

IP Restoration vs. WDM Protection: Is There an Optimal Choice?

Andrea Fumagalli and Luca Valcarenghi, University of Texas at Dallas

Abstract

The exponentially growing number of Internet users armed with emerging multimedia Internet applications is continuously thirsty for more network capacity. Wavelength-division multiplexing networks that directly support IP — the so-called IP over WDM architecture — have the appropriate characteristics to quench this bandwidth thirst. As everyday life increasingly relies on telecommunication services, users become more and more demanding, and connection reliability is currently as critical as high capacity. Both IP and WDM layers can fulfill this need by providing various resilient schemes to protect users' traffic from disruptions due to network faults. This article first reviews the most common restoration and protection schemes available at the IP and WDM layers. These schemes may be present concurrently in the IP over WDM architecture, with the resilient mechanism of each connection specifically chosen as a function of the overall cost, application requirements, and management complexity. The second part of the article describes a versatile heuristic based on simulated annealing that may be adopted to optimize the concurrent use of IP restoration and WDM protection schemes in the same (mesh) network. The proposed heuristic allows varying the percentage of traffic protected by the WDM layer and that of traffic relying on IP restoration, taking into account topology constraints and network cost minimization. An additional feature of the proposed heuristic is the potential to trade solution optimality for computational time, thus yielding fast solutions in support of interactive design.

The skyrocketing growth of data-centric traffic is rapidly depleting today's Internet resources. Internet service providers anticipate that Internet traffic will grow between 50 and 300 percent yearly. The aggregate bandwidth required in the backbone networks will shortly surpass hundreds of terabits per second. In this scenario optical networks will play a key role in providing novel, high-capacity, cost-effective transport networks compatible with existing transport and switching solutions, such as synchronous digital hierarchy (SDH)/synchronous optical network (SONET), asynchronous transfer mode (ATM), and Internet Protocol (IP).

In the optical layer (OL),¹ wavelength-division multiplexing (WDM) yields a cost-effective way to harness the fiber bandwidth by providing several parallel channels or *lightpaths*, each operating at a transmission rate compatible with current electronic rates.² A lightpath is an end-to-end path of light that may be established between a node pair (e.g., IP routers) that may not be physically adjacent (i.e., not connected by a single

fiber hop), thus increasing the logical network connectivity. To be of practical use, however, the advantages of high speed, high network connectivity, and optical transparency³ offered by a WDM network must be accompanied by adequate network survivability features.

In general, a network is referred to as *survivable* if it provides some ability to recover ongoing connections disrupted by the catastrophic failure of a network component, such as a line interruption or node failure. The design of survivable networks is obtained by exploiting *restoration* and/or *protection* schemes. Upon network failure, a restoration scheme dynamically looks for backup paths of spare capacity in the network. A protection scheme reserves, in advance, dedicated backup paths and wavelengths in the network. The former scheme is commonly available at higher layers (e.g., the IP layer). The latter scheme is commonly used at the transport (e.g., WDM) layer.

This work was partially supported by Alcatel CRC of Richardson (TX) and the TxTEC Consortium.

¹ The optical layer is defined by International Telecommunication Union.

² Dense WDM (DWDM) systems with transmission capacity of 1.6 Tb/s (160 channels at 10 Gb/s each) on a single fiber were recently announced.

³ Transparency is the capability of supporting multiple transmission rates and modulation formats.

This article addresses the issue of survivability in IP over WDM networks. The article is organized in two parts.

The first part presents a survey of available restoration and protection schemes that may be used in networks with arbitrary (mesh) topology.⁴ The advantages of the presented resilient schemes are discussed, leading to the conclusion that, very likely, IP over WDM networks will require multiple coexisting resilient schemes in the same network.

The second part of the article illustrates a heuristic approach that constitutes a framework for achieving combined optimization of the IP restoration and WDM protection schemes concurrently present in the same mesh networks. The heuristic is designed to optimally solve what we call the IP-aware wavelength minimum mileage (IWM) problem. The IWM problem consists of minimizing the overall network cost, that is, the wavelength mileage for both working and spare lightpaths in the OL plus the *resilience cost* of the IP layer. The latter cost may be associated with either a performance-related quantity — an equivalent cost representing the quality of service (QoS) degradation due to IP rerouting — or an abstract parameter representing the willingness of the designer to provide more or less reliability at either network layer. In the example discussed here, the IWM problem is solved by assuming that the traffic routed on any line of the mesh can be made resilient to faults by using either an IP restoration scheme or a WDM protection scheme. For each line, the proposed heuristic chooses the resilient scheme that overall yields the minimum network cost. The heuristic is based on the combined use of:

- A simulated annealing (SA) algorithm designed to select which lines of the mesh are protected at the optical layer (the others require IP restoration)
- A shortest path algorithm applied to route the lightpath demands on an auxiliary graph, which is derived from the mesh and takes into account the cost of both restoration and protection schemes

Among other advantages, the SA algorithm allows us to exploit the trade-off between solution optimality and computational speed. The optimization framework is thus suitable for the interactive design of survivable IP over WDM mesh networks. In addition, due to SA versatility, the cost function under optimization can easily be modified by the designer to take into account various protection schemes. For the sake of demonstration, two protection schemes are considered in the article: the dedicated-path-switched WDM self-healing ring (DP-WSHR) and the shared-line-switched WDM self-healing ring (SL-WSHR), showing that in either case cost-effective solutions are found with the proposed heuristic. By means of the auxiliary graph, it is possible to optimally deal with design constraints that may arise due to network layout, management complexity, and/or technology limitations. Indeed, the percentage of traffic and network lines that are protected in either layer may be arbitrarily chosen by the designer, thus providing a flexible approach that can address a variety of distinct design scenarios.

