Ring Speed Restoration and Optical Core Mesh Networks

Ahmet A. Akyamaç, Subir Biswas, Jean Labourdette, Sid Chaudhuri, Krishna Bala
Tellium Inc., 2 Crescent Place, Oceanport, NJ 07757 USA

Abstract

Optical mesh networks using intelligent optical cross-connects are the choice for the next generation optical core. Mesh networks can be designed for fast and guaranteed recovery from single failures and can be designed to protect traffic more efficiently via shared mesh restoration. In this paper, we present a mesh restoration architecture for the optical core. The architecture provides for all the advantages of mesh networks and achieves ring-like restoration performance. Furthermore, the protocol allows for two types of restoration depending on the application needs: dedicated mesh protection that guarantees sub-50 ms. restoration latency and shared mesh restoration that can restore services within 50 ms. to a couple hundred ms., which is sufficient for the majority of the voice-band and other streaming services.

1. Introduction

While SONET rings offer sub-50 ms. restoration, rings are not the best answer for building point-to-point connections over long distances. It has already been accepted in the industry that a mesh topology using Optical Cross-connects (OXC's) can provide a scalable and capacity-efficient solution for designing next generation optical networks [1]. Mesh networks provide many advantages such as significant bandwidth and cost savings, dynamic provisioning in the presence of unpredictable traffic, fast setups and interconnections, and smaller footprint [2]. Contrary to rings, the capacity of a mesh network can be increased on a link by link basis, making the network more flexible. Furthermore, mesh networks can be designed for fast and guaranteed recovery from single failures and can be designed to protect traffic more efficiently [3].

In this paper, we present a mesh restoration architecture for the optical core. The architecture provides for all the advantages of mesh networks and achieves ring-like restoration performance. Furthermore, the protocol allows for two types of restoration depending on the application needs: dedicated mesh protection that guarantees sub-50 ms. restoration latency and shared mesh restoration that can restore services within 50 ms. to a couple hundred ms., which is sufficient for the majority of the voice-band and other streaming services. We also present measurements of dedicated and shared mesh restoration times through experiments conducted in lab and field environments and software based protocol simulation. The measurements demonstrate ring-like restoration performance.

This paper is organized as follows: In Section 2, we present a description of the dedicated mesh protection and shared mesh restoration architectures. In Section 3, we present measurements from experiments in lab and field environments and from our simulation software. We conclude in Section 4.

2. Network Model and Restoration Architecture

We consider a network model that consists of optical cross-connect (OXC) switches connected by fibers. The fibers contain multiple optical channels (wavelengths) that carry lightpaths. The lightpaths carry end-to-end traffic between switches and are restorable against link or node failures. We assume the restoration process is performed at the OXC level, which is comprised of Tellium's Aurora Optical Switches (AOS). The AOS can be used to construct mesh topologies and its StarNet operating system offers dedicated mesh protection and shared mesh restoration protocols [4]. In the following, we describe the concepts of shared risk groups and path diversity, followed by the dedicated mesh protection and shared mesh restoration protocols.

2.1 Shared Risk Groups

Fibers are carried through cables, which in turn pass through conduits. A conduit may contain multiple cables, which in turn may contain multiple fibers carrying traffic for different source and destination switches. Failures of multiple optical channels are usually due to fiber, cable or conduit cuts. Consider the 6-node optical network of Fig. 1. Each cylinder in the figure represents a conduit. Optical channels across the two links connecting two distinct pairs of nodes traverse the same conduit. If the conduit and the fibers it contains are accidentally severed all the optical channels inside the conduit fail. The concept of shared risk optical group (SRG) expresses the risk relationship that associates all the optical channels with a single failure [5]. An SRG may consist of all the optical channels in a single fiber, of the optical channels through all the fibers wrapped in the same cable, or of all the optical channels traversing the same conduit. A fiber may traverse multiple conduits; hence channels may be associated with more than one SRG. If an SRG fails, all channels in that SRG fail.

