IP-oriented control of unidirectional-path-switched-ring–based transport networks

Vishal Sharma

Metanoia, Incorporated, 1600 Villa Street, Mountain View, California 94041
v.sharma@ieee.org

Abhimanyu Das and Charles Chen

Mahi Networks, Incorporated, 1039 N. McDowell Boulevard, Petaluma, California 94954
adas@mahinetworks.com; cchen@mahinetworks.com

Received 7 October 2002; revised manuscript received 4 February 2003

An important requirement in the IP-based control of time-division multiplexing (TDM) optical transport networks is to utilize the in-built protection capabilities of synchronous optical network (SONET) unidirectional path-switched rings (UPSRs) and to automate the UPSR-protected path setup in mixed mesh–ring networks. This requires modifications to existing IP signaling and routing protocols and new processing rules at the network nodes. Here we leverage IP routing and signaling and multiprotocol label switching (MPLS) fast-reroute techniques for accurately advertising UPSR ring topologies to remote nodes and dynamically establishing UPSR-protected paths across a transport network. Our proposal also makes a NUT1-like (NUT, nonpreemptible unprotected traffic) feature possible in UPSRs, which allows for efficient utilization of UPSR protection bandwidth. We achieve this by encoding UPSR-specific information in the open shortest-path-first (OSPF) link state advertisements and in signaling messages of the Resource Reservation Protocol (RSVP) with TE extensions. In addition, we modify the signaling and routing state machines at the nodes to interpret and process this information to perform UPSR topology discovery and path computation. The uniqueness of our proposals is that the algorithms and the rules specified here allow for existing IP-based protocols [such as those within the generalized MPLS (GMPLS) framework, which currently applies to mesh networks] to be efficiently adapted for this context while still achieving our objective of exploiting UPSR-protection capabilities.

OCIS codes: 060.0060, 060.4250.

1. Introduction

As IP-based signaling and routing protocols are adopted for the dynamic control of optical time-division multiplexing (TDM) networks, it is imperative that they account for the large installed base of synchronous optical network (SONET) UPSR (unidirectional path-switched rings) [1] and BLSR (bidirectional line-switched rings) [2] rings. Thus, in mixed mesh–ring networks, the IP-based control protocols must allow for the automatic establishment of SONET channels while utilizing SONET ring protection capabilities. In this paper, we take a pragmatic approach and focus on SONET UPSRs. This is motivated by the interworking possibilities offered by UPSRs today (which is much more difficult, if not nonexistent, in BLSRs), which makes mixed mesh–ring networks with UPSRs good initial

candidates for the application of dynamic IP control to transport networks. As such, fully solving the topology distribution and path setup problem for UPSRs is a valuable first step.

There are two important issues to consider when IP protocols are used to manage and control legacy SONET ring networks. The first is to advertise the transport ring topology by use of IP routing protocols in a way that allows for path computation at each network node. The first issue involves advertising information about both the clockwise and the counterclockwise fibers of a ring and enabling a remote node to distinguish between working and protection bandwidth. The second is to establish dynamic ring-protected paths by use of IP-based signaling protocols. This involves automating the establishment *simultaneously* of both the working and the protection path. We solve both these issues in this paper.

We begin with a brief background of UPSR protection and of the IP-based protocols that are defined for the dynamic control of transport networks.

1.A. Unidirectional Path-Switched Ring Architecture

Ring topologies are by far the most widely deployed SONET network topology, and a common ring protection–restoration scheme in use today is UPSR protection. A UPSR is a survivable, closed-loop, transport architecture that protects against fiber cuts and node failures by providing duplicate, geographically diverse paths for each circuit [1]. Adjacent nodes on the ring are connected by use of a single pair of optical fibers, which form two counterrotating rings carrying traffic in opposite directions [1] (see Fig. 1). Thus, working traffic travels in one direction (say, clockwise) on one fiber while a protection path is provided in the opposite direction over the other fiber. A source typically sends traffic in both directions around the ring, so a UPSR can be used to provide a fully protected end-to-end path on a ring. Protection paths are set up and reserved when the working path is set up. In UPSR networks, the destination node on the ring monitors transmission on both fibers and performs a protection switch to the alternate path if it detects degraded (or loss of) transmission. Thus, switching between fibers is immediate, and no communication is needed with the transmitter.

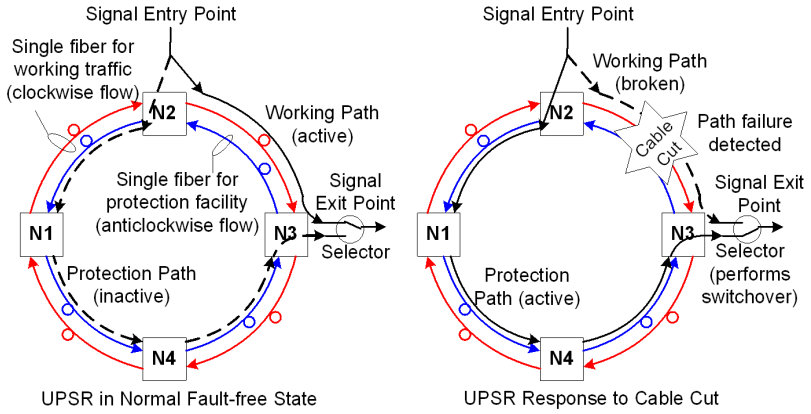

Fig. 1. SONET UPSR protection switching architecture.

The UPSR is an economical choice for most access and smaller metro applications, because its protection switching mechanism is much simpler than that of two-fiber (2F) or four-fiber (4F) BLSRs, and, unlike the 4F BLSR, it requires only two fibers to operate.

