Restoration in Layered Architectures with a WDM Mesh Optical Layer

Georgios Ellinas, Eric Bouillet, Ramu Ramamurthy, Jean-François Labourdette, Sid Chaudhuri, Krishna Bala

Tellium Inc.
2 Crescent Place
Oceanport, NJ 07757
{gellinas, ebouillet, ramu, jlabourdette, sc, kbala@tellium.com }

Abstract

This work provides an overview and a comparison of various techniques used to restore lightpaths in a layered architecture. The comparison takes into consideration figures of merit such as speed of restoration and redundant capacity. The objective is to optimize the performance of the network (including fast end-to-end service survivability, minimizing the spare resources required) under all circumstances, using available resources and implemented mechanisms.

1 Introduction

It is widely accepted that optical cross-connects (OXCs) will be used to implement the next generation mesh optical networks [1]. Optical network equipment vendors are currently implementing the next generation of optical switching systems capable of switching hundreds of lightpaths, each carrying millions of voice calls, or thousands of video streams. Optical network architectures not only provide transmission capacities to higher transport levels, but also the intelligence required for fast lightpath provisioning and fast and efficient failure restoration [2,3,4]. The emergence of intelligent optical network elements are instrumental in making such optical architectures a reality today.

Optical network architectures can suffer failures and breakdowns, either due to accidental fiber cuts and operation mistakes, or equipment malfunctions such as switch, card or component failures. Two competing approaches are being proposed for providing the appropriate recovery mechanisms in these circumstances. In the *peer-to-peer* approach [5,6,7], interweaved optical and higher layer equipment act in symbiosis under the same control plane. In the *overlay* approach [5,6], optical and higher layer domains are two separate entities with individual control planes, exchanging management services through a standard interface. The peer-to-peer approach relies on a unified bandwidth management protocol to reassign bandwidth away from defective areas in the network and reestablish the interrupted data services. In the overlay approach, each layer independently relies on its own restoration mechanism in a manner that is independent and transparent to one another.

A single network layer or a combination of multiple layers can be used for failure restoration in a layered architecture. The aim is to provide service protection against a variety of failure conditions while restoring all failures quickly and with the minimum amount of protection capacity. Failure restoration involving multiple layers can enhance end-to-end survivability of the service by having each layer's protection scheme supplement

each other. It is important to note however, that multilayer protection may not be required or may be difficult to implement because of race conditions and complex escalation strategies and interlayer protocols.

Section 2 presents the layered network architecture and describes each layer. Section 3 examines how restoration can be addressed in such architecture and Section 4 addresses restoration specifically in the optical layer. Section 5 discusses whether multilayer restoration is necessary, and addresses possible escalation approaches in a network architecture where multiple layers can be used to restore a service. Concluding remarks follow in Section 6.

2 Overview of a Layered Network Architecture

In this section we review the fundamental parts that constitute a network and its functionality. It goes without saying that many architectures exist or have been suggested and it is not the intention of this paper to enumerate them exhaustively (see [5,8,9,10] for further information and useful references on this topic). However we observe that all the proposed architectures repose on a common denominator. It is this generic model that we present here. The model consists of three superimposed layers. Each layer provides well-defined services to its superjacent layer while concealing implementation details to it. As shown in Figure 1, from top to bottom the layers are (1) Service layer, (2) Logical (Electrical) layer and (3) Optical layer.

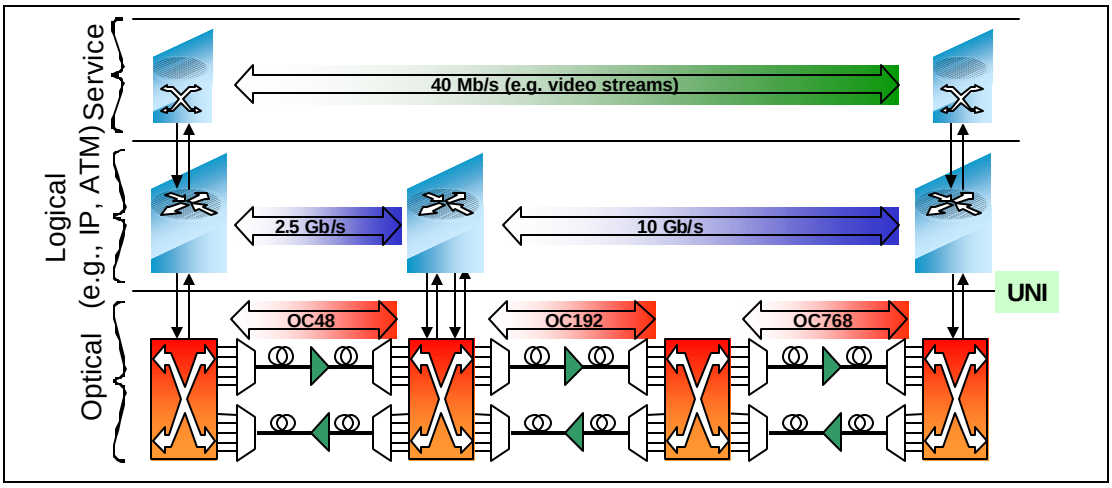

Figure 1. LayeredArchitecture

(a) Service Layer

In the service layer, clients such as edge or service routers or MSPPs located in a provider Point-of-Presence (PoP) represent users and the data communication among them. Using a graph representation, a node corresponds to a client who emits and receives data, and a link represents a service or a two-way data stream between clients. Link attributes in this layer correspond to minimum QoS requirements, which transpose into bit

rates, jitter, and bit-propagation or round-trip delay constraints. Service Level Agreements (SLA) for instance are negotiated and crafted in this layer.

