

Addressing Transparency in DWDM mesh survivable networks

Sid Chaudhuri, Eric Bouillet, Georgios Ellinas
Tellium, Inc. 2 crescent place, Oceanport, NJ 07757

Abstract: We identify the penalties incurred in all-optical networks on routing and restoration operations. We show that these penalties can be avoided if we replace the transparent switches with opaque switches interfaced by transponders.

Dense Wavelength Division Multiplexed (DWDM) mesh networks that route optical connections using wavelength selective optical cross-connects (OXC) have been proposed to implement next generation networks. In a DWDM network, optical paths on specific wavelengths are set up through the action of OXC. All Optical Switches (AOS) are OXC that perform elementary functions such as wavelength routing and switching with all operations executed in the optical domain, thereby providing end-to-end transparent “clear” channels (optical connections). Because of the absence of wavelength conversion the constraint of wavelength continuity applies to all optical connections: a signal must remain on the same wavelength from source to destination. The restriction imposed by wavelength continuity can be avoided in an opaque network built as point-to-point DWDM systems with transponders and switches connecting the DWDM channels. The switch fabric in this scenario can be either optical or electrical. In a switch with Optical/Electronic interface (OE switch) the incoming optical signal is converted to an electronic signal, routed and converted back to the optical domain, possibly on a different wavelength.

Following a wave of timely technological breakthroughs, optical network vendors are now announcing a variety of all optical gears capable of redirecting light without electronic intervention. However, there is much ambiguity about the applications and benefits of this technology. All optical OXC are simpler, they may circumvent electronic bottlenecks and be relatively inexpensive, but they lack important functionality such as wavelength conversion, performance monitoring, fault isolation and keep-alive signal generation. OE switches on the other hand, have all the desired functionality.

Previous analyses have principally underlined the impact of lack of wavelength conversion on blocking probabilities - or conversely, on capacity requirement - with most of the emphasis being put on the design of efficient routing and wavelength selection algorithms to minimize this impact. Results reveal that full benefit can be obtained by adding wavelength conversion capability to 10%-20% of the nodes [1,3], and heuristics for allocation of wavelength convertible OXC have also been proposed [2,4]. The general consensus is that the gain attained with wavelength conversion is usually modest, giving an apparent cost-effectiveness advantage to all-optical solutions. Yet, little effort has been devoted on assessing other impediments associated with all-optical architectures (e.g., network operation and survivability issues). This paper examines these issues from both the EON and AON viewpoints and shows that due to its limited leverage for protection mechanisms, the AON architecture requires up to twice as much resources to provide equivalent tolerance against equipment failures.

Figure 1(a) illustrates an opaque switch and Figure 1(b) illustrates an all-optical switch. Architectures based on the opaque switch are the most versatile, supporting the complete range of protection schemes, ranging from span protection to mesh shared restoration. This is the only architecture in place today that allows standards based multi-vendor inter-operability. Another variation of the architecture is with the WDM long-haul optics being integrated within the optical switch. This eliminates two short reach interfaces on one lightpath through a node. Architecture based on all-optical switches may be less expensive and capable of switching higher speed signals, because they are oblivious to data rate. On the other hand, a restriction generally referred to as the *wavelength continuity* constraint is imposed. All optical switches also suffer from limited signaling and monitoring abilities. These hindrances reduces the scope of recovery scenarios for this architecture essentially to 1+1 dedicated protections, thus forfeiting more cost-effective protocols such as mesh shared restoration.

Consider the example of Figure 2. The figure depicts a 4-node AON with the routing of 3 bi-directional demands {AB, AD, DC}. In the example each demand pair shares a fiber, and so must be at a different wavelength. Even though each fiber is carrying two demands, the solution requires 3 wavelengths. Had the number of wavelengths been limited to two per fiber, then either more fibers would have been required or one of the demands would have been blocked. If on the other hand the OXC could do wavelength conversion, two wavelengths would have sufficed. Now note that survivability issues are not addressed in this example. We examine this aspect next.

There exist three categories of protections. (1) In local span protection the OXC closest to the failure attempt to restore the lightpaths on an alternate circuit around the failure. This protection scheme is faster and yields higher availability, since only the disabled portion of the path needs to be bypassed. (2) In end-to-end dedicated protection the ingress and the egress OXC of the failed optical connection attempt to restore the signal on a predefined backup path that is edge-disjoint, or diverse, from the primary path. Path diversity guarantees that primary and backup lightpaths will not simultaneously succumb to single failures. This approach requires extraordinary amount of spare capacity. However, the backup paths remain “live” in permanence, thus saving crucial path-setup latencies when recovery takes place. (3) In shared mesh restoration, backup paths can share capacity if the corresponding primary

paths are edge-disjoint. Compared to 1+1 dedicated protection, this scheme allows considerable saving in terms of resource utilization. In addition, the backup resources can be utilized for lower priority traffic in normal network conditions. However, recovery is slower than dedicated protections, essentially because it involves signaling and path-setup procedures to reallocate the spare capacity.

