Cost vs. Capacity Tradeoff with Shared Mesh Protection in Optical-Bypass-Enabled Backbone Networks

Jane M. Simmons
Monarch Network Architects, Holmdel, NJ 07733

Abstract: Shared mesh protection based on cross-connecting pre-deployed protection subconnections is well suited for the optical-bypass-enabled networks currently being deployed. This scheme poses a cost-versus-capacity tradeoff, which is investigated through studies on several realistic backbone networks.

OCIS codes: (060.4510) Optical communications; (060.4250) Networks

1. Introduction

Optical-bypass technology is finally being deployed in carrier backbone networks on a large scale. Such systems allow traffic transiting a network node to remain in the optical domain rather than undergoing optical-electrical-optical (O-E-O) conversion. The combination of optical bypass and long-reach transmission eliminates a significant amount of O-E-O terminal equipment and reduces the operational requirements. While the amount of equipment needing to be deployed in the network is reduced, which potentially lowers the failure rate [1], providing protection for the network traffic is still critical. Dedicated 1+1 protection, with one working (i.e., primary) path and one hot standby path, is clearly one protection option. However, if efficient use of network capacity is important, and restoration times on the order of hundred(s) of milliseconds are tolerable, then some type of shared protection scheme is desirable.

There are numerous possible shared protection architectures as described in [2], [3]. For the core optical systems currently being deployed, it is essential that the protection scheme be compatible with the new technology. First, the scheme should be able to take advantage of optical bypass. If the amount of bypass is not large enough, the cost premium for the all-optical network elements will not be justified. Maintaining a significant level of optical bypass for both the working and protection capacity is desirable. Second, the scheme ideally should be path-based, such that the same end-to-end protection mechanism is triggered independent of where in the path the failure occurs. Fault isolation can take longer in systems with optical bypass because the signal is not electronically terminated after each hop. Thus, fault-independent protection schemes are favored. Third, the protection scheme must be compatible with the underlying physical system. In long-reach systems, amplifier transients that arise when the power level on a link is suddenly changed are a potential problem that can cause error bursts. Protection schemes where, for example, lasers are rapidly turned on or off, or connections are rapidly switched in the optical domain, may result in undesirable network transients.

In this paper, we assume the shared protection scheme is based on pre-deploying protection 'subconnections' that are cross-connected at the edge as needed at the time of failure in order to form an end-to-end protection path. (Pre-deploying subconnections in an optical-bypass-enabled network and 'stitching' them together on-demand to form connection paths was described in [4].) Individual subconnections can traverse multiple links, where the intermediate nodes are optically bypassed. In addition, the protection subconnections are always 'lit', so that the scheme avoids issues with power transients.

The operation of this protection scheme will be described in more detail in Section 2. Along with its many operational advantages, the scheme poses an interesting tradeoff of network cost versus network capacity. As will be illustrated in Section 3, the shorter the protection subconnections, the more opportunities there are for sharing, leading to lower capacity requirements. However, shorter subconnections require more terminal equipment (i.e., transponders, edge cross-connect ports), resulting in higher cost. This paper explores the cost versus capacity tradeoff in more detail, using the results of a network study performed on several realistic backbone networks. The network study and its results are covered in Section 4. Variations on the protection scheme are discussed in Section 5.

2. Cross-Connect-Based Shared Mesh Protection

Fig. 1 is used to illustrate the operation of the shared mesh protection scheme analyzed in this paper. It is assumed that each network node is equipped with a core switch that is capable of optical bypass. This core switch can be, for

example, an all-optical switch or a multi-degree optical add/drop multiplexer (OADM-MD), as will be addressed below. Many, if not all, of the nodes are also equipped with an edge cross-connect. In contrast to the core switch that enables optical bypass, the edge cross-connect can be a more traditional OEO-based switch, or it could be, for example, a small MEMS-based switch.

The first step is to deploy protection subconnections such that the appropriate level of protection is provided for the network traffic. Fig. 1 depicts three protection subconnections as indicated by the thick solid lines: A-D, D-E, and D-G. Each subconnection occupies one wavelength on the links it traverses; it is assumed all subconnections are bi-directional. Each subconnection terminates in a transponder card at its two endpoints; the transponders feed into an edge cross-connect at the node. Assuming that the subconnections are shorter than the system optical reach, there is no need for transponders at any intermediate point of the subconnection. For example, note that the A-D subconnection optically bypasses nodes B and C. (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. Optical reach is the distance an optical signal can travel before the signal quality degrades to a level that necessitates regeneration.)

