Internet Draft
Expiration: September 2001

Nasir Ghani
James Fu
Dan Guo
Xinyi Liu
Zhensheng Zhang
Sorrento Networks Inc

Paul Bonenfant
Leah Zhang
Antonio Rodriguez Moral
Murali Krishnaswamy
Photuris Inc

Dimitri Papadimitriou
Alcatel

Sudheer Dharanikota
Raj Jain
Nayna Networks

Architectural Framework for Automatic Protection Provisioning In Dynamic Optical Rings

draft-ghani-optical-rings-01.txt

Status of this memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups ay also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt.

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

Distribution of this memo is unlimited.

Abstract

Given the large installed base of ring fiber-plants and the extensive experience operators have gained in operating SONET/SDH ring networks, optical rings are becoming increasingly important. As such, optical rings will play a crucial role in the migration from existing TDM-based SONET/SDH architectures to more dynamic lightpath provisioning paradigms. To date, various optical ring concepts have been tabled, proposing multi-services support and mirroring the fast protection switching capabilities of existing SONET/SDH rings. Nevertheless, the emerging MPL(ambda)S/GMPLS framework for optical networks is largely based upon (optical) mesh routing concepts. Clearly, there is a strong need to formalize a more comprehensive architectural framework for optical rings and ensure its proper integration within the emerging

MPL(ambda)S/GMPLS architecture. Along these lines, the various optical ring schemes are summarized and their associated dynamic provisioning concerns detailed.

Table of Contents

List of Acronyms

ADM: Add-drop multiplexer
APS: Automatic protection switching
AS: Autonomous system (routing domain)
BER: Bit error rate
BLSR: Bi-directional line-switched ring
BPSR: Bi-directional path-switched ring
COPS: Common open policy service
CR-LDP: Constraint-routing label distribution protocol

DPRING: Dedicated protection ring
DRI: Dual ring interconnection
DWDM: Dense wavelength division multiplexing
DXC: Digital cross-connect
EXC: Electronic cross-connect (electronic cross-point switch)
FIS: Failure indication signal
FRS: Failure recovery signal
GFP: Generic framing protocol (for data over SONET/SDH)
GMPLS: Generalized multi-protocol label switching
IED: Integrated edge device
IGP: Interior gateway protocol
ILM: Incoming label map
IPoRPR: IP over resilient packet ring
LMP: Link management protocol
LOF: Loss of framing
LOL: Loss of light
LOS: Loss of signal
LSA: Link state attribute
LSP: Label switched path

LSR: Label switch router (also lambda switch router)
MEMS: Micro-electro-mechanical systems
MPLS: Multi-protocol label switching
NHLFE: Next-hop label forwarding entry
NMS: Network management system
NNI: Network-to-network interface
O-ADM: Optical add-drop multiplexer
O-BLSR: Optical bi-directional line-switched rings
O-BPSR: Optical bi-directional path-switched rings
OC-n: Optical carrier
OCh: Optical channel
OMS: Optical multiplex section
OPU: Optical payload unit
OSC: Optical supervisory channel
OSPF: Open shortest path first protocol
OXC: Optical cross-connect switch
PDH: Plesiochronous digital hierarchy
PML: Protection merge LSR
PMTG: Protected MPLS traffic group
PSL: Protection switch LSR
PXC: Photonic cross-connect switch
RNT: Reverse notification tree
RPR: Resilient packet ring
RSVP: Resource reservation protocol
RWA: Routing and wavelength assignment
SDH: Synchronous digital hierarchy
SHR: Self-healing ring
SNC: Sub-network connection
SNCP: Sub-network connection protection
SPRING: Shared protection ring
SONET: Synchronous optical network
SRLG: Shared risk link group
STM: Synchronous transfer module
TCP: Transport control protocol

TDM: Time division multiplexing
TLV: Type length value (field)
UNI: User network interface
VPOR: Virtual private optical ring
UPSR: Uni-directional path-switched ring
WDM: Wavelength division multiplexing
WRS: Wavelength-routing switch

1. Introduction

Many networks today are based upon fiber-ring architectures, as evidenced by the proliferation of SONET/SDH rings all the way from the long-haul backbone to the metropolitan and regional areas. Most larger backbone rings represent significant investments on the part of service providers, and expectedly will have longer lifetimes. Additionally, in the regional metro space, hierarchical SONET/SDH ring architectures are also very commonplace. For example, at the access-side, smaller (optical carrier/synchronous transfer module) OC-3/STM-1 (155 Mb/s) tributary rings are used to aggregate and groom traffic from enterprise customers. These rings are then connected to larger granularity OC-12/STM-4 (622 Mb/s) and possibly OC-48/STM-16 (2.5 Gb/s) rings spanning larger metropolitan distances. Metropolitan rings are then used to feed into even larger regional (and possibly long-haul) fiber-ring topologies with increased bit rates, such as OC-192/STM-64 (10 Gb/s). As a result, ring architectures will clearly play a major role in the evolution of optical networks.

Given this large, entrenched base of ring topologies, currently many operators are planning for a migration to equivalent dynamic optical ring architectures. Dynamic optical rings can be defined as fiber

rings with dynamic lightpath provisioning capabilities (such as routing, add/drop, and protection). These optical wavelength routing rings, commonly also referred to as optical add-drop ring multiplexer (O-ADM) rings, will form the mainstay architecture for most metro/ regional and even long-haul networks, helping operators ease their transition to future optical (mesh or hybrid ring-mesh) networks. Since many operators have significant experience in deploying and maintaining SONET/SDH rings, future optical analogs of such time-division multiplexing (TDM) ring switching are of great transitional value. Here, wavelength channels (as opposed to TDM circuits) undergo bypass, add, or drop operations at ring network elements [MARCENAC]. Optical rings will allow operators to immediately leverage their current fiber topologies and avoid lengthy fiber-expansion costs (i.e., associated with deploying mesh networks). Furthermore, ADM-based ring architectures are well-known for their operational simplicity and inherently fast protection switching capabilities, and perhaps, this is the main reason for the widescale acceptance of SONET/SDH technology. Network operators have become well-accustomed to the fast, timely recovery capabilities provided by SONET/SDH automatic protection switching (APS) schemes, such as uni-directional path switched rings (UPSR)/1+1 sub-network connection protection (SNCP) and bi-directional line switched rings (BLSR)/multiplex section shared protection rings (MS/SPRINGs) [GR1230],[T1.105.01],[G.841]. These architectures can achieve service recovery within 50 ms after a fault event, via

detailed electronic frame monitoring and fast protection switchover signaling provisions.

Meanwhile, recently there have also been significant developments in extending the ubiquitous multi-protocol label switching (MPLS) framework to the optical networking domain, namely "IP over optical" via MPL(ambda)S [AWDUCHE],[GHANI1],[RAJAGOPALAN] and more recently, generalized MPLS (GMPLS) [ASHWOOD1],[XU]. Nevertheless, given its origins from (mesh) IP packet routing networks, this framework as it stands today, is largely geared to support dynamic optical mesh networks. Conversely, no standards exist for optical rings and most offerings do not provide dynamic channel routing (add-drop) capabilities, relying instead upon proprietary, static solutions. Now given the abundance and strategic importance of ring fiber-plants (as detailed above), it is crucial to extend the existing MPL(ambda)S/GMPLS framework to provision dynamic optical ring networks. Although some may state that rings are special cases of meshes (technically speaking), the various intricacies of ring networks require special attention in the MPL(ambda)S/GMPLS framework. As most long-haul optical networks continue to migrate towards mesh-based GMPLS/MPL(ambda)S setups, along with increasingly MPLS-based "client" router networks, intermediate metro/regional networks (largely ring-based) must also evolve to a similar architecture. Such a uniform provisioning framework will permit true optical services provisioning across all network/geographic domains.

In particular, the MPL(ambda)S/GMPLS framework must address ring channel provisioning and protection switching functions. Undoubtedly, optical (ring) solutions must provide equivalent, or improved, capabilities in order to replace TDM rings in a timely manner. Since each fiber (or wavelength) in an optical network can now carry a much higher degree of multiplexed traffic, APS capabilities are even more crucial. This report details an architectural framework for optical rings, representing a logical, structured evolution (expansion) from existing SONET/SDH (TDM) ring paradigms. Optical ring equivalents of SONET/SDH protection schemes are presented and detailed provisioning issues outlined within the context of the broader MPL(ambda)S/GMPLS framework.

