

On Determining the Optimal Optical Reach for a Long-Haul Network

Jane M. Simmons, *Senior Member, IEEE*

Abstract— Systems that enable optical bypass, where traffic transiting a node can remain in the optical domain as opposed to undergoing costly O-E-O conversion, are gradually being accepted in carrier networks. An important factor in determining the cost effectiveness of such systems is the optical reach, the distance an optical signal can travel before needing to be regenerated. Longer optical reach results in a smaller number of required regenerations and hence less equipment and lower operating costs. In order to achieve longer reach, however, more expensive equipment such as amplifiers and transponders is typically needed. As the optical reach continues to increase, the cost benefit provided by reduced regeneration is eventually offset by the more expensive system equipment, leading to a concave curve of total network cost versus optical reach. We analyze four representative North American long-haul networks over a range of assumptions to determine the optimal optical reach from a cost perspective. In such networks, an optical reach in the range of 2,500 to 3,500 km yields the minimum, or close to the minimum, total capital cost over a wide range of assumptions, while representing a good tradeoff between decreased operating costs and increased initial network cost.

Index Terms— Optical networks, optical reach, optical bypass, regeneration, ultra long-haul networks

I. INTRODUCTION

IN traditional long-haul network systems, all traffic that enters a network node is converted from the optical domain to the electronic domain, regardless of whether the traffic is destined for that node. More recently, systems enabling “optical bypass” have been developed, which allow traffic transiting a network node to remain in the optical domain, rather than undergoing costly optical-electronic-optical (O-E-O) conversion. Many prior studies have described the potential benefits of optical bypass [1]–[3], including reduced capital costs, due to the elimination of a significant amount of O-E-O terminal equipment, and reduced operating costs, due to faster provisioning and reduced space, power, and sparring requirements.

North American long-haul carriers are gradually moving towards deploying networks with various degrees of optical bypass. There are two necessary ingredients for such systems: network elements that enable optical bypass, and extended

optical 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. Even though a node may be equipped with optical bypass equipment, some of the traffic transiting that node may still need to undergo O-E-O conversion for purposes of regeneration. (All-optical regeneration is currently not commercially viable.) Clearly the optical reach has to be at least twice the typical link length in order to achieve a significant amount of optical bypass in the network [3]. Many factors affect the optical reach: the type of amplification, the launched power of the signal, the modulation format of the signal, etc. In general, the longer the optical reach, the less regeneration is required in the network. However, achieving longer optical reach typically requires more expensive equipment, such as amplifiers and transponders [4]. (We use the term transponder to represent a combination transmitter/receiver card that has a short reach interface on the client side and a WDM signal on the network side.) At some point, as the optical reach continues to increase, the cost benefit provided by reduced regeneration is offset by the more expensive system equipment, leading to a concave curve of total network cost versus optical reach (i.e., the curve has a minimum-cost point). In addition, the initial capital costs of the network (i.e., network costs prior to any demands being provisioned) monotonically increase as the optical reach increases. Thus, increasing the optical reach of a system beyond a certain point may be counterproductive.

This paper examines what the optimal optical reach is for optical-bypass-enabled long-haul networks. It analyzes several networks that are representative of North American backbone networks, as described in Section II. Section III discusses how the optical reach affects the number of regenerations. It is shown that in a range of networks with very different topologies and demand sets, increasing the reach from 1,000 km to 3,000 km eliminates roughly 90% of the regenerations. Section IV examines total network capital cost as a function of optical reach. Based on the cost assumptions used in this section, an optical reach in the range of 2,000 km to 3,500 km results in a network capital cost that is within 5% of the minimum for all of the networks studied.

TABLE I.
TOPOLOGICAL PROPERTIES OF THE FOUR MODELED NETWORKS

	Number of Nodes	Number of Links	Average Nodal Degree	Average Link Length (km)	Total km of Network Fiber
Network 1	16	20	2.5	950	19,000
Network 2	25	38	3.0	700	26,500
Network 3	35	46	2.6	675	31,000
Network 4	55	78	2.8	400	31,200

In Section V, we examine the sensitivity of these results with respect to pricing and traffic pattern assumptions. It is shown that an optical reach in the range of 2,500 to 3,500 km yields the minimum, or close to the minimum costs, over a wide array of assumptions. Operational costs decrease as the reach increases, while the initial capital costs increase. Thus, the optimal operating point depends on the relative importance of these factors.

The study shows that the trend of system vendors to target an optical reach of beyond 5,000 km is not justified from a cost perspective. For example, at 6,000 km reach, the capital costs are 10% to 20% higher than the minimum, based on the cost assumptions used; the initial capital costs are roughly 70% higher. Increasing the reach from 3,000 km to 6,000 km results in a relatively small decrease in the number of regenerations, and hence, little improvement in operational costs. Even with the distance-independent traffic model considered in Section V, which should benefit from longer reach, 6,000 km is still not optimal.

II. NETWORK MODELS

In order to study the effect of optical reach in a long-haul network, we examine four typical North American backbone networks. These networks are representative of backbone networks in terms of both topology and traffic set, and the results provided here should be applicable to backbone carrier networks in general.

A. Topology and Demand Set

The topologies of the four networks are summarized in Table 1. Nodes are the add/drop and/or switching locations in a network. Links are the physical connections between nodes. The degree of a node is the number of incident links at that node; the average nodal degree in a network equals twice the number of links divided by the number of nodes.

