INVITED PAPER

Optical Network Management and Control

This article discusses optical network management, control, and operation from the point of view of a large telecommunication carrier.

By Robert D. Doverspike, *Fellow IEEE*, and Jennifer Yates, *Member IEEE*

ABSTRACT | While dense wavelength division multiplexing equipment has been deployed in networks of major telecommunications carriers for over a decade, the capabilities of its networking and associated network control and management have not caught up to those of digital cross-connect systems and packet-switched counterparts in higher layer networks. We shed light on this situation by examining the current structure of the optical layer, its relationship to other network technology layers, and current network management and control implementations. We provide additional insight by explaining how a combination of business and technical perspectives has driven evolution of the optical layer. We conclude by exploring activities to close this gap in the future.

KEYWORDS | Network control; network layers optical layer; network management

NOMENCLATURE

B-DCS	Broadband digital cross-connect system.
BoD	Bandwidth on demand.
CCAMP	Common control and measurement plane.
CMIP	Common management information protocol.
CLI	Command line interface.
CMISE	Common management information service.
CO	Central office.
CORBA	Common object request broker architecture.
DARPA	Defense Advanced Research Projects Agency.
DCS	Digital cross-connect system.
DWDM	Dense wavelength division multiplexing.
EMS	Element management system.
E-NNI	External network-to-network interface.
EVC	Ethernet virtual circuit.
FEC	Forward error correction.
FEC	Forwarding equivalence class (used in MPLS).
FXC	Fiber cross connect.
Gb/s	Gigabits per second.
IETF	Internet Engineering Task Force.
GMPLS	Generalized multiprotocol label switching.
GUI	Graphical user interface.
IOS	Intelligent optical switch.
ITU-T	International Telecommunication Union-Telecommunication Standardization Sector.
MIB	Management information base.
MPLS	Multiprotocol label switching.
MPLS-TE	MPLS-traffic engineering.
Muxponder	Multiplexer + transponder.
NE	Network element.
NMS	Network management system.
OIF	Optical Internetworking Forum.
OMS	Optical mesh service.
OSPF	Open shortest path first.
OSS	Operations support system.
OT	Optical Transponder
OTN	Optical transport network.
PCE	Path computation element.
PMD	Polarization mode dispersion.
QPSK	Quadrature phase shift keying.
REN	Research and education network.
ROADM	Reconfigurable optical add/drop multiplexer.
SNMP	Simple network management protocol.
SONET	Synchronous Optical NETwork.
SRLG	Shared risk link group.
TDM	Time division multiplexing.
TL1	Transaction language 1.

Manuscript received July 21, 2011; revised November 24, 2011 and December 26, 2011; accepted December 27, 2011.
R. D. Doverspike is with AT&T Labs Research, Middletown, NJ 07932 USA (e-mail: rdd@research.att.com).
J. Yates is with AT&T Labs Research, Florham Park, NJ 07932 USA (e-mail: jyates@research.att.com).

Digital Object Identifier: 10.1109/JPROC.2011.2182169

W-DCS Wideband digital cross-connect system.
XML Extensible markup language.

I. INTRODUCTION

The phrase "optical network management and control" cuts a broad swath in the telecommunications industry; consequently, our first task is to clearly define the bounds of this paper. First, the term *optical* itself tends to be used very broadly. For example, a popular interpretation is to classify any equipment with an optical interface as "optical equipment." This broader definition would include a large class of equipment that supports electrical-based cross-connection, such as SONET/SDH DCSs. In fact, today, because of the rapid evolution of small form optics, virtually all telecommunications equipment can support optical interfaces. Therefore, in this paper, we will confine ourselves to a more strictly defined *optical layer*, which consists of DWDM equipment and its supporting fiber network. We define this more precisely later.

Second, network management and control is addressed in a broad range of bodies, such as standards organizations, forums, research collaborations, conferences, and journals. The choice of network management and control strategy will vary for each telecommunications carrier (*carrier* for short) depending on its needs and, for a large network carrier, will not be exclusively dependent on optical network management choices developed in these bodies. Therefore, rather than venture into these much broader areas, we focus on a realistic context within which the optical layer is structured and operated in today's large telecommunications carriers. However, in the last sections, we briefly discuss potential future impact of key standards and ideas. Critical to this context are two concepts: network layering and restoration. In large telecommunications carriers, the optical layer is a slave to its higher layer networks. For example, virtually all demand for optical-layer connections comes from links of higher layer (overlay) networks. This relationship between the layers is intrinsically coupled and depends heavily on which layers provide restoration.

To aid in this understanding, we include historical perspectives of how the optical layer evolved to its present configuration. Perhaps most importantly, we include a discussion of the business context, which is important to explain the tradeoffs and priorities that led to the current implementations of network management and control. Finally, once we have described the current state of the optical layer, we will discuss R&D activities for the future evolution of the optical layer and its network control and management.

Section II provides background on the context within which the optical layer operates. Section III discusses the evolution and structure of today's optical layer. Section IV branches into today's network management and control. Section V explores current research into evolution of the optical layer, including our assessment of its most likely evolution path.

II. NETWORK SEGMENTS AND LAYERS

A. Network Segments

Fig. 1 illustrates how we conceptually segment a large national terrestrial network. Large telecommunications carriers are organized into metropolitan (*metro*) areas and