2. Framework for Optical Rings

Many of the proposed concepts in optical ring networks derive their origins from those in SONET/SDH ring networks. Here it is assumed that readers are familiar with basic SONET/SDH constructs, albeit a brief introduction is presented in the Appendix (Section 4). Due to the inherent transparency of optical switching technologies, optical rings can develop significantly upon existing TDM SONET/SDH rings. Namely, the concept of "transparent" optical rings is envisioned, permitting a full range of protocols/bit-streams being carried in their native format, e.g., SONET/SDH, ATM, IP, GFP, ESCON, SDL, Gigabit Ethernet, etc. Note that there are also standardization efforts to define generic mappings/encapsulations for data protocols, such as the generic framing procedure (GFP) [GFP] for the optical payload unit (OPUk), as defined

in ITU-T G.709 [G709]. Fundamentally optical ring network elements must perform the optical "equivalents" of TDM ADM channel operations: pass-through, add, drop, and fast APS functions [ARIJIS],[MARCENAC]. Expectedly, "active" operations are implied here, otherwise the principles of dynamic wavelength provisioning, and hence MPL(ambda)S/GMPLS are largely inapplicable. Specifically, TDM circuit/timeslots are now replaced by wavelength-based lightpath entities. These requirements can be met by either using optical add-drop multiplexer (O-ADM) or optical cross-connect (OXC)/wavelength routing switch (WRS) nodes. With regards to the latter, purely optical nodes can also be further classified as photonic cross-connect (PXC) nodes. The latter types (OXC, PXC) are well-suited to larger inter-ring connection (i.e., switching) applications, which require added (mesh) spatial switching capabilities. Note that the terms O-ADM, OXC, and PXC in the context herein are used to refer to complete ring node systems (and not just node sub-systems as may be done in a more detailed context).

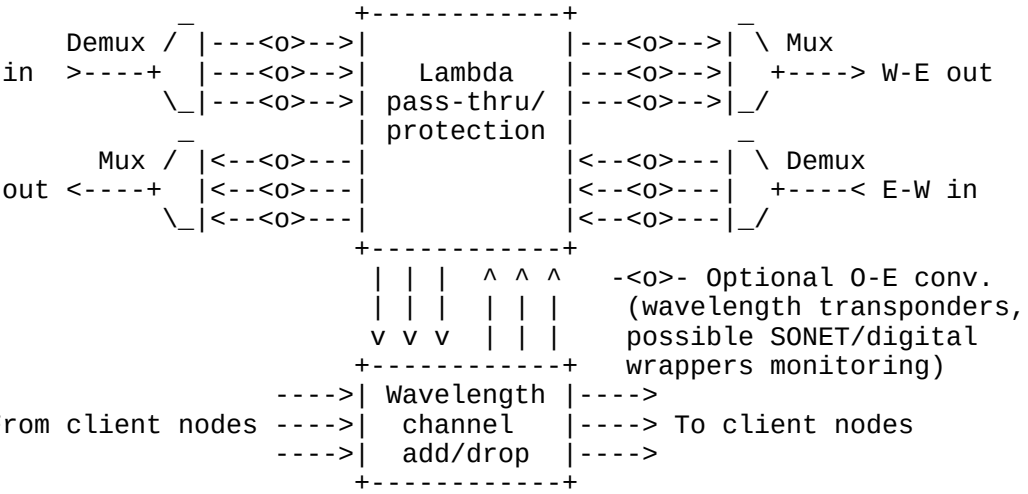

Figure 1: Sample optical ring node (2-fibers shown)

A generic overview of a two-fiber optical ring device is shown in Figure 1 and can easily be extended for four-fiber rings. Optical demux (mux) devices split (combine) wavelength channels (wavelength bands) from incoming (outgoing) fibers and connect to a wavelength channel (band) add/drop/protection unit. This stage can be implemented using a variety of techniques, such as optical switches (e.g., micro-electro-mechanical systems (MEMS), bubble, thermo-optic, etc), digital/electronic cross-point switches (DXC/EXC), or simpler 2x1 switching devices. The add/drop channels help to form the access stage of a ring node and this is where signals are mapped/de-mapped onto/from wavelength transmitters/receivers. Optionally, O-E designs can perform edge SONET/SDH (or digital wrappers [G.709]) payload mapping at the access stage.

Overall, optical ring nodes can exhibit many different levels of

functionality. For example, purely optical add-drop/switching fabrics are incapable of performing wavelength conversion but offer true signal format transparency. Conversely, EXC-based designs using opto-

electronic (O-E) conversion techniques will not have wavelength interchange restrictions, but will reduce signal format transparency. Therefore, as a tradeoff, "hybrid" designs are also possible, using EXC switches or tunable lasers to offer partial wavelength conversion capabilities for selected wavelengths and/or links. Moreover, numerous studies have shown that partial wavelength conversion capabilities yield network blocking probabilities close to those achieved with full wavelength conversion switches (i.e., O-E based). Hence, for the foreseeable future, optical networks will comprise of non-conversion and conversion-capable devices. Note that all-optical wavelength conversion techniques are also being actively researched, but commercial components are not yet available. Given all these variations of optical ring nodes, is important to define an optical ring framework that, to the extent possible, is independent of implementation and encompasses all (or as many of) these possibilities.

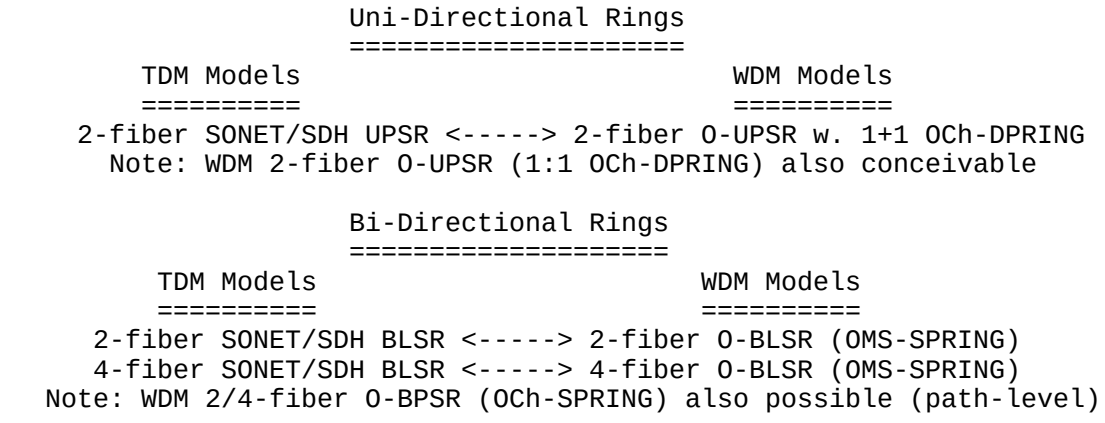

Figure 2: Mapping between SONET (SDH) and optical ring architectures

To date, the ANSI T1X1 and ITU-T SG15 have been most active with regards To work/proposals for optical ring architectures, e.g., see [CHEN], [CVIJETIC1-2],[SOULLIERE]. Although this work represents a good starting point, detailed standards (comparable to SONET UPSR, BLSR) are yet to emerge. Overall, optical ring proposals are classified into two major types, namely dedicated protection rings (DPRING) and shared protection rings (SPRING), and this delineation is re-used here to define the conceptual framework. The general relationship between SONET/SDH and proposed/emerging WDM (optical) shared/dedicated ring architectures is shown in Figure 2, and details are discussed subsequently. More specific provisioning (signaling) requirements are treated in Section 3. Note that the terms optical channel and lightpath are used in an interchangeable manner to represent wavelength circuits. Furthermore, the prefix "O" is used to identify "optical" ring concepts, in order to clearly discern them from existing TDM ring (SONET/SDH) schemes.