Resilient Schemes at the IP and WDM Layers

In the two-layer IP over WDM architecture, each layer can provide its own independent resilient scheme. Restoration/protection schemes may be implemented at the IP layer (possibly using the multiprotocol label switching, MPLS, sublayer) and the optical (WDM) layer (OL).

While the resilience concept is naturally embedded in the IP layer, in which the actual path used to route a packet from source to destination is dynamically found and maintained by the routers, protection techniques at the optical layer are just emerging, driven by the need for coarse and fast protection schemes. One can notice the increasing importance of WDM layer resilient schemes by the amount of work recently published on this topic [1–7].

In this section we review the most common restoration and protection schemes, discussing their advantages and disadvantages.

IP Layer Restoration Schemes

In general, resilient schemes available at the network layer, such as IP (IP/MPLS), have the capability to recover multiple faults and operate at small traffic granularity.⁵ However, these schemes are generally slow as they require online processing upon failure occurrence.

Dynamic routing and MPLS protection switching [8] are schemes currently considered to achieve network survivability at the IP (IP/MPLS) layer. These schemes are discussed in the next sections.

IP Dynamic Routing — With dynamic routing, reachable active routers are found dynamically, thus adapting IP routing to possible network faults. This task is accomplished by exchanging, between adjacent routers, control messages that are used to update the routers' routing tables, thus enabling IP packets to be dynamically rerouted around link and node failures. This protocol guarantees networkwide survivability, independent of the underlying physical network.

The faults in the network can be detected by the routers either explicitly or implicitly [9]. Explicitly, faults are detected at the local level and signaled to the neighboring routers through regular exchange of routing protocol control messages, such as Internet Control Message Protocol (ICMP). Otherwise, the expiration of timers on background messaging, such as KEEPALIVE and HELLO messages, implicitly signals the presence of faults. Once a router detects a line fault, it recalculates the affected routes and updates its routing tables. Then it propagates the occurred changes to its neighboring routers using UPDATE messages such as Open Shortest Path First (OSPF) link state advertisement (LSA) or Border Gateway Protocol-4 (BGP-4). The dynamic routing protocol efficiently uses spare network resources and is flexible to topological changes. However, it is usually slow, from tens of seconds to minutes, and its behavior is unpredictable.

Some enhancements of the protocol have been proposed [10] to overcome the former drawbacks.


One approach is equal cost multipath forwarding (ECMF), in which the router relies on more than one path for transmitting packets sharing a common destination. This solution is already implemented in some routing protocols, such as OSPF and Intermediate System-Intermediate System (IS-IS), by maintaining multiple next-hop entries for the same destination within each router's routing table. This way, in case of failure, a fraction of packets are guaranteed to flow to the destination until the router routing table is updated with the recalculated routes.

A second approach partitions the network into multiple areas, as defined in hierarchical link state routing protocols such as OSPF and IS-IS. This way routing updates are confined to the affected area, minimizing the network reconfiguration convergence time.

Another proposed approach decreases the failure detection time by either reducing the HELLO timeout, so HELLO mes-

⁴ Arbitrary mesh topologies are expected to be the customary fabric of backbone and metro networks due to their high connectivity and design flexibility.

⁵ Granularity is determined by the protocol traffic unit at that particular layer.



■ Figure 1. Classification of optical layer (WDM) resilience schemes. Acronyms in parenthesis are equivalent terms for a particular resilience scheme.

sages are sent more frequently, or rapid rate pinging through ICMP ECHO requests [11]. In either case, this approach increases control traffic overhead.

MPLS Protection Switching — MPLS protection switching is an alternative approach to circumvent the latency drawback of dynamic routing. MPLS protection switching is enabled through a hierarchy of paths called label-switched paths (LSPs) obtained by prepending a stack of labels or tags to packet headers. The protection entities can be set up either dynamically or in a prenegotiated way.

Protection entities, dynamically set up, restore traffic based on failure information, bandwidth allocation, and optimized reroute assignment. LSPs crossing a failed line or label-switched router (LSR) are disrupted at the point of failure and must be reestablished using reservation signaling.

Pre-negotiated protection consists of working LSPs that have preestablished protection paths. The preestablished protection path is node and link disjoint from the working path. The network resources of the protection path can be either reserved beforehand (unused unless failure occurs), or dynamically allocated to low-priority traffic that is allowed to use them in the absence of network failure. Preemption of low-priority traffic may be necessary upon network failure.

Both MPLS protection switching techniques, dynamic and pre-negotiated, can be performed on a line basis (by rerouting a portion of the LSP around the failed line, or *link rerouting*) or a path basis (by rerouting the entire LSP, or *edge-to-edge rerouting*). In general, dynamic protection increases resource utilization but requires longer restoration times than preestablished protection. Link rerouting is faster than edge-to-edge

rerouting because, in the latter case, the failure notification must reach the head-ends of all the LSPs affected by the failure. Link rerouting may be completed in about 50 ms but is not well suited for handling node failure.


WDM-Layer Resilient Schemes

Both the optical channel (OCh) section and optical multiplex section (OMS) — OL sublayers — feature *dynamic restoration* and *preplanned protection*. Like in the higher layers, restoration schemes are more efficient from a capacity viewpoint but relatively slow (common restoration completion time is on the order of seconds or minutes). Protection schemes guarantee service restoration completion times of hundreds, tens, and even fractions of milliseconds.