While all four networks have an average nodal degree in the range of 2.5 to 3.0, which is typical of North American long-haul networks, they cover a range of network densities. At one extreme is the very sparse Network 1, which represents the express layer of a two-tier architecture, where only nodes that generate a large amount of traffic and geographically strategic nodes (i.e., junction sites) are included in the express layer [5]. At the other extreme, is the much more dense Network 4, which includes nodes with a much wider range of traffic levels.

Table 2 summarizes the demand sets of the four networks; all four networks are assumed to have a system rate of

TABLE II.
PROPERTIES OF THE DEMAND SETS OF THE FOUR MODELED NETWORKS

	Number of Demands	Percentage of Traffic Protected	Average Distance of Primary Path (km)	Average Distance of Secondary Path (km)
Network 1	350	25%	1,700	4,000
Network 2	375	55%	1,450	2,650
Network 3	350	40%	1,700	3,000
Network 4	14,000	45%	1,300	2,400

OC-192. Networks 1, 2, and 3 have only OC-192 demands; it is presumed that these demands represent wavelength services or subrate demands that have already been groomed into OC-192s. Network 4 includes demands that range from DS-3s to OC-192s, with roughly 50% of the demands (on a capacity basis) being subrate demands. Thus, a significant amount of grooming is required in Network 4.

The last two columns in Table 2 indicate the average path lengths of the routed demands, i.e., the distance from source to destination averaged over all routed demands. For Network 4, the average routed distances were weighted based on the service rate of the demand. For demands requiring protection, dedicated 1+1 client-layer protection was used; the primary and secondary paths were chosen to be maximally link- and node-disjoint.

B. Network Elements

In all of the network designs, it was assumed that network elements supporting optical bypass are available. At degree-two nodes (i.e., nodes with two incident links), OADM's were typically deployed. The OADM functionality is illustrated in Fig. 1; transponders are needed only for the add/drop traffic at a node (including regenerations), with the transiting traffic remaining in the optical domain. It is assumed that the OADM operates on a per-wavelength granularity; it can selectively drop any combination of wavelengths. It is also assumed that the OADM supports wavelength reuse. A wavelength that is dropped from one link of the OADM does not pass through the OADM, thus allowing the same wavelength to be added on the other link.

At nodes of higher degree, OADM-MultiDegree's (OADM-MDs) (also called Extended OADM's) were deployed. These devices provide the same optical bypass flexibility for transiting traffic as an all-optical switch; the difference is that, in contrast to an all-optical switch, the add/drop traffic does not pass through the switching portion of the device, which typically results in a lower cost element. (Thus, edge




Fig 1. Optical Add/Drop Multiplexer (OADM).




Fig. 2. A degree-3 OADM MultiDegree (OADM-MD).

configurability is not supported with an OADM-MD.) For implementation details of this network element, see [6] and [7]. An OADM-MD of degree D is functionally similar to $D(D-1)/2$ OADMs. In the networks considered in this study, almost all nodes had degree four or less. A degree-3 OADM-MD is functionally illustrated in Fig. 2. To enable both wavelength reuse and transit wavelength switching, the optical bypass element is assumed to be able to block or pass any desired wavelength.

In all of the network designs in this study, the most cost effective routing, i.e., the routing that minimizes regeneration in the network, was used for the particular topology, demand set, and optical reach assumptions. OADMs and OADM-MDs were only deployed at nodes where their cost was justified by the presence of optical bypass. For example, if a degree-2 node had no optical bypass, either due to the traffic pattern or due to a limited optical reach, then two back-to-back optical terminals would be placed at this node as opposed to an OADM (two optical terminal are expected to be lower cost than an OADM). As another example, consider Node A in Fig. 3, which has a nodal degree of three. Assume that the traffic pattern and optical reach justify optical bypass in the direction between Nodes C and D, but not between Node B and either Node C or Node D. Rather than deploy a degree-3 OADM-MD at Node A, an OADM is deployed along the links to Nodes C and D, and an Optical Terminal is deployed on the link to Node B. Note that in practice, one may still choose to deploy a degree-3 OADM-MD at Node A, e.g., to accommodate future traffic changes, however, this was not done in the present study.

C. Network Capacity

Some optical-bypass systems offer an optical reach versus capacity tradeoff; i.e., the optical reach can be increased, but at the expense of lower capacity. For example, a system may attain longer optical reach by launching the optical signal at a higher power level, but in order to mitigate the signal nonlinear impairments resulting from the increased power




Fig. 3. At Node A, it is assumed that optical bypass is advantageous in the direction between Node C and Node D, but not between Node B and either Node C or Node D. Thus, rather than deploy a degree-3 OADM-MD at Node A, a combination of an OADM and an Optical Terminal is deployed.

level, the system may have wider channel spacing, resulting in smaller system capacity.

In all of the networks studied, the largest load on any link was approximately 80 wavelengths, with each wavelength at an OC-192 rate. It was assumed that this capacity could be supported with any of the optical reaches considered in this study (80 wavelengths can be achieved with 50GHz channel spacing in the C-band, which is not technologically aggressive). The average load over all links in the network ranged from about 40 to 50 wavelengths for the four networks. This represents roughly a 50% increase over the average capacity utilized in today's backbone networks [8], and is in line with carrier traffic forecasts for the next few years. The average link load in a given network remained close to constant as the optical reach varied, even for the case of Network 4, where a significant amount of grooming was necessary to carry the substrate demands.

If required network capacities were to significantly increase beyond these numbers, then trading off reach for capacity would be an important consideration.

