All-Optical Networking – Evolution, Benefits, Challenges, and Future Vision

Adel A. M. Saleh, *Fellow, IEEE* and Jane M. Simmons, *Fellow, IEEE*

Invited Paper

***Abstract*— While all-optical networking had its origins in the research community a quarter of a century ago, the realization of the vision has not had a straight trajectory. The original goal of the *all-optical network* was based on keeping the data signals entirely in the optical domain from source to destination to eliminate the so-called *electronic bottleneck*, and to allow arbitrary signal formats, bit-rates, and protocols to be transported. The latter property is referred to as *transparency*. When all-optical networks were finally commercialized around the turn of the century, however, a modified reality emerged; the quest for transparency was replaced by the more pragmatic objective of reducing the network cost and energy consumption. Moreover, especially for networks of large geographical extent, electronics were still present at some (relatively few) points along the data path, for signal regeneration and traffic grooming. This modified vision captures the state of today's networks, though terms like *all-optical* and *transparent* are still used to describe this technology. However, continued advancements are bringing back some aspects of the original transparency vision. In this paper, we review the evolution of all-optical networking, from the early vision to its present vibrant state, which was made possible by great advances in optical transmission and all-optical switching technologies. We describe the numerous benefits afforded by the technology, and its relative merits and drawbacks compared to competing technologies, sometimes referred to as *opaque*. We also discuss the remaining challenges and future directions of all-optical networking. While all-optical solutions permeate today's access, metro, and core networks, this paper focuses on the core.**

***Index Terms*— All-optical networks, core networks, flexible spectrum, OEO networks, optical reach, regeneration, ROADM, transparency, wavebands, wavelength selective switches**

I. Introduction

All-optical networking is founded on the premise of maintaining a network connection in the optical domain from its source to its destination, thereby removing the intermediate electronics, which tend to be more costly and less scalable than optics. First proposed a quarter of a century ago, the field of all-optical networking has undergone several embodiments as researchers have probed the practical realities and limitations of optical technology. While in its present state it has eliminated much, but certainly not all, of the electronics along a data path, it nevertheless represents a major paradigm shift that has profoundly affected the design, operation, and economics of networks. All-optical solutions permeate today's access, metro, and core networks; however, the focus of this paper is on core (or national-scale) networks.

To appreciate the benefits and challenges of all-optical networking, it is instructive to review the state of transmission and switching in the optical layer in the early 1990s. Section II outlines the architecture of these legacy networks, highlighting the drawbacks that spawned the all-optical vision. Section III discusses very early research into key principles for realizing all-optical networks. The state of all-optical networking was further advanced in the 1990s, largely by government-funded programs; some of the major innovations produced by this research are summarized in Section IV.

All-optical networks were finally commercialized in the late 1990s. The commercial reality was somewhat different from the early vision, as practical limitations needed to be addressed. Section V details the first commercialization efforts and why all-optical networks did not successfully remove all of the intermediate electronics in a data path. (This did not stop these networks from being referred to as 'all-optical'.)

Support for all-optical networking was not universal, primarily due to concerns over the complexity of operating such networks. Alternative solutions for addressing some of the drawbacks of legacy networks were developed, chiefly among them being photonic integration; this is covered in Section VI.

All-optical networks have persevered, however, with new optical technologies, architectures, and algorithms that have improved the performance and simplified the operation of such networks. Today, most commercial and government core networks include all-optical networking technology. The current state of optical technology is discussed in Section VII.

Finally, Section VIII discusses where all-optical networking is heading, including a return to some of the earlier vision.

II. Legacy Networks vs. All-Optical Networks

By the early 1990s, two key innovations had already been realized in optical transmission. First, multiple channels of light, or wavelengths, could be multiplexed together onto a single fiber, giving rise to wavelength division multiplexing (WDM). Second, the erbium-doped fiber amplifier (EDFA) was developed to amplify all wavelengths carried on a fiber, rather than needing one amplifier per wavelength. Together, these two innovations dramatically increased the network capacity in a cost-effective manner.

While optical transmission benefited from these technologies, complementary switching breakthroughs were largely absent from the network nodes. (The nodes of a core network are typically the major cities at which traffic is sourced, termi-

Adel A.M. Saleh was with DARPA, and is now at UCSB, Santa Barbara, CA. Jane M. Simmons is with Monarch Network Architects, Holmdel, NJ. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

nated, and switched.) Figure 1(a) depicts a simplified version of a typical nodal architecture in this timeframe. The node in the figure is assumed to sit at the intersection of three WDM fiber routes; i.e., a *degree-three* node. (Each fiber line shown in the figure actually represents two fibers, one for each direction of transmission.) The incoming WDM fiber is terminated in a demultiplexer (demux) that separates the WDM signal into its constituent wavelengths. Each wavelength is then received by a *transponder* (TxRx), which converts the optical signal to the electronic domain, cleans up the signal, and then converts it to an optical signal suitable for short intra-nodal transmission. The signal is then converted back to the electrical domain for processing by an electronic switch (e.g., a Synchronous Optical Network (SONET) switch, IP router, etc.). The process is reversed for the outgoing traffic. Eventually, the traffic passes through another transponder as it exits the node, which converts the optical signal to a wavelength suitable for WDM transmission. (Note that a transponder includes both a transmitter and receiver.) The outgoing wavelength assigned to a connection does not have to be the same as the incoming wavelength. The outgoing wavelengths are then combined in a WDM multiplexer (mux) onto an outgoing fiber. This architecture is known as *optical-electrical-optical* (OEO), due to its transitions between the two domains.

It is important to distinguish between two types of traffic that enter a node: *bypass* traffic and *drop* traffic. *Bypass traffic* transits the node on its way to its ultimate destination. *Drop traffic* carries data that is actually destined for the node; i.e., it is 'dropped' from the optical layer to a higher layer (e.g., IP) at this node (this usage of the term 'drop' is unrelated to the notion of 'dropping a packet', which is an undesirable event). Similarly, *add* traffic is the traffic sourced by a higher layer at the node. Regardless of the traffic type, in the OEO architecture of Fig. 1(a), transponders are needed at a node for *every* incoming wavelength and for *every* outgoing wavelength. Thus, a cross-country connection, as exemplified in Fig. 1(b), would pass through tens of transponders.

The large number of required transponders yields nodal solutions that are expensive, require a lot of power and space, and that present reliability issues. It also requires the electronic switch fabric, which is difficult to scale, to grow in concert with the network traffic. The recognition of this nodal 'electronic bottleneck' was one impetus for the all-optical vision, where a network connection is maintained in the optical domain from source to destination.

The concept of an all-optical node is illustrated in Fig. 2(a), where the key component is the all-optical switch. Such a device has all-optical interfaces, and switches the various wavelengths in the optical domain. Thus, those wavelengths carrying bypass traffic can remain in the optical domain through the node. This is referred to as *optical bypass*. Transponders are needed only for the wavelengths that are dropped or added at this node, as shown in Fig. 2(a).