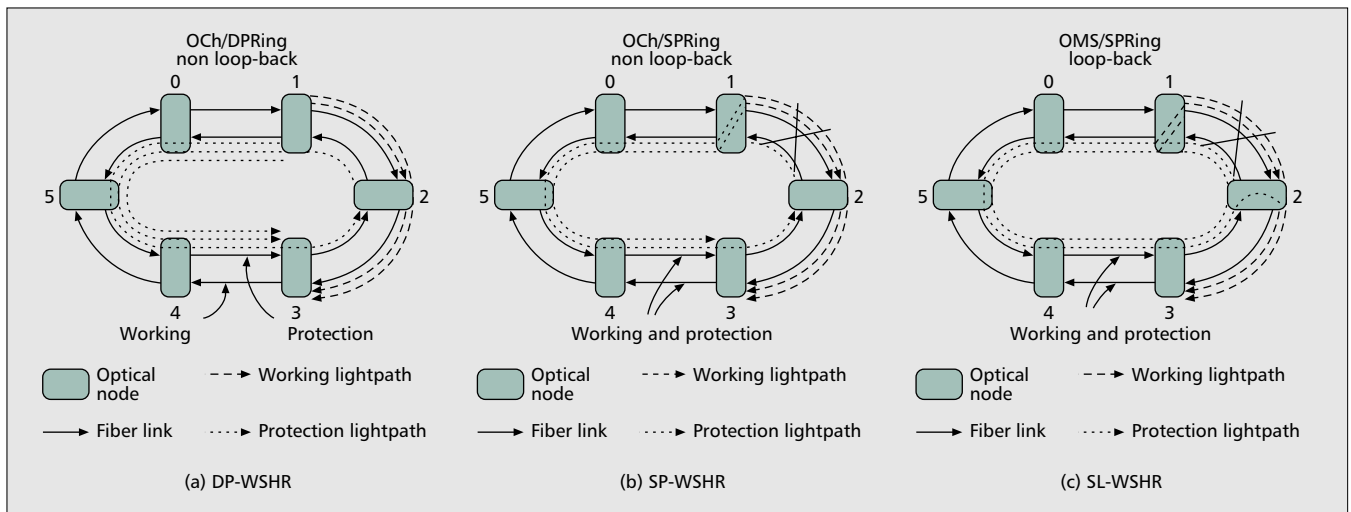
The main difference between OCh and OMS resilient schemes is represented by the granularity at which the layers operate. OCh resilient schemes protect individual lightpaths, thus allowing selective recovery of optical line terminal (OLT) failures. OMS resilient schemes work at the aggregate signal level, thus recovering all lightpaths present on the failed line concurrently [12]. OCh resilient schemes require more spare devices (e.g., transponders, multiplexers, demultiplexers) than their OMS counterparts because each working lightpath must be demultiplexed and multiplexed at every node in order to be individually switched [12].

In this section we provide a brief survey of the most common OL restoration and protection schemes summarized in Fig. 1. For further details the interested reader is referred to [3, 12].

The following terminology is used in this article. A *working lightpath* is the default path of light established between a source-destination pair absent a network failure. A *protection*



■ Figure 2. Protection wavelength allocation in WDM mesh networks: a) DPP; b) SPP; and c) multi-WSHR covering.



■ Figure 3. Protection wavelength allocation in WDM self-healing rings; a) DP-WSHR; b) SP-WSHR; c) SL-WSHR.

wavelength is a spare channel that may be used in case of network failure. A concatenation of protection wavelengths along a path forms a *protection lightpath*.

Optical Layer Restoration — The OCh path restoration scheme requires that, upon network failure, every affected working lightpath be replaced by a protection lightpath. The protection lightpath is computed using either a centralized [12] or distributed [3, 13] approach. With the former approach a control node keeps record of the network status, computes the protection lightpaths, and passes the information onto the network nodes. With the latter approach, source and destination nodes dynamically search for the protection wavelengths required to reestablish the disrupted lightpaths.

The OMS line restoration scheme requires that an alternative route be found to locally divert all affected working lightpaths around the faulty line. This scheme is triggered by the failed line end nodes which make use of a distributed algorithm to dynamically discover the alternative route.

Optical Layer Protection — OCh and OMS protection schemes are also called, respectively, *path protection* and *line protection*. In both schemes, protection (or spare) resources can be either *dedicated*, in which case the spare resource is dedicated to a single working lightpath, or *shared*, in which case the same spare resource may be used to provide protection to multiple working lightpaths. OCh and OMS protection schemes are available in *mesh* and *ring* network topologies.

Dedicated protection is commonly referred to as 1 + 1 or 1:1 protection. In 1+1 protection the source node transmits on both the working and protection lightpaths simultaneously. The destination keeps monitoring both lightpaths, dynamically choosing the signal with the best performance (e.g., signal-to-noise ratio). Therefore, if degradation of the signal is detected on the working lightpath, the destination immediately switches over the protection lightpath. In 1:1 protection, transmission occurs on the working lightpath only, while the protection lightpath may be used to transmit low-priority traffic. Upon failure of the working lightpath, both the source and destination switch over the protection lightpath, preempting the low-priority traffic.

Shared protection, also referred to as 1:N, allows spare wavelengths to be shared by a number of working lightpaths. In case of a fault, each disrupted transmission is switched over the protection wavelengths. This operation requires some signaling to notify the network nodes of the new transmission paths and to make sure that the protection wavelengths on

the various fibers are correctly interconnected to form the required protection lightpaths. Once the spare resource is used to protect a working lightpath, it will not be available to protect the other working lightpaths until the original working lightpath is reestablished.

Specific protection schemes based on these concepts are described below.

Dedicated path protection (DPP), depicted in Fig. 2a, is typically applied to mesh networks. DPP can be fast (1 + 1 DPP just requires that the receiving node switch to the protection lightpath) and robust in the face of multiple faults (they must not occur simultaneously in lines belonging to the working and protection lightpath of the same connection). It requires a low degree of management complexity, but does not efficiently use protection resources.

