Dynegy's National Long-Haul Optical Mesh Network Utilizing Intelligent Optical Switches

Pambos Charalambous*, Chad Dennis*, Georgios Ellinas[§], Eric Bouillet[§], Jean-François Labourdette[§], Sid Chaudhuri[§], Mikhail Morokhovich[§], David Shales[§]
*Dynegy Communications, 1000 Louisiana, Houston, TX 77002
[§]Tellium Inc., 2 Crescent Place, Oceanport, NJ 07757

Abstract: We present Dynegy's long-haul optical mesh network utilizing Tellium's intelligent optical switches. Shared mesh restoration is used to protect against failures and a re-optimization algorithm is implemented to regain portion of the protection capacity used.

Wavelength Division Multiplexed (WDM) networks that route optical connections using intelligent optical cross-connects (OXCs) are firmly established as the core constituent of next generation networks. With connection rates reaching tens of Gb/s, preventing and repairing failures is increasingly becoming an integral part of the network design process. In end-to-end dedicated (1+1) mesh protection (Figure 1), the ingress and egress OXCs of the failed connection attempt to restore the signal on a predefined backup path that is disjoint from the primary path. Path diversity guarantees that primary and backup paths will not simultaneously succumb to the same failure. Diversity of routes in an optical network is defined using the notion of Shared Risk Groups [1]. A set of optical channels that have the same risk of failure is called a Shared Risk Group (SRG). For example, all the channels that are multiplexed onto a WDM fiber-link is a shared risk group since the failure of the fiber-link simultaneously affects all the channels that are carried over that fiber. SRGs are configured by the network operator with the knowledge of the physical fiber plant of the optical network. Even though dedicated protection requires large amount of capacity, the backup path remains "live" in permanence, thus saving crucial path-setup latency when recovery takes place. In shared mesh restoration (Figure 2), backup paths can share capacity if the corresponding primary paths are mutually diverse. Compared to dedicated protection, this scheme allows considerable saving in terms of capacity required [1]. In addition, the backup resources can be utilized for lower priority preemptible traffic in normal network operating mode. However recovery is slower than dedicated protection, essentially because it involves signaling and path-setup procedures to establish the backup path. In particular, we note that the restoration time will be proportional to the length of the backup path and the number of hops, and if recovery latency is an issue this length must be kept under acceptable limits. However this constraints may increase the cost of the solutions, as it is sometime more cost-effective to use longer paths with available shareable capacity than shorter paths where shareable capacity must be reserved. This tradeoff can be handled by an appropriate cost model in the route computation algorithm [2].