With all-optical nodes, the resulting picture of an end-to-end connection is as shown in Fig. 2(b), where transponders are needed only at the endpoints. Clearly, this technology enables the removal of a tremendous amount of electronics from the network, and the attendant cost, power, space, and reliability burdens. Furthermore, all-optical core switches, by operating on wavelengths, are more scalable than their electronic counterparts. Another benefit of the all-optical vision is the provisioning time for a new connection is greatly decreased, as equipment needs to be installed only at the endpoints, as opposed to at every node along the path.

While providing many advantages, all-optical networks do pose some technical and operational challenges. As these challenges were gradually uncovered, new algorithms and operational paradigms were developed to maintain the viability and advantages of all-optical networking. This evolutionary development process permeates the discussion below.

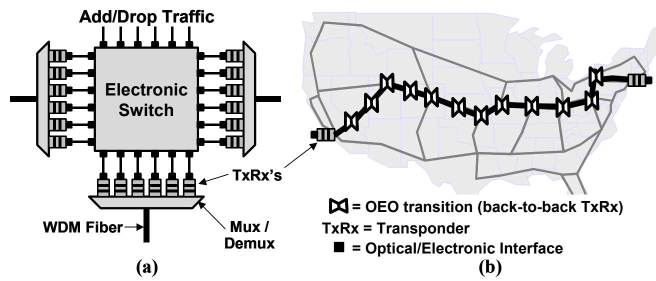

Fig. 1. (a) A degree-three OEO nodal architecture. (b) A cross-country single-wavelength connection, which requires tens of intermediate OEO transitions (i.e., back-to-back transponders). In both figures, as well as in Fig. 2, each fiber line shown is actually two fibers, one for each direction of transmission. Also, optical amplifiers at the node and along the fiber path (typically inserted every 60 to 100 km) are not shown.

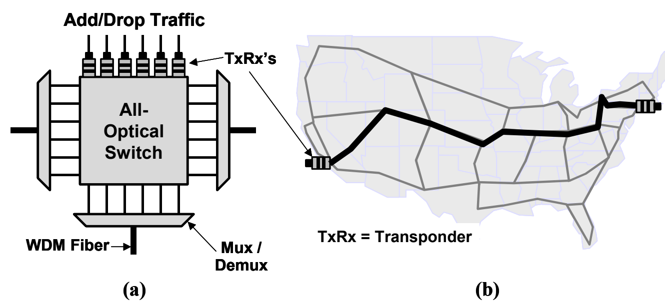

Fig. 2. (a) A degree-three all-optical nodal architecture; the WDM multiplexers/demultiplexers (mux/demux) are typically integrated with the all-optical switch. (b) A cross-country single-wavelength connection, requiring transponders only at the end points.

III. Early Fundamentals of All-Optical Networking

Early research probed important fundamental properties of all-optical networks, many of which, though not all, proved prescient when such networks were finally commercialized.

A. Transparency

Section II enumerated several benefits of all-optical networking, yet they were not the chief motivating factors for initial research in this area. The prime driver of much of the early research was the desire for *transparency*. It was envisioned that by maintaining connections in the optical domain, an all-optical network would be transparent [1], similar to a piece of glass; i.e., any type of traffic could be simply passed through the network. This would provide the ultimate future-proof network that could accommodate any new service or transmission innovation. This notion was captured by referring to an all-optical network as an 'optical ether' [2].

There are several types of transparency. *Digital transparency* refers to a network being compatible with digital signals of any bit-rate, modulation format, or protocol. *Analog transparency* refers to a network that can carry a range of analog signals and deliver them with acceptable fidelity. Finally, s*pectral transparency* refers to a network that offers flexibility in the placement of wavelengths within the optical spectrum.

It must be noted that despite the early emphasis on transparency, it still has not been realized to a large degree in today's core optical networks. However, recent developments may nudge networks in that direction. This topic is revisited in Section VIII.

B. Wavelength-Routing Switches and Wavelength Reuse

The earliest work on all-optical networks typically assumed that the core optical devices in the network nodes were passive splitters, combiners, and broadcasting stars. For example, a 1x2 'passive' (i.e., wavelength independent) splitter takes an optical signal at its input, and sends a copy of that signal to both outputs. If the optical signal is WDM, all wavelengths comprising the signal are transmitted from the input to both outputs, with no means of controlling the individual wavelengths.

Such passive networks support a relatively small number of connections for a given number of wavelengths. This is illustrated by the simple example of Fig. 3(a). Assume that there are only two wavelengths in the WDM signal. Node A is transmitting data to node C on wavelength #1, and transmitting data to node D on wavelength #2. Node B is equipped with a passive broadcasting star; thus both wavelengths are sent to nodes C and D, and also to E even though A is not communicating with this node. (Nodes C and D are required to have an optical filter to receive the proper wavelength.) If E wants to transmit to D, there are no available wavelengths on the fiber to do so, and the call request would be blocked.

This limitation prompted the recognition that scalable all-optical networks require all-optical wavelength-routing switches at the nodes, which are capable of routing individual wavelengths [3][4]. Such a wavelength-routing capability allows for *wavelength reuse*, where multiple connections on disjoint fibers are carried on the same wavelength. The advantage of wavelength reuse can be seen in Fig. 3(b). Node B is equipped with a wavelength-routing switch, such that node A transmits only wavelength #1 to node C and only wavelength #2 to node D. Now, if E wants to transmit to D, it can use wavelength #1, as shown by the dashed line.

All-optical wavelength-routing switches are still one of the key enabling technologies of today's all-optical networks.

C. Two-layer Architecture

In the OEO architecture of Fig. 1(a), the core switch is electronic and typically performs functions other than directing wavelengths from input ports to output ports. For example, an IP router examines each IP packet carried on an incoming wavelength, selects the outgoing router port, and repackages the packets onto outgoing wavelengths. Thus, the IP router is both routing and *grooming* the traffic (grooming is the repackaging of traffic at intermediate nodes of a data path).

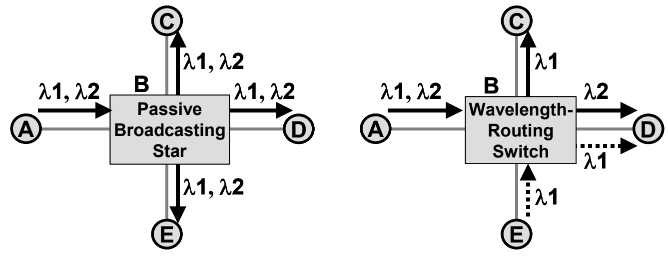

Fig. 3. (a) With a passive broadcasting star at Node B, all wavelengths sent from A are transmitted to C, D, and E, limiting future connections. (b) With a wavelength-routing switch, only those wavelengths meant for C and D are routed there, allowing for wavelength reuse (e.g., λ1).

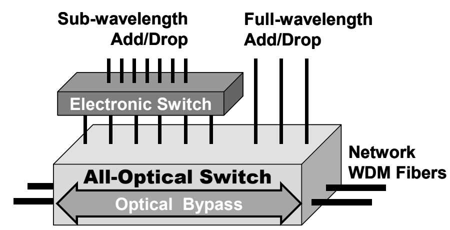

Fig. 4. Two-layer switch hierarchy, with the electronic switch serving as a relatively small edge switch.