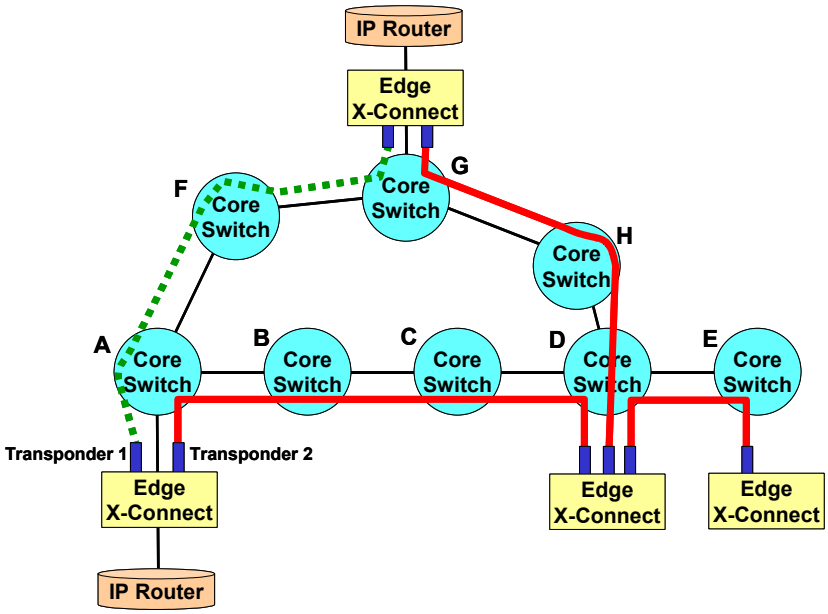

Fig. 1. The three protection subconnections are indicated by the thick solid lines. The working path between A and G is shown by the dotted line. If there is a failure anywhere along this working path, the edge cross-connect at D will concatenate the A-D and the D-G subconnections to form the protection path.

The subconnection transponders are always 'on'; i.e., the protection wavelengths are always lit. When a failure occurs that disrupts a working connection, one or more edge cross-connects are configured such that the appropriate subconnections are concatenated to form an end-to-end protection path. Once the protection path is established, the failed connection switches over to that path. For example, consider the working A-G connection, as indicated by the dotted line in the figure. If this connection fails, the edge cross-connect at Node D will connect the A-D subconnection to the D-G subconnection to form the A-G protection path (assuming no other connection is already using these protection subconnections). In addition, the edge cross-connects at Nodes A and G will be reconfigured such that the client traffic is directed to the protection path (e.g., the IP router at Node A will be cross-connected to Transponder 2). After the failure is restored, the connection reverts back to its original path, and the subconnections can be used to protect other connections.

This protection scheme has several advantageous features. It is path-based so that the same protection path is formed regardless of the fault location. The protection subconnections are always lit, thus, there is no issue with transients. There is no need to tune any transponders to a different wavelength when a failure occurs; the scheme can be used in a system that does not have tunable transponders. There also is no need to reconfigure the core switch elements; only the edge cross-connects need to be configured at the time of failure. In addition, because the subconnections are pre-deployed, there is no need for real-time impairment-based routing; i.e., at the time of failure, there is no need to analyze the physical layer to determine where regeneration is needed along the protection path. This simplicity in operation should translate into a more rapid recovery time. The scheme is also advantageous in that it provides protection against transponder failures in addition to link and node failures.

Furthermore, this scheme is compatible with an OADM-MD type of core switch, as opposed to being compatible with just the more flexible all-optical switch. The OADM-MD and the all-optical switch are illustrated at a functional level in Fig. 2(a) and Fig. 2(b), respectively. All-optical switches are typically architected using more scalable technology, thereby allowing for a larger switch fabric. The key differentiating property is that the add/drop traffic passes through the all-optical switch fabric, which allows a transponder to access any of the network ports. In the OADM-MD, the add/drop traffic does not enter the switch fabric; the transponders on its add/drop ports are tied to a particular network port. For example, in Fig. 1, if the core switch at Node A is an OADM-MD, then Transponder 1 is tied to the port to/from link A-F. This restriction is not an issue with the cross-connect-based protection scheme described here because at the time of failure, the edge cross-connect connects the client (e.g., an IP router) to Transponder 2, such that the protection path can go out on link A-B. For technology and/or cost reasons, some networks have been deployed with OADM-MDs as opposed to all-optical switches. Thus, compatibility with such core switches can be important.

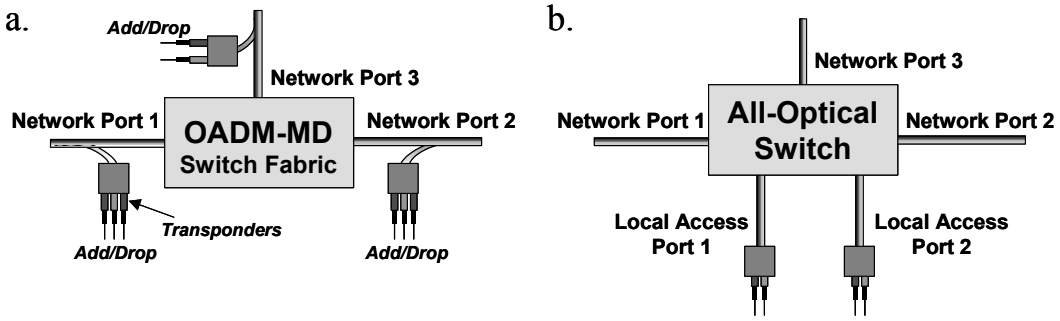

Fig. 2. a) OADM-MD with three network ports. The add/drop traffic is tapped off of the network ports; each transponder can access only one network port. b) All-Optical Switch with three network ports and two local access ports. The add/drop traffic passes through the switch fabric allowing any transponder to access any network port.