(b) Logical Layer

Also known as electrical or digital layer, the logical layer aggregates services into large transmission "pipes" and assures their proper routing from PoP to PoP with prescribed QoS. Using a graph representation, a logical node corresponds, for example, to an IP core router, an ATM backbone switch or a digital cross-connect (DCS), and a logical link connects the ports of two adjacent nodes. Capacity of a link, and data processed by the nodes are expressed in units of bits per second (b/s) in increments of DS-3 (45 Mb/s) to OC-192 (10 Gb/s). The logical layer may consist of several interconnected sub-networks, either for scalability reasons, as it is easier to manage several smaller networks than a large network (hierarchical decomposition), or because the sub-networks belong to several independent carriers (multi-vendors, two-tier networking) or employ different technologies (e.g., IP versus ATM). In either case, boundaries and proper network interfaces within the logical layer delimit the sub-networks and their respective domains of operation.

The logical layer fulfils several roles: (1) it maintains a consistent topological view of the layer, (2) it manages the address space, (3) it routes streams on request, and (4) it polices the traffic to ensure a fair share of capacity among data streams and to guarantee each individual's QoS. The first part, also called topology discovery, is achieved by way of a Neighbor Discovery Protocol (NDP) in conjunction with the Open Shortest Path First (OSPF) protocol[1]. NDP operates in a distributed manner through in-band signaling to construct local port-to-port connectivity databases at each node. OSPF completes the topology discovery by assembling and globally disseminating pieces of information collected by NDP, plus additional information such as link states, to the logical plane. Logical nodes have only a few tens of ports, and with the exception of very small networks a full connectivity featuring one link between every pair of node is not probable. Instead, services may have to be routed in the logical layer through one or more transit nodes to the desired destination using CR-LDP/RSVP explicit routing and bandwidth reservation protocols. The computation of a logical path must satisfy a set of constraints, such as round-trip delays and spare bandwidth, defined in the service layer in accordance to prescribed QoS. Note that the failure of a logical link or logical node is detected by NDP, and advertised by OSPF. That is, the layer has the primitives to detect a failure and resume interrupted services.

(c) Optical Layer

The optical layer offers and manages the capacity required to transport traffic between clients in the logical layer. Figure 2 depicts an example of a logical network (two IP routers) linked to an optical network (four optical switches). Optical switch ports are either: (1) add/drop-ports, interfacing the optical layer to the client's logical layer, or (2) network ports, interconnecting optical switches. Using our graph representation, nodes are

optical switches, and links are bundles of bi-directional optical channels between pairs of optical switches. An optical channel is a wavelength that connects the network ports of adjacent optical switches. A link in the logical layer is realized by way of optical channels in tandem forming a lightpath (circuit) between the end-nodes of that link.

The optical layer faces the same challenges, and conceptually even borrows solutions from the logical layer. For instance, it relies on Generalized MPLS (GMPLS) [11-13], also formerly known as MPLambdaS (an extension of MPLS) to encompass all types of architectures, including wavelength-oriented traffic engineering and management. It also relies on Neighbor Discovery Protocol (NDP)/Link Management Protocol (LMP) [6,7,14] and Open Shortest Path First (OSPF) protocol [15] to create and publicize the network's topological views. Differences that set apart the optical layer from its logical counterpart are among others: (1) routing in the optical layer is exclusively circuit oriented, (2) circuit set-up and tear-down is done at a much slower time scale and (3) the bandwidth granularity of the logical layer is much lower than the granularity of the optical layer.
In the overlay approach the layers work individually, with the client logical layer leasing resources from the optical layer. The User Network Interface (UNI) harmonizes communication of control messages between the two domains. The Optical Interworking Forum (OIF) is currently specifying UNI requirements and is working on their standardization [16]. In addition, an optical carrier will normally acquire network components from several vendors. A suite of protocols is being developed in the Internet Engineering Task Force (IETF) to allow for the seamless interaction between the various network components. As part of the suite, the Link Management Protocol (LMP), for example, is used to maintain control channel connectivity, verify component link connectivity and isolate link, fiber or channel failures within the network [17].

