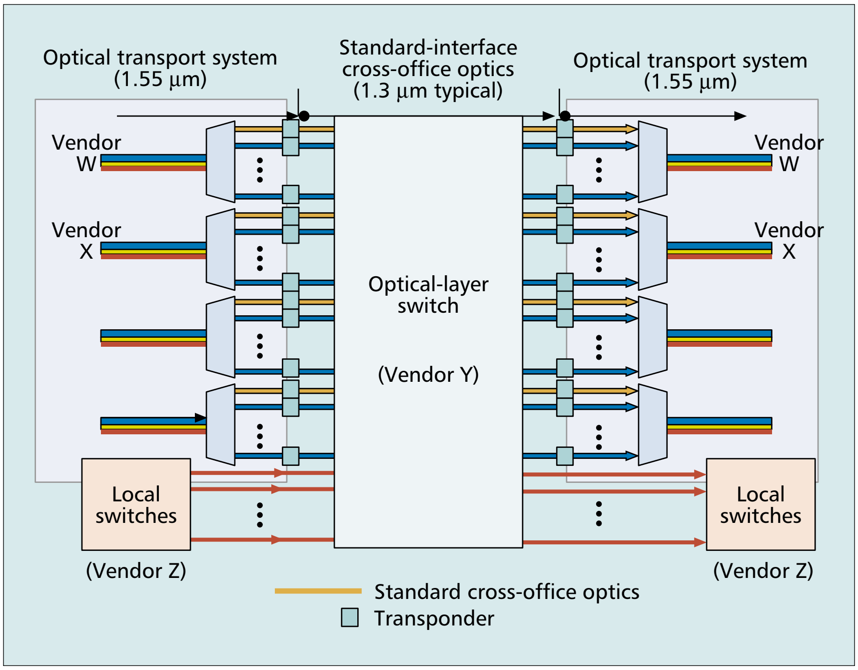

FIGURE 2. Multivendor interoperability and wavelength translation as by-products of opaque network architectures.

route of a lightpath through a transparent network. The PMD requirement becomes an issue at bit rates of 40 Gb/s and higher, and SNR can potentially constrain the number of spans for a lightpath to 3 [8]. Thus, routing algorithms in transparent networks should explicitly include PMD and SNR constraints. However, the challenge of performance engineering continental-scale transparent reconfigurable wavelength-routed networks remains severe and, in networks that push limits, remains unsolved despite some attempts to formalize the routing problem [8].

It is apparent that a number of key carrier requirements — dynamic configuration, wavelength conversion, multivendor interoperability of transport equipment (WDM), low network-level cost — would be very hard to meet in a transparent network architecture. Therefore, an opaque network solution will remain for now the only practical and cost-effective way to build a dynamic, scalable, and manageable core backbone network.

OPAQUE NETWORK ARCHITECTURE

Even though the opaque network solution may be more expensive in terms of equipment costs when the core network capacity increases significantly, it offers the following key ingredients for a large-scale manageable network:

- No cascading of physical impairments. This eliminates the need to engineer end-to-end systems (only span engineering is required) and allows full flexibility in signal routing.
- Multivendor interoperability using standard intra-office interfaces. Figure 2 demonstrates an example of how different client, WDM system, and switch vendors can operate together when the switch uses standard cross-office optics.
- Wavelength conversion enabled. Opaque switches that use standard interface cross-office optics, such as the one shown in Fig. 2, can utilize network capacity for service without any restrictions, and additional significant cost savings can be offered by sharing restoration capacity in mesh architectures.
- The network size and length of the lightpaths can be large, since regeneration and retiming are present along the physical path of the signal.
- Link-by-link network evolution. This permits link-by-link incorporation of new technology, as the network is partitioned into point-to-point optical links.

Having shown that transparent core mesh network architectures are likely to remain unrealistic for quite some time, we now turn our attention to opaque network architectures. Today's architecture contains opaque switches (with an electronic switch fabric) in an opaque network (with transponders present in the WDM system). This architecture is shown in Fig. 3a. The interfaces to the fabric are opaque interfaces, with transceivers present at all interfaces to the switch, which subsequently enable the switch to access the SONET/SDH overhead bytes for control and management functions. These transceivers provide support for fault detection and isolation, performance monitoring, connection verification, neighbor/topology discovery and signaling, as well as support for implementing the network routing and restoration protocols.

The opaque switch approach was, however, faced with a number of challenges when confronted with the (unrealistic) traffic growth projections from just a few years ago. It would eventually reach scaling limitations in signal bit rate (e.g., OC-192), switch matrix port count (e.g., 1000 × 1000 switch fabric), and NE cost. These were the key motivations behind the attempt to develop large port-count (256 to 1000 ports and beyond) transparent switches. Note, though, that the opaque switches would still have remained in the network architecture in order to provide some key network functions: grooming and multiplexing, SLA verification, and control and management.[3]

Figure 3b shows a transparent switch architecture. This switch architecture has transparent interface cards and no opaque transceiver (TR) cards on its add/drop ports. Therefore, it has no direct access to the overhead bytes for control and signaling. The optical switch fabric is bit-rate-independent and accommodates any data rates available. The advantage of such a switch architecture is that for an $N \times N$ architecture there are N interfaces/ports to the switch fabric regardless of the type of interfaces. No data-rate-specific interface cards are used, so no replacement is needed when the switch operates at higher data rates (or different formats), provided the optical power budget is sufficient for that rate. This is in contrast to the opaque switches where the number of ports depends on the type of port. The drop-side ports are connected to a client (e.g., IP router, opaque OXC, or ATM switch) that provides SONET/SDH termination through its opaque ports.[4]

The promise of optical switching was that, unlike integrated electronic switches, an optical switch fabric's complexity is a flat function, independent of the bit rate of the signals it handles. Transparent switches were expected to be cheaper in terms of switching fabric and interface card cost than opaque switches. This would have resulted in significant cost reduction to network operators because a large amount of the traf-

[3] If grooming and multiplexing functions are not required, it is possible to provide SLA verification, and control and management functions via a transparent switch with O/E interfaces for the drop ports.