While all-optical switches are well suited for routing wavelengths, they are not ideal for performing grooming. Functions such as buffering data or processing packets are difficult to perform with optics. This limitation gave rise to a two-layer switch architecture, where an all-optical switch processes wavelengths, and an electronic switch operates at a sub-wavelength granularity [5].

Figure 4 illustrates such a two-layer switch hierarchy. The all-optical switch is the core switch, through which all traffic enters and exits, whereas the electronic switch is an adjunct 'edge' switch. Only those wavelengths carrying traffic that requires grooming are dropped from the all-optical switch to the electronic switch. While this architecture does not eliminate electronic switches, the size of the electronic switch is greatly reduced as compared to the architecture of Fig. 1(a).

D. Wavelength Continuity Constraint

The resulting network embodiment of the all-optical paradigm is a collection of *lightpaths*. As first defined in [6], a lightpath is a connection between any two network nodes that remains in the optical domain and which is carried on the same wavelength from source to destination. This highlights an important challenge of all-optical networks, namely the *wavelength continuity constraint*. In order to establish a lightpath, the same wavelength must be available (i.e., not used by any other connection) on all links over which the lightpath is routed. Thus, in attempting to establish a new lightpath, a scenario may arise where all links of the lightpath have sufficient available capacity, but where no one wavelength is available on every link of the path, leading to the new lightpath being blocked. This is in contrast to OEO networks where a wavelength can be arbitrarily chosen for each link over which a connection is routed. Thus, it can be expected that the level of blocking in an all-optical network would be somewhat higher than that in an OEO network.

Reference [6] observed that the task of assigning wavelengths to a set of lightpaths using the minimal number of wavelengths is equivalent to the well-known graph-coloring problem, which is NP-complete. Thus, heuristic algorithms are needed for efficient wavelength assignment in all-optical networks, a step that is not needed in OEO networks. Two such algorithms were presented in [6], and were shown to be effective in achieving blocking levels similar to an OEO network; these algorithms are now known as First-Fit and Most-Used, and are still in use in network design tools today [7].

Note that early work on all-optical networks considered the possibility of changing the wavelength of a connection while in the optical domain, thereby eliminating the wavelength continuity constraint. However, even today, all-optical wavelength conversion has not been realized on a large scale, largely due to cost, and the fact that it is limited to simple, spectrally inefficient modulation formats.

IV. Early Research Programs

Significant progress in the field of all-optical networking transpired during the early to mid 1990s due to large research efforts in the US, Europe and Japan by individual companies and by government-supported cross-industry consortia. Some of these efforts are highlighted here.

A. ONTC and AON

The Optical Network Technology Consortium (ONTC) [8] and the All-Optical Networks (AON) consortium [9][10] were launched in the U.S. in 1993 and were supported by the Defense Advanced Research Projects Agency (DARPA). Their goal was to advance the architecture and technologies of core all-optical networks to enable advanced applications for the telecommunications and computer industries as well as for defense networks.

Both consortia identified the need to handle both full wavelength services and sub-wavelength circuit or packet services. To accomplish this, ONTC pursued an architecture consisting of a configurable all-optical WDM core, interconnecting finer-granularity electronic switches, whereby a sub-wavelength connection hops through a number of such switches from source to destination. This action enables traffic grooming at the edge of the network to increase the network efficiency, similar to what was described in Section III.C.

Conversely, AON pursued a pure all-optical WDM architecture consisting of highly dynamic, all-optical tree-networks at the edge, interconnected by a configurable all-optical core. Sub-wavelength circuit and packet services were provided via time-division multiplexing in the optical domain, using fast-tunable WDM transmitters. The time synchronization needed for this approach was made possible by the tree topologies of the edge networks, which enabled simple scheduling.

B. MWTN

At roughly the same timeframe, the Pan European RACE (Research in Advanced Communications for Europe) program formed the Multi-Wavelength Transport Network (MWTN) consortium [5][11]. To a large extent, this consortium pioneered the layered core network architecture used in today's networks, as partially depicted in Fig. 4.

C. Optical Path Layer Technologies Thrust

Similar research was being carried out in Japan [12], which also pursued a layered network architecture. One of the contributions was to distinguish between end-to-end optical paths containing no wavelength conversion and those containing wavelength conversion. The former is called a wavelength path (WP), and the latter a virtual wavelength path (VWP). Various novel architectures were devised for all-optical cross-connects for both WP and VWP path types.

D. MONET

Another DARPA-supported consortium on all-optical networks, which was launched in the U.S. in 1996, is the Multi-wavelength Optical Networking (MONET) consortium [13][14]. Of significance is the use of "networking" (rather than simply "networks") in the program name to emphasize the dynamic nature of the network. Thus, connections in the optical layer were required to be established and torn down in a dynamic manner, not only to perform fast optical-layer restoration from network failures, but also to enable dynamic bandwidth-on-demand across a national-scale network. MONET utilized the layered electrical/optical architecture discussed above, and included research on both the core (or long-distance) network and the regional/metro-core (or local-exchange) networks. A testbed was built, which demonstrated establishing dynamic all-optical connections originating at a local-exchange network, going through a long-distance network, and terminating in another local-exchange network.

Interestingly, a small economic study done within the MONET program indicated that some savings in the cost of the network would result due to reducing the number of transponders. Though, at that time, this result was not given sufficient attention, it turned out later that MONET greatly underestimated the cost savings, and, as discussed below, this aspect of all-optical networks was ultimately the most relevant consequence of this technology.

V. Commercialization

In the mid 1990s, the all-optical vision first entered commercial networks in the form of a nodal element known as an optical add/drop multiplexer (OADM). An OADM is similar to the general all-optical switch shown in Fig. 2(a) in Section II, except that it is specifically designed for degree-two nodes. As with the all-optical switch, traffic that is transiting the node remains in the optical domain, and transponders are needed only for the add/drop traffic. These early OADMs, however, were not flexible, as they allowed only for the add/drop of specific wavelengths.

The first core all-optical *system*, which addressed amplification, transponders, and flexible all-optical switching, was developed by Corvis Corp. in the late 1990s. The first commercial all-optical system deployment was the Broadwing network in 2000 [15]. This groundbreaking system elucidated many of the practical aspects of all-optical networking, as described next. (Several vendors have since developed core all-optical systems, including Alcatel-Lucent, Ciena, Ericsson, Huawei, and Nokia Siemens.)

A. Optical Reach

An important consideration in optical networks is the *optical reach*, which is the distance an optical signal can travel before the signal quality degrades to a level that necessitates regeneration. Regeneration 'cleans up' the signal, and is typically accomplished by passing the optical signal through two back-to-back transponders. This allows the signal to be re-amplified, re-shaped, and re-timed (i.e., '3R' regeneration). As shown in Fig. 1(a), legacy OEO networks naturally perform the regeneration function for all traffic transiting the node; in this figure, an electronic switch is used to direct the signal between incoming and outgoing transponders.

Legacy networks based on EDFA amplifier technology had an optical reach of approximately 500 km. For a link longer than 500 km, it was necessary to regenerate the traffic at intermediate sites along that link (a link is the fiber route connecting two adjacent nodes). Thus, in addition to regenerating traffic at every node in an OEO network, there may be dedicated regeneration sites located along some of the links, further exacerbating problems with system cost and power.

