Routing and restoration architectures in mesh optical networks

Georgios Ellinas
Eric Bouillet
Ramu Ramamurthy
Jean-François Labourdette
Sid Chaudhuri
Krishna Bala
Tellium Inc.
Oceanport, NJ USA

ABSTRACT
This paper provides an overview of various techniques used to provision lightpaths in a layered architecture that utilizes a dedicated control and management plane for each layer. It also reviews and compares a set of OXC-based protection and restoration architectures, focusing only in the optical domain. The comparison takes into consideration figures of merit such as economic aspects, availability of restoration, and speed of restoration to sieve out the most appropriate protection schemes in conformance to each scenario.

1 Introduction

Dense Wavelength Division Multiplexed (DWDM) mesh networks that route optical connections using optical cross-connects (OXCs) have been proposed as the means to implement the next generation optical networks [1]. Following a wave of timely technologic breakthroughs, optical network equipment vendors are now announcing a variety of optical switching systems capable of exchanging and redirecting several terabits of information per second. The dimensions of the proposed switches are colossal, ranging from a few tens to several thousand ports with each single port capable of carrying millions of voice calls, or thousands of video streams. Optical network architectures as we envision them now not only provide transmission capacities to higher transport levels, such as inter-router connectivity in an IP-centric infrastructure, but also provide the intelligence required for efficient routing and fast failure restoration in core networks [2-4]. This is possible due to the emergence of optical network elements that have the intelligence required to efficiently manage such networks.

The network architectures under consideration in this paper contain opaque (OEO) switches (with an electronic switch fabric) in an opaque network (with transponders present in the WDM systems). The interfaces to the switch fabric are opaque interfaces, which means that transceivers are present at all interfaces to the switch, and these transceivers provide an OE (input) and EO (output) conversion of the signal. The presence of the transceivers at the edges of the switch fabric enables the switch to access the signal's overhead bytes for control and signaling functions. The opaque transceivers provide support for fault detection and isolation, performance monitoring, connection verification, neighbor/topology discovery and signaling, as well as support for implementing the network routing and restoration protocols.

Optical network architectures are exposed to multifarious risks of breakdowns, either due to human-induced mishaps such as accidental fiber cuts and operation mistakes, or equipment malfunctions such as switch and laser failures. Two competing approaches are being proposed for providing the appropriate recovery mechanisms that guarantee service flow continuity in these circumstances. In the *peer-to-peer* approach [5-7], interweaved optical and higher layer equipment act in symbiosis under the same control plane. In the *overlay* approach [5,6], optical and logical 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. In this presentation the focus is on restoration in the optical domain assuming the overlay approach. We compare different lightpath on-line routing techniques, as well as different OXC-based protection and restoration approaches and we derive tradeoffs and discuss reliability versus complexity and economic considerations.

In section 2 of this document we describe the network architecture and specifically the optical layer in that architecture. Section 3 deals with the objectives and constraints that must be considered when selecting a route in the optical layer. Section 3 also discusses the advantages and disadvantages of centralized versus distributed routing algorithms. Section 4 addresses OXC-based protection/restoration in an optical mesh network. It defines (and differentiates) the notions of Shared Risk Group and node switch failures in the optical layer and provides a comparative overview of existing and proposed OXC-based protection and restoration mechanisms, with details on costs, restoration speed, level of protection, and implementation complexity for each mechanism. Experiments and results are presented in Section 5, followed by concluding remarks in Section 6.

2 Network Architecture

Many network architectures exist or have been suggested and it is not the intention of this paper to enumerate them exhaustively (see [5,8-10] 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.

This paper focus is only on the optical layer, which 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] 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 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 [16]. In addition, an optical carrier will normally acquire network components from several vendors. A suite of protocols is being developed to allow for the

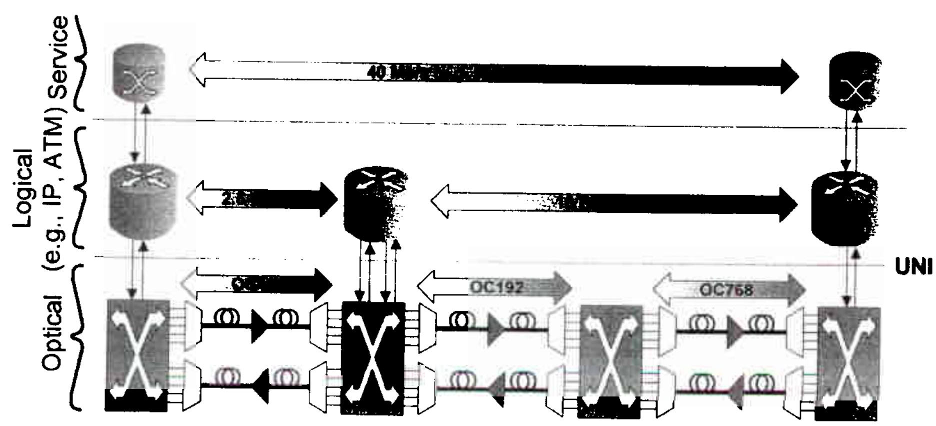

Figure 1: Layered architecture.

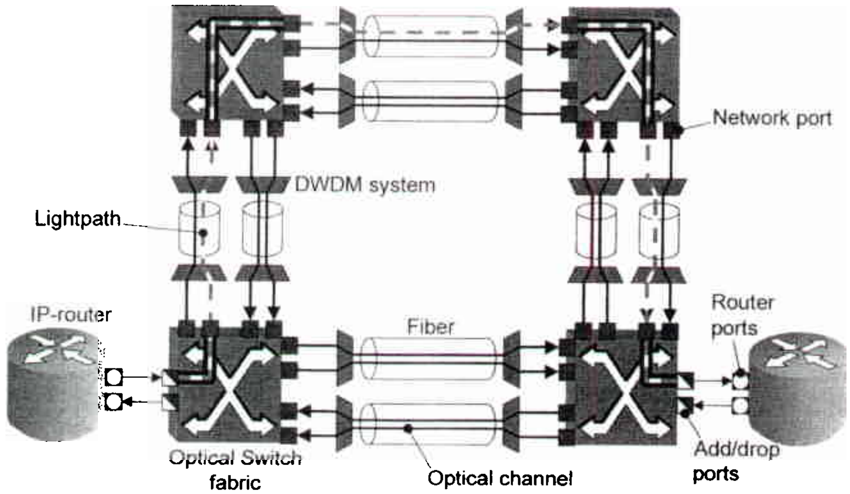

Figure 2: Optical layer.

seamless interaction between the various network components. As part of the suite, 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].

3 Routing in the Optical Layer

3.1 Route computation

Routing in a multi-layered architecture is preferably done separately in the logical and in the optical layer [18]. It is the latter that is addressed here. When solving the route computation problem, several metrics need to be considered. Depending on the allotted budget and the desired QoS, each metric either enters as a parameter in the algorithm's objective function to be minimized, or used as constraint to eliminate solutions that do not meet practical limits [19]. Some of these metrics are:

1. Cost: The use of optical channels entails a cost. It is henceforth important to ensure that the cumulated cost does not exceed the client's budget.
2. Bit-rate: The bit-rate of all the selected optical channels along the route must meet the lightpath's bandwidth prescribed by the client layer. If the switch implements grooming operations, then the capacity of one optical channel may be shared among several lightpaths of lower capacity.
3. Bit-propagation delays—same as cost. Requires a link length or link delay attribute.

The objective of the shortest path algorithm is to either preserve spare capacity by using a minimum number of optical channels, or to find the solution that incurs the minimum cost. Therefore, among the link attributes that are necessary to achieve routing operations, are (a) the cost per optical channel, (b) the list of available optical channels in the link, (c) their bit-rates, (d) the grooming capabilities of the link, and (e) the propagation delay. Other factors that must be considered are:

4. Resilience (e.g., link failure versus node failure, single failure versus multiple failures).
5. Maximum restoration-time required to restore a service if a failure occurs.
6. Complete or partial knowledge about the network-state and link attributes. We address this issue in the next section.

It is important to note that the routing problem addressed in this paper is on-line routing in an opaque network architecture that utilizes OEO switches, where wavelength conversion is a byproduct of the switch architecture.

