

Optical Networking: Past, Present, and Future

Joseph Berthold, *Senior Member, IEEE, Member, OSA*, Adel A. M. Saleh, *Life Fellow, IEEE, Fellow, OSA*, Loudon Blair, *Member, IEEE*, Jane M. Simmons, *Senior Member, IEEE*

Invited Paper

Abstract—Over the past 25 years, networks have evolved from being relatively static with fairly homogeneous traffic to being more configurable and carrying a heterogeneous array of services. As the applications are ultimately the driver of network evolution, the paper begins with a brief history of circuit, packet, and wave services, along with the development of the corresponding transport layers. The discussion then moves to the evolution of network-node architecture, with an emphasis on the optical-electrical-optical and optical-bypass paradigms. Scalability and cost-effectiveness in meeting network demands are two key factors in the discussion. The evolution of networking equipment, along with the development of the optical control plane, has facilitated a configurable optical layer. The enabling technologies, along with their ramifications, are discussed. Finally, the paper speculates on how capacity might evolve in the future, to handle the undoubtedly new services that are on the horizon.

Index Terms—configurability, dynamic networks, network services, network evolution, network transport layers, optical bypass, optical network elements, optical networking

I. INTRODUCTION

OVER the last 25 years, we have witnessed a dramatic increase in the capacity of commercial optical networks, which has led to an even more dramatic revolution in network services. This network capacity increase was made possible by continuing increases in the capacity of fiber optic transmission systems and by similar increases in the capacity of switching systems. Since 1983, transmission system capacity has grown by more than four orders of magnitude, while switching system capacity has grown by a still respectable three orders of magnitude.

More recently, we have seen the commercial deployment of optical network elements that can set up and reconfigure wavelengths under software control. Now that the long held vision of dynamic optical networking has made measurable progress, what are the anticipated next steps, and what possibilities are still out of reach, patiently awaiting technology breakthroughs?

In Section II, we review how telecommunication network services have evolved and, in Section III, we show how these are mapped into network transport layers. These two sections

provide a view of the current network evolution path, and the context in which optical networking can contribute. The remaining sections then specifically focus on the optical networking layer. Section IV starts by looking at how network nodal elements have evolved. Section V then explores network agility and Section VI addresses how increases in network capacity may be accommodated. The paper concludes with Section VII.

II. NETWORK SERVICES

A. Network Service Evolution

It is interesting to review just how much commercial networking services have changed over the last 25 years. While the traffic in today's networks is dominated by video and data, with voice representing a small percentage of the total, the opposite was true in the mid 1980s. Then, networks were built to carry voice traffic, and the few data services that were available over global networks were carried on a voice-optimized infrastructure. Early applications of data communications included time-shared computing systems and transaction processing systems. Computer makers developed proprietary applications and protocols to support these applications. Later, open networking systems enabled personal communications via email, global information access, and distributed computing.

Digital transmission technology over fiber optics provided near error-free transmission, making it possible to greatly simplify data communications protocols. Frame Relay Service exploited these innovations, and is one early example of a commercially successful wide area data networking service built upon international standards. In the government and education communities, a new set of data networking protocols, including the Internet Protocol (IP) and others, led to interoperable global-scale data networking.

While data networking capabilities of that era were primitive, many visionaries within the networking community saw a future where networks would emerge to support integrated voice, data, and video services, delivered over fiber optic access lines, with service bandwidth four orders of magnitude greater than existed at that time.

This multi-service vision has made great strides towards fulfillment. Voice, video, and data service bundles are available from a number of service providers. The Internet provides global information access, and much of the traffic on

network backbones is composed of rich media, especially video. Traffic growth continues at a rate of about 50% per year, without yet showing signs of slowing [1].

B. Network Services Today

After this brief historical overview, we now consider network services today, primarily from the point of view of two broad categories of customers: residential and business.

1) Residential Services

The largest bandwidth drivers for residential users are high-speed Internet access and video entertainment. Internet access, in addition to supporting email and providing access to information, is increasingly used to deliver video content, either relatively short video clips or the download of full-length programs for subsequent viewing.

Video broadcast program material is primarily delivered to a subscriber for immediate viewing, or for storage on a personal video recorder for later viewing. Even though today's broadcast video service offerings include several hundred channels, efficiencies in video coding allow these to be carried in one or a very few gigabits per second (Gb/s) Ethernet channels. With simple digital replication at network branch points, copies of a broadcast stream can be delivered from a video head-end to distant serving offices at relatively low cost, with total bandwidth independent of subscriber population.

There is an increase in video-on-demand services, where program content can be selected and the delivery controlled by the subscriber. If video services change from a broadcast delivery model to a personalized content delivery model, the magnitude of traffic in metro networks will scale with the number of subscribers simultaneously using the service. Given the data rate for high-definition digitized video is approximately 6 to 8 Mb/s, the traffic loads in metro networks supporting video-on-demand can be two orders of magnitude larger than required for digitized voice.

The main conclusion to draw is that optical network capacity requirements for video content delivery in the metro are very uncertain, driving carriers to choose architectures that are highly scalable, yet tolerant of uncertainty as to the magnitude, location and timing of traffic.

2) Large Business and Government

Small businesses' networking requirements are generally met by telephony and Internet access services. Medium and large businesses and governments have additional needs, and often have dedicated communications and information technology staff. For the remainder of our discussion of business services we will concentrate on large businesses and governments, since their bandwidth needs are very large, and their networking requirements are quite different from residential and small business customers.

Large enterprises and governments consider information technology as a strategic asset, and often build and operate their own voice and data networks for reasons of guaranteeing performance, enhancing security, and to be in a better position to control the evolution of their networks in support of the business. Their networks need to support a variety of applica-

tions, including voice, private intranet services, Internet access, remote database access, data replication, as well as business continuity/disaster recovery. While many enterprises are moving to the use of IP with many of their applications, this does not necessarily imply that they will procure only IP services from network service providers. It is very common for a large enterprise to procure a broad mixture of services that include wireline voice, public Internet access, IP virtual private networks (VPNs), Ethernet or Multiprotocol Label Switching (MPLS) services at Layer 2, time division multiplexing (TDM) private lines, Ethernet private lines, and wave services. It is also common for enterprises with special needs, such as financial institutions, to lease optical fiber and build dense wavelength division multiplexing (DWDM) networks, procuring and operating the transmission equipment themselves. Governments, at multiple levels and in multiple agencies, often build and operate their own optical networks. Large enterprises and governments represent a substantial traffic load on both metro and core optical networks.

C. Services at Multiple Network Layers

With the general trend to migrate applications onto an IP-based protocol stack, it is important to consider whether this implies that service providers will offer only IP services in the future. In this section, we discuss three general categories of service offerings below the IP network layer, and the needs they fulfill. These categories include Layer-2 packet services, TDM services, and wave services.

1) Layer-2 Packet Services

Packet transport is the natural choice for supporting bursty data services, as it allows sharing of network links by many users and applications. Substantial statistical multiplexing factors, often more than 20:1, are used in today's data networks, especially near the network edge. Layer-2 packet services can be thought of as virtual wires interconnecting network-layer devices, such as IP routers. Many business customers, especially large businesses and governments, prefer to procure a Layer-2 service rather than an IP service from a network service provider, as discussed earlier. The customer might then build and operate a private network using these Layer-2 links as switch interconnects.

Frame Relay, Asynchronous Transfer Mode (ATM), MPLS, and Ethernet are examples of Layer-2 packet services. While Frame Relay and ATM are in decline, MPLS and Ethernet services are growing. There is currently much interest in further extending Ethernet to support a wide-area Layer-2 commercial service [2]. Some of the prime attractions of Ethernet are its ubiquity in customer networking equipment, the low cost of Ethernet semiconductor and optical components, and the excellent quality of standards specifications, which result in a marketplace with many interoperable products. Finally, there has been a great deal of work in the industry to reach agreement on service descriptions, with strong participation among global service providers. The Metro Ethernet Forum (MEF) has been instrumental in reaching these agreements.

The MEF defines a number of Ethernet services, including E-Line service and E-LAN service. E-Line service is a point-to-point service whereas E-LAN is a multipoint-to-multipoint service. E-Line service can be used as a replacement for Frame Relay, ATM, or low-speed private line services that may have been used to support data services. E-LAN service can provide an extension of a business' local area network across a metro area. They offer potential benefits of cost and security over the alternative of a public Internet service.

2) TDM Services

TDM services provide guaranteed bandwidth, low latency and low jitter, as well as security benefits. They are simple to understand and to manage, and are the mainstay of high capacity customer managed networks with stringent performance and security requirements.

Low data-rate TDM services, with rates below the 50/155 Mb/s container sizes of Synchronous Optical Network (SONET)/Synchronous Digital Hierarchy (SDH) [3], are often considered as candidates for replacement by circuit emulation packet services [4]. As many of these circuit services carry packet flows, either data or packetized voice, such a transition is logical if it can be done economically. There continue to be applications that have stringent timing requirements, and we will have to wait and see if circuit emulation technology proves adequate. For circuit services at 155 Mb/s and above, circuit emulation is possible in principle, but so far, the economics have not been adequate.

All indications are that TDM circuit services will be an important service provider offering for the foreseeable future.