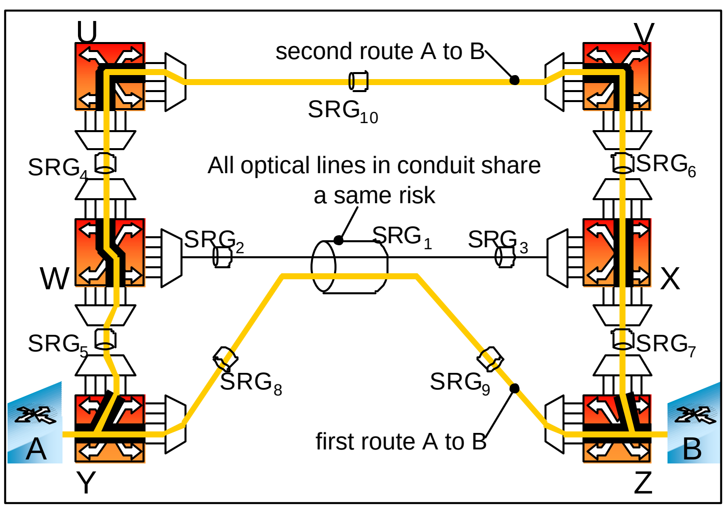

Figure 1: Shared risk optical groups (SRG's)

2.2 Node or SRG Failure

Resilience against failures is achieved by (primary) lightpaths that are protected by diverse (or disjoint) backup lightpaths. We consider two types of protection: (1) SRG failure resilient, and (2) node failure resilient. Resilience against SRG failures is achieved by way of path diversity, as shown in Fig. 2(a), where the paths are SRG disjoint. Such a scheme guarantees against a single SRG failure. For protection against node failures, it is necessary to provision routes that are SRG-and-node disjoint, as shown in the example of Fig. 2(b). However, node-diverse paths consume more resources than the less conservative SRG-diverse scheme pictured in Fig. 2(a) [6].

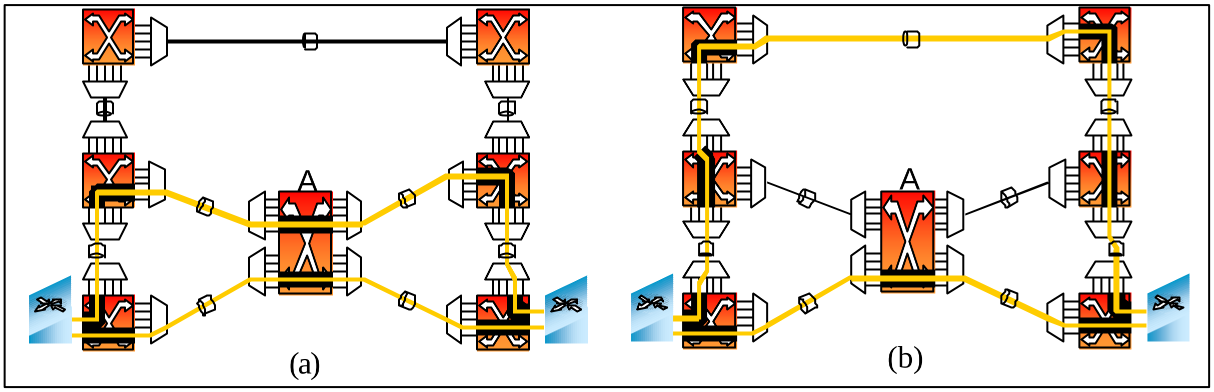

Figure 2: (a) SRG disjoint paths **(b)** SRG-and-node disjoint paths

2.3 Dedicated Mesh Protection

Dedicated mesh protection provides a fast and guaranteed UPSR/path switching ring-like restoration protocol over a mesh, as illustrated in Fig. 3. The network consists of four switches (A to D) and two lightpath demands (AB and CD) routed across an eight node optical network (S to Z). The primary and backup paths for each lightpath are SRG disjoint (for link disjoint routing) or SRG-and-node disjoint (for node disjoint routing). This path diversity guarantees that primary and backup paths will not be simultaneously affected by the same failure. During normal operation, both paths carry the optical signal and the egress node selects one of the two copies. The backup path is pre-allocated (remains *live*), thus saving crucial path-setup latency during restoration (for dedicated mesh protection, this is equivalent to switching to the alternate copy). This is the fastest restoration scheme since for every lightpath one device is responsible for all the necessary failure detection and restoration functions. This protocol guarantees sub-50 ms. restoration latency and is very well suited for applications with extremely low latency tolerance. But it is also the most capacity-intensive since the protocol uses full protection capacity redundancy [6].

