IEEE Communications Magazine, vol. 44, no. 11, November 2006, pp. 88-94.

Network Design in Realistic 'All-Optical' Backbone Networks

Jane M. Simmons

ABSTRACT

Optical-bypass technology is finally being deployed in carrier backbone networks on a large scale. The reality, however, is that the resulting networks are not truly all-optical; all connections cannot be carried end-to-end solely in the optical domain. While a significant amount of regeneration can be eliminated through the deployment of new ultra-long-haul technology, a small amount is still required. We study the optical network design problem for realistic backbone networks, with a focus on the impact of regeneration.

INDEX TERMS

All-optical networks, network design, optical reach, regeneration, routing, ultra-long-haul technology, wavelength assignment

I. INTRODUCTION

In all-optical networks, traffic is carried end-to-end in the optical domain, without any intermediate optical-electrical-optical (O-E-O) conversion. The promise of such networks is the elimination of a significant amount of electronic equipment, as well as added capabilities, such as the ability to transport any type of data format through the network. All-optical systems give rise to the well-known wavelength continuity constraint. If a signal is transmitted in the optical domain over a number of hops, it must be carried on the same wavelength on each of the hops (assuming all-optical wavelength conversion is not present). Thus, routing connections and assigning wavelengths to them are important aspects of the design process.

Network design for all-optical networks has been addressed in numerous previous studies, as summarized in [1]. The assumption in much of this work is that the network is truly all-optical, where all intermediate O-E-O conversion is eliminated. However, now that 'all-optical' technology is finally being deployed in carrier backbone networks, the reality is that these systems are only 'mostly all-optical', as defined in [2]. Some amount of intermediate O-E-O conversion is deployed in these networks for purposes of regenerating the signal. The fact that regeneration is not completely removed from the network has important ramifications for optical network design. The impact of regeneration, as well as other characteristics of practical backbone networks, is the focus of this paper. For ease of exposition, we will continue to use the term 'all-optical', however, these networks might better be described as 'optical bypass enabled'.

II. BACKGROUND

All-optical networks require two necessary ingredients: network elements that are capable of optical bypass and optics capable of extended reach. The network elements must have the ability to selectively drop traffic at a node, while allowing the remaining traffic to transit the node in the optical domain. Examples of such network elements are *optical add drop multiplexers* (OADMs) and *all-optical switches*. *Optical reach* is the distance an optical signal can travel before the signal quality degrades to a level that necessitates regeneration. Many factors affect the optical reach; e.g., the type of amplification, the launched power of the signal, and the modulation format of the signal. Compared to the 500 km optical reach of traditional systems, the optical reach of long-haul systems currently being deployed is on the order of 2,000 to 4,000 km. This increased reach, in combination with optical-bypass network elements, eliminates a significant amount of the regeneration in a backbone network. However, given that the longest connections in a North American backbone network are on the order of 8,000 km, clearly some regeneration is still required.

The fact that there is regeneration has two major implications. First, the cost of provisioning a circuit may consist of more than just the transponders at the two connection endpoints; intermediate regenerators may be required, which will impact the network cost. (A transponder is a combination transmitter/receiver card that has a short reach interface on the client side and a WDM-compatible signal on the network side.) We are referring to the capital cost of the network here; however, note that regenerators also add to the operational costs, as they require installation and maintenance. Thus, from a cost perspective, it is in most cases desirable to minimize regeneration. Secondly, regeneration typically affords the opportunity for wavelength conversion [3], [4]. For example, if regeneration is accomplished via back-to-back transponders, then the two transponders can be tuned to different wavelengths to accomplish wavelength conversion. In addition, all-optical regenerators, which may be commercially viable in the future, typically are capable of wavelength conversion as well [5]. The fact that wavelength conversion is essentially achieved ‘for free’ when regenerating can be used to simplify network design. Various regeneration architectures are discussed in [6].

In this paper, we consider the network design problem for today’s realistic backbone networks. We assume that we are provided with a topology and a set of demands, with the goal of carrying all demands while minimizing the capital costs of the network. This is the problem that is typically posed by carriers as they evaluate the equipment of various system vendors. We present several design aspects that should be considered in order to generate optimal, or close-to-optimal, results. While the focus is on network design with a fixed demand set, many of the points considered are applicable for the real-time operation of a network, where demands are continually added and removed.