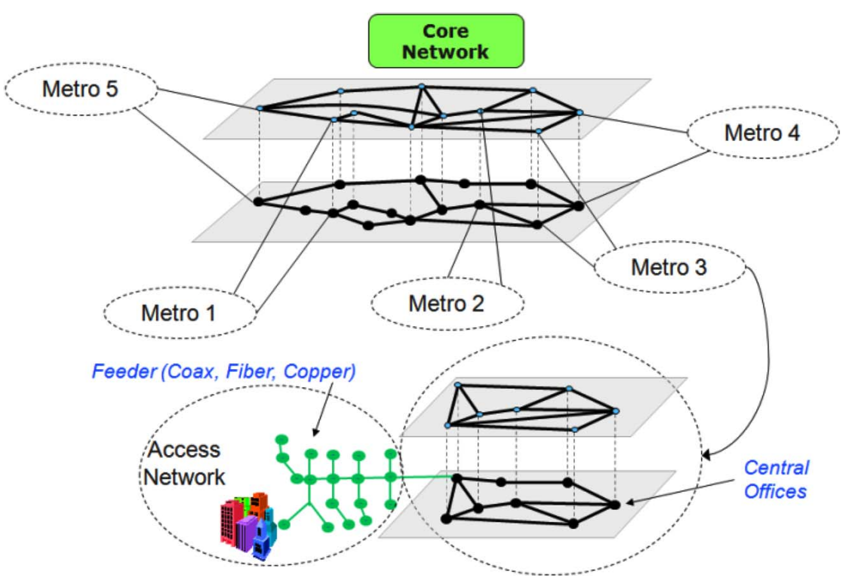

Fig. 1. *Terrestrial network layers and segmentation.*

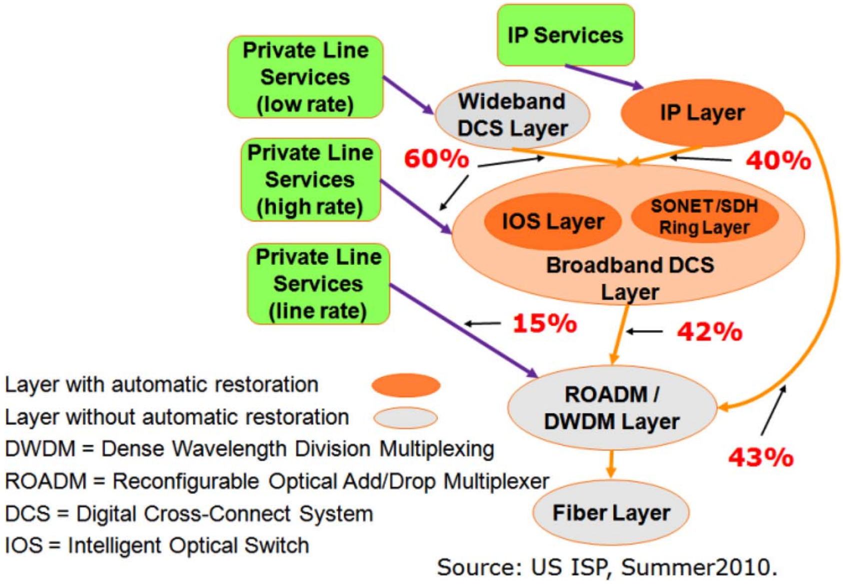

Fig. 2. *Simplified depiction of core-segment network layers.*

place the majority of their equipment in buildings called COs. Almost all COs today are interconnected by optical fiber. The *access* segment of the network refers to the portion between a customer location and its first (serving) CO. Note that the term "customer" could include another carrier. The *core* segment interconnects metro segments. Networks are further organized into network layers that consist of *nodes* (switching or cross-connect equipment) and *links* (logical adjacencies between the equipment), which we can visually depict as network graphs vertically stacked on top of one another. Links (capacity) of a higher layer network are provided as point-to-point demands (also called traffic, *connections*, or *circuits*, depending on the layer) in lower layer networks. See [10] and [11] for more details about the networking and business context of this segmentation.

B. Network Layers

Fig. 2 (borrowed from [16]) is a depiction of the core network layers of a large carrier. It consists of two major types of core services: IP (or colloquially, Internet) and *private line*. IP services are provided by the IP layer (typically routers) while private line services are provided through three different circuit-switched layers: 1) a W-DCS layer for low rate private line services (1.5 Mb/s); 2) a B-DCS layer for intermediate rate private line services (45–622 Mb/s), which in turn is composed of the IOS layer (technically an intelligent broadband DCS layer) and/or the SONET ring layer; and 3) the ROADM layer for high rate private line services (generally, 2.5 Gb/s and up).

Space does not permit us to describe these layers and technologies in detail. We refer the reader to [10] and [14] for background. As one observes, characterizing the traffic and use of the optical layer is not simple because virtually all of its circuits transport links of higher layer networks. In large carriers, many of these higher layer networks are owned by (internal to) the carrier, as shown in Fig. 2. Furthermore, the highest rate (line rate) private line services that route directly onto the optical layer usually emanate from links of packet networks of other carriers or large business customers who transport these links by leasing circuits (private lines). For example, many small regional carriers (usually subsidized by government or academia) called RENs lease private lines to interconnect their switches or computers. A key takeaway is that the design characteristics of packet networks drive most of the management and control of the optical layer. We return to this important observation in Section IV.

As expressed earlier, many in the industry sweep up the equipment that constitutes the nodes of the upper layer networks of Fig. 2 (such as DCSs) into a broader definition of "optical" equipment. We do not attempt to cover network management and control for all these different types of equipment in this paper. Instead, we focus the definition of optical layer to include legacy point-to-point DWDM systems and newer ROADMs, plus the fiber layer over which they route. We note that because of the ability to concentrate technology today, many vendors enable combinations of these different technology layers into different plug-in slots of the same "box" (e.g., a DWDM optical transponder on a router platform). Although we could address each of these combinations, for simplicity we will restrict the above definition to standalone optical-layer equipment. Furthermore, we concentrate on the core segment of the network; however, we provide a brief discussion of the metro segment later.

III. EVOLUTION AND STRUCTURE OF TODAY'S OPTICAL LAYER

A. Early DWDM Equipment

DWDM equipment was first deployed to relieve fiber exhaust in core carrier networks in the mid-1990s. Much of this work was pioneered by researchers at Bell Labs (e.g., see [25]). The first DWDM equipment was deployed with *optical transponders* (or simply called *transponders*) to support some pre-SONET interfaces, but soon after mostly supported SONET and SDH. The first DWDM equipment were configured in *point-to-point* (or *linear*) configurations. That is, client signals enter the transponder at a DWDM terminal (say, location A) via a standard intraoffice wavelength (typically 1.3 μm). The optical signal is regenerated, that is, detected, converted to electronic form, and transmitted by a laser at a fixed wavelength defined by a channel-grid (usually in the 1.55-μm range), and then, using a form of wavelength grating, multiplexed with other signals at different wavelengths into a multiwavelength signal over an optical fiber. Terminal and intermediate optical amplifiers are used to transport the multiplexed signal as far as possible, yet still meet signal quality requirements for all constituent channels. At a matching DWDM terminal at the far end (location Z), the process is reversed, where the line signal is finally demultiplexed into its constituent channels and signals.

The incoming (demultiplexed) signal on each channel at location Z is received by its associated transponder and then transmitted to its client interface at the intraoffice wavelength. A similar set of equipment and process occurs in the reverse direction of transmission (from Z to A). Generally, in carrier-based networks, the two-way signals are grouped into side-by-side ports on an interface card. All signals entering the DWDM terminal at A and Z are multiplexed or demultiplexed together. These early point-to-point systems have no intermediate add/drop, enabled 4–16 wavelengths per fiber and sometimes had their shelves organized consistent with service and protection interfaces of SONET/SDH linear systems or rings. In core networks using mesh restoration, the service and protection halves of these DWDM systems tended to be used in a standalone mode.

B. Reconfigurable Optical Add/Drop Multiplexer (ROADM)

Today legacy point-to-point DWDM systems still carry older circuits and sometimes are used for segments of new circuit orders, especially lower rate circuits. However, most large carriers now augment their optical layer with ROADMs. In contrast to a point-to-point DWDM system, a ROADM can interface multiple fiber directions (or *degrees*). This has encouraged the development of more flexibly tuned transponders (called *nondirectional* or *steerable*) and the ability to perform a remotely controlled optical cross connect (e.g., "through" *wavelength-selective cross connects*). See [14] and [31]. A ROADM can optically (i.e., without electrical conversion) cross connect the constituent signals from two different fiber directions without fully demultiplexing the aggregate signal (assuming they have the same wavelength). This is called a *transit* or *through* cross connection. Or, it can cross connect a constituent signal from a fiber direction to an end transponder, called an *add/drop* cross connection. All ROADM vendors provide a CLI for communication with a ROADM and an EMS that enables communication with a group of ROADMs. These network management and control systems are used to allow personnel to perform optical cross connects. Thus, because of the ability to remotely cross connect wavelengths, ROADMs begin to add connection management features more akin to DCS equipment in upper layer networks.

C. Provisioning in Today's Optical Layer

Before we discuss the network management and control of optical-layer networks, it is helpful to understand today's optical *circuit provisioning process* in large carrier networks. While the circuit provisioning process is more highly automated in the higher layer networks, it is a combination of automated and manual steps in the optical layer. First, we give a few preliminaries. The fiber interconnections between equipment within a single CO use fiber patch cords that are organized via an *optical patch panel*. For example, when installation personnel install a high-speed card or plug-in in an IP-layer router, they usually fiber its ports to ports on the patch panel. They do a similar procedure when installing a ROADM transponder. At some point during circuit provisioning, an order is issued to cross connect the router ports to the (client) ports of a transponder. Possibly the same personnel perform this request by manually fibering jumpers between the appropriate ports on the patch panel itself. We note that there exists a type of automated patch panel, which we call an FXC. See [14]. If an FXC is deployed, then the installation personnel must still fiber the transponder ports and client equipment to the FXC, but when the provisioning order is given, the FXC can cross connect its ports under remote control. However, today, there are few FXCs deployed in large carriers; therefore, in this section, we will assume the patch panel dominates, but return to the FXC in our last section.