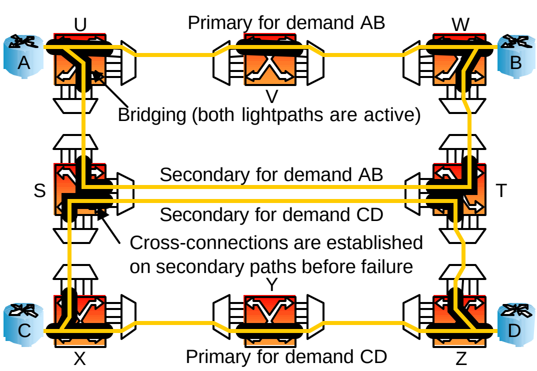

Figure 1. Dedicated Mesh (1+1) Protection

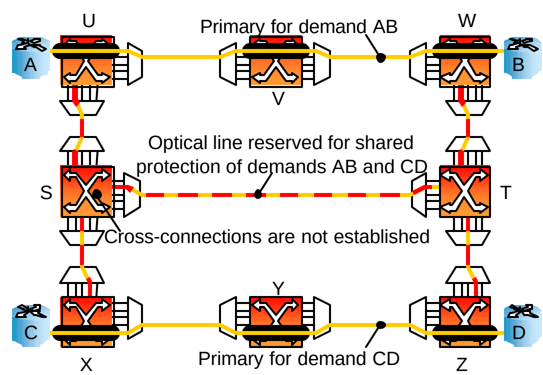

Figure 2. Shared Mesh Restoration

A cost model assigns costs to links in the network that represents some cost of using the channel in a lightpath route (e.g. fiber-mileage). The quality of the lightpath route is the sum of costs of all channels in the route. For the cost model we define the length of a path as the sum of the predefined weights of the edges that constitute it. The metric or policy used for weighting the edges is different for primary paths and backup paths. For primary paths it is the real cost of using the edges. For backup path it is a function of the primary path [3]. A backup edge e is assigned: (1) Infinite weight if it intersects with an SRG of the primary path. (2) Weight w_e if new capacity is required to provision the path, and (3) Weight $s_e = \varepsilon w_e$ $0 \le \varepsilon \le 1$, if the path can share existing capacity reserved for pre-established backup paths. The cost of a primary and its protection is then the sum of their respective lengths. The underlying idea here is to encourage "sharing", whereby existing capacity can be reused for provisioning multiple backup paths. The ratio s_e to w_e can be adjusted for the desired level of sharing. For smaller values of s_e, backup paths will be selected minimizing the number of non-shareable edges (weights w_e) in view, eventually leading to arbitrary long paths (as expressed in number of hops) that consist uniquely of shareable edges (weights s_e). For larger values of s_e, routing is performed regardless of sharing opportunities and backup paths will end-up requiring substantially more capacity (see [2] for more details). The routing of each lightpath will attempt to minimize the total cost of all channels in the lightpath route. A user-defined cost is assigned to fiber-links that reflects the real cost of using a channel on that fiber.

Provisioning of shared mesh restored lightpaths in Dynegy's live network, that utilizes Tellium's intelligent Aurora Optical Switches, was achieved by calculating the working and backup paths using the weight assignment as described above. Dynegy's network is on the order of 50 nodes and 70 trunks, and is carrying shared mesh restored demands amounting to several hundred Gigabits of service. A number of tests have been conducted for this network as outlined below. All tests were successful and restoration time targets were met.

Test Scenarios:

- Switching traffic to the backup path on all shared lightpaths provisioned in the network.
- Restoration upon single failure: verify that a successful lightpath protection switching takes place upon a single failure.
- Performance tests (switching time measurements) on the most and least stressed links of the network, on an express, local and hybrid link: verify that lightpath switching time complies with specifications.
- Wavelength Management System (WMS) lightpath re-provisioning between two locations: verify that WMS will trigger the re-provisioning of a new backup path when both primary and backup paths are disconnected between two locations.
- Restoration upon a node failure: verify lightpath restoration during a node failure at a site.
- System start-up after a node failure at a site: monitor the system recovery after the node failure.

During operations, requests for services are received and provisioned using an online routing algorithm that takes all of the information available at the time of the request to make the appropriate routing decision. Demand churn and network changes such as the addition/deletion of new links and/or capacity, causes the routing to become sub-optimal, thereby creating opportunities for improvements in network bandwidth efficiency. Re-optimization seizes on these opportunities and offers the network operator the ability to better adapt to the dynamics of the network. Re-routing backup paths only is an attractive way to regain some of the protection bandwidth and reduce backup path length while avoiding any service interruption. Table 1 summarizes the results for the re-routing of backup paths in Dynegy's live network during a maintenance window. For this exercise 65 backup paths were re-optimized and tested in 5 hours. The port count of Table 1 consists of ports used for the protection channels only. The table shows the quantities measured before and after re-optimization. We observe that backup path re-optimization saves 31% of the protection ports which in turn translates to 20% savings of the total number of ports. Worth noticing for this scenario, is also the reduction in protection latency measured as the average number of channels traversed by the protection paths, which decreases from 5.87 to 4.88 hops. This in turn reduces the backup latency (restoration time) by 25.61% (calculated using the average length of the backup path in miles before and after the reoptimization).

Scenario	Backup port count			Avg. backup hops		Max backup hops	
Name	Before	After	% save	Before	After	Before	After
Dynegy Network	310	214	30.97%	5.87	4.88	15	10

Table 1. Partial Re-optimization

References

[1] G. Ellinas, E. Bouillet, R. Ramamurthy, J.F. Labourdette, S. Chaudhuri, K. Bala, "Routing and Restoration Architectures in Mesh Optical Networks", to appear in Optical Networks Magazine.

[2] E. Bouillet, J. Labourdette, R. Ramamurthy, S. Chaudhuri, "Enhanced Algorithm Cost Model to Control tradeoffs in Provisioning Shared Mesh Restored Lightpaths", in Proc. OFC 2002, Anaheim, CA, March 2002.

[3] R. Ramamurthy, J. Labourdette, S. Chaudhuri, R. Levy, P. Charalambous, C. Dennis, "Routing Lightpaths in Optical Mesh Networks with Express Links", in Proc. OFC 2002, Anaheim, CA, March 2002.