3) Wave Services

Wave services originated as virtual fiber services supported on wavelengths of DWDM transmission systems. Wave services may meet a need for very high bandwidth circuits, or meet a requirement for protocol and bit-rate transparency for the data carried by the service. Early implementations were bit transparent, limiting manageability to what could be accomplished with non-intrusive monitoring. Currently, wave services often achieve a combination of transparency and manageability by employing Optical Transport Network (OTN) "digital wrapper" technology standardized by the ITU-T. OTN is discussed in Section III.

With continued progress to higher channel rates in DWDM systems, today's wavelengths are tomorrow's fractional wavelengths. Even though the IEEE has just begun standardization of a 100 gigabit Ethernet (100 GbE) signal in the 802.3ba Task Force [5], some customers already see the need for higher-rate service offerings that might require a number of wavelengths to deliver the required bandwidth.

D. Dynamic Services

Dynamic networking services give customers control of their network resources so they can reconfigure connections and change bandwidth as they need. While public IP network services are inherently dynamic, meaning data can be exchanged anywhere across the global Internet, it is difficult to meet stringent service guarantees over the public Internet.

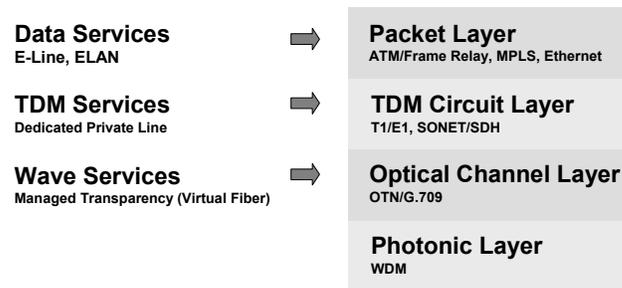


Fig. 1. Mapping of network services to network transport layers.

Dynamic customer control of lower-layer public services, which can provide quality-of-service guarantees, has been quite limited. The introduction of distributed control-plane technology into TDM circuit switched networks has enabled a new class of user controlled high-bandwidth dynamic circuit services. Service providers and network equipment vendors around the world have worked diligently over the past five years to develop protocols and implementation agreements to support multi-vendor dynamic connection establishment, and have engaged in several worldwide testing events to advance interoperability [6]-[8]. It can be anticipated that dynamic connection establishment functionality will be extended from the TDM circuit layer up to Layer 2, as demonstrated in the recent OIF interoperability event, and down to the wavelength layer, enabling an extension of customer-controlled network services in both of these directions.

III. NETWORK TRANSPORT LAYERS

Network infrastructure is often described as a layered arrangement, where each layer plays a distinct role. For purposes of discussion here, we break the infrastructure down into four layers (see Fig. 1). The first three layers are packet transport, TDM circuit transport, and optical channel transport. These three transport layers align with the three service types discussed in II-C above, and can therefore serve as a service delivery platform in addition to providing transport of aggregated traffic. The fourth layer is the photonic path, representing an all-optical connection in the network. Photonic paths are not currently common service offerings, but they offer significant infrastructure benefits in their ability to reduce cost, space, and power dissipation. Not all four layers need be present in every network implementation.

A. Packet Transport Layer

With most services moving to IP, why should a carrier deploy a packet transport layer in addition to the IP networking layer? The answer to this question lies in the requirement for the network to support multiple services simultaneously, where each service may require stringent, yet different, performance with respect to packet loss, transmission latency, and delay variation or packet jitter. Carriers require a packet transport layer to support traffic engineering, provide guaranteed quality of service parameters, and support service monitoring including performance management and fault management.

MPLS [9] was developed during the late 1990s to support the IP layer with such a packet transport capability. It has now been widely accepted within carriers' core networks as a mechanism to provide improved traffic management for IP traffic flows. Over the past few years, with the promise of low cost, Ethernet has become attractive as a packet transport layer for packet access and packet service connectivity in the metropolitan area.

A challenge that network operators have faced as they started introducing Ethernet outside of the local area network (LAN) is that legacy Ethernet does not support the carrier-grade features required by the carriers to manage the end-to-end health of the service. This has resulted in an intensive packet operations, administration and maintenance (OAM) standardization effort within the standards bodies resulting in OAM standards such as those being developed within the IEEE [10] and ITU-T [11].

Further, as carriers began to understand these management challenges, transport-oriented connection-oriented packet transport techniques such as Transport MPLS (T-MPLS) [12] and Provider Backbone Bridging – Traffic Engineering (PBB-TE) [13] have grown in popularity. The intent with these methods is that the connection-oriented packet tunnels operate at a level in the network that is equivalent to SONET/SDH, providing similar predictable and managed connectivity to the appropriate data clients. The connection-oriented tunnels are agnostic to higher layer services and they provide any-to-any flexibility and fine-granularity of connectivity that allows for efficient use of the underlying network resources.

B. TDM Circuit Transport Layer

The SONET/SDH standards developed during the late 1980s and early 1990s fundamentally changed the world of optical fiber communications by defining a common set of optical signals, operations procedures, and a multiplexing hierarchy. SONET/SDH provided the first opportunity for carriers to offer consistent end-to-end optical communications across multiple network domains, administered by different network operators and implemented using different vendors' equipment. Self-healing ring protection mechanisms provided a robust and highly available transport substrate upon which the newly-forming Internet and broadband services (including voice, data and video) could be offered.

As a managed transport network layer in the early to mid 1990s, SONET/SDH acted as a "server" for all of the existing service "clients" at that time. Service clients included TDM voice circuits for both residential and business customers, as well as TDM private line, Frame Relay, and ATM in support of broadband data services for business customers.

To accommodate increasing capacity demands due to exponential growth in data traffic, two approaches were used: (i) the line rate of SONET/SDH systems evolved in multiples of four from 155 Mb/s (OC-3/STM-1) to today's 40 Gb/s (OC-768/STM-256), and (ii) multiple SONET/SDH systems were overlaid on top of one another. Both approaches resulted in new overbuilds that required the dedicated use of multiple optical fibers for connectivity and the costly addition of electronic regenerators between add/drop locations.

C. Optical Channel Transport over DWDM

Driven by a desire to augment the rising capacity demands and to reduce the cost of scaling SONET/SDH connectivity, network operators introduced wavelength division multiplexing (WDM) to their networks in the mid 1990s. Despite this increase in traffic capacity, the network's fundamental dependence on the SONET/SDH architecture as the underlying server layer did not change.

Over the last ten years, networking protocols have evolved and we have seen the growth in popularity of technologies like IP/MPLS, Carrier Ethernet, Fiber Channel, ESCON, etc. For reasons mostly related to cost, these technologies have been built directly over dark fiber or WDM as well. In each of these scenarios, WDM was typically used as a capacity extension to the higher-layer protocols. In some cases, such as SONET/SDH, the WDM component was managed within the context of the powerful SONET/SDH management capability. In the case of some of the other technologies (e.g., Ethernet or Fiber Channel), the data protocols were not suited to wide area carrier-grade management. Therefore, the carrier's choice was limited to operating these other technologies over SONET/SDH (which turned out to be CAPEX intensive) or not managing the service very well (resulting in OPEX challenges).

Another perspective to note is that, by enhancing the network capacity with WDM, a new server layer was created – an unmanaged WDM layer – beneath all of these different technologies. Unfortunately, at this stage in the network's evolution, the quality of network management was dependent on the capabilities of the higher-layer protocols, some of which were good and some of which were (and continue to be) poor. This point is important because, as traffic has increased and as capacities have grown, there has been an increasing demand for wave services, and growth in a diverse set of data-service protocols. Because of this, there clearly becomes a need for the WDM layer to become an autonomous networking layer in its own right and support its own management capability.

Today, service demands on the network are becoming increasingly complex. For reasons such as security, client independence, or simply wholesale-service delivery, many network customers are demanding service transparency for a wide array of standard and more specialized communications protocols. Such transparency today may be supported using dedicated fiber or wavelengths, but these connections provide simple transport with limited monitoring capabilities. OAM techniques are typically based on the capability of each client protocol. With significant growth in demand, there is a need to manage these multi-service demands in a manner that maintains the client-signal integrity in addition to being able to offer wavelengths themselves as self-contained managed service offerings. Further, as the network evolves, legacy SONET/SDH or ATM networks will not be decommissioned immediately. Instead, it will be necessary to operate concurrent networks with different technologies and slowly migrate these services from the old to the new networking paradigms.

1) OTN/G.709 "Digital Wrapper"

ITU's OTN standard [14] with ITU-T G.709 framing structure [15] (sometimes called "Digital Wrapper") provides the necessary management capability for this new WDM server

layer and offers a common managed foundation upon which to build all transport technology clients. OTN/G.709 was standardized in the early 2000s. The forward error correction (FEC) portion of the standard (described in ITU-T G.975 [16]) was quickly accepted as a means to extend the distance of high-speed optical signals before requiring electrical regeneration. However, the OAM management capabilities were not widely adopted, until recently. Recognizing that multi-vendor interoperability has always been challenging, the ability of OTN/G.709 to carry different protocols transparently without affecting content, control channels, or timing makes it an ideal platform upon which to evolve the network infrastructure without undue risk. Using G.709 framing, with its associated rich OAM capabilities, effectively levels the playing field across all of the different transport technologies by providing a consistent managed view for all transport services. OTN/G.709 provides a new convergence layer for optical networks without compromising the resilient operations and management capabilities we have come to expect from SONET/SDH. Instead, SONET/SDH has now become one of the many clients of the WDM network.