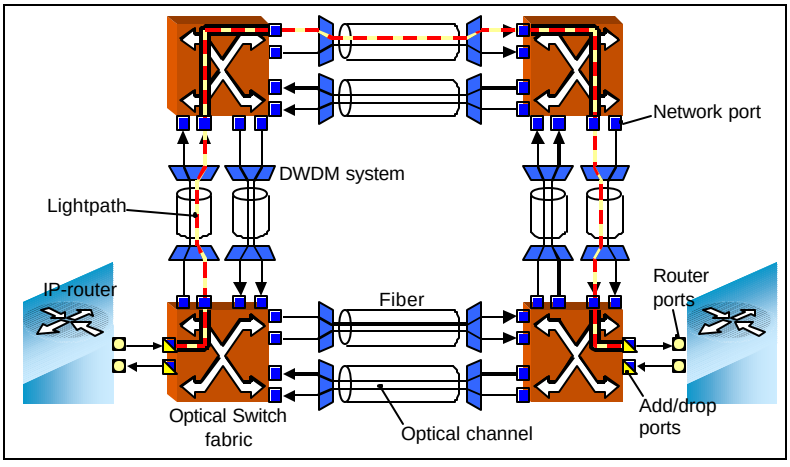

Figure 2. Optical layer

[1] In this paper the assumption is of a single network running IP-centric Multi Protocol Label Switching (MPLS) protocols.

3 Failure Restoration in a Layered Architecture

Restoration is invoked upon failure of one or more elements along paths carrying end-to-end services. Both logical and optical layers may implement an autonomous recovery scheme, and both may react to the same defect. Failures are either one of two types: (1) logical, such as a malfunctioning IP-router, or (2) optical, for instance a fiber cut. There are thus four possible scenarios, depending on the origin of the failure, and the layer that provides the restoration [18]:

1. Failure and restoration in the optical layer as shown in Figure 3(b) based on the original routing shown in Figure 3(a). Figure 3(a) illustrates the connectivity in the optical layer and the resulting connectivity in the logical layer during normal mode of operation. Figure 3(b) is an example of optical failure restored in the optical layer. The affected lightpath is restored away from the failure using optical capacity that was reserved for this purpose. The operation is transparent from the logical layer, which remains unchanged.
2. Failure and restoration in the logical layer as shown in Figure 3(c). This figure illustrates a logical failure (ATM switch or IP router failure) restored in the logical layer. After failure the service is re-routed using the remaining capacity of the logical layer. The operation is transparent to the optical layer.
3. Optical failure repaired in the logical layer as shown in Figure 3(d). This figure illustrates an optical failure restored in the logical layer. If the optical layer fails to restore the optical failure after a certain time-lapse, the logical layer can restore the service on a different logical path, using for instance implicit LSP protection in MPLS.
4. Logical failure repaired in the optical layer as shown in Figure 3(e). Unlike any of the previous protection schemes, restoring a logical failure with leverage from the optical layer involves reconfiguration with creation of new connections in the logical layer and the optical layer. This type of restoration may be necessary if after logical failure the remaining capacity in the logical layer is insufficient to re-route all the affected services. Additional logical capacity can be created with the provisioning of new lightpaths.

A fifth and most realistic situation consists of optical failures repaired simultaneously and independently in both layers. Since this is a combination of scenarios mentioned above, it is not considered here.

Restoration in an IP-centric logical layer is accomplished by Multi-Protocol Label Switching (MPLS) [19]. MPLS enables a hierarchy of Label Switched Paths (LSPs) to be defined by pre-pending a stack of labels or tags to packet headers. Upon an optical or electrical failure occurrence, packets along a given disrupted LSP can be routed to a predefined restoration LSP by modifying the label maps of the routers at the end-points of the original LSP [20]. In a similar manner, restoration of optical failures in the optical layer is also achieved by way of redundancy. Studies indicate that restoration in the optical layer requires substantially more spare capacity, depending on the diligence and the quality of the protection, yet overall the solution is more economical due to lower cost per units of capacity [20].

MPLS offers undeniable potentials for fast restoration. The principal advantage of MPLS is its ability to recover indiscriminately from failures in the logical layer or the optical layer as suggested in Figure 3(c) and Figure 3(d). However, a single failure may affect thousands of LSPs, and trigger an avalanche of alarms and corrective actions. The resulting amount of signaling can be orders of magnitude higher than in the optical layer, which is able to switch hundreds of LSPs multiplexed into a single wavelength at once. Also in MPLS restoration, primary and backup LSPs must not succumb together to a malfunction in the logical or in the optical layer. In order to satisfy the second condition, the logical layer must explicitly inquire about the risk relationship between the lightpaths that compose its logical connectivity and compute the LSPs, primaries and respective backups, accordingly. Proposed specifications for the UNI interface allow the logical layer to request lightpaths that are disjoint from selected subsets of pre-established lightpaths. However, this approach yields lower availability than other approaches that allow the optical layer to decide on the restoration mechanisms. Such requests may thus sometime be impossible to realize even if the capacity is available in the optical layer. Another strategy is to rely on NDP, OSPF and IP self-routing properties to advertise and correct failures in the logical configuration, but then the restoration time is not as attractive in terms of restoration speed as it would be with predefined restoration LSPs.