3.2 Centralized vs distributed routing

The procedure to route a lightpath consists of two tasks: (1) route selection, and (2) channel selection. Route selection involves computation of the primary and backup paths from the ingress port to the egress port across the mesh optical network. Channel selection deals with selecting individual optical channels along the primary and backup routes. The problems of selecting a route together with selecting channels on the route are closely coupled and if an optimal solution is sought both problems should be solved simultaneously. In practice, depending on the view of the network and the amount of information available at each node, both tasks can be accomplished in either a centralized, or in a distributed manner. In the centralized approach, the joint problems of route computation and channel assignment for a lightpath are solved from a Centralized Management System (CMS) that has access to the complete network state of the optical cross-connects, including the topology and lightpath databases. In the distributed approach, lightpaths are routed by the Control Module (CM) of the ingress OXC using its local databases. The local database at the switch contains the summarized topology disseminated by a link-state protocol such as OSPF, and those provisioned demands whose primary or backup paths traverse that switch. Routing is therefore performed first and is in the most part oblivious to the network-wide configuration of the lightpaths; channels are assigned next.

In centralized provisioning a request for a lightpath is established through the UNI or the Network Management System (NMS) and sent to the CMS. The CMS computes the route, assigns the optical channels along the route, and send a request to establish the lightpath to the ingress CM of the route. The message sent by the CMS to the CM contains the description of the route and the port numbers of each optical channel. The ingress CM configures the associated OXC to create a new connection for the lightpath in accordance with the information provided by the CMS, or returns an error message with the updated state of the switch to the CMS if it cannot create the connection. The CM forwards the request to the next CM on the route and waits for lightpath-setup confirmation from this CM. The lightpath is established and the CMS returns "success" to the UNI if all CMs along the route successfully complete the cross-connection. In distributed provisioning, the UNI request is sent directly to the ingress CM, which acts as a substitute for the CMS. The CM uses its local database to compute the route, but does not have all the necessary information to select the optical channels along the route except for the optical channels adjacent to it. Instead, each CM selects the adjacent optical channel. It then executes the cross-connection, submits the request to set-up a lightpath to the next CM along the route and waits for lightpath set-up confirmation from this CM.

Delegating routing operations to the CMs (distributed routing) enhances network scalability. The downside of distributed routing is that it is possible that the route is sub-optimal or even cannot satisfy all the constraints (although some techniques can be employed to estimate the state of the network using appropriate partial information [20]). Such a solution will lead to increased network capacity requirements and higher network cost [21]. Multiple routing attempts may also be necessary before a feasible path is found.

Finally, note that strategy employed in routing, and the strategy used in the restoration mechanism as described in the next section are independent. It is possible for instance to provision a lightpath in a centralized manner, while relying on a distributed mechanism to protect it.

4 Restoration in the Optical Layer

4.1 Shared risk groups

Failures of multiple optical channels are usually due to fiber or cable cuts. Consider the 6-node optical network of Figure 3(a). 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 Group (SRG) expresses the risk relationship that associates all the optical channels with a single failure [19,22]. 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. Since a fiber may run through several conduits, an optical channel may belong to several SRGs. By default in this paper all the optical channels between one node-pair belong exclusively to one SRG. The routing algorithms exploit SRG maps to discover SRG-diverse routes so that after any conduit is cut, there is always at least one viable route remaining for restoration [19]. For instance, in Figure 3(a) the subtraction of any SRG and the optical channels that traverse it, affects at most one of the two shown routes from A to B.

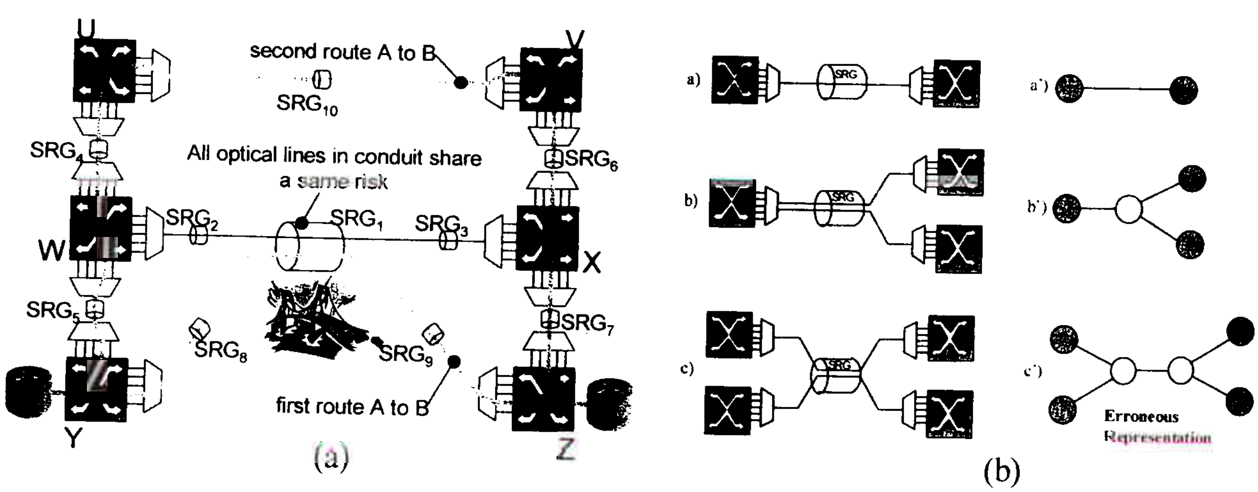

Figure 3: (a) Shared risk groups (b) SRG classification.

For obvious reasons, a network topological view alone does not encompass the notion of SRGs. With the exception of the default case, there exists no simple way to automatically generate this information. The network operator must manually provide it.

All SRGs can be expressed as one or a combination of three possible primary types. We describe them in Figure 3(b). The default and most conventional type, is type (a) in the figure, which associates an optical channel risk failure with a fiber cut. Another type of SRG very likely to be encountered, is type (b). This type is typical of fibers terminating at a switch and sharing the same conduit into the office; that is a conduit cut would affect all the optical channels terminating at the switch. Types (a) and (b) can be characterized in a graph representation as pictured by (a′) and (b′). For instance, the removal of the edge in (a′), or the middle node in (b′), disconnects all the nodes which is tantamount to an SRG failure in each case. Using these elementary transformations it is possible to model the network as a graph onto which established shortest-path operations can be applied. Type (c) SRG is the most difficult kind to model and provide diverse routing for. It occurs in a few instances, such as for example fibers from many origins and destinations routed into a single submarine conduit, or dense metropolitan areas. Contrary to type (a) and (b), there is no convenient way to graphically represent type (c) SRGs and their presence can increase dramatically the complexity of the SRG-diverse routing problem. A naïve representation of the type (c) SRG would be to present it as graph (c′). Such a representation however is erroneous, as it introduces additional paths not present in the original network topology, which could lead us to routing computations that are not physically feasible.

Note that if an arbitrary set of links can belong to the same SRG, then the problem of finding SRG-diverse primary and backup routes between a pair of nodes in the network is NP-complete (see Appendix 7.1). Essentially, the difficulty of 1 + 1 SRG-diverse routing arises because the architecture allows SRGs to be defined in arbitrary and impractical ways which forces an algorithm to enumerate (potentially exponential number of) paths in the worst-case (unless P = NP).

4.2 Node vs SRG failure

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 explained in Section 4.1 and illustrated in Figure 4(a). Some level of node failure protection can be realized by way of a redundant switch fabric. Note that this architecture does not recover from severe events, such as electrical fires, that affect both switching fabrics. If protection against this type of failure is also desired, it is necessary to provision routes that are SRG-and-node disjoint, as shown in the example of Figure 4(b). However, node-diverse paths consume more resources than the less conservative SRG-diverse scheme pictured in Figure 4(a).

4.3 Taxonomy of protection/restoration schemes

In this section, we classify the existing protection/restoration types. For each type of protection/restoration we briefly explain its distinctive characteristics such as tradeoffs of cost versus efficiency, speed of recovery, and other features that set it apart from other types without discussing the implementation details.