III. PROBLEM SIZE

When developing network design algorithms it is important to consider the size of the problem. All-optical backbone networks generally have less than 100 nodes. (Nodes are the add/drop and/or switching locations in a network.) In North America, the typical average nodal degree (i.e., number of links incident on a node) is between 2.5 and 3. The size of the demand set provided by carriers depends on whether the traffic requires grooming or not (and of course depends on the size of the carrier). If the traffic is at the line rate (no grooming needed), there are typically a few hundred to a couple of thousand demands to be routed. If the traffic is subrate, such that grooming is needed, there could be tens of thousands of demands to be groomed and routed.

The number of wavelengths per fiber is determined by the equipment of the particular system vendor. This is typically on the order of 80 wavelengths. For very large designs, where both the C and L spectral bands are needed, the number of wavelengths can be on the order of 160 wavelengths.

IEEE Communications Magazine, vol. 44, no. 11, November 2006, pp. 88-94. © 2006 IEEE

When carriers issue network design studies to evaluate technology, they typically project demand sets over several years, where a design must be done for each year. Furthermore, they often request that the designs be done under several different protection assumptions (e.g., unprotected, dedicated client protection, shared mesh restoration). Thus, although these designs are performed 'off-line', the run-time to perform the network design is still very important. It is desirable that the network design process be completed on the order of a couple of minutes or less, thereby requiring scalable, efficient algorithms. Based on the problem size as given above, techniques that tend not to scale well, such as integer linear programming, generally will not meet these requirements.

IV. SINGLE-STEP VS TWO-STEP ROUTING AND WAVELENGTH ASSIGNMENT

Routing and wavelength assignment are important aspects of the design process for all-optical networks. Before discussing routing and wavelength assignment algorithms in more detail, we first consider whether these design steps should be handled separately or at the same time. One approach is to first route all of the connections, and then assign wavelengths to them as a separate step. With this strategy, it is possible that no wavelength is available for the route that is found. Another approach is to combine the steps so that routing is tied to a particular wavelength. This latter approach is typically accomplished by starting with a particular wavelength and reducing the network topology to only those links on which this wavelength is available. The routing algorithm is then run on this pruned topology. If a suitable route cannot be found, another wavelength is chosen and the process run through again on the corresponding topology. With this combined approach, it is guaranteed that there will be a free wavelength on any route that is found.

The presence of regeneration impacts the efficacy of the combined approach however. Consider the backbone network shown in Fig. 1, with a connection from New York to Los Angeles. Such a connection will typically require at least two regenerations with current ultra-long-haul technology. Assume that for the path shown, these regenerations occur at Chicago and at Denver. When assigning wavelengths to this path, it is necessary to find a free wavelength from New York to Chicago, from Chicago to Denver, and from Denver to Los Angeles, as opposed to finding a single free wavelength from New York to Los Angeles. The wavelengths on the three segments comprising the path do not need to be the same due to the wavelength conversion that can occur concomitantly with regeneration.

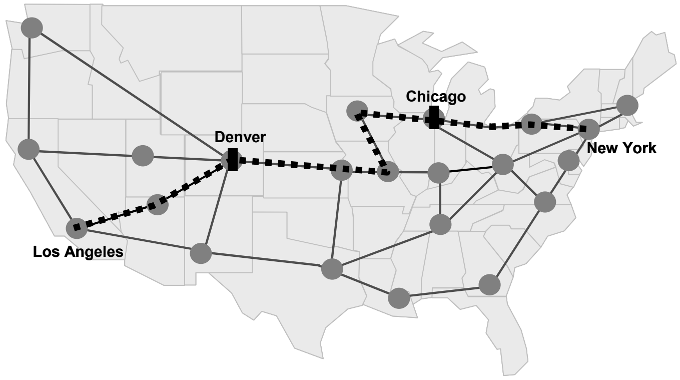

Fig. 1. Assume the path shown between New York and Los Angeles requires regeneration in Chicago and Denver. The wavelengths assigned to the three resulting path segments do not need to be the same, thus simplifying the wavelength assignment problem for this connection.

Using the combined routing and wavelength assignment procedure actually makes the problem unnecessarily more difficult because it implicitly searches for a wavelength that is free along the *whole* length of the path. One strategy would be to select ahead of time where the regenerations are likely to occur along a connection, and apply the combined routing and wavelength assignment approach to each segment individually. However, the route that is ultimately found could be somewhat circuitous and require regeneration at different sites, so that the process may need to be run through again.