1.B. *Generalized Multiprotocol Label Switching Protocol Suite*

Generalized multiprotocol label switching (GMPLS) extends the MPLS [3] concept of label-switched paths and its traffic-engineering capabilities to the control of TDM, lambda, and fiber-switched networks [4,15]. GMPLS aims to provide a single, unified control plane architecture for multiple switching layers by adapting existing MPLS signaling [5] and IP routing protocols [6] for non-IP transport networks [7]. This requires several modifications to the network elements and the IP protocols. First, it requires that the transport network elements have IP-based control channels for interelement communication. Second, it requires an extension of the IP signaling protocols within the GMPLS protocol suite to instantiate lambda/TDM circuits in addition to IP label-switched paths (LSPs) and an extension of IP routing protocols to advertise link and node properties and other constraints important in TDM transport networks (such as end-point switching and link-protection capabilities). A reader interested in understanding these issues in the context of SDH (synchronous digital hierarchy)/SONET networks is referred to the paper by Berstein *et al.* [8]. Several of these extensions are being pursued in standards bodies, such as the Internet Engineering Task Force (IETF) [9,10]. The basic elements of the GMPLS architecture are illustrated in Fig. 2.

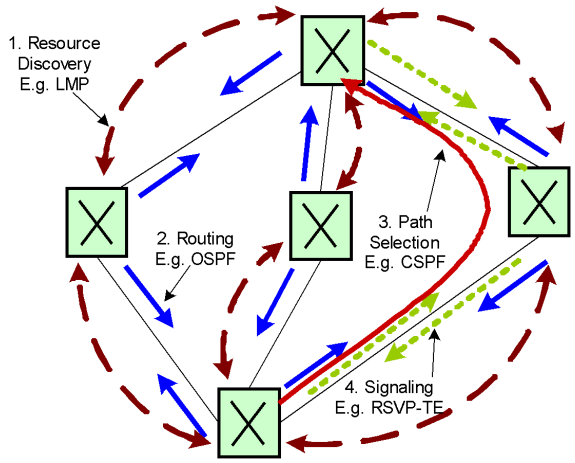

Fig. 2. GMPLS: basic architectural components.

When IP-based protocols are applied from data networks to IP-controlled transport networks, a key difference is the following. In the packet domain, the forwarding of data (IP data packets) and control information (IP signaling and routing protocol packets) is inherently on the *same* channel. In the transport domain, however, there is a natural separation of the control and forwarding planes. Further, the forwarding plane is circuit switched, whereas the control plane is packet switched. To allow for resilient IP-controlled transport networks, we therefore need to look at not only the data channels but also the control channels. In this paper we address only the protection of the data channels in the forwarding plane. The issues of control plane redundancy, while extremely important, are very different and are not dealt with here.

2. Motivation and Related Research

The current optical transport network in North America has a preponderance of SONET rings, more than 100,000 [11] at last reckoning, amounting to billions (10^9) of dollars in capital investment. There are also an equal number of SDH rings deployed across the world, representing a similar investment. A large fraction of these rings in the access and metro environments are UPSRs (or their SDH counterparts). Therefore any attempt to automate the provisioning and control of optical transport networks must take this large installed base into account and must be able to interwork seamlessly across the deployed UPSR infrastructure.

The GMPLS framework proposed within the IETF aims to automate the provisioning of paths in optical networks, but so far it has remained focused on mesh networks. Thus extensions to existing RSVP (Resource Reservation Protocol)-TE [10]/CR-LDP signaling and OSPF (optical shortest-path-first)-TE [12]/ISIS-TE routing protocols for dynamic path computation and path establishment in TDM and dense DWM (DWDM) networks are geared at mesh topologies (perhaps because mesh topologies, being conceptually similar to packet network topologies, allow for a more natural application of IP protocols).

Clearly, it is not sufficient to automate the provisioning of paths merely over mesh topologies. In the absence of an integrated solution, path provisioning over rings continues to require the mostly manual TL1-based (Transaction Language 1, a language used to communicate with TDM switching equipment) configuration used today. The value of an automated control plane solution that can incorporate this existing equipment base, therefore, is tremendous.

Our proposals for dynamic IP-based configuration of UPSRs greatly enhance the management and operation of UPSR-based transport networks. This is especially important considering the time-consuming, static per-node configuration (using management systems) that is largely prevalent today. Another advantage of our proposal is that it allows for the provisioning of a type of nonpreemptible unprotected traffic (NUT) [13] on UPSR rings, which is not possible with a traditional UPSR configuration. This is an important feature, usually associated with the more complex 2F-BLSR and 4F-BLSR networks.

The primary motivation behind our research is to develop a control plane solution that incorporates SONET UPSRs. By solving a very practical problem, our research makes it easier for service providers to move toward adopting an automated control plane for their transport networks. To the best of our knowledge, we are not aware of any research that specifically addresses the issue of leveraging IP routing and signaling protocols to control *TDM ring networks*. As will be evident in the remainder of this paper, applying IP protocols to ring topologies requires a rethinking of the protocol extensions in light of the properties of rings.

For example, GMPLS signaling protocols [9] have defined the protection object/TLV (type, length, value) to specify the type of protection (1+1, 1:1, unprotected, or enhanced, for example) desired by a LSP at *each hop* along its path. The protocol allows for a source to set multiple link flag bits in the protection object to indicate the type of protection that is acceptable for the LSP. This allows for LSP setup in networks where links may offer different levels of protection. However, the protection object by itself is not sufficient to signal and establish an LSP with UPSR protection. For that it is also necessary to extend the signaling protocol to establish (on a UPSR) the protection segment of an LSP in conjunction with its working segment, and to correlate the two segments.