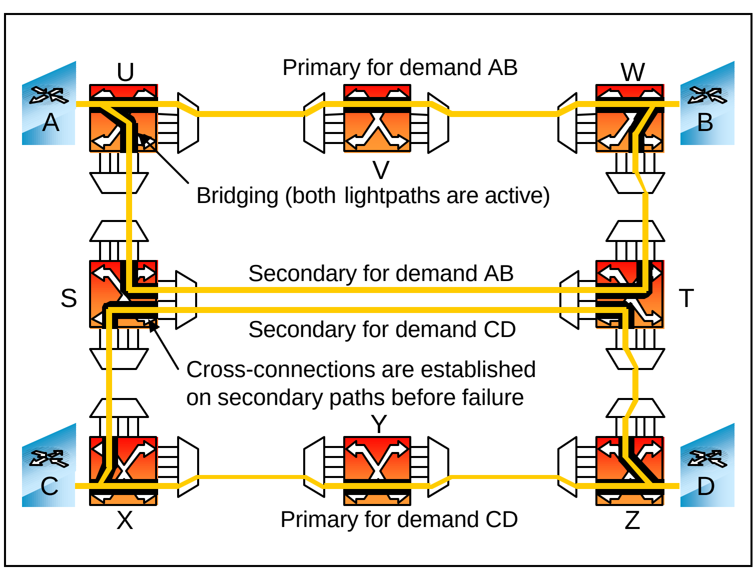

Figure 3. Dedicated Mesh Protection

2.4 Shared Mesh Restoration

Shared mesh restoration provides a capacity-shared BLSR/MSSPRING ring-like restoration protocol in which pre-computed backup paths for multiple primary paths can share protection capacity and therefore, can reduce the restoration cost. In this protocol, backup paths are pre-defined but the cross-connections along these paths are not created until a failure occurs. Shared mesh restoration is illustrated in Fig. 4.

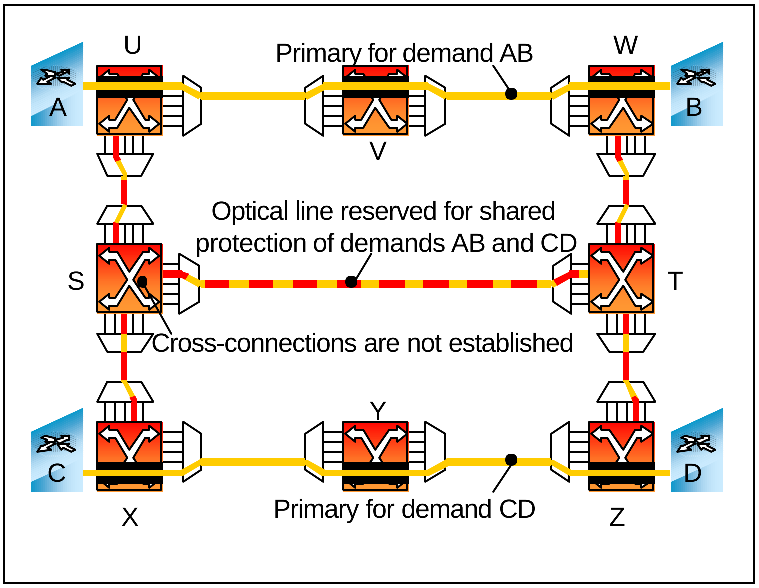

Figure 4. Shared Mesh Restoration

In shared mesh restoration the backup paths can share capacity if the corresponding primary paths are mutually diverse (SRG diverse or SRG-and-node diverse). The backup path is reserved (but not *live* since multiple lightpaths can be sharing it). Hence, recovery may be slower than dedicated mesh protection since it involves signaling and path setup to establish the cross-connections on the backup path during restoration. Compared to dedicated mesh protection, this scheme allows considerable saving in terms of capacity required [7]. Although sharing reduces the restoration speed, the protocol can restore services within 50 ms. to a couple hundred ms., which is sufficient for the majority of the voice-band and other streaming services.

3. Restoration Performance

Tellium has conducted several studies to measure dedicated and shared mesh restoration times through experiments in lab and field environments and software based protocol simulation. In the following we report restoration latencies observed in lab experiments and those estimated through simulation. Data from field experiments are proprietary in nature and hence are not reported here.