The fourth scenario implies a minimum of synergy between the restoration architectures deployed in each layer; the optical layer does not know a-priori the logical connectivity of the client and hence cannot take the initiative to restore a logical failure. Both layers however could coordinate their effort to resume interrupted services, with the optical layer getting directives from the logical layer. In particular, the logical layer could provision spare capacity in the optical domain and reclaim some of it upon failure of one of the routers in order to create new logical connections and balance the load on the surviving routers. The feasibility of this scheme is subject to UNI specifications.

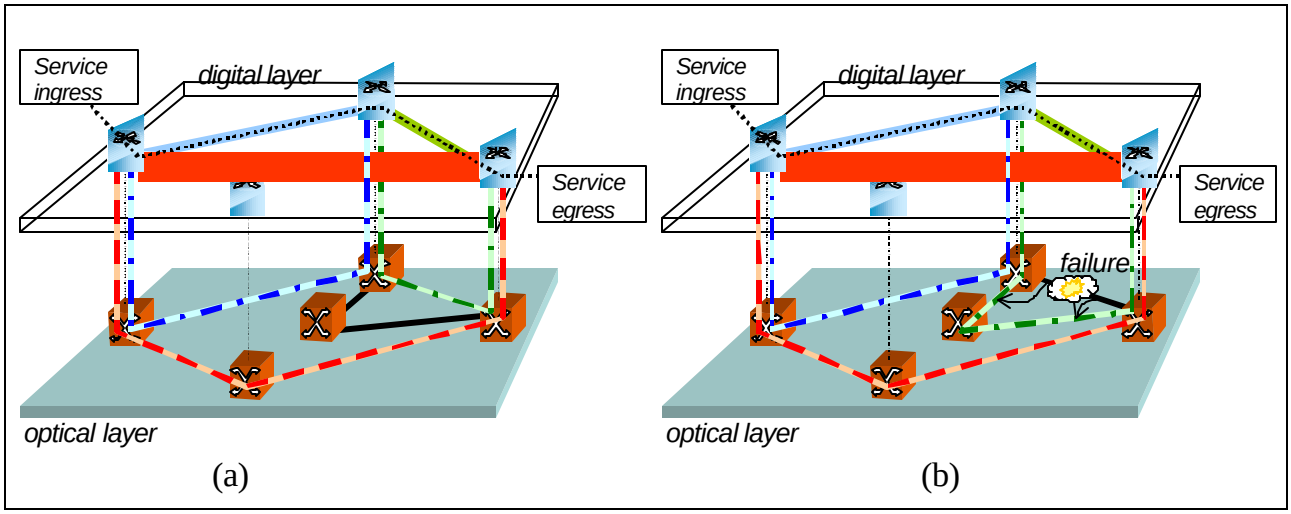

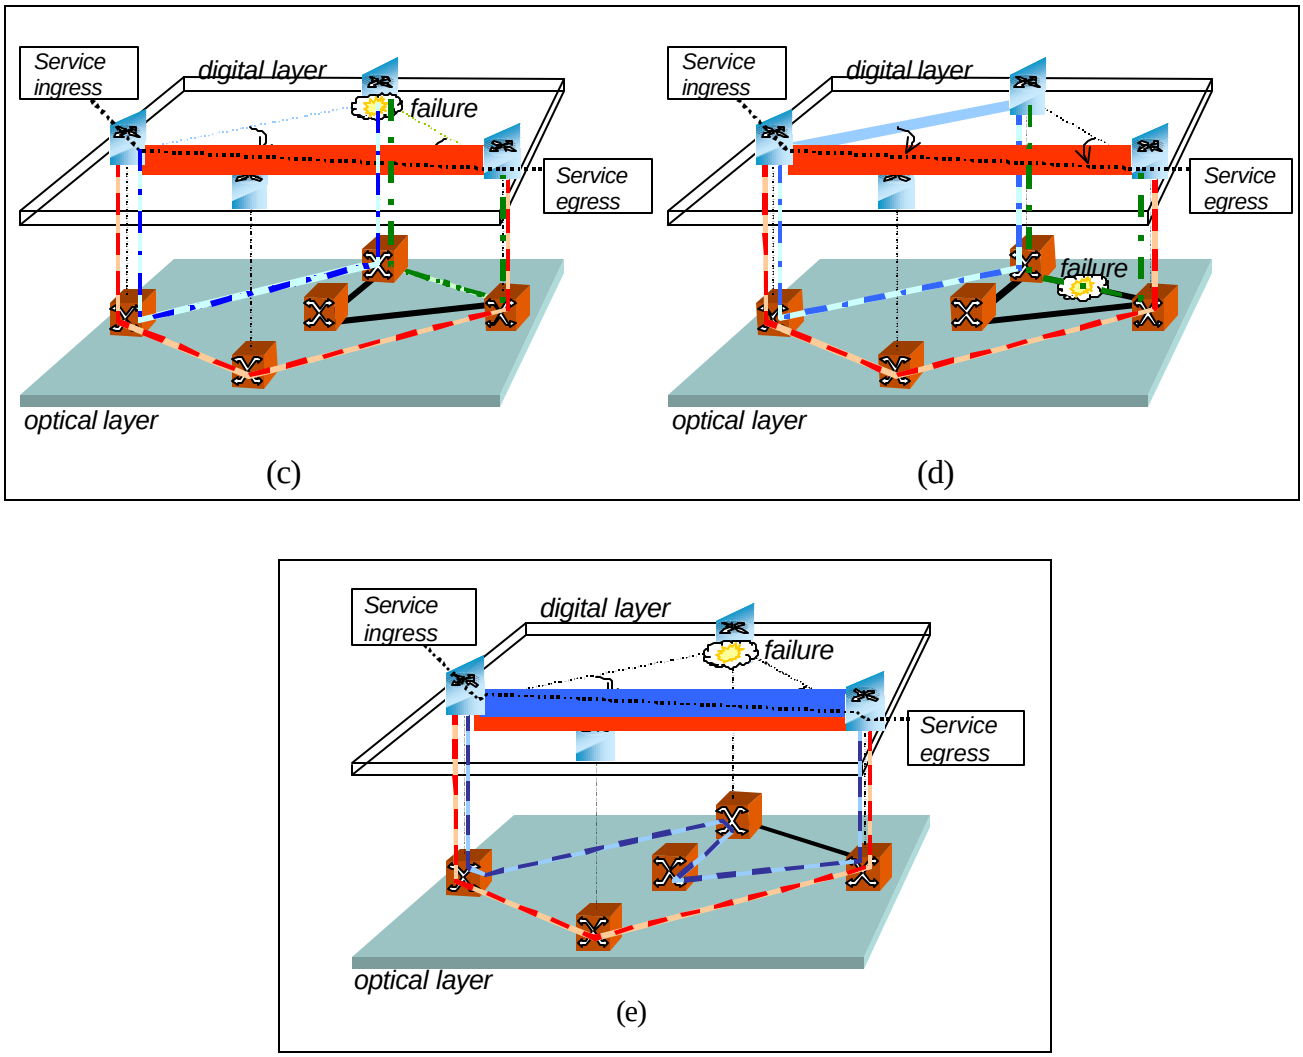

Figure 3: (a) Routing before a failure occurs (b) Failure and restoration in the optical layer (c) Failure and restoration in the logical layer (d) Optical failure restoration in the logical layer (e) Logical failure restoration in the optical layer

To summarize, although both scenarios one and three mentioned above address the same problem of recovering from failure in the optical layer, the first, which recovers the failure in the layer where it occurs, is preferable in terms of cost, and speed [21,22]. The same is also true with the second over the fourth scenario. In addition, because the preferred mechanisms are confined within their own layers, that helps simplify the restoration approach, and avoid architectural complexities and interdependence of mixed-layer approaches.

4 Restoration in the Optical Layer

In end-to-end OXC-based path protection, the ingress and egress nodes of the failed optical connection attempt to restore the signal on a predefined backup path, which is link-disjoint from the primary path. Path diversity

guarantees that primary and backup lightpaths will not simultaneously succumb to a single failure. There are two sub-types of path protection: (1) 1+1 dedicated protection, and (2) mesh restoration.

4.1. Dedicated Protection

Dedicated 1+1 protection is illustrated in Figure 4. The network consists of four logical nodes (A to D) and two demands (AB and CD) accommodated across an eight node optical network (S to Z.) The provisioning algorithm of this architecture computes and establishes simultaneously the primaries and their link-disjoint protection paths. During normal operation mode, both paths carry the optical signal and the egress node selects one of the two copies. In the example of Figure 4, all the optical channels on primary and secondary paths are active. In particular, the configuration reserves two optical channels between nodes S and T for protection. This is the fastest restoration scheme since for every lightpath one device at the termination of the lightpath is responsible for all the necessary failure detection and restoration functions. But it is also the most exigent in terms of resource consumption.