One way to achieve this would be to allow the *focal nodes* [16] of UPSR rings (for example, nodes 1 and 3 in Fig. 4 below) to split a single connection request into two LSP establishment requests within a given UPSR ring, one to establish the working segment of the LSP and the other to establish the protect segment. Thus in Fig. 4, assuming that a

unidirectional connection is desired between nodes 5 and 2, N3 would be responsible for splitting the original request coming from N5 into two requests. The first would establish the working segment through nodes N4, N1, and N2, and the second would establish the protect segment through nodes N3 and N2. This suggests an adaptation of the RSVP-TE fast-reroute techniques available for LSP setup in the IP domain [14] and is indeed an approach that we discuss further in Section 6.

Similarly, even though the current GMPLS routing enhancements to OSPF-TE have defined a link-protection-type sub-TLV, this is not adequate to convey information about UPSR links such as whether they correspond to working or protect fibers and the specific UPSR to which they belong. So additional enhancements are needed to advertise and process UPSR link LSAs (link state advertisements) with OSPF-TE.

In the following sections we focus on how GMPLS RSVP-TE signaling and the corresponding processing rules may be extended to signal UPSR-protected LSPs and on how GMPLS OSPF-TE processing may be extended to advertise links belonging to UPSRs so that remote network nodes may build the topology of UPSRs in the network. Note that although we use OSPF-TE and RSVP-TE as examples, our proposals are generic and can be applied to other routing and signaling protocols as well, such as IS–IS (intermediate system to intermediate system) and CR-LDP (constraint-based routing label distribution protocol).

3. Illustrative Example

Before we get into the protocol-specific details of our solution, it will be useful to illustrate our ideas with an example. We consider how an enterprise's connection would be provisioned over a network of UPSR rings by a service provider that uses GMPLS to manage private-line customers. The process of end-to-end provisioning of a protected circuit requested by the customer is illustrated in Fig. 3. [For simplicity, in the following we assume that a unidirectional circuit is required from CPE1 to CPE2 (CPE, customer provisioning equipment).]

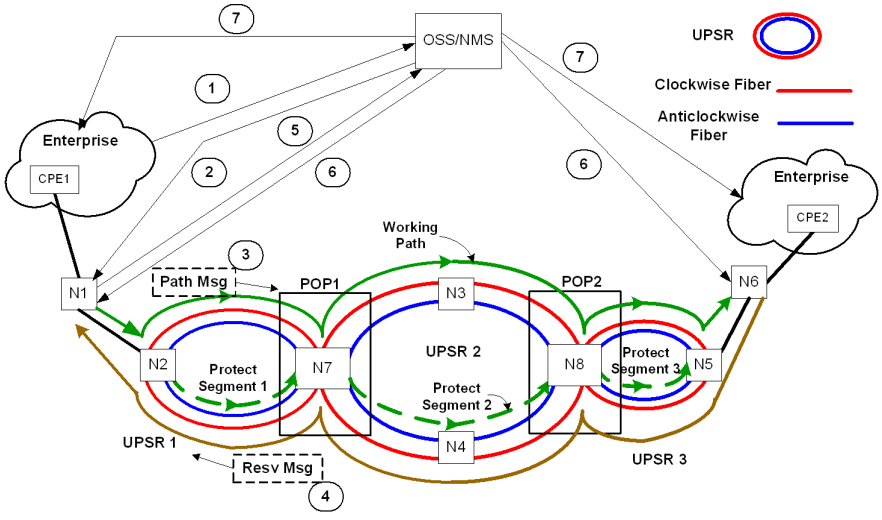

Fig. 3. Illustration of how path setup would work from one facility of an enterprise customer to another across a carrier's SONET UPSR infrastructure, when the carrier implements the extensions proposed in this paper.

The process of requesting and provisioning the circuit would work as follows:

- Step 1: The customer requests the service provider's network management system (NMS) or operations support system (OSS) for a TDM circuit with the required bandwidth and protection parameters. The required circuit must extend from the CPE in one enterprise domain (CPE1) to the CPE in the remote domain (CPE2).
- Step 2: The NMS informs the first node connected to the CPE (node N1) to initiate IP-based signaling across its network to establish a UPSR-protected circuit to the node connected to the remote CPE (node N6).
- Step 3: N1 initiates a signaling request (a RSVP-TE path message with our UPSR enhancements) destined to N6 to configure both the working and the protect segments of the UPSR-protected circuit. Since the service provider runs IP-based routing with our UPSR enhancements, N1 knows the complete topology of the UPSR network and has sufficient information to route the signaling requests to N6 over the appropriate clockwise or counterclockwise UPSR links. Thus the RSVP-TE path message travels from N1 through UPSRs 1, 2, and 3, en route to N6. Also, within each ring, appropriate signaling is initiated to simultaneously establish the protect segment, which protects the segment of the working path traversing each ring. Thus, protect segment 1 protects the portion of the working path routed over UPSR 1, and so on.
- Step 4: Node N6, upon receiving the signaling request from N1, responds with a RSVP-TE Resv message, which travels back toward N1 as shown in Fig. 3. The Resv messages establish the cross connections needed for the requested circuit on their way back to node N1.
- Step 5: Upon receipt of a Resv message at N1, node N1 informs the NMS/OSS about the successful creation of the circuit between nodes N1 and N6 in the service provider's network. All that remains at this point is for the service provider's NMS to configure the cross connects at nodes N1 and N6 to connect to their respective CPEs.
- Step 6: The NMS configures the cross connects at nodes N1 and N6 to connect the appropriate slots on the links between CPE1 and N1 and between CPE2 and N6, respectively, to the circuit just established between nodes N1 and N6.
- Step 7: After (possibly) a period of testing and validation of the TDM circuit just established, the NMS informs the two CPEs that the circuit request was successful and that the circuit is available for carrying user traffic.