In order to take advantage of the all-optical paradigm, it is necessary to have a system with *extended* optical reach. This is illustrated in Fig. 5. It is assumed that each of the nodes is equipped with an all-optical switch or an OADM; the length of each link is 1,000 km. Consider a connection between nodes A and E. If the optical reach of the system were less than 2,000 km, then optical bypass cannot occur despite the presence of nodal elements that are capable of switching in the optical domain; the signal would need to be regenerated at nodes B, C, and D. (Note that regeneration is performed only at network nodes unless the optical reach is less than the link distance.) If the reach were 3,000 km, the signal could optically bypass B and C, but would still need to be regenerated at D. An optical reach of 4,000 km is required to maintain the connection in the optical domain from A to E.

This example illustrates that both extended optical reach and all-optical switching elements are required to achieve optical bypass in the core. Initial all-optical systems had an optical reach on the order of 3,000 km. This extended reach was accomplished by using: Raman amplifiers instead of EDFA amplifiers, where distributed Raman amplification uses the fiber itself to amplify the optical signal, so that the rate of signal degradation is less steep as compared to EDFA amplification [16]; advanced modulation techniques in the transponders; high quality lasers and filters; dispersion compensation along the fiber; and powerful forward error correction (FEC).

While 3,000 km reach is significantly greater than the legacy 500 km reach, it is not long enough to allow cross-continental connections to remain in the optical domain from source to destination. Any connection that extends further than the reach requires regeneration. This was the first practical realization that commercial all-optical networks would not be truly all-optical. A more accurate name for these networks would in fact be 'optical bypass enabled' [7].

Nevertheless, studies have shown that with an optical reach of 3,000 km, more than 90% of the regenerations are eliminated from a continental-scale network as compared to a legacy network with 500 km reach [7][17]. This significantly reduces the cost, power, and space requirements of a network, thus achieving that part of the vision. However, the presence of some intermediary transponders, which are typically specific to a particular protocol, modulation format, and bit-rate, deviated from the pure transparency vision.

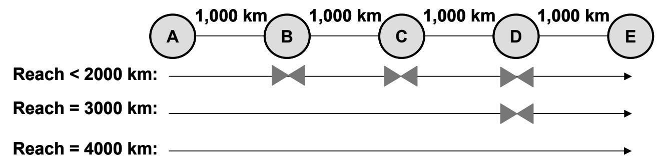

Fig. 5. Even if all nodes are equipped with switches capable of optical bypass, regeneration may be required if the optical reach is not long enough.

It is technologically possible to produce a system with an optical reach of more than 8,000 km, such that all regenerations would be eliminated from a cross-continental scale network. However, to do so would require very costly technology; e.g., complex transponders, elaborate amplification schemes, etc. It has been shown that the cost savings resulting from removing the residual amount of regeneration in a network would not justify the very high cost of this technology [18]. In fact, the analysis of [18] demonstrated that when the extra cost of extended-reach technology is compared to the savings generated by the removal of regenerations, the optimal optical reach in a core network is approximately 2,000 to 2,500 km, which is the reach attained by most commercial 'all-optical' systems today.

Furthermore, as noted in Section III.C, there are network functions best performed in the electronic domain, such as traffic grooming. Any wavelength that needs to undergo grooming would need to be processed electronically, e.g., in a SONET switch or an IP router, such that the end-to-end connection would not remain in the optical domain regardless of the reach. (Note that, typically, only a small fraction of the traffic at each node needs grooming; hence, as mentioned in Section III.C, the required size of the electronic switch is much smaller than the core all-optical switch).

B. Optical Amplifier Transients

An important operational issue in all-optical networks is their susceptibility to optical amplifier transients. Such transients result when there is a sudden change in the power level on a fiber, as may occur when connections are brought up or down. Transmission systems, whether EDFA or Raman based, typically have dynamic controls to dampen such power variations. However, excursions in the signal power level, even if brief, are undesirable as they lead to error bursts. Furthermore, in the presence of optical bypass, transients on one link may have a ripple effect, causing errors to propagate.

Transients due to wavelengths being brought down by a failure are unavoidable. However, transients caused by a system operation, e.g., bringing up a new connection, or restoring service after a failure, are generally unacceptable. Thus, for all-optical networks, it is important that operational procedures be implemented to avoid, or at least minimize, transients, while still taking advantage of optical bypass.

One effective mechanism devised in response to this challenge makes use of *pre-deployed subconnections* [19]. A subconnection typically spans multiple links, and is terminated at either end in a transponder, with optical bypass at the intermediate nodes. The transponders are always lit, even if the subconnection is not carrying traffic. When a new connection is needed, the appropriate subconnections are concatenated together in the electrical domain to form the end-to-end path. Because the subconnections are pre-lit, there is no change in power level on the fiber, thereby avoiding problems with transients. This methodology has been shown to be capacity-efficient, without forfeiting the benefits of optical bypass [20].

C. Fault Localization

One of the benefits of electronically processing the data at every node of an OEO network is detailed error monitoring can be performed, which is important for rapidly locating a failure. Rapid fault localization is more challenging in networks with optical bypass because electronic processing does not occur at most nodes. This is especially true when just one connection, rather than an entire fiber, has failed. One solution is to deploy *optical performance monitors* (OPMs) throughout the network, to check on the health of the optical signals [21]. However, most OPMs are limited in what signal characteristics they can monitor, and tend to be costly. Another approach is to employ failure-independent protection schemes, where the same protection mechanism is triggered for a failed connection regardless of the failure location. Such mechanisms typically take the form of path protection, where the backup path is completely disjoint from the original path. Switchover to the backup path occurs as soon as a connection endpoint detects a failure; fault localization can occur on a much longer timescale. Path protection is the predominant form of protection in today's core networks.

D. Network Design Tools

Software tools are needed for network design, whether the networks are OEO or all-optical. Determining how traffic should be groomed or how traffic should be protected is best calculated using efficient algorithms embodied in a design tool. With all-optical networks, however, there are two additional functions that are usually handled in a design tool.

First, algorithms are needed to determine where to regenerate each connection. For simplicity, optical reach is quoted in terms of a distance (e.g., 2,000 km), but in reality it is based on numerous factors: amplifier type, fiber type, fiber distance and loss, the number of switches optically bypassed, to name just a few. Each vendor's all-optical system has its own engineering rules that must be met. In addition, the topology and fiber plant of each carrier's network impact regeneration decisions. Furthermore, the regeneration decisions made for one connection must be valid regardless of the other connections that may be added or removed in the future.

Despite the underlying complexity, it is possible to come up with relatively simple rules that are suitable for real-time calculation of required regeneration sites [7]. For example, the impact of relevant factors on the *optical-signal-to-noise-ratio* (OSNR) can be tallied, with regeneration required for a connection when the OSNR drops below a prescribed threshold.

Another important design-tool function was touched on in Section III.D: the assignment of wavelengths to connections. Many effective wavelength-assignment algorithms have been developed, beginning with those in [6]. Note that one of the more important aspects of this process is to take advantage of regeneration allowing the wavelength of a connection to be changed, if needed. Furthermore, if a connection would be blocked due to wavelength contention, an extra regeneration may be added for purposes of wavelength conversion [22].