III. REGENERATION

A. Quantitative Analysis

This section discusses quantitatively how increasing the optical reach of the system affects the amount of required regeneration in a network. For each of the four reference networks, the optical reach was varied from 1,000 km to 6,000 km, in 500 km increments, and a minimal-cost network design was performed. For comparison purposes, a conventional system with 600 km optical reach and pure O-E-O nodes (i.e., the only network elements used were optical terminals) was also studied for each of the four networks. In any of the scenarios, if the optical reach was less than the length of a link, then one or more dedicated regeneration sites with back-to-back optical terminals were added along the link.

The number of regenerations versus optical reach for each of the four networks is plotted in Fig. 4. Compared to a conventional O-E-O-based system with 600 km reach, increasing the reach to 1,000 km and enabling optical bypass reduces the number of regenerations by roughly 50% in all of

the designs. In the region of relatively short optical reach, the number of regenerations is approximately a linear function of optical reach. Thus, this type of decrease is expected.

Furthermore, increasing the optical reach from 1,000 km to 3,000 km eliminates roughly 90% of the remaining regenerations; at a reach of 3,500 km, this factor is approximately 95%. As can be seen from all four graphs, the additional benefit in increasing the optical reach beyond the range of 3,000 or 3,500 km is relatively small, indicating that the associated decrease in operating costs would be small.

B. Regeneration Architecture

There are several architectural options for handling regeneration at a node. As indicated earlier, all-optical regeneration is not currently commercially viable, and thus only regeneration in the electronic domain is considered.

One option for implementing regeneration is to simply connect two back-to-back transponders, as illustrated in Fig. 5. The short reach interfaces of the two transponders are connected via a patch panel. The advantage of this scheme is its relatively low cost; the disadvantage is the network is not totally reconfigurable without requiring manual intervention. For example, the figure illustrates two regenerated connections: one in the East/West direction, and the other in the South/West direction. Manual intervention is required to take down these connections and establish a regenerated connection in the East/South direction.

Another option is to pass the regenerated traffic through an

optical switch, as shown in Fig. 6 (an optical switch that switches client signals is sufficient for this function; an all-optical switch is not necessary). In this application, the optical switch is acting as an edge or adjunct switch. The signal requiring regeneration enters the node and is dropped by the network element to a transponder that feeds into the optical switch. The signal can be directed to any of the other transponders connected to the switch; thus, it can exit the node on any of the links. This allows a greater degree of remote configurability compared to the patch panel option, but it does incur the cost of the switch. Depending on the type of optical switch and its internal loss, short reach interfaces may be needed at all switch ports, adding further to the cost.

A third option is to use special regenerator cards as opposed to back-to-back individual transponders; this is illustrated in Fig. 7. Functionally, a regenerator card is identical to back-to-back transponders; the major difference is short reach interfaces are not needed for the regenerator card, thus eliminating some of the cost. Of the three architectures described here, this is the lowest cost, but it is also the least configurable. Regenerator cards cannot be used in conjunction with an optical edge switch. Also, they cannot be used to add/drop traffic that is sourced/sunk at the node. The fact that the card is different from the normal transponder card leads to additional inventory issues as well. This is especially problematic if fixed-wavelength transmitters are used and the input and output wavelengths for the regeneration card are permitted to be different.

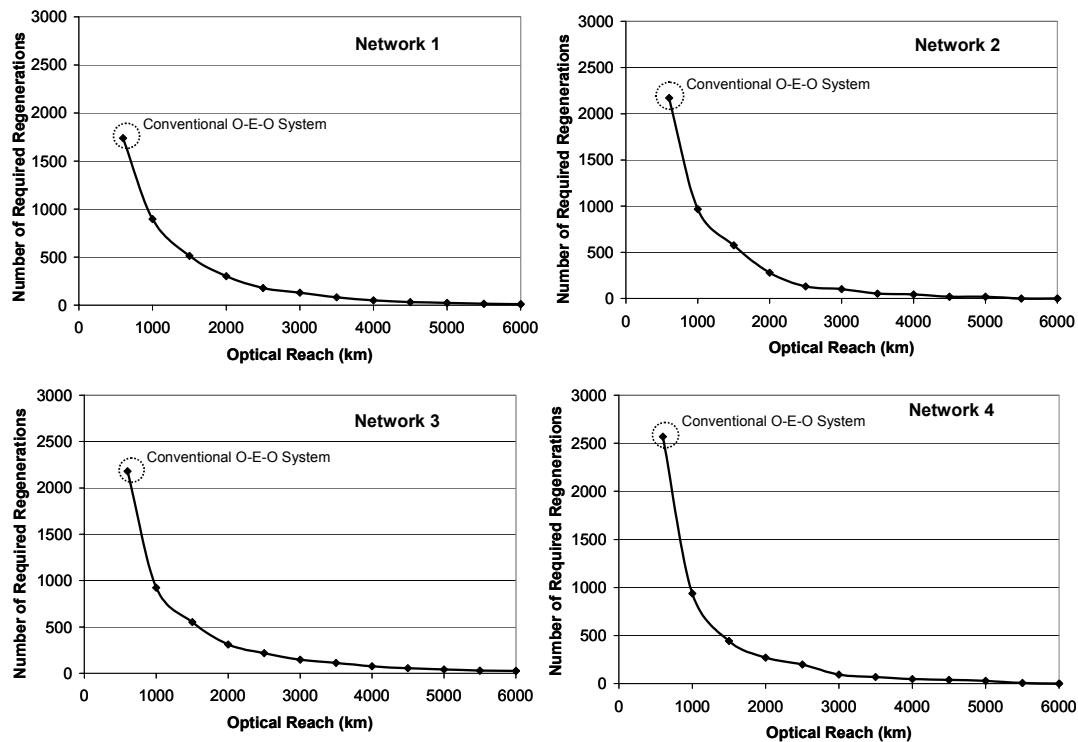


Fig. 4. The number of required regenerations in each of the four reference networks is plotted against optical reach. As a reference point, the number of regenerations required in a conventional O-E-O based system is also shown. The curve rapidly decreases as the reach increases from 1,000 km to 2,000 km, and then starts to flatten out. At a reach of 3,000 km, roughly 90% of the required regenerations have been eliminated, compared with 1,000 km reach.