Observe that in the process just outlined, apart from the communication between the NMS and nodes N1 and N6, all NMS-based interaction within the service provider's network is completely eliminated. That is, the service provider no longer needed to configure cross connects on a per-node basis. Rather, the entire process of creating a protected UPSR trail or circuit within the provider's network was fully automated. This is the benefit of the algorithms and protocol enhancements proposed in the remainder of this paper. Note also that the mechanisms we describe apply only to the interaction between the nodes within the service provider's domain; the interaction of the CPE with the service provider's network is not part of our mechanisms.

4. Enhancements to Routing Protocols and Path Computation

The goal of the routing enhancements is for every node in a mixed mesh–ring network to discover the complete topology of the network's UPSR rings. This means that, for each link, every node must be able to ascertain information such as

- whether the link belongs to a UPSR, and if so
- the particular UPSR to which the link belongs, and
- whether the advertised link corresponds to the working or the protect components of the UPSR (which is applicable when the provider uses a strict working fiber/protect fiber designation for the clockwise and the counterclockwise fibers of the UPSR).

We observe that carriers typically use one of two options when provisioning UPSRs:

1. The first is to designate, a priori, one fiber to carry only working traffic (the working fiber) and the other to carry only protection traffic (the protect fiber). This is helpful in network planning and in the provisioning and tracking of fibers and circuits.
2. The second is to avoid explicitly preassigning UPSR fibers as working or protect but rather to allow either one of the two fibers to be selected to carry the working or protect path of any circuit. In this case, the same fiber could carry working paths and protect paths for different UPSR circuits. This is useful in creating bidirectional UPSR circuits, where a carrier would like to route the working segments of both the forward and the reverse circuits on the short path (and hence the protect segments in both directions on the long path). So each fiber would carry the working path for one direction of the circuit and the protect path for the other direction of the circuit. This is helpful in keeping the working circuits and the protection circuits on the same fiber pairs (facilities).

In this section we focus on the first option, which is the more involved one. We discuss the second option in Section 5, as it follows easily from our discussion here.

We also note that unused bandwidth on the protection channels of a UPSR can be used to route low-priority traffic, if it can be discovered by remote network nodes using the routing enhancements discussed in Subsection 4.A. The unused protect bandwidth is usable by NUT.

Nonpreemptible unprotected bandwidth is bandwidth on a segment of the protect fiber for which the corresponding bandwidth on an appropriate segment of the working fiber has been used to create an unprotected circuit, which will, therefore, never make use of its corresponding timeslots on the protect fiber. Thus, traffic using this bandwidth is unprotected, but nonpreemptible, akin to NUT in 2F and 4F BLSRs.

4.A. Enhancements to OSPF-TE Subobjects

We propose enhancements to the subobjects in GMPLS OSPF-TE [14] and propose a set of rules that every network node must follow to derive a consistent network topology from the OSPF LSAs. The enhancements to OSPF-TE are as follows:

- The OSPF-TE LSA must have a field to indicate the underlying *protection technology* to which a link belongs. Currently, this could be linear, UPSR, 2F-BLSR, 4F-BLSR, SNCP (subnetwork connection point), or MS-SPRING (multiplex-section shared-protection ring). This ensures that a receiving node always knows the type of ring to which a link described by the LSA belongs. It also provides extensibility, since new types of rings (such as optical rings) can also be incorporated easily by means of defining new code points for this field.
- An optional “Ring ID” field must be added to the protection subobject and used in OSPF-TE LSAs describing *UPSR links* (shorthand for “links on UPSR rings”) so that links belonging to a given UPSR ring may be readily identified. This ensures

that when multiple UPSRs are superimposed over the same set of physical nodes, this attribute enables remote nodes to distinguish easily between the links of different UPSRs.

- UPSR link components that lie on the working and the protect fibers are distinguished by the link-protection-type sub-TLV. Link components on the working fiber are advertised with the "enhanced" link-protection type, and those on the protect fiber with the "unprotected" link-protection type.

- Finally, the OSPF-TE LSA must also allow the bandwidth available for routing NUT traffic to be discovered by remote nodes. The easiest way to do this would be for the node originating a protect link to advertise the specific channels on it that are available for routing NUT traffic (by advertising, for example, a channel bit map to indicate which of the free channels may be used to route NUT traffic). However, for scalability, we propose to advertise only a single number, namely, the *maximum contiguous bandwidth* available for NUT traffic. (The method to calculate the maximum contiguous bandwidth is explained in Subsection 4.B.) The switching capability specific information field in the interface switching capability descriptor in GMPLS-OSPF-TE [12] can be used to carry this value.

A difference between IP-controlled UPSRs and normal IP networks is that IP assumes that reverse traffic between the end points of a link can be sent back over the *same* link. In a UPSR with the working/protect fiber distinction, however, the reverse direction of a UPSR "link" lies on a counterclockwise fiber and *may not* be used for sending working traffic in the reverse direction. Thus, even though each advertised UPSR link is physically bidirectional, the "reverse path" for sending traffic back to the source node need *not* be over the same link but rather over the remaining part of the ring. Thus in creating the network topology from the OSPF-TE LSAs, a remote node must keep the logical unidirectional nature of the clockwise (and counterclockwise) components of a link in mind.