[4] Note that integrating the opaque interfaces on the drop-side interfaces of the transparent switch can also provide the opaque function through inband signaling.

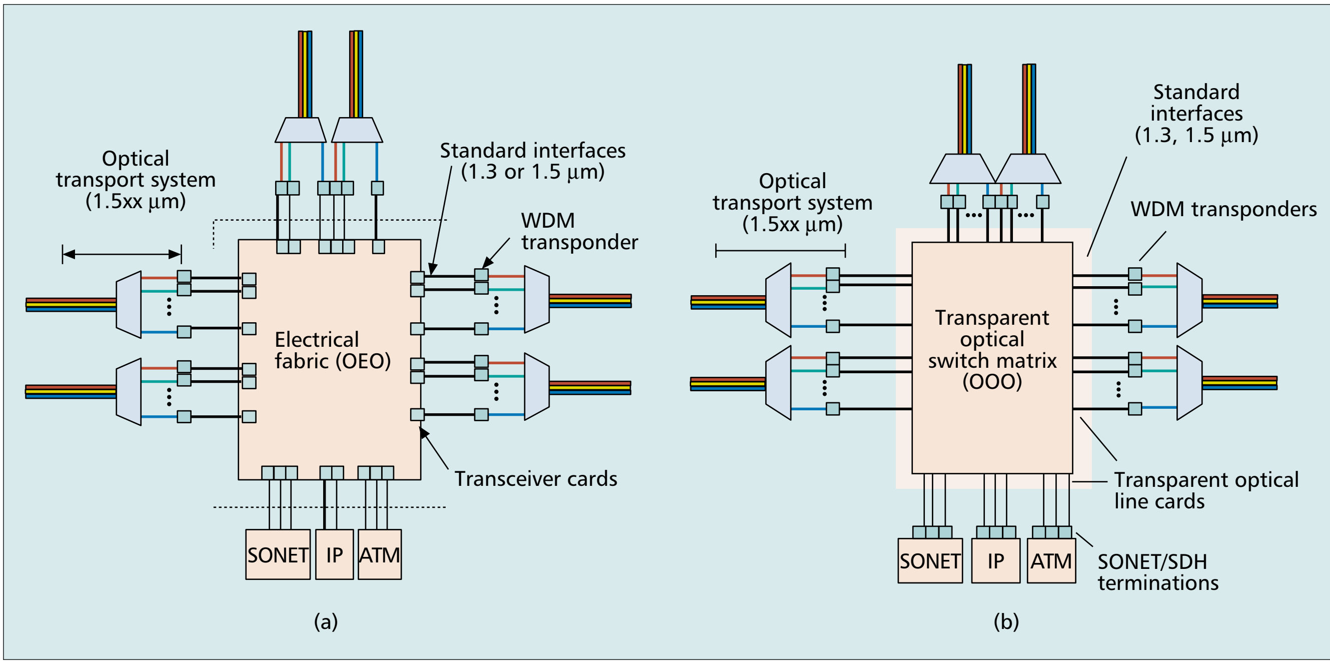

FIGURE 3. *a) Opaque switch architecture; b) transparent switch architecture.*

fic that passes through an office would be able to bypass the opaque switch (typically approximately 75 percent through-to-total ratio). This would in turn eliminate about 75 percent of the network element's OE interfaces, and thus something approaching 75 percent of its cost, power, and footprint.[5] Transparent switches essentially would have helped relieve the demand for opaque switch ports and reduce the cost of transporting lightpaths.

Since the transparent switch fabric is bit-rate- and data-format-independent, the switch matrix can scale more easily than electrical switch fabrics to potentially 1000, 4000, or even 8000 ports that can accommodate up to 40 Gb/s per port. For these reasons, as bit rates rise, it was thought that optical switch fabrics would eventually prevail. Even though in early stages of development, the crossover point at which the cost of opaque switch fabric was still going to be cheaper than transparent switch fabric appeared to be at OC-48, it soon moved to OC-192 and beyond. But under today's more realistic traffic growth scenario, and given the lack of deployment of 40 Gb/s WDM systems and the continued decline in price of OEO components, the need for and promise of transparent switches appeared to have moved beyond the foreseeable future.

Interface	VSR	SR-1	SR-2	IR-1	Hi-P VSR	LR-1
Standards body	OIF	Bellcore	Bellcore	Bellcore	OIF	Bellcore
Length (km)	0.6	12	20	24	0.6	48
Nominal wavelength (nm)	1290–1330	1290–1330	1530–1565	1290–1330	1290–1330	1290–1330
Attenuation (dB)	4+1	6+1	6+2	11+1	11+1	22
Optical path penalty (dB)	1	1	2	1	1	
Unit power consumption (W)	6	7.5		11	6	13

Table 1. *Cross-office standard interfaces (10 Gb/s).*

Besides the demise of several of the drivers for high-port-count transparent switches, important challenges remain to be solved even in an opaque architecture. The main challenge to architectures that use transparent switches is providing the control and management functionalities that are readily available when we have access to the electrical signal and consequently to the SONET/SDH overhead bytes. The section that follows focuses on the network control and management functions for an opaque network with opaque or transparent switches.