The main drawback of this shared protection scheme is the cost of the edge cross-connect. An alternative scheme based on pre-deployed protection subconnections that does not rely on an edge cross-connect was proposed in [5]. This scheme uses the combination of tunable regenerator cards and all-optical switches to concatenate the protection subconnections. (Regenerators are similar to back-to-back transponders with the intermediary short reach interfaces eliminated.) At the time of failure recovery, the scheme requires retuning of some transponders and regenerators and requires reconfiguration of the all-optical switches at the endpoints of the connection (the scheme is not compatible with an OADM-MD). In addition, it requires that the regenerators be turned off on the first and last protection subconnections of the path while the transponders at either end of the original connection are reconfigured to light these same subconnections. The process of simultaneously bringing down the regenerator power while bringing up the transponder power on that same wavelength needs to be done very carefully in order to avoid transients due to power fluctuations. Or, these operations can be done in series; however, in order to avoid transients, the operations need to be performed slowly, which will increase the restoration time. Overall, the scheme of [5] avoids the cost of the edge cross-connect (plus uses regenerator cards, which are less expensive than back-to-back transponders) but requires greater operational complexity.

In the next section, we investigate the tradeoff of cost versus capacity that arises from shared mesh protection based on pre-deployed protection subconnections. This analysis applies to both the cross-connect-based scheme described above, as well as the scheme proposed in [5].

3. Cost vs. Capacity Tradeoff

The fewer hops traversed by a protection subconnection, the more opportunity there potentially is to share the subconnection, leading to more efficient use of capacity. However, each subconnection requires a transponder and a port on the edge cross-connect at each of the two endpoints. Thus, deploying shorter subconnections leads to more subconnections being deployed, which results in higher cost. This tradeoff is more clearly illustrated in the network example shown in Fig. 3. It is assumed that each of the nodes in this figure is equipped with a core switch capable of optical bypass.

In this figure, there are three working connections, as shown by the dotted lines: A-E, A-J, and K-E. In Fig. 3(a), there are four protection subconnections to protect this traffic, as shown by the thick solid lines: A-G, G-E, G-J, and K-A. If, for example, connection K-E fails, then the protection path is formed by using the edge cross-connect at Node A to concatenate the K-A and the A-G subconnections and using the edge cross-connect at Node G

to concatenate the A-G and the G-E subconnections. There are a total of eight transponders needed for the four protection subconnections in this example. If one counts each hop of a protection subconnection as having unit capacity, then a total of seven units of capacity are required for protection.

Contrast this design with the protection subconnections deployed in Fig. 3(b). The working connections are the same; however, there are only three protection subconnections to protect the traffic: A-E, A-J, and K-A. There is clearly less sharing achieved with this configuration; e.g., the A-J subconnection is used to protect just the A-J working connection. The protection subconnections require a total of only six transponders, however, the total protection capacity is now nine units. Thus, using longer subconnections reduces the network cost but increases the amount of protection capacity needed. In a realistic network, such tradeoffs occur on a larger scale, as will be quantified in the next section.

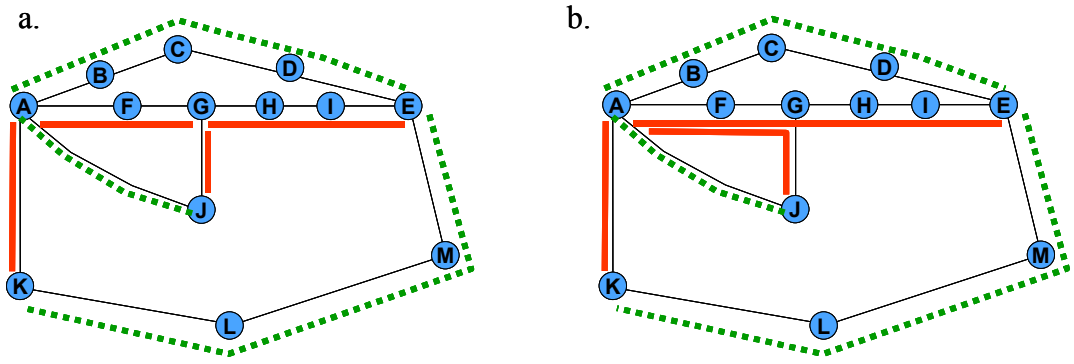

Fig. 3. The three working paths are shown by the dotted lines. The protection subconnections are indicated by the thick solid lines. (a) Four protection subconnections are deployed, requiring 8 transponders and occupying 7 units of capacity. (b) Three protection subconnections are deployed, requiring only 6 transponders but occupying 9 units of capacity.

4. Network Study

In order to quantify the tradeoff of cost versus capacity, we analyzed subconnection-based shared mesh protection in four different backbone networks. The analyzed networks are representative of North American backbone networks in terms of both topology and traffic set. The networks covered a range of sizes as indicated by the topological statistics shown in Table 1.

Table 1. Topological Properties of the Backbone Networks Used in the Study

	Network 1	Network 2	Network 3	Network 4
Number of Nodes	18	32	55	74
Number of Links	23	43	73	95
Average Link Length (km)	1000	750	450	400
Average Nodal Degree	2.6	2.7	2.7	2.6
Percentage of Nodes with Degree 2	56%	50%	51%	60%

All demands in the traffic sets were assumed to require shared mesh protection; the demands were added one-by-one to the network, with no knowledge of future demands. Several designs were performed for each network, where the order in which the demands were added was varied. The results presented below represent the averages for each network. Enough demands were added such that the resulting capacity requirement on the most heavily loaded link was on the order of 80 to 100 wavelengths. All demands were at the line rate (i.e., no traffic grooming was needed).