2.1 Dedicated Path Protection Rings (DPRING)

Dedicated protection rings are relatively simple in design and usually associated with two-fiber uni-directional (path-switched) O-ADM rings, O-UPSR/2. These rings can implement "edge-to-edge" wavelength channel

protection, and are therefore more commonly termed as optical channel DPRING (OCh-DPRING) [ARIJIS]. Note that the term "edge-to-edge" is chosen here, referring to a "sub-network" connection (SNC) entity, since it is most germane to a single ring (domain) and not necessarily a complete "end-to-end" client connection, see [XUE]. Both the 1+1

(non-signaled) and 1:1 (signaled) protection switching paradigms can be used herein. Each fiber in a OCh-DPRING carries wavelength channels in counter-propagating directions, with one fiber each for working and protection channels. The 1+1 OCh-DPRING solution is similar to SONET UPSR rings, with bi-directional connections consuming wavelength resources on all fibers, i.e., permanent head-end bridging. This OCh-DPRING scheme is shown in Figure 3 for a uni-directional channel. Note that an all-optical OCh-DPRING will likely require the same wavelength value on the working and protection path (i.e., unless ingress traffic bridging is done onto two separate wavelength transmitters). Since receiver-based switchovers are performed, no complex signaling protocols are required for 1+1 optical protection unless 1+1 bi-directional switching is employed (see [MANCHESTER1]). However, there is normally an added power penalty when performing optical head-end bridging, e.g., if a single laser's output signal is head-end bridged [SOULLIERE].

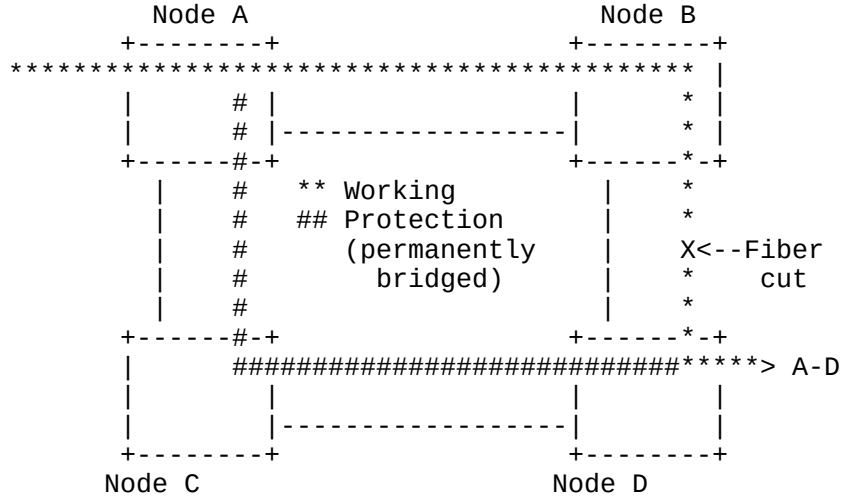

Figure 3: 1+1 wavelength path protection (2-fiber OCh-DPRING)

Additionally, signaled 1:1 protection is also conceivable for the OCh-DPRING, essentially re-using protection wavelengths for lower-priority traffic, i.e., head-end switching [SOULLIERE]. This requires an optical APS signaling protocol that has yet to be specified, a major task. However, note that overhead bytes have been proposed in the ITU for the OMS level, as per G.709 [G.709], and these can be used for conveying APS signaling. Although 1:1 channel protection improves upon idle resource utilization here, it still has limited spatial wavelength re-use and is rather disruptive (i.e., full ring/path switch can affect many users, albeit lower pre-emptable priority). The 1:1 OCh-DPRING structure is shown in Figure 4, where the lower-priority lightpath C-D occupies a protection wavelength/span for lightpath A-D. Overall, however, signaled protection is mostly proposed for more advanced shared

(line, path) ring architectures, Section 2.2. Note that the co-existence of both 1+1 and 1:1 OCh protection mechanisms in the same two-fiber DPRING may also be possible (i.e., since the underlying fiber/ wavelength plan is the same). However this issue requires further, more careful investigation of non-homogeneous ring behaviors.

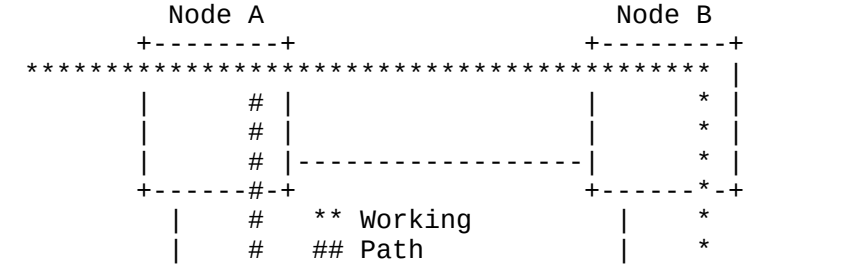

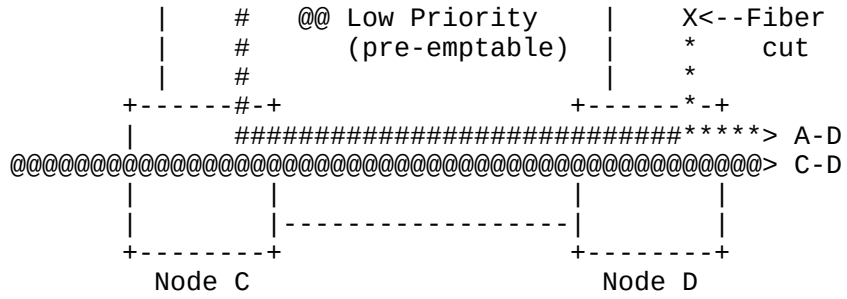

Figure 4: 1:1 wavelength path protection (2-fiber OCh-DPRING)

Note that depending upon the ring node's fault detection mechanism, switchover signaling can be actuated using a variety of methods (see Section 2.4, Section 3.2). For example, translucent designs using "inband overhead" monitoring (as defined by SONET B1-bytes [GR1230] or digital wrappers defect indicator bytes [G.709]) can detect progressive signal degradation and estimate bit-error rate (BER) values, etc. Inband overhead byte monitoring can also be used to detect progressive signal degradation. Alternatively, for transparent optical rings, optical monitoring techniques, such as power or signal-to-noise ratios (SNR) can be used to detect fiber (or wavelength) faults, Section 2.4.

In summary, the OCh-DPRING scheme requires full (100%) protection resource overhead and cannot achieve spatial re-use, somewhat akin to SONET UPSR rings. Hence, the OCh-DPRING scheme is best suited for hubbed traffic demands, where wavelength counts (and not spatial traffic demand distributions) are the dominant factors.

2.2 Shared Protection Rings (SPRING)

Shared protection ring (SPRING) architectures are designed to improve upon spatial resource utilization over UPSR designs. These rings are derived from SONET BLSR rings and are usually more complex, requiring active signaling for fast recovery. Overall, two shared ring schemes have been proposed, namely at the optical multiplex section (OMS) and the optical channel (OCh) level, respectively. In all such schemes, bi-

directional connections between two endpoint nodes must traverse the same set of intermediate nodes. Details are now presented (see also [ARIJIS]).

2.2.1 Optical Multiplex Section-Shared Protection Rings (OMS-SPRING)

Fiber cuts are one of the most common faults in ring networks, and given the increased multiplexing of DWDM systems, it is very desirable to also protect at the fiber span (i.e., OMS) level. Since per-channel protection switching can involve excessive signaling for large channel counts, fiber (i.e., optical line) protection can be much more scalable. Fiber protection basically provides an alternate fiber path between two adjacent nodes experiencing a fiber cut, and usually also requires signaling between the two end-points of fiber cut. Fiber protection is best applied to "fiber-rich" four-fiber rings, although two-fiber schemes are also possible. However, carefully note that line protection requires fiber fault detection and isolation capabilities, unlike end-to-end channel protection. A variety of OMS shared protection rings are possible, termed OMS-SPRING, and details are presented.

