ure 1) and that there be restoration mechanisms to recover connections passing through a failed border node. Such restoration mechanisms are yet to be addressed, but must be designed to operate across a range of control domain control planes and transport technologies.

Transport layer connection restoration provides reliable links to higher layer services. However, failures involving higher layer devices, such as port failures and node failures, cannot be recovered using transport network restoration. Large IP networks today utilize IP layer protocols to re-route traffic around failed routers, interfaces and links. Transport layer failure recovery is often used, but is not necessary if IP layer mechanisms exist. Due to the unreliable nature of IP routers today, ISPs often use a dual router architecture, in which each central office (CO) contains two routers designed so that if one fails, traffic is automatically re-routed via the other. Studies have shown that utilizing transport layer failure recovery is significantly more expensive with a dual router architecture than using unprotected transport connections and allowing IP to re-route traffic upon failure [5]. This is due to the fact that little additional capacity is required to handle SRLG failures beyond that required to handle router failures. If routers were more reliable – eliminating the need for two routers at each CO - then optical layer restoration can be more cost effective than IP layer recovery. An integrated approach may provide the most cost effective solution [5]. For IP, at least, building a reliable router is an important requirement for transport layer restoration to become imperative. Thus routers need first to develop mechanisms for hitless software upgrades, non-stop routing, 1:N interface protection and so on.

4. Conclusions

Fast mesh restoration is critical in designing an intelligent transport network. However, there are numerous research, implementation and political challenges that remain to be addressed in developing cost-effective, secure and rapid mesh restoration.

5. References

[1] B. Cortez, "The Emerging Intelligent Optical Network: Now a Reality," *OFC* 2002.
[2] E. Mannie, editor, "Generalized Multi-Protocol Label Switching (GMPLS) Architecture," Internet Draft, 2002.
[3] P. Sebos *et al.,* "Effectiveness of Shared Risk Link Group Auto-Discovery in Optical Networks," *OFC* 2002.
[4] R. Doverspike and J. Yates, "Challenges for MPLS in Optical Network Restoration," *IEEE Commun. Mag.*, Feb. 2001.
[5] S. Phillips *et al.*, "Network Studies in IP/Optical Restoration," *OFC* 2002.

FQ2 **11:00 AM**

Local Optimization of Shared Backup Channels in Optical Mesh Networks

E. Bouillet, J. Labourdette, G. Ellinas, S. Chaudhuri, *Tellium, Oceanport, NJ, Email: ebouillet@tellium.com*.

This paper illustrates the complexity of assigning backup-channels in shared-mesh-protected optical networks. We propose a distributed recurring method to solve this problem, and show that substantial savings are achievable.

1. Introduction

Wavelength Division Multiplexed (WDM) networks that route optical connections using intelligent optical cross-connects (OXCs) is firmly established as the core constituent of next generation networks. With connection rates reaching tens of Gigabits/s, preventing and repairing failures is increasingly becoming an integral part of the network design process. In this work we consider two categories of end-to-end path restoration as supported in Tellium Aurora Optical Switch™ (see also [3,4]). Other categories include line protection and re-provisioning[7], these are not considered here. 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, or diverse, from the primary path. Path diversity guarantees that primary and backup paths will not simultaneously succumb to the same failure. This approach requires large amount of capacity, that is more than the working capacity since backup paths are longer than working paths. However 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 (1+1) mesh protection, this scheme allows considerable saving in terms of capacity required[3]. In addition, the backup resources can be utilized for lower priority preemptible traffic in normal network operating mode. However recovery is slower than dedicated (1+1) mesh protections, essentially because it involves signaling and path-setup procedures to establish the backup path.

They are two different policies to assign the protection channels: (1) pre-assign the protection channels to each backup-path before failure occurrence, or (2) rely on the protection mechanism to select the channels from a pool of reserved channels after failure occurrence[6]. Although more cost-efficient, approach (2) requires time-consuming inter-node communication to agree on the channel assignment, and to it we prefer approach (1), which can achieve sub-200ms restoration times in large networks[2]. However the gain in restoration time is paid for with additional complexity to fill up the backup-to-channel lookup tables at each node during provisioning (when speed is less of an issue.) We show in this paper that this operation is tantamount to a graph-coloring problem. In particular we show how a first fit based assignment can be easily improved using a graph-coloring algorithm. We finally discuss an application of this algorithm to migrate service protections from 1+1 mesh protection to shared mesh protection.