All network nodes in the designs were assumed to be equipped with core switches capable of optical bypass. The optical reach of the system was assumed to be 3,000 km.

Alternative-path-routing was used for routing the working connections. Up to five paths were considered for the working path of each connection; a path was deemed acceptable if its length was less than 1,000 km or if its length was no more than 50% greater than the shortest possible working distance for that connection, and if the path required no more than the minimum possible number of regenerations for that connection. Of these possible paths, as each demand was added to the network, the working path was selected that resulted in the minimum additional needed protection capacity.

As demands were added, the design algorithm ensured enough protection subconnections were deployed such that recovery was possible from any single link or node failure. In order to illustrate the tradeoff of cost versus capacity, a methodology was used where certain network nodes were designated as protection hubs. The hubs served as sites where the protection capacity was chopped to form protection subconnections; i.e., protection

capacity routed through the hubs was electronically terminated and fed into an edge cross-connect (while most protection capacity was broken at the hubs, more efficient results were obtained by allowing a small amount of optical bypass at the hubs). If the distance between hubs for a particular subconnection was longer than the optical reach, then regeneration was added such that shorter subconnections were formed. Increasing the number of hubs resulted in shorter subconnections, which in turn led to more sharing, but more cost. Thus, varying the number of nodes selected as protection hubs illuminated the cost versus capacity tradeoff. Note that other methodologies could be used to illustrate this tradeoff (e.g., [6]), however, the approach of selecting protection hubs is operationally very simple. (The notion of designating protection hubs is similar to the selection of 'high-level' nodes in the hierarchical shared mesh protection scheme of [7].)

Hub nodes were selected based on several factors. One key factor was the degree of the node (i.e., the number of links incident at the node); higher degree favored the selection of the node as a hub because this allowed concatenation of subconnections entering the node from all directions. Nodes that generated a lot of traffic were favored as hubs because many protection subconnections would start/end at these nodes anyway (to protect the node's own traffic) and thus represented good junction sites. In addition, nodes that were topologically located such that they were good regeneration sites were also favored as hubs; optical reach would cause many subconnections to terminate at these nodes anyway. Note that the hub nodes only affected the protection traffic. The working paths were regenerated solely based on path distance.

As each demand was added to the network, the design process reused existing protection capacity as much as possible. Two working connections were allowed to share a protection subconnection if the two working paths had no links or nodes in common. One restriction that was enforced was that a particular protection subconnection could not be shared by more than five working paths. As was shown in [8], this represents a good tradeoff between capacity efficiency and susceptibility to a second failure.

4.1. Network Study Results

The results of the network study are illustrated in the graphs shown in Fig. 4. These graphs plot the normalized total number of transponders needed in the network versus the normalized total required capacity in the network. (Any regeneration was implemented as two transponders; the network capacity is measured in wavelength-km.) The totals shown are for the combination of the working and the protection traffic. The number of required transponders gives an idea of the network cost. Thus, these graphs are representative of the cost versus capacity tradeoff exhibited by the shared mesh protection scheme.

The percentage of network nodes that were selected as protection hubs is indicated at each point along the graph. (Slightly different percentages were selected for the different networks, based on the number of nodes in the network with degree-3 or higher.) As expected, as the number of protection hubs increases, the number of required transponders increases, but the required capacity decreases. The amount of possible capacity savings varied by network size. For the very sparse Network 1, there was a difference of approximately 5% in capacity as the number of hubs was varied. However, for the much larger Network 4, there was a 20% difference in capacity. Similarly, the range of the required number of transponders varied by network size. For Network 1, the difference in number of transponders varied by up to 25% over the range of designs. For Network 4, it varied by up to 40%. (This is not to say that the total network costs vary by this large percentage; there are clearly other capital costs in a network. In a full network, as a rough approximation, transponders represent somewhere around 50% of the total cost.) These results indicate that selecting the operating point is more critical as the network size increases.

Assuming that networks will eventually fill with traffic, decreasing the capacity requirements will ultimately lower the overall network cost by postponing the investment in a network upgrade. Thus, it is desirable to select an operating point where both the required number of transponders and capacity are relatively low. Selecting 15 to 20% of the nodes to be protection hubs produced relatively good designs for the various networks. In this range, both the capacity and the required number of transponders were no more than ~10% greater than their respective minimums.

The graphs in Fig. 4 also plot the average amount of optical bypass in the network at each of the design points. As the number of hubs increases, the amount of optical bypass decreases. However, the amount of optical bypass is still very substantial even with a high degree of sharing. At the point where 20% of the nodes are protection hubs, the average optical bypass achieved in Network 1 is roughly 40%; for the much more dense Network 4, it is about 65%. (Shorter link distances provide more opportunity for nodes to be bypassed, assuming the same optical reach.) Clearly, optical bypass technology can be effective in reducing the amount of overall required terminal equipment, even in a network with shared mesh protection [9].

The largest edge cross-connect needed in each of the network designs was on the order of about 200x200, however, roughly 90% of the nodes required an edge cross-connect of size smaller than 128x128.