For example in Fig. 4, node 3 will advertise router LSAs with three links, link(3,4), link(3,2), and link(3,5). Link(3,4) will describe the working segment of the UPSR link between nodes 3 and 4. Link(3,2) will describe the protect segment of the UPSR link between nodes 3 and 2, and link(3,5) will describe a normal, linear, unprotected link between nodes 3 and 5. The protection technology subobject and link-protection-type subobject for link(3,4) will be UPSR and enhanced, respectively, whereas for link(3,2) they will be UPSR and unprotected. For link(3,5) the subobjects will be linear and unprotected, respectively. Note that the clockwise (working) segment of every UPSR link is advertised by the node at one end of the link, whereas the counterclockwise (protect) segment is advertised by the node at the other end.

To send working traffic from N2 to N3, one may use the clockwise (working) component 2→3 of the UPSR link between nodes N2 and N3. To send working traffic back from N3 to N2, however, one may either use the same clockwise fiber, and go via the clockwise components of the links between nodes N3, N4, N2, and N1, in the order 3→4→1→2, or one may use the counterclockwise (protect) component 3→2 of the link between nodes N3 and N2. Therefore, a bidirectional working-traffic link between nodes N2 and N3 may be composed logically of the clockwise component 2→3 of the link between nodes N2 and N3 and the clockwise components 3→4, 4→1, and 1→2 of the links between nodes N3, N4, N2, and N1. Alternatively, it may be composed logically of the clockwise component 2→3 and the counterclockwise component 3→2 of the link between nodes N2 and N3.

4.B. Enhancements to Rules for Advertising and Processing Link State Assignments

Finally, the changes to the rules for advertising and processing LSAs are as follows:

- The working and protect components of a UPSR link must each be advertised in a separate LSA by the node from which the component emanates. Thus in Fig. 4, for the UPSR link between nodes 2 and 3, the clockwise component 2→3 will be advertised by node 2, whereas the counterclockwise component 3→2 will be advertised by node 3.

- All LSAs describing working or protect components of a UPSR link must have the UPSR code point in the new OSPF-protection technology subobject.

- The working component of a UPSR link must be advertised with the enhanced link-protection type in the protection subobject, whereas the protect component must be advertised with the unprotected link-protection type (for the case in which the provider explicitly designates working and protect fibers on the UPSR).

- The maximum contiguous bandwidth available for routing NUT traffic must be calculated and advertised by a node and can be done as follows. For each outgoing protect link, the node calculates the timeslots used by all the unprotected TDM circuits (LSPs) routed on the corresponding working link (the reverse direction of the outgoing link). The same timeslots on the protect link would then be available for routing NUT traffic, less the timeslots being used by any NUT circuits that may already be configured over the outgoing link. The node accounts for this and advertises the largest block of contiguous bandwidth available for use by NUT traffic on that outgoing protect link.

The protection technology subobject, the link-protection-type field, and the Ring ID together enable a remote node to completely identify the working and protect components of a UPSR's links. The topology that a remote node uses for calculating paths via shortest-path-first (SPF)/constraint SPF (CSPF) for protected traffic (or traffic that needs to be routed over the protect fiber) must use only the LSAs corresponding to working components of UPSR links. However, when paths are computed for LSPs able to be routed over the protect components of UPSR links, a node may use LSAs that describe both the working and the protect components of UPSR links.

Note that by advertising and processing the LSAs for the protect component of UPSR links separately, we enable a remote node to discover unused bandwidth on the protection channels, which can be used to route low-priority traffic.

Given these changes in LSA processing at the sending and receiving nodes (and some changes to signaling to be described in Section 6), a node can now treat the task of setting up a UPSR-protected TDM trail as equivalent to setting up a working LSP, which is routed over the working fibers on the UPSRs in its path, and detour LSPs that are initiated by every ingress focal node and routed over the protect fibers on the UPSRs. The intermediate links taken by the working LSP can be explicitly specified at the source by means of consulting its TE database. The path of the detour LSPs need not be specified and can be calculated by the respective focal nodes. Once the focal nodes on the path of the UPSR-protected TDM trail can correlate the working and protect/detour LSPs, the end result is the same as if one had set up a single, manually provisioned, UPSR-protected TDM circuit.

5. Unidirectional Path-Switched Ring without Explicit Working/Protect Fiber Distinction

In this section we focus on the case in which a carrier does not explicitly define a working fiber and a protect fiber on a UPSR, and either one of the two fibers could be used to carry the working and protect path of any circuit.

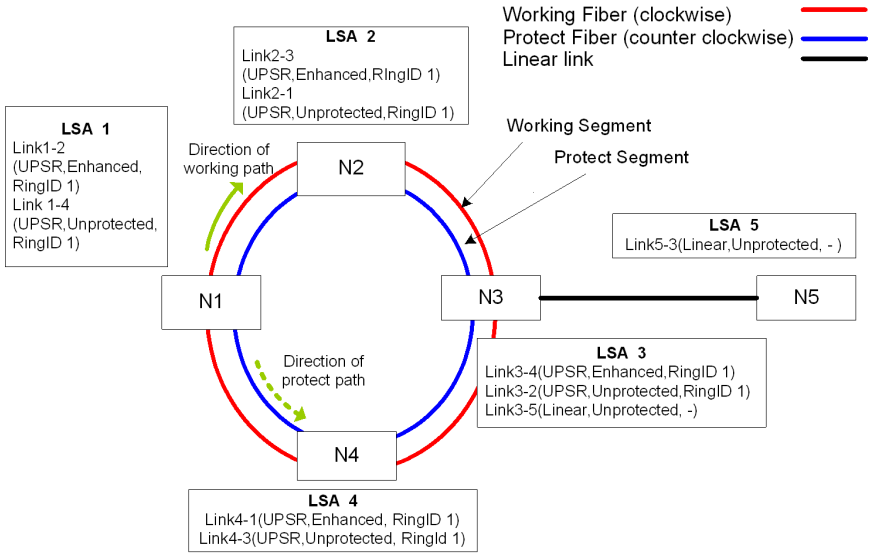

Fig. 4. OSPF-TE LSA advertisements to enable correct UPSR topology inference at remote nodes.