In this section we describe briefly the procedure used in our experiments. First we make the following assumptions: both architectures have identical topologies and are compared under identical, finite, demand patterns. EON implements shared mesh restoration, while AON implements 1+1 dedicated protection. The simulator proceeds as follows:

1. The architectures are given, with initial capacity C_e and number of fiber-pairs F_e per link, where C_e is expressed in units of idle channel pairs. For AON architectures, F_e represents the number of times a wavelength can be reused in a single link for different lightpaths (once per fiber-pair).
2. Demands arrive in finite sequence, in arbitrary order. Alternatively, existing demands can disconnect and release allotted resources. The type of each request is a random decision that depends on the *churn*, defined as the probability that the request corresponds to a demand termination.
3. We compute a pair of edge-disjoint lightpaths, one primary and one protection, to accommodate each demand set-up request. Paths are computed in conformity with wavelength continuity constraint if the inspected architecture is AON. In the case of EON architecture, the protection path may share channels with pre-established protection paths. If the computation succeeds, the new demand is routed to the first channels available on the designed paths otherwise the demand is lost.
4. The algorithm stops once all the demands have been processed and either activated or blocked.

We exercised the simulator described in the previous section on a 50 node, 88 edge infrastructure, with 2019 OC48s. We measure two figures of merit, *blocking percentage*, defined as the blocked demands expressed in percents of the offered demands and *network lifespan* defined as the time before the first blocking event occurs. In practice a carrier will add capacity to accommodate a new demand request instead of blocking it. From a carrier perspective it is therefore more advantageous to maximize the network lifespan, time at which new capacity must be added, rather than minimizing the blocking percentage.

In the first set of experiments we assume no protection against failure and attempt to capture the effect of wavelength continuity constraints. We first consider the case without churn. Assuming wavelength conversion everywhere, the network capacity is dimensioned off-line such that the blocking does not exceed 5% of the offered demand, independently of the order of arrival. With this capacity configuration, both architectures are then simulated for 20 demand scenarios. Results are shown in Figure 3. The plots correspond to the High-Low values of blocking percentage. We observe that the blocking percentage is 3.7 to 4% for AON and 2 to 3.3% for EON. These results are consistent with previous findings [3]. More important, the network lifespan is in the range 67-81% for AON, compared to 88-94% for EON; that is at least 7%, and up to 27% more of the total demand before the first blocking occurs.

We repeat the experiment with churn 0.1. Intuitively, the larger the churn, the more likelihood there is of fragmentation in wavelength assignments, making it less probable that the same wavelength will be available throughout the network for longer paths. The simulation confirms this. As illustrated in Figure 4, AON's lifespan varies from 49% to 68% of the total demand, well under the 78-95% span observed in the EON case.

In the next set of experiment we investigate the benefits of shared mesh protections versus dedicated protection which is the only possible class of protections for AON architectures. The procedure used is the same as before. Results are summarized in Figure 5. With shared mesh restoration, EON's lifespan is in the range 95% to 100%. With dedicated protection, AON accommodates between 45% and 50% of the demands before it experiences blocking.

It is true that AOS technology bypasses electronic switching and eliminates the needs for transponders. By way of this design, they consume less power, do not require as much heat dissipation, and occupy less floor space than opaque switches. They also profit from scalable port-counts. However, in addition to wavelength continuity constraints, which as demonstrated earlier restrict their ability to provision capacity efficiently and perform all range of protection mechanisms, they suffer from other drawbacks. (1) Signal degradation detection and isolation are not easily implemented. (2) Lack of visibility of overheads means that a separate outband signal network must be built and managed. (3) Neighbor discovery protocols to create port-state databases necessary for network auto-discovery and topology dissemination are complex without inband overhead access. (4) Optical impairment accumulation (cross-talk, noise, etc.) tend to restrict their geographical and nodal size.

In this paper we demonstrate that transparent OXCs may be short of important functionality, and usability necessary to build robust and cost-effective networks. We arrive at our conclusion mainly by quantifying the amount of resources required to accommodate predefined OC48 traffic forecasts in the proposed architectures. We show that carriers opting for all-optical solutions will have to overbuild their infrastructures and surmount tremendous management obstacles in order to be competitive.