VI. OEO Strikes Back

Despite the progress made in all-optical networks, their acceptance was not universal in the 1990s and early 2000s. Some of the initial skepticism was targeted at the technology. It was assumed that the accumulation of optical impairments would render all-optical systems impractical. In addition, there was pushback regarding the management and operational issues of all-optical networks, which were described in Section V. To emphasize that OEO systems did not pose the challenges of transparent systems, such networks were referred to as *opaque* [23]. (Networks that are mostly all-optical but with some OEO are sometimes called *translucent* [24]).

One of the main areas of contention focused on the impact of the wavelength continuity constraint. As discussed above, if a signal optically bypasses a node, then it must be carried on the same wavelength both entering and exiting the node. This interdependence of wavelengths from one link to another does not exist in OEO networks. However, numerous studies have shown that, if intelligent wavelength-assignment algorithms are used, wavelength contention is not an issue. The analysis of [25] demonstrated that a *small* amount of wavelength conversion is sufficient to achieve blocking levels similar to OEO networks. As pointed out in [22], the small amount of regeneration (and grooming) that is required in an all-optical network provides the opportunity to accomplish the desired level of wavelength conversion.

Another concern was that all-optical architectures are less amenable to multi-vendor deployments, because all-optical interfaces are not standardized (as opposed to electrical interfaces, which are). While strictly true, this did not turn out to be an important point in practice. If needed, OEO regeneration can be used at the boundary between all-optical systems from different vendors, a strategy known as *islands of transparency* [17].

One concern that was borne out regarded the economics of all-optical networking. In fact, all-optical networks are not an economic choice for *all* core networks. As pointed out in Section V.A, the technology to achieve extended optical reach is more costly. For example, Raman amplifiers cost more than EDFA amplifiers; transponders that are capable of 2,500 km reach are more expensive than those that are capable of only 500 km reach. Thus, the cost effectiveness of an all-optical solution depends on the amount of regeneration that is removed, which in turn is dependent on the level of traffic and the traffic pattern. All-optical technology would likely not be economically justified for a network with a relatively low amount of traffic or with relatively short connections.

However, with the Internet burgeoning in the 2000 timeframe, most carriers were faced with an exploding level of traffic. While not all carriers bought into the all-optical vision, the electronic burden of legacy OEO networks was not sustainable. Thus, if the OEO paradigm were to be maintained, it would need to be implemented with different technology, as described next.

A. Photonic Integrated Circuits (PICs)

Optical systems traditionally have been assembled using discrete components, e.g., each transponder may be a separate 'pizza-box' sized card that is plugged into a chassis. This contributes to the cost, power, space, and reliability issues associated with OEO technology. However, in the 2005 time-frame, PIC technology, where numerous WDM components are monolithically integrated on a chip, was introduced to core optical networks by Infinera Corp. [26]. PICs, by miniaturizing and integrating traditional optical components, offer an alternative means of implementing the OEO vision. For example, ten lasers, a multiplexer, and several control components may be integrated on a single PIC transmitter chip.

With PIC technology, the OEO paradigm illustrated in Fig. 1(a) remains; however, the cost, power, space, and reliability burdens are significantly reduced. In that sense, it represents a viable alternative technology to optical bypass. However, one bottleneck that the PIC approach has thus far not addressed is switching, particularly at large network nodes. Core switching remains in the electronic domain [26], similar to Fig. 1(a). Thus, the scalability issues (e.g., in terms of cost, size and power) of core electronic switches are not eliminated.

A more scalable option might be to combine micro-electro-mechanical-system (MEMS) technology [27] with PIC technology to address the switching bottleneck, as shown in Fig. 6. (The trick is to find a compact, cost-effective way to combine, or integrate, these two technologies.) The MEMS switch itself is all-optical, thereby taking advantage of the scalability of optics. While the switch is all-optical, this architecture is still essentially OEO. To be more precise, Fig. 1(a) can be considered OEO-E-OEO, whereas Fig. 6 can be considered OEO-O-OEO. If grooming is desired at the node, then a two-layer switch hierarchy can be used, similar to Fig. 4, with a *small* electronic grooming switch at the edge.

Note that the optical signals switched by the MEMS switch in Fig. 6 are the intra-nodal transmission signals (typically referred to as 'short reach' or 'very short reach'). As a comparison, the all-optical switch of Fig. 2(a) switches the individual wavelengths comprising the WDM long-haul signal.

VII. All-optical Networks Persevere

Despite some of the early doubts regarding all-optical networks, and the competition from PIC/OEO technology, all-optical systems are the predominant technology in today's core networks. The 'perseverance' of all-optical networks is in part due to the development of several enabling technologies that have simplified the operation of these networks.

One of these technologies is the *tunable transponder*, which can be tuned to any of the system wavelengths. As their costs have decreased, tunable transponders have displaced *fixed transponders*, which are capable of generating/receiving only one particular wavelength. The ability to remotely set the transponder to whichever wavelength has been assigned to a

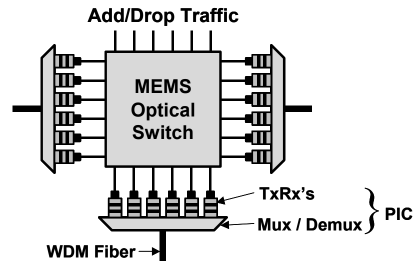

Fig. 6. A scalable, non-blocking OEO-O-OEO network node architecture, which combines PIC and MEMS technologies.

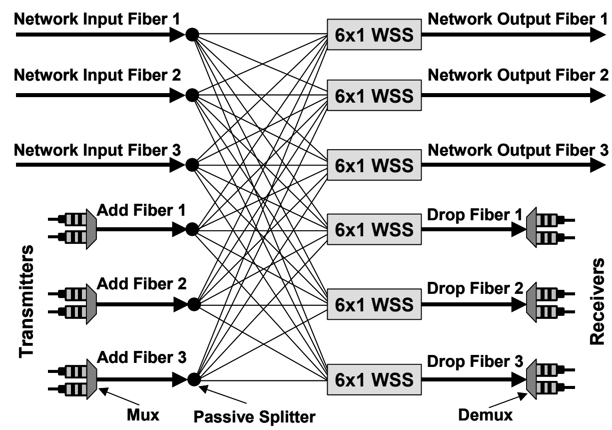

Fig. 7. WSS-based all-optical switch, or ROADM, for a degree-three node.

connection simplifies the network operation. Tunability also alleviates transponder inventory and sparing issues, whether in an all-optical or OEO network.

Other operational benefits have resulted from improved amplifier-transient control [28] and more advanced network design tools. However, the development of greatest impact, with respect to both scalability and operational flexibility, was the all-optical switch based on the 'wavelength-selective-switch' (WSS). Early all-optical switches were designed with a layered architecture, where a separate switching plane existed for each wavelength [9]-[14]. While suitable for early 4- to 16-wavelength systems, a layered approach is untenable for today's 80-wavelength systems. Switch architecture was transformed in 2003 with the introduction of the 1xN WSS [29], which is capable of taking a WDM input and directing *any* of the constituent wavelengths to *any* of the N output ports (an Nx1 WSS performs the reverse operation). (Actually, the concept of a WSS-based all-optical switch was suggested early on in [30].)