Before entering into the details of the protection/restoration taxonomy, we present in Table 1 six possible protection/restoration approaches (this is the extension of a similar table presented in [23]), the details of which are presented later in this paper. The table enumerates the three components managed during restoration. The components are the alternate route around a failure, the channels used along that route, and the embedding of the route into the optical switches. Each category indicates the dependence of each component on the origin of the failure. Components that do not depend on the failure may be assigned before the failure occurs. For components that are assigned after failure occurrence, the table also distin-

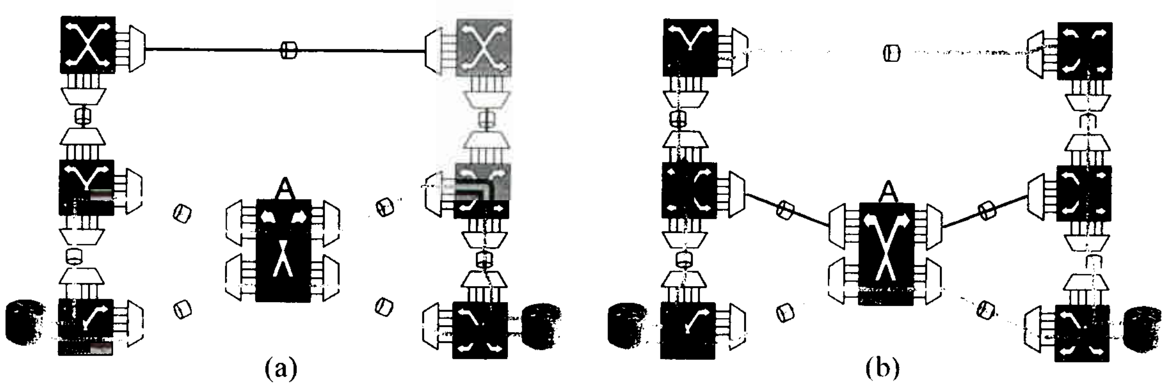

Figure 4(a): SRG-diverse paths (b) Node-diverse paths.

Category	Restoration Route		Channel Assignment on Restoration Route		Cross-connect on Restoration Route	Failure Specific?
	Computed	Assigned	Computed	Assigned		
1 + 1 Dedicated Protection (Cat. 1)	Before	Before	Before	Before	Before	No
Shared Mesh Restoration with Pre-assigned Channels (Cat. 2)	Before	Before	Before	Before	After	No
Shared Mesh Restoration with Reserved, not Pre-assigned Channels (Cat. 3) [26]	Before	Before	After	After	After	Yes (channel only)
Shared Mesh Restoration with Preplanned Maps (Routes and Channels) (Cat. 4)	Before	After	Before	After	After	Yes (route and channel)
Shared Mesh Restoration with Preplanned Maps (Routes only) (Cat. 5)	Before	After	After	After	After	Yes channel only)
Reprovisioning (Cat. 6)	After	After	After	After	After	Yes (route and channel)

Table 1: Different protection/restoration categories and their dependence on failure origins.

guishes between pre-computed components and components determined in real-time. The former depends on the ability of the optical network to perform rapid fault isolation and select the pre-computed components from a look-up table or map based on the origin of the fault.

In this paper, with a few exceptions, we only consider the case of pre-computed restoration paths independent of the failure (categories 1 and 2 in Table 1). For the case where the restoration paths are computed in real-time after a failure is detected and localized, see [24,25]. In addition, we highlight architectures where the restoration path is the same for all possible failures (as opposed to architectures where a different pre-computed path may exist for different failures).

Starting from the highest level of this hierarchy, protections against failures can be classified into two main categories, (1) *local span protection,* and (2) *end-to-end lightpath protection* or *path protection.*

4.3.1 Local span protection

In OXC-based local span protection, whenever a failure is detected, the optical nodes closest to the failure attempt to reroute the lightpaths through alternate circuits around the failure. This protection scheme theoretically yields faster and higher availability, since most of the time only the disabled portion of the path is bypassed. On the downside, the alternate routes differ for each failure and are difficult to anticipate. The theoretically preferred approach to address this problem is to simulate all possible failures and create directive maps stored in the optical switches that assign a pre-computed switch configuration for each failure scenario (categories 4, 5 in Table 1). However, the generated maps may become very large and the concept does not scale well with the size of the network. The approach also entails lengthy computations whenever a new lightpath is provisioned since it must account for every failure scenario in order to populate the map. For this reason, spare capacity is usually not reserved ahead of time. Instead, the restoration routes are computed on the fly upon a failure event (category 6 in Table 1). This becomes an issue if the failure disrupts many parallel optical channels and sets off a cascade of real time recovery procedures at the optical switches adjacent to the failure. In either case, this protection scheme relies on the ability of the network to isolate the failure. Furthermore, link-based restoration schemes require more restoration capacity than path-based schemes [27,28].

4.3.2 Path protection

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 SRG-disjoint, or diverse, from the primary path [29,30]. Path diversity guarantees that primary and backup lightpaths will not simultaneously succumb to a single failure. Unlike local span protection, secondary routes are provisioned with the primaries and thus the restoration does not involve further real-time path computations. Also, the detection of the signal degradation or signal failure takes place at the add-drop ports of the path's egress switches, where access to the overhead bytes is a byproduct of the port's design. Another advantage of OXC-based path protection is that the restoration processing is distributed among ingress and egress OXC nodes of all the lightpaths involved in the failure, compared to OXC-based local span protection where a comparable amount of processing is executed by a smaller set of OXC nodes adjacent to the failure. In the following, we will only consider the cases where the protection path is failure-independent and is thus the same for all types of failures. By way of this restriction, the restoration paths may be computed and assigned before a failure occurrence. There are two sub-types of path protection: (1) 1 + 1 dedicated protection, and (2) shared mesh restoration.

4.3.2.1 Dedicated protection

(a) Protect against single SRG failures

Dedicated 1 + 1 protection is illustrated in Figure 5. 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 SRG-disjoint protection paths. During normal operation mode, both paths carry the optical signal and the egress node selects one of the two copies. This is the fastest restoration scheme since for every lightpath one device is responsible for all the necessary failure detection and restoration functions. But it is also the most exigent in terms of resource consumption. The problem of finding SRG-diverse routes is trivial if SRGs are of type (a) or (b), or a combination of both, since they can be easily represented in a graph model as indicated earlier in Section 4.1. There exist optimal algorithms (in terms of minimizing the capacity required) to

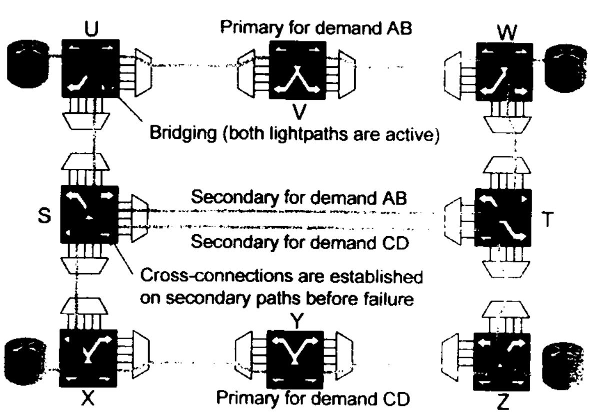

Figure 5: Dedicated 1 + 1 protection.

solve this problem [31]. For the case of type (c) SRGs the problem is provably NP-complete and pseudo-optimal solutions must be obtained, using enumerative approaches. The problem is also NP-complete if the selected routes must respect a set of independent constraints, such as maximum round trip-delay (or alternatively maximum path length expressed in geographical distance units).

(b) Protect against single node failures

If protection against node failure is also desired, then primary and backup paths must be node-disjoint in addition to SRG-disjoint. As explained earlier, node protection may require more protection capacity. However, experiments indicate that the two types of protection use comparable amount of capacity. The problem of finding node diverse routes is equivalent to the problem of link-diverse path routing. The same algorithm is capable to solve either problem, using the transformation pictured in Figure 6 to represent nodes as directed links, by introducing an in-node (shown as ⊕) and an out-node (shown as ⊖) [32]. Clearly, if a pair of paths is edge-disjoint in Figure 6(b), then it is also node-disjoint in Figure 6(a). In this example there is no edge-disjoint path from A to Z in the graph representation, and thus no node-disjoint pair of paths in the network.