Breaking the problem into separate steps for routing, selection of regeneration sites, and wavelength assignment avoids this issue, in addition to running more quickly. As will be discussed later, the amount of wavelength contention that generally results from operating this strategy with real backbone networks is very small. The next three sections of the paper treat routing, regeneration, and wavelength assignment as three related but separate steps in the network design process.

V. ROUTING

The traffic sets provided by carriers generally require no more than a single fiber pair per link, based on the large capacity of current technology (over 1 Tbps per fiber). A single set of bi-directional amplifiers, with a fixed maximum capacity, is deployed on each link, and can be viewed as a fixed cost. (While lighting multiple fiber pairs per link may be possible, it implies more amplifier costs, larger switches, and is operationally not as appealing to carriers.) To a first approximation, the nodal switching equipment can also be considered a fixed cost. Network nodes are equipped with switching equipment based on their nodal degree; the switching fabrics of ultra-long-haul systems are typically deployed with a fixed maximum capacity.

Each connection in a demand set requires a transponder at its two endpoints. Thus, for a given demand set, the chief variable cost is the number of intermediate regenerators required along the connection paths. This assumes the design does not require processing such as grooming or shared protection, the results of which can also affect the network cost. (While grooming and shared protection are important in some network designs, they are not the focus of this paper. However, these topics will be touched on briefly in relation to regeneration in Section X.) Overall, with these assumptions, the design goal is to minimize regeneration, subject to carrying all demands on a single fiber pair.

It is assumed that regeneration occurs only at network nodes (placing regenerators in amplifier huts is more difficult to manage; it also may be infeasible due to physical space constraints). We also assume that regeneration is determined on a per-connection basis. In other strategies, certain nodes are designated as regeneration sites, and all paths routed through these nodes are regenerated, whether needed or not; this strategy generally leads to a more costly network. (Designating sites for regeneration does make sense in some scenarios, however. For example, if there are multiple vendors providing network equipment, with the various systems being incompatible, designated O-E-O sites can serve as buffers between the systems.)

We also assume that any node is capable of regeneration. As previously described, regeneration can be accomplished via back-to-back transponders, where these are the same transponders used for adding/dropping traffic at a node. Thus, carriers generally do not restrict where regeneration can occur.

Note that in the real-time operation of a network, transponders may be available only at certain sites, and thus regeneration is limited to these sites. This real-time problem is studied in [7], [8].

A. Minimum Regeneration Path

Finding the shortest path for a connection is not guaranteed to minimize the amount of regeneration. Consider the example shown in Fig. 2, with a connection from A to F, and assume the optical reach is 2,000 km. Path 1 is the shortest path for this connection, with a length of 3,500 km. Based on the path length and the optical reach, it is expected that one regeneration would be required. However, because regenerations occur only at nodal sites, two regenerations are actually required (e.g., the regenerators could be at Nodes C and E). Path 2, while longer, with a length of 3,800 km, requires just one regeneration (at Node H).

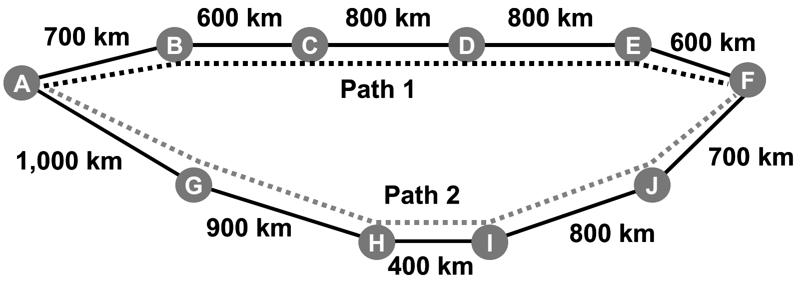

Fig. 2. For a connection from Node A to Node F, Path 1 is shorter (3,500 km) than Path 2 (3,800 km). However, for a system with an optical reach of 2,000 km, Path 2 requires only one regeneration whereas Path 1 requires two regenerations (assuming regeneration occurs only at nodes).

Similarly, finding the shortest dual path for a protected connection does not necessarily minimize regeneration. In Fig. 3, assume a 1+1 protected connection is required between Nodes A and D, and assume the optical reach is 2,000 km. The shortest dual path between A and D that is link- and node-disjoint is: A-B-C-D and A-E-F-D. These two paths have a combined distance of 6,300 km, and require a total of three regenerations (at Nodes C, E and F). However, the minimum-regeneration dual path is: A-E-C-D and A-B-D. This protected path is a total of 6,600 km, but requires a total of only two regenerations (at Nodes E and B).