Another important operational issue associated with a transparent switch that must be taken into consideration is power budget management. Because of the relatively high insertion loss of the optical switch fabrics and the resulting loss from input to output port (between 3–6 dB fiber-to-fiber insertion loss that includes connectors, misalignments, and all impairments with path loss uniformity of less than 1 dB), traditionally deployed cross-office optics cannot be supported with a transparent switch. Therefore, such architectures require higher-cost cross-office optics or new low-cost ones currently being worked on in the Optical Interworking Forum (OIF) [9]. Table 1 shows the current and proposed power requirements for various cross-office interfaces (10 Gb/s signals).

[5] *Studies show that 90 percent of OEO switches' cost resides in the electronic part, especially the transponders, one of which is required at every port to convert the signal from the optical to the electronic domain and vice versa. Furthermore, transponders consume a substantial amount of power, generate heat that must be dissipated, and entail larger floor space occupancy — a combination of factors that ultimately leads to even higher operational costs. For example, the per port power consumption for an opaque interface at OC-192 rates is typically around 50 W and could fall to approximately 6–10 W for transparent switch interfaces.*

Network Control and Management for Opaque Networks

With Opaque Switches

Access to the SONET/SDH overhead bytes at the opaque interface cards is a key enabler of network control and management functionalities as outlined below.

- It allows an opaque switch to perform in-band signaling and provisioning functions. Access to the overhead bytes also allows the switches to run dedicated (1+1) protection and shared mesh restoration protocols.
- It allows an opaque switch to run automated neighbor and topology discovery protocols. The ability of the network to autonomously create and maintain its resource databases is the fundamental building block for an efficient, flexible, and manageable network. The neighbor discovery protocol allows the network to create and maintain the critical port-state and topology databases. An automated neighbor discovery protocol is required in order to perform node-port assignment discovery between client and switch architectures and inter-office node-port associations between two neighboring switch architectures. The network topologies are then created automatically by, for example, a management system that utilizes the associations found during the neighbor discovery process.
- It allows an opaque switch to perform fault detection and performance monitoring. Lightpath-based restoration or switch fabric protection switching can be triggered by a detected failure condition. Fault isolation in such architecture relies on the alarms generated by the interface cards after a failure is detected.
- It enables the switch to provide service assurances such as performing connection verification and control to avoid misconnections.
- It allows an opaque switch to generate an unequipped signal (keep-alive) at every idle transceiver on the switch's network side to prevent alarms in other equipment connected to the switch. This feature is required in switch architectures that implement shared mesh restoration since the channels have been provisioned but do not carry any signal until a failure event occurs.

With Transparent Switches

The lack of laser transmitters and access to the overhead bytes in a transparent switch pose a number of challenges in creating a seamless interoperable and manageable network. Network control and management features are collectively very difficult to achieve in a transparent switch without forfeiting the economies the switch was designed to extract. In order to address these challenges, an opaque function is required that may be provided by either deploying opaque cards on the drop side of a transparent switch or relying on the opaque function located at WDM transponders and/or the client equipment, effectively using the equipment as proxies.

Automatic Port/Neighbor and Topology Discovery — Automatic port/neighbor discovery and topology discovery are key aspects of a service provider's requirements. The Link Management Protocol (LMP) standard [10] has been proposed to automatically discover node-port associations between the client and transparent switches, and between two neighboring transparent switches. LMP handles transparent switches by using dedicated opaque cards temporarily or the opaque interfaces on the clients and out-of-band signaling to discover connectivity between switches. After LMP is executed and the node-port associations are ascertained, the network topology can be created automatically by a centralized management system or in a distributed way. If LMP cannot be implemented, a proprietary automated procedure can be designed for port/neighbor discovery. In order to automate this process, the transparent switch has to be equipped with a small number of opaque interface cards providing electrical signal generation functions. These cards conduct neighbor/port discovery in an automated process by establishing a connection between the opaque interface cards of two adjacent switches, as shown