4.3.2.2 Shared mesh restored lightpaths

(a) Protect against single SRG failures

As in dedicated protection, in shared mesh restoration backup paths are predefined, but the cross-connections along these paths are not created until a failure occurs. During normal operation 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 SRG-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. Two lightpaths, or their protection, are said to be mutually compatible, if they are not affected by the same failure. If not, they are incompatible. Figure 7(a) (for normal mode) and Figure 7(b) (for restoration mode) picture an example of shared mesh restoration. The demand and the network are the same as in dedicated protection. The dashed lines represent reserved channels. Using the routing of Figure 7(a), demands AB and CD are compatible with respect to SRG-failures and thus their protection paths share a single optical channel in link S-T, one less than in dedicated protection. Upon failure as depicted in Figure 7(b), the ingress and egress nodes of the disconnected paths (X and Z in example) emit a request to the switches along the protection paths (S and T in example) to establish the cross-connections for that path. Once the cross-connections are established, each ingress and egress node restores

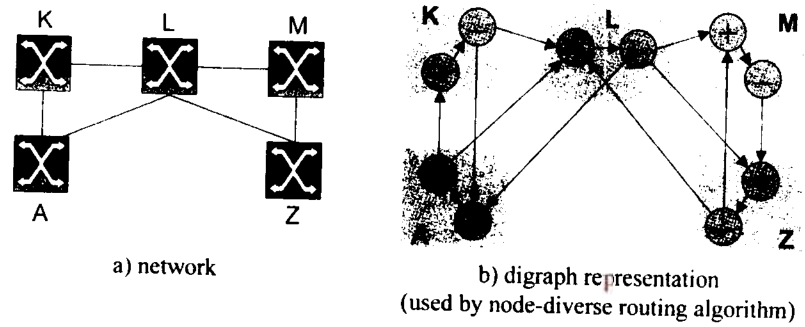

Figure 6: Network and equivalent graph representation.

the connection to the new path. This architecture requires fewer resources than in dedicated protection. However, the restoration involves 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.

Note that in the previous section reserved channels were assigned to the restoration paths regardless of the origin of the failure. They are actually two policies to assign reserved channels to restoration paths. A *failure dependent* strategy assigns the reserved channels in real time after failure occurrence on a first-come, first-served basis depending on availability (category 3 in Table 1) [33]. We assume that a proper spare channel-provisioning scheme ensures enough restoration channels so that all lightpaths can be restored. A *failure independent* strategy assigns the reserved channels at the time of lightpath provisioning prior to a failure occurrence (category 2 in Table 1). The advantage of the second approach is that during lightpath restoration the switches on the protection paths immediately and individually cross-connect to predetermined channels, based on the identifier of the lightpath being restored.

If the channels are not pre-assigned, adjacent switches on the protection path must agree on the channels to be used for restoring the lightpath before establishing a cross-connection. This approach requires a handshake protocol and inter-switch signaling, which may be time consuming and inadequate if restoration speed is an issue. However, it is possible that in very rare pathological cases, the failure dependent approach reserves less channels for an equivalent level of protection. Consider for instance Figure 8. The figure represents four primary lightpaths routed in an optical network and their respective protection paths outlined with dotted lines. In this example the optical switches are not pictured. The lightpaths are pair-wise intersecting over six different fibers. The protection paths are all intersecting into one single fiber and may share reserved channels. A primary fiber cut in this example interrupts only two lightpaths, and thus two reserved channels are enough to restore the paths on a first-come first-serve basis. If the channels are predetermined in a failure independent manner, four reserved channels, one for each protection path, are necessary to accommodate all existing combinations of lightpath pairs. However, experimental results indicate that in most situations both approaches are equivalent and the potential saving even if it exists may not justify the additional complexity and processing time of the failure dependent approach [26].

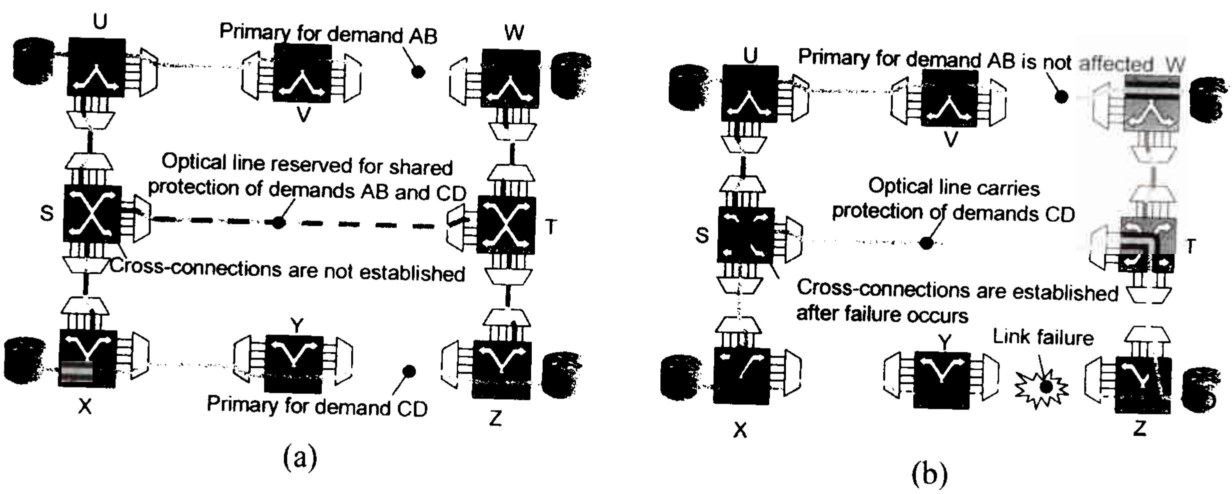

Figure 7: Shared Mesh Restoration: (a) Network connections before a failure occurs (b) Network connections after a failure occurs.

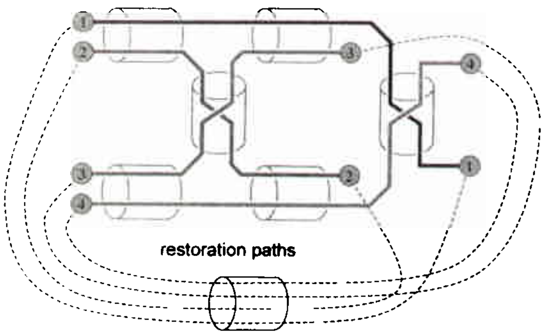

Figure 8: Pair-wise full interference (optical switches are not pictured).

Given an optical network with default SRGs, and a set of already established shared mesh restored lightpaths, the problem of finding a feasible primary and shared-backup path for a new lightpath request is NP-complete. The proof of the NP-completeness is shown in Appendix 7.2.

(b) Protect against single node or SRG failures

In shared mesh restoration, node-diversity between primary and backup paths does not guarantee full protection against node failures. Consider the example of Figure 9. The figure illustrates two primary bi-directional lightpaths (solid lines) and their corresponding backup paths (dotted lines) for demands (***d,h***) and (***b,f***). In this example SRGs consist exclusively of all the optical channels in a link. The primary paths intersect on node ***i***, while both secondary paths traverse links ***e-f*** and ***e-d***. The primary paths are SRG-disjoint, and so according to the sharing rules, the secondary paths may share reserved channels in their mutual links. If node ***i*** fails, reserved channels are allotted on a first-come first-serve basis and one demand is lost. This is an acceptable outcome if there is no commitment to restore either demand after a node failure. Otherwise the secondary paths cannot share channels if restoration is also required in the event of a node failure for both demands.

If it is desired that one demand only, say (***h,d***), survives a node failure, then in theory the secondary paths may share reserved channels. However, in practice if node ***i*** fails, contention occurs for the reserved channels during restoration, and a first-come first-serve policy does not enforce systematic restoration of (***h,d***), the demand that must be restored. This problem can be averted by assigning a higher priority to (***h,d***), and if necessary preempting (***b,f***) in the recovery phase. Preemptive, priority-centric protocols entail processing overheads thereby increasing the restoration time, and are not elaborated in this document. If the restoration protocols do not allow preemption (or priorities), the routing algorithms must implement a more conservative redefinition of the sharing condition in order to avoid these pitfalls. As a result of these enhanced sharing-conditions, the routing algorithm can guarantee the appropriate protection level for a lightpath.