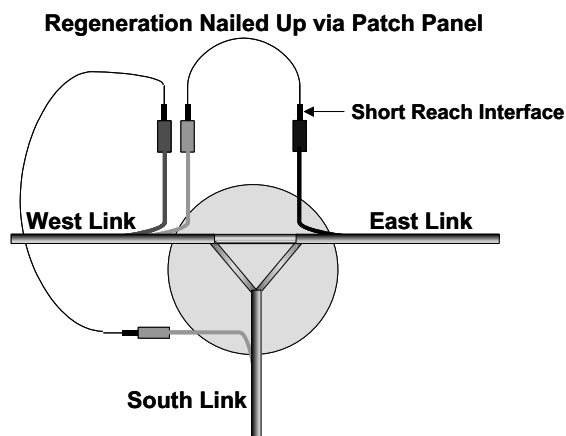


Fig. 5. A simple method of architecting a regeneration is to connect the short reach interfaces of two transponders via a patch panel. The figure illustrates two connections with regeneration: one connection in the East/West direction, the other in the South/West direction. Manual intervention is required to modify the connection patterns.

In all of the network designs performed for this study, it is assumed that the first option described above, i.e., directly connecting back-to-back transponders, is used for regeneration. Carriers typically evaluate capital costs based on a fixed demand set, thus, configurability of the design is not measured. Note that with this architecture, the designs as a function of optical reach are not entirely equivalent. Longer reach results in less regeneration, and thus less 'patch paneling'; in effect, a more configurable network.

Regenerator cards are not widely deployed in networks with extended optical reach, and thus were not used for this study. The use of regenerator cards would decrease the cost benefits of increasing the optical reach because regeneration would be less costly.

C. Taking Advantage of Regeneration

In addition to restoring the optical signal, regeneration can also serve other purposes. One important function is wavelength conversion [9]. When a signal undergoes O-E-O regeneration, there is no requirement that the outgoing wavelength be identical to the incoming wavelength, thereby providing an opportunity to wavelength convert at no extra cost. It has been shown in numerous studies that sparse wavelength conversion is sufficient for achieving utilization levels close to those of networks with full wavelength conversion [10]. It was shown in [11] that the small amount of regeneration that occurs in networks with extended optical reach provides enough opportunities to wavelength convert such that high network efficiency can be achieved.

Regeneration affords a good opportunity to groom traffic as well. Once the signal is converted to the electronic domain, it can be directed to a grooming switch to more efficiently pack the traffic. Ideally, grooming algorithms take into account optical reach when selecting optimal grooming sites.




Fig. 6. In this architecture, all add/drop at the node is directed to an optical edge switch. This provides flexibility to remotely configure a connection.




Fig. 7. In this architecture, a regenerator card is used, eliminating the short reach interfaces for regeneration. This is cost effective, but inflexible; it can lead to inventory problems if fixed-wavelength transmitters are used.

In the network designs described in this paper, regeneration was strategically combined with wavelength conversion, and for Network 4, with grooming, to produce efficient networks.

IV. COST ANALYSIS

Increasing the optical reach of a system generally increases the cost of the networking equipment. For example, to achieve an optical reach of greater than about 1,500 km, some amount of Raman amplification is likely needed, as opposed to more conventional, lower-cost erbium amplification. To increase the optical reach further may require using more pumps or higher-powered pumps, further adding to the cost of the amplifiers. In addition, more dispersion compensation or more precise gain flattening filters may be necessitated with longer optical reach. Transient control in the amplifiers is another area that needs to be addressed with long-reach systems, which adds further cost to the system. In addition to the amplifier cost, transponder costs typically increase as the reach increases due to the need for more complex modulation schemes, more precise lasers and filters, and more powerful error-correcting coding.