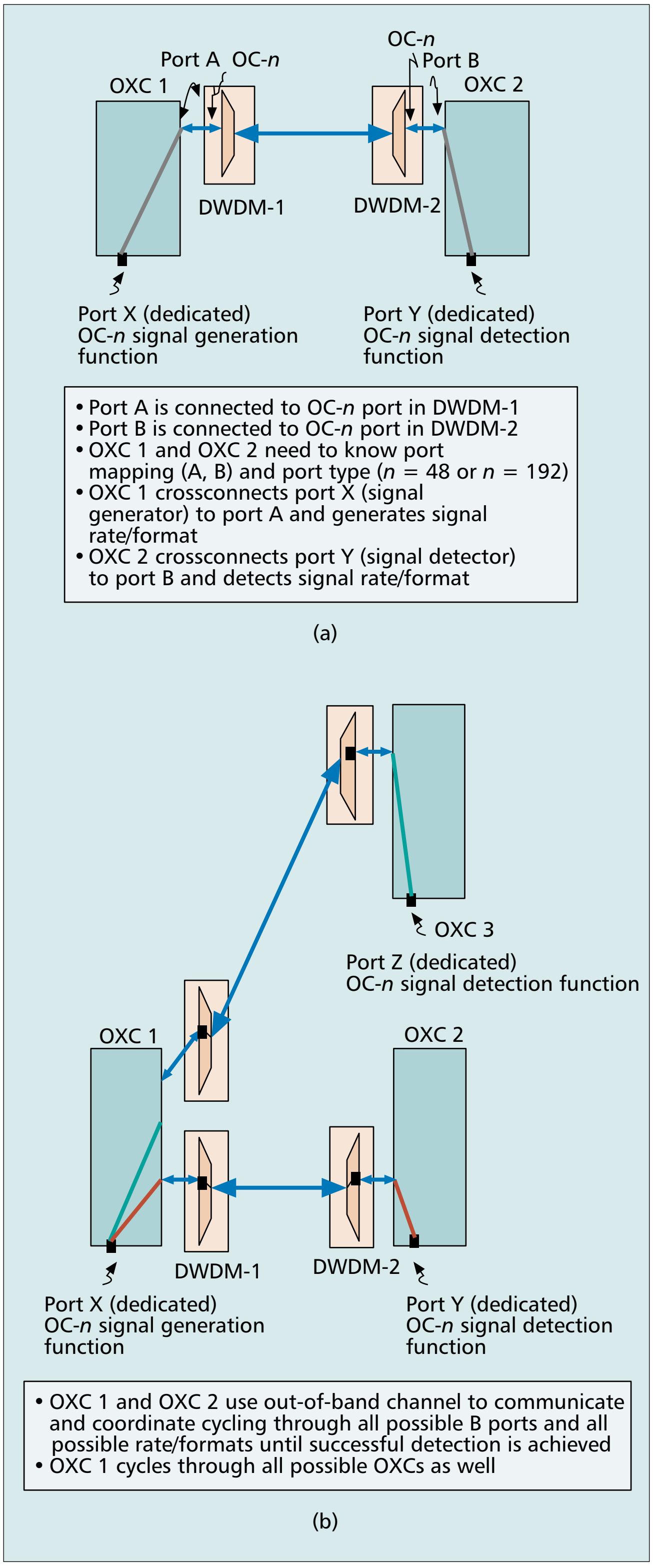

FIGURE 4. a) Use of opaque cards for automated neighbor and port discovery; b) cycling through all possible OXCs and ports

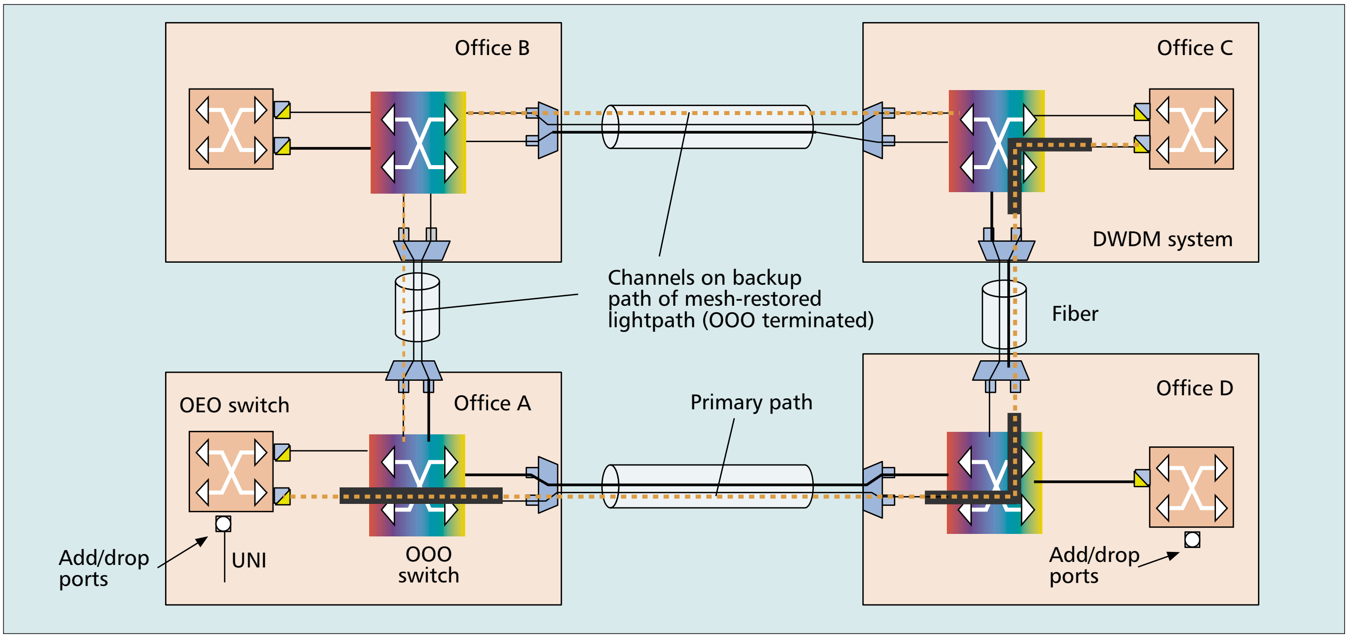

FIGURE 5. Shared mesh restoration example.

in Fig. 4a. When a transparent interface port on the switch is connected to a WDM transponder, the port status is (manually) changed to a new state called *ready-to-discover*. Upon such change of port status or at predefined times, the automated neighbor/port discovery process will be executed.

In a *centralized approach*, an element management system (EMS) would instruct a switch (e.g., OXC_i) to connect the OC-*n* signal generation port to one of the ready-to-discover ports. It then instructs another switch (OXC_j) to set up crossconnects between the OC-*n* signal detection port and the ready-to-discover ports, one after the other, until the signal sent by OXC_i is detected (Fig. 4b). If such a signal is detected, OXC_i reports this to EMS, which derives the neighbor/port adjacency relationship and informs the two nodes of it. If the appropriate signal is not detected on any of OXC_j's ready-to-discover ports, EMS instructs another switch, say OXC_k, to carry out a similar procedure. This procedure is repeated until a signal from OXC_i is detected, yielding a neighbor/port association or all the ready-to-discover ports at all switches have been tried. If no association is discovered, the port on OXC_i remains in the ready-to-discover state.