2) *Optical Ethernet*

As discussed in Section III.A, a new and important client to the WDM network is Ethernet. Increased data traffic flows have resulted in a significant shift toward the use of Ethernet as a common transport layer for packet networks with both 1GbE and 10GbE becoming the preferred interfaces of choice for data devices.

In particular, the 10GbE LAN PHY standard [17] has seen widespread acceptance. Network operators are now being challenged to transport 10GbE LAN PHY transparently outside of the LAN across the wide area network (WAN) in a manageable and interoperable manner. While a full rate 10GbE signal does not fit within a SONET/SDH 10-Gb/s frame, it can be carried “digitally wrapped” within an over-clocked OTU-2 frame [18]. In addition, because of its powerful OAM attributes, OTN/G.709 is now recognized as a key enabler for a robust and manageable evolution from traditional SONET/SDH networking to Ethernet-based packet transport.

Today, standardization activities within the IEEE’s P802.3ba 40Gb/s and 100Gb/s Ethernet Task Force and the ITU’s Study Group 15 that are looking at new 40-Gb/s and 100-Gb/s transmission speeds for Ethernet are being careful to consider support for transparent multiplexing of lower-rate OTN signals, including a way to carry 10GbE LAN in a fully transparent fashion.

D. *Photonic Layer*

The optical transmission functions that define photonic paths across a network have not traditionally been considered a network layer in the same manner as the above (they are sometimes referred to as Layer 0). Instead, the optical transmission and switching technologies form a network substrate upon which all other network functions are built. The following sections explore how the photonic layer is evolving into an autonomous networking layer in its own right.

IV. NETWORK NODAL ELEMENTS

The network nodes are the sites in the network that source/terminate and switch traffic. Over the past 25 years, there has been a general trend in nodal evolution to handle traffic at coarser granularities, while reserving finer-traffic processing for only the traffic that requires it. This has enabled scalable nodal growth of several orders of magnitude.

In the 1980s, backbone nodal equipment was based on Digital Cross-connect Systems (DCS) that operated on the asynchronous T1/E1 digital carrier hierarchy. (A T1 carries 1.5 Mb/s, or 24 64-kb/s voice channels; a T3 carries 45 Mb/s. The respective data signals are known as DS-1 and DS-3. In Europe, the hierarchy is based on an E1, or 2.0 Mb/s.) For example, AT&T employed the Digital Access and Cross-Connect System (DACCS) series of equipment.

In a typical North American architecture, all T3s entering an office were fully terminated on a ‘3/3’ DS-3 cross-connect, with DS-3 ports and DS-3 switch granularity (e.g., a DACS III). Traffic that needed further processing at the DS-1 level was sent to a ‘3/1’ cross-connect with DS-3 ports and DS-1 switch granularity (e.g., a DACS IV). As traffic levels grew to tens of thousands of DS-3s, terminating every DS-3 at every intermediate node of its path became very costly.

The nodal architecture significantly changed with the development of synchronous SONET/SDH technology, specifically the add/drop multiplexer (ADM). With synchronous technology, portions of a signal can be dropped (or inserted) without terminating the remainder of the signal. For example, a SONET ADM may operate on an STS-48 carrying 48 DS-3s, where it drops only those DS-3 tributaries that require further processing at the node, with the remaining traffic passing through the ADM. Thus, the DCS layer is needed only for those DS-3s that actually need to be processed at a node, providing a more scalable architecture.

ADMs are ‘degree-two’ devices, and thus support linear and ring configurations. The canonical architecture of the mid-1990s consisted of numerous SONET/SDH rings overlaid on the network topology. A given ring structure might support tens of wavelengths, giving rise to the ‘stacked-ring’ architecture, where each wavelength-ring was essentially an independent entity [19]. Typically, one ADM was deployed on each wavelength at each ring node. Any communication between the ADMs generally occurred at the DS-3 level, either with DCS equipment or with manual DS-3 patch-panels.

More flexibility was provided in the late 1990s with the introduction of large SONET/SDH cross-connects, which enabled mesh topologies. Such cross-connects, or ‘optical switches’, were capable of fast, mesh-based protection switching. Furthermore, when integrated with the control plane (discussed further in Section V), they opened up the possibility of auto-provisioning of substrate traffic.

However, this model still required that all wavelengths entering a node be electronically terminated (i.e., prior to being delivered to the SONET/SDH equipment). As traffic levels grew, deploying the corresponding electronic terminating equipment became very costly. To alleviate the costs to a degree, manual wavelength-bypass was employed, where if all the traffic carried on a wavelength did not need processing at a

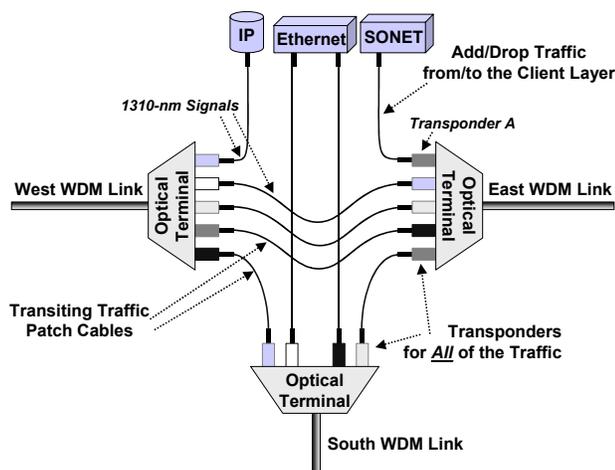


Fig. 2. O-E-O architecture at a degree-three node. Transponders are required for all traffic entering the node. Note that this configuration is not configurable; e.g., without manual intervention, transponder *A* can be used only for add/drop traffic, and only for traffic adding/dropping to/from the East link.

node, it was passed through the node in the optical domain, via a patch-panel. However, such an arrangement was not automatically configurable, and was prone to manual errors.

This gave rise to optical network elements that could enable optical-bypass, where only the wavelengths that need to be processed at a node are electronically terminated; the remainder of the wavelengths stay in the optical domain. These elements also can provide automated reconfigurability.

As moving from an architecture where all traffic entering a node is electronically terminated to one where the bulk of the traffic remains in the optical domain has profound effects on the economics and operation of networks, these two contrasting models are probed in more detail in the next two sections.

A. O-E-O Architecture

The architecture where all traffic entering a node is converted from the optical domain to the electrical domain, and back to the optical domain is known as the optical-electrical-optical (O-E-O) paradigm [20]. This is illustrated in Fig. 2 for a degree-three node (i.e., a node with three incident links). The WDM signal on each of the incoming fibers is terminated on an optical terminal, which demultiplexes the signal into its constituent wavelengths. Each wavelength is directed to a separate transponder that converts the WDM-compatible signal to the electrical domain and then to a standard 1310-nm optical signal. Similarly, in the outgoing direction, the optical terminal multiplexes the signals from each of the associated transponders, thereby generating a WDM signal.

Two types of nodal traffic should be distinguished, as indicated in the figure: traffic that is transiting the node on its way to its final destination, and traffic that is added/dropped at the node. To support a transiting signal, two transponders are interconnected in the node to form the through-path. The architecture shown in Fig. 2 is non-configurable, where a patch-cable is used for interconnection. (Automated reconfigurability is discussed in Section V.A.2.) A pair of interconnected transponders can operate at different wavelengths,

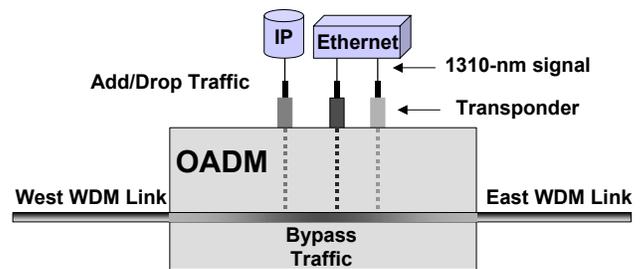


Fig. 3. Functional illustration of an OADM.

so that wavelength conversion can occur as the signal traverses the node. The add/drop traffic communicates with the client layer (e.g., SONET/SDH, IP, Ethernet) via a 1310-nm signal.

The O-E-O architecture does provide several benefits. First, the process of passing through back-to-back transponders restores the quality of the optical signal; i.e., all transiting traffic is regenerated. Second, converting all traffic to a well-defined 1310-nm signal provides multi-vendor support; e.g., different vendors can potentially provide the transmission system for each of the links at a node. Third, the electronics of the transponder provides the opportunity for detailed performance monitoring (e.g., the SONET/SDH overhead bytes can be inspected at each node along the path), allowing most failures to be readily detected and localized. Finally, the wavelength conversion inherently supported at each node implies that wavelengths can be assigned independently on each link.

The disadvantages of the O-E-O architecture revolve around scalability. As the level of traffic grows, utilizing two transponders for every transiting wavelength at a node can be very costly. Furthermore, the number of required transponders poses challenges in physical space, power requirements, and heat dissipation. Additionally, the amount of equipment that must be deployed to support a new connection slows down the provisioning process.