2. Shared Mesh Protection Provisioning using Vertex Coloring

We use the term Shared Risk Optical Group (SROG) to indicate a group of optical equipment that share a common risk of failure. Two mesh restored protection paths are "compatible" and may share a protection channel if their respective primary paths are SROG disjoint. Although only single SROG failures are considered here, the description of the algorithm can easily be transposed to protect against node failure as well: replace SROG by node where it applies. Otherwise they are said to be "conflicting". Given a group of protection paths traversing a common link, the problem is to assign the minimum number of protection channels to the paths in the link in accordance to the rules of sharing. Typical online provisioning algorithm assigns protection channels on a first-come first-serve basis and reserve new channels when sharing is not possible with present protection channels. In this approach the number of protection channels depends ultimately on the order of arrival of the protection

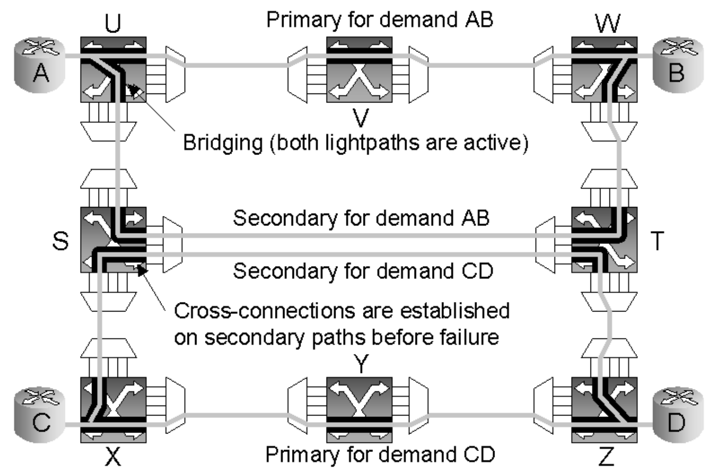

Figure 1. Dedicated (1+1) Mesh Protection

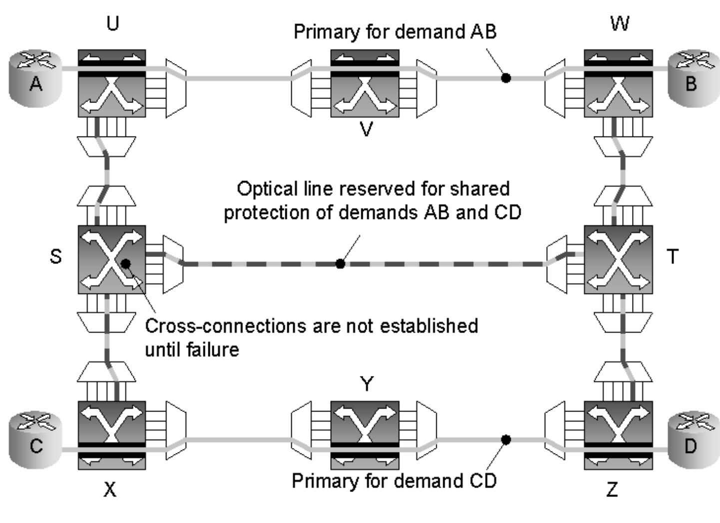

Figure 2. Shared Mesh Protection

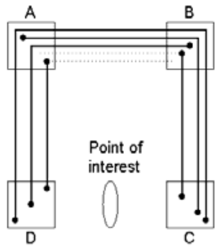

Fig (3a) Network consists of 4 nodes, 4 links and 5 lightpaths.