The provisioning of shared mesh restored paths that guarantees protection against node failures entails more conservative rules of sharing. Assuming that (a) p_i and p_j denote two primary lightpaths, (b) S_i and S_j denote the SRG sets traversed by lightpaths p_i and p_j respectively, (c) N_i and N_j denote the node sets traversed by lightpaths p_i and p_j respectively, (d) $M_i \subseteq N_i$ and $M_j \subseteq N_j$ denote critical node sets (these are nodes for which restoration is required in case of failure—typically all nodes along the path, except originating and terminating nodes) and (e) Primaries p_i and p_j must be restored if a single failure occurs in $\{S_i, M_i\}$ and $\{S_j, M_j\}$ respectively:

Then p_i and p_j are "compatible" and their secondary paths may share channels if:

$$\text{(a) } S_i \cap S_j = \varnothing$$

$$\text{(b) } M_i \cap N_j = \varnothing \text{ and } M_j \cap N_i = \varnothing$$

Otherwise p_i and p_j are said to be "conflicting".

Rules (a) and (b) express the condition that to be compatible, two primaries must be SRG and critical node-disjoint, that is they must not be affected by the same failure for which recovery is required.

Note that the current tenet in network mesh restoration is that protection against single link failures protects against node failures for nodes of degree up to and including 3. This is true for dedicated protection, but not for shared mesh restoration, as demonstrated in the example of Figure 10.

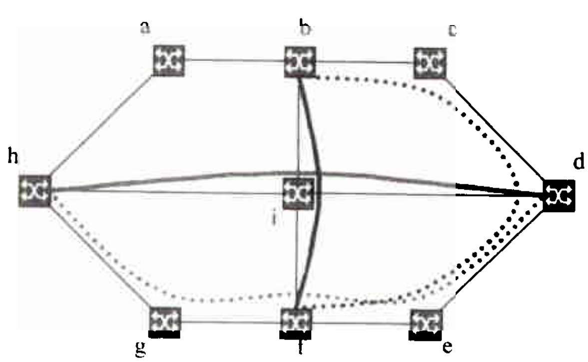

Figure 9: Shared mesh restoration, node disjoint.

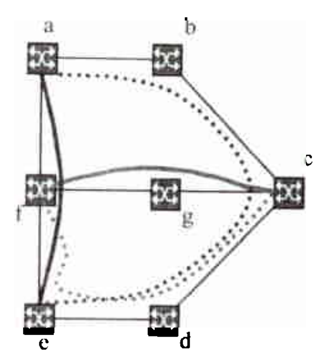

Figure 10: Example of node failures for nodes of degree up to and including 3.

Protection Architecture	SRG type (Section 4.1)	Complexity of on-line Routing	Restoration Resources (% of working capacity)[1]	Speed of Protection
1 + 1 SRG Failure protected	a), b)	Polynomial	100%-170%	Very fast
	c)	NP-complete	>100%	Very fast
1 + 1 SRG and node failure protected	a), b)	Polynomial	100%-175%	Very fast
	c)	NP-complete	>100%	Very fast
SRG failure shared mesh restored	a), b), c)	NP-complete	40%-70%	Fast
SRG failure node restored	a), b), c)	NP-complete	40%-80%	Fast

Table 2: Summary of different architectures and their complexity (centralized routing).
[1]All the percentages are based on a specific set of experiments and may change for different networks and demand sets.

The example depicts two primary bi-directional lightpaths for demands (***a,e***) and (***f,c***). Their respective restoration paths traverse links ***d-e*** and ***c-d***. Node ***f*** is the source of (***f,c***) and an intermediate node of (***a,e***). Even though node ***f*** has degree 3, the sole provisioning of mesh restoration for link failures would not protect the lightpaths against a failure of this node, since node *c* may initiate a request to reserve protection channels on edge ***c-d*** and/or ***d-e***, and thereby prevent the restoration of (***a,e***). If such level of protection is required, the rules of sharing stipulated above must be enforced with node ***f*** in the critical node list of demand (***a,e***).

4.4 Architectures and solutions—summary

In the following table, we summarize the protection architectures presented earlier in this section. For each protection architecture and SRG type, we indicate the complexity of the lightpath provisioning operation (Polynomial or NP-complete), the speed of protection as well as the cost of the service expressed in amount of resources used by the protection mechanism (in percentage of effective capacity). As previously mentioned, *on-line routing* in an opaque network architecture that utilizes OEO switches is utilized for this work.

5 Experiments and Results

5.1 Resource utilization of each restoration architecture for centralized routing

All simulation experiments were run on two networks. N17 is a 17-node, 24-edge network that has a degree distribution of (8,6,1,2) nodes with respective degrees (2,3,4,5). N100 is a 100-node, 137-edge network that has a degree distribution of (50,28,20,2) nodes with respective degrees (2,3,4,5). These are realistic architectures representative of existing topologies. It is assumed that these architectures have infinite link capacity, and SRGs comprise exclusively of all the optical channels in individual links. Five network robustness scenarios are considered: no protection; dedicated protection with provisioning to recover from single link failures; dedicated protection with provisioning to recover from single link or node failures; shared mesh restoration to recover from link failure; and shared mesh restoration to recover from single node failures. In N17, demand is uniform, and consists of two bi-directional lightpaths between every pair of nodes. This amounts to 272 lightpaths. In N100, 3278 node-pairs out of 4950 possible node pairs are connected by 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 11 and Figure 12. The quantities shown on the charts are averages over a series of 10 experiments using various demand arrival orders. These results indicate that dedicated link-disjoint and node-disjoint approaches consume approximately the same total amount of capacity. This is expected since dedicated protection against link failures protects against node failures as well for nodes up to and including degree 3, and these nodes constitute a majority of all the nodes in the networks. This property does not apply to shared mesh restoration, and in lightpaths that must be protected against node failures, even for nodes of degree 3. Because of this and other reasons whose interpretation is not within the scope of this presentation, the provisioning of shared mesh restoration to recover from node failures requires (in comparison to the dedicated schemes) more resources than to recover from link failures. The relative difference however is negligible considering the benefits of protecting the network against node failures.

5.3 Centralized vs distributed routing

In the next set of experiments we compare centralized and distributed routing. We use the same network as before, but different order of demands. Here we assume that distributed routing has access to topological information and link utilization, but does not know about existing lightpaths, and thus cannot derive lightpath compatibility information. The reserved channels are thus

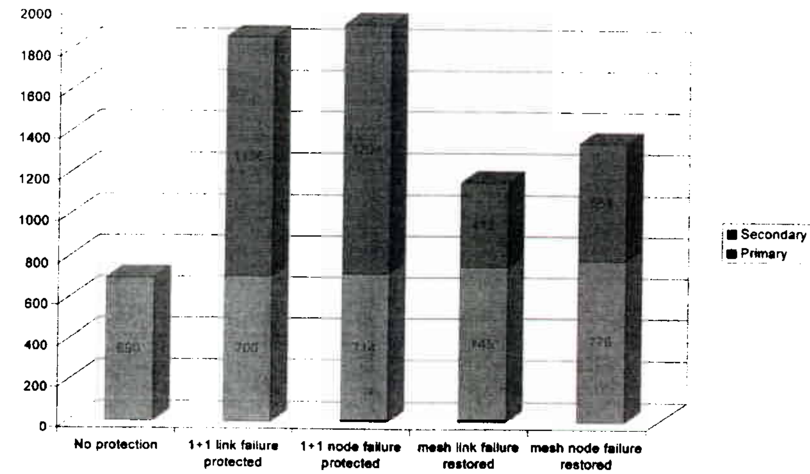

Figure 11: Comparison of capacity usage for different protection architectures (17 nodes).