The prototypical WSS-based all-optical-switch architecture is shown in Fig. 7, for a (degree-three) node with three network fibers and three add/drop fibers. The WDM signal on each input fiber is broadcast by a passive splitter to the six 6x1 WSSs. Each WSS selects which wavelengths from each input fiber are directed to its associated output fiber.

This switch architecture is very scalable. For an 80-wavelength system, the design shown in Fig. 7 is the equivalent of a 480x480 switch, which is larger than what is available today using competing all-optical technology (e.g., MEMS). Furthermore, the architecture is *directionless*: a transponder on any add/drop fiber can access any of the network fibers. This allows the client attached to the transponder (e.g., an IP router) to establish a connection in any direction.

As WSSs have grown in size, it is also feasible to use a small number of WSSs as the mux/demux on the add/drop fibers [31]. The wavelength-selective property allows any wavelength to be transmitted/received from/to any port of the mux/demux, a feature that is known as *colorless*. Colorless switches are especially important with tunable transponders; otherwise, a transponder would need to be moved to a different mux/demux port if it is tuned to a different wavelength.

Another benefit of the architecture of Fig. 7, which is known as the *broadcast-and-select* architecture, is that it supports multicast, where an input wavelength is directed to multiple output fibers. A final advantage is that it can be upgraded from, say, a degree-three switch to a higher degree switch, without affecting the existing network traffic.

One potential drawback is that wavelength contention can occur on the add and drop fibers. For example, assume an IP router connects to transponders only on Add Fiber 1, and assume this router needs to establish one connection on Network Fiber 1 and one on Network Fiber 2. If there is only one available wavelength on these two network fibers and it is the same wavelength, then only one of the connections can be established, due to wavelength contention on the add fiber. Thus, the architecture of Fig. 7 is not *contentionless*. To remedy this, an external fiber cross-connect could be added such that the IP router can access transponders on any of the add/drop fibers. However, with good wavelength assignment algorithms, such contention does not occur often in practice.

Note that flexible all-optical switches have come to be called *reconfigurable OADMs* (ROADMs) (even though the term 'OADM' formerly implied a degree-two device).

VIII. Future Directions

All-optical networks continue to evolve in response to ever-increasing levels of traffic, with a growing emphasis on network flexibility, in addition to the ongoing concerns of cost and power savings.

A. Configurability and Dynamic Networking

Networks today are mostly quasi-static, with connection setup typically requiring on-site manual involvement, and connections often remaining established for months or years. As a first move away from this relatively fixed environment, networks are becoming configurable, where connections can be established remotely through software control, assuming the necessary equipment is already deployed in the network. Configurability takes advantage of flexible network elements such as ROADMs and tunable transponders.

The next step in this evolution is dynamic networking, where connections can be rapidly established and torn down, without the involvement of operations personnel. Dynamic networking delivers bandwidth where and when it is needed, providing significant cost benefits to the end-user. Research is underway in DARPA's Core Optical Network (CORONET) program to enable connection setup times on the order of 0.1 to 1 sec. [31]. This will take advantage of advances such as very fast switching and distributed network control.

B. Wavebands

ROADMs typically operate on the granularity of a wavelength; i.e., each wavelength can be independently directed to any port. An alternative approach that is gaining renewed interest is *waveband*-based switching [32][33][34]. A waveband is a set of wavelengths that are switched as a single unit. By switching on a waveband granularity, as opposed to switching individual wavelengths, many switch ports can be saved, leading to lower cost and power consumption.

Switching solely at the waveband level does result in somewhat less efficient networks, due to the coarser granularity of control. Thus, it makes sense to deploy waveband switching as part of a hierarchical multi-granular switch, which includes both a band-level and a wavelength-level switch [35][36]. If individual wavelengths need to be dropped from a waveband, or if wavelengths need to be packaged (i.e., groomed) into different wavebands, then the affected wavebands are demultiplexed, with the constituent individual wavelengths then processed by the wavelength-level switch. Note that wavelength conversion may be needed to move a wavelength from one waveband to another. Clearly, wavebands only make sense if the traffic level is high enough that the bulk of the wavebands are switched solely in the band-level switch.

Actually, wavebands were used in the first commercial all-optical system [15] (without requiring inefficient guardbands between the wavebands), but the concept did not catch on at that time. However, with traffic levels steadily increasing, this approach may be revisited.

C. Flexible Spectrum

To address the growing need for capacity, systems generally increase the number of wavelengths on a fiber and the line-rate of each wavelength. For example, early WDM systems were composed of 4 to 16 2.5-Gb/s wavelengths; more recently, transmission systems have been deployed with 80 40-Gb/s wavelengths, with 100 Gb/s wavelengths on the horizon. One of the benefits of greater line-rate is that historically the cost-per-bit/s has decreased. For example, a 10 Gb/s transponder is approximately twice the cost of a 2.5 Gb/s transponder. Conversely, one drawback is that the bitrate of the network services has not increased at the same pace, thus requiring more traffic grooming to keep the wavelengths efficiently packed. Since grooming is currently performed in the electronic domain, it has become a prime area of concern for carriers with respect to cost and power.

One proposal for addressing the mismatch of line-rate and service-rate is to deploy lower-rate, but much more closely spaced, wavelengths [37]. For example, instead of deploying 100 Gb/s wavelengths with 50 GHz spacing, deploy 10 Gb/s wavelengths with 5 GHz spacing. The overall system capacity remains the same, but the wavelength line-rate is better matched to the service-rate, resulting in much less required grooming. There will be major technical challenges in packing wavelengths so closely (e.g., cross-phase modulation in the fiber and adding/dropping individual channels), but there is evidence that this is feasible [38]. For cost effectiveness, PIC technology would be used to generate a comb of closely spaced wavelengths, which could be treated as a waveband.

Extending this concept, ideally the wavelengths can be any standard rate, from say 10 Gb/s to 100 Gb/s, or even up to 1 Tb/s, to best match the service that is being transported. One can then pack these variable-rate wavelengths into *fixed-sized* wavebands [37] (using efficient packing algorithms). This flexible-spectrum system can then be combined with a hierarchical waveband/wavelength switch architecture.

An alternative, non-banded, approach was proposed in [39], where each wavelength is 'elastic' and grows or shrinks in bandwidth as needed; i.e., a 'gridless' system. Switching is performed at the granularity of a wavelength. One complexity of such a system is that a wavelength may be optically bypassed from a link carrying one set of frequency allocations to a link with a completely different allocation (i.e., the spectrum may be 'sliced' up differently on each link). Thus, algorithms are needed to avoid spectral conflicts. It may be challenging to maintain high spectral utilization. Note that this complexity is avoided with the waveband-based scheme because the spectrum is partitioned uniformly into wavebands on each link.

D. Flexible Transponders and Switches

Flexible spectrum is one manifestation of greater network transparency. To exploit transparency further, one can use transponders that are capable of various line-rates and modulation formats. The initial impetus for flexible transponders was to trade off optical reach and capacity [40][41]. Thus, a single transponder could be capable of, for example, 100 Gb/s with 1,500 km reach or 10 Gb/s with 3,000 km reach, with the different operational configurations controlled by software. In another manifestation, a single transponder could be capable of producing a 100 Gb/s signal suitable for 50 GHz spacing and 1,500 km reach, or a 100 Gb/s signal suitable for 75 GHz spacing and 2,500 km reach. This flexibility potentially will allow carriers to better utilize capacity and regeneration resources.