Order	Demand
1	AD
2	CD
3	BC
4	AC
5	BD

Table (3b) Order in which lightpaths are routed

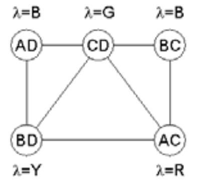

Fig (3c) Online coloring for link D-C requires 4 colors (B,G,R,Y)

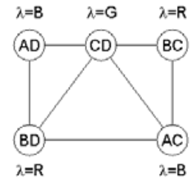

Fig (3d) Optimized coloring for link D-C requires 3 colors (B,G,R)

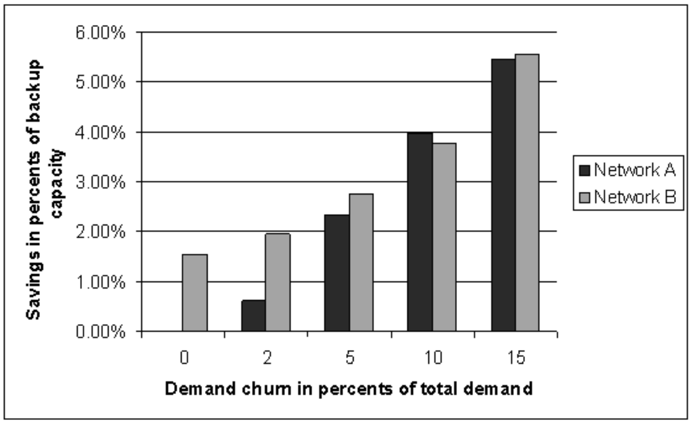

Figure 4.

paths. Since the order cannot be determined in advance, an optimization algorithm must be invoked at regular intervals to reassign the channels. In this write-up we show that finding the optimum assignment is equivalent to solving a vertex-coloring problem.

The allocation of protection channels is tantamount to a vertex-coloring problem: given the set of all restoration paths that intersect on a given link, represent every path as a vertex, and connect with an edge every pair of vertices whose corresponding paths are conflicting. Assign a distinctive color to each protection channel, and allot a protection channel to each path, that is color the vertices. Clearly, two vertices cannot be allotted the same color if they are connected by an edge, since the corresponding restoration paths are conflicting and cannot share a channel. The objective is to minimize the number of protection channels (respectively number of colors) required to accommodate all backup paths (respectively color all vertices), while avoiding conflicts.

This problem is known to be NP-hard, however there are many heuristics that can be used to compute sub-optimal solutions. A vertex-coloring algorithm that offers a good tradeoff between quality and runtime complexity is DSATUR[1].

Example

Consider the example of Figure 3 above (3a through 3d.) The figure illustrates five lightpaths {AD,CD,BC,AC,BD} and their protections, routed in a 4-node ring network. All the protections traverse link e_{CD}. The demands are provisioned following the sequence indicated in Table 3b. If we use a typical online shared mesh protection provisioning, and apply the graph representation presented earlier to e_{CD}, we obtain the "coloring" shown in Figure 3c. Even though a single failure in this example affects at most three primaries, this coloring consumes 4 colors, indicating that 4 protection channels are required. An optimized coloring yields the solution shown in Figure 3d, which consumes only 3 colors. Comparing 1c and 1d, we observe that a new channel (R) should have been allotted to the protection path of demand (BC) instead of sharing channel (B) with the protection of demand (AD). This solution however is not considered because not optimal when the third demand is being provisioned (that is {AD,CD,AC} are routed and {BD} has not yet arrived) since at that time it would consumes 3 channels instead of 2.

3. Implementation and Applications

The optimized channel reassignment is a low priority procedure. It can be a program thread running in background, or at regular intervals. The information necessary to accomplish this task is available locally in every OXC and independent of non-adjacent OXCs. Thus each OXC can run a copy of the algorithm in a distributed manner, locally and independently of other OXCs. A change in the allocation of a protection channel needs only to be propagated to its end-points. Since protection channels are "booked" and actually not cross-connected until a restoration occurs, the task amounts to no more than modifying and exchanging sharing databases between pairs of nodes. For every OXC-pair connected by at least one optical line, the OXC with highest IP address is delegated to perform the task.