Dedicated line protection (DLP) reserves protection wavelengths between the end nodes of each line utilized by working lightpaths. Similar to OCh DPP, DLP may require more spare capacity allocation than other schemes, but can be even faster than DPP. Indeed, in the case of 1:1 protection, the restoration completion time is faster than 1:1 DPP because signaling is confined within the area around the faulty line.

Dedicated-path-switched WSHR (DP-WSHR) — also called *optical unidirectional path-switched ring (OUPSR)* or *OCh dedicated protection ring (OCh/DPRing)* — is the DPP equivalent for ring networks (Fig. 3a). DP-WSHR restoration is fast (on the order of a millisecond or a fraction of a millisecond), but the total protection wavelength mileage is equal to or larger than that required in the other OCh and OMS ring protection schemes.

Optical unidirectional line-switched ring (OULSR) is similar to DP-WSHR. It utilizes two counter-rotating fibers, one for working lightpaths and the other for protection lightpaths. The difference is that, in the OULSR case, all the lightpaths passing through the failed line are jointly switched over the protection fiber. With respect to DP-WSHR, this scheme requires the same wavelength mileage but, due to its line switching, employs less expensive devices while achieving similar restoration completion times.

In **shared-path protection (SPP)** (also termed 1:N path protection), protection wavelengths are shared by a number of line and node-disjoint working lightpaths (Fig. 2b). SPP achieves more efficient utilization of spare resources than DPP, at the price of more complex control (signaling is required between source and destination nodes) and longer restoration completion time (on the order of 100 ms). In addition, if two or more faulty connections share the same protection wavelengths, only one can be recovered.

Shared-line protection (SLP) (also termed 1:N line protection) applies the SPP technique locally to the faulty line. In this case, it is possible to make use of shared protection resources among different failure scenarios, thus yielding better resource utilization than achieved in DLP. The restoration completion time is generally faster than that of SPP because of locally limited signaling.

Shared-path WSHR (SP-WSHR) — also termed *optical bidirectional path-switched ring (OBPSR)* or *OCh shared protection ring (OCh/SPRing)* — represents the SPP equivalent in ring networks. This scheme's peculiarity is *nonloopback* switching, depicted in Fig. 3b. In SP-WSHR, in case of failure, each working lightpath is switched to the protection lightpath at its source node. Therefore, the recovered traffic reaches the destination node only along the protection lightpath. SP-

WSHR is the most efficient among the WSHR protection techniques in terms of spare resource utilization, but it requires complex control and signaling. Restoration time may also be affected for the same reason.

Bidirectional **shared line-switched WSHR (SL-WSHR)** is physically implemented with either two fibers (optical two-fiber/BLSR, O-2F/BLSR) or four fibers (optical four-fiber/BLSR, O-4F/BLSR). In either case, working lightpaths and protection wavelengths may be carried using both directions of propagation. Its peculiarity is *loopback* switching, depicted in Fig. 3c. Upon failure, the working lightpaths are switched, at one failure end, to the protection wavelengths of the counter-rotating fiber. When they reach the other failure end they are looped back along their original working wavelengths to reach their destination nodes. For each direction of propagation the number of protection wavelengths is determined by the largest number of working lightpaths in any line flowing in the opposite direction (defined as the *ring load*). The SL-WSHR scheme is simple and fast (completion switching time on the order of tens of milliseconds) because the switching mechanism is promptly activated upon fault detection without requiring any further signaling.

Coordination between IP and WDM Layers

Coordination between resilient schemes, possibly at distinct layers, is required to avoid multiple schemes concurrently activated upon a single network fault. This coordination is commonly achieved by resorting to *escalation strategies* that sequentially activate the different resilient schemes starting from either the lowest or highest network layer. Escalation strategies are governed by either explicit messaging between the different layers or arbitrarily setting failure detection and restoration completion times.

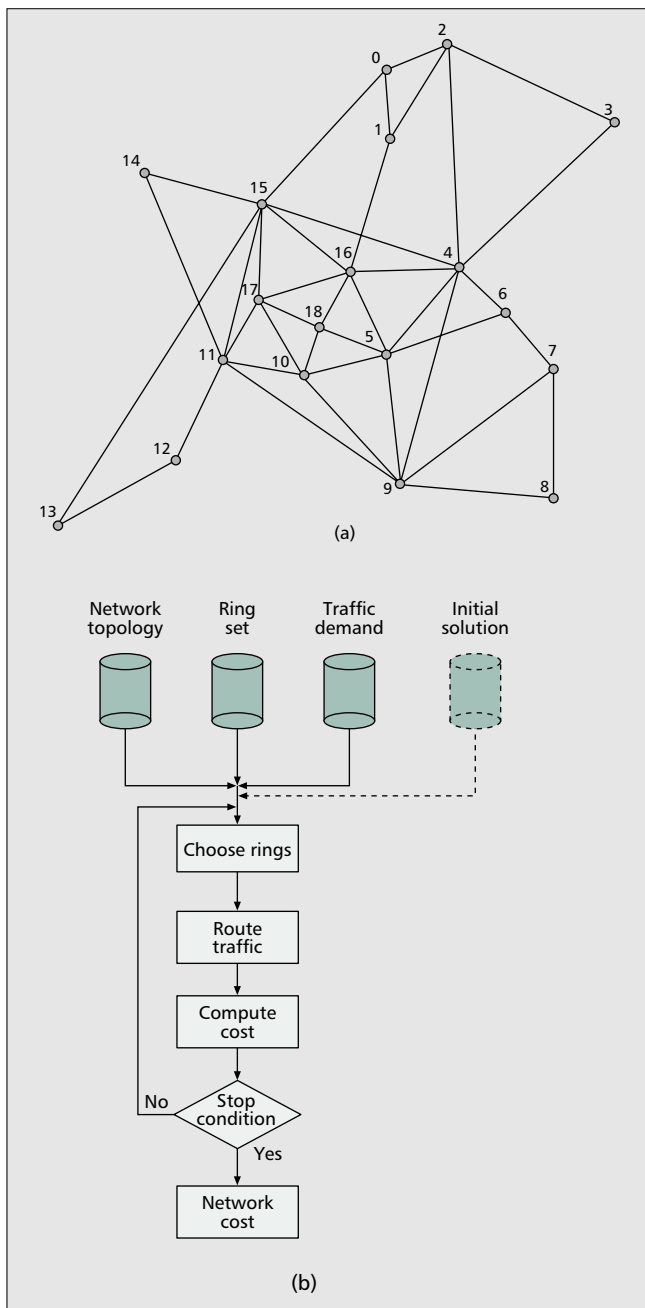
Table 1 summarizes some of the fault detection and restoration completion times at the IP/MPLS and WDM layers. It must be noted that MPLS resilient schemes are not yet mature nor supported by experimental data. Similarly, OCh/OMS restoration schemes are not yet standardized. On the contrary, IP dynamic routing and OCh/OMS protection techniques are already either used in commercial networks or tested in field trials. Their characteristics are thus well known.