We list four broad categories of provisioning steps in the core segment. In many cases, a circuit order may require steps from all four categories.

1) Manual: installation personnel visit CO, install cards and plug-ins, and fiber them to the patch panel.
2) Manual: installation personnel visit CO and cross connect ports via the patch panel.
3) Semiautomated: provisioners request optical cross connects via a CLI or EMS.
4) Fully automated: an OSS is fed a circuit path from a network planner or planning tool and then

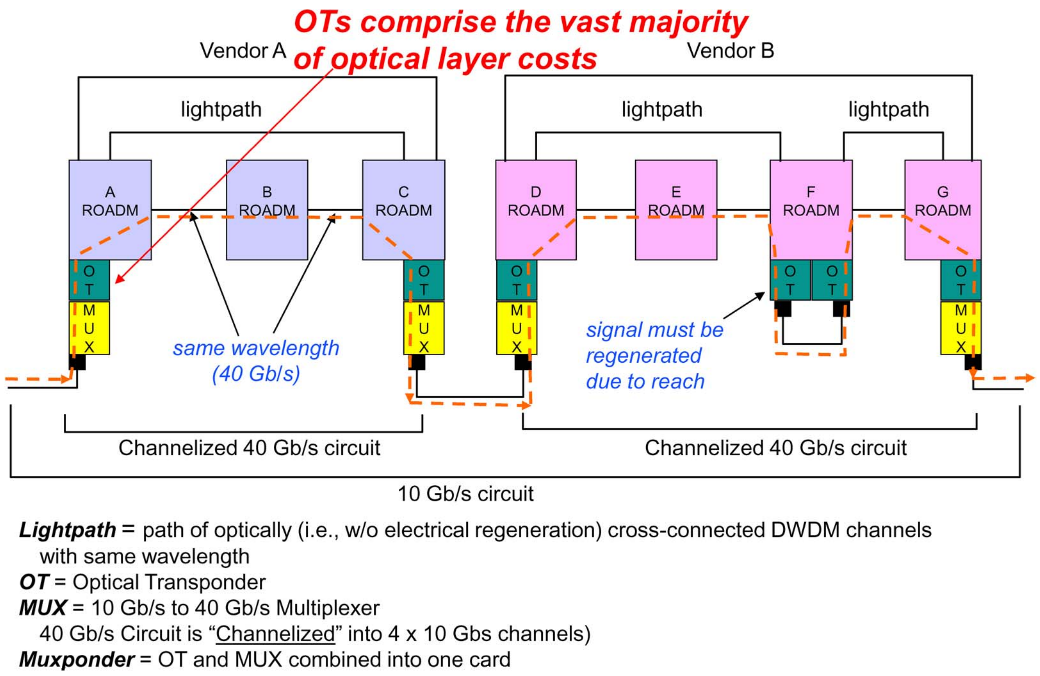

Fig. 3. ***Path of 10-Gb/s circuit over two 40-Gb/s circuits.***

automatically sends optical cross-connect commands to the CLI or EMS.

Carriers are mostly doing category 3) today.

Fig. 3 depicts a realistic example within the optical layer of Fig. 2, where a 10-Gb/s circuit is provisioned between ROADMs A-G. For example, this circuit might transport a higher layer link between two routers which generate the client signals at ROADMs A and G. There are two *vendor subnetworks* in this example, where a vendor subnetwork is defined to be the topology of vendor ROADMs (nodes) from a given equipment vendor plus their interconnecting links (fibers). This is also called a *domain* in many standards organizations. A *lightpath* is a path of optically cross-connected DWDM channels, i.e., with no intermediate optical–electrical–optical (OEO) conversion. Because DWDM systems from different vendors do not generally support a handoff (interface) between lightpaths, for a circuit to cross vendor subnetworks requires add/dropping through transponders. The ROADMs in this example support 40-Gb/s channels/wavelengths. Another complicating factor in today's networks is the evolution of the top signal rate over the years. In this example, we need to multiplex the 10-Gb/s circuit into the 40-Gb/s wavelengths. DWDM equipment vendors provide a combo card, colloquially dubbed a *muxponder*, which provides both TDM (dubbed "mux" in Fig. 3) and transponder functionality.

To provision our example 10-Gb/s circuit, we must first provision two 40-Gb/s *channelized* circuits (i.e., they provide 4 × 10-Gb/s subchannels), one in each subnetwork (A-C and D-G). Furthermore, because of optical reach limitations, the 40-Gb/s circuit must demultiplex at F and thus traverse two lightpaths in the second subnetwork. This requires interconnection between the ports of the two transponders at ROADM F. This process is accomplished by a combination of steps from the four categories mentioned above. To illustrate, once the cards and ports are installed [category 1)], a step of category 2) is required at ROADM F. The optical cross connects between A-B-C, D-E-F, and F-G are steps of category 3) [or 4)]. Once the two 40-Gb/s channelized circuits are brought into service, two 10-Gb/s circuits are provisioned (A-C and another D-G), which can be done by a step of category 3) [or 4)]. Finally, the client signal is interconnected to the muxponders at A and G [category 2)] and the two subnetwork circuits are interconnected via the muxponder ports at C and D [category 2)]. Note that, strictly speaking, this example uses a mixture of three different types of cross-connect technology: manual fibering (e.g., at node F), remote controlled optical cross connect (e.g., at node B), and electrical TDM (e.g., assigning the 10-Gb/s circuit to a channel of the channelized 40-Gb/s circuit at A). Such is the nature of today's optical layer.

Effectively, the above implies the optical layer itself consists of multiple sublayers, each with routing procedures and provisioning processes. Fig. 4 shows an example of five layers to support the provisioning of two 10-Gb/s circuits. In fact, many optical-layer networks support a 2.5-Gb/s muxponder, for which we must add yet another sublayer. An interesting observation from Fig. 4 is that because of the logical links created at each layer, sometimes