A byproduct of the optimized channel reassignment is that it can be used to migrate the protection paths of mesh dedicated protections to shared mesh protections if desired. By changing their protection type to shared mesh protections, we allow the thread to apply the channel reassignment optimization to these services. The algorithm does not optimize the routes of the backup paths however, and the resulting solution is thus not as efficient as a re-optimization algorithm that re-routes the backup paths to maximize sharing[5].

4. Experiments

For our experiments we compare the benefits of local protection channel optimization on two realistic core mesh networks. Network A consists of shared-mesh capable optical switches in 46 cities interconnected by 75 fiber-trunks and loaded with 570 lightpaths. Network B consists of 61 switches, 88 fiber-trunks, and 419 lightpaths. For each network, we provision all the demands in sequence, using various values of demand churns (expressed in percent of the total demand), and perform a local channel optimization after all the demands are routed. We measure the amount of protection channel required before and after optimization and report the saving in % of total backup capacity in Figure 4. Our measurements indicate that as the demand churn increases, the number of protection channels that can be freed becomes substantial.

5. Conclusion

This document proposes a distributed method that rearranges the allocation of shared channels reserved for restoration, with objective to minimize the number of allotted channels. This algorithm can be implemented as an independent background process to supplement existing provisioning algorithms. It is effective to correct sub-optimality inherent to a first fit based provisioning, or seize on improvement opportunities that are brought forth by demand churn.

6. References

[1] Daniel Brélaz, "New Methods to Color the Vertices of a Graph," Communications of the ACM, Vol 22, Num 4. April 1979.
[2] A. Akyamac, S. Biswas, J.F. Labourdette, S. Chaudhuri, K. Bala, "Ring Speed Restoration and Optical Core Mesh Networks," NOC'02, Darmstadt Germany, Jun. 2002.
[3] J.F Labourdette, E. Bouillet, R. Ramamurthy, G. Ellinas, S. Chaudhuri, K. Bala, "Routing Strategies for Capacity Efficient and Fast Restorable Mesh Optical Networks," Photonic Network Communications, pp. 219-235, July-Dec 2002.
[4] G. Ellinas, et al., "Routing and Restoration Architectures in Mesh Optical Networks," Optical Networks Magazine, Jan-Feb 2003.
[5] E. Bouillet, P. Mishra, J.F. Labourdette, K. Perlove, S. French, "Lightpath Re-optimization in Mesh Optical Networks," NOC'02, Darsmstadt Germany, Jun. 2002.
[6] S. Datta, et al, "Efficient Channel Reservation for Backup Paths in Optical Mesh Networks," IEEE GLOBECOM 2001, San Antonio, TX, Nov. 2001.
[7] J. Doucette, W.D. Grover, "Comparison of Mesh Protection and Restoration Schemes, and the Dependency on Graph Connectivity," 3rd International Workshop on the Design of Reliable Communication Networks, Budapest, Hungary, Oct. 2001.

FQ3 **11:15 AM**

Capacity Requirements for Network Recovery from Node Failure with Dynamic Path Restoration

G. Shen, W. Grover, *TRLabs and Department of Electrical and Computer Engineering, University of Alberta, Edmonton, AB, Canada, Email: gshen@trlabs.ca*.

Node failure is not as frequent as span failure but recent events have emphasized its importance in network planning. We study the effects on capacity design if full or partial recovery from node failures is provided using failure-specific path restoration.

1. Introduction

Most studies of restorable networking consider span failures as the primary class of failure scenario. It is, however, often noted that because of its end-to-end orientation, a path restoration mechanism has an inherent ability also to respond to node failures. The spare capacity that ensures 100% span restorability is not necessarily adequate to ensure any particular target level of recovery from a node failure, however. Pre-planned shared backup path protection (SBPP) [1] does inherently protect transiting flows against node loss if primary and backup paths are all node disjoint. But SBPP also generally requires more spare capacity than dynamic path restoration and, due to its fixed pre-planned nature, has an inherently lower availability against to dual failure scenarios. It is of interest, therefore, to consider how much extra spare capacity an adaptive path restorable network needs to support node recovery, beyond that needed for span restorability. Other studies [2, 3, 4] have considered node recovery issues but to our knowledge the specific questions we ask, and the particular mechanism [6] and capacity design model [5] we consider are novel