Two-fiber OMS-SPRING line (fiber) protection schemes, termed herein as O-BLSR/2, are very similar conceptually to SONET BLSR/2 designs. For example, to permit resource sharing and (intra-fiber) coordination between working/protection channels, these rings require a wavelength numbering/assignment scheme to effect a grouping between working and

protection channels. This essentially creates two "virtual fibers" from each physical fiber, albeit each with only half the number of wavelengths. The rules for such a partitioning are somewhat similar to those for timeslot partitioning in SONET BLSR/2 rings. Specifically, each fiber has an equal number of working and protection wavelengths traveling in the same direction, and the working wavelength group in a given fiber corresponds to the protection wavelength group in the other fiber. This implies that opposing directions will be routed on the same side (i.e., through common nodes) but use different fibers. Note that the intra-fiber wavelength plan requirement of O-BLSR/2 rings complicates ring node design (e.g., if O-E based or all-optical based). Also, for all-optical rings, added wavelength numbering qualifications are required to enforce the wavelength-continuity constraint. In particular, the actual wavelength values within each group (working, protection) have to match each other, thereby precluding the need for wavelength conversion upon switchover events. For example, the first (W/2) fiber wavelengths can be assigned for working channels, whereas the last (W/2) wavelengths can be assigned for protection channels. This condition can be relaxed for translucent (O-E) node designs, where wavelength values can also be interchanged upon protection switching.

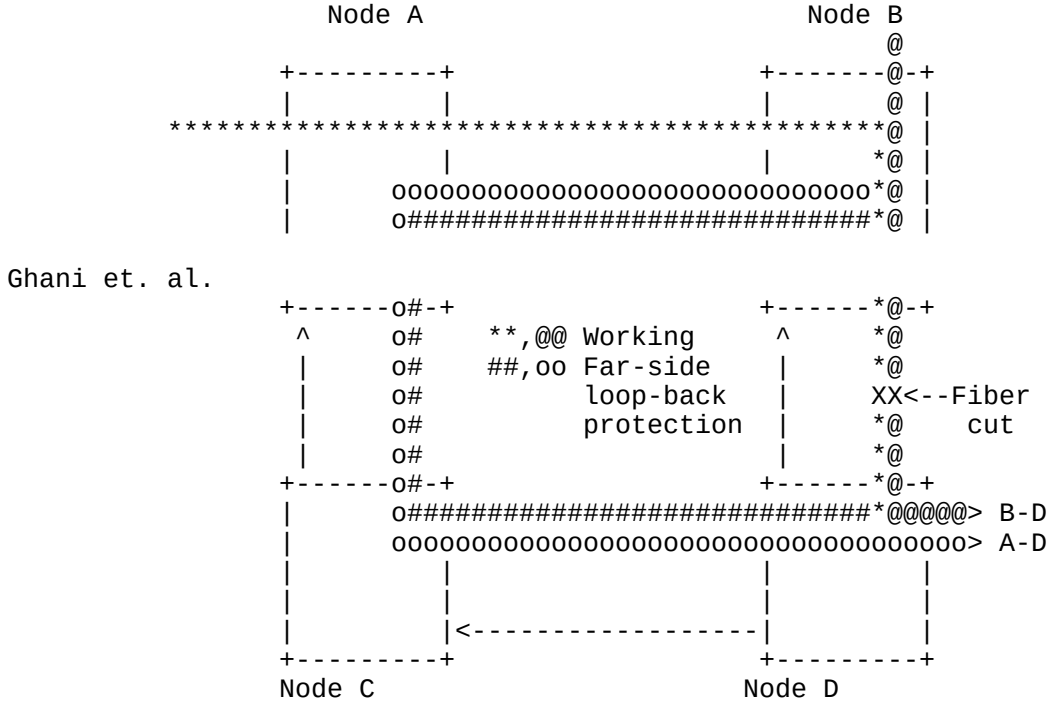

Figure 5: Loop-back span protection (O-BLSR/2)

Two-fiber OMS-SPRING (O-BLSR/2) protection is possible using "loop-back" protection (i.e., conceptually similar to SONET BLSR/2 rings), namely, all failed fiber wavelengths are re-routed on to the protection fiber route on the counter-propagating side of the ring, see Figure 5. Again, wavelength continuity concerns may arise for all-optical rings. As expected, protection signaling is done on the far-side. Albeit an alternative, optical loop-back protection, however, is not very attractive since it increases the distance and transmission delay of the restored channels (nearly doubling path lengths, as per SONET BLSR/2). Related analog degradations and other parameters will likely further hinder applicability, especially in all-optical rings. Furthermore, loop-back protection fully consumes protection fiber resources and limits recovery to single fiber-cut faults at any given time. As a result, it is unlikely that two-fiber OMS loop-back schemes (O-BLSR/2) will see much favor in practical settings.

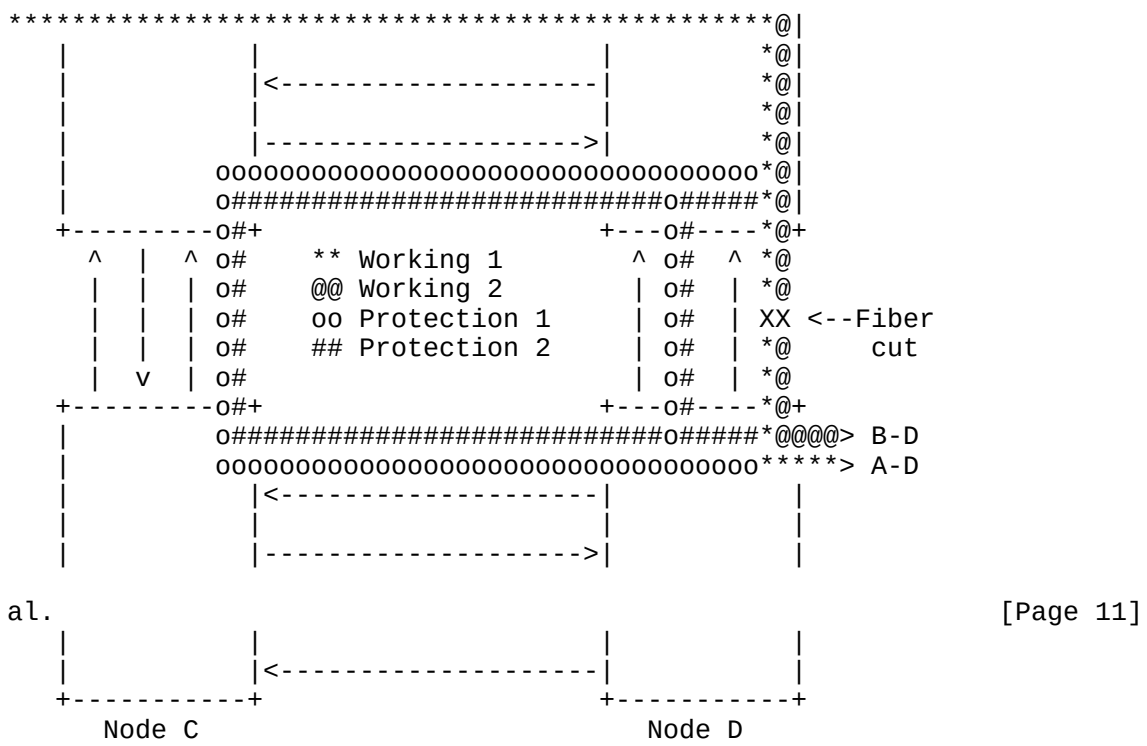

Figure 6: Span protection, loop-back and near-side (O-BLSR/4)

Span switching is much more attractive for four-fiber rings since protection fibers have the same wavelength directionality as working fibers. By logically extending SONET BLSR/4 architectures, four-fiber OMS-SPRING schemes can also be defined to protect against fiber cuts, termed O-BLSR/4. Clearly, these rings can implement loop-back (i.e., far-side) span protection by simply re-routing all failed working wavelengths on a fiber onto their associated, counter-propagating protection fiber (like O-BLSR/2). This will extend the channel routes as shown in Figure 6 (e.g., lightpath A-D protection via A-B-A-C-D, lightpath B-D protection via B-A-C-D). Far-side loop-back switching is especially attractive if all working-side fibers are cut (e.g., conduit fault), but again suffers from increased analog degradations.

Unlike two-fiber rings, four-fiber OMS-SPRING designs also enable more interesting and less-disruptive "near-side" protection switching. This form of protection switching is largely designed to protect against fiber (and channel) failures, and not node failures. One of its main purposes is to relegate protection signaling actions to the failed side of the ring (i.e., working, near-side). In other words, the "dual-directional" nature of fiber diversity of the four-fiber ring is exploited to maintain the same edge-to-edge node route between working and protection paths. Near-side protection switching is a generic concept that can be applied on both the line and path levels. The line-level case is illustrated for the O-BLSR/4 scheme in Figure 6, where the failed wavelengths are routed to the same-direction protection fiber on the near-side (only single direction shown for two working lightpaths A-D, B-D, traversing the outer working fiber). Overall, near-side line switching improves resource efficiencies since it does not disrupt traffic along the whole (long-side) protection route, as per loop-back techniques. However, near-side switching is less robust since it can only protect against working fiber faults, and not those that may also affect near-side protection fibers.