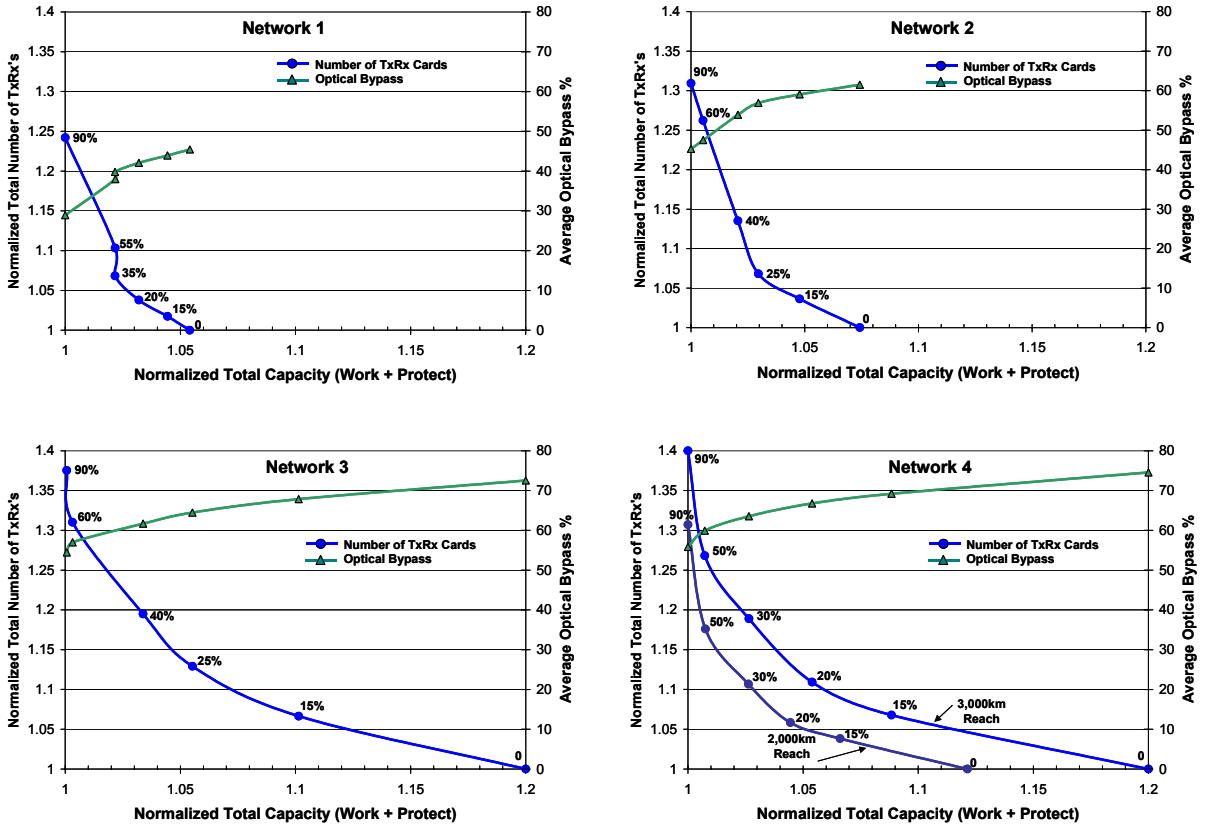

Fig. 4. The curves with the circular markers indicate the normalized total number of transponders (TxRx's) versus the normalized total capacity, for each of the four networks. The labels indicate the percentage of nodes that were selected as protection hubs. The curves with the triangular markers indicate the average optical bypass in the network. All curves are based on an optical reach of 3,000 km, except for the TxRx vs. Capacity curve with a 2,000 km reach shown for Network 4.

4.2. Effect of Optical Reach

The results presented above were based on an optical reach of 3,000 km, which is in-line with the reach of recently deployed systems, e.g., [10], [11]. With higher transmission rates, e.g., 40 Gb/s, the optical reach will likely be reduced. Shorter reach will result in increased regeneration, which will naturally allow more sharing of the protection capacity. To illustrate this effect, we generated designs for Network 4 with an optical reach of 2,000 km, using the same methodology as above. The resulting normalized transponder vs. capacity curve is shown in Fig. 4. The tradeoff between cost and capacity is less severe as compared with a 3,000 km reach. Selecting 15 to 20% of the nodes to be hubs again resulted in a good design point, with both the transponders and the capacity being no more than ~5% greater than their respective minimums.

4.3. Comparison with Dedicated 1+1 Protection

Designs with dedicated 1+1 protection were also generated for each of the four networks. Shared mesh protection, regardless of the operating point, is much more efficient in the use of capacity. Depending on the network, approximately 40 to 60% of the capacity can be saved by deploying shared mesh protection as opposed to dedicated protection. The cost savings, however, are not as clear. The number of required transponders was anywhere from 5% to 25% higher with dedicated protection. This does not directly translate into higher cost, however, because the dedicated protection transponders do not have to terminate in edge cross-connects. Thus, it is possible to have dedicated protection designs that are actually lower cost than a shared mesh design, depending on the architecture.