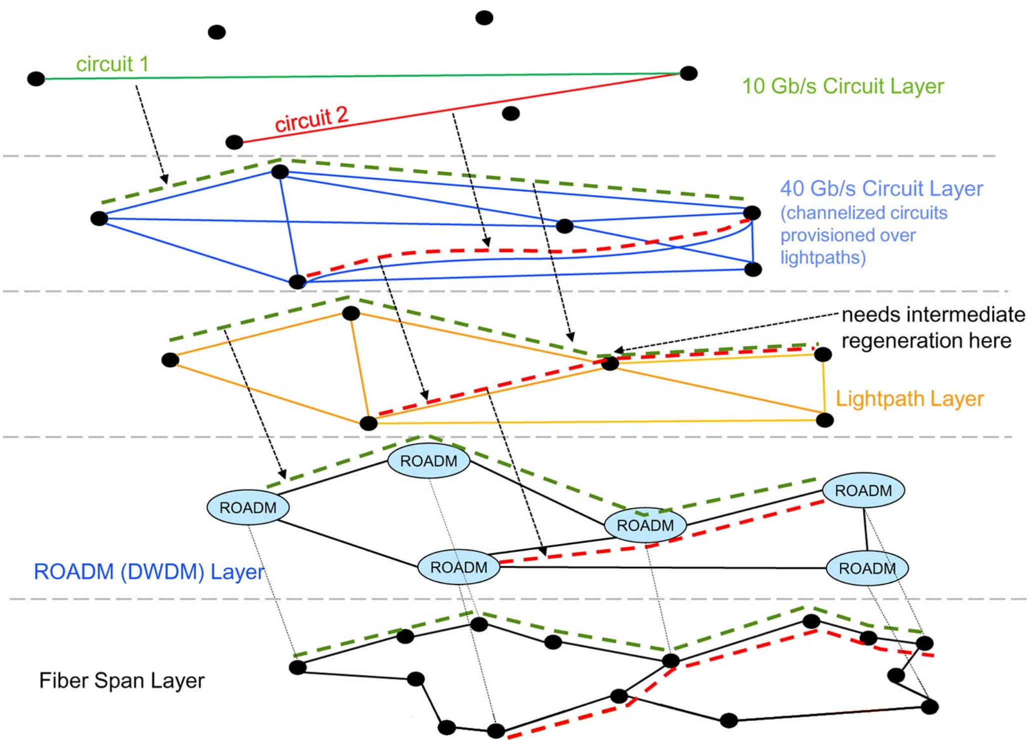

Fig. 4. *Sublayering within optical layer.*

links at a given layer appear to be diversely routed, when in fact they converge over segments of lower layer networks. We discuss this very important point in Section IV.

IV. MANAGEMENT AND CONTROL IN TODAY'S OPTICAL LAYER

The ITU-T has defined various areas of network management. Here, we will confine ourselves to the principal areas of *configuration* management (installing or removing equipment, making their settings, and bringing them in or out of service), *connection* management (effecting cross connects to enable end-to-end connections or circuits), and *fault* management (reporting and analyzing outages and quality of signal). The area of performance management is also relevant, but applies more to packet networks; therefore, here for simplicity we will lump relevant aspects of optical performance management into the area of fault management. In the previous section, we discussed provisioning, which is a combination of configuration management and connection management.

A. Legacy DWDM Systems

Clearly, the control plane and network management capabilities of early DWDM systems were simple or nonexistent. Although there were hybrid systems that also contained cards with electrical fabrics, they had no optical cross-connect fabrics and therefore no purely optical connection management functionality.

Thus, configuration management and fault management were the predominant network management functionalities provided in early systems. Virtually all the fault management (alarms) of these systems are based on SONET/SDH protocols from the client signals. The few exceptions are alarms for amplifier failures, which are based mostly on loss of power (DB attenuation). Also, instead of providing sophisticated and automatic optical signal analysis features, because the DWDM links were usually coupled with SONET rings or linear systems with inline protection, maintenance personnel could put the constituent SONET rings or chains into protection mode and then put test analyzers on the DWDM signal.

Legacy point-to-point DWDM systems were generally installed with simple text-based network management interfaces and a standardized protocol. An example is Bellcore's TL1 [2]. TL1 enabled a simple interface to an OSS. The SONET/SDH standard specifies fault management associated with the client signals, such as alarms and performance monitoring. However, for DWDM systems, there is usually an internal communications interface, usually provided over a low rate sideband wavelength (channel). Besides enabling communication between the NEs, this channel is used to communicate with the inline amplifiers. The protocol over the internal communications channel is proprietary.

B. ROADMs

A few EMSs (even sometimes just one) are often used to control the entire vendor subnetwork, even if the network is scattered over many different geographical regions. Even though the ROADMs have a CLI, most carriers prefer to interface to the ROADM via the EMS because of the more sophisticated GUI and tailored visualization of ROADM settings and state. Furthermore, the EMS provides an interface to an OSS, typically called a *northbound* interface using protocols such as CMISE, SNMP [3], CORBA, or XML [36]. Also of interest is that many EMSs use TL1 for their internal protocol with their NEs because it simplifies the implementation of an external TL1 network management interface for those carriers who require it. Most ROADMs today internally use the OTN signal standard for setting up subnetwork circuits. Firmware or software in the transponders is used to encapsulate client signals of different types (e.g., SONET, SDH, Ethernet, Fibre Channel) into the internal OTN signal rates. We will cover OTN more in Section V.

Today there is a wide variation in capability across different ROADM EMSs. Some EMSs can automatically route and cross connect a circuit between a pair of specified transponder ports. Here, the EMS chooses the links and the wavelength, sends cross-connect commands to the individual NEs, monitors status of the circuit request, and reports completion to the northbound interface. Other EMSs operate only on a single NE basis.

In contrast to upper layers networks, signal quality complicates the optical layer. For example, provisioning a new circuit requires tuning the transponder laser, balancing power in the amplifiers, and other *settling* of the signal. Furthermore, as show in Figs. 3 and 4, optical reach is an important issue and sometimes intermediate regeneration is needed to support a circuit. Because computing optical reach is a very complicated optical problem and is dependent on specific, proprietary vendor technology, most vendors also produce a coordinated NMS. The NMS has two main functions: 1) assist planners in the engineering aspects of building or augmenting vendor ROADM subnetworks over existing fibers and locations and 2) simulate the paths of circuits over a deployed vendor subnetwork, taking into account requirements for signal quality. As the reader may have quickly surmised, this requires that for every circuit request, the provisioner must consult an NMS for each segment of the path that crosses a vendor subnetwork. For example, say a carrier installs vendor-A DWDM equipment for regional transport (connecting smaller groups of metro areas) and vendor-B DWDM equipment for long haul (between major cities). Thus, even with just two vendors, many circuits whose endpoints are in smaller metros will route through three segments corresponding to vendor subnetworks A-B-A.

Armed with the path, wavelength, and regeneration information produced by the NMS for each segment, the provisioner then enters the request into a provisioning OSS. The OSS produces an order document (form) for each equipment installation and cross-connect specification, segment by segment. The disposition of each cross connect then depends on its step category defined in the previous section: category 2) is sent to a workforce management organization, category 3) is sent to a provisioning center whose personnel enter commands to the EMS or CLI, and category 4) step is automatically sent to the northbound interface of the appropriate EMS.

Not surprisingly, the time today required to provision a circuit in the optical layer can be long. To summarize the reasons:

1) the NMS/EMS interaction can be laborious;
2) there may be no flow through from OSS to EMS (via northbound interface);
3) many portions of the circuit order require manual steps, such as manual cross connection (patch panel) due to intermediate regeneration or crossing of vendor subnetworks;
4) even with semiautomated or fully automated cross connection (which is an order of magnitude faster than above), optical signal settling times can be long compared to cross-connect speeds in higher layer networks.

We will discuss some of the business context that led to this evolution in Section V.

Finally, fault management is similar to that of the point-to-point DWDM system, except that all newer ROADM internally use OTN encapsulation of the circuits and, as a result, the alarms identify affected slots and ports in terms of the OTN termination-point information models and alarm specifications. Other alarm specifications are used for the client side of the optical transponder (e.g., SONET, SDH, Ethernet).

C. Integrated Interlayer Network Management

We revisit two of the key network characteristics highlighted in the introduction, namely network layering and restoration. Because today restoration is typically performed at higher layer networks, outages that originate at lower layers are more difficult to diagnose and respond. For example, an outage or performance degradation of a DWDM amplifier or a fiber cut can sometimes affect ten or more links in the IP layer, while the failure of an intermediate tranponder may affect only one IP-layer link and be hard to differentiate from outage of an individual router port. Thus, the most effective approach to network management must model the complex relationship of the layers.

IP backbones have traditionally relied on IP-layer reconvergence mechanisms, (generally called internal gateway protocols), such as OSPF [20] or more explicit restoration protocols such as MPLS fast reroute and MPLS-TE [21]. All of these protocols have been designed and standardized within the IETF.