2.2.2 Optical Channel-Shared Protection Ring (OCh-SPRING)

Bi-directional SONET rings have only considered line protection since individual channel (time-slot) failures within the ring were considered relatively rare. However, when considering the types of failures that

optical rings may be required to protect against, yet another type of SPRING design may emerge as a more viable alternative to a simpler optical analog of two- and four-fiber multiplex-section SPRING's, namely the optical channel-level SPRING (OCh-SPRING). Specifically, for the case of both all-optical and transponder-based (i.e., translucent) optical rings, the need arises to protect against isolated opto-electronic (source or intermediate wavelength transmitter/receiver) failures that will affect only a single optical channel entity at a time. This implies the need for a protection architecture that performs OCh-level switching, based upon OCh-level indications, independently for

each optical channel on the ring. As the OMS-SPRING switches based upon OMS-level failure indications, and switches all optical channel entities as a group within the OMS, it is incapable of protecting lightpath connections independently of one another based upon OCh-level failure indications. A protection architecture that can protect optical channels individually based upon per-channel failure indications is the OCh-SPRING [MANCHESTER1], as facilitated by per-wavelength routing/processing/monitoring capabilities. In other words, this is essentially the optical bi-directional path switched ring (O-BPSR) concept for shared protection rings. Again, two variants of the OCh-SPRING are possible, namely for two- (O-BPSR/2) and four-fiber (O-BPSR/4) rings.

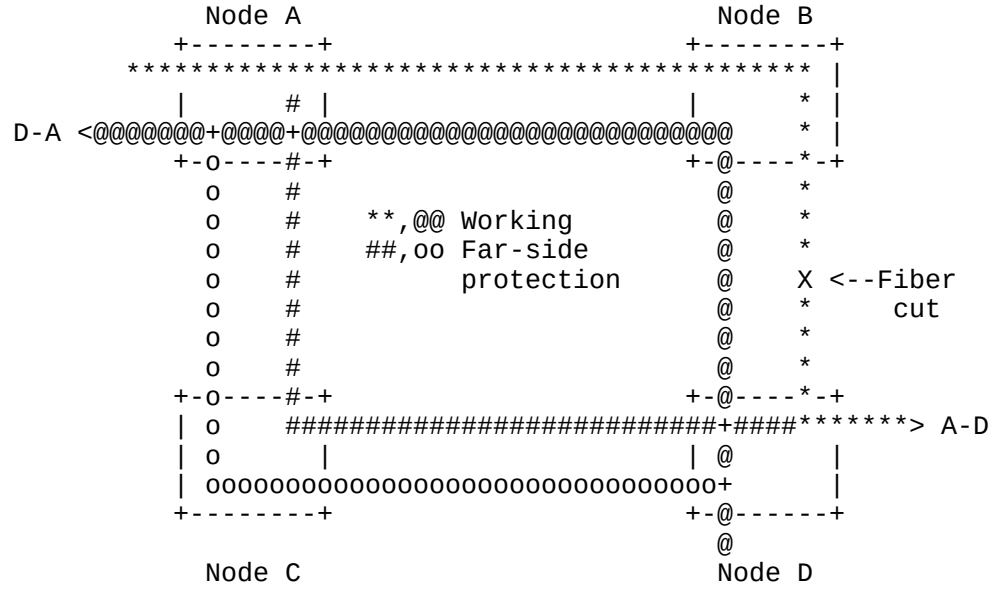

Figure 7: Bi-directional path protection schemes (O-BPSR/2)

The O-BPSR/2 proposal is based upon 1:1 wavelength protection (i.e., signaled switchover) and utilizes the same wavelength plan as its line-switching counterpart, O-BLSR/2. The scheme implements full bi-directional edge-to-edge switching on the "far-side" of the ring, i.e., on the side away from the fault. For example, lightpath A-D (D-A) re-routed from A-B-D (D-B-A) to A-C-D (D-C-A), Figure 7. Again, depending upon the translucency level, wavelength continuity may be required along all edge-to-edge routes. Channel switchovers are performed for channels in both directions, regardless of which one actually failed. This is more beneficial in case of both working and protection fiber cuts on the working side, e.g., conduit cuts. Lightpath faults are detected by downstream nodes (see Section 2.4), which then effect switchover actions via expedited upstream signaling along the far-side (albeit no standards are defined yet). Clearly, far-side edge-to-edge path switching will be the most disruptive, since (lower-priority) traffic and fast signaling are required on the opposite side of the ring. However, far-side switching can protect against intermediate node failures. It should, however, be noted that signaling latencies will dictate maximum ring sizes (node count limits) for all

edge-to-edge ring switching schemes.

By utilizing the SPRING wavelength plan, O-BPSR/2 solutions also significantly improve spatial resource sharing over their UPSR counterparts, especially for "non-hubbed" traffic demands [ARIJIS]. Furthermore, differing levels of protection resource sharing can also be allowed. For example, obviously, idle protection wavelengths can be used to carry lower-priority pre-emptable traffic (termed 1:1 shared). Furthermore, protection wavelengths (on the far-side) themselves can be shared between multiple working channels. This achieves a "1:N" protection resource multiplexing effect for each wavelength (hop), and not just the complete protection path. This feature improves resource efficiency significantly, especially for all-optical rings without wavelength conversion, and will yield reduced call setup blocking probabilities. Additionally, customers transporting relatively low-priority (cost) traffic may be satisfied with pre-emptable lightpath connections, namely ones that can be dropped (squelched) in the favor of allocating resources to higher-priority connections (e.g., at call setup request or protection switching times). These multiple levels of protection/sharing (e.g., 1:1 dedicated, 1:1 shared, 1:N shared, pre-emptable) will allow operators to define several differing classes (or grades) of service. Further variations to the O-BPSR/2 framework are for future study.

Four-fiber optical path-switched rings (O-BPSR/4) have also been defined and can provide more advanced capabilities. These rings also require a wavelength numbering plan, and it is best to choose one that mirrors the counterpart four-fiber OMS-SPRING scheme (and therefore, conceptually parallel to four-fiber SONET ring time-slot assignments). Specifically, due to increased fiber resources, there is no need for intra-fiber wavelength partitioning, and therefore, two counter-rotating fibers (i.e., all wavelengths) can be reserved for working and protection traffic, respectively. Differing directions of a bi-directional connection are therefore routed on different working fibers between the same ring nodes. O-BPSR/4 protection schemes are largely variations of 1:1 protection switching scheme, as illustrated for a single direction in Figure 8. Furthermore, strong conceptual parallels exist with O-BLSR/4 line-switching concepts with regards to protection routing. In the most straightforward case, ubiquitous far-side path switching can be implemented, with both paths (of a bi-directional circuit) being switched over on to their corresponding protection fiber routes on the opposite side of the ring (as per OCh-SPRING O-BPSR/2). Far-side path switching can protect against failure of all standby resources on the working side (i.e., complete multi-fiber ring conduit cut).