3.1 Dedicated Mesh Protection

One class of lab experiments for dedicated mesh protection was conducted using a network of four AOS switches in test labs on site. For this setup, 16 parallel OC-48 (STM-16) lightpaths with two hop primary paths and two hop backup paths were set up (similar to a four node ring configuration). Experiments were conducted in which a link on the primary path was failed whereby all the lightpaths were simultaneously failed. This experiment was repeated 50 times and a sample protection latency was measured for each experiment. Fig. 5 shows the distribution of the restoration latencies observed. As seen in Fig. 5, the latency for dedicated mesh was found to be consistently less than 50ms and typically between 30-40ms.

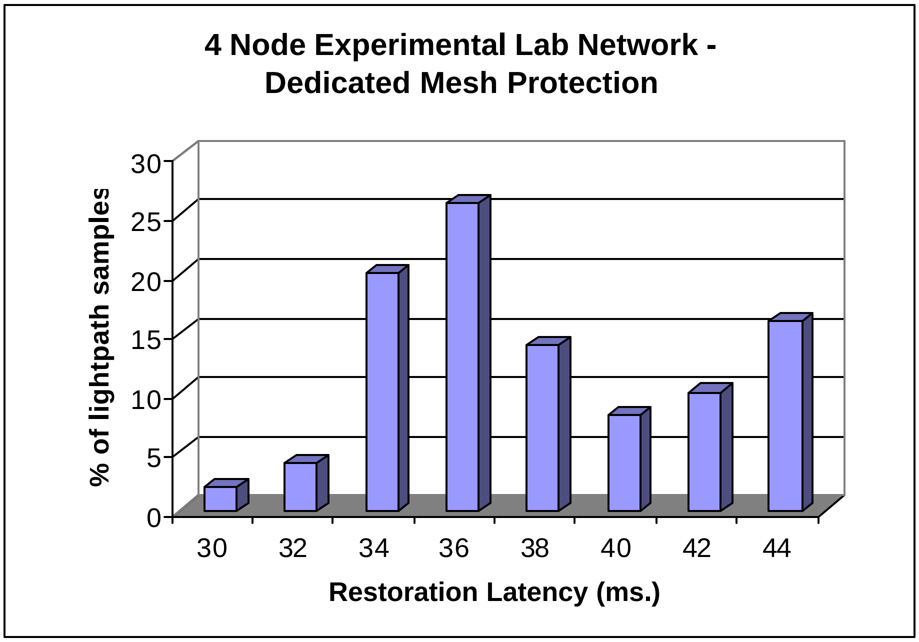

Figure 5. Dedicated mesh restoration latency in 4-node experimental lab network

3.2 Shared Mesh Restoration

For shared mesh restoration, one class of lab experiments was conducted using a network of four AOS switches in test labs on site. Similar to the previous setup, parallel OC-48 (STM-16) lightpaths with two hop primary paths and two hop backup paths were used (similar to a four node ring configuration). Experiments were conducted in which links were failed whereby all the lightpaths were simultaneously failed. The studies involved varying the number of lightpaths failed and observing the final restoration latencies. This experiment represents the worst case scenario in that all lightpaths follow the same route and the restoration process is performed for all failed lightpaths by the same set of end nodes. This provides results that are worse than typical mesh networks in which the end nodes of the failed lightpaths would typically be distributed across the network. Fig. 6 shows the maximum observed shared restoration latency versus the number of lightpaths simultaneously failed.

A second class of studies was performed for significantly larger networks, which are not feasible to implement in the lab due to space and equipment constraints. These studies were conducted using StarNet Modeler [8], which is a mesh network restoration modeling, simulation and visualization tool developed by Tellium. StarNet Modeler incorporates the same route computation module used in the AOS network and contains detailed state machines to model the StarNet mesh restoration protocols. It is calibrated using measurements from the AOS and experimental results from the lab testbed. StarNet Modeler is used by Tellium and its customers to model what-if scenarios involving failures and repairs to estimate expected restoration performance. In the following, we present sample results for two such representative studies.