In this case, all the enhancements to the OSPF-TE subobjects outlined in Subsection 4.A carry over, with the exception that there are no longer designated “working” and “protect” fibers on the ring, so every link is advertised with only the enhanced link-protection type.

Also, the method described above to calculate the bandwidth available for carrying NUT traffic on the outgoing protect links at a node now extends to all of a node’s outgoing links. So for each outgoing link, a node must calculate the timeslots used by the unprotected TDM circuits (LSPs) routed over the reverse (incoming) direction of the link. The corresponding timeslots on the outgoing link would then be available for routing NUT traffic, less the timeslots used by any NUT circuits that may already be configured over the outgoing link.

6. Enhancements to Signaling Mechanisms and Protocols

The fundamental goal of the signaling enhancements is to set up a TDM circuit by use of a single LSP setup from the source. In other words, the signaling should be able to establish both the working and protect segments of the LSP over each UPSR that the circuit crosses. This is accomplished by allowing the focal nodes on each intermediate UPSR to spawn a detour LSP [14], which sets up the protect segment corresponding to the working segment of the LSP. In essence, this helps localize the effect of failures, just as is the case for SONET channels established by configuration, since failures on intermediate UPSRs can now be handled locally.

Since the objective is to enable RSVP-TE to configure the cross connects at every node through which a TDM circuit passes, an alternative would be to initiate two explicitly routed LSPs from the source itself, one of which takes the working segments over the UPSRs en route while the other takes the protect segments. Even though this option requires fewer changes to the GMPLS signaling protocols, it is not a very robust solution. For one, the failure of any node along the working path would cause a switchover to the protection

segments in all UPSRs over which the trail was routed. This is clearly unacceptable and goes against the very essence of UPSRs, which is to localize the effect of failures. Second, for a TDM trail passing through two UPSRs, if the working segment on the first ring failed, and the protect segment on the second ring failed, both the working and protect LSPs would be torn down, thus bringing down the TDM trail, even though a path using the working segment in the first ring and the protect segment in the second ring would still be available to service the LSP.

Our signaling proposals work within the framework of the GMPLS protocols while using some RSVP-TE fast-reroute enhancements from the packet domain, but they still integrate SONET UPSR rings (and their protection capabilities) into the network.

6.A. *Detour Link-Switched Path Enhancements*

We follow the same basic strategy as outlined in the detour draft [14], where there is a fast-reroute object in the primary LSP and a detour object in the detour LSP. As in the detour draft [14], the detour LSP is initiated at the points of local repair (PLRs), which are the ingress focal nodes in UPSRs, and is merged back at the merge points, which are the egress focal nodes in UPSRs. The fast-reroute proposal, however, has no provision for a source to specify the PLR and the merge points, a functionality that is needed in the UPSR case to allow UPSR hub nodes for correctly setting up working and protect paths. (Recall that the fast-reroute proposal is designed for packet LSPs, where the objective is to protect as many nodes/links by local repair as possible.) To enable this, we propose to use two bits (the PLR bit and the MERGE bit) in the explicit route object (ERO) subobject of the RSVP path message, which allows every node specified in the ERO at the source to be marked as a MERGE or a PLR node, as needed.

6.B. *Resource Reservation Protocol State Machine Enhancements at Focal Nodes*

Our signaling enhancements are designed to minimize any change in normal processing of GMPLS RSVP-TE messages. The only changes necessary in RSVP-TE processing are at the UPSR focal nodes.

The ingress focal node on receiving a TDM circuit setup request checks for the presence of the fast-reroute object to determine whether the TDM circuit is to be UPSR protected. If this object is present, the node further checks the ERO object to see whether it is a designated PLR and, if so, locates the very next merge point, which is the egress focal node on the current ring. Having done so, the node consults its TE database and constructs an ERO for the detour LSP that it uses to establish the protect path over the UPSR. The original LSP path request continues over the working path specified in its ERO to set up the working UPSR circuit. Upon receiving labels from both the working and detour LSPs, the ingress focal node configures its UPSR hardware to enable swapover for UPSR protection.

The egress focal node, on receiving path messages from both the working and detour LSPs, merges them and forwards only one LSP request downstream of itself, as specified in the detour draft [14]. Upon receiving a Resv message from its downstream neighbor, the node sends Resv messages back individually for the working and detour LSP setup requests over appropriate fibers of the UPSR. It also configures its UPSR cross connects to enable UPSR protection on the two incoming links specified by the two LSP requests.