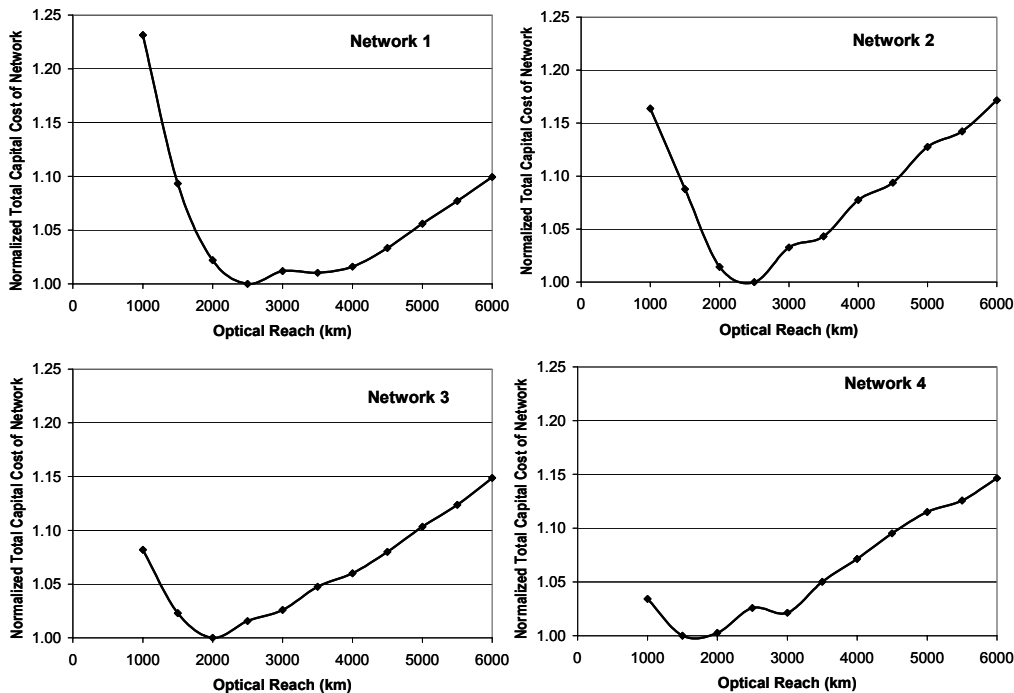


Fig. 8. The normalized total cost of the network is plotted versus optical reach for each of the four reference networks. The minimum cost point ranges from 1,500 km to 2,500 km, depending on the network.

A. Cost Model

In this paper, it is assumed that amplification and transponder costs increase by a factor of $P\%$ for every doubling of optical reach. Our reference optical reach is 1,000 km (there is little optical bypass in a long-haul network with shorter reach). Thus, for an optical reach of R , the cost increase factor is $(1+P/100)^{\log_2(R/1000)}$.

In all of the designs, the bulk of the cost was in transponders and amplification. The ratio of the cost of a bi-directional amplifier to a transponder was assumed to be 3.5.

It is assumed for a given reach that there is only one type of transponder card. For longer reach systems, it may be worthwhile to have two types of transponders, one for longer connections and a lower-cost transponder for shorter connections. This option leads to inventory and sparing issues, especially if the transponders are not tunable, and was not considered in this study.

The cost of the network elements (OADMs, OADM-MDs, and Optical Terminals) is broken into two major pieces. First, there is typically optical amplification at the input and/or output of the network element. It is assumed that this portion of the network element increases in cost at the same rate as the in-line amplification (i.e., increase of $P\%$ for every doubling in reach). The remainder of the network element, e.g., the portion that performs the actual add/drop of traffic, is assumed to be constant in cost as the optical reach increases. One could argue that elements with lower noise figure or flatter filtering, etc., are needed as the reach increases, which would lead to higher costs. However, since the complexity of the

element remains about the same, it is assumed the cost increases would be relatively small. (The ‘body’ of the network elements represented about 10% of the total cost of the network.)

As the number of transponders at a node increases, more nodal equipment bays, and possibly more nodal internal amplification, filters, etc. are needed. For simplicity, this cost is assumed to be equally spread over all transponders, and the amortized cost is included in the transponder cost.

Reference [4] examined the cost structure of first-generation ultra-long-haul (ULH) equipment, and concluded that the cost ratio between ULH systems and conventional reach systems for components such as amplifiers and transponders was roughly 2.5. At the time of the study, the optical reach of ULH systems was in the range of 2,500 to 3,000 km, representing a four to five-fold increase over conventional reach. Assuming the same model of a $P\%$ cost increase for every doubling of reach yields a P of about 50%.

Given that long-reach systems have significantly matured, and the cost of optical components for such systems have decreased, it is assumed in this section that P equals 25%. This number is consistent with more recent evidence. Section V examines how different values of P affect the results.

B. Network Capital Costs

As described in Section III, a minimal-cost network design was performed for each of the four reference networks as the optical reach was varied from 1,000 km to 6,000 km, in 500 km increments. Fig. 8 plots the normalized total network capital cost of the designs versus optical reach. The capital