The first study was for a mesh reference network of 17 nodes with 224 OC-48 (STM-16) services using SRG disjoint routing. This network represents a hypothetical N. American carrier network with an average degree of 3.1. In Fig. 7, we show the restoration latency

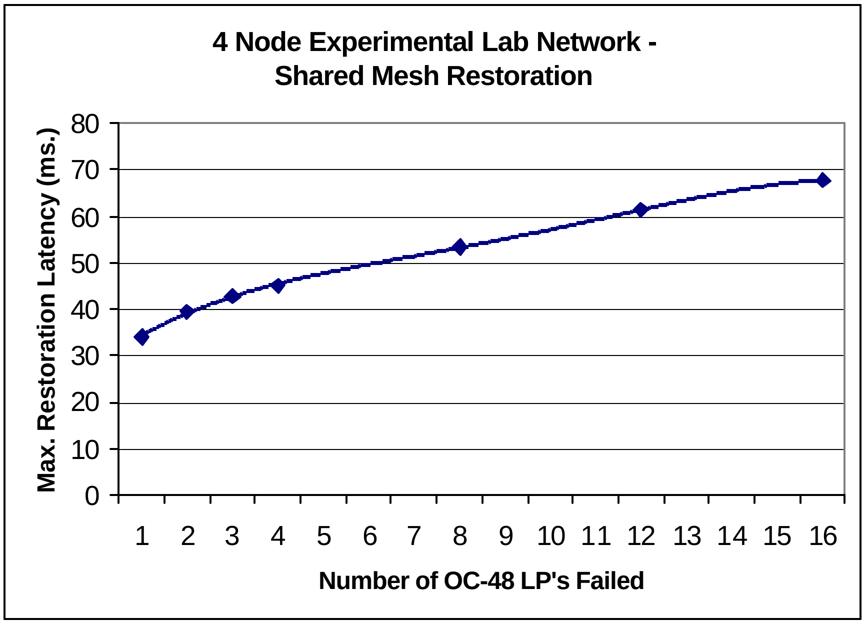

Figure 5. Shared mesh restoration latency in 4-node experimental lab network

distribution following a link failure simultaneously affecting 35 OC-48 (STM-16) lightpaths. As can be seen from Fig. 7, the majority of the lightpaths were restored within 80-140 ms. and the overall restoration process was completed at 155 ms.

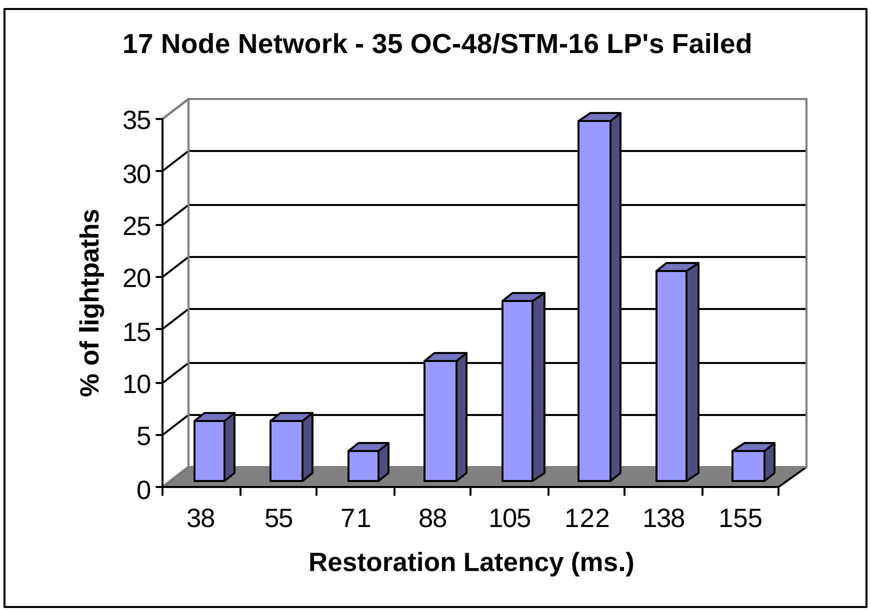

Figure 7. Shared mesh restoration latency distribution for 35 failed lightpaths

The second study was for a mesh reference network of 50 nodes with 910 OC-48 (STM-16) services using SRG disjoint routing. This network represents a hypothetical pan-European network with an average degree of 3.44. Fig. 8 shows the restoration latency distribution following a link failure simultaneously affecting 56 OC-48 (STM-16) lightpaths. Not all end points of these lightpaths are the same and are distributed across the network. As can be seen from Fig. 4, the majority of the lightpaths were restored within approximately 80-150 ms., well below 200 ms., with the overall restoration process being completed at 161 ms. These results are typical of what we have observed using StarNet Modeler for customer networks.

