

Hierarchical Restoration in a Backbone Network

Jane M. Simmons, *AT&T Labs*
100 Schulz Dr., Red Bank, NJ, 07701
jane@research.att.com, (732) 345-3220, (732) 345-3040 (fax)

1. Introduction

One promising approach to provisioning and restoration in long-haul Wavelength Division Multiplexed (WDM) networks is to deploy a mesh of optical cross-connects that operate on individual wavelengths. Such cross-connects provide the ability to reconfigure the network topology in response to changing network conditions; this includes allowing configurable optical bypass of a node. Here, we describe a cross-connect-based, *hierarchical* restoration scheme for backbone networks that reduces the required port-count of the cross-connects, that enables a significant amount of optical bypass, and that scales well with increasing traffic demand.

Hierarchical network architectures generally involve partitioning the network nodes and links into multiple levels, where the intricate details of one layer are hidden from the layers above it. Each successive layer operates on a simpler abstraction of the real network, allowing more tractable management of large networks. Such schemes have typically been used to simplify network routing.[1] In our proposed two-level hierarchical restoration scheme, a small number of the network nodes are designated as high-level nodes, with the remaining nodes being low-level nodes. The high-level nodes serve as the anchor points for failure recovery; i.e., recovery takes place between high-level nodes, as opposed to the nodes at either end of the failure.

There are a multitude of virtues in implementing a hierarchical restoration scheme. First, using the high-level nodes as the anchor points for failure recovery results in large bundles of recovery traffic that can be expressed between high-level nodes without needing to be terminated at intermediate low-level nodes. This allows optical bypass to be implemented at the low-level nodes, resulting in the elimination of a significant amount of electronic terminating equipment. (An analysis of the potential benefit of optical bypass in rings with a hierarchical architecture was presented in [2].) The expressing of large bundles of traffic also reduces the required cross-connect port-count. Furthermore, as traffic demands increase, hierarchical restoration enables the use of coarser granularity, and thus more cost-effective, cross-connects. Thus, the approach scales well with traffic demand. Additionally, using the high-level view for restoration greatly simplifies rerouting decisions, resulting in rapid failure recovery.

Hierarchical restoration generally works more efficiently if the hierarchical structure is also used to route traffic under non-failure conditions, however, this is not a requirement for hierarchical restoration to operate.

2. Operation of a Hierarchical Restoration Scheme

Consider the portion of a mesh network shown in Figure 1A, where we assume each node is equipped with a cross-connect. First, we examine failure recovery assuming a 'flat' mesh-based architecture, where all nodes are considered equal peers. When a failure occurs, the recovery traffic may need to be partitioned into finer granularity units, each of which follows a different recovery path. For example, assume Link 1 fails. It may not be possible to find a single recovery path that can accommodate all of the traffic that normally flows from Node C to Node B over this link (protection capacity is often provisioned this way to reduce the amount of overall required protection capacity). A possible recovery strategy may be to send a portion of the traffic over the path C-F-E-A-B and a portion over the path C-F-I-H-A-B. Next, consider recovering from a failure of Link 4 (independent of whether Link 1 has failed); one of the protection paths might be G-F-I-J-K. Thus, the protection link between Nodes F and I is utilized if either Link 1 or Link 4 fails.

Focusing on Node F, we see that some protection traffic may arrive from Node C and need to be split onto Links 3 and 5. Also, Node F needs to be able to take protection traffic from Node G and send it out on Link 5. This is illustrated in Figure 2. Typically, the amount of protection traffic that remains intact going through Node F will be less than a full wavelength. Thus, the protection traffic must be electronically terminated to allow grooming, so that the protection wavelengths can be more fully utilized. For example, grooming at Node F may allow just one protection wavelength to be deployed on Link 5 as opposed to two.

The need to groom the protection traffic at most, or all, of the nodes leads to a significant amount of required terminating equipment for the protection links. There is very little opportunity to optically bypass a node, thus, the cross-connects at each node need to be sized to handle all of the protection traffic that is routed through it. Furthermore, cross-connect ports are needed to allow the protection traffic to be dropped and added. These requirements arise because of the many 'ad-hoc' protection paths that exist in the network.




Figure 1 A. A portion of a mesh backbone network, with all nodes being equal peers. Restoration typically occurs between the nodes on either side of a failure, and over multiple paths. Recovery from the failure of Link 1 is shown in the figure. B. A high-level abstraction of the corresponding portion of the network shown in A. Restoration occurs between high-level nodes.

Next, consider imposing a hierarchical architecture on the network shown in Figure 1A. We designate Nodes A, K, and S as high-level nodes, with the remaining nodes being low-level nodes. High-level nodes will typically be those nodes with relatively high degree that serve as major intersection points of the network, although it is not necessary to designate all high-degree nodes as high-level. In a backbone network, it is reasonable to designate on the order of 10% of the nodes as high level. The high-level view corresponding to Figure 1A is shown in Figure 1B; the details of the lower level are hidden in this high-level perspective. From a top-level abstraction, the network can be viewed as being comprised of interconnected high-level nodes, thus greatly simplifying restoration decisions.

Note that in Figure 1A, there are three node disjoint (and hence link disjoint) paths between each of the high-level node pairs. In a long distance backbone network, it is common to be able to find three such paths. However, the restoration scheme depends on the existence of only two diverse paths, not three.

Along the paths between any pair of high-level nodes, we distinguish between express traffic and local traffic (relative to this pair). We define express traffic as traffic whose source and destination do not lie along the paths between the two nodes. Conversely, local traffic is traffic that has a source and/or destination lying along the paths between the high-level nodes.