5. Variations of the Protection Scheme

5.1. Break at Hubs Only When Needed

In the design strategy described above, protection capacity was electronically terminated at the protection hubs when first deployed. Another strategy is to assume that the protection hubs represent *potential* junction sites for the protection capacity; protection capacity is not broken at a hub until it is needed to enable sharing of the capacity. For example, referring back to Fig. 3(a), assume the first two added demands are A-E and K-E. At that point, there

is no need to divide up the A-E protection subconnection. If the A-J demand is added next, and Node G is designated as a protection hub, then the A-E subconnection is broken into the A-G and G-E subconnections, such that the A-G subconnection can be shared to protect the A-J demand. Operationally, this means that transponders are added after a protection subconnection has already been established. However, the protection capacity does not carry live traffic most of the time, thus, there should not be a problem with reconfiguring a subconnection that already has been deployed.

This strategy was used to generate designs for the same four networks as above. The graphs of normalized total transponders versus normalized total capacity are shown in Fig. 5 for each of the networks. The curves that were plotted in Fig. 4 are also shown on these graphs for comparison purposes. As expected, if the subconnections are only broken at the hubs when needed, fewer transponders are needed. This is especially true when the percentage of hubs is very high, where approximately 10% of the transponders can be saved while still requiring the same amount of capacity.

Note, however, that this strategy would not be practical if the traffic is highly dynamic, with connections being rapidly added and deleted. There would not be time to add in the transponders as needed to divide up a protection subconnection; the transponders would need to be pre-deployed as was assumed in Section 4.

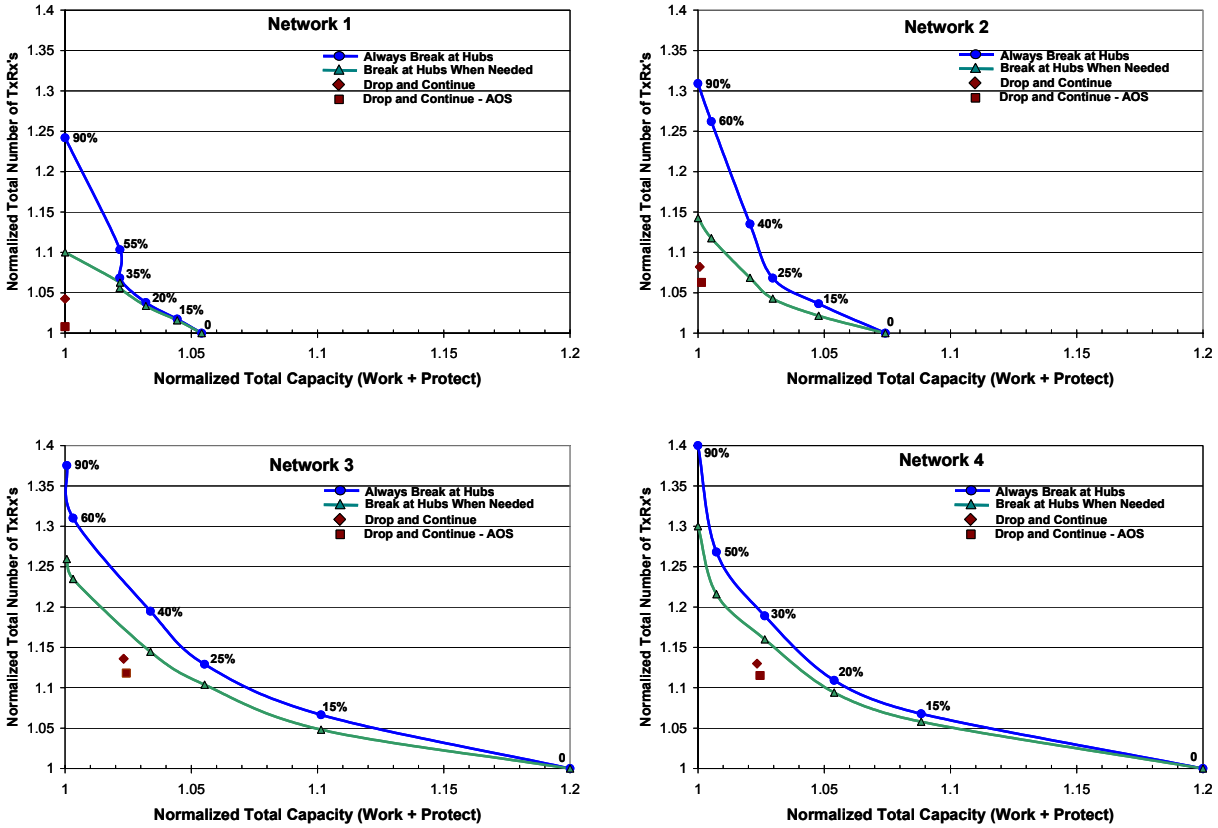

Fig. 5. The curves with the triangular markers indicate the normalized transponder vs. capacity curve for the strategy where protection capacity is broken at the hubs only when needed. The curves from Fig. 4, indicated by the circular markers, are included again here for comparison purposes. Breaking the protection paths at the hubs only when needed clearly requires fewer transponders for the same amount of required capacity. The results from using a drop-and-continue strategy are also shown, including results that take advantage of an all-optical switch (AOS). All designs used a 3,000 km optical reach.