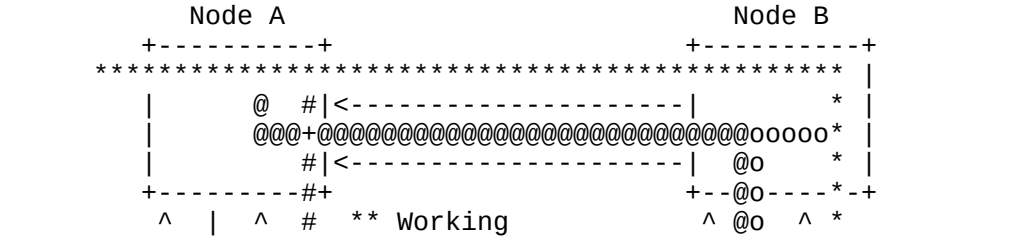

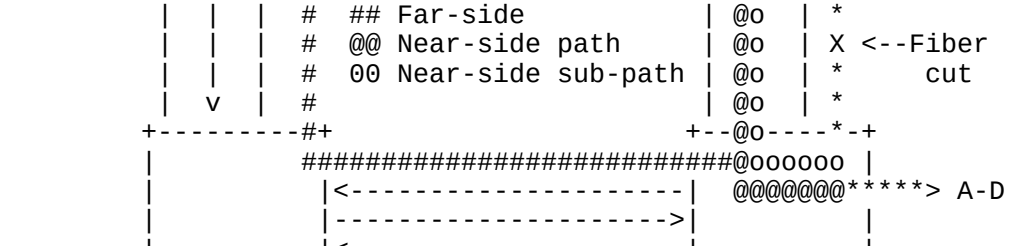

Figure 8: Path protection schemes (O-BPSR/4), one-direction shown

Furthermore, four-fiber ring near-side protection switching concepts (Section 2.2.2) can also be applied on a path-level. In fact, more variations are possible, namely edge-to-edge and intermediate near-side path switching. The first, O-BPSR/4 edge-to-edge near-side path switching, routes both the working and protection lightpaths from the working fiber on to the protection fiber in the same direction, see Figure 8. Meanwhile, to minimize the drop-rate of possible low-priority connections using the protection wavelengths (and to an extent, to also reduce signaling overheads), intermediate near-side path switching can be considered. This form of protection switching only performs partial working path re-routing, as illustrated for lightpath A-D in Figure 8, where the failed segment B-D is switched to the associated wavelength on the protection set of the second fiber. This largely limits protection signaling to the two adjacent ring nodes, but cannot overcome node failures. Due to the bi-directionality requirement, both channel directions are switched regardless of if one or both failed. For both forms of (O-BPSR/4) near-side switching, all-optical nodes will have to ensure that wavelength continuity considerations are met. Note that O-BPSR/2/4 concepts can also be applied at the wavelength band level and this can be studied further.

Depending upon the optical ring node designs, protection resource sharing can also be achieved for four-fiber rings (and hence multi-level service definitions). For example, some have proposed a straightforward fiber protection implementation using 2x1 fiber switches before any mux/de-mux stages (Figure 1). This implementation precludes complimentary wavelength-level processing capabilities (such as pass-through, add, drop), and hence will hinder wavelength sharing on protection fibers (more restrictive). Clearly, in order to share wavelengths on the protection spans and improve resource utilization (i.e., for OMS/OCh-SPRING O-BLSR/4), per-wavelength processing is required for both working and protection fiber channels. This essentially means that a fiber cut can also be handled by multiple channel-level re-routing actions, although implementation concerns can be more challenging. Here, "batch" control commands (to switch multiple wavelengths) can be developed, since all wavelengths on a failed span are re-routed along a common route. Furthermore, sharing protection resources will require larger add/drop or switching fabrics (Figure 1). Clearly, "full-blown" four-fiber rings can support many more users of any given service category, as compared to two-fiber ring schemes.

In general, operators may also want to provision multiple (ring) protection schemes off of the same fiber infrastructure. In this regard, a generic limitation of fiber protection is that it treats all wavelengths (channels) in a fiber equally, and therefore alone it cannot achieve (channel-specific) service differentiation. However, span protection can co-exist with channel protection if a priority mechanism is used to "arbitrate" between the two recovery mechanisms. Various such mechanisms are conceivable, either signaling or non-signaling based (for possible further study). For two-fiber UPSR schemes (O-UPSR/2), span protection is not applicable for 1+1 channel protection. However, (signaled) bi-directional OMS/OCh-SPRING schemes (i.e., those using the O-BLSR/2 or O-BLSR/4 wavelength plans) can support both mechanisms, with idle protection spans carrying lower-priority traffic. As an example, co-existence between channel (O-BPSR) and line (O-BLSR) protection mechanisms can be achieved in the protection signaling specification via an appropriate "priority" mechanism. Typically, span protection should be done first since it represents "lower-level" (or more coarse) recovery. This can be

achieved by inhibiting all channel failure message responses and only responding to fiber/span failure messages. Further details and intricacies are outside the current scope and require careful future considerations.

2.3 Signaling Channel Architectures

Optical rings allow for significant latitude in signaling channel architectures, and two overall categories are possible, namely in-band and out-band signaling. In-band signaling requires the use of overhead framing bytes (as reserved in a SONET/SDH or OCh (digital wrapper) frame header) that are reserved on data channels. Such mechanisms are best suited for O-E based optical node designs, where edge client signals are mapped into synchronized electronic frames that already contain the required signaling bytes. Alternatively, out-band signaling can be used to more clearly decouple the data and control planes. Out-band signaling can be done using a dedicated control wavelength, commonly termed as the optical supervisory channel (OSC), or even via a physically separate, out-of-band network (such as an Ethernet LAN). Note that some have termed the OSC approach as in-band also, since the control wavelength (typically 1510 nm) "physically" resides in the fiber itself. However, as far as data-control channel interaction is concerned, there is no interaction and hence this approach is termed as out-band. Note that the OSC channel will require appropriate hardware support (filters, receivers, laser transmitters, etc). Recently efforts are beginning to emerge for defining a broad range of OSC standards, see [FREDETTE],[SZERENYI].

In general, an out-band OSC-based approach is more attractive to some since it allows for genuine service-transparent optical ring paradigms, also stated in [SOULLIERE]. Specifically, this approach utilizes the same fiber plant, precluding limitations with a completely external out-of-band signaling network, yet still permitting true client wavelength (payload) transparency. However, out-band signaling systems need to ensure adequate bandwidth levels for increasingly large data

wavelengths counts (in the hundreds). As a result, further considerations are needed for the out-band OSC channel approach (see T1X1 generic proposal [SZERENYI]). For example, some have proposed using sub-rate (synchronous) TDM circuit streams to partition and guarantee OSC bandwidth to all data wavelengths, typically for SONET/SDH in-band signaling transport. Others have proposed (asynchronous) packet signaling on the OSC channels. In either case, whenever fast recovery guarantees are required, some form of bandwidth scheduling, be it TDM or packet scheduling (possibly with priority drop mechanisms), will likely be required on the OSC channels. This introduces added, but necessary, complexity concerns. Additionally, signaling channel robustness is also of concern and here, backup control channel provisions are also being considered, see [LANG],[FREDETTE].

2.4 Fault Detection and Isolation

The ability to quickly detect, and preferably localize, fault events is crucial to achieving fast service recovery. So far, the above discussions have focused more upon switchover actions, and assume that fault detection (possibly localization) is already done. Now a key differentiating aspect of optical networks, unlike SONET/SDH networks, is that more variations of fault detection and localization mechanisms can be utilized (as will be detailed subsequently). In order to allow for full flexibility, it is therefore preferable that network-level optical (ring) fault recovery, notification, and detection/isolation mechanisms be clearly separable and independent of each other (more detailed discussion in Section 3.2.2). A review of the various monitoring solutions is now presented, see also [CEUPPENS],[GHANI2], [MANCHESTER1].

Many first-generation and even current-generation WDM systems simply re-use existing SONET/SDH schemes to detect and isolate channel faults inside the core optical (ring) network. These solutions include re-using B1 byte monitoring and loss of framing (LOF)/loss of signal (LOS) alarm information. Such solutions have been commonly referred to as opto-electronic (O-E) and/or frame-monitoring schemes [GHANI2], [CUEPPENS], since they require that all monitored data wavelengths be "opaque" or "translucent". The digital wrappers approach, which represents a counterpart to SONET/SDH framing, also essentially embodies a similar O-E based solution, e.g., forward/reverse defect indicator (FDI/RDI) bytes, etc [G.709]. Since most operators are quite familiar with SONET/SDH overhead monitoring, O-E type schemes have one definite advantage, namely, well-defined standards. This permits faster vendor interoperability (albeit not considering proprietary usages of various unused overhead bytes). However, opaque monitoring represents some serious limitations. First of all, per-channel electronic overheads usually pose increased systems costs and power requirements. More importantly, such designs are largely unscalable to very large, ultra-dense WDM systems, and generally inhibit evolutions to truly transparent networks [BHANDARI]. Furthermore, O-E monitoring requires mappings for all client payload types. Now although well-defined encapsulations exist for IP, ATM, and Frame Relay protocols, further extensions may be necessary, e.g., for new gigabit Ethernet standards, ESCON, cable video signals, etc. Note however, that O-E monitoring may be suitable for monitoring out-band control channels, since these

are electrically terminated at each node.