B. Optical-Bypass-Enabled Architecture

As indicated previously, these scalability issues led to the development of *optical-bypass* technology, where transiting traffic remains in the optical domain, thus eliminating the need for transponders for this traffic. Initially, this capability was available only for degree-two configurations, with the Optical Add-Drop Multiplexer (OADM), introduced commercially in the mid-1990s [20]. An OADM is functionally illustrated in Fig. 3. Note that transponders are required only for the add/drop traffic. Optical bypass was extended to higher-degree nodes with the development of the multi-degree OADM (OADM-MD) and the all-optical switch (AOS), which were introduced commercially in the 2000 timeframe. The OADM-MD and AOS are functionally illustrated for degree-three nodes in Fig. 4 and Fig. 5, respectively, where both elements support optical bypass in all directions through the node. (These two elements differ in the amount of configurability they provide, as discussed in Section V.A.2.)

The initial optical-bypass vision assumed a pure O-O-O architecture for all transiting traffic, where a connection remains in the optical domain from its source to its destination. The

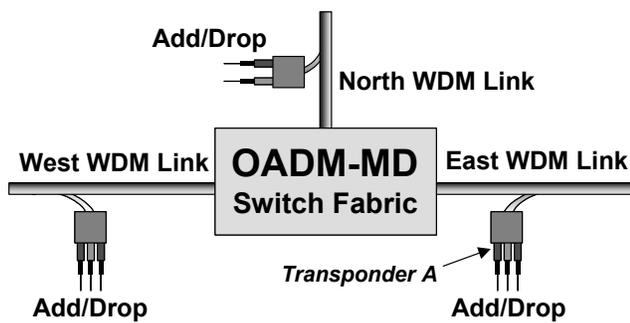


Fig. 4. Functional illustration of a degree-three OADM-MD. Transponder A can be used only for traffic adding/dropping to/from the East link.

commercial reality of optical-bypass-enabled networks, however, is somewhat different. First, in U.S. and pan-European backbone networks, the optical reach of commercially available systems is not long enough to carry all connections in the optical domain end-to-end. (The optical reach is the distance an optical signal can travel before it requires regeneration.) Moreover, studies have shown that increasing the optical reach to eliminate all need for regeneration in a backbone network is not likely to be cost effective [21], [22].

Furthermore, regardless of the geographic tier, traffic may require electrical processing at a small number of intermediate nodes for functions such as grooming (i.e., bundling low-rate demands onto higher-rate wavelengths) and shared protection. As line-rates increase, typically more grooming is needed, which reduces the opportunities for optical bypass somewhat. Shared protection may be implemented with electronic edge cross-connects (usually at higher-degree nodes) that switch the spare capacity in response to a failure. The small amount of electronics that is required for these functions can be beneficial, as discussed below in relation to wavelength assignment.

While not truly all-optical, optical-bypass-enabled networks typically eliminate more than 60% of the nodal transponders as compared to an O-E-O architecture, providing significant benefits in cost, space, power, and heat dissipation [23], [24]. (Even with traffic that requires grooming and/or shared protection, a significant amount of optical bypass can still be realized, e.g., [25].) Optical bypass also enables more scalable switching architectures, as discussed in Section V.A.2. Furthermore, with fewer transponders required along a path, installation costs and deployment times are markedly lower; this typically leads to higher availability rates as well [26].

Clearly, the benefits of optical bypass are dependent on the network topology and the traffic set, where traffic that transits more intermediate nodes benefits more from this technology. Additionally, since optical-bypass elements come with a price premium as compared to optical terminals, the level of traffic must be high enough to realize a net capital-cost savings.

Removing electronics does create another set of challenges, however. Because the signal is not regenerated at every node, taking advantage of optical-bypass-enabling network elements necessitated extended-reach transmission. Whereas legacy transmission systems have an optical reach on the order of 600 km, backbone networks with extended-reach technology typically have an optical reach in the range of 2,000 km to 4,000

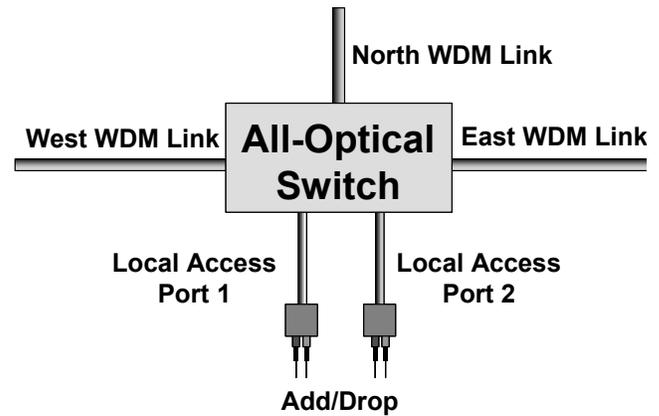


Fig. 5. Functional illustration of a degree-three all-optical switch. The add/drop traffic can access any of the network links.

km. Achieving this extended reach required technology such as advanced optical amplifiers (e.g., Raman), novel modulation formats, stronger FEC, and various techniques to combat optical impairments. (Note that with a reach of 600 km, legacy networks often required regenerations at sites along a link, not just at the link endpoints. The combination of extended reach and optical bypass eliminates as much as 90% of the total transponders as compared to legacy systems.) With the smaller geographic extent of a metro or regional network, an optical reach consistent with a high level of optical bypass is generally on the order of 1,000 km or less.

The elimination of the transponders for transiting traffic also removed the associated wavelength conversion capability. Thus, a signal enters and exits a node on the same wavelength, which implies that the assignment of a wavelength on one link can affect the availability of that wavelength on several other links. Much research has gone into the development of wavelength-assignment algorithms, resulting in relatively simple and effective schemes [27]. Furthermore, the small amount of O-E-O conversion that is still required in the network provides an opportunity to convert the wavelength of the signal. This reduces the number of consecutive links on which a common available wavelength must be found. With the combination of intelligent algorithms and the wavelength conversion that naturally occurs in a network (e.g., due to sparse regeneration, grooming, or shared protection), many studies have shown that the loss of network efficiency due to wavelength contention in optical-bypass-enabled networks is very small [28]-[30].

To compensate for the lack of node-by-node electronic monitoring, various optical performance monitoring (OPM) techniques have been developed [31], [32]. For example, optical signal parameters such as power level, wavelength accuracy, and optical-signal-to-noise-ratio (OSNR) can be monitored; more advanced OPMs can monitor the Q-factor of the signal, to detect a wider range of failures.

The one remaining challenge of optical bypass is standardizing the transmission system properties such that multi-vendor architectures can be more readily supported.

Overall, optical bypass has addressed many of the scalability challenges presented by the O-E-O architecture.

The advantages in both capital costs and operating costs have led several major carriers worldwide to adopt this technology [33]-[35].

C. Metro vs. Backbone Nodal Architecture

While technologies such as WDM and optical bypass have gradually migrated from backbone networks to metro networks, there are several differences in the implementations. For example, while the number of wavelengths on a fiber in backbone networks is typically 80 or more, the corresponding number in metro networks is 40 or less. The topology of the tiers tends to be different as well, resulting in different requirements for the optical network elements. Backbone networks are generally mesh-based whereas many metro networks are composed of interconnected rings (although even metro networks are currently migrating to mesh). Degree-four, or smaller, network elements are sufficient for most backbone nodes, whereas a metro network may require up to degree-eight network elements (e.g., at an intersection of four rings).

The electronic switching technology is also different in the two network tiers. While the backbone network makes use of large core IP routers and SONET/SDH grooming switches, the metro is typically equipped with access IP routers, smaller SONET/SDH boxes, and Ethernet switches. Links interconnecting core IP routers often consist of one or more wavelengths carried directly over the same DWDM transmission systems used to interconnect core SONET/SDH switches and provide wave services. While DWDM interfaces are sometimes available both in IP routers and SONET/SDH switches, service providers often choose to require an O-E-O demarcation between these network elements to simplify multi-vendor interoperability. In the metro, there is a trend among transport product suppliers, driven by some service providers, towards switching convergence, where a single box would provide circuit switching, packet switching, as well as optical wavelength switching. It is still an open question as to whether the majority of service providers will move from their current practice of deploying a modular transport architecture, with differing vendors for each equipment layer, in favor of a converged architecture and fewer vendors.

D. Future Nodal Evolution

The nodal architecture has evolved from legacy O-E-O deployments to a pure O-O-O vision (that never was commercially realized) to the current hybrid solution. It is interesting to consider the evolution of nodal architecture going forward. One possibility is that it reverts back to the O-E-O paradigm, where the issues of transponder size, and to some degree cost, are mitigated through photonic integrated circuits (PIC) [36]. With PIC technology, multiple transponders are integrated onto a small chip, such that capacity is installed in bulk at the nodes. With O-E-O termination at every node, issues such as wavelength contention are removed (however, as noted earlier, wavelength assignment has already been addressed through intelligent algorithms). While an improvement over an O-E-O architecture with discrete transponders, it is not clear how PIC technology addresses the challenge of heat dissipation as the network scales in capacity. Furthermore, pro-

viding full automated configurability at a node would require a very large switch, which poses a significant challenge.