It is worth commenting on the emerging, promising MPLS schemes. Fault detection in MPLS can be achieved by exploiting different existing techniques, such as SONET/SDH-based framing alarms, OL fault detection techniques, and IP-based KEEPALIVE, HELLO, and ICMP messages, which guarantee detection times ranging from a few tens of milliseconds to hundreds of seconds. MPLS may resort to control plane signaling, such as Constraint-Based Routing Label Distribution Protocol (CR-LDP) or modified Resource Reservation Protocol (RSVP), to perform traffic restoration with time intervals ranging from a few dozen milliseconds to hundreds of seconds.

In the rest of this article we focus on well-established IP and WDM resilient schemes — IP dynamic routing and OCh/OMS WSHR protection — and illustrate a framework for jointly designing the resilient schemes at both layers.

Survivable IP over WDM Mesh Networks Based on WSHR

Mesh topologies offer high connectivity which greatly improves network reliability and load balancing for congestion control in the Internet. The connectivity degree can be increased further by means of lightpaths established at the optical layer. Although current IP/MPLS reliability schemes



■ Figure 4. The pan-European network and solution approach.

Scheme	Detection time	Restoration completion time
IP dynamic routing	100 ms–180 s	1–100 s
MPLS fast (link) rerouting	0.1–100 ms	50–100 ms
MPLS edge-to-edge rerouting	100 ms–180 s	1–100 s
OCh and OMS restoration	> 100ms	≥ 50 ms
DP-WSHR, DL-WSHR, DPP	1–10 ms	10 μs–10 ms
SP-WSHR, SL-WSHR, SPP	1–10 ms	1–100 ms

■ Table 1. Service restoration times of IP and WDM resilience techniques [3, 10, 11].

offer varying levels of restoration/protection granularity, they cannot provide all the necessary functionality needed by the next-generation Internet (i.e., fast restoration). It is thus expected that both IP and WDM layers will each provide some degree of survivability against network faults.

An optical mesh example is represented by the pan-European backbone network, shown in Fig. 4a. Nodes are equipped with OXCs consisting of opto-mechanical switches, and wavelength conversion capabilities. OXCs provide the routing capabilities for establishing lightpaths throughout the network, while the opto-mechanical switches provide the survivability functions at the OMS sublayer. Nodes are connected by bidirectional optical fiber lines, with the number of wavelengths in each direction determined by the traffic demands and the protection scheme being used.

For demonstration purposes, the optical protection schemes considered in the rest of the article are WSHR-based. A WSHR protection scheme can be adopted in a mesh network by overlaying a set of rings on the physical mesh topology. The set of rings, also called the *ring cover*, is then used to design multiple WSHRs that protect the traffic routed on the covered lines. We term this approach, illustrated in Fig. 2c, the multi-WSHR scheme.

A complete covering of the mesh is possible when the topology is 2-connected,⁶ since this property ensures that every node in the mesh can be connected to at least one ring. When using a multi-WSHR scheme, the protection scheme of a single WSHR is applied to each individual ring that covers the mesh, with the following additional constraints:

- Every working lightpath carried on a line covered (protected) by two or more rings must be protected by only one ring.
- Rings can share fibers but not protection wavelengths.

An important feature of multi-WSHR is the fact that each ring is managed independent of the others. Upon a fault, the required signaling and control is thus confined within the faulty ring. The multi-WSHR scheme is resilient to multiple faults occurring on lines covered by distinct rings. Furthermore, since working lightpaths can span over multiple rings, only the nodes in which working lightpaths are switched from one protection ring to another require OXC capabilities. The remaining nodes need only optical add/drop multiplexing capabilities.

The IP-Aware Wavelength Minimum Mileage Problem and Solution

The availability of two resilient schemes in the same network — IP dynamic routing and WSHR protection — poses the question of which scheme must be used to protect which

traffic. The option of using both resilient schemes for all traffic may not be the most cost-effective. A more careful design may consist of a hybrid solution in which only part of the traffic is protected by WSHR, while the remaining part relies only on IP restoration.

An efficient heuristic is described to determine how both resilient schemes may be optimally combined to provide a cost-effective survivable network. The problem solved by the heuristic is called IP-aware wavelength minimum mileage (IWM) problem. In the IWM problem the total wavelength mileage in the WDM layer is minimized jointly with an arbitrarily defined IP layer cost. The cost of the

IP layer may be defined to take into account various factors, including performance degradation determined by the IP restoration scheme, such as traffic disruption and congestion originating from the interaction of rerouted traffic with regular traffic.