An important change that we make in the RSVP processing rules relative to the detour draft [14] is that every ingress focal node, upon receiving an LSP path message with the fast-reroute object set for the first time (which indicates to the node to spawn off a detour LSP), does not forward a Resv message upstream (toward the source) until it gets Resv messages from both the working and detour LSPs. Alternatively, until this node receives the first Resv message from both the working and detour LSPs, any PathError message it receives for either LSP is forwarded upstream. This is necessary so that if the detour

LSP is unable to set up the protect path over a UPSR ring, the end-to-end LSP setup fails as well. Stated more formally, if a node detects that it is an ingress focal node/PLR and spawns a detour LSP request, then it should send the first Resv message back upstream only when it receives a Resv message back from both the working and detour LSPs. This check is applied only for the first Resv message that the ingress focal node has to send back. Once the LSP is up (that is, both the working and protect paths have been successfully created) a node failure on either the working or the protect LSP segment on the ring, but not both, should not tear down the overall LSP. Thus, after the LSP is up, a PathError message received at the ingress focal node over either the working or the protect segment alone should not be forwarded upstream, unless a PathError has been received from both segments of the LSP. Figure 5 illustrates the setting up of the working and protect paths over a UPSR topology.

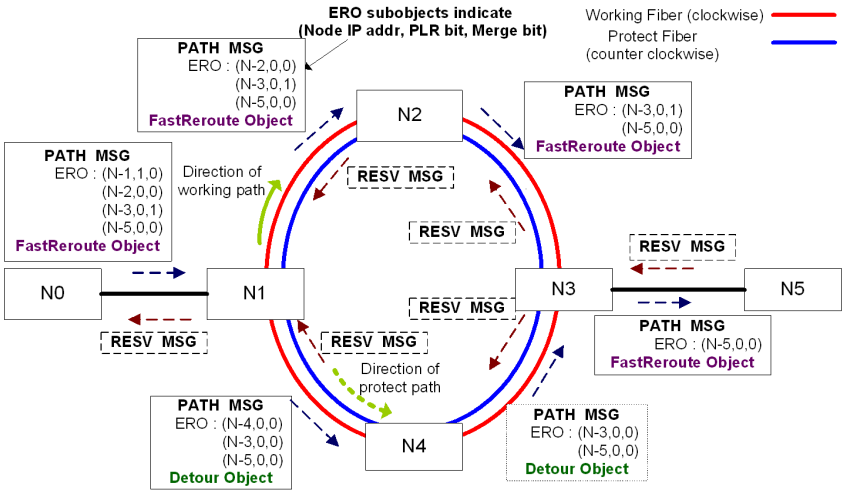

Fig. 5. Setup of working and detour LSPs over UPSR ring topology, from N0 to N5 (PATH and Resv message flow). Note that the protect fiber is not shown above, for simplicity.

6.C. *Bidirectional Protected Label-Switched Path Setup Using Upstream Label Objects*

A bidirectional UPSR-protected LSP can be setup by inclusion of an UPSTREAM label in the detour LSP setup. Note that the usual paradigm for the UPSTREAM label does not apply when the carrier makes a strict working fiber/protect fiber distinction. The UPSTREAM label in the usual PATH message in GMPLS [14] is used to set up the reverse working path for a bidirectional LSP. In other words, the UPSTREAM label establishes a working path in the reverse direction to that corresponding to the direction of the PATH message (which is the direction of the forward working path).

However, when a provider explicitly defines a working fiber and a protect fiber on a UPSR (see Fig. 6), the UPSTREAM label in the RSVP PATH message traveling on the working fiber (for setting up the working TDM LSP) will set up the protect LSP corresponding to the working path in the reverse direction. Similarly, the UPSTREAM label included in the detour object (for setting up the protection LSP for the forward working path) will set up the working LSP for the reverse working path.

When the service provider makes no distinction between the working and protect fibers the situation is simpler and conforms to the usual paradigm for the UPSTREAM label object. In this case, the UPSTREAM label in the RSVP PATH message sent to establish the working segment of the TDM circuit sets up the working LSP corresponding to the TDM circuit in

the reverse direction. Likewise, the UPSTREAM label included in the detour object (that is used to establish the protection segment corresponding to the working segment) also sets up the protection segment corresponding to the TDM circuit in the reverse direction.

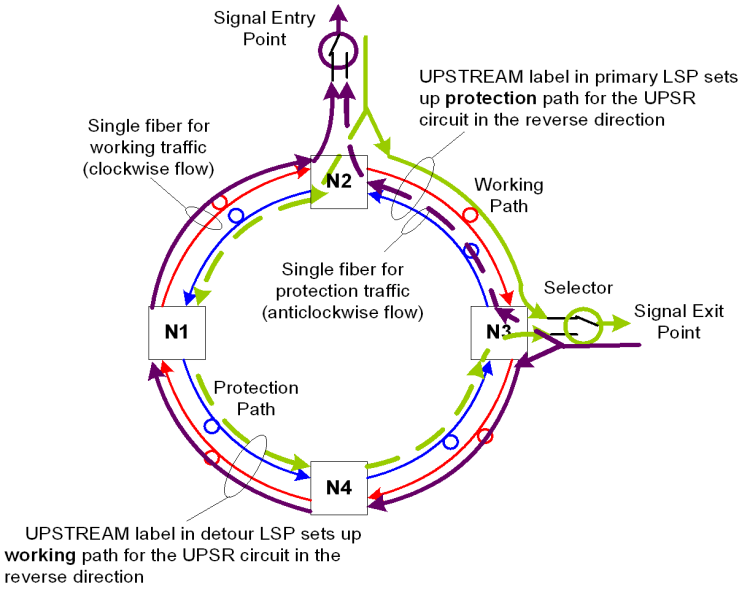

Fig. 6. Illustration of bidirectional TDM LSP setup on a UPSR. Note the differences with the usual paradigm for the UPSTREAM label for bidirectional LSPs.