Transponder flexibility is only half of the picture; the all-optical switches must be compatible with signals of arbitrary bandwidth as well. To a limited degree, this flexibility exists in today's ROADMs. The filter passbands of most ROADMs are compatible with signals from 10 Gb/s to 100 Gb/s, as long as the signals fit in a 50 GHz passband. This allows the system line-rate to be upgraded without replacing the ROADMs. However, more flexibility would be needed to accommodate the gridless system described above. Gridless ROADMs are feasible, using, for example, liquid crystal (LC) technology, where an array of LC pixels can be configured to create passbands of arbitrary shape and bandwidth [42].

E. Reducing Power Consumption in Future Networks

With network traffic demand continuing to grow at a fast pace, the corresponding increase in power consumption is becoming of real concern. While the current techniques of all-optical networking do reduce power consumption, there is still a need for additional power-reduction innovations.

The current techniques for realizing long-reach systems with 100 Gb/s per wavelength require intensive electronic processing at the transponders, which consumes high power. If the same techniques are used in future systems with higher bit rates per wavelengths (e.g., 400 Gb/s or 1 Tb/s), the situation will be greatly exacerbated. To control the power in such future systems, all-optical processing techniques that preserve the optical reach are needed.

Largely due to cost reasons, signal regeneration and wavelength conversion are currently accomplished by using two back-to-back transponders, which involves OEO conversion and consumes large power. As the bit rate per wavelength increases, performing these functions in the optical domain starts to be more attractive, but this will be challenging given the complex modulation formats required in the future.

As mentioned earlier, the continually increasing bitrate per wavelength requires more traffic grooming, which is performed today in power-hungry electronic circuit or packet switches. Thus, an emerging vision to reduce power is to employ *all-optical grooming*, in either the frequency or time domain, to reduce or eliminate electronic grooming. In frequency-domain optical grooming, the hierarchical waveband/wavelength approach described in Section VIII.C could be used, with the switches equipped with all-optical wavelength converters to repack wavelengths into wavebands as needed. Two time-domain optical grooming techniques being widely researched are optical burst switching (OBS) and optical packet switching (OPS). In OBS, bursts of data are electronically buffered at the edge of the network. A control packet is transmitted a short time ahead of the data burst to schedule the required resources at intermediate nodes. However, scheduling the bursts while maintaining high network utilization efficiency is difficult. In OPS, the electronic switch fabric and buffer of current switches are replaced by their optical analogs. However, the size of optical buffers appears to be quite limited. Much more research is needed before any of the above all-optical grooming techniques emerge.

Finally, as network traffic continues to grow, the capacity of a single fiber will be exceeded. Using a multi-fiber system is not scalable in either cost or power consumption. Rather, it has been suggested to use multi-core-, or multimode-fiber [43]. The big challenges are maintaining the extended optical reach needed for an all-optical solution, controlling the processing power needed to separate the signals in the various cores or modes, and realizing optical amplifiers and ROADMs that operate simultaneously on all of the signals in the fiber.

IX. Conclusions

The early vision of all-optical networking stressed exploiting the transparency of optics and the virtual elimination of electronics from the network. After a long journey full of challenges, pragmatic realizations, and major technological advances, all-optical networking has reached its current vibrant state, achieving great savings in cost and power consumption, while enabling ease of network operation and graceful upgradability. More advances are still needed to continue to cope with the explosive growth of future networks.

REFERENCES

[1] D. B. Payne and J. R. Stern, "Transparent single-mode fiber optical networks," *J. Lightw. Technol.*, vol. LT-4, no. 7, July 1986, pp. 864-869.

[2] M. C. Brain and P. Cochrane, "Wavelength-routed optical networks using coherent transmission," *ICC*, Philadelphia, Jun. 12-15, 1988, pp. 26–31.

[3] G. R. Hill, "A wavelength routing approach to optical communications networks," *Infocom*, New Orleans, Mar. 27-31, 1988, pp. 354-362.

[4] T. E. Stern, "Linear lightwave networks: how far can they go?," *Globecom*, San Diego, Dec. 2-5, 1990, vol. 3, pp. 1866-1872.

[5] S. Johansson, M. Lindblom, P. Granestrand, B. Lagerstrom, and L. Thylen, "Optical cross-connect system in broad-band networks: system concept and demonstrators description," *J. Lightw. Technol.*, vol. 11, no. 5/6, May/June 1993, pp. 698-694.

[6] I. Chlamtac, A. Ganz, and G. Karmi, "Purely optical networks for terabit communication," *Infocom,* Ottawa, Apr. 23-27, 1989, vol. 3, pp. 887-896.

[7] J. M. Simmons, *Optical Network Design and Planning*, Springer, New York, 2008.

[8] C. A. Brackett, A. S. Acampora, J. Sweitzer, G. Tangonan, M. T. Smith, W. Lennon, K.-C. Wang, and R. H. Hobbs, "A scalable multiwavelength multihop optical network: a proposal for research on all-optical networks," *J. Lightw. Technol.*, vol. 11, no. 5/6, May/June, 1993, pp. 736-753.

[9] S. B. Alexander, R. S. Bondurant, D. Byrne, V.W. S. Chan, S. G. Finn, R. Gallager, B. S. Glance, H. A. Haus, P. Humblet, R. Jain, I. P. Kaminow, M. Karol, R. S. Kennedy, A. Kirby, H. Q. Le, A. A. M. Saleh, B. A. Schofield, J. H. Shapiro, N. K. Shankaranarayanan, R. E. Thomas, R. C. Williamson, and R. W. Wilson, "A precompetitive consortium on wide-band all-optical networks," *J. Lightw. Technol.*, vol. 11, no. 5/6, May/June, 1993, pp. 714-735.

[10] I. P. Kaminow, C. R. Doerr, C.Dragone, T. Koch, U. Koren, A. A. M. Saleh, A. J. Kirby, C. M. Ozveren, B. Schofield, R. E. Thomas, R. A. Barry, D. M. Castagnozzi, V. W. S. Chan, B. R. Hemenway, D. Marquis, S. A. Parikh, M. L. Stevens, E. A. Swanson, S. G. Finn, R. G. Gallager, "A wideband all-optical WDM network," *J. on Selected Areas in Comm.,* vol. 14, no. 5, May 1996, pp. 780-799.

[11] S. Johansson, "Transport network involving a reconfigurable WDM network layer - a European demonstration," *J. Lightw. Technol.*, vol. 14, no. 6, June 1996, pp. 1341-1348.

[12] K. Sato, S. Okamoto, and H. Hadama, "Network performance and integrity enhancement with optical path layer technologies," *J. on Selected Areas in Comm.,* vol. 12, no. 1, Jan. 1994, pp. 159-170.

[13] A. A. M. Saleh, "Overview of the MONET, Multiwavelength Optical Networking, program," *Optical Fiber Communications (OFC)*, pp. 240, Feb. 25 – Mar. 1, 1996.