Why do IP backbones usually rely on IP-layer reconvergence instead of lower layer restoration? The answer lies in the historical reliability of router hardware, protocols, and required maintenance procedures, such as software upgrades. As a consequence, to achieve sufficient network availability, IP backbones were typically designed with sufficient spare capacity to restore the network from the potential outage of an entire router, whether due to hardware/software failure or maintenance activity. Therefore, the majority of fiber outages and other optical-layer failures can be restored without significant additional capacity beyond that required for the potential (single) router outages. However, effectively planning this capacity requires detailed knowledge of the lower layer outage modes—how all the IP links are routed over DWDM systems, fibers, etc. The industry models these relationships via a generic concept called the SRLG. Restoration capacity planning then involves detailed analysis of all of the potential SRLG outages and appropriate capacity allocations to achieve the desired target for network availability.

Most large routers today provide the ability to "bundle" multiple physical link (interfaces) between adjacent routers into one "logical" link, which is then advertised as one link by the interior gateway protocol. With IP routing protocols that do not take into account link capacity (e.g., OSPF—but note a capacity-sensitive version called OSPF-TE has been defined), losing a significant number of component links of a link bundle (but not all), would normally result in the normal traffic load on this link being carried on the remaining capacity, potentially leading to significant congestion. How can this happen? Because of the multiple layering, as the link bundle grows over time (by adding additional links), it is possible that some links in the bundle are routed over different optical-layer paths than others. In recent years, router technologies have been adapted to handle such scenarios, shutting down the remaining capacity in the event that the link capacity drops below a certain threshold. However, determining what that threshold should be across all possible failure scenarios, and then ensuring sufficient capacity elsewhere in the network is complicated.

Routers will detect outages which occur anywhere on a link, be it due to a port outage of the router at the remote end of the link, an optical amplifier failure, or fiber cut. The router cannot readily distinguish—however, it will reroute traffic accordingly and generate traps to inform operations personnel. However, the IP and optical layers are typically managed by very distinct work groups or even via an external carrier (e.g., leased private line). In the event of an optical-layer outage, the alarm notifications would also be created to the optical maintenance work groups. Thus, without sophisticated alarm correlation mechanisms between the events from the two different layers, there can be significant duplication of troubleshooting activities across the two work groups. Efficient correlation of alarms generated by the two different layers can ensure that both work groups are rapidly informed of the issue, but that only the optical-layer group need necessarily respond as they would need to activate the necessary repair. See [34] for a more in-depth discussion of this approach.

D. Metro Segment

In contrast to the core segment, metro networks have considerably smaller geographical diameter. Also, many carriers use a single DWDM vendor in a given metro area. Thus, intervendor (domain) routing and intermediate regeneration are often not issues. On the other hand, in contrast to the core segment, ROADMs usually are installed in only a portion of the COs of a large metro. Thus, a circuit path can involve complex access provisioning on distribution/feeder fiber followed by long sequences of patch panel cross connects in COs. These hurdles have blunted the business driver for more automatic connection management in the optical layer of metro areas. For example, if a circuit requires 15 manual cross connects over direct fiber and only one section of automated cross connection over ROADMs, it is hard to prove the business case for the ROADM segment since overall cost is not highly impacted. Length constraints prevent us from delving into more detailed metro issues.

V. FUTURE EVOLUTION OF THE OPTICAL LAYER

Armed with an understanding of the current environment of the optical layer in the core network segment, we are now prepared to discuss potential paths forward for network management and control. However, requoting from the introduction, a wide range of network management protocols exists and a large carrier's choice is based on its individual needs. To avoid a lengthy discussion on the various management protocols and their specifics, we will provide a general perspective and summarize the salient observations from the previous sections, along with business perspectives.

A. Network Control and Management Gap

We summarize the following observations about the optical layer in today's carrier environment.

1) The optical layer can require many manual steps to provision a circuit, such as NMS/EMS circuit design coordination, crossing vendor subnetworks, and intermediate regeneration because of optical reach limitations.
2) Even the fully automated portions of provisioning an optical-layer circuit are significantly slower than its higher layer counterparts.
3) Evolution of the optical layer has been heavily motivated to reduce costs for interfaces to upper layer switches. This has resulted in a simple focus to increase "rate and reach."
4) Restoration is provided via higher network layers and, thus, planning, network management, and

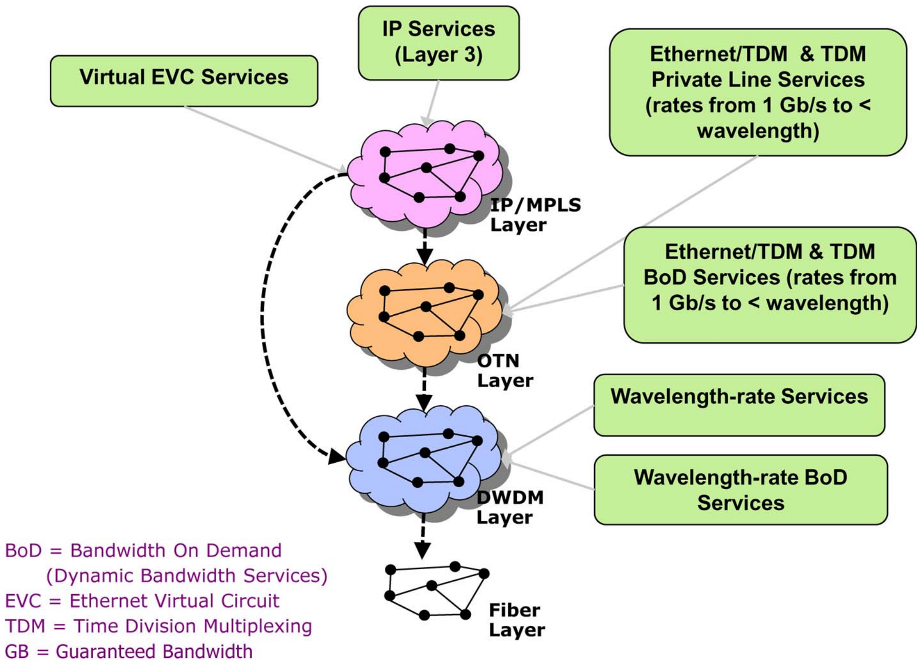

Fig. 5. *Potential future core network architecture.*

restoration must work in a more integrated fashion across the layers.

5) No large-scaled dynamic services have been implemented that would require rapid connection management in the optical layer.

Given observations 3)–5), it has been hard to justify a business case to evolve optical-layer technology and network management capabilities to enable provisioning times akin to those of DCS layers or even faster (flow routing) via MPLS tunnels in routers. In fact, glancing again at Fig. 2, we notice that except for the very highest rate private line services (which only consume a small portion of optical-layer capacity), the optical layer is basically a slave to the other internal upper layers, notably the IP layer, which historically has been the most rapidly growing layer. Thus, demand for the optical layer (from links of higher layer networks) is not akin to phone calls or web access requests, but results from a slower network design process. Furthermore, we observe that one of the main historical business drivers for evolution of the optical layer has been to support cost reduction of the interfaces on IP-layer routers, which have followed a steady improvement from economy of scale for well over a decade. This has resulted in a simple focus (some might say a "frenzy") to increase "rate and reach" in DWDM equipment.

As a result of all these observations, a gap has formed between the network management and operations of today's optical layer and the dynamic and automatic nature of its higher layer networks. Up until now, many in the industry have ignored this gap or assumed it would be bridged soon, yet, this gap has persisted for over a decade. This persists because, as we have pointed out, optical-layer evolution is not only influenced by technology evolution, but business perspectives, as well. For example if, in contrast to observation 5), demand for a high-volume, rapid, and dynamic optical-layer connection service had manifested, then carriers would have proved this in their internal business cases and this gap would have been bridged much more quickly.