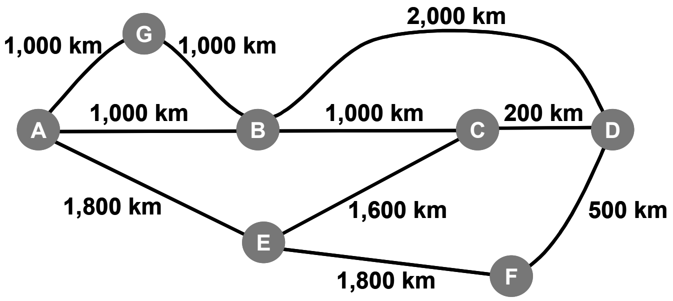

Fig. 3. Assume the optical reach is 2,000 km. For a protected connection from Node A to Node D, the combination of diverse paths A-B-C-D and A-E-F-D is the shortest (6,300 km), but requires three regenerations. The combination of diverse paths A-E-C-D and A-B-D is longer (6,600 km), but requires two regenerations.

As demonstrated by these examples, simply finding the shortest path or the shortest dual path will not necessarily find the lowest cost path. Furthermore, as carriers typically project a level of demand that pushes the boundary of existing technology, it is important to be able to efficiently pack connections in a network. Thus, some amount of load balancing is generally required when routing. Consequently, we would like to find a *set of paths* that meets the minimum regeneration for a particular source/destination pair.

One strategy is to first generate a large number of paths, not necessarily completely disjoint, for each required source/destination combination. This can be done using a k-shortest path algorithm (e.g., [9]), where k is say 10. In addition, one can run a shortest path algorithm, where a link (or link sequence) that is expected to have high load is removed from the topology. The goal is to find many paths that use a diversity of links. Once the path list is produced, the minimal number of regenerations among these paths is determined. While this is not guaranteed to be the minimal possible number of regenerations, it is highly likely that one of the paths does meet this goal. The path set can then be narrowed down to only those paths that meet this minimal regeneration. A small number of these paths are then selected from this set to serve as the actual candidate path set for the given source/destination pair.

Using about three candidate paths per source/destination pair generally yields good load balancing. There is no requirement that these candidate paths be totally diverse from each other. However, it is desirable that they be diverse with respect to the links that are most likely to be congested in the network (doing a quick preliminary routing of all connections on their shortest path gives a good idea of where congestion will likely occur). Note, however, that picking *all* paths to avoid the congested links may not be the best strategy because it may simply shift the congestion. When selecting the candidate paths, all other factors being equal, the paths with fewest hops should be chosen, to simplify the wavelength assignment process. Note that in scenarios where the network will end up being very heavily loaded, it may be necessary to include a path in the candidate path set that has an extra regeneration in order to avoid a highly loaded link.

As described above, it is envisioned that this path generation process occur prior to the routing of the demands. One of the candidate paths is then selected for each connection based on the state of the network at the time that connection is added to the network design (various criteria for selecting a path are possible; e.g., see [1]). If at the time a connection is added to the network there is no capacity available on any of the pre-selected paths, then a new path can be searched for based on the network links that do have available capacity. Clearly, the order in which the connections are routed affects the network design. For example, the routing order can be based on the distance between source and destination. Or, another effective scheme is 'round-robin' ordering, where the second occurrence of a particular source/destination pair is not routed until the first occurrence of each of the other source/destination pairs in the traffic set have been routed.

Rather than pre-selecting candidate paths prior to routing, it is also possible to dynamically search for a path as each connection is added to the network. This allows a tighter coupling between the current state of the network and the route selection process. One drawback with this approach is that a different path could end up being used for each instance of a given source/destination pair. This has the effect of decreasing the network 'interference length', which can potentially lead to more contention in the wavelength assignment process [10]. (As defined in [10], interference length is the average number of hops shared by two paths that have at least one hop in common.) In addition, the dynamic routing strategy is more time consuming.