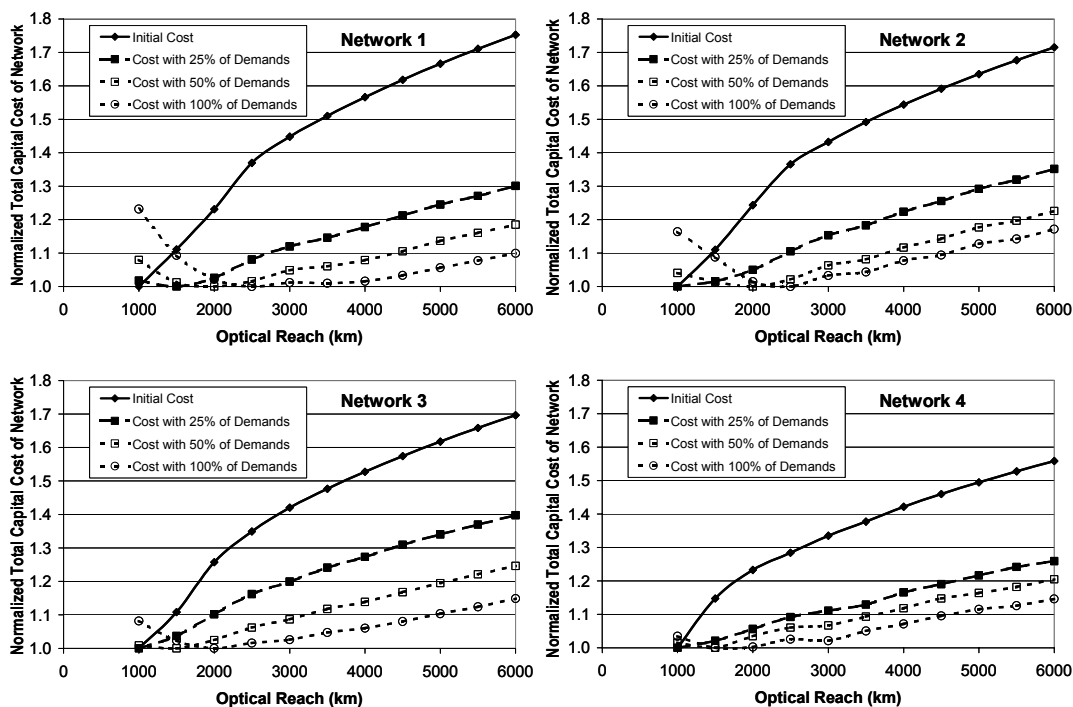


Fig. 9. The solid lines indicate the normalized initial cost of the network, which includes the cost of the network elements and the amplifiers. The initial costs steadily go up as the reach increases. The dashed lines indicate the normalized network cost at different network loads (for each load, the costs are normalized to the minimum-cost design for that scenario). Higher load shifts the minimum cost point to the right, until the system becomes fully loaded.

cost includes the cost of amplifiers, network elements, and transponders. For Network 4, it also includes the cost of grooming and muxing.

For both Networks 1 and 2, the minimal cost occurs at an optical reach of 2,500 km. For Network 3 the minimal point is 2,000 km. Network 4, which includes a significant amount of grooming, achieves a minimal cost at 1,500 km reach. The minimal-cost reach tends to decrease as the network density increases. However, for all four of the networks, an optical reach between 2,000 and 3,500 km yields a total cost that was within 5% of the minimal cost. Thus, it is possible to select a system reach that is cost effective over a range of networks.

An optical reach of 1,000 km yields a cost that is between 5% and 25% higher than the minimum, with the cost heavily dependent on the network. A reach of 1,000 km is relatively cost effective for the dense, grooming network of Network 4; however, it is not effective for the sparse topology of Network 1. At the other end of the spectrum, an optical reach of 6,000 km yields costs between 10% and 20% higher than the minimal cost.

The costs shown in Fig. 8 are for the total network after all demands have been provisioned. It is also important to consider the initial cost of the network, i.e., the costs before any demands are added. This represents the initial capital investment required, prior to generating any revenue. In the networks studied, the initial network costs represent about 30% of the total network capital costs at 1,000 km reach, and roughly 50% to 60% of the total network capital costs for reaches over 2,500 km.

The solid-line plots in Fig. 9 indicate the normalized initial network costs for each of the four reference networks. All of the initial-cost components increase at the rate of $P\%$ for every doubling in reach except for the ‘bodies’ of the network elements; thus, the only differences among the initial-cost plots shown in Fig. 9 are due to the number and type of network elements. At 3,000 km reach, the initial costs are approximately 40% higher than that for 1,000 km reach; at 6,000 km reach, this factor increases to roughly 70%.

Clearly, the minimum initial network cost occurs at 1,000 km, the shortest reach considered in this study. However, while the initial cost of the network increases with optical reach, the marginal cost of adding a demand to the network decreases because of less regeneration. Thus, as more demands are added to the network, the minimum cost point occurs at a longer optical reach. This can be seen in Fig. 9, which shows the normalized costs when 25%, 50% and 100% of the demands are added to the network (the 100% load curve is identical to that of Fig. 8). For each of these scenarios, the costs are normalized to the minimum-cost design for that scenario. As shown in the figure, as the load increases, the minimum-cost optical reach gradually shifts to the right. With 25% of the traffic, the minimum cost point is still 1,000 km for Networks 2, 3, and 4, and is 1,500 km for Network 1. Thus, from a cost perspective, extended optical reach is not beneficial at low loads. At 50%, the minimum cost points for Networks 1 through 4 are, respectively: 2,000 km, 2,000 km, 1,500 km, and 1,500 km. This compares with the respective