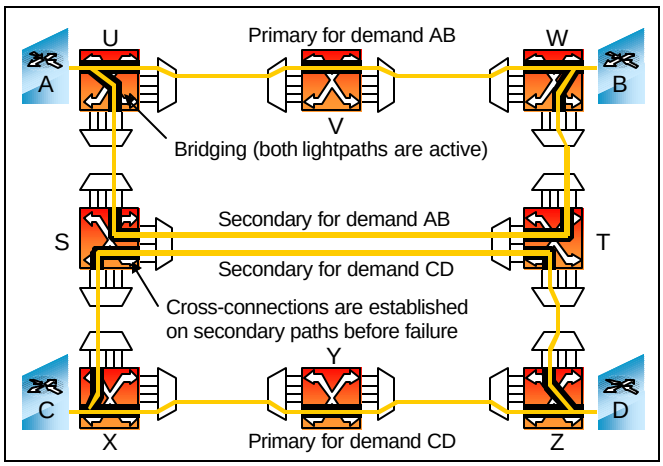

Figure 4: Dedicated 1+1 protection

4.2. Mesh Restored Lightpaths

As in dedicated protection, in mesh restoration backup paths are predefined, but the cross-connections along these paths are not created until a failure occurs. During normal operation modes the spare optical channels reserved for protection are not used. Since the capacity is only "soft reserved", the same optical channel can be shared to protect multiple lightpaths. There is a condition though that two backup lightpaths may share a reserved channel only if their respective primaries are link-disjoint, so that a failure does not interrupt both primary paths. If that happened, there would be contention for the reserved channel and only one of the two lightpaths would be successfully restored. Figure 5(a) (for normal mode) and Figure 5(b) (for restoration mode) picture an example of mesh restoration. The dashed lines represent reserved channels. In this case, the protection paths for demands AB and CD share a single optical channel in link S-T, one less than in dedicated protection.

However, the restoration involves a bit more processing to signal and establish the cross-connections along the restoration path. There is thus an evident trade-off between capacity utilization and recovery time.

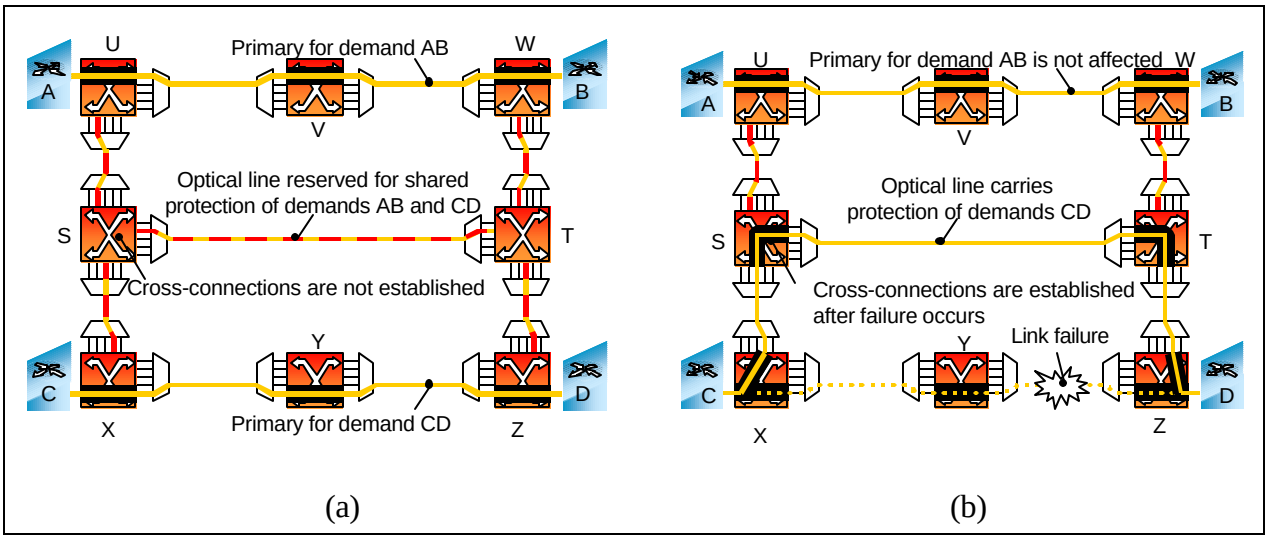

Figure 5. Mesh Restoration: (a) Network connections before a failure occurs (b) Network connections after a failure occurs

Simulation experiments were run on a 100-node (N100), 137-edge network that has a degree distribution of (50,28,20,2) nodes with respective degrees (2,3,4,5). It was assumed that this architecture has infinite link capacity. Three network robustness scenarios were considered: no protection; dedicated protection with provisioning to recover from single link failures; and mesh restoration to recover from link failure. In N100, 3278 node-pairs out of 4950 possible node pairs are connected by one bi-directional lightpaths. Requests for lightpaths arrive one at the time (on-line routing) in a finite sequence and in an order that is arbitrary but common to each scenario to ensure a fair comparison. Figures of merit are capacity requirements separated into their primary and restoration parts, and expressed in units of bi-directional OC-48 channels. Results are presented in Figure 6. The quantities shown on the chart are averages over series of 10 experiments using various demand arrival orders. These results clearly demonstrate the advantages of shared mesh restoration over dedicated protection in terms of redundant capacity required to protect against all single link failures [23].