VI. SELECTION OF REGENERATION SITES

Once the route has been chosen for a connection, the next step is selecting the regeneration sites, if needed. For a given connection that is longer than the optical reach, there may be several regeneration options. Consider Path 1 illustrated in Fig. 2. As discussed above, for an optical reach of 2,000 km, two regenerations are required. The regenerations can occur at Nodes C and D, or at Nodes C and E. Selecting the regeneration sites determines the segments for which wavelength assignment needs to be performed. The regeneration points should be selected such that the resulting segments best align with the segments already created by the network design, as this helps to minimize wavelength assignment conflicts. If during the wavelength assignment process it is not possible to find a wavelength that is free along an entire segment, then different regeneration sites can be considered for the corresponding connection so that different segments are created.

In the discussion above, it was assumed that regeneration is determined simply based on the distance the optical signal has traveled. In reality, the determination of when an optical signal needs to be regenerated is based on a variety of factors, such as optical signal-to-noise-ratio (OSNR), amount of accumulated chromatic dispersion, as well as several non-linear impairments. The exact regeneration rules will heavily depend on the underlying technology of the particular system, as well as the characteristics of the carrier's network, e.g., spacing between amplifiers, the loss between amplifiers, the fiber type. Various link metrics can be used to capture the performance penalty of a link, where this metric can be used instead of distance in the routing process to better ensure that minimum regeneration paths are found. (Also note that the network elements themselves contribute to the performance penalty; the effective penalty of a link needs to be increased based on the elements at its endpoints.) For example, in a system where the performance is dominated by the OSNR, and assuming there is no net gain or loss on each link, then link noise figure, in linear units, can be used as an appropriate routing metric. (The noise figure of a link is defined as the ratio of the OSNR at the start of the link to the OSNR at the end of a link.)

Nevertheless, in the initial stages of network design, the details of the fiber plant are generally not provided, and distance-based regeneration is often used as a rough approximation.

VII. WAVELENGTH ASSIGNMENT

After regeneration sites are chosen for all connections, wavelength assignment needs to be performed for each of the resulting segments. As discussed previously, wavelengths can be assigned independently to the segments comprising a particular connection. Wavelength assignment is not particularly onerous in practical backbone networks because the number of hops comprising a segment is not large. Table 1 provides the statistics of four network topologies and demand sets that are representative of North American backbone networks (all demands are at the line rate in these scenarios). The last three columns indicate the average number of hops per segment, assuming an optical reach of

2,000, 3,000, and 4,000 km, respectively. This number increases as the optical reach increases, but remains relatively small even with 4,000 km reach.

Table 1. Average Hops per Segment as a Function of Optical Reach

	Number of Nodes	**Number of Links**	**Average Link Length (km)**	**Number of Demands**	**Average Path Distance (km)**	**Average Hops Per Segment (2,000 km reach)**	**Average Hops Per Segment (3,000 km reach)**	**Average Hops Per Segment (4,000 km reach)**
Network 1	16	20	950	400	2,150	1.4	1.9	2.4
Network 2	25	38	700	400	1,875	1.8	2.5	2.7
Network 3	35	46	675	400	2,071	2.4	3.0	3.3
Network 4	55	70	470	1,200	2,107	3.1	4.0	4.7

A wavelength assignment strategy similar to Relative Capacity Loss [11] was used for each of the network designs listed in Table 1. All but one of the designs achieved 100% wavelength utilization, while achieving minimal regeneration; i.e., if there were L connections routed on the most heavily load link, only L wavelengths were needed in the wavelength assignment phase. For Network 4, with 4,000 km optical reach, approximately 1% more wavelengths were needed as compared to L. In this scenario, if the capacity supported by the system were limited to L, then a small number of additional regenerations could be added to alleviate the wavelength contention issues, as described in [3], [4].

These results indicate that separating the routing and wavelength assignment process can still produce very efficient network designs. Furthermore, the algorithms run quickly. On a 1.6GHz PC, the designs for Networks 1, 2, and 3 ran in less than 5 seconds; the designs for Network 4 ran in approximately 15 seconds.

A. Wavelengths with Different Optical Reach

Depending on the underlying technology of the ultra-long-haul system and the type of fiber installed in a carrier's network, all wavelengths may not have the same optical reach. For example, some fiber types have regions of very low dispersion leading to more non-linear impairments, which may result in the wavelengths in this region having reduced optical reach relative to the rest of the system spectrum. These wavelengths typically represent a small percentage of the overall wavelengths. In the wavelength assignment process, it is best to preferentially use these wavelengths when encountering segments that are short; i.e., assign a wavelength with the shortest reach that is still greater than the distance of the segment. Otherwise, these wavelengths may be left to the end of the assignment process, when extra regenerations may be needed to chop the remaining segments into even shorter segments.