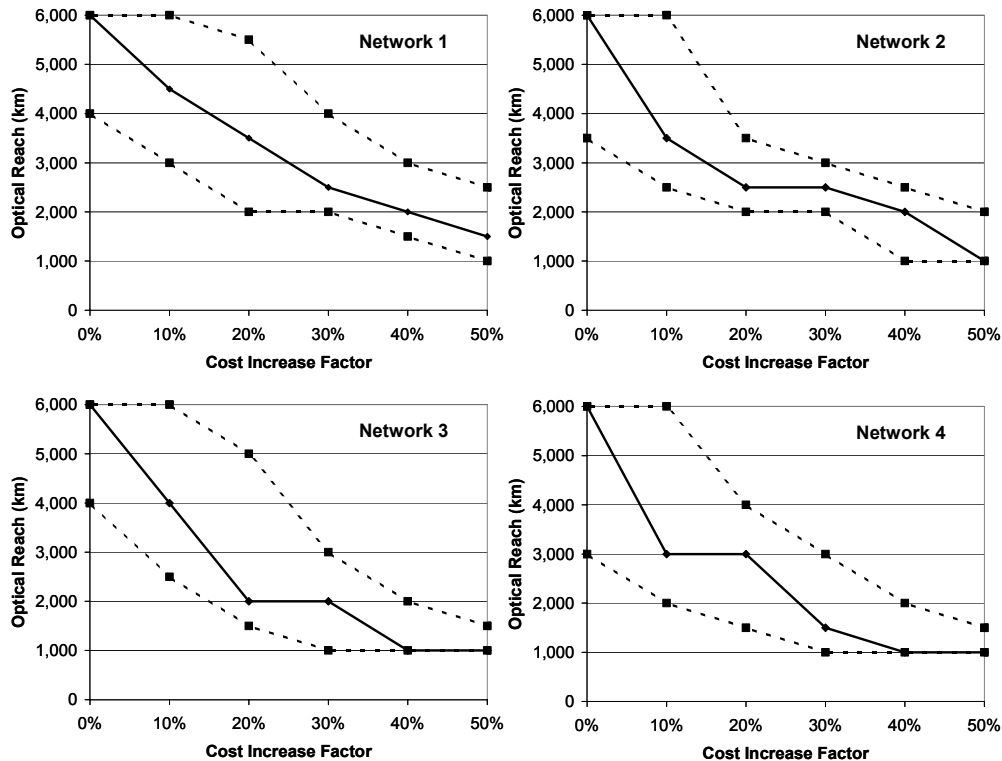


Fig. 10. The cost increase factor P is varied from 0% to 50%. The solid line in each of the graphs indicates the optical reach that achieves the minimum cost for a given value of P . Optical reaches that fall between the dashed lines achieve a cost that is within 5% of the minimum.

minimum cost points of 2,500 km, 2,500 km, 2,000 km, and 1,500 km in the 100% scenario.

The shifting to the right of the optimal point continues only until the system is fully loaded. As was discussed in Section II.C, the maximum capacity required in each of the four networks with 100% of the demands added is about 80 OC-192 wavelengths, which is assumed to be accommodated in the C-band. This is probably the capacity limit of the C-band for very long reach (say greater than 4,000 km). Thus, higher demand would likely require lighting the L-band or lighting more fiber, either of which would be relatively more expensive to do at longer reach. Thus, we do not expect that increasing the load beyond the scenarios considered here would shift the minimum cost point much further to the right.

C. Network Operational Costs

Operational costs of a network are in part related to the number of regenerations in a system. First, provisioning time is affected, as regenerations require deployment of more equipment along the connection path. A larger number of regenerations also leads to requirements for more physical space, power, heat dissipation, etc. In addition, sparring costs increase if more regeneration, and hence more equipment, is needed. These costs depend on many factors that are very difficult to quantify. Thus, while it is expected that fewer regenerations in a network lead to lower operating costs, we do not quantify this benefit in absolute cost terms.

The optical reach of a system also has an impact on network availability. As discussed earlier, the transponders in an extended reach system are typically more complex due to advanced modulation and coding schemes. Thus, the failure rate of an individual transponder is likely to increase as the optical reach increases. However, for a given end-to-end connection, longer reach leads to fewer regenerations, and hence fewer transponders, along the connection path. For a system with a reach of about 3,000 km, it has been shown that the network availability is improved compared to conventional reach; i.e., the increased failure rate of the transponders is more than compensated for by the reduction in number of transponders along the path of a connection [12]. However, the benefits are likely to diminish if the reach is increased much beyond 3,000 km. As was shown in section III, the number of regenerations decreases very slowly beyond this range, while it is expected that the complexity of the transponder, and hence, the failure rate, would continue to increase. At some point, as the reach continues to increase, the network availability may in fact start to decrease.

V. SENSITIVITY ANALYSIS

A. Sensitivity to Cost Increase Factor

The results shown in the previous section are dependent upon the assumption that transponder and amplifier costs increase by a factor of 25% for every doubling in optical

reach. In this section, a range of cost increase factors is considered.

For each of the four networks, the cost increase factor P was varied from 0% to 50%. The solid line plotted in the four graphs of Fig. 10 indicates the optical reach at which the minimal capital costs are attained for each value of P . The region between the dashed lines on each graph indicates the range of optical reaches that yields a cost that is within 5% of the minimum cost. It is assumed that being within 5% of the minimum is acceptable to a carrier, and that other factors would then be considered to select a system architecture.

An optical reach of 2,500 to 3,500 km is within the 5% range for all four networks for many of the values of P . This indicates that such a reach is a 'robust' choice as component costs change.

It is only if P is 10% or less that a reach of 6,000 km is within the 5% range. And it is only for a P of 0% (i.e., no price increase) that a reach of 6,000 km is within the 5% range and 2,500 km to 3,500 km is not within this range for all four networks. Thus, it is very difficult to justify a reach of 6,000 km based on capital costs. This, combined with the fact that operational costs will decrease very little after 3,000 km, indicates that 6,000 km optical reach is not an optimum choice. (Note that for P equal to 0%, the cost of the network decreases with reach until zero regenerations are present. For some of the networks, a reach beyond 6,000 km would have yielded a slightly lower cost network, however, 6,000 km was the maximum reach considered in this study.)

Throughout the study, we assumed that the system cost steadily increases with optical reach due to, for example, higher quality components, more dispersion compensation, etc. being required. It is possible there may also be spikes in the cost when major changes in technology are needed to increase the reach. For example, if increasing the reach beyond a certain distance necessitates the use of RZ (return to zero) technology as opposed to NRZ (non return to zero), the transponder cost may jump. If such a spike occurs in the range of 2,500 km to 3,500 km optical reach, and the spike is significant, then the optimal cost point would likely occur right before that cost spike is incurred.