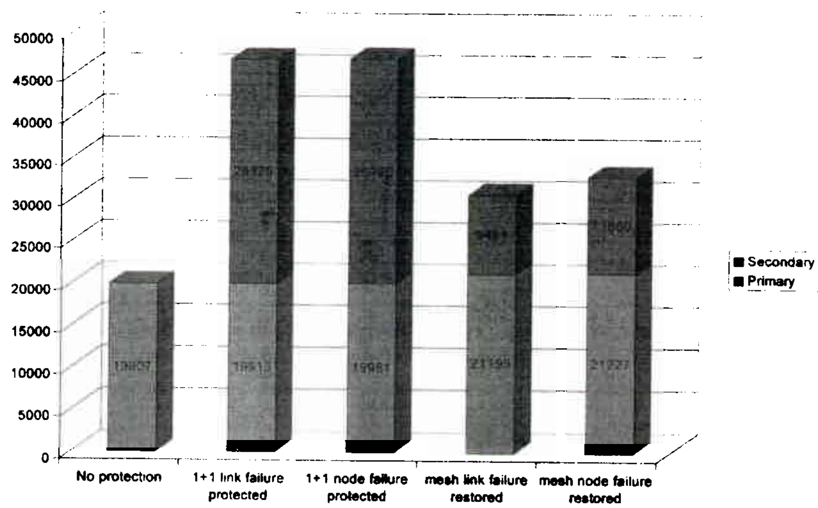

Figure 12: Comparison of capacity usage for different protection architectures (100 nodes).

assigned after the routes are computed—as opposed to centralized routing, where routes can be computed to maximize the sharing and minimize the allocation of new optical channels on primary and backup paths.

Results for the N17 network are shown in Figure 13. Experiments indicate that distributed routing of shared mesh restored lightpaths incurs a capacity penalty of 12% to 17% over centralized routing. Similar results for the N100 network can be observed in Figure 14.

Note that in shared mesh restoration, provisioning of protection paths sometimes requires longer paths that consist exclusively of reserved channels, rather than shorter paths where new channels must be reserved. Several experiments were performed to evaluate the effect of additional hops allowed on backup path over shortest-hop alternate path. Various tradeoffs between the back-up path length and the protection capacity for shared-mesh protection (by limiting the length of the back-up path, or changing the cost of using shared channels) are reported in [34].

6 Conclusion

In this paper we have discussed various on-line routing techniques to provision lightpaths in a layered architecture that uses dedicated control and management plane for each layer. Focusing our attention on the OXC-based restoration in the optical domain only, we reviewed and compared a set of protection and restoration architectures. In order to conduct these comparisons we first described multiple scenarios based on architecture characteristics (e.g., Shared Risk Groups), the type of failures for which protection was required (e.g., node or link) and whether the provisioning of lightpaths was centralized or distributed. We then compared available protection/restoration mechanisms taking into consideration figures of merit

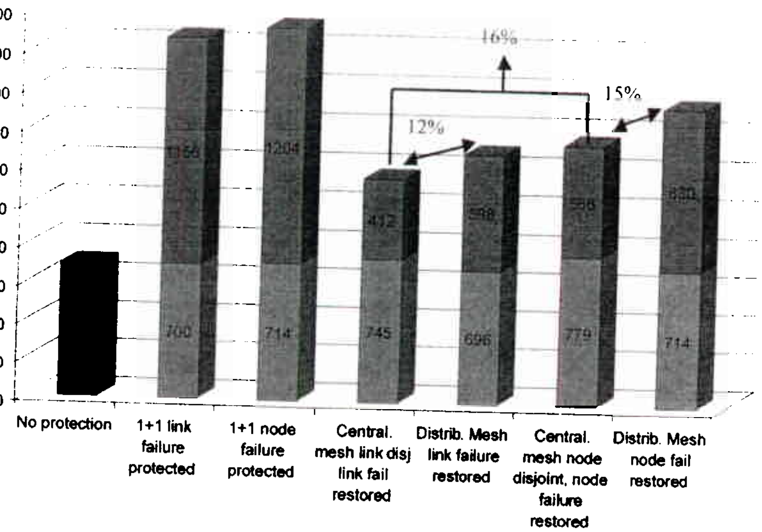

Figure 13: Comparison of centralized Vs distributed routing for different protection architectures (17 nodes).

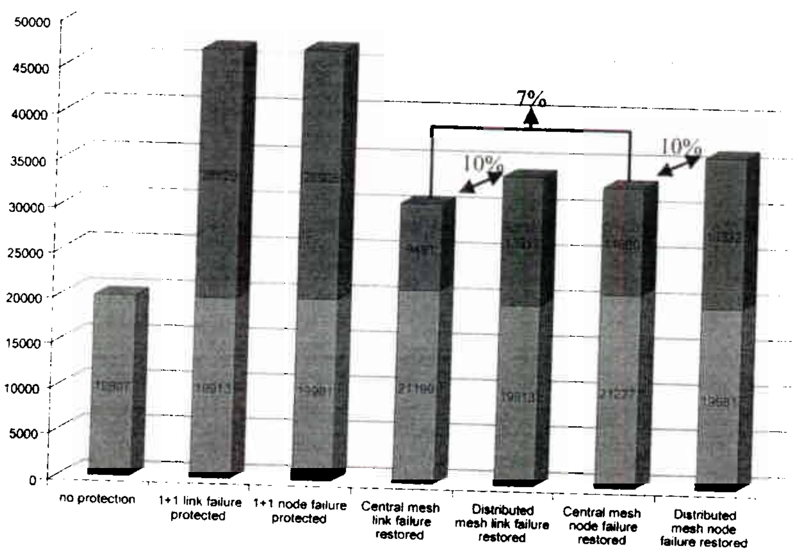

Figure 14: Comparison of centralized Vs distributed routing for different protection architectures (100 nodes).

such as redundant capacity and speed of restoration to sieve out the most appropriate protection schemes for each failure scenario. Experimental results show that the amount of resources used protecting node failures using the dedicated (1 + 1) protection mechanism is only marginally higher than (1 + 1) dedicated protection against SRG failures, and the difference may not justify the implementation of a dedicated protection mechanism that protects against SRG failures only. If minimizing the cost of the service is the main objective, then shared mesh restoration, which allows channel sharing among the restoration-routes, is more appropriate. Shared mesh restoration has also the additional advantage over dedicated protection that it promotes preemptible services that use idle restoration channels, giving way to higher priority services when a failure occurs. However, unlike dedicated protection, shared mesh restoration against node failures consumes substantially more resources than restoration against link failures. Thus, if shared mesh restoration is utilized, there is a trade-off in terms of enhanced network resilience versus network capacity and subsequently network cost.

Other important aspects that has been put aside in this paper, and are good candidate topics for future work in the area, are the constraint of round-trip delays and simulation of restoration times under various load conditions. For a fair comparison with other types of protection, it is important to limit the length of the restoration paths in order to respect prescribed round-trip delays and provide equivalent level of services on primary and restored paths. Our experiments indicate that restoration paths with unrestricted lengths may reserve an arbitrary large number of channels, although on the average the length does not increase much compared to the shortest-hop alternate restoration paths [36].

7 Appendix

7.1 Diverse routing with general SRGs

SRG Diverse Routing Problem: Given an optical network topology with general SRGs, and given node A and node Z, is there a feasible 1 + 1 SRG diverse route from A to Z?

SRG Diverse Routing Problem is NP-complete:

1) SRG Diverse Routing ∈ NP. Given a pair of diverse routes from node A to node Z we can ensure that their edges do not share any SRGs.
2) We can apply the 3-SAT problem expression [35] to the SRG Diverse Routing problem as follows: The 3-SAT boolean expression consists of a set of clauses $\{C_1, C_2, \ldots, C_N\}$ where each clause is a disjunction of 3 literals, e.g., $C_1 = \{x_1, \underline{x}_2, x_3\}$ ($\underline{x}$ indicates not x). The problem seeks to find an assignment of boolean values to the variables that satisfies all the clauses. Given a 3-SAT boolean expression we derive an optical network topology as indicated in Figure 15.

To each variable x, we associate two edges (each with available unassigned channels), one labeled x and one labeled $\underline{x}$, as illustrated in Figure 15. To each clause C = {x,y,z}, we associate three edges (each with available channels), labeled Cx, Cy, Cz. There is an SRG defined between edge labeled Cx, and $\underline{x}$. There is an SRG diverse route between A and Z iff the 3-SAT expression is satisfied. If there is an SRG disjoint pair of routes between A and Z, then one of the routes would have to traverse the top part of the graph (through edges labeled x), and the diverse route would have to traverse the bottom part (through edges labeled Cx). If the route traverses through the edge labeled x then the variable x is assigned the value of 1, and if it traverses through the edge labeled $\underline{x}$, the variable x is assigned the value of 0 (the route has to traverse one or the other edge). This assignment of boolean values to the variables must satisfy each clause because, given clause C = {x,y,z} the diverse route must traverse one of the edges Cx, Cy, or Cz. Say it traverses Cx, then the boolean assignment to variable x must satisfy clause C (because of the way SRGs are defined). If there is a satisfying assignment of boolean values to variables then there is an SRG diverse pair of routes. If variable x takes value 1, then the primary path traverses through edge x, and if x takes value 0, the primary path traverses through edge $\underline{x}$. For each clause C = {x,y,z}, the backup path traverses through one of the edges Cx,Cy,Cz, whichever is satisfied.