B. Wavelength Assignment for Bi-directional Connections

Carriers typically provide demand sets where all connections are bi-directional; i.e., a connection from Node A to Node B implies a connection from Node B to Node A. We will assume that both directions of the connection are routed over the same path and regenerated at the same sites, thereby yielding identical segments in the two directions. An interesting question is whether the same wavelength should be assigned to both directions of the segment.

From a network management point of view, it may be most expedient to simply assign the same wavelength to both directions of the segment. However, there are scenarios where assigning different wavelengths can improve the efficiency of the network. Consider the very simple four-node topology of Fig. 4, and assume that the fiber capacity is just two wavelengths, and assume that the optical reach is longer than any of the possible connection paths. Assume that one bi-directional connection is established between A and C, and another between A and D. In Fig. 4a, Wavelength 1 is assigned to both directions of the A-C connection, and Wavelength 2 is assigned to both directions of the A-D connection. With this wavelength assignment in Fig. 4a, it is not possible to add a bi-directional connection between C and D without wavelength converting at Node B. Whether this conversion is achieved through regeneration or, in the future, with all-optical wavelength conversion, an extra cost will be incurred. In Fig. 4b, where different wavelengths are assigned to the two directions of connections A-C and A-D, it is possible to add the C-D connection without any need for intermediate wavelength conversion (assuming that different wavelength are assigned to the two directions of C-D, as shown)

While this is just a small example, this type of situation does arise in the design of real networks. Even when the number of wavelengths is large, there are scenarios where using different wavelengths in the two directions of a connection can result in a lower-cost network, due to less wavelength contention. It occurs most commonly when there are degree-3 nodes, with a lot of bypass traffic in all three directions through the node. (Furthermore, note that there are protection schemes where it is mandatory that different wavelengths be used in the two directions of a connection, e.g., some optical-layer shared ring protection schemes.)

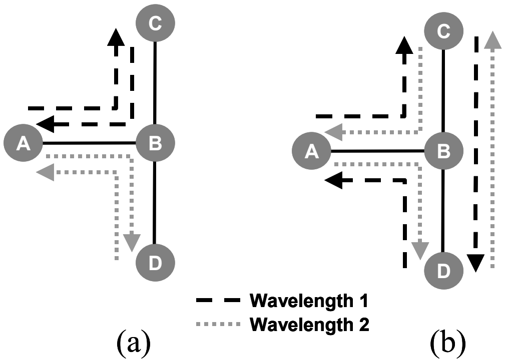

Fig. 4. Assume that there are only two wavelengths in this small network. (a) Wavelength 1 is assigned to both directions of the A-C connection, and Wavelength 2 is assigned to both directions of the A-D connection. If a connection between C and D is now added, that connection must undergo wavelength conversion at Node B in order to avoid wavelength conflicts with the existing connections. (b) Different wavelengths are assigned to the two directions of connections A-C and A-D. The connection between C and D can be added without any wavelength conversion, as shown.

VIII. NETWORK ELEMENTS

Network elements that enable optical bypass are typically more expensive than traditional O-E-O network elements. For example, at a degree-3 node, one degree-3 all-optical switch is generally more expensive than three optical terminals. The extra cost of the optical-bypass equipment is offset by the resulting reduction in number of regenerations at the node. Clearly, there is a minimum number of regenerations that must be eliminated at a particular node to justify the cost of the optical bypass equipment. After performing the routing and regeneration stages, each node can be checked to see if this threshold has been met. If not, some or all of the optical bypass capabilities can be removed at a node to lower the cost. For example, if a degree-2 node does not have enough traffic optically bypassing it, the OADM can be replaced by two back-to-back optical terminals. As another example, a degree-3 node may have significant bypass in one direction through the node but not the other two. In this case, the degree-3 all-optical switch could be replaced by the combination of an OADM and an optical terminal, where the OADM is oriented to accommodate the direction with the bypass traffic. (In practice, one may still choose to deploy full bypass at a node even if it is not justified by the cost, e.g., to cost-effectively accommodate future traffic changes.)

Removing bypass equipment from a node will add a regeneration in any connections that were previously bypassing the node. This additional regeneration may mean that a regeneration at another site for this connection can now be removed. Thus, after any equipment changes have been made, it may be worthwhile to run through the design process again to optimize for the current equipment.