Alternatively, networks may move to a more agile optical layer (as discussed in Section V), and may migrate more functions to the optical domain, to take advantage of the trend of optics to scale better than electronics as the traffic level and the wavelength bit-rate increase. For example, optical-grooming techniques such as optical burst switching [37] potentially may be implemented in edge networks to reduce the electronic-grooming requirements at the edge-core interface [38]. Optical-grooming schemes tend to rely on scheduling or collision management, and thus are better suited for small geographic areas (e.g., metro networks) with low latency. As a second example, all-optical regenerators may be used instead of back-to-back transponders (or instead of electronic regenerator cards) to regenerate, and convert the wavelength of, optical signals [39], [40]. Currently, much of the all-optical regeneration technology that has been proposed is capable of only 2R regeneration, where a signal is reamplified and reshaped, but not retimed. To clean up the timing jitter, it may still be necessary to make use of a limited amount of electronic regeneration, which is capable of retiming as well.

Electronic-based and optical-based technologies clearly have their associated strengths; thus, networks are likely to remain a mix of both O-E-O and O-O-O technologies.

V. NETWORK AGILITY

Networks have traditionally been fairly static, with connections remaining established for months, or even years. For example, a typical lifetime for a DS-1 circuit in the 1980s/1990s was on the order of one to two years. In this timeframe, circuit configurability was typically provided via the DCS layer. In addition to being used for circuit provisioning, the DCS layer could also be used for dynamic restoration from a failure. For example, AT&T's FASTAR mesh-based restoration system, developed in 1992, reroutes individual failed DS-3s using the DACS III cross-connects [41].

However, as networks have grown in size, operating solely on fine levels of traffic such as a DS-3 is not scalable. Thus, reconfigurability at coarser granularities, e.g., a wavelength, has become a priority. Furthermore, the reconfigurability needs to be rapid, as motivated by both new services and by restoration requirements, as detailed next.

The process of provisioning a new wavelength has historically been slow, requiring much up-front planning and manual intervention at several sites in the network. In addition to the time and cost involved with any network modifications, the manual nature of the process left it vulnerable to operator errors. While this may have sufficed with relatively static and predictable traffic patterns, as the nature of services has moved towards bursty data connections, the need for greater optical-layer agility has grown as well. For example, with the development of new services such as grid computing and e-science applications [42], [43], where high bit-rates pipes are required for relatively short periods of time, the network must be rapidly reconfigurable to deliver the bandwidth to where it is needed.

Network survivability is also undergoing a transformation that demands rapid configurability of the optical layer. Over the past ten years, the trend was for carriers to implement optical-layer protection with dedicated 1+1 architectures, where two diverse paths are established for a connection, and the destination selects the better of the two signals. With 1+1 protection, optical-layer reconfigurability is not required for failure recovery. While relatively simple to implement, and capable of providing rapid recovery (e.g., tens of milliseconds), 1+1 protection is inefficient with respect to capacity. Providing a diverse protection path for each connection can almost triple the total amount of required network bandwidth [20].

To reduce the spare capacity requirements, the industry has gravitated towards shared mesh restoration, where working paths that are routed diversely are eligible to share restoration capacity. To enable this sharing, the network must be capable of reconfiguring the restoration capacity in response to whichever path has failed. This dynamic capability affords a roughly 50% to 70% savings in spare capacity as compared to 1+1 protection, along with greater resilience to multiple concurrent failures. Clearly, rapid reconfigurability is desirable (but shared restoration still takes on the order of hundreds of milliseconds or more to recover from a failure).

Electronic switching can provide the desired configurability. However, due to the potential cost and scalability issues, the focus has moved towards providing an agile physical optical layer, as is the emphasis of the remainder of this section.

A. Configurable Optical Networks

Configurable optical networks have been realized through a combination of hardware and software advances. The technology must allow a wavelength to be brought up, torn down, or rerouted through automated means, without affecting other existing traffic. Furthermore, the equipment must be capable of being configured remotely through software.

1) Tunability

One of the developments that greatly improved the flexibility of optical networks is tunability. Initially, transponders were capable of transmitting a fixed frequency, which posed deployment problems in optical-bypass-enabled networks where the choice of wavelength for a connection is important. Fixed-wavelength transponders also posed inventory issues (even in O-E-O networks), especially as the number of supportable wavelengths on a fiber grew to 80 or more.

In roughly the 2000 timeframe, most system vendors started to support transponders with tunable lasers capable of tuning to any wavelength in the transmission spectrum. In some architectures, the receive-side of the transponder is equipped with a filter to select a particular wavelength from a WDM signal. Transponders with tunable filters were commercially available more recently [44].

To take full advantage of tunable transponders, it is necessary that the chassis slots in which the transponders are deployed be capable of supporting any wavelength. First, consider optical network elements with multi-wavelength add/drop ports; i.e., a WDM signal is present on the add/drop port, thereby requiring a multiplexer/demultiplexer architecture. One common mux/demux architecture is based on arrayed waveguide gratings (AWGs), where each slot is tied to a

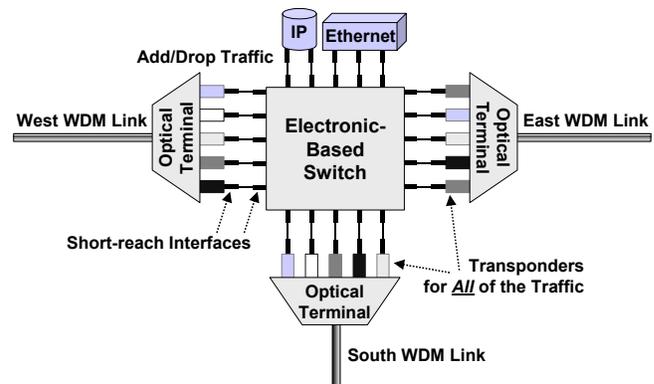


Fig. 6. A degree-three O-E-O nodal architecture, where an electronic switch is added to provide automated configurability.

particular wavelength [19]. This arrangement negates some of the benefits of tunable transponders, as a transponder would need to be physically moved to a different slot if its frequency changes. A more flexible scheme makes use of passive combiners and splitters for the mux/demux function. This architecture supports any wavelength in any slot; however, it does incur additional loss, which typically necessitates additional amplification on the add/drop path. Additionally, the transponder receiver must be equipped with a filter to select the desired wavelength. Another flexible-slot solution is a single-wavelength add/drop architecture based on a $1 \times N$ wavelength selective switch (WSS), where the switch can direct any wavelength to any of N ports [45]. However, WSSs are currently limited in size, which caps the maximum number of add/drop wavelengths.

2) Nodal Configurability

The optical network elements support varying degrees of configurability. The O-E-O architecture shown in Fig. 2 is clearly not configurable without manual intervention. The path through the node is dependent on the optical terminals in which the interconnected transponders are deployed. Furthermore, the add/drop traffic is tied to a particular link at the node; for example, any client service using transponder *A* in the figure can add/drop only from the East link. Note that the transponders are effectively partitioned into those that support transiting traffic and those that support add/drop traffic, where manual intervention is required to change the apportionment.

Adding a core switch to the node, as shown in Fig. 6, addresses these configurability issues. With the addition of a switch, traffic can pass through the node in any direction, add/drop traffic can access any network link, and any transponder can be used for either transiting or add/drop traffic. The switch shown in Fig. 6 is electronic-based, with short-reach interfaces on all of the ports. The switch could also be photonic, where the optical signal is switched (e.g., using micro-electro-mechanical-systems (MEMS) technology [46]), thereby eliminating the short-reach interfaces on the switch ports. In either case, as discussed previously, the O-E-O architecture requires that transponders be deployed for all traffic entering the node. An electronic-based core switch adds to the cost, space, power, and heat problems discussed previously.

Furthermore, as traffic levels grow, the required size of the switch becomes problematic. For example, consider a degree-four node with 80 wavelengths per fiber, where up to 50% nodal add/drop must be supported. This requires a switch size of 480x480, which is beyond what today's electronic and MEMS technologies can achieve cost effectively.

The situation is exacerbated when a grooming switch or an IP router is used as the core nodal switch. This unnecessarily burns grooming/routing resources for traffic that is transiting the node as well as for wave services that are adding/dropping at the node. The size and cost of this architecture is likely untenable as networks evolve.

In contrast, optical bypass provides configurability in a more scalable fashion. The initial OADMs that were commercially available in the mid 1990s were not configurable. The carrier needed to specify up-front which wavelengths would drop at a node, and the OADM would remain fixed in this configuration. However, by 2000, fully configurable OADMs were available, where any wavelength could be readily configured as transiting or add/drop, without affecting any of the existing traffic on the OADM. OADMs with this flexibility were later christened as reconfigurable-OADMs or ROADMs.

This same flexibility applies to an OADM-MD and an AOS. In addition, these higher-degree elements support core configurability, where the path through the node can be automatically configured. However, these two elements differ with respect to edge configurability. With the OADM-MD, the transponders are tied to a particular network link; e.g., as illustrated in Fig. 4, transponder *A* can add/drop only to/from the East link. Contrast this with the AOS of Fig. 5. Here, the add/drop traffic passes through the switch fabric, such that the transponders have access to any network port. Thus, the AOS supports edge configurability (i.e., 'direction independence'), whereas the OADM-MD does not.

It is worth commenting on the technology of these two elements. OADM-MDs were originally based on wavelength-blocker technology, where the number of wavelength blockers required in a degree- D OADM-MD scaled quadratically with D [47], [48]. Current OADM-MDs are typically built using WSSs, where the number of required WSSs scales linearly with D [49]. This improved scalability supports a larger D , such that a subset of the D ports can be used for the add/drop traffic; this affords the add/drop traffic the same level of configurability as the network traffic. With this configuration, the element is an AOS. Note that an AOS can also be based on MEMS technology, where both the network traffic and the add/drop traffic pass through the switch [50]. However, as indicated earlier, at large nodes, the required MEMS switch size is likely beyond what can currently be supported.