2.1 Express Traffic Recovery

Assume there is a link failure somewhere along one of the paths between high-level Nodes A and K; e.g., assume that Link 1 in Figure 1A fails. We assume Nodes A and K monitor the health of all links between them, or receive notification of their status from a centralized controller (i.e., the scheme can work in either a decentralized or centralized fashion). Consider express traffic that normally is routed from Node K, over D, C, and B, to Node A. When Node K learns of the Link 1 failure, its cross-connect immediately shunts this express traffic over one, or both, of the remaining express paths between K and A (i.e., K-G-F-E-A and K-J-I-H-A), as opposed to routing the traffic until the point of failure is reached. Node A does the same for the affected express traffic in the reverse direction. If there are multiple wavelengths worth of traffic to restore, then some of the recovery wavelengths are sent over the path K-G-F-E-A, and some are sent over the path K-J-I-H-A. (Splitting the traffic over multiple paths typically leads to less overall required protection capacity. If there are only two diverse paths, then all of the recovery traffic is routed over the sole remaining path.) The most important point is that for any failure along the paths between A and K, the recovery of express traffic takes place between nodes A and K, as opposed to the two nodes directly adjacent to the failure. None of the express recovery wavelengths need to terminate at any low-level nodes between A and K.

2.2 Local Traffic Recovery

Next, consider the recovery of local traffic, for example in Figure 1A, traffic that originates at Node D, passes through C, and terminates at Node B. When Link 1 fails, the traffic is routed on protection links from C to K; from there it can be express routed to Node A over the two remaining express paths, and from Node A it is routed to Node B. Thus, the recovery takes place with a combination of local protection fiber (e.g., from C to K) and express protection fiber. The local protection fiber needs to be terminated at low-level nodes.

As another example of local traffic recovery, consider traffic that originates at C and is routed via B and A to Node H. When Link 1 fails, the traffic is routed on protection links from C to K; at that point there are several options. One choice is to send all of the protection traffic directly to H from K. Or, the traffic can be express routed to A using both the top and middle express paths, and from A be sent to H.




Figure 2. Node F must split the protection capacity coming from Node C, and merge one portion of it with the capacity coming from Node G. In order to perform this grooming, Node F must electronically terminate the protection links.




Figure 3. A hierarchical cross-connect architecture. Implementing optical bypass and handling some traffic on the fiber-level cross-connect relieves the port-count requirements of the wavelength-level cross-connect.

3. Hierarchical Cross-connect Architecture

Using the high-level nodes as anchor points for recovery results in large bundles of protection traffic that can express through the low-level nodes; i.e., there is a very large amount of optical bypass at the low-level nodes. Correspondingly, the amount of required electronic terminating equipment can be significantly reduced, resulting in large cost savings. In addition, the scheme provides several opportunities for decreasing the port-count requirements of the cross-connects. If the traffic demand is large enough, whole protection fibers (as opposed to individual wavelengths) can be expressed between high-level nodes. This enables a hierarchical cross-connect architecture, as shown in Figure 3, where both a fiber-level and a wavelength-level cross-connect are used.

Switching traffic on a fiber-level cross-connect reduces the port-count of the wavelength-level cross-connect, i.e., these fibers are not demultiplexed into their constituent wavelengths and delivered to the finer grain cross-connect. Furthermore, at low-level nodes, where entire fibers may bypass a node, it is possible to run the fiber through the node without any terminations at all; this somewhat limits the flexibility of the architecture, but reduces the size of the fiber-level cross-connect.

If a fiber contains some wavelengths that need to drop at a node, then that fiber is delivered to a drop port by the fiber-level cross-connect and demultiplexed. If some of the demultiplexed wavelengths bypass the node, then they can be immediately routed to the multiplexer without entering the wavelength-level cross-connect (or a wavelength add/drop multiplexer can be used). Thus, the wavelength-level cross-connects in low-level nodes need to be sized to handle the local protection traffic, but not the express protection traffic. If the local traffic through some low-level nodes is small enough, it may be worthwhile to completely remove the cross-connect from these nodes (assuming, in the event of a failure, it is acceptable to not recover the small amount of local traffic, and assuming the cross-connect is not needed for provisioning).

To add flexibility, wavelengths that bypass a node can be switched through the wavelength-level cross-connect. However, they do not need corresponding add and drop ports on the cross-connect, so a port-size benefit is still realized.

Note that as the number of wavelengths per fiber increases, the number of nodes where an entire protection fiber bypasses a node is likely to be small. In this scenario, rather than using a fiber-level cross-connect, it may be more advantageous to deploy a ‘band level’ cross-connect, that operates on fixed bands of wavelengths.

4. Conclusions

We have presented a cross-connect based, two-level hierarchical architecture for providing restoration in a backbone network. By expressing large bundles of protection traffic between high-level nodes, this scheme reduces the amount of terminating equipment and the cross-connect sizes that are required in the network. As traffic demand increases, the amount of express traffic will correspondingly increase, which will amplify the cross-connect and bypass advantages; thus, a hierarchical restoration scheme scales very well with demand.

References

- [1] Bertsekas, D. and Gallager, R., *Data Networks*, 2nd Edition, Prentice-Hall, 1992, pp. 379-382.
- [2] Gerstel, O., Ramaswami, R., and Sasaki, G., “Cost Effective Traffic Grooming in WDM Rings,” *Infocom '98*, March 29-April 2, 1998, San Francisco, CA, pp. 69-77.