IX. TOPOLOGY

While carriers provide the topology of their network along with the projected demand set, it is generally not mandatory in the design process that all links in the given topology be included in the final design. Long reach affords a greater opportunity to save cost by removing one or more links from the topology. By removing a link, the cost of the amplifiers along that link, along with the cost of the switch port at each of the two link endpoints, can be eliminated. The tradeoff is that some connections may need to be routed over a longer path; however, if the optical reach of the system is large enough, there may be little extra regeneration as a result of the longer paths. The cost of the extra regeneration needs to be compared to the cost savings obtained by removing the amplifiers and the switch ports to determine if the topology modification makes sense economically. The capacity and the protection capabilities of the resulting network need to be considered as well.

X. CONCLUSION

We have looked at network design in today's realistic backbone networks, with an emphasis on the impact of regeneration. The presence of regeneration, as opposed to a purely all-optical network, makes route selection somewhat more important, as the choice of route can directly impact the network cost. Conversely, the fact that wavelength conversion can occur for free when regenerating makes the wavelength assignment problem simpler. Results over a range of real networks have demonstrated that treating routing and wavelength assignment as separate steps can produce cost-effective, efficient designs.

Regenerating in the electronic domain also can be combined with other aspects of network design. For example, in shared mesh restoration, some of the segments that are created by regeneration can be used as the units of shared protection capacity. The segments are converted to the electronic domain at either endpoint, and thus can be passed through an electronic switch to be configured as needed for protection purposes. Or, if subrate demands are being routed, then regeneration can be combined with grooming. Passing a signal through an electronic grooming switch also regenerates the signal; thus, the grooming sites for a demand should be chosen as much as possible to coincide with needed regeneration for the demand.

The optical reach of today's ultra-long-haul systems is clearly shorter than the length of the longest demands. An interesting question is whether it would be worthwhile increasing the optical reach in order to eliminate all regeneration. Increasing the reach would require more expensive components; e.g., higher-power amplifier pumps, more precise lasers and filters, better forward error correcting chips, etc. It is unlikely that the increased system costs would be offset by the decrease in regeneration, as explored in more detail in [6]. At least from a capital cost perspective, it is probably not worth increasing the optical reach to eliminate all regeneration. However, there may be advantages from an operational standpoint.

References

[1] H. Zang, J. Jue and B. Mukherjee, "A review of routing and wavelength assignment approaches for wavelength-routed optical networks," *Optical Networks Magazine,* Vol.1, No.1, Jan. 2000, pp. 47-60.

[2] A. A. M. Saleh, "Defining all-optical networking and assessing its benefits in metro, regional and backbone networks," *OFC'03*, Atlanta, GA, Mar. 23-28, 2003, Paper WQ1.

[3] W. Van Parys, et al., "Quantifying the benefits of selective wavelength regeneration in ultra long-haul WDM networks," *OFC'01*, Mar. 19-22, 2001, Paper TuT4.

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

[5] J. Leuthold, J. Jaques, S. Cabot, "All-optical wavelength conversion and regeneration," *OFC'03*, Atlanta, GA, Mar. 23-28, 2003, Paper WN1.

[6] J. M. Simmons, "On determining the optimal optical reach for a long-haul network," *IEEE Journal of Lightwave Technology*, Vol. 23, No. 3, Mar. 2005, pp. 1039-1048.

[7] T. Carpenter, et al., "Cost-conscious impairment-aware routing," *OFC'04*, Feb. 22-27, 2004, Los Angeles, CA, Paper MF88.

[8] X. Yang and B. Ramamurthy, "Dynamic routing in translucent WDM optical networks: the intradomain case," *IEEE Journal of Lightwave Technology*, Vol. 23, No. 3, Mar. 2005, pp. 955-971.

[9] J. Hershberger, M. Maxel, and S. Suri, "Finding the k shortest simple paths: a new algorithm and its implementation," *Fifth Workshop on Algorithm Engineering and Experiments*, Baltimore, MD, Jan. 11, 2003.

[10] R.A. Barry and P.A. Humblet, "Models of blocking probability in all-optical networks with and without wavelength changers," *IEEE Journal of Selected Areas in Communications*, Vol. 14, No. 5, June 1996, pp. 858-867.

[11] X. Zhang and C. Qiao, "Wavelength assignment for dynamic traffic in multi-fiber WDM networks," *ICC'98*, Oct. 12-15, 1998.