It is possible to add edge configurability to an OADM-MD through the addition of an edge switch (either electronic or photonic), where only the add/drop traffic passes through this switch. Using the same example of a degree-four node with 80 wavelengths per fiber and up to 50% add/drop, the edge switch would need to be 320x320. This compares to 480x480 with a central switch fabric for both express and add/drop traffic.

By providing the core switching in the optical domain, the optical-bypass-enabled elements also enable more scalable

grooming/routing architectures. Only the traffic that actually needs to be groomed/routed at a node needs to pass through the grooming-switch or router. For example, an architecture where IP routers are deployed over optical-bypass elements was shown in [51] to provide significant cost and operational benefits.

Another benefit of optical bypass with respect to configurability is that the amount of equipment needed to support a new connection is significantly smaller as compared to an O-E-O architecture. Thus, to support the same level of configurability, many fewer transponders need to be pre-deployed.

3) IP over a Configurable Optical Layer

A configurable optical layer has ramifications for the higher electronic layers, especially in an IP-over-optical environment. The optical layer can be reconfigured to deliver bandwidth where needed by the IP layer; however, the impact on the IP layer must be considered. First, existing capacity between IP routers can be increased through the provisioning of additional wavelengths that are routed over the same path as the existing capacity (i.e., another wavelength is added to the trunk between two routers). This has minimal effect on the IP layer as the router adjacencies remain intact. Second, the capacity between two adjacent routers may be expanded or shifted by routing wavelengths over new paths. This can affect parameters such as the latency between the routers; this new state information would need to be disseminated in the IP layer. A much more disruptive operation is when the optical layer is reconfigured such that IP router adjacencies are changed. This change in the IP-layer topology can lead to convergence issues with the IP routing protocol; thus, the consequences of such changes need to be taken into account.

A configurable optical layer also allows for dynamic cross-layer restoration, where the optical layer is reconfigured at the time of failure to provide an alternate path for the IP traffic. This has been shown to be more cost effective than IP-layer protection alone, mainly due to a reduction in the number of required IP router ports [52], [53].

4) Software Control

An integral component of a configurable network is the software that automated remote configuration of the network equipment. As traffic demand grew, so did the complexity associated with configuring the network. This complexity exposed a critical weakness with the network provisioning processes – a significant dependence on manual operations – that resulted in long network configuration times. In the late 1990s control automation software was introduced into the optical network for the first time, as the optical control plane [54].

The optical control plane comprises a set of software applications that reside on each network element to distribute decision-making processes. These functions automate many of the traditional manual processes that were performed when configuring the network. When automated, the time taken to reconfigure the optical network was substantially reduced from months to seconds.

The optical control plane typically supports four capabilities: discovery, routing, path computation, and signaling. Discovery is the mechanism where each network element learns

the network address and networking capabilities of its adjacent neighbors. The significance of discovery is that network inventory, configuration and resource utilization data is now automatically generated, resulting in accurate database entries. In fact, with discovery, the network becomes the database of record, allowing network management systems to quickly generate accurate topological graphs without the challenge of database entry keystroke errors. Routing is the method used to disseminate this resource and configuration knowledge to all other network elements in the network. Part of the routing function includes path computation, which uses the global knowledge received from all other network elements to calculate appropriate paths for network connections (usually using a shortest path algorithm such as Dijkstra [55]). Finally, signaling is the mechanism whereby network elements communicate the setup or tear down of optical connections across the network. Because signaling enables the rapid reconfiguration of network resources, it became a key enabler for the successful commercialization of fast mesh restoration in the optical network. Modern optical control plane protocols now perform this function and restore service within hundreds of milliseconds.

For the last few years, optical control plane functions have been the subject of intense standardization activities within the ITU and the IETF. The ITU has developed a suite of recommendations under the title Automatically Switch Optical Network (ASON) [56] that primarily focuses on the architecture and requirements for control plane enabled networks. Definitions for the Optical User-Network Interface (O-UNI) and Network-Network Interface (NNI) are examples of control plane interfaces that have been developed within this framework. Within the IETF, IP/MPLS protocols have been enhanced with traffic engineering extensions in support of Generalized Multiprotocol Label Switching (GMPLS) [57]. To ensure interoperability between different vendors' software implementations, the Optical Internetworking Forum (OIF) has produced interoperability agreements [58] for both O-UNI and Exterior NNI (E-NNI) interfaces.

To date, standards activities have focused on optical control plane for electrically switched networks using SONET/SDH or OTN. However, with the significant rise in popularity of reconfigurable all-optical network devices such as multi-degree ROADMs, work on a control plane for purely optical switched network is under consideration within the above standards bodies. Because of the analog nature of optical propagation, this work is still in the early stages of definition. Consequently, today's control plane implementations for the all-optical layer remain proprietary.

B. Dynamic Networks

Moving forward, the requirements of evolving services are driving the network from being configurable to being dynamic. In the configurable model, a human generally initiates the provisioning process, e.g., through the use of a planning tool. In a dynamic model, not only is the provisioning process automated, but it also is completely under software control. The higher layers of the network automatically request bandwidth from the optical layer, which is then reconfigured accordingly. Services may be provisioned and brought down in

seconds, or possibly sub-seconds. This rapid response time is needed for future applications such as distributed computing, or future bandwidth-intensive, real-time, collaborative applications involving people, display devices, computers, and storage devices, all distributed over several network nodes. Establishing all-to-all connectivity among the nodes involved would be too expensive. However, providing connectivity only when needed, in a dynamic way, can result in great cost savings.

VI. NETWORK CAPACITY

There are generally two approaches for increasing the fiber capacity: increase the number of wavelengths supported on a fiber or increase the bit-rate of each wavelength. Historically, both approaches have been used. In the late 1980s, quasi-WDM was implemented, where just two wavelengths (at 1310 nm and 1550 nm) were supported on a fiber, each one carrying rates of tens of Mb/s. By the mid-1990s, 8 to 16 2.5-Gb/s wavelengths were supported on a fiber, where the wavelengths were located in the 1500-nm region of the spectrum, with wavelength spacing on the order of 100 GHz to 400 GHz. This rapidly increased to 80 to ~200 10-Gb/s wavelengths by year 2000, with 25-GHz to 50-GHz wavelength spacing. Current benchmarks are 40 to 80 40-Gb/s wavelengths per fiber. Thus, in roughly a 20-year span, the capacity per fiber has increased by more than four orders of magnitude.

Increasing the fiber capacity beyond these numbers will require advanced modulation formats, e.g., multi-level amplitude and/or phase modulation [59], [60]. Assuming that higher capacity is possible, it is interesting to consider whether it is more desirable to increase the number of wavelengths or increase the bit-rate of each wavelength.

A. Increasing the Bit-Rate of a Wavelength

Thus far, the transport bit-rates have followed the SONET/SDH hierarchy, with each successive bit-rate representing a four-fold increase; e.g., 2.5 Gb/s to 10 Gb/s to 40 Gb/s. Historically, one of the advantages of increasing the bit-rate has been cost. With each quadrupling of the bit-rate, the cost of the associated transponders has increased by a factor of 2 to 2.5, yielding a steadily decreasing cost per bit/sec. Furthermore, the power and space requirements per bit/sec have decreased as well with increasing bit-rate, thus improving the network operating costs. Given the challenges posed by high-speed electronics, however, it is not clear that this trajectory will continue. While the cost target for a 40-Gb/s transponder is 2.5 to 3 times more than a 10-Gb/s transponder, the current cost is actually more than four times greater (i.e., increased cost per bit/sec). Moving beyond 40 Gb/s will be even more difficult. Originally, the industry favored 160 Gb/s as the next transport rate; however, the standards bodies have now settled on 100 Gb/s instead to coincide with the Ethernet hierarchy.

A second advantage of increasing the bit-rate relates to switching. Photonic switches are generally limited in the number of supported ports; e.g., 320x320. For a given level of traffic, increasing the bit-rate decreases the number of wavelengths entering a node (assuming the wavelengths are well packed), resulting in a smaller required switch size. Thus,

from a switching perspective, increasing capacity through increased bit-rate is more readily scalable.

One of the disadvantages of higher-rate wavelengths is that the signals are more susceptible to optical impairments such as chromatic dispersion, polarization mode dispersion (PMD), intra-cross-phase modulation (IXPM), and intra-four-wave mixing (IFWM) [61], [62]. Thus, the optical reach decreases as the bit-rate increases. For example, a system may have a 4,000-km reach for 2.5-Gb/s signals, a 3,000-km reach for 10-Gb/s signals, and a 2,000-km reach for 40-Gb/s signals. Shorter reach translates to more regeneration, which may nullify the benefits achieved through lower-cost-per-bit/sec transponders [38].

Another disadvantage to increasing the wavelength bit-rate stems from the fact that the required bit-rate of most services is less than that of a full wavelength. Thus, these subrate services need to be bundled together, i.e., groomed, in order to pack the wavelengths efficiently. Grooming is currently implemented in the electronic domain, typically using fine-granularity O-E-O switches. For a given traffic set composed of subrate demands, a higher line-rate requires that the traffic undergo more grooming to achieve the same, or close to the same, level of wavelength efficiency. This translates to larger grooming switches and more O-E-O terminating equipment, all of which adds to the cost, space, and power requirements.