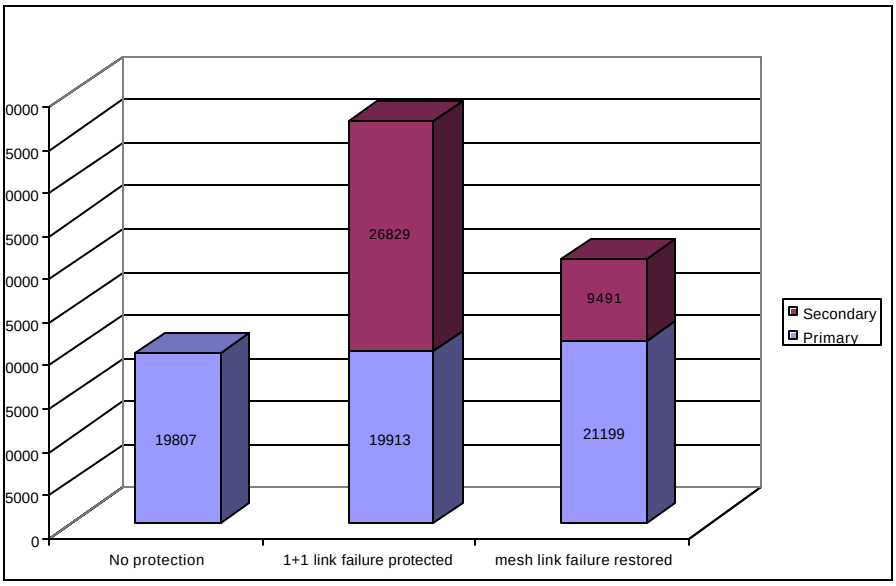

Figure 6. Comparison of capacity usage for different protection architectures (100 nodes)

5 Multilayer Protection

Notwithstanding the architectural complexities and interdependence of mixed-layer restoration approaches, this section investigates how more than one layer can offer protection in a multilayer network. The goal of such an approach would be to use the protection capabilities of each layer in order to provide additional survivability capabilities in the network.

A number of different multilayer protection strategies can be implemented in a layered architecture. As pointed in the previous section, restoring the network at the layer where the failure occurred is a fast approach better suited for the recovery of failures in that layer. In addition, protection and restoration in different layers may be mixed together for the best overall result. For instance, fast optical protection architecture for fiber and OXC failures can be supplemented by service-based restoration at the logical layer. In this case, the optical layer can offer bulk recovery of the services while the logical layer can offer finer restoration granularity.

5.1. Escalation Strategies

If a multilayer protection approach is adopted, an escalation strategy has to be provided to coordinate the protection processes of the different layers. The absence of an escalation strategy can create race conditions between the protection mechanisms with unpredictable (and potentially catastrophic) results. This section identifies the issues associated with the escalation strategies, and presents a comparison of the different escalation options.

The escalation strategies can include either a parallel or a sequential activation of restoration mechanisms. In the parallel approach, restoration mechanisms from different layers are trying to restore the same failure simultaneously, which will result in a very fast restoration time. However, the different restoration mechanisms must be coordinated so as not to obstruct each other or compete for the same restoration resources. In the sequential case, restoration mechanisms from different layers attempt to restore the failure one layer at a time. One sequential approach could be to wait until one layer has failed to restore the services or a fixed time interval has passed before the restoration process is taken over by another layer [24,25].

Typically, the sequential approach will be slower in terms of overall restoration times, but it is more easily implemented. Such a strategy only needs to predetermine the order in which the layers will attempt restoration and when the transition from one layer to another takes place. The type of failure will usually determine the layer where the sequential restoration process will start. For example, if a fiber is cut, restoration can start at the optical layer, where it can achieve fast restoration of bulk traffic. On the other hand, if an IP router fails, restoration can start at the logical layer. This will allow for finer granularity of the restoration process but it will be much slower.

A hold-off timer function can be used in the sequential approach to mark the transition from the protection mechanism of one layer to the protection mechanism of another layer. Although this is a simple approach, it can result in a cascade effect and potentially introduce considerable delay to the restoration process. Another sequential escalation strategy can use a diagnostics method to ascertain that a protection mechanism will not be successful before the hold-off timer expires and subsequently hand over the restoration control to the next layer. Even though this scheme will result in much faster restoration times, this is a much more complex approach to implement. Other escalation strategies include signaling approaches (e.g., through UNI or through the management system) to coordinate the passing of the protection responsibilities to another layer.

Take for example a multilayer network consisting of IP, Gigabit Ethernet (GbE) and WDM layers. Restoration in such a network can take place at the WDM and/or IP layers, as GbE does not support restoration capabilities. If the WDM layer only is used for restoration, fast restoration can be utilized to quickly restore services during a network failure condition at the optical layer. If the IP layer only is used for restoration, router-based rerouting will result in longer restoration times but it will require less redundant capacity, because of the finer restoration granularity in this layer. In general, using the IP layer

only for restoration may not meet the service's QoS requirements. However, for some data services such as e-mail and ftp restoration at the IP layer only may be sufficient.

If a two-layer strategy is employed, a parallel escalation strategy may suffice. Because of the different time scales between the restoration mechanisms in the WDM and IP layers, the WDM layer will restore the traffic before the IP layer even detects the error. A sequential approach on the other hand would require coordination between the optical layer and the routing protocols used in the IP router-based rerouting. The best approach for such a network would be to use the optical layer for the restoration of time-sensitive services and use IP router-based rerouting for the restoration of less critical and time-insensitive services.