More precisely, the WDM layer cost is defined as the total wavelength mileage (i.e., miles of working and protection wavelengths in the WSHRs) required in the network to support the given set of traffic demands. The IP layer cost is defined as the sum of the wavelength mileage of the traffic streams that are not protected by the WSHR scheme and a penalty factor, $K_{IP}^{s,d}(m, n)$. The penalty factor $K_{IP}^{s,d}(m, n)$ is normalized to the cost of one wavelength mile.

In its most general form, penalty $K_{IP}^{s,d}(m, n)$ is a function of the particular lightpath, identified by the source, s , and destination, d , of the network line, identified by the line end node pair (m, n) . High penalties may be assigned to special lightpaths (lightpaths with particular requirements, e.g., maximum delay, minimum required bandwidth) or to particular network lines.

Alternatively, penalty $K_{IP}^{s,d}(m, n)$ may be used to enforce a specific type of resilient scheme on a selected set of network lines independent of the routed lightpaths, in which case it may be simplified as $K_{IP}(m, n)$.

Penalty $K_{IP}^{s,d}(m, n)$ may also be utilized to determine the fraction of traffic that must be protected by the WSHR scheme, independent of the lightpath and network line under consideration. In this case the penalty is simply characterized by K_{IP} , which becomes a tuning factor the designer may use to achieve the intended distribution of resilient schemes in the network. This application is further considered in the article.

In solving the IWM problem, we assume that the designer defines a matrix of traffic demands in terms of lightpaths between node pairs. The traffic demands are based on the estimated network load. For the WDM layer, two WSHR protection schemes are considered: DP-WSHR and SL-WSHR. Practical design constraints, such as maximum ring size, and maximum number of rings covering the same line and the same node, must also be taken into account in order to match available technology performance and desired network management complexity.

The IWM problem consists of three subproblems:

- The RC subproblem: Select the (multi-WSHR) ring cover that will provide the protection scheme in the optical layer.
- The WL subproblem: Select the working lightpath(s) for each traffic demand.
- The SW subproblem: Select which ring protects which working lightpath, and decide which part of the lightpath is protected by the optical and which by the IP layer

Applying Simulated Annealing to Solving the IWM Problem — The survivable mesh network under design is represented by a graph $G(V, E)$ in which V is the set of vertices or nodes of the network and E is the set of directed arcs representing the network lines connecting the nodes. The solution of the IWM

⁶ A graph is 2-connected if it cannot be disconnected by removing less than two edges from the graph.

ILP			SA			
# of paths	Time	λ -miles	Average time (s)	Minimum λ -miles	Average λ -miles	Maximum λ -miles
1 (SP)	1 min	5,711,890	57	5,155,140	5,280,980	5,403,890
2	2 min	5,037,850	124	5,042,390	5,172,120	5,282,900
8	4 min	5,030,680	223	4,959,940	5,100,360	5,225,400
8	5 min	5,018,250	283	4,949,780	5,082,670	5,161,220
8	5 hr	4,515,350	18331	4,823,110	4,848,464	4,908,660

■ Table 2. ILP results vs. SA results for the pan-European network assuming nonuniform lightpath demands and SL-WSHR.

problem consists of choosing, from a given set of candidate rings $C = \{c_1, \dots, c_i, \dots, c_R\}$ a subset $C' = \{c'_1, \dots, c'_i, \dots, c'_m\}$ that minimizes the cost function, defined as

$$f_{tot}(d_{mn}, \lambda_{mn}^{sd}), \quad (1)$$

where λ_{mn}^{sd} indicates the number of lightpaths from node s to node d routed on arc $(m, n) \in E$, and d_{mn} is the length of arc $(m, n) \in E$.

The equation representing the cost function f_{tot} may have different formulations, depending on the WSHR protection technique under consideration. Its general expression consists of one cost, f_w , that represents the total cost (working+protection) of optically protected lightpaths and another cost, f_{IP} , representing the cost of the lightpaths protected at the IP layer:

$$f_{tot} = f_w + f_{IP}. \quad (2)$$

Function f_w depends on the specific WSHR technique utilized in the multi-WSHR protection scheme and assumes the two formulations, respectively for DP-WSHR and SL-WSHR, computed (as f_1) in [5].

Function f_{IP} is calculated as follows:

$$f_{IP} = \sum_{(m,n) \text{ not OL protected}} \sum_{sd} (\lambda_{mn}^{sd} (d_{mn} + K_{IP} \cdot p)), \quad (3)$$

where p is a normalization factor that depends on the physical network characteristics (in our example $p = \sum_{m,n} d_{mn}$).

The algorithm designed to solve the IWM problem operates as shown in the flow chart of Fig. 4b. Input to the algorithm consists of the network topology $G(V, E)$, the set of candidate rings C , and the demand matrix $D(i, j)$, which denotes the number of lightpath demands required for each node pair. An initial solution consisting of a subset of rings C_0 (subset of C) may be defined by the user. This initial solution is used by the algorithm to start the optimization process and performs incremental runs of the algorithm to improve previously found solutions. If not otherwise specified, C_0 is empty.

The basic approach of the optimization algorithm consists of selecting a subset of rings from set C along which the working lightpaths are routed with the aim of minimizing cost function f_{tot} . The routing is performed utilizing an auxiliary graph called the *ring interconnection graph* (RIG) [5] which indicates the various routing costs, that is, the cost of the optical protection when a ring is available on the chosen line or the cost of the IP restoration, determined by K_{IP} , when a ring is not available on the chosen line. The cost of the chosen solution is computed and used to determine the suitability of the chosen subset of rings. This operation is repeated until a stop condition is met (e.g., expiration of the allowed CPU time for the optimization process). The returned solution is the one with the lowest cost among the solutions explored by the algorithm.