However, a higher bit-rate does provide an advantage with respect to the grooming of bursty traffic. With a low bit-rate, a small number of flows are carried on a wavelength, where the wavelength can only be partially filled to account for the burstiness of the traffic. With a higher bit-rate, more flows are packed together, allowing the system to take advantage of statistical multiplexing. As a percentage of the wavelength capacity, less spare capacity needs to be reserved to accommodate the traffic burstiness, thereby resulting in an overall higher rate of network efficiency.

B. Increasing the Number of Wavelengths

Rather than increasing the bit-rate, one can boost the fiber capacity by increasing the number of wavelengths while decreasing the spacing between channels and maintaining, or possibly lowering, the bit-rate. For example, in the future, this might imply implementing a 1000x10-Gb/s system as opposed to a 100x100-Gb/s system. The biggest advantage of this approach is the reduced need for electronic grooming. More services are likely to be delivered to the network that are already at the rate of a wavelength, thereby requiring no grooming. For those services with rates less than that of a wavelength, an appreciably smaller amount of electronic grooming is required to achieve efficiently packed wavelengths.

It is difficult to predict the overall effect on optical reach with high-wavelength counts. While the bit-rate may be lower, which lowers the susceptibility to the impairments mentioned earlier, the closer channel spacing may lead to problems with impairments such as crosstalk and cross-phase modulation (XPM). Additionally, functions such as receiving a signal may be more difficult with the tight spacing. Techniques such as coherent detection at the receiver may be needed.

While this approach may not appear to offer the advantages in cost, power, and space historically realized through higher-speed transponders, and may present scaling challenges for wavelength-level switches, there are alternative means of addressing these issues. To address the first set of issues, one can consider the PIC approach mentioned in Section IV.D, where several transponders are integrated on a chip. For example, one chip could include ten 10-Gb/s transponders. A goal would be to have this chip be no more costly, and require no more space or power, than a single 100-Gb/s transponder. Note that, in principle, PIC technology does not imply an O-E-O architecture; however, supporting tunability with integration may be challenging, which would limit its efficacy in a network with optical bypass. Additionally, meeting the tight spacing between channels (e.g., 10-GHz) may be difficult with PIC technology. As an alternative, hybrid integration of individual, miniature, low-cost components may be a more suitable solution rather than monolithic integration.

With respect to the challenges of switching, one solution is to implement waveband switching, where groups of wavelengths are treated as a single unit for switching purposes, thereby reducing the size of the switch fabric. Waveband switching has been considered in the past for economic reasons, especially in the cost-sensitive metro area. One objection that has been raised to coarse switching is that it may reduce the flexibility of the network somewhat. However, studies have shown that, with good algorithms, the network efficiency lost due to waveband switching as opposed to wavelength switching is small [63]. Furthermore, if the options are switching ten 10-Gb/s wavelengths as a unit versus one 100-Gb/s wavelength, the switching granularities are the same.

Moreover, one can improve the flexibility of waveband switching by implementing a hierarchical switch architecture (e.g., [64], [65]), with waveband grooming. In a two-level switch hierarchy, the bulk of the switching occurs at the waveband level, but a small amount of switching is also supported at the wavelength level to improve the network efficiency. The wavelength-level switch allows the wavebands to be groomed. For example, the frequency of a specific wavelength can be converted, which moves the signal to another waveband. Conversion may also be possible at the waveband level, where the frequencies of all wavelengths comprising a band are shifted as a unit [66], [67]. It is expected that these functions of wavelength and waveband conversion will ultimately be performed in the optical domain. Essentially, the paradigm of increased number of wavelengths, at possibly lower rate, trades O-E-O wavelength grooming for all-optical waveband grooming.

C. Flexible Bit-Rate Wavelengths

One of the drawbacks of a system composed of purely low bit-rate wavelengths is that inverse multiplexing is required for the services that require a higher rate. With inverse multiplexing, multiple wavelengths are used to carry a single service; e.g., four 10-Gb/s wavelengths are used to carry one 40-Gb/s service. To avoid the added operational complexity of this solution, an alternative approach is to support a mix of bit-rates on a single transmission system; e.g., a range from 10-Gb/s to 100-Gb/s wavelengths. This allows the bit-rate of the

wavelength to better match the service that is being transported. Thus, low bit-rate services can undergo minimal grooming, whereas high bit-rate services can be carried without inverse multiplexing. Depending on the distribution of the wavelength bit-rates, the switches may be able to operate on a wavelength granularity. However, if wavebands are used, then one can consider variable sized wavebands that adjust depending on the system configuration [68].

VII. CONCLUSION

Clearly, networking has undergone tremendous changes over the past 25 years. This has been driven by the push-pull of services and networking capabilities. With a chiefly homogeneous service (i.e., voice), network evolution was centered on providing more capacity in a cost-effective manner. With respect to transmission, this meant more wavelengths and higher line-rates. With respect to switching, this generally meant switching traffic at coarser granularities when possible.

However, the surge in data, and, more recently, video services, has taken networking in different directions. Providing large-capacity static pipes was no longer sufficient to meet the needs of the application layer. This led to the introduction of a configurable optical layer and the control plane, which has further encouraged the growth of more dynamic services.

In addition to covering the past and present, this paper has also speculated on future evolution; for example, more widespread adoption of OTN technology to simplify network operations, a greater role for Ethernet, increased use of optics to achieve cost efficiencies, and a mix of wavelength line-rates (possibly in conjunction with wavebands) to better match the service granularities. Whatever direction networks follow, it is clear that steady traffic growth, with a greater diversity of services, will continue. To meet this growth, carriers will continue to seek technologies that provide cost, scalability, and operational advantages.

REFERENCES

- [1] A. Odlyzko – MINTS. Available: <http://www.dtc.umn.edu/mints/home.html>
- [2] "Ethernet Services Definitions - Phase I," Metro Ethernet Forum, 2004, MEF 6. Available: <http://www.metroethernetforum.org/PDFs/Standards/MEF6.pdf>.
- [3] C. A. Siller, Jr. and M. Shafi, *SONET/SDH: A Sourcebook of Synchronous Networking*, New York: IEEE Press, 1996.
- [4] "Circuit emulation service definitions, framework and requirements in metro Ethernet networks," Metro Ethernet Forum, 2004, MEF 3. Available: <http://www.metroethernetforum.org/PDFs/Standards/MEF3.pdf>.
- [5] <http://grouper.ieee.org/groups/802/3/ba/index.html>
- [6] "2007 Worldwide Interoperability Demonstration: On-Demand Ethernet services across global optical networks," <http://www.oiforum.com/public/downloads/OIF-White-Paper-2007.pdf>
- [7] "2005 OIF Worldwide Interoperability Demonstration: Setup of Ethernet services over global optical networks," http://www.oiforum.com/public/downloads/2005InterOpDemoWhitePaper_FINAL.pdf
- [8] "2004 OIF Worldwide Interoperability Demonstration," <http://www.oiforum.com/public/downloads/Demo-WPaper-FINAL.pdf>
- [9] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol Label Switching Architecture," RFC 3031, January 2001.
- [10] Connectivity Fault Management, IEEE 802.1ag, 2007.
- [11] ITU-T Y.1731, "OAM functions and mechanisms for Ethernet based networks," 2006.
- [12] ITU-T G.8110.1, "Architecture of Transport MPLS (T-MPLS) Layer Network," 2006.
- [13] Provider Backbone Bridge Traffic Engineering, IEEE 802.1Qay, Draft 2.0, 2008.
- [14] ITU-T G.872, "Architecture of Optical Transport Networks," 2001.
- [15] ITU-T G.709, "Interfaces for the Optical Transport Network," 2003.
- [16] ITU-T G. 975, "Forward error correction for submarine systems," 2000.
- [17] 10 Gigabit Ethernet, IEEE 802.3ae, 2002.
- [18] ITU-T G.sup43, "Transport of IEEE 10G Base-R in Optical Transport Networks (OTN)," 2006.
- [19] R. Ramaswami and K. N. Sivarajan, *Optical Networks: A Practical Perspective*, 2nd Edition, Morgan Kaufmann, San Francisco, 2001.
- [20] J. M. Simmons, *Optical Network Design and Planning*, Norwell, MA: Springer, 2008.
- [21] J. M. Simmons, "On determining the optimal optical reach for a long-haul network," *J. Lightw. Technol.*, vol. 23, no. 3, pp. 1039-1048, Mar. 2005.
- [22] H. P. Sardesai, Y. Shen, and R. Ranganathan, "Optimal WDM layer partitioning and transmission reach in optical networks," *Optical Fiber Commun. (OFC)*, Anaheim, CA, Mar. 6-11, 2005, Paper OTuP4.
- [23] A. A. M. Saleh, "Transparent optical networking in backbone networks," *Optical Fiber Commun. (OFC)*, Baltimore, MD, Mar. 7-10, 2000, Paper ThD7.
- [24] M. Sharma, et al., "Value of agile transparent optical networks," *15th Annual Meeting of the IEEE LEOS*, Nov. 10-14, 2002, pp. 293 - 294.
- [25] J. M. Simmons, "Cost vs. capacity tradeoff with shared mesh protection in optical-bypass-enabled backbone networks," *Optical Fiber Comm./National Fiber Optic Engineers Conf. (OFC/NFOEC)*, Anaheim, CA, Mar. 25-29, 2007, Paper NThC2.
- [26] B. Manseur and J. Leung, "Comparative analysis of network reliability and optical reach," *National Fiber Optic Engineers Conf. (NFOEC)*, Orlando, FL, Sep. 7-11, 2003.
- [27] H. Zang, J. P. Jue, and B. Mukherjee, "A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks," *Optical Networks Magazine*, vol. 1, no. 1, pp. 47-60, Jan. 2000.
- [28] S. Subramaniam, M. Azizoglu, and A. K. Somani, "All-optical networks with sparse wavelength conversion," *IEEE/ACM Trans. Netw.*, vol. 4, no. 4, pp. 544-557, Aug. 1996.
- [29] W. Van Parys, et al., "Quantifying the benefits of selective wavelength regeneration in ultra long-haul WDM networks," *Optical Fiber Commun. (OFC)*, Anaheim, CA, Mar. 19-22, 2001, Paper TuT4.
- [30] J. M. Simmons, "Analysis of wavelength conversion in all-optical express backbone networks," *Optical Fiber Commun. (OFC)*, Anaheim, CA, Mar. 17-22, 2002, Paper TuG2.
- [31] D. C. Kilper, et al., "Optical performance monitoring," *J. Lightw. Technol.*, vol. 22, no. 1, pp. 294-304, Jan. 2004.
- [32] A. E. Willner, "The optical network of the future: Can optical performance monitoring enable automated, intelligent and robust systems?," *Optics and Photonics News*, pp. 30-35, Mar. 2006.
- [33] P. Hofmann, et al., "DWDM long haul network deployment for the Verizon GNI nationwide network," *Optical Fiber Commun./National Fiber Optic Engineers Conf. (OFC/NFOEC'05)*, Anaheim, CA, Mar. 6-11, 2005, Paper OTuP5.
- [34] M. Bortz, "Broadwing's experience with optical network planning and deployment," *Optical Fiber Commun./National Fiber Optic Engineers Conf. (OFC/NFOEC'06)*, Anaheim, CA, Mar. 5-10, 2006, Paper NThH1.
- [35] K. Tse, "AT&T's Photonic Network," *Optical Fiber Commun./National Fiber Optic Engineers Conf. (OFC/NFOEC'08)*, San Diego, CA, Feb. 24-28, 2008, Paper NMC1.
- [36] D. F. Welch, et al., "The realization of large-scale photonic integrated circuits and the associated impact on fiber-optic communication systems," *J. Lightw. Technol.*, vol. 24, no. 12, pp. 4674-4683, Dec. 2006.
- [37] Y. Chen, C. Qiao, and X. Yu, "Optical burst switching: A new area in optical networking research," *IEEE Network*, vol. 18, no. 3, pp. 16-23, May/June 2004.
- [38] A. A. M. Saleh and J. M. Simmons, "Evolution toward the next-generation core optical network," *J. Lightw. Technol.*, vol. 24, no. 9, pp. 3303-3321, Sep. 2006.
- [39] O. Leclerc, et al., "Optical regeneration at 40 Gb/s and beyond," *J. Lightw. Technol.*, vol. 21, no. 11, pp. 2779-2790, Nov. 2003.
- [40] J. Leuthold, J. Jaques, and S. Cabot, "All-optical wavelength conversion and regeneration," *Optical Fiber Commun. (OFC)*, Los Angeles, CA, Feb. 22-27, 2004, Paper WN1.