B. Technology Evolution of the Optical Layer

Optical and WDM transport technology has undergone impressive technological advancement in the past 15 years. As previously described, DWDM technology started with a few wavelengths, low bit rates, and limited point-to-point networking. Today, ROADM systems are being deployed with rates of 100 Gb/s, 80 wavelengths, and lightpaths with 1000–1500-km reach. This has been enabled by technologies such as coherent detection (very high rate signal processing that allows more sophisticated detection of different optical pulses) and various forms of QPSK (enables a larger set of symbols by varying characteristics of the optical pulse). Besides rate and reach improvements, coherent detection dispels many previously awkward or expensive methods to overcome optical impairments, such as PMD and thus enables transport over a wider variety of fiber types. See [15] and [33].

If we examine [16], we find that the historical explosive growth of intercity IP traffic is leveling off. Also, the economy of scale for higher rate packet-switch interfaces is flattening. Thus, the principal drivers for higher "rate" wavelengths will not be as intense as in the past. The top-rate interface on packet switches has steadily evolved in steps, e.g., 155 Mb/s, 622 Mb/s, 2.5 Gb/s, 10 Gb/s, 40 Gb/s, and 100 Gb/s. DWDM channel rates have matched. The long-term effect is that just as we maximized the reach at a given wavelength rate, up popped the need for the next higher router interface rate and then its associated optical reach decreased. This suggests that as the frenzy for increased maximum rate quells, the need for intermediate regeneration should eventually mitigate.

We note that one side effect of the newer coherent detection technologies is that lightpath settling times have increased, which contributes to the network management gap. This is another example of business context driving the current network management and control environment: namely, driving down interface costs (both IP layer and optical layer) was deemed a greater priority than decreasing provisioning times.

C. Advent of the OTN Layer

As SONET and SDH have run out of gas, the OTN technology has emerged [17]. The OTN protocol stack was originally proposed to standardize the overhead channels and use of forward error correction (FEC) in optical networks. This was a key technology advancement to enable the evolution of rate and reach mentioned above. Since then, it has evolved into a multiplexing hierarchy, an internal transport protocol for DWDM, and container/encapsulation mechanism for different signal formats. Therefore, similar to how DCSs evolved to automatically cross connect lower rate channels among higher rate SONET or SDH interfaces, the OTN switch is a form of DCS that has recently emerged to cross connect lower rate channels among higher rate interfaces. However, another business question has emerged: If *OTN switches* provide all the network management functionality (and more) of their previous DCS counterparts, what is the motivation to bridge the optical-layer management and control gap?

Fig. 5 shows potential, future core architecture. In this architecture, lower rate private line services have migrated to EVC services in the IP/MPLS layer. Private line services at 1 Gb/s or higher route over the OTN layer, whose lowest signal rate is 1.2 Gb/s. Private line service at the highest rate routes directly over the ROADM layer. Note that the links of the IP layer have the option of routing over the OTN layer or directly onto the DWDM layer. This option is discussed more in the next section.

D. Advanced Network Management and Control Capabilities

In Fig. 5, note that we divide private line traffic into two categories: traditional and BoD. Although BoD has been a popular study and topic of publication for years, few carriers have implemented full-fledged services for DCS layers, let alone the optical layer, as we noted in observation 5) in Section V-A. For example, the authors of this paper pioneered AT&T's OMS from its first proof of concept (in early 2000s) up until its service launch in 2005, which was, at the time, one of the first truly long-distance high-rate BoD services. See [9] and [30]. However, adhering to the narrower definitions of this paper, we note that although OMS uses the term "optical," it is actually provided by the IOS layer. As mentioned previously, the IOS layer is an intelligent broadband DCS layer. However, of relevance here, OMS was enabled because of the sophisticated network management and control capabilities of the IOS layer. Once a customer has his customer premise equipment connected via the access/metro segments (a "pipe") to the IOS in the core CO, he/she can set up circuits on-demand between any of his interfaces at the various locations, up to the pipe capacity. Furthermore, the IOS layer provides extra channels for restoration and therefore the extra capacity needed for BoD demand can share the restoration channels, which is key to its successful business case.

Clearly, given the previous description of the today's optical layer, extending BoD to the optical layer is more challenging, both from technical and business contexts. We cannot fully cover the publications addressing optical-layer BoD, but note that CORONET [7] is a project that addresses this problem and is sponsored by DARPA. The principal goals of CORONET are a dynamic core optical layer, wherein circuits can be rapidly provisioned under a highly distributed control plane. CORONET Phase I addressed network architecture, protocols, and design [5], [6]. While the OTN switch was not defined at the beginning of Phase I, as of the writing of this paper, CORONET Phase II is underway and is addressing the role of the OTN layer and practical commercial implementation of these goals. Activities include realistic cost studies of different architectural alternatives for interrelationship of the layers in Fig. 5.

E. Methods for Fully Automated Provisioning

Putting aside business case justification for now, from the previous sections, we observe that if we want to advance the current state of the art in optical-layer network management and control to similar levels as its higher layer networks, then we must overcome the manual provisioning steps described earlier. We now describe a sequence of technologies and tools in the R&D phase to accomplish this feat. The most time-consuming manual steps [categories 1) and 2) in Section III-C] involve fiber interconnection. These steps arise from three major causes: 1) wiring of customer equipment (via metro/access segment) to the end transponders; 2) interconnection of circuits between vendor subnetworks; and 3) intermediate regeneration. Two key ideas to automate these steps are the use of the FXC, discussed earlier, and transponder pooling. Today, to limit costs, most carriers tend to install and interconnect transponders per individual circuit order, rather than installing and fibering sharable pools of transponders. See [12] and [4] for optimization algorithms for sizing and placing pools of transponders. Both of these concepts are key components of the CORONET project [32]. Beyond initial service provisioning, the ability to switch a circuit (via the FXC) to a spare transponder is also needed to enable rapid restoration: both to provision a circuit over an alternate restoration path that crosses two or more lightpaths and to perform "hitless" rerouting (normalization) of a circuit path after repair of an outage [35].

The next longest category of manual steps is the interactions of provisioning/planning personnel with the NMS and EMS. The main purpose of the NMS is to theoretically route (also called "design") a circuit over a path of lightpaths (including selection of spare wavelengths) and intermediate transponders (if needed) to ensure that adequate spare channel capacity exists and that signal quality is provided. As described previously, multiple vendor subnetworks greatly exacerbate delays in the provisioning process. The authors and collaborators have derived and implemented a process in AT&T's network to automate the NMS portion of the provisioning step. The key idea for this process is to request that each vendor NMS precalculate a *reachability* matrix which specifies the pairs of ROADMs between which lightpaths can be established (i.e., where no intermediate regeneration is needed), then, build a sophisticated network-wide optical-layer routing tool. The tool uses the reachability matrix to construct a graph of logical edges that represent where potential lightpaths in each vendor subnetwork can be created. Other edges are added to the graph to model the cost and how vendor subnetworks can be interconnected via transponders. Circuits are then routed over this augmented graph to minimize cost or achieve fiber-layer diversity objectives. Such a tool is described in [26].

Once the NMS interactions are automated, we next must turn attention to the manual interaction of the provisioner with the EMSs. This then brings up the question of which control plane protocol to use for the ROADMs and EMS. This issue is discussed in the next section.