To get around the limitations of opaque monitoring, various vendors have proposed optical monitoring schemes using non-intrusive signal tapping setups. These solutions are particularly germane to monitoring non-SONET/SDH payload types, and analyze such parameters as fiber/wavelength (optical signal) power levels, optical signal-to-noise ratios (O-SNR), Q-factors, etc (see [GHANI2],[CUEPPENS]). For example, power monitoring can detect fiber cuts in under 10 ms, and this is capable of meeting the most stringent of recovery requirements. Power monitoring is also termed as loss of signal (LOS)/loss of light (LOL) fault, and can trigger various protection actions, such as 1+1 receiver switchovers or fault/alarm messaging. However, although optical monitoring is of high interest to vendors and service providers alike, the current lack of standards (and to an extent, advanced features) is hindering its widescale adoption. Most current all-optical solutions simply perform line power-level monitoring, and are therefore best-suited for O-BLSR support. Although per-wavelength power-level monitoring can also be done, this approach is not cost-effective at all for large channel counts (i.e., hundreds of wavelengths, as per DWDM). Nevertheless, such per-wavelength monitoring capabilities inside the network core will be needed in order to support transparent O-BPSR schemes, in the absence of any O-E (SONET) frame monitoring. Although optical monitoring resources (such as spectrum scanners) can be shared between multiple fibers/wavelengths to control costs, the resulting fault detection times will be much longer. Further adding in transient switching times (milliseconds range), achieving the "50 ms" SONET/SDH recovery time ceiling may prove difficult. Regardless, since optical component technologies are continually undergoing rapid improvements and miniaturization, it remains to be seen if these concerns, indeed, may be mitigated in the foreseeable future. Moreover, some network designers may actually want to use optical monitoring techniques to complement capabilities in opaque networks, e.g., observe transponder performance/behaviors to predict failure conditions.

A more timely and cost-effective alternative may be to perform "edge" channel (OCh) fault isolation, as suggested in [BHANDARI]. Specifically, no channel-level monitoring is performed inside the network (between

the edge points), thereby precluding excessive (expensive) O-E conversions or OCh-level optical monitoring. Instead, fault isolation is only done at the channel edge points. This can be achieved using a variety of techniques, implemented in the appropriate receiver/interface cards (either optical power monitoring or electronic frame monitoring after O-E conversion). All that is required is that the channel protection sub-path be "dis-joint" (Section 3.1.2) from the working paths. This approach is very attractive in all-optical networks (both ring and mesh), where operators request service transparency with "SONET-like" recovery times. Also, this solution is well-suited for "edge-to-edge" channel protection schemes, such as those detailed for O-UPSR/2 or O-BPSR (far-side, edge-to-edge near-side) setups. Note here that the "edge" regions can either comprise single rings (i.e., SNC portion) or a series of rings (or hybrid ring-meshes), forming a larger (optical) sub-domain (also see discussions in Section 3.3.1).

3. Dynamic Provisioning Issues

Recent developments have extended the MPLS protocol framework from the packet/flow switching domain to the optical lightpath switching domain. Termed multi-protocol lambda switching, MPL(ambda)S [AWDUCHE],[GHANI1], [RAJAGOPALAN], this work draws analogies between labels and wavelengths and intends to re-use/extend signaling and resource discovery protocols for the optical domain. Optical nodes (such as cross-connects or add-drop multiplexers) use IP addressing schemes and run extended MPLS routing and generalized signaling protocols, i.e., lambda switch routers. More importantly, recent proposals for a generalized MPLS (GMPLS) [ASHWOOD1] framework are furthering this trend, extending basic MPLS concepts to provision more generalized label switched path (G-LSP) entities, e.g., TDM circuits via SONET ADM’s, lightpaths via wavelength cross-connects, etc. In parallel, there has been a lot of focus on defining LSP recovery schemes for MPLS networks, albeit, mostly at the packet flow level [DOVOLSKY],[OWENS1],[KINI2]. Possibly, these schemes can also be investigated for their potential applicability to "optical LSP" (i.e., lightpath) protection. However, in general, due to the IP-centric origins of the MPLS framework, the above work is generally tailored for mesh (optical) networks, even though its generic nature does not preclude specialized, topology-specific applications or extensions.

Given all the variations of optical rings (Section 2), it is very advantageous to develop a comprehensive provisioning framework and align it with the larger MPL(ambda)S/GMPLS architecture. In particular, two signaling mechanisms are required for optical rings:

- First, signaling is required for dynamic ring configuration and lightpath provisioning operations, such as setup/takedown. These mechanisms must specify both the working and protection entities of the lightpaths and incorporate all of the intricacies of (ring) protection switching mechanisms. Ring resource management will also be a critical part of the provisioning stage.

- Second, a signaling mechanism is required to perform the automatic protection switching (APS) actions, as determined by related ring protection schemes, Section 2.

An initial look at these two crucial topics is now presented and is intended to serve as basis for further, more defining work.

3.1 Channel Setup Requirements

>From an operator’s point of view, ring networks will likely interface to (or even migrate into) mesh networks in the near future (e.g., metro rings to regional/long-haul mesh). Given the likely adoption of MPL(ambda)S/GMPLS type protocols for optical mesh provisioning, it is

prudent to choose likewise for ring networks, thereby enabling an even closer interworking. For optical ring channel setup/takedown, the overall provisioning capabilities developed under the ubiquitous MPL(ambda)S/GMPLS frameworks are quite applicable. Namely extensions to MPLS signaling protocols are already being proposed to handle the specifics of optical lightpath routing [ASHWOOD2-3],[KOMPELLA1-3],[YU]. However, provisioning ring lightpaths (working, protection) will require

added considerations, and some of these are now considered more closely.

3.1.1 Signaling Extensions

At the core of ring channel provisioning is the concept of a service definition, as commonly extended through various means, e.g., either from a element/network management system (EMS/NMS) or via an "optical user network interface" (O-UNI). Since the latter approach has been the focus of many standardization efforts, discussions herein will give it closer consideration. Recently, many "O-UNI" definitions have been tabled for optical networks, proposing various new "signaled" interfaces [ARVIND],[ABOULMAGD],[MCADAMS],[XUE] along with expanded features in the MPLS LSP setup messaging [YU],[KOMPELLA3]. In short, service definitions supply the signaled information "attributes" for subsequent channel setups. Channel setup, in turn, implies the more general category of routing and wavelength assignment (RWA) and policy control (Section 3.1.3). Setup information usually includes many details, such as the channel framing type (e.g., SONET/SDH, digital wrappers, IEEE Ethernet, etc), bit-rate (2.5 Gb/s OC-48/STM-16, 10 Gb/s OC-192/STM-64, 1.0 Gb/s 10 Gb/s Ethernet, etc), protection type (shared, dedicated, enhanced, unprotected), and priority (non-pre-emptable, pre-emptable), etc, see [ABOULMAGD],[ASHWOOD1]. Provisions have also been suggested for indicating lightpath diversity levels (e.g., node, link, etc), see [XUE],[ABOULMAGD]. By and large, these generic attributes also apply to ring networks, although their detailed usages and applications require further considerations.

The above-mentioned service definition "attributes" need to be "mapped" into appropriate signaling messages in order to setup the lightpath channels (e.g., as per RSVP-TE, CR-LDP signaling [ASHWOOD2-3], [KOMPELLA3]). Here, a key step in this mapping will be to first translate the desired (requested) user lighpath attributes into appropriate ring channel request types, i.e., as per the various types of optical rings (Section 2). Specifically, users may request various channel priority or protection types (amongst other attributes), and these must be translated to the appropriate channel types given the underlying ring specifics. For example, a user request for a non-pre-emptable, non-shared protected channel in a O-UPSR/2 (two-fiber) setup may be translated into a simple 1+1 working/protection channel request. Alternatively, the same request in a O-BPSR/4 (four-fiber) setup may be resolved as a 1:1 working/protection channel request. Such mappings are required before any ligthpath routing can be performed (Section 3.1.3). Overall, the mapping of (signaled) channel attributes from user requests to the exact ring lightpath types is very implementation-specific and hence should not be the subject of standardization (i.e., vendor-value add feature).