In a *distributed approach*, OXC_i communicates with all other OXCs to determine if they have ports in the ready-to-discover state. If they have, the corresponding OXC goes through the procedure of crossconnecting its signal detection port to every ready-to-discover port, each in turn. If the signal is correctly received, OXC_j communicates its identity and that of the port to OXC_i. In turn, OXC_i communicates the corresponding information to OXC_j. Both switches then create the appropriate entry into their neighbor/port adjacency database. Such updates are then communicated to EMS. Both the centralized and distributed procedures could be carried out in parallel at all switches with ready-to-discover ports to speed up the discovery process. Furthermore, both procedures may also require trying different opaque transceiver speeds as well, using the technique described above. However, as this process happens infrequently, there are no real-time requirements on creating these adjacency relationships.

Lightpath Provisioning and Network Protection and Restoration — A-to-Z provisioning and signaling is done in a centralized or distributed approach and can be triggered by an operator (point-and-click A-to-Z provisioning) or user–network interface/generalized multiprotocol label switching (UNI/GMPLS) signaling through the opaque function (at the client or the add/drop ports) [11, 12].

Network protection and restoration can be provided in two different ways. Protection and restoration can be supported entirely through the opaque clients of the transparent switches. In this case, the transparent switches are not involved in the protection and restoration process, and all lightpaths and shared backup channels effectively terminate on the opaque clients. In such a scheme, dedicated protection times of 50 ms and shared mesh restoration times of approximately 200 ms can be achieved. Alternatively, protection and restoration can be supported entirely within the transparent switch-based network, with lightpaths and shared backup channels terminating on transparent switches. The restoration crossconnects are then performed at the transparent switches upon appropriate triggering (e.g., signal degrade or signal fail conditions) by the opaque function through a control link. Note that optical performance monitoring (OPM) can take place at the transparent switch interfaces in the form of optical power monitoring, but electrical performance monitoring can only take place at the opaque endpoints of a lightpath, since there is no signal visibility on transparent switches.

Figure 5 shows a network of interconnected transparent switches that have client architectures (opaque switches) connected to their add/drop ports, and the channels on the backup path of the shared mesh restored lightpaths terminate at the transparent switches. In this example, the transparent switches take an active part in the restoration by reconfiguring their switch fabrics in order to route the backup path.

Restoration triggering and signaling in transparent-based restoration are expected to differ from the corresponding functionality used in opaque-based restoration. An out-of-band signaling scheme will now be required to support transparent-based restoration. A number of restoration triggers can also be used. For example, the failure can be detected at the WDM system, and a communication protocol between the OXC and WDM systems [13, 14] can be used to trigger restoration at the switch. The failure can also be detected at the client's interface, and an out-of-band control channel between the transparent switches and the client drop nodes can be used to trigger restoration at the switches. An out-of-band signaling channel would also be required between the lightpath end nodes (transparent switches) and the client drop nodes to support some form of local protection between the client and a transparent switch within an office, complementing the network protection/restoration scheme. A handshaking protocol between the client drop nodes and switch end nodes is also required in order to accommodate the interplay between network restoration and the switch-client node protection protocol.

Transparent-based restoration is a lower-cost approach than opaque-based restoration, as it saves a number of opaque interface ports. However, it may result in slower restoration (possibly on the order of seconds), and it requires new out-of-band signaling channels between transparent switches, and between transparent switches and client equipment. Note that the latter will also require vendor cooperation and new standards definition.

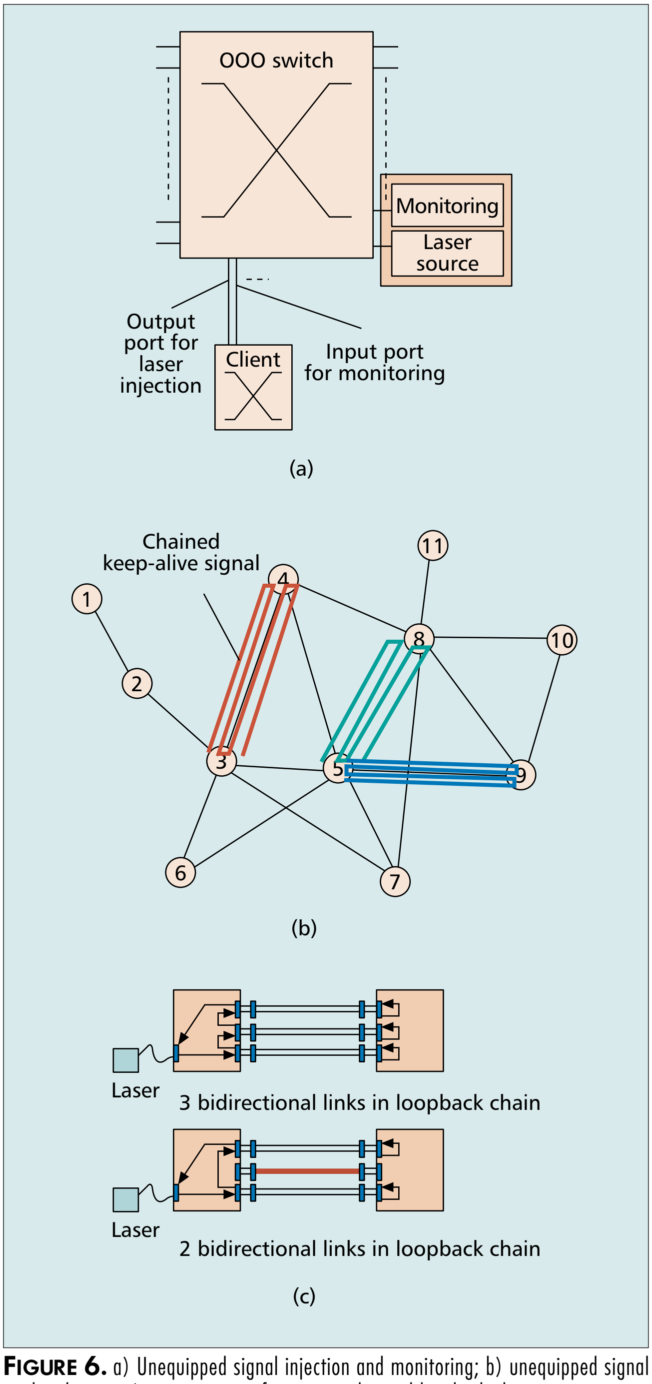

FIGURE 6. a) Unequipped signal injection and monitoring; b) unequipped signal distribution; c) maintenance of unequipped signal loopback chain.