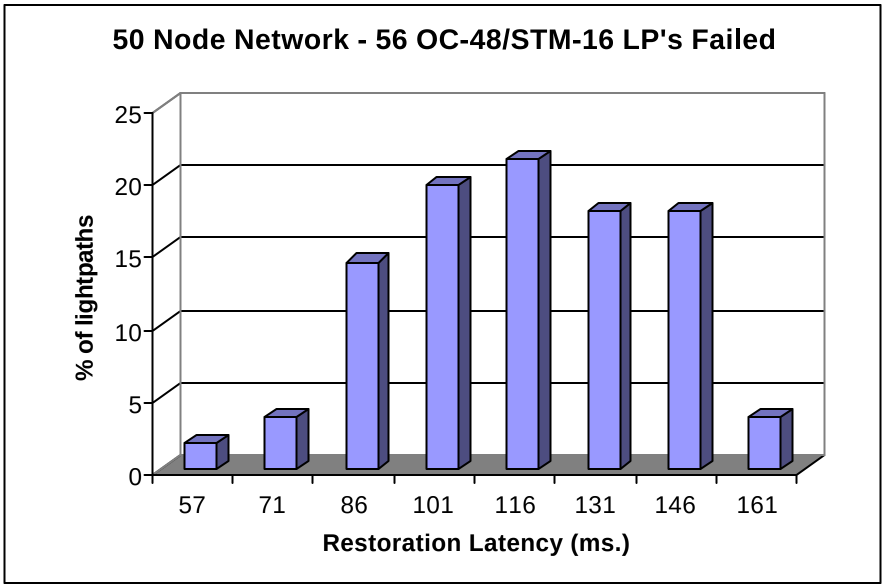

Figure 8. Shared mesh restoration latency distribution for 56 failed lightpaths

4. Conclusion

In this paper, we presented a mesh restoration architecture for the optical core that provides for all the advantages of mesh networks while achieving ring-like restoration performance. The restoration protocol allows for two types of restoration depending on the application needs: dedicated mesh protection that guarantees sub-50 ms. restoration latency and shared mesh restoration that can restore services within 50 ms. to a couple hundred ms., which is sufficient for the majority of the voice-band and other streaming services. We also presented measurements of dedicated and shared mesh restoration times through experiments conducted in lab and field environments and software based protocol simulation, which demonstrate ring-like restoration performance.

Contact E-mail: aakyamac@tellium.com, Tel: +1(732) 483-3073, Fax: +1(732) 483-3357

5. References

[1] J. Schallenberg, "Is 50ms Restoration Necessary?", proceedings of the *IEEE Bandwidth Management Workshop IX*, Dec. 2001.
[2] S. Samieian, Z. Bogdanowicz, R. Ramamurthy, S. Chaudhuri, K. Bala, "Capacity Requirements of Ring, 1+1 and Mesh Configurations", in *Lightwave Magazine*, Aug. 2001.
[3] R. Ramamurthy, Z. Bogdanowicz, S. Samieian, D. Saha, B. Rajagopalan, S. Sengupta, S. Chaudhuri, and K. Bala, "Capacity Performance of Dynamic Provisioning in Optical Networks", in *IEEE Journal of Lightwave Technology*, vol. 19, issue 1, pp. 40-48, Jan. 2001.
[4] "StarNet Operating System: Features and Benefits", *http://www.tellium.com/products/starNetOS.html*.
[5] S. Sengupta, R. Ramamurthy, "From Network Design to Dynamic Provisioning and Restoration in Optical Cross-Connect Mesh Networks: An Architectural and Algorithmic Overview". in *IEEE Network*, vol. 15, issue 4, Jul./Aug. 2001, pp. 46-54.
[6] G. Ellinas, E. Bouillet, R. Ramamurthy, J. Labourdette, S. Chaudhuri, K. Bala, "Routing and Restoration Architectures in Mesh Optical Networks", to appear in *Optical Networks Magazine*.
[7] J. Labourdette, E. Bouillet, R. Ramamurthy, G. Ellinas, S. Chaudhuri, K. Bala, "Routing Strategies for Capacity-Efficient and Fast-Restorable Mesh Optical Networks", to appear in *Photonic Network Communications* special issue on *Routing, Protection and Restoration Strategies and Algorithms for WDM Optical Networks*.
[8] S. Biswas, S. Chaudhuri, S. Datta, K. Hua, S. Sengupta, "StarNet Modeler: A Hybrid Modeling System for Path Restoration in Optical Networks", in *OPNETWORK 2001*, Wash. DC, Aug. 2001.