Once the user request is properly mapped (on to the ring) and its lightpath route computed (Section 3.1.3), various MPLS LSP signaling capabilities can be exploited for the actual setup [ASHWOOD2-3]. Clearly, one such feature is explicit route (ER) signaling, which can explicitly indicate the required path and reserve resources.

Specifically, ER inserts the complete route specification in appropriate route specification objects (i.e., explicit route fields in RSVP-TE PATH or CR-LDP LABEL_REQUEST messages). Additionally, bi-directional channel

setup provisions have also been considered [ASHWOOD2-3],[GUO1], helping ensure that both uni-directional lightpaths of a bi-directional LSP/G-LSP traverse the same set of nodes. In conjunction with shared risk link group (SRLG) "disjointness" information (Section 3.1.2), this signaling feature is directly applicable to O-BLSR/2/4 and O-BPSR/2/4 setups (i.e., where bi-directional channels must traverse the same set of ring nodes). Sample proposals for lightpath setup signaling using appropriately defined TLV objects are presented in [KOMPELLA3],[YU] and the additional related references therein. Any extensions to the setup signaling message (object) types for ring channel provisioning need further study.

3.1.2 Resource and State Dissemination

In addition to the above setup information requirements, provisioning algorithms (Section 3.1.3) need to know the existing static topological details and available dynamic resource levels (as detailed in [CHIU], [BERNSTEIN]) in order to compute ring routes. Consider the first requirement. Examples of basic static topological information are the number of fibers, ring nodes, and their connectivity. For fiber elements, information is required to indicate the link type (transparent, service-aware), the number and location of supported wavelength channels (e.g., ITU-T grid spacing, offsets, guard bands), related analog metrics (loss, dispersion figures), etc. Meanwhile, for ring node elements, many (static) details are pertinent. Examples include the ring configuration type (O-UPSR, O-BLSR, O-BPSR, or multiple), number of fiber ports (e.g., incoming, outgoing, add, drops), fiber port protection type (1+1 protected or unprotected), type of ports supported (e.g., transparent, opaque), performance monitoring capabilities (e.g., optical, electrical, per-channel, per-span), signal regeneration (e.g., 1R, 2R, 3R), wavelength conversion capabilities (e.g., none, partial/selected, full), protection switching capabilities (e.g., per-channel, per-fiber, per-conduit), etc. Since ring schemes are intricately associated with the directionality and protection association (working, protection) of fibers or wavelength groups inside fibers, this information must also be incorporated.

In traditional data networks, interior gateway protocols (IGP) are used to disseminate static topology and dynamic resource information. Recent additions for supporting opaque link state attribute (LSA) definitions (RFC 2370) will help further facilitate extensions to "non-data" routing applications. More recently, many proposals have tabled extensions thereof for optical networks, and in fact, many of the above-discussed requirements (for static topology and dynamic resource information) have already been proposed within the context of mesh-routing MPL(ambda)S/GMPLS networks [KOMPELLA1-3]. For example, IGP provisions have been considered to indicate wavelength conversion capabilities and dynamic link-level resource (wavelength) utilizations/ levels. Such active resource updates are vital for dynamic ring RWA algorithms. Delineations between different link-level resource classes have also been proposed (i.e., active, free, reserved, pre-emptable wavelength sets), see [KINI1-2]. The actual control/ specification of wavelength plans can be done either statically (via

the NMS) or dynamically (e.g., based upon changing network topologies). As an application here, such resource class delineations can be leveraged to control intra-fiber wavelength plans (e.g., per O-BLSR/2, O-BPSR/2 schemes).

In addition, recently the concept of a shared risk link group (SRLG) definition has also been proposed to help identify risk associations between various entities, see [RAJAGOPALAN],[PAPADIMITRIOU1]. By using this information, adequate resource "disjointness" can be introduced into the constraint-based path computation (routing, Section 3.1.3) phase, thereby reducing simultaneous lightpath failures (e.g., between working and protection paths). Recently, a detailed, comprehensive

treatment of the SRLG concept has been presented in [PAPADIMITRIOU1], in order to formalize the link between risk groups and route computation. Here, two different hierarchical resource inference/diversity models are defined, namely physical (e.g., wavelength, fiber, conduit, etc) and logical (or geographical, i.e., node, zone, region). An encoding scheme is also presented for encoding/summarizing SRLG identifiers (e.g., between logical boundaries) along with possible mechanisms for risk assignment. Collectively, SRLG information and the associated lightpath risk derivation mechanisms are crucial for service provisioning in optical networks, given the high levels of traffic multiplexing and also resource co-location (e.g., wavelengths in fibers, fibers in conduits, etc), see also Section 3.1.3.

Overall, many of the above-described informational (IGP) extensions are also very applicable to optical ring networks. As these concepts mature, along with their usage definitions, their application herein will indeed be highly practical. Additional specifications or applications of such augmented LSA's are for future study.

3.1.3 Constraint-Based Routing/Path Computation

During the channel (i.e., LSP/G-LSP) setup phase, lightpath route computation is performed by utilizing the available network information (e.g., topology, ring-type, resource levels, risk groups, etc). Specifically constrained routing/path computation is required, and this can be deemed as a subset of the more generic constraint-based routing paradigm [GHANI1]. Here, the constraints are now more specific to optical parameters (e.g., topologies, wavelengths, converters, amplifiers, etc) and policies (e.g., as per SLA requirements). As an aside, generic policy control (management issue) can also be implemented (in addition to the compute-centric RWA processes) in order to enforce user SLA guidelines. A sample application using the well-defined, generic common open policy service protocol (COPS RFC 2748) is presented in [GHANI3].

To date, much detailed research has been done on the subject of constraint-based lightpath routing, technically termed as the routing and wavelength assignment (RWA) [ZANG] and/or virtual topology design [DUTTA] problem. In performing lightpath channel routing, typically, there are two sub-problems which have to be resolved, namely lightpath route computation and subsequent wavelength selection [DUTTA]. Overall,

many types of RWA algorithms have been proposed in the literature, ranging from complex optimization-type formulations, to constrained shortest-path methods, to various simplified heuristics. Usually, RWA algorithms aim to optimize specific objectives, subject to various broad constraints such as hop counts, wavelength re-use (in given network segments), propagation delays, protection/priority levels, residual resources, revenues, risk probabilities, etc. The objectives could either be resource oriented, namely high resource efficiency, or performance oriented, such as low blocking probability, etc. Moreover, many ring-specific lightpath routing algorithms have also been researched, see [MARCENAC] and related references. Clearly, any ring lightpath RWA algorithms will be very tightly coupled with the actual ring types (uni-directional, bi-directional), wavelength plans, wavelength conversion capabilities, and various other specific considerations. For example, in O-UPSR/2 rings for a 1+1 protected channel request, two "disjoint" paths must be found between the source and destination nodes, along with necessary permanent bridging/receiving resources at the endpoints. Alternatively, in O-BLSR/4 rings for a 1:1 shared long-side protected channel request, the RWA schemes will only need to search the long-side of the ring for protection channel routes, and this can include any assigned protection wavelengths. Note that the actual computation phase can be implemented in a variety of ways, such as distributed shortest-path/heuristic computations (e.g., specific renditions of Dijkstra

algorithms) or via centralized route/policy control servers.

Note that for (ring) protection schemes, further RWA considerations are required. Specifically, at setup time "joint" RWA algorithms are necessary for resolving the routes and associated wavelengths for both the working and protection (sub)paths, see [DOSHI] for sample proposal. These computations can (should) further utilize SRLG-based information to ensure adequate resource/risk diversity between working and protection channels, see [PAPDIMITRIOU1] Appendix. For example, (ring) protection paths require shared-resource (i.e., risk) separation from working entities, i.e., "disjoint". In this context, an entity can be a full edge-to-edge lightpath (as per O-BPSR/2/4 near/far-side and O-UPSR/2), a portion of a lightpath (i.e., sub-path as per O-BPSR/4 intermediate near-side), or a complete fiber span (as per O-BLSR/2/4). Moreover, SRLG definitions can be used to effect inter-fiber delineation between working and protection fibers (for the case of O-UPSR/2 and O-BLSR/4 rings), i.e., working and protection SRLG identifiers. Generic discussion of routing diversity (dis-jointness) is also presented in [DOVOLSKY],[OWENS2],[XUE].

Overall, it is