[14] R. E. Wagner, R. C. Alferness, A. A. M. Saleh, and M. S. Goodman, "MONET: multiwavelength optical networking," *J. Lightw. Technol.*, vol. 14, no. 6, June 1996, pp. 1349-1355.

[15] Business Wire, "Broadwing Communications to use Corvis technology for enhanced wavelength services," May 10, 2000. Available at: findarticles.com/p/articles/mi_m0EIN/is_2000_May_10/ai_61994410/

[16] K. Rottwitt and A. Stentz, "Raman amplification in lightwave communication systems," in *Optical Fiber Telecommunications IV A*, I. Kaminow and T. Li, Editors, San Diego: Academic Press, 2002, pp. 213-258.

[17] A. A. M. Saleh, "Transparent optical networking in backbone networks," *OFC 2000*, Baltimore, MD, Mar. 5-10, 2000, ThD7, pp. 62-64.

[18] J. M. Simmons, "On determining the optimal optical reach for a long-haul network," *J. Lightw. Technol.*, vol. 23, no. 3, Mar. 2005, pp. 1039-1048.

[19] J. M. Simmons, A. A. M. Saleh, and L. Benmohamed, "Extending Generalized Multi-Protocol Label Switching to configurable all-optical networks," *National Fiber Optic Engineers Conference (NFOEC'01)*, Baltimore, MD, July 8-12, 2001, pp. 14-23.

[20] J. M. Simmons, "Cost vs. capacity tradeoff with shared mesh protection in optical-bypass-enabled backbone networks," *OFC/NFOEC'07*, Anaheim, CA, Mar. 25-29, 2007, Paper NThC2.

[21] D. C. Kilper, R. Bach, D. J. Blumenthal, D. Einstein, T. Landolsi, L. Ostar, M. Preiss, and A. E. Willner, "Optical performance monitoring," *J. Lightw. Technol.*, vol. 22, no. 1, Jan. 2004, pp. 294-304.

[22] J. M. Simmons, "Analysis of wavelength conversion in all-optical express backbone networks," *OFC'02*, Anaheim, CA, Mar. 17-22, 2002, Paper TuG2.

[23] K. Bala, R. R. Cordell, and E. L. Goldstein, "The case for opaque multiwavelength optical networks," *LEOS Summer Topicals*, Aug. 7-11, 1995, Paper WA4, pp. 58-59.

[24] B. Ramamurthy, H. Feng, D. Datta, J. P. Heritage, and B. Mukherjee, "Transparent vs. opaque vs. translucent wavelength-routed optical networks," *OFC'99*, San Diego, Feb. 21-26, 1999, Paper TuF2, vol. 1, pp. 59-61.

[25] S. Subramaniam, M. Azizoglu, and A. K. Somani, "All-optical networks with sparse wavelength conversion," *IEEE/ACM Transactions on Networking*, vol. 4, no. 4, Aug. 1996, pp. 544-557.

[26] D. F. Welch, F. A. Kish, R. Nagarajan, C. H. Joyner, R. P. Schneider, Jr., V. G. Dominic, M. L. Mitchell, S. G. Grubb, T.-K. Chiang, D. D. Perkins, and A. C. Nilsson, "The realization of large-scale photonic integrated circuits and the associated impact on fiber-optic communication systems," *J. Lightw. Technol.*, vol. 24, no. 12, Dec. 2006, pp. 4674-4683.

[27] M. C. Wu, O. Solgaard, and J. E. Ford, "Optical MEMS for lightwave communication," *J. Lightw. Technol.*, vol. 24, no.12, Dec. 2006, pp. 4433-4454.

[28] X. Zhou, M. Feuer, and M. Birk, "Fast control of inter-channel SRS and residual EDFA transients using a multiple-wavelength forward-pumped discrete Raman amplifier," *OFC/NFOEC 2007*, Anaheim, CA, March 25-29, 2007, Paper OMN4.

[29] S. Mechels, L. Muller, G. D. Morley, and D. Tillett, "1D MEMS-based wavelength switching subsystem," *IEEE Commun. Mag.*, vol. 41, no. 3, pp. 88-94, Mar. 2003.

[30] S. Okamoto, A. Watanabe, and K. Sato, "A new optical path cross-connect system architecture utilizing delivery and coupling matrix switch," *IEICE Trans. on Comm.*, vol. E77-B, no. 10, pp. 1272-1274.

[31] A. L. Chiu, 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, Jan. 2012, pp. 1-14.

[32] X. Cao, V. Anand, and C. Qiao, "Waveband switching in optical networks," *IEEE Trans. Commun*., vol. 41, Apr. 2003, pp. 105–112.

[33] R. Parthiban, R. S. Tucker, and C. Leckie, "Waveband grooming and IP aggregation in optical networks," *J. Lightw. Technol*., vol. 21, no. 11, Nov. 2003, pp. 2476–2488.

[34] K. Sato and H. Hasegawa, "Optical networking technologies that will create future bandwidth-abundant networks," *J. Opt. Commun. Netw,* vol. 1, no. 2, Jul. 2009, pp. A81-A93.

[35] K. Harada, K. Shimizu, T. Kudou, and T. Ozeki, "Hierarchical optical path cross-connect systems for large scale WDM networks," *OFC'99,* San Diego, CA, Feb. 21-26, 1999, Paper WM55.

[36] A. A. M. Saleh and J. M. Simmons, "Architectural principles of optical regional and metropolitan access networks," *J. Lightw. Technol.*, vol. 17, no. 12, Dec. 1999, pp. 2431-2448.

[37] J. Berthold, A. A. M. Saleh, L. Blair, and J. M. Simmons, "Optical networking: past, present, and future," *J. Lightw. Techol.*, vol. 26, no. 9, May 1, 2008, pp. 1104-1118.

[38] N. Alic, Y. Myslivets, and S. Radic, "Non-coherent high spectral efficiency long haul waveband transmission," *IEEE Photon. Tech. Lett.*, vol. 24, no. 2, Jan. 15, 2012, pp. 113-115.

[39] M. Jinno, H. Takara, B. Kozicki, Y. Tsukishima, Y. Sone, and S. Matsuoka, "Spectrum-efficient and scalable elastic optical path network: architecture, benefits, and enabling technologies," *IEEE Commun. Mag.*, vol. 47, no. 11, Nov. 2009, pp. 66-73.

[40] A. Bocoi, M. Schuster, F. Rambach, M. Kiese, C.-A. Bunge, and B. Spinnler, "Reach-dependent capacity in optical networks enabled by OFDM," *OFC/NFOEC*, San Diego, Mar. 22-26, 2009, Paper OMQ4.

[41] A. Klekamp, O. Rival, A. Morea, R. Dischler, and F. Buchali, "Transparent WDM network with bitrate tunable optical OFDM transponders," *OFC/NFOEC'2010*, San Diego, Mar. 21-25, 2010, Paper NTuB5.

[42] T. A. Strasser and J. L. Wagener, " Wavelength-selective switches for ROADM applications," *J. Selected Topics in Quantum Electronics*, vol. 16, no. 5, Sep./Oct. 2010, pp. 1150-1157.

[43] A. R. Chraplyvy, "The coming capacity crunch," *ECOC* Plenary Talk, Vienna, Austria, Sept. 20-24, 2009, Paper 1.0.2.