F. Potential Impacts of Standards Organizations

The three standards organizations and their subgroups that influence the DWDM and optical-layer network management and control the most are the IETF, ITU, and the OIF. We briefly describe their efforts as they relate to the optical layer. However, much of the work of these organizations is directed at the DCS layers; reiterating our earlier definition, the "optical layer" in this paper is confined to DWDM equipment and its supporting fiber network. Therefore, although most major DWDM equipment manufacturers contribute to and attend these standards bodies, for the reasons described earlier there is still a major gap between the standards and deployment in DWDM equipment in carrier networks, especially for connection management (i.e., fully automated provisioning discussed in the previous section).

Of particular impact is ITU Study Group 15. This is because, as described earlier, most recently deployed DWDM equipment uses OTN for its multiplexing hierarchy, internal signal formatting, FEC, and data other communications. See ITU standards G.709 and G.798 [17]. Therefore, most optical vendors incorporate ITU fault management objects and specifications into their equipment models and internal MIB. These objects mostly manifest via alarms and notifications sent from the EMS to the northbound interface.

The most salient of the connection management control plane approaches for the optical layer is GMPLS [1], [8] derived in the IETF CCAMP working group [24]. However, some of the major issues identified earlier (e.g., manual cross connection, optical reachability limitations, and intervendor subnetworks) were not completely addressed by the original GMPLS signaling protocols. For example, optical routing and reachability issues are being addressed in the IETF PCE [23]. In addition, there are many research projects and proposals that address how, via standards bodies, to model impairments and incorporate their impact to reachability constraints in routing. For example, see [27]. Interdomain subnetwork communication is being addressed in the OIF via an E-NNI protocol [29]. Some advanced ideas for utilizing the emerging capabilities for nondirectional, colorless (tunable) transponders and beyond, such as dynamically changing the wavelength/channel spacing and rate, are explored in the EO-NET project [12].

While it is outside our focus in this paper to discuss all related standards, ideas, and proposals in the literature, we briefly discuss PCE because it may be well suited to the complex routing and provisioning problems mentioned above. A PCE is defined by the IETF (RFC 4655 [22]), as "An entity (component, application, or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints." For example, a PCE could communicate with different vendor subnetworks (or domains), could store and update reachability information associated with each subnetwork, compute complex capacity-sensitive wavelength assignment optimization, and interact with distributed, inter-NE provisioning protocols. For example, GMPLS includes signaling (RSVP-TE) and routing (OSPF-TE). PCE for the optical layer will support different models, such as PCE-based signaling or GMPLS-based signaling or a hybrid of both.

Another extremely complex need, which we illustrated in Fig. 4, that perhaps could be accomplished through PCE, is the ability to diversely route groups of connection requests which require an offline and graph-based knowledge of fiber-layer routes (i.e., an SRLG database that includes how upper layer links route over it). All these capabilities (and more) are needed in the core network of a large carrier. Currently, there are no standardized PCE implementations that have been implemented in large carrier networks. However, AT&T has implemented a planning system that incorporates all these mentioned features that planners use on a daily basis to route circuits (connections) through the DWDM layer [26]. AT&T is exploring the feasibility to extend this capability to a PCE implementation that interacts with potential standardized control planes of next-generation ROADMs with nondirectional, colorless (tunable), and (possibly) FXC-like capabilities.

G. Business Case for Optical-Layer Evolution

After over a decade of technical development, while optical-layer capacity, connectivity, cost improvements, and signal quality have enjoyed great advancement, optical management and control has evolved more slowly. We have shown this is clearly not due to lack of R&D, both in advanced network architectures and protocols. Thus, the next step in this evolution is to prepare a business case that will meet the economic criteria expected by network planning and finance organizations of large telecommunications carriers. Given the many demands for resources in a large telecommunications carrier, ideas such as FXCs, transponder pooling, faster circuit tuning/settling, better routing tools, and optical-layer restoration would most likely have to result in cost savings and/or revenue opportunities to be broadly adopted. The authors feel that most of these advances will eventually be implemented because of 1) the leveling of core IP traffic growth (and thus the lack of historically frenzied need for wavelength rate increase); 2) continued decline in transponder costs and prices; and 3) advancements in DWDM technologies. However, the key variable will be the rate of this implementation, which will hinge on the ability to prove the business cases. ■

Acknowledgment

The authors would like to thank D. Brungard for her guidance in optical standards and Dr. P. Magill for his guidance in optical systems evolution.

REFERENCES

[1] P. Ashwood-Smith, Y. Fan, A. Banerjee, J. Drake, J. Lang, L. Berger, G. Bernstein, K. Kompella, E. Mannie, B. Rajagopalan, D. Saha, Z. Tang, Y. Rekhter, and V. Sharma, *Generalized MPLS—Signaling Functional Description*, IETF Internet draft, Jun. 2000.

[2] Bellcore, *Operations Application Messages—Language for Operations Application Messages*, TR-NWT-000831. [Online]. Available: http://telecom-info.telcordia.com/

[3] U. Black, *Network Management Standards SNMP, CMIP, TMN, MIBs, and Object Libraries*, A. Bittner, Ed. New York: McGraw-Hill, 1995, ISBN: 007005570X.

[4] S. Chen, I. Ljubic, and S. Raghavan, "The regenerator location problem," *Networks.*, vol. 55, no. 3, 2010, pp. 205–220.

[5] A. Chiu, G. Choudhury, G. Clapp, R. Doverspike, J. W. Gannett, J. G. Klincewicz, G. Li, R. A. Skoog, J. Strand, A. Von Lehmen, and D. Xu, "Network design and architectures for highly dynamic next-generation IP-over-optical long distance networks," *J. Lightw. Technol.*, vol. 27, no. 12, pp. 1878–1890, Jun. 2009.

[6] A. Chiu, A. G. Choudhury, G. Clapp, R. Doverspike, M. Feuer, J. W. Gannett, J. Jackel, G. T. Kim, J. G. Klincewicz, T. J. Kwon, G. Li, P. Magill, J. M. Simmons, R. A. Skoog, J. Strand, A. Von Lehmen, B. J. Wilson, S. L. Woodward, and D. Xu, "Architectures and protocols for capacity-efficient, highly-dynamic and highly-resilient core networks," *J. Opt. Commun. Netw.*, vol. 4, no. 1, pp. 1–14, Jan. 2012.

[7] *DARPA CORONET Project*. [Online]. Available: http://www.darpa.mil/Our_Work/STO/Programs/Dynamic_Multi-Terabit_Core_Optical_Networks_(CORONET).aspx

[8] R. Doverspike and J. Yates, "Challenges for MPLS in optical network restoration," *IEEE Commun. Mag.*, vol. 39, no. 2, pp. 89–97, Feb. 2001.

[9] R. Doverspike and J. Yates, "Practical aspects of bandwidth-on-demand in optical networks Panel on Emerging Networks Service Provider Summit, Anaheim, CA, Mar. 2007.

[10] R. Doverspike and P. Magill, "Chapter 13 in Optical Fiber Telecommunications V B," in *Commercial Optical Networks, Overlay Networks and Services*. Amsterdam, The Netherlands: Elsevier, 2008.

[11] R. Doverspike, K. K. Ramakrishnan, and C. Chase, "Chapter 2 in Guide to Reliable Internet Services and Applications," in *Structural Overview of Commercial Long Distance IP Networks*, C. Kalmanek, S. Misra, and R. Yang, Eds. 1st ed. New York: Springer-Verlag, 2010.

[12] EO-NET, *Project on Elastic Optical Networks*. [Online]. Available: http://www.celtic-initiative.org/Projects/Celtic-projects/Call7/EO-Net/eonet-default.asp

[13] C. V. Saradhi, R. Fedrizzi, A. Zanardi, E. Salvadori, G. M. Galimberti, A. Tanzi, G. Martinelli, and O. Gerstel, "Traffic independent heuristics for regenerator site selection for providing any-to-any optical connectivity," in *Proc. Conf. Opt. Fiber Commun./Nat. Fiber Opt. Eng. Conf.*, Los Angeles, CA, Mar. 2010, pp. 1–3.