- [41] C.-W. Chao, G. Fuoco, and D. Kropfl, "FASTAR platform gives the network a competitive edge," *AT&T Technical Journal*, vol. 73, no. 4, pp. 69-81, Jul./Aug. 1994.
- [42] I. Foster and C. Kesselman, Editors, *The Grid 2: Blueprint for a New Computing Infrastructure*, 2nd Edition, San Francisco: Morgan Kaufmann, 2004.
- [43] G. Karmous-Edwards, "Today's optical network research infrastructures for e-science applications," *Optical Fiber Commun./National Fiber Optic Engineers Conf. (OFC/NFOEC'06)*, Anaheim, CA, Mar. 5-10, 2006, Paper OWU3.
- [44] P. Duggan, "Tunable filters provide flexibility for agile networks," *WDM Solutions*, Feb. 2003, pp. 9-12.
- [45] K. Ennser, et al., "Reconfigurable add/drop multiplexer design to implement flexibility in optical networks," *Int. Conf. on Transparent Optical Networks (ICTON) 2006*, pp. 74-77.
- [46] M. C. Wu, O. Solgaard, and J. E. Ford, "Optical MEMS for lightwave communication," *J. Lightw. Technol.*, vol. 24, no. 12, pp. 4433-4454, Dec. 2006.
- [47] I. Tomkos, et al., "Ultra-long-haul DWDM network with 320x320 wavelength-port 'Broadcast & Select' OXCs," *European Conf. on Optical Commun. (ECOC)*, Copenhagen, Denmark, Sep. 8-12, 2002, vol. 5, pp. 1-2.
- [48] A. R. Pratt, et al., "40 x 10.7 Gbit/s DWDM transmission over a meshed ULH network with dynamically reconfigurable optical crossconnects," *Optical Fiber Commun. (OFC)*, Atlanta, GA, Mar. 23-28, 2003, Paper PD9.
- [49] S. Mechels, et al., "1D MEMS-based wavelength switching subsystem," *IEEE Communications Magazine*, vol. 41, no. 3, pp. 88-94, Mar. 2003.
- [50] V. Kaman, et al., "Comparison of wavelength-selective cross-connect architectures for reconfigurable all-optical networks," *Int. Conf. on Photonics in Switching*, Crete, Greece, Oct. 16-18, 2006, pp. 1-3.
- [51] R. Batchellor and O. Gerstel, "Cost effective architectures for core transport networks," *Optical Fiber Commun./National Fiber Optic Engineers Conf. (OFC/NFOEC'06)*, Anaheim, CA, Mar. 5-10, 2006, Paper PDP42.
- [52] M. Pickavet, et al., "Recovery in multilayer optical networks," *J. Lightw. Technol.*, vol. 24, no. 1, pp. 122-134, Jan. 2006.
- [53] C. Chigan, G. W. Atkinson, and R. Nagarajan, "Cost effectiveness of joint multilayer protection in packet-over-optical networks," *J. Lightw. Technol.*, vol. 21, no. 11, pp. 2694-2704, Nov. 2003.
- [54] G. Bernstein, B. Rajagolopan and D. Saha, *Optical Network Control, Architecture Protocols and Standards*, Reading: Addison-Wesley Professional, 2003.
- [55] E. W. Dijkstra, "A note on two problems in connection with graphs," *Number Math*, vol. 1, 1959, pp.269 – 271.
- [56] ITU-T G.8080, "Architecture for the automatic switched optical networks (ASON)", 2003.
- [57] IETF RFC3945, "E. Mannie, Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", 2004.
- [58] <http://www.oiforum.com/public/impagreements.html>
- [59] M. Suzuki and I. Morita, "High spectral efficiency for large-capacity optical communication systems," *Int. Conf. on Transparent Optical Networks*, Wroclaw, Poland, July 4-8, 2004, pp. 171 – 176.
- [60] E. Ip and J. Kahn, "Carrier synchronization for 3- and 4-bit-per-symbol optical transmission," *J. Lightw. Technol.*, vol. 23, no. 12, Dec. 2005, pp. 4110-4124.
- [61] H. Kogelnik, et al., "Polarization-Mode Dispersion," in *Optical Fiber Telecommunications IV B*, I. Kaminow and T. Li, Editors, New York: Elsevier Science, 2002.
- [62] P. Bayvel and R. Killely, "Nonlinear Optical Effects in WDM Transmission," in *Optical Fiber Telecommunications IV B*, I. Kaminow and T. Li, Editors, New York: Elsevier Science, 2002.
- [63] P. Bullock, C. Ward, and Q. Wang, "Optimizing wavelength grouping granularity for optical add-drop network architectures," *Optical Fiber Commun. (OFC)*, Atlanta, GA, Mar. 23-28, 2003, Paper WH2.
- [64] K. Harada, et al., "Hierarchical optical path cross-connect systems for large scale WDM networks," *Optical Fiber Commun. (OFC)*, San Diego, CA, Feb. 21-26, 1999, Paper WM55.
- [65] A. A. M. Saleh and J. M. Simmons, "Architectural principles of optical regional and metropolitan access networks," *J. Lightw. Technol.*, vol. 17, no. 12, pp. 2431-2448, Dec. 1999.
- [66] N. Chi, et al., "All-optical wavelength conversion and multichannel 2R regeneration based on highly nonlinear dispersion-imbalanced loop mirror," *Photon. Technol. Lett.*, vol. 14, no. 11, pp. 1581-1583, Nov. 2002.
- [67] Y.-K. Huang, et al., "Simultaneous all-optical 3R regeneration of multiple WDM channels," *18th Annual Meeting of the IEEE LEOS*, Sydney, Australia, Oct. 23-27, 2005, Paper MM2.
- [68] R. Izmailov, et al., "Non-uniform waveband hierarchy in hybrid optical networks," *IEEE INFOCOM 2003*, San Francisco, CA, Mar. 30-Apr. 3, 2003, vol. 2, pp.1344-1354.