5.2. Drop-and-Continue

Many of the core switching elements that provide optical bypass are also capable of drop-and-continue, where a signal is both dropped at a node and transmitted all-optically to the next node. This feature can be employed to implement a variation of the shared mesh protection scheme discussed above. Consider the example shown in Fig. 6(a). The working paths are the same as in Fig. 3. Note that the A-E subconnection also drops at Node G. This allows this subconnection to be concatenated to the G-J subconnection to provide protection for the A-J connection. When used to protect A-J, only the portion of the A-E subconnection that runs between A and G is utilized.

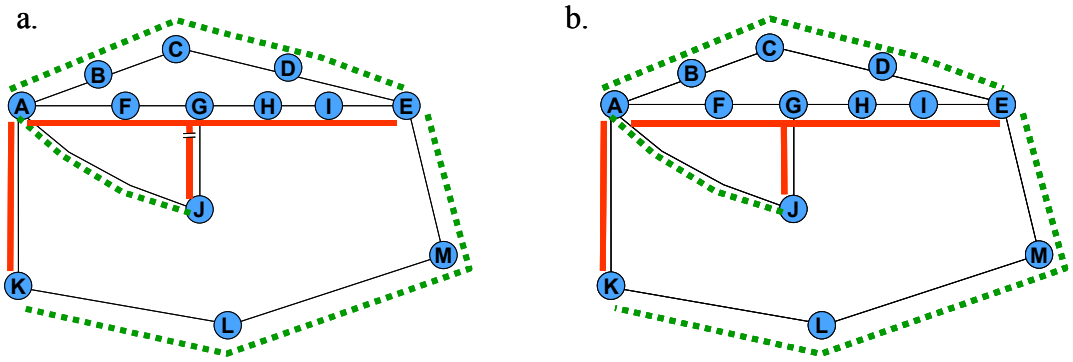

Fig. 6. The working paths, which are the same as in Fig. 3, are indicated by the dotted lines. The protection subconnections are indicated by the thick solid lines. (a) Three protection subconnections are deployed, requiring 7 transponders and occupying 7 units of capacity. The A-E subconnection includes a drop-and-continue at Node G. (b) Two protection subconnections are deployed, requiring only 5 transponders and occupying 7 units of capacity. One subconnection spans both A-E and G-J; this subconnection all-optically branches at Node G.

Fig. 7 provides more detail of the architecture of Node G. Just one transponder is needed at Node G for the A-E subconnection. Overall, the architecture of Fig. 6(a) requires 7 transponders and 7 units of capacity, compared with Fig. 3(a) that requires 8 transponders and 7 units of capacity. (Note that the strategy of Fig. 6(a) only provides a benefit when just one portion of a subconnection is needed for sharing. If the G-E portion of the subconnection needed to be shared as well, then one would end up with separate A-G and G-E subconnections, as in Fig. 3(a).)

While the drop-and-continue scheme can potentially provide a better operating point, it is operationally more complex. Referring to Fig. 7, the transponder at Node G on the A-E subconnection is normally kept off. When the working path of A-J fails, this transponder must be turned on. At the same time, it is necessary for the core switch at Node G to block the signal on the A-E subconnection coming from Node E (otherwise it would interfere with the protection traffic being carried from G to A). This procedure of bringing up the transponder at G while simultaneously reconfiguring the switch at G to block the signal from E must be done very carefully so that the power level remains close to constant from G to A. (These operational challenges are similar to those of the scheme in [5], although no retuning of wavelengths is needed.)

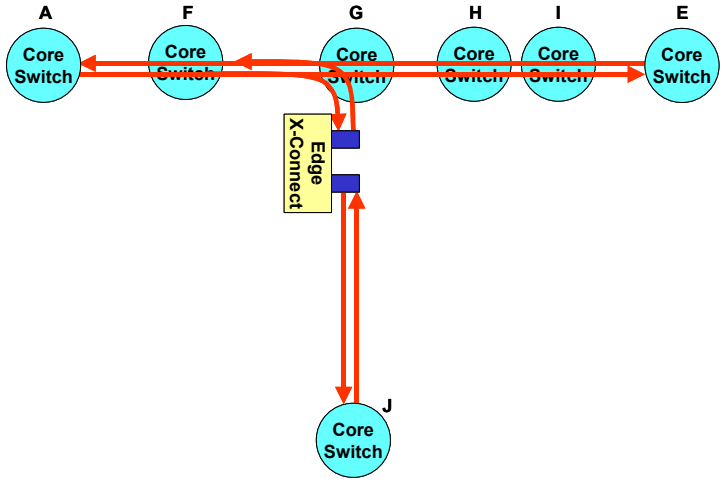

Fig. 7. Node G of Fig. 6(a) is shown in more detail here. The A-E subconnection drops at Node G, but also continues on to Node E. The transponder at G on the A-E subconnection is normally off. It must be turned on when there is a failure on the A-J working path.