6 Conclusion

This paper describes multilayer networks and how a single layer utilizes its protection capabilities in order to restore services after a failure occurs. Specifically, restoration in the optical layer is described and dedicated as well as mesh lightpath-based restorations are presented. A number of escalation strategies are also presented for the case where a combination of different restoration methods from different layers are to be used. From the discussion it is evident that the simplest solution is to allow the layer where the failure occurred to restore for that failure. This will help in avoiding functional duplication, race conditions, as well as complex interaction and inter-layer coordination problems. Even though allowing for a multilayer restoration mechanism can enhance the overall survivability of the network, complex escalation techniques will violate the layering principle by violating inter-layer independence. Since the optical layer is the lowest layer in the transport hierarchy, using the restoration mechanism in that layer is ideal for fast recovery of critical services during failure conditions such as fiber cuts and optical switch failures. In contrast, IP router-based rerouting is a best effort approach more suitable to the restoration of less critical and time-insensitive data.

7 References

[1] T.E. Stern, K. Bala, Multiwavelength Optical Networks: A Layered Approach, Prentice Hall, May 1999.

[2] D. Benjamin, R. Trudel, S. Shew, E. Kus, "Optical Services over an Intelligent Optical Network", IEEE Communications Magazine, pp. 73-78, September 2001.

[3] E. Varma, S. Sankaranarayanan, G. Newsome, Z. Lin, H. Epstein, "Architecting the Services Optical Network", IEEE Communications Magazine, pp. 80-87, September 2001.

[4] A. McGuire, S. Mirza, D. Freeland, “Application of Control Plane Technology to Dynamic Configuration Management”, IEEE Communications Magazine, pp. 94-99, September 2001.

[5] “IP over Optical Networks: A Framework”, draft-ietf-ipo-framework-00.txt, IETF draft, January 2001.

[6] B. Rajagopalan, D. Pendarakis, D. Saha, S. Ramamurthy, and K. Bala, “IP over Optical Networks: Architectural Aspects”, IEEE Communications Magazine, September 2000.

[7] G. Bernstein, J. Yates, D. Saha, “Control and Management of Optical Transport Networks”, IEEE Communications Magazine, October 2000.

[8] G. Hjálmtýsson, J. Yates, S. Chaudhuri and A. Greenberg, "Smart Routers – Simple Optics: An Architecture for the Optical Internet", IEEE/OSA Journal of Lightwave Technology, December 2000.

[9] G. Bernstein, J. Yates and D. Saha, "IP-Centric Control and Management of Optical Transport Networks", IEEE Communications Magazine, pp. 161-167, October 2000.

[10] K. Sato, S. Okamoto, “Photonic Transport Technologies to Create Robust Backbone Networks”, IEEE Communications Magazine, pp. 78-87, August 1999.

[11] “Generalized Multi-Protocol Label Switching Architecture”, draft-many-gmpls-architecture-00.txt, IETF draft, August 2001.

[12] “Generalized MPLS - Signaling Functional Description”, draft-ietf-mpls-generalized-signaling-05.txt, IETF draft, January 2001.

[13] “Generalized MPLS Signaling - RSVP-TE Extensions”, draft-ietf-mpls-generalized-rsvp-te-04.txt, IETF draft, January 2001.

[14] “Link Management Protocol”, draft-ietf-mpls-lmp-02.txt, IETF draft, September 2001.

[15] “OSPF version 2”, IETF RFC 2328, April 1998.

[16] “User Network Interface (UNI) 1.0 Signaling Specification”, OIF contribution OIF2000.125.6, September 2001.

[17] “Link Management Protocol (LMP) for WDM Transmission Systems”, IETF draft, June 2001.

[18] P. Demeester, M. Gryseels, A. Autenrieth, C. Brianza, L. Castagna, G. Signorelli, R. Clemente, M. Ravera, A. Jajszczyk, D. Janukowicz, K.V. Doorselaere, Y. Harada, “Resilience in Multilayer Networks”, IEEE Communications Magazine, pp. 70-76, August 1999.

[19] S. Makam et al, “Framework for MPLS Based Recovery”, draft-ietf-mpls-recovery-frmwrk-03.txt, IETF draft, July 2001.

[20] R. D. Doverspike et al. “Transport Network Architectures in an IP World”, Proc. Infocom’2000, pp. 305-314, 2000.

[21] J. Manchester, P. Bonenfant, C. Newton, “The Evolution of Transport Network Survivability” IEEE Communication Magazine, pp. 44-51, August 1999.

[22] R. Doverspike, S. Phillips, J. Westbrook, “Transport Network Architectures in an IP World”, Proc. Infocom’2000, 2000.

[23] E. Bouillet, G. Ellinas, R. Ramamurthy, J. Labourdette, S. Chaudhuri, K. Bala “Routing and Restoration Architectures in Mesh Optical Networks”, to appear as an invited paper in the Optical Networks Magazine.

[24] “Planning for Service Survivability in Broadband Multilayer Networks”, Bellcore SR-4317, August 1997.

[25] R. Batchellor, “Coordinating Protection in Multiple Layers”, OIF Contribution, oif99.038.0, April 1999.