More specifically, to select the rings of the cover an SA approach is used [5]. Other heuristic approaches, such as Genetic Algorithm or Tabu Search, may be equally suitable to perform this task.

While selecting the ring cover, some practical (user-defined) design constraints are also taken into account:

- N_r : the maximum number of rings allowed to cover the same link
- N_v : the maximum number of rings allowed to connect a node

Only subsets of rings that satisfy the specified constraints are considered by the SA algorithm.

Numerical Results

Some numerical results are discussed to illustrate the features of the proposed heuristic approach, that is, the possibility:

- To trade solution optimality for solution time
- To optimize different WSHR protection techniques
- To study the distribution of the optimal partition of the network lines between WDM protection and IP restoration

The results are obtained after carefully tuning some of the SA parameters. A complete description of the tuning procedure can be found in [5]. The advantage of using the SA heuristic when solution time is limited is demonstrated by the results summarized in Table 2. These experiments consider the application of a multi-WSHR scheme based on SL-WSHRs to the pan-European backbone network depicted in Fig. 4a. The maximum number of nodes allowed in the rings is 6, which gives a set of 63 candidate rings. The lightpath demands are symmetric but nonuniform [5], and all traffic is protected in the OL (i.e., K_{IP} is infinite) with $N_r = 2$ and $N_v = 4$. The computational time is limited from a few minutes to 5 hr (on a Pentium 200 MHz machine). Table 2 compares the results obtained by both an integer linear programming (ILP) formulation presented in [6] and the SA approach. The values in the leftmost column indicate the number of alternative paths between each node pair from which the ILP solver may choose to optimally route the working lightpaths.⁷ The SA heuristic outperforms the ILP solver's best found result when computational time is limited to one minute or less.⁸

The versatility of the proposed heuristic approach is shown in Table 3. The setup for these experiments is the same as utilized previously. The last three columns report the *total fiber mileage* required to guarantee optical protection of all traffic. The fiber mileage is simply obtained from the wavelength mileage cost on a link-by-link basis, and depends on the number of wavelengths per fiber (eight in the experiment) and the particular WSHR scheme being utilized. Three different WSHR schemes have been considered: DP-WSHR, shared fiber SL-WSHR, and non-shared-fiber SL-WSHR. The last two schemes are conceptually similar to O-2F/BLSR and O-4F/BLSR, respectively, with a few significant differences: the number of fibers in the ring is not strictly limited to two or four, and the fibers of a line may be shared by multiple rings. Results indicate that multi-WSHR based on DP-WSHR is approximately 15 percent more expensive than SL-WSHR. The shared fiber and non-shared SL-WSHR solutions require the same wavelength mileage, but the

⁷ Value 1 means that there is only one path, which corresponds to solving the WL subproblem first, and then the RC and SW subproblems.

⁸ The ILP solver solutions summarized in Table 2 are the best found solutions in the limited computational time and may not correspond to the optimal problem solution.

former needs 3 percent less fiber mileage.

Table 4 demonstrates that by varying the value of K_{IP} , it is possible to shift the percentage of working lightpaths and network lines protected, respectively, at the IP and WDM layers. These simulations are run on the network depicted in Fig. 2c characterized by line lengths varying from a few hundreds to thousands of miles.

A symmetric and uniform traffic demand of five lightpaths between every node pair is considered. Physical constraints $N_r = 2$ and $N_v = 4$ are assumed. From the last column of Table 4 we see that, for decreasing values of K_{IP} (i.e., decreasing cost of IP restoration), the percentage of working wavelength miles protected at the IP layer increases. This trend continues until all lightpaths are not protected at the optical layer and fully rely on IP layer resilience capabilities. For $K_{IP} = 0$, the restoration at the IP layer is assumed to carry no additional cost beyond the working wavelength miles.

Conclusions

This article reviews the most common restoration and protection schemes available at the IP and WDM layers. It is envisioned that both layers will jointly cooperate to provide the necessary resilient schemes to protect user traffic from disruption due to network failures. A heuristic approach, based on simulated annealing, is presented to optimally choose the resilient scheme for each individual line of an IP over WDM mesh network. To illustrate the applicability of the proposed approach with a practical design case, the cost function under optimization is defined as the sum of the total provisioned WDM bandwidth and the cost of the IP restoration scheme characterized in the form of a penalty factor, K_{IP} . This penalty factor may be used to consider various design factors such as blocking probability of the restoration scheme, restoration latency, and IP router processing capability. The proposed heuristic proves to have the flexibility to deal with different multiwavelength self-healing ring protection schemes, such as DP-WSHR and SL-WSHR, without requiring significant changes in the optimization technique. Only the cost function needs to be adjusted to the particular protection scheme under consideration.

More study is required in this relatively young field of survivable IP over WDM networks, and several additional critical factors must be considered in the design of such networks. These factors include the application's requirements (e.g., not every application allows relatively slow IP restoration schemes), other protection schemes (e.g., end-to-end path protection), and legacy constraints (e.g., nodes in the network may support only few resilient schemes).