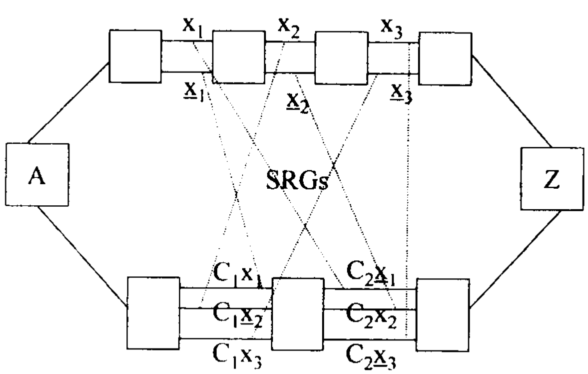

Figure 15: Optical network corresponding to the 3-SAT expression $C_1 = \{x_1, \underline{x}_2, x_3\}$, $C_2 = \{\underline{x}_1, x_2, \underline{x}_3\}$.

7.2 Shared mesh restored lightpath routing

Shared Mesh Restored Lightpath Routing Problem: Given an optical network topology (with default SRGs), and a set of already provisioned shared mesh restored lightpaths P, and given node A and node Z, is there a feasible shared mesh restored (edge-failure restored) lightpath from A to Z?

Shared Mesh Restored Lightpath Routing Problem is NP-complete:

1. Shared Mesh Restored Lightpath Routing ∈ NP. Given a feasible pair of primary and backup paths for a shared mesh restored lightpath, from node A to node Z we can ensure that the primary and backup paths are edge-diverse, and that the rules of sharing are not violated on the backup path.
2. We can apply the 3-SAT expression to Shared Mesh Restored Lightpath Routing as follows: The 3-SAT boolean expression consists of a set of clauses $\{C_1, C_2, \ldots, C_N\}$ where each clause is a disjunction of 3 literals, e.g., $C_1 = \{x_1, \underline{x}_2, x_3\}$ ($\underline{x}$ indicates not x). The problem seeks to find an assignment of boolean values to the variables that satisfies all the clauses. Given a 3-SAT boolean expression we derive an optical network topology as indicated in Figure 16.

To each variable x, we associate two edges (with available unassigned channels), one labeled x and one labeled $\underline{x}$, as illustrated in Figure 16. To each clause C = {x,y,z}, we associate three edges (with no available channels, and one shareable channel), labeled Cx, Cy, Cz. The set of provisioned shared mesh-restored lightpaths is defined as follows: For each clause C = {x,y,z}, there are three shared mesh-restored lightpaths defined, the first primary path has one edge $\underline{x}$, and backup path has three edges, the middle one of which is edge Cx. There is a feasible shared mesh-restored lightpath between A and Z iff the 3-SAT expression is satisfied. If there is a feasible shared mesh-restored lightpath between A and Z, then the primary route would have to traverse the top part of the graph (through edges labeled x), and the backup route would have to tra-

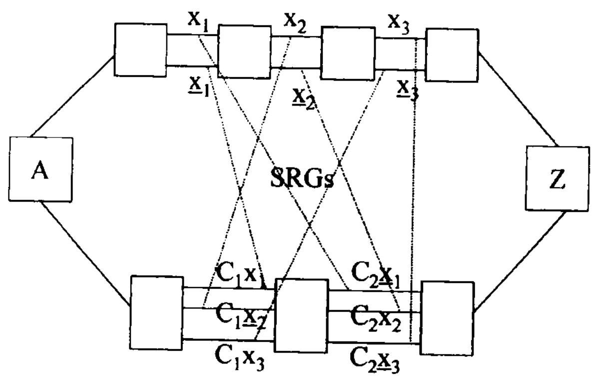

Figure 16: Optical network corresponding to the 3-SAT expression C1 = $\{x_1, \underline{x}_2, x_3\}$, C2 = $\{\underline{x}_1, x_2, \underline{x}_3\}$. Six lightpaths are already provisioned, P1 = $\{B\underline{x}_1C, BFC_1x_1GC\}$, P2 = $\{C\underline{x}_2D, CFC_1\underline{x}_2GD\}$, P3 = $\{Dx_3E, DFC_1x_3GE\}$, P4 = $\{Bx_1C, BGC_2\underline{x}_1HC\}$, P5 = $\{C\underline{x}_2D, CGC_2x_2HD\}$, P6 = $\{Dx_3E, DGC_2\underline{x}_3HE\}$.

verse the bottom part (through edges labeled Cx). If the route traverses through the edge labeled x then the variable x is assigned the value of 1, and if it traverses through the edge labeled $\underline{x}$, the variable x is assigned the value of 0 (the route has to traverse one or the other edge). This assignment of boolean values to the variables must satisfy each clause because, given clause C = {x,y,z} the diverse route must traverse one of the edges Cx, Cy, or Cz. Say it traverses Cx, then the boolean assignment to variable x must satisfy clause C (because there is already a provisioned path whose primary path uses edge $\underline{x}$, whose backup path uses edge Cx). If there is a satisfying assignment of boolean values to variables then there is a feasible shared mesh-restored lightpath from node A to node Z. If variable x takes value 1, then the primary path traverses through edge x, and if x takes value 0, the primary path traverses through edge $\underline{x}$. For each clause C = {x,y,z}, the backup path traverses through one of the edges Cx, Cy, Cz, whichever is satisfied.

In the common case, when diverse routes with available unassigned channels exist (and can be determined by standard diverse routing algorithms), a feasible shared mesh-restored lightpath exists as well (because in the worst case, all unassigned channels can be used on the backup path with no sharing). However, in the case that there is no 1 + 1 diverse path between node A and node Z the difficulty arises. This is because in this worst-case, the routing of already provisioned shared mesh-restored lightpaths can lead to complicated "sharing" situations which makes it very difficult (short of exhaustive exponential enumeration of primary paths) to determine if a feasible solution exists. As mention above, if there is at least one unassigned channel on each link, a feasible pair of (default) SRG diverse primary and backup routes can be found in polynomial time. However, in the optimization version, we want to minimize the total number of unassigned channels used. This can be shown to be NP-hard by reduction from MAX-SAT [35] through a similar construction.

8 References

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

[2] D. Benjamin, R. Trudel, S. Shew, and 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, and H. Epstein, "Architecting the Services Optical Network", IEEE Communications Magazine, pp. 80-87, September 2001.

[4] A. McGuire, S. Mirza, and 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, and 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 and 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] G. Ellinas, E. Bouillet, R. Ramamurthy, J. Labourdette, S. Chaudhuri, and K. Bala, "Restoration in Layered Architectures with a WDM Mesh Optical

Layer", Annual Review of Communications, International Engineering Consortium, June 2002.

[19] J. Strand, A. Chiu, and R. Tkach, "Issues for Routing in the Optical Layer", IEEE Communications Magazine, pp. 81-87, February 2001.

[20] Eric Bouillet, Jean-François Labourdette, Georgios Ellinas, and Ramu Ramamurthy, 'Stochastic Approaches to Compute Mesh Restored Lightpaths in DWDM Optical Network Architectures", Proc. Infocom'2002, New York City, NY, June 2002.

[21] S. Chaudhuri, R. Ramamurthy, and S. Sengupta, "Comparison of Centralized and Distributed Provisioning of Lightpaths in Optical Networks", Proc. OFC'01, Anaheim, CA, March 2001.

[22] 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", IEEE/OSA Journal of Lightwave Technology, Vol. 19, no 1, pp. 40-48, January 2001.

[23] R. Doverspike and J. Yates, "Challenges for MPLS in Optical Network Restoration", IEEE Communication Magazine, February 2001.

[24] C. W. Chao, P. M. Dollard, J. E. Weythman, L. T. Nguyen, and H. Eslambolchi, "FASTAR—A robust system for fast DS3 restoration", Proc. Globecom'91, 1991.