B. Sensitivity to Traffic Pattern

The traffic pattern also plays a role in determining the optimal optical reach. If the average path length is very long, then one expects very long reach to be more beneficial. More regeneration is removed, which offsets the increase in system cost. It has been postulated that as Internet usage increases, traffic will become more distant independent and the average path length will increase [2]. It is not clear that this effect will materialize in a significant fashion because of distributed server technology. Nevertheless, in this section we consider a distance-independent traffic model.

We take Network 2, with 25 nodes, and assume that the demand set is all-to-all; i.e., each of the 25 nodes has a connection to each of the other 24 nodes. 25% of these demands are arbitrarily chosen to be 1+1 client-layer

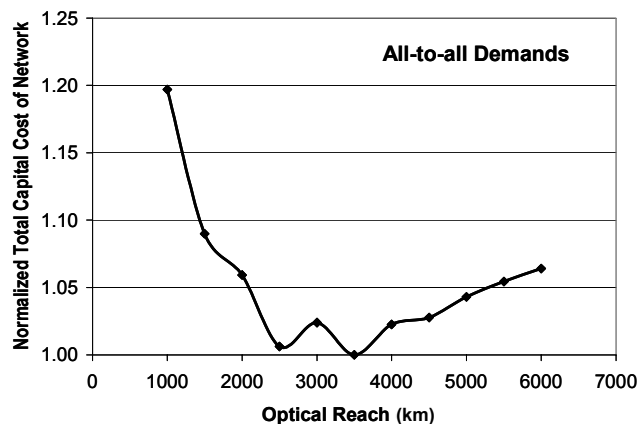


Fig 11. This plot shows the cost of Network 2 versus optical reach, with all-to-all traffic. Compared to the original demand set, all-to-all traffic requires longer average demand paths. The point of minimal cost increases to 3,500 km, from 2,500 km for the original demand set.

protected. (We chose this amount of protected traffic to yield a network where the most heavily loaded link had roughly 80 wavelengths, to match the load of the designs described earlier.) A minimum capital cost network was designed for this demand set; the average primary path distance was 2,550 km and the average secondary path distance was 4,000 km (compared to 1,450 km and 2,650 km, respectively, for the original demand set). With the cost increase factor equal to 25%, the cost plot as a function of optical reach is shown in Fig. 11. The minimal point is at 3,500 km, compared to 2,500 km for the original demand set; thus, as expected, longer optical reach provides relatively more benefit when the demand paths are longer. However, note that 2,500 km and 3,000 km reach are still within 5% of the minimum cost.

VI. CONCLUSIONS

This paper has studied the impact of increased optical reach on the cost of a network, where it is assumed that the cost of system equipment increases as the optical reach increases. In four networks that ranged from sparse to dense, and with a variety of demand sets, a reach of 2,500 km to 3,500 km yielded minimal, or close to minimal, capital costs over a wide range of assumptions. It was also shown that a reach of 3,000 km eliminates roughly 90% of the regenerations required in a network, as compared to a 1,000 km reach. Thus, operational costs will exhibit only a small decrease if the reach is increased beyond 3,000 km.

Some system vendors are targeting an optical reach of beyond 5,000 km. Based on the results of this study, this long of a reach is not justified from a cost perspective. In almost all scenarios, it produces a more costly network design.

It is concluded that for the long-haul networks and demand sets considered here, an optical reach in the range of 2,500 to 3,500 km provides the best tradeoff for capital costs, initial network costs, and operations costs. The optimal operating point depends on the relative importance of these factors.

REFERENCES

- [1] A. A. M. Saleh, "Transparent optical networking in backbone networks," *OFC'00*, March 5-10, 2000, ThD7.
- [2] R. E. Wagner, et al., "The potential of optical layer networks," *OFC'01*, March 17-22, 2001, TuT3.
- [3] A.A.M Saleh, "Defining all-optical networking and assessing its benefits in metro, regional and backbone networks," *OFC'03*, March 23-28, 2003, WQ1.
- [4] A.F. Wallace, "Ultra long-haul DWDM: Network economics," *OFC'01* Presentation, March 17-22, 2001, TuT1.
- [5] J.M. Simmons, "Economic and architectural benefits of hierarchical backbone networks," *LEOS Summer Topicals*, July 24-28, 2000.
- [6] I. Tomkos, et al., "Ultra-long-haul DWDM network with 320x320 wavelength-port 'Broadcast & Select' OXCs," *ECOC'02*, Sept. 8-12, 2002.
- [7] A. Pratt, et al., "40 x 10.7 Gbit/s DWDM transmission over a meshed ULH network with dynamically reconfigurable optical crossconnects," *OFC'03*, March 23-28, 2003, PD09.
- [8] M. Fuller, "Analysts see hope for U.S. long-haul," *Lightwave Europe Magazine*, June, 2003.
- [9] W. Van Parys, et al., "Quantifying the benefits of selective wavelength regeneration in ultra long-haul WDM networks," *OFC'01*, March 19-22, 2001, TuT4.
- [10] S. Subramaniam, M. Azizoglu, and A.K. Somani, "All-optical networks with sparse wavelength conversion," *IEEE/ACM Transactions on Networking*, vol. 4, pp. 544-557, August 1996.
- [11] J.M. Simmons, "Analysis of wavelength conversion in all-optical express backbone networks," *OFC'02*, March 17-22, 2002, TuG2.
- [12] B. Manseur and J. Leung, "Comparative analysis of network reliability and optical reach," *NFOEC'03*, September 7-11, 2003.