Designs using the drop-and-continue feature of Fig. 6(a) and Fig. 7 were generated for each of the four networks. It is assumed that 'drop' transponders can be added after a subconnection has been established (e.g., the drop transponder at Node G could be added after the A-E subconnection is deployed). In addition, the 'break only when needed' strategy described in the previous section was used. Approximately 15% of the nodes were set as protection hubs; the 'drops' could be added at nodes that were not hubs. The diamond-shaped points in the graphs of Fig. 5 indicate the normalized required number of transponders and required capacity for this protection strategy. For the same number of protection hubs, drop-and-continue increases the number of required transponders (because drops

are added at non-hubs), but reduces the capacity requirements (because more sharing can occur). It produces better results than increasing the number of hubs in order to achieve the same reduction in capacity; i.e., for the same capacity requirements, drop-and-continue reduces the number of transponders by about 5%. A carrier would need to consider whether this strategy offers enough of a benefit to justify the additional complexity. (It is possible that more savings could be realized with different traffic patterns or different design strategies.)

If the core switches are all-optical switches, as opposed to OADM-MDs, another option is possible. With an all-optical switch at Node G in Fig. 7, the 'drop' transponder at Node G could be used to form an A-G subconnection *or* a G-E subconnection. At the time of failure, the all-optical switch at G would need to be configured to form the desired subconnection. This provides more flexibility with the drop transponders. However, the capacity requirements may increase because a working path that requires the A-G subconnection for protection must be totally disjoint from a working path that requires the G-E subconnection (because there is only one drop transponder at G that can be utilized at the time of a failure). If there were two transponders at Node G, such that the A-G and G-E subconnections are independent, then the associated working paths do not need to be disjoint.

Network designs were performed based on the drop-and-continue strategy combined with all-optical switches, again with ~15% of the nodes as protection hubs. The results are shown in Fig. 5 by the square-shaped points. Compared to the original drop-and-continue results, the number of required transponders is reduced by about 2%. The required capacity increased slightly.

One can take the drop-and-continue strategy one step further. Rather than having any transponders at Node G on the A-E protection subconnection, there could be one large branching subconnection that includes both A-E and G-J. This is illustrated in Fig. 6(b). By reconfiguring the switch at Node G, this large subconnection can be used as an A-E, A-J, or J-E subconnection. The advantage of this scheme is that no transponders are needed at G for this subconnection. However, more capacity will likely be required as all the working paths that use these three subconnections must be mutually disjoint. Furthermore, the same wavelength must be used on all six links traversed by the subconnection (assuming all-optical wavelength converters are not present). Implementing this on a large scale in a network may lead to problems with wavelength contention. Designs were not performed using this strategy.

6. Conclusion

Shared mesh protection based on cross-connecting pre-deployed protection subconnections is operationally and functionally compatible with optical-bypass-enabled networks. We have investigated the cost versus capacity tradeoff with this protection scheme through network designs on four realistic backbone networks, using a strategy where a subset of the nodes were selected as protection hubs. The cost versus capacity tradeoff was more pronounced as the networks increased in size. For all of the networks, selecting 15 to 20% of the nodes to be protection hubs produced relatively good operating points.

Variations of the protection scheme were studied that took advantage of the drop-and-continue capabilities that are often provided by core optical switches. While these schemes can produce benefits in reducing cost or capacity, they are operationally more difficult to implement.

It is important to note that the same scheme of pre-deploying subconnections and concatenating them together to form end-to-end paths also can be used for rapidly provisioning traffic in a dynamic network. A deployed subconnection could be used for protection or for provisioning of new traffic, depending on the state of the network. This method is also compatible with the GMPLS paradigm, with the subconnection ID serving as the 'label', as described in [4]. Such a scheme should be operationally simple for both provisioning and protection and also efficient with respect to capacity.

7. References

[1] B. Manseur and J. Leung, "Comparative analysis of network reliability and optical reach," *NFOEC'03*, September 7-11, 2003.

[2] G. Ellinas, et al., "Routing and restoration architectures in mesh optical networks," *Optical Networks Magazine,* Jan/Feb. 2003, pp. 91-106.

[3] W. Grover, Mesh-based Survivable Transport Networks: Options and Strategies for Optical, MPLS, SONET and ATM Networking, Upper Saddle River, NJ, Prentice-Hall, 2003.

[4] J.M. Simmons, A.A.M. Saleh, and L. Benmohamed, "Extending Generalized Multi Protocol Label Switching to configurable all-optical networks," *NFOEC'01*, Baltimore, MD, July 8-12, 2001.

[5] G. Li, A. Chiu, and J. Strand, "Failure recovery in all-optical ULH networks," *Design of Reliable Commun. Networks*, Oct. 16-19, 2005.

[6] J. Weston-Dawkes and S. Baroni, "Mesh network grooming and restoration optimized for optical bypass," *NFOEC*, Dallas, TX, Sept. 2002.

[7] J.M. Simmons, "Hierarchical restoration in a backbone network," *OFC'99*, San Diego, CA, Feb. 21-26, 1999, TuL2.

[8] R. Ramamurthy, "Limiting sharing on protection channels in mesh optical networks," *OFC'03,* Atlanta, GA, Mar. 23-28, 2003, TuI3.

[9] R. Ranganathan, et al., "Express lightpaths and shared protection in optical mesh networks," *ECOC'02*, Copenhagen, Sept. 8-12, 2002.

[10] P. Hofmann, et al., "DWDM long haul network deployment for the Verizon GNI nationwide network," OFC/NFOEC'05, Anaheim, CA, Mar. 6-11, 2005, OTuP5.

[11] M. Bortz, "Broadwing's experience with optical network planning and deployment," *OFC/NFOEC'06*, Anaheim, CA, Mar. 5-10, 2006.