[25] W. D. Grover, in "Distributed restoration of the transport network", in Telecommunications Network Management into the 21st Century—Techniques, Standards, Technologies and Applications", S. Aidarous, T. Plevyak, IEEE Press, 1994.

[26] S. Datta, S. Sengupta, S. Biswas, and S. Datta, "Efficient Channel Reservation for Backup Paths in Optical Mesh Networks", IEEE Proc. Globecom'01, San Antonio, TX, November 2001.

[27] R. Iraschko, M. H. MacGregor, and W. D. Grover, "Optimal Capacity Placement for Path Restoration in STM or ATM Mesh Survivable Networks", IEEE/ACM Trans. on Networking, vol. 6, no 3, pp. 325-336, June 98.

[28] R. R. Iraschko, M. H. MacGregor, and W. D. Grover, "Optimal Capacity Placement for Path Restoration in Mesh Survivable Networks." In Proc. ICC'96, pp. 1568-1574, Dallas, TX, June 1996.

[29] S. Ramamaurthy and B. Mukherjee, "Survivable WDM mesh networks. Part I—Protection", IEEE Proc. Infocom'99, Vol. 2, pp. 744-751, 1999.

[30] S. Ramamaurthy and B. Mukherjee, "Survivable WDM mesh networks. Part II—Restoration", IEEE Proc. ICC '99, Vol. 3, pp. 2023-2030, 1999.

[31] J. W. Suurballe, "Disjoint Paths in a Network", Networks, Vol. 4, pp. 125-145, 1974.

[32] G. Chartrand and L. Lesniak, Graphs & Digraphs, Chapman and Hall, 1996.

[33] B. Doshi, *et al.*, "Optical Network Design and Restoration" Bell-Labs Technical Journal, Vol. 4, No. 1, pp. 58-84, January-March 1999.

[34] J. Labourdette, E. Bouillet, R. Ramamurthy, G. Ellinas, S. Chaudhuri, and K. Bala, "Routing Strategies for Capacity-Efficient and Fast-Restorable Mesh Optical Networks", Photonic Network Communications—special issue on "Routing, Protection, and Restoration Strategies and Algorithms for WDM Optical Networks", vol. 4, no. 3/4, pp. 219-235, 2002.

[35] M. Garey and D. Johnson, Computers and Intractability. A Guide to the Theory of NP-Completeness, W. H. Freeman and Company, 1979.

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

Georgios Ellinas
Tellium Inc.
2 Crescent Place
Oceanport, NJ 07757
Tel: 732-923-4100
Fax: 732-728-9862
Email: gellinas@tellium.com

Georgios Ellinas received the Ph.D. degree in Electrical Engineering from Columbia University, New York in 1998. From 1993 to 2000 he was also with the Optical Networking Research program at Telcordia Technologies, performing research for both the Optical Networks Technology Consortium (ONTC) and Multiwavelength Optical Networking (MONET) projects. In 1998, he joined Telcordia Technologies as a full-time Sr. Research Scientist in the Optical Networking Research group. In 1999 he served as an Adjunct Assistant Professor at Columbia University and in 2000 he served as an Adjunct Assistant Professor at the University of Maryland, teaching a course on Multiwavelength Optical Networking. He is currently a Sr. Network Architect with Tellium Inc., working on provisioning and restoration algorithms in mesh networks, as well as on the architecture design of Tellium's MEMS-based optical switch. His research interests are in the areas of optical architectures, routing and wavelength assignment algorithms, fault protection/restoration techniques and packet switching in multiwavelength optical networks.

Dr. Ellinas has authored and co-authored more than 40 journal and conference papers. He is also the holder of 9 U.S. Patents on optical networking and he has more than 25 patent applications currently pending.

Eric Bouillet
Tellium Inc.
2 Crescent Place
Oceanport, NJ 07757
Tel: 732-923-4100
Fax: 732-728-9862
Email: ebouillet@tellium.com

Dr. Eric Bouillet is a senior network architect at Tellium. In this role, he works on the design of optical networks and optimization of lightpath provisioning and fault restoration algorithms.

Before joining Tellium, Eric was a member of the technical staff in the Mathematical Sciences Research center in Bell Labs, Lucent Technologies, Murray Hill, New Jersey, working in the areas of routing in ATM and Optical networks.

Eric holds an MS and a Ph.D. in Electrical Engineering from Columbia University. He also holds a joint degree from l'Ecole Nationale Supárieure des Télécommunications ENST Paris and EURECOM Sophia Antipolis. He has authored and co-authored several journal and conference papers and has a number of U.S. and international patents pending in the area of optical networking.

Ramu Ramamurthy
Tellium Inc.
2 Crescent Place
Oceanport, NJ 07757
Tel: 732-923-4100
Fax: 732-728-9862
Email: ramu@tellium.com

Ramu Ramamurthy is a senior network architect at Tellium, where he works on the design of algorithms and protocols for dynamic provisioning and restoration in optical networks. Prior to joining Tellium, he was a research scientist at Telcordia Technologies, where he worked on network control and the management of IP/WDM optical networks. Ramu holds a B.Tech. degree from the Indian Institute of Technology, Madras, and MS and Ph.D. degrees from the University of California, Davis.

Jean-François Labourdette
Tellium Inc.
2 Crescent Place
Oceanport, NJ 07757
Tel: 732-923-4100
Fax: 732-728-9862
Email: jlabourdette@tellium.com

Dr. Jean-François Labourdette is manager of network routing and design at Tellium, a position he has held since October 2000. In this position, he is responsible for Tellium routing architecture and algorithms, and network dimensioning and design studies.

Previously, he was a manager of multi-service globalization planning at AT&T, where he was responsible for network, service and operations planning for AT&T international data services (Frame Relay, ATM, and MPLS IP VPN). Earlier at AT&T, he was a system engineer in the routing planning group, working on dynamic call routing for AT&T switched network and facility routing and re-arrangement for AT&T network.

Jean-François holds a Ph.D from Columbia University, where he was the recipient of the 1991 Eliahu I. Jury Award for best dissertation. He is a senior member of the IEEE and a member of the Optical Society of America.

Sid Chaudhuri
Tellium Inc.
2 Crescent Place
Oceanport, NJ 07757
Tel: 732-923-4100
Fax: 732-728-9862
Email: sc@tellium.com

Dr. Sid Chaudhuri is director of network architecture at Tellium, a position he has held since March 2000. In this role, he led the development of the mesh restoration and IP-centric control architecture implemented in the company's Aurora Optical Switch. He also led the development of Tellium's StarNet Software Design Tools.

Prior to joining Tellium, Sid worked for AT&T Laboratories Research where he developed intelligent optical network architectures using SONET, DWDM and optical switching technologies and spearheaded AT&T's Core Transport Network architecture. Previously, he was a distinguished member of the technical staff at Bell Laboratories, where he developed SONET/SDH transport and cross-connect system architecture.

Sid is president and chairman of the Optical Internetworking Forum (OIF), where he led the group in the first set of interoperability agreements.

Sid holds a Ph.D. from the University of Pittsburgh, where he was an Andrew W. Mellon fellow.

Krishna Bala
Tellium Inc.
2 Crescent Place
Oceanport, NJ 07757
Tel: 732-923-4100
Fax: 732-728-9862
Email: kbala@tellium.com

Dr. Krishna Bala is a co-founder and chief technical officer of Tellium. He is the lead system architect for Tellium's Aurora Optical Switch and the Full Spectrum All-Optical Switch. From 1997 to 1999, Krishna was Tellium's manager of optical cross-connect product development. He was a member of the team that launched Tellium's optical switching product line and established the feasibility of key technologies. Krishna is a pioneer in the field of optical switching and networking, in which he is the holder of several patents. He is also a co-author of the book "Multiwavelength Optical Networks: A Layered Approach." (Prentice Hall 1999).

Prior to joining Tellium, Krishna was a senior scientist in the Optical Networking Group of Bellcore, a research and telecommunications services provider, from 1992 to 1997. While at Bellcore, he was responsible for local exchange optical networking architecture development and analysis and received the Bellcore President's Award for outstanding achievement. Krishna holds a Ph.D. in Electrical Engineering from Columbia University. His thesis at Columbia was one of the first in the area of optical wavelength routing.

Krishna is chairman of the signaling working group of the Optical Internetworking Forum that is an industry group that enables the creation of optical networking standards.