Fault Detection and Fault Localization — Fault detection, other than loss of light (LOL) that is monitored at the transparent interfaces of the switch, takes place in the opaque interfaces located at the lightpath endpoints. A failure must be detected and isolated at the level of a replaceable unit. The fault isolation need not be an instantaneous process. As long as the failure is detected (at the lightpath endpoints) and a restoration mechanism is triggered quickly, fault isolation can be a slower process. With transparent switches, fault localization can take place at the management system by correlating the alarm information generated by the switches. In some cases, fault localization may require alarms generated by the WDM systems as well. In that case, one could essentially use the transponder's access to the electrical signal as a proxy for opaque interfaces in support of control and management functions. In addition, the usage of sequential loop-back (as will be explained in the following section) can support fault localization, when it is not possible via alarm correlation. Note that, if a communication channel between the switches and the WDM systems is implemented, fault localization is expected to be a simpler process [13, 14].

Unequipped Signal Generation, Distribution, and Maintenance — In a network with interconnected OXCs, to be capable of providing shared mesh restoration, the provisioned mesh-restored channels (when not in use) require the presence of an unequipped (keep-alive) signal. This is true because the lack of an unequipped signal results in a) alarms generated at the WDM systems that have knowledge of provisioned channels but detect no light on those channels, b) lack of monitoring of the restoration channels to ensure availability when/if a failure occurs, and c) increased restoration time if a failure occurs, due to the additional time required to turn on the WDM lasers and perform power adjustments and equalization. Unequipped signal generation becomes an issue in transparent switches, since the (transparent) interface ports cannot inject a keep-alive signal. This issue can be addressed by filling all idle channels (unused and reserved for restoration) with unequipped signals, using a limited number of lasers at the drop side of the optical switches and having each signal propagate along several idle channels, looping back and forth between the switches (Figs. 6a–c). The header of the keep-alive signal will contain the address of the originating node followed by the address of the looped-back node as the ID of the keep-alive signal.

The proposed scheme equips transparent switches with additional opaque transceivers so that all provisioned ports between two switches can be looped back through the fabric, thereby creating an unequipped signal on the ports. This approach requires as many opaque transceivers as roughly half the average node degree per switch.

There are two ways to inject unequipped signals into a network (Fig. 6a). One is to use the opaque functionality at the add/drop ports of each node to generate and monitor the keep-alive signal. The advantage of this approach is that there is no need for separate laser sources, monitors, and the software to manage them. Another approach is to use separate laser sources and monitors, and to connect them directly to the transparent switch. This approach saves one client port, but loses the corresponding advantages as well.

Keep-alive signals are sent to all the shared mesh restoration channels not in use, and to all the working channels not in use. Consider a keep-alive network that consists of all those not-in-use channels, as in Fig. 6b. The two node addresses forming the keep-alive signal ID can be used to establish an order in the network. The keep-alive signal can be sent from the node with the smaller ID to the node with the larger ID in a looped-back chain, as shown in Fig. 6b. The signal traverses only one span from the source node and always comes back to the same node from which it was sent (Fig. 6c). Each node knows all the channels that are not used, and are attached to it. Among those channels, it also knows to which channels it should send its own keep-alive signal, and on which channel it should expect the keep-alive-signal to be sent by a neighboring node. Finally, the ports on a switch connected through multiple links to another switch need to be crossconnected appropriately as part of the loopback chain to avoid cycles that would not carry the unequipped signal. This can be accomplished by having the crossconnect detect the signal entering the switch (OPM functionality) and then complete the loopback crossconnect according to an increasing order of port ID. This technique allows the switches to autonomously set the loopback crossconnects, since after each crossconnect one of the switches will detect a signal on one of its provisioned ports.

Loopback chains would change when a new port pair is added, used for provisioning service, or used for restoration of a shared mesh restorable service.

A complication in this scheme arises when a channel in the loopback chain is seized for shared mesh restoration. When this happens, the loopback chain could be open for the time it takes to reconfigure the switch fabric and create a new chain. Soak times of more than 10 ms in the WDM equipment will ensure that during this time alarms will not be raised and the WDM lasers will not shut down.

If dedicated opaque cards are not present, an out-of-band communication scheme between the OXCs and WDM systems could be used to work around the keep-alive signal generation issues by suppressing alarms and keeping the WDM lasers on even in the absence of keep-alive signals for provisioned but not in use channels [13, 14].

Role of Transparency in the Near Future

Even though the use, in the core, of transparent switches that are cost effective at very high bit rates is not currently justified, there still exist some niche applications in today's networks that could use a small number of transparent switches.