References:

- [1] K. Bala, E. Bouillet, G. Ellinas, IEEE/OSA OFC97, pp 120-121, February 1997.
- [2] S. Subramaniam, M. Azizoglu and A. Somani, OFC98, February 1998.
- [3] S. Subramaniam, M. Azizoglu, and A.K. Somani, IEEE INFOCOM 96, San Fransisco, CA, 1996.
- [4] H. Harai, M. Murata, H. Miyahara, Journal of Lightwave Technologies, vol. 17-4, April 1999.

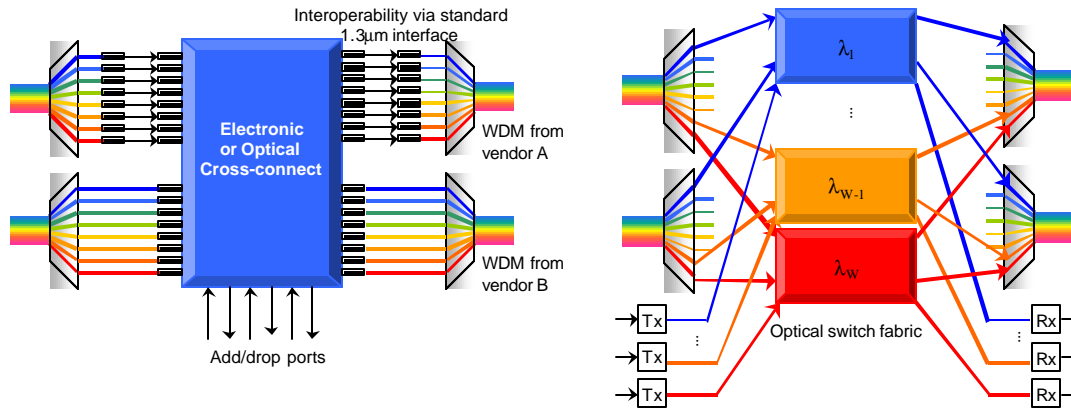


Figure 1a. and 1b.




Figure 2.

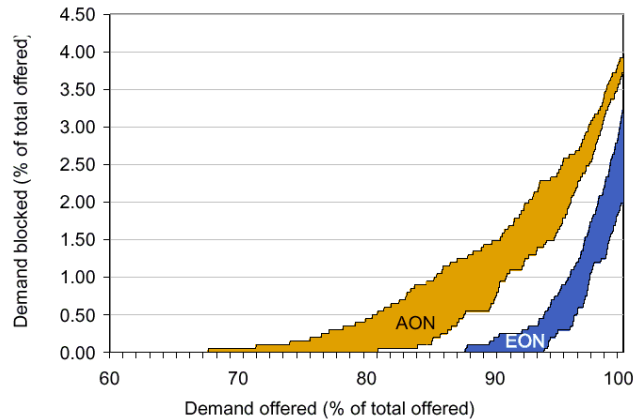


Figure 3.

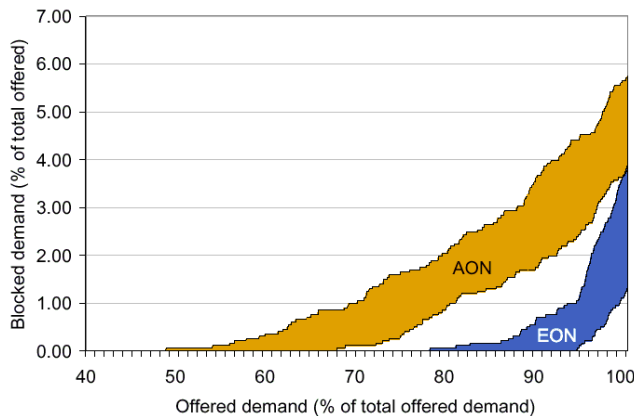


Figure 4.

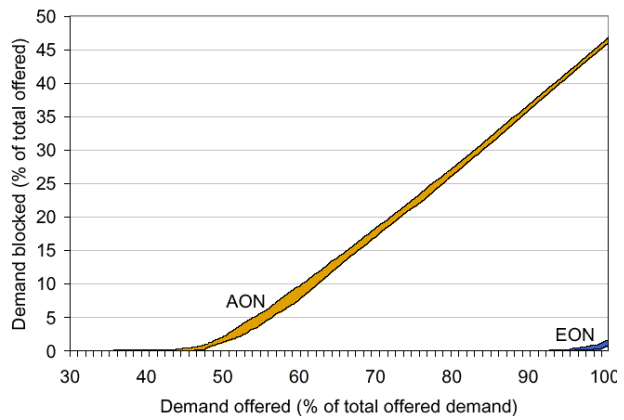


Figure 5.