References

- [1] J. B. Slevinsky, W. D. Grover, and M. H. MacGregor, "An Algorithm for Survivable Network Design Employing Multiple Self-Healing Rings," *Proc. IEEE GLOBECOM '93*, Nov.-Dec. 1993, vol. 3, pp. 1568-73.
- [2] L. Crochat and J. Y. Le Boudec, "Design Protection for WDM Optical Networks," *IEEE JSAC*, vol. 16, no. 7, Sept. 1998.
- [3] S. Ramamurthy and B. Mukherjee, "Survivable WDM Mesh Networks, Part II -Restoration," *IEEE Int'l Conf. Commun., ICC '99*, Vancouver, Canada, June 6-10, 1999, IEEE, vol. 3, pp. 2023-30.
- [4] B. V. Caenegem et al., "Dimensioning of Survivable WDM Networks," *IEEE JSAC*, vol. 16, no. 7, Sept. 1998, pp. 1146-57.
- [5] A. Fumagalli and L. Valcarenghi, "Fast Optimization of Survivable WDM Mesh Networks Based on Multiple Self-Healing Rings," *All-Opt. Networking 1999: Architecture, Control, and Mgmt. Issues*, Boston, MA, 19-21 Sept. 1999, vol. 3843, pp. 44-55.

WSHR protection	Wavelength costs			Fiber costs		
	Min	Ave	Max	Min	Ave	Max
DP-WSHR	5,547,690	5,766,082	5,910,070	711,310	739,911	759,380
Shared fiber SL-WSHR	4,855,710	4,957,815	5,055,350	620,470	637,270	650,170
Non-shared fiber SL-WSHR	4,855,710	4,957,815	5,055,350	639,380	659,206	675,920

Table 3. Total cost for the Pan-European network with three types of WSHR protection schemes: DP-WSHR, shared fiber SL-WSHR (similar to the O-2F/BLSR), non-shared fiber SL-WSHR (similar to the O-4F/BLSR).

K_{IP}	f_w	Total λ -mi, working	Total WDM λ -mi, prot.	IP λ -mi, working	IP λ -miles work/ total λ -miles work
0.20	975,200	455,800	519,400	0	0.00
0.14	972,200	453,500	518,700	8000	0.02
0.12	956,700	452,600	504,100	13500	0.03
0.10	904,700	459,000	445,700	43500	0.09
0.08	723,000	445,200	277,800	196000	0.44
0.06	630,400	440,200	190,200	269000	0.61
0.00	432,500	432,500	0	432500	1.00

Table 4. K_{IP} variation ($p = 23,100$).

- [6] A. Fumagalli et al., "Survivable Networks Based on Optimal Routing and WDM Self-Healing Rings," *Proc. IEEE INFOCOM '99*, Mar. 21-25 1999.
- [7] G. Mohan and K. Somani, "Routing Dependable Connections with Specified Failure Restoration Guarantees in WDM Networks," *19th Annual Joint Conf. IEEE Comp. and Commun. Soc., INFOCOM 2000*, Tel Aviv, Israel, Mar., 26-30 2000, vol. 3, pp. 1761-70.
- [8] T. M. Chen and T. H. Oh, "Reliable Services in MPLS," *IEEE Commun. Mag.*, vol. 37, no. 12, Dec. 1999, pp. 58-62.
- [9] J. Anderson et al., "Protocols and Architectures for IP Optical Networking," *Bell Labs Tech. Journal*, vol. 4, no. 1, Jan.-Mar. 1999, pp. 105-24.
- [10] C. Metz, "IP Protection and Restoration," *IEEE Internet Comp.*, vol. 4, no. 2, Mar.-Apr. 2000, pp. 97-102.
- [11] G. Hjálmtýsson et al., "Simple IP Restoration for IP/GbE/10GbE Optical Networks," *OFC 2000*, Baltimore, MD, Mar. 5-10 2000, PD36.
- [12] O. Gerstel and R. Ramaswami, "Optical Layer Survivability: a Services Perspective," *IEEE Commun. Mag.*, vol. 38, no. 3, Mar. 2000, pp. 104-13.
- [13] S. Bandyopadhyay, Sengupta A., and A. Jaekel, "Fault-Tolerant Routing Scheme for All-Optical Networks," *All-Opt. Networking 1998: Architecture, Control, and Mgmt. Issues*, Boston, MA, Nov. 5-8, 1998, vol. 3531, pp. 420-31.

Biographies

ANDREA FUMAGALLI [M] (andrea@utdallas.edu) received his Ph.D. from the Politecnico di Torino, Italy, in 1992. From 1987 to 1989 he worked as a consultant for the national telephone research center in Italy, CSELT, where he was involved in the design and implementation of advanced telephone services supported over a fiber-based packet-switching ring network, LION. Between 1990 and 1992 he was a visiting researcher at the University of Massachusetts, where he defined and studied the concept of switched delay lines used as a storage method for all-optical packet switching. From 1992 to 1998, he was on the faculty of the Electrical Engineering Department at the Politecnico di Torino. In 1994 and 1995, on leave of absence from the Politecnico di Torino, he was at the University of Massachusetts where he was involved in a number of projects sponsored by NSF and DARPA. Since 1997, he has been associate professor of electrical engineering at the University of Texas at Dallas. Primarily dealing with optical network architectures and protocols, his research has been supported by national and international funding agencies.

LUCA VALCARENGHI [SM] (lucav@utdallas.edu) received his Laurea degree from Politecnico di Torino, Italy, in 1997 and his M.S. in electrical engineering from the University of Texas at Dallas (UTD) in 1999. He is currently enrolled in the Ph.D. program in electrical engineering — telecommunications at UTD where he is a research assistant and member of the Center for Advanced Telecommunications Systems and Services (CATSS). He received the TxTEC fellowship in 1998 and 1999. He is a feature editor of *SPiE/Baltzer Optical Networks Magazine* and on the Organizing Committee of OptiComm 2000. His main research interests are optical network design, analysis, and optimization; artificial intelligence optimization techniques; next-generation Internet; and IP over WDM networks.