Based on the challenges outlined in the previous sections, and the forecast traffic demand and near-term applications that will be supported, it is anticipated that in the near future transparency will be first limited in metropolitan area networks, utilizing reconfigurable optical add/drop multiplexers (ROADMs) [15], and some ultra-long-haul applications in the core, utilizing a small number of wavelength-selective crossconnects (WSXCs)/ROADMs on high-capacity routes. ROADMs can be utilized in MANs at central offices and customer locations in much the same way that SONET's introduction created a need for large numbers of SONET ADMs. They provide network flexibility and can be used to manage continually changing traffic patterns and customer service requirements. WSXCs can also be used in ultra-long-haul applications in the core network in a completely transparent manner. Even though these network elements allow for end-to-end bit rate and data format transparency, as previously outlined, they face a number of challenges. However, in the near future it is possible that these network elements could be utilized in a few predetermined and nonreconfigurable high-capacity routes to provide end-to-end transparency between fixed end nodes.

Conclusion

The current state of affairs in terms of network deployment, applications, and traffic demand does not justify the large-scale use of transparent switches in today's networks. Some niche applications do exist, but can mostly be addressed using a number of small transparent switches. Provided the traffic grows and bit rates increase substantially, there may emerge a need for an additional network layer utilizing transparent optical switches. In the meantime, the deployment of transparent network elements is expected to remain limited to WSXC and ROADM architectures on high-capacity routes.

While completely transparent core mesh networks are still far off, even transparent switches in opaque networks face technological as well as control and management challenges. Although most of these issues can be addressed via clever innovation as well as standardization efforts, transparent switches complemented by an opaque function will not be ready for deployment in the network until all the challenges addressed in this article are successfully resolved.

Furthermore, we anticipate that opaque switches will always remain for the embedded service base even after transparent switches are eventually introduced in the network. These opaque switches will provide the grooming and multiplexing functions, as well as some of the necessary control and management functions, and will scale and decrease in cost with rapid progress in electronics.

References

[1] P. Green, "Optical Networking Update," *IEEE JSAC*, vol. 14, no. 5, June 1996, pp. 764–79.

[2] T. Stern and K. Bala, *Multiwavelength Optical Networks: A Layered Approach*, Prentice Hall, May 1999.

[3] R. D. Doverspike, S. Phillips, J. R. Westbrook, "Future Transport Network Architectures," *IEEE Commun. Mag.*, vol. 37, no. 8, Aug. 1999, pp. 96–101.

[4] K. Bala *et al.*, "The case for opaque multiwavelength lightwave networks," *IEEE/LEOS Summer Topical Meeting*, 1995.

[5] B. Mukherjee, "WDM Optical Communication Networks: Progress and Challenges," *IEEE JSAC*, 18, no. 10, Oct. 2000, pp. 1810–24.

[6] E. Bouillet *et al.*, "Addressing Transparency in DWDM Mesh Survivable Networks," *Proc. OFC '01*, Anaheim, CA, Mar. 2001.

[7] B. Ramamurthy *et al.*, "Impact of Transmission Impairments on the Teletraffic Performance of Wavelength-routed Optical Networks," *IEEE/OSA J. Lightwave Tech.*, vol. 17, Oct. 1999, pp. 1713–23.

[8] J. Strand, A. Chiu, and R. Tkach, "Issues for Routing in the Optical Layer," *IEEE Commun. Mag.*, Feb. 2001, pp. 81–87.

[9] J. Ojha, "OIF Very Short Reach (VSR) Interface," OIF White Paper, http://www.oiforum.com.

[10] "Link Management Protocol (LMP)," IETF draft, draft-ietf-ccamp-lmp-09.txt, June 2003.

[11] "User Network Interface (UNI) 1.0 Signaling Specification," OIF cont. OIF2000.125.6, Sept. 2001.

[12] "Generalized MPLS — Signaling Functional Description," IETF draft, draft-ietf-mpls-generalized-signaling-05.txt, Jan. 2001.

[13] "Link Management Protocol (LMP) for Dense Wavelength Division Multiplexing (DWDM) Optical Line Systems," IETF draft, draft-ietf-ccamp-lmp-wdm-02.txt, Mar. 2003.

[14] "Optical Link Interface Requirements," IETF draft, draft-ietf-ccamp-oli-reqts-00.txt, Feb. 2002.

[15] J. Bayne and M. Sharma, "Broadcast-and-Select OADM Enables Low-Cost Transparency," *Lightwave Mag.*, Dec. 200.

Additional Reading

[1] G. Ellinas *et al.*, "Transparent Optical Switches: Technology Issues and Challenges," *IEC Annual Rev. Commun.*, vol. 56, Nov. 2003.

Biographies

GEORGIOS ELLINAS (gellinas@ccny.cuny.edu) holds B.S., M.S., and Ph.D. degrees in electrical engineering from Columbia University. He is currently an associate professor of electrical engineering at City College of the City University of New York. Before joining City College, he was a senior network architect at Tellium Inc. In this role he worked on lightpath provisioning and fault restoration algorithms in optical mesh networks, and the architecture design of the MEMS-based all-optical switch. He also served as a senior research scientist in Telcordia Technologies' Optical Networking Research Group. He performed research for the Optical Networks Technology Consortium (ONTC), Multiwavelength Optical Networking (MONET), and Next Generation Internet (NGI) projects from 1993 to 2000. He also served as an adjunct assistant professor at Columbia University and the University of Maryland, teaching courses on multiwavelength optical networking in 1999 and 2000, respectively. He was awarded a Fulbright fellowship from 1987 to 1991 for undergraduate studies at Columbia University, and has authored and co-authored more than 50 journal and conference papers. He is also the holder of 12 patents on optical networking, and has more than 20 U.S. and international patent applications currently pending.