[14] M. D. Feuer, D. C. Kilper, and S. L. Woodward, "Chapter 8 in Optical Fiber Telecommunications V B," in *ROADMs and Their System Applications*. Amsterdam, The Netherlands: Elsevier, 2008.

[15] C. Fludger, T. Duthel, D. van den Borne, C. Schulien, E.-D. Schmidt, T. Wuth, J. Geyer, E. De Man, G.-D. Khoe, and H. de Waardt, "Coherent equalization and POLMUX-RZ-DQPSK for robust 100-GE transmission," *J. Lightw. Technol.*, vol. 26, no. 1, pp. 64–72, Jan. 2008.

[16] A. Gerber and R. Doverspike, "Traffic types and growth in backbone networks," in *Proc. Conf. Opt. Fiber Commun./Nat. Fiber Opt. Eng. Conf.*, Los Angeles, CA, Mar. 2011, pp. 1–3.

[17] S. Gorshe, *A Tutorial on ITU-T G.709 Optical Transport Networks (OTN) Technology White Paper*, 2010. [Online]. Available: http://www.elettronicanews.it/ 01NET/Photo_Library/775/PMC_OTN_pdf.pdf

[18] *International Telecommunication Union (ITU)—Sector T*. [Online]. Available: http://www.itu.int/ITU-T/

[19] *Internet Engineering Task Force (IETF)*. [Online]. Available: http://www.ietf.org/

[20] IETF RFC 2328, *OSPF Version 2*, Apr. 1998. [Online]. Available: http://www.ietf.org/rfc/rfc2328.txt

[21] IETF RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnel*, May 2005. [Online]. Available: http://www.ietf.org/rfc/rfc4090.txt

[22] IETF RFC 4655, *A Path Computation Element (PCE)-Based Architecture*, Aug. 2006. [Online]. Available: http://tools.ietf.org/html/rfc4655

[23] IETF, *Path Computation Element (PCE) Subgroup*. [Online]. Available: http://datatracker.ietf.org/wg/pce/charter/

[24] IETF, *CCAMP Subgroup*. [Online]. Available: http://datatracker.ietf.org/wg/ccamp/charter/

[25] I. Kaminow and T. Koch, Eds., *Optical Fiber Telecommunications III*. New York: Academic, Apr. 14, 1997.

[26] G. Li, A. Chiu, R. Doverspike, D. Xu, and D. Wang, "Efficient routing in heterogeneous core DWDM networks," in *Proc. Conf. Opt. Fiber Commun./Nat. Fiber Opt. Eng. Conf.*, San Diego, CA, Mar. 2010, pp. 1–3.

[27] R. Martinez, C. Pinart, F. Cugini, N. Andriolli, L. Valcarenghi, P. Castoldi, L. Wosinska, J. Comellas, and G. Junyent, "Challenges and requirements for introducing impairment-awareness into the management and control planes of ASON/GMPLS WDM networks," *IEEE Commun. Mag.*, vol. 44, no. 12, pp. 76–85, Dec. 2006.

[28] *Optical Internetworking Forum (OIF)*. [Online]. Available: http://www.oiforum.com/

[29] Optical Internetworking Forum (OIF), *External Network-Network Interface (E-NNI) OSPF-Based Routing—1.0 (Intra-Carrier) Implementation Agreement*, OIF-ENNI-OSPF-01.0, Jan. 17, 2007. [Online].

Available: http://www.oiforum.com/public/documents/OIF-ENNI-OSPF-01.0.pdf

[30] *Optical Mesh Service*. [Online]. Available: http://www.business.att.com/wholesale/Service/ data-networking-wholesale/long-haul-access-wholesale/optical-mesh-service-wholesale

[31] P. Palacharla, X. Wang, I. Kim, D. Bihon, M. Feuer, and S. Woodward, "Blocking performance in dynamic optical networks based on colorless, non-directional ROADMs," in *Proc. Conf. Opt. Fiber Commun./Nat. Fiber Opt. Eng. Conf.*, Los Angeles, CA, Mar. 2011, pp. 1–3.

[32] J. Strand, "Integrated route selection, transponder placement, wavelength assignment, and restoration in an advanced ROADM architecture," *J. Lightw. Res.*, May 2011.

[33] M. G. Taylor, "Coherent detection method using DSP for Demodulation of signal and subsequent equalization of propagation impairments," *IEEE Photon. Technol. Lett.*, vol. 16, no. 2, pp. 674–676, Feb. 2004.

[34] J. Yates and Z. Ge, "Chapter 12 in Guide to Reliable Internet Services and Applications," in *Network Management: Fault Management, Performance Management, and Planned Maintenance*, C. Kalmanek, S. Misra, and R. Yang, Eds., 1st ed. New York: Springer-Verlag, 2010.

[35] X. Zhang, M. Birk, A. Chiu, R. Doverspike, M. D. Feuer, P. Magill, E. Mavrogiorgis, J. Pastor, S. L. Woodward, and J. Yates, "Bridge-and-roll demonstration in GRIPhoN (globally reconfigurable intelligent photonic network)," in *Proc. Conf. Opt. Fiber Commun./Nat. Fiber Opt. Eng. Conf.*, San Diego, CA, Mar. 2010, pp. 1–3.

[36] J. Zuidweg, *Next Generation Intelligent Networks*. Norwood, MA: Artech House, 2002.

ABOUT THE AUTHORS

Robert D. Doverspike (Fellow, IEEE) received the undergraduate degree from the University of Colorado, Boulder and the M.S. and Ph.D. degrees from Rensselaer Polytechnic Institute (RPI), Troy, NY.

He began his career with Bell Labs and, upon divestiture of the Bell System, went to Bellcore (now Telcordia). Later, he returned to AT&T Labs (Research) where he is now Executive Director of Network Evolution Research. He has made extensive contributions to the field of optimization of multilayered transmission and switching networks and pioneered the concept of packet transport in metro and long distance networks. He also pioneered work in spearheading the deployment of new architectures for transport and IP networks, network restoration, and integrated network management of IP-over-optical-layer networks. He has over 1500 citations to his books and articles over diverse areas/publications such as Telecommunications, Optical Networking, Mathematical Programming, IEEE Magazine, IEEE Communications Society, Operations Research, Applied Probability, and Network Management.

Dr. Doverspike holds many professional leadership positions and awards, such as INFORMS Fellow, member of the Optical Society of America (OSA), co-founder of the INFORMS Technical Section on Telecommunications, Optical Fiber Communications (OFC) Network Technologies and Applications Committee, Design of Reliable Communications Networks (DRCN) Steering Committee, and Associate Editor of the *Journal of Heuristics*.

Jennifer Yates (Member, IEEE) received the B.E. (honors) and B.Sc. degrees from the University of Western Australia, Perth, Australia and the Ph.D. degree from The University of Melbourne, Melbourne, Australia.

She is an Executive Director of Research at AT&T Labs, Research, leading Research's Network and Service Management Department. The department is inventing and prototyping future service and network management capabilities focused on mobility, IP and IPTV services, and driving these technologies to large-scale deployment across AT&T networks. Her recent research has focused on service quality management and advanced data mining to detect, troubleshoot, and resolve service and network issues. Her earlier work focused on IP and optical network integration, ranging from network management and control, network reliability, IP control of optical networks (GMPLS), and IP and optical network integration. She was instrumental in AT&T's industry-leading optical mesh service deployment, which made the much touted "optical bandwidth on demand" a commercial reality.