7. Conclusion

In this paper we have presented a comprehensive proposal to enable IP-based automated topology/resource discovery and path computation and have automated path establishment for UPSR transport networks. By suitably modifying the existing GMPLS signaling and routing protocols for mesh networks, and adapting the concepts of RSVP-TE fast reroute, we allow for the setting up of UPSR-protected LSPs in mixed mesh–ring networks. In addition, by enabling a NUT-like feature for UPSRs, we also provide a mechanism to use the protection bandwidth of a UPSR more efficiently.

Our analysis of the mechanisms and the protocol changes proposed indicates that from a technical or engineering perspective, it is indeed feasible to interwork UPSR and mesh-based networks in a dynamic setting and to take advantage of UPSR protection. However, in this paper we provide only an architectural solution that illustrates how service providers may create and operate a UPSR-based dynamic transport network. Actual deployment requires additional analysis by the service providers regarding traffic patterns, network management, security and cost issues, which we have not addressed in this paper. Other practical factors pertinent to deployment may also need to be considered, despite several important advantages of our proposal that we highlighted in Section 6. For instance, the TDM equipment in a carrier network would need to be enhanced with a standardized IP-based control plane. Its feasibility would depend upon each provider's equipment and network configurations.

Several directions of future research are possible from here. One is to extend our proposal to the case of dual-ring interworking (DRI) between UPSRs or between UPSRs and mesh networks, since dual-node interconnection using drop-and-continue avoids single

points of failure. Another is to look at the issue of end-to-end protection of TDM circuits, that is, protection and traffic engineering across areas and domains. Yet another is to consider how protection is handled in the case of multicast. With applications such as video conferencing and web casts, this will become increasingly important. It would also be useful to look at control plane redundancy, which is an important topic in its own right and one that we did not cover in the current paper. Finally, one may see how these notions can be generalized to apply to BLSRs.

Acknowledgments

We acknowledge the insightful comments and suggestions provided by the Associate Editor, Mark Allen, and by the anonymous reviewers, which helped us to refine certain aspects of our scheme.

References and Links

[1] GR-1400-CORE, "SONET dual-fed unidirectional path switched ring (UPSR) equipment generic criteria," Issue 2 (Bellcore, January 1999), http://www.telcordia.com/.

[2] GR-1230-CORE, "SONET bi-directional line switched ring (BLSR) equipment generic criteria," Issue 4 (Bellcore, December 1998), http://www.telcordia.com/.

[3] G. Swallow, "MPLS advantages for traffic engineering," IEEE Commun. Mag. (December 1999), pp. 54–57.

[4] D. Awduche and Y. Rekhter, "Multiprotocol lambda switching: combining MPLS traffic engineering control with optical crossconnects," IEEE Commun. Mag. (March 2001), pp. 111–116.

[5] A Bannerjee, J. Drake, J.P. Lang, B. Turner, D. Awduche, L. Berger, K. Kompella, and Y. Rekhter, "Generalized multi protocol label switching: an overview of signaling enhancements and recovery techniques," IEEE Commun. Mag. (July 2001), pp. 144–151.

[6] A. Banerjee, J. Drake, J. P. Lang, B. Turner, K. Kompella, and Y. Rekhter, "Generalized multi protocol label switching: an overview of routing and management enhancements," IEEE Commun. Mag. (January 2001), pp. 144–150.

[7] G. Bernstein, J. Yates, and D. Saha "IP-centric control and management of optical transport networks," IEEE Commun. Mag. (October 2000), pp. 161–167.

[8] G. Bernstein, E. Mannie, and V. Sharma, "Framework for MPLS-based control of optical SDH/SONET networks," IEEE Netw. (July/August 2001), pp. 20–26.

[9] L. Berger, ed., "Generalized MPLS—signaling functional description," Work in Progress, draft-ietf-mpls-generalized-signaling-09.txt (Internet Engineering Task Force, August 2002),http://www.ietf.org/internet-drafts/draft-ietf-mpls-generalized-signaling-09.txt .

[10] L. Berger, ed., "Generalized MPLS signaling—RSVP-TE extensions," Work in Progress, draft-ietf-mpls-generalized-rsvp-te-09.txt (Internet Engineering Task Force, September 2002), http://www.ietf.org/internet-drafts/draft-ietf-mpls-generalized-rsvp-te-09.txt.

[11] J. Duffy, "Intelligent services make MANs hot," Network World (8 May 2000), http://www.nwfusion.com/news/2000/0508infra1.html.

[12] K. Kompella and Y. Rekhter, eds., "OSPF extensions in support of generalized MPLS," Work in Progress, draft-ietf-ospf-gmpls-extensions-09.txt (Internet Engineering Task Force, December 2002), http://www.ietf.org/internet-drafts/draft-ietf-ccamp-ospf-gmpls-extensions-09.txt.

[13] Nonpreemptible unprotected traffic, a feature commonly offered in more expensive and more complex 2F and 4F BLSR systems.

[14] P. Pan, D. H. Gan, G. Swallow, J. P. Vasseur, D. Cooper, A. Atlas, and M. Jork, "Fast reroute extensions to RSVP-TE for LSP tunnels," Work in Progress, draft-ietf-mpls-rsvp-lsp-fastreroute-01.txt (Internet Engineering Task Force, November 2002), http://www.ietf.org/internet-drafts/draft-ietf-mpls-rsvp-lsp-fastreroute-01.txt.

[15] Y. Xu, P. N. Lamy, E. L. Varma, and R. Nagarajan, "Generalized MPLS-based distributed control architecture for automatically switched transport networks," Bell-Labs Tech. J. **6**(1) (January–June 2001), pp. 13–32.

[16] We define a focal node on a UPSR ring to be a node that either originates or terminates a TDM LSP (or TDM channel/circuit) or one that sits at the intersection of two or more UPSR rings.