JEAN-FRANCOIS LABOURDETTE [SM] is manager of system engineering at Tellium, a position he has held since January 2003. In this position he is responsible for all network element and network management system engineering activities. Before that, beginning in October 2000, he was manager of network routing and design at Tellium, where he was responsible for Tellium's routing architecture and algorithms, network dimensioning, and design studies. Previously, he was a manager of multiservice globalization planning at AT&T, where he was responsible for network, service, and operations planning for AT&T international data services (frame relay, ATM, and MPLS IP VPN). Earlier at AT&T he was a system engineer in the routing planning group, working on dynamic call routing for AT&T's switched network and facility routing and rearrangement for the AT&T network. He received his undergraduate degree in electrical engineering from l'Ecole Nationale Supérieure des Télécommunications, Brest, France. He holds a Ph.D. from Columbia University, where he was the recipient of the 1991 Eliahu I. Jury Award for best dissertation. He is a member of the Optical Society of America. He has more than 40 publications in international conferences and journals, including a Best Paper Award (list of publications available at http://www.edas.info/S.cgi!author=labourdette). He has served on the Technical Program Committees of numerous international conferences and on panels for the National Science Foundation.

JAMES WALKER is president and founder of JayWalker Technical Consulting. His experience includes 16 years in Bell Labs-Research where he specialized in optical telecommunications applications of micro-electromechanical systems (MEMS) and two years as director of advanced technologies at Tellium, Inc. where he was jointly responsible for leading a group of world-class researchers developing a large-port-count all-optical crossconnect based on 3D MEMS technology. He and his Bell Labs colleagues are credited with the formation of a platform of MEMS devices based on moving dielectric films called Mechanical Anti-Reflection Switch (MARS) technology. Network elements based on this platform include high-speed opto-mechanical modulators, variable optical attenuators, dynamic dispersion compensators, and dynamic gain equalizing filters. He was also instrumental in the development of a gallium arsenide on silicon integration technology known as opto-electronic VLSI that has formed the basis of high-speed backplane optical interconnects and spawned Aralight Inc. He is also a member of the Technical Advisory Board of Flex Micro, Inc. He earned his Bachelor's and Master's degrees in electrical engineering from Rutgers University in 1984 and 1989, respectively. He holds over 49 U.S. patents, has authored over 100 journal and conference papers, has given numerous invited papers at technical conferences, and has chaired several international technical conferences.

SID CHAUDHURI is director of network architecture at Tellium, a position he has held since March 2000. In this role he led the development of the mesh restoration and IP-centric control architecture implemented in the company's Aurora Optical Switch. He also led the development of Tellium's software design tools. Prior to joining Tellium, he worked for AT&T Laboratories Research where he developed intelligent optical network architectures using SONET, DWDM, and optical switching technologies, and spearheaded AT&T's Core Transport Network architecture. Previously, he was a distinguished member of technical staff at Bell Laboratories, where he developed SONET/SDH transport and crossconnect system architecture. He served as president and chairman of the Optical Internetworking Forum (OIF) and as chairman of its physical and link layer group, where he led the group in the first set of interoperability agreements. He holds a Ph.D. from the University of Pittsburgh, where he was an Andrew W. Mellon fellow.

LIH Y. LIN [SM] joined the Electrical Engineering Department at the University of Washington as an associate professor in 2003. She received her Ph.D. in electrical engineering from the University of California at Los Angeles in 1996, with thesis topics on high-power high-speed velocity-matched distributed photodetectors and micromachined integrated optics. From 1996 to 2000 she worked in AT&T Labs-Research as a senior technical staff member. Her main research activities were in micromachined technologies for optical switching and lightwave communication system. In March 2000 she joined Tellium, Inc. as director of optical technologies to work on high-port-count MEMS optical crossconnects. She is a member of OSA. She has served on the technical program committees of and co-chaired various technical conferences, including the International Optical MEMS Conference, CLEO Pacific Rim, IEEE LEOS Annual Meeting, OSA Annual Meeting, and OSA Photonics in Switching Topical Meeting. She is now on the steering committee of the International Optical MEMS Conference. She was also an editor of the *Journal of Lightwave Technology* Special Issue on Optical MEMS and Its Future Trends.

EVAN GOLDSTEIN is currently an affiliate professor at the University of Washington. From 2000 to 2002 he was the director of optical networking systems at Tellium. Prior to joining Tellium, he held various research positions with AT&T Labs, Bellcore, and Bell Labs. He has co-authored well over 120 publications covering various aspects of lightwave systems, focusing on erbium doped fiber amplifiers, lightwave system performance, architectural issues in WDM networks, and the application of lightwave micromachines as core-transport switches. He has served on the technical program committees of all major conferences in the optical communications field. He is also a frequent lecturer and short-course instructor at research institutions and conferences around the globe. He has been a fundamental contributor to the field of WDM communications since the inception of research in the area in the mid-1980s. He holds a B.A. in philosophy from Antioch College, M.A. and M.Phil degrees in philosophy from Columbia University, and B.S., M.S., and Ph.D. degrees in electrical engineering from Columbia University.

KRISHNA BALA is co-founder and chief technical officer of Tellium. From 1997 to 1999 he was Tellium's manager of optical crossconnect product development. He was a member of the team that launched Tellium's optical switching product line and established the feasibility of key technologies. He is a pioneer in the field of optical switching and networking, in which he is the holder of several patents. He is also a co-author of the book *Multiwavelength Optical Networks: A Layered Approach*. Prior to joining Tellium, he was a senior scientist in the Optical Networking Group of Bellcore, a research and telecommunications services provider, from 1992 to 1997. While at Bellcore he was responsible for local exchange optical networking architecture development and analysis, and received the Bellcore President's Award for outstanding achievement. He holds a B.S. in electrical engineering from Bombay University and a Ph.D. in electrical engineering from Columbia University. His thesis at Columbia was one of the first in the area of optical wavelength routing. He also served as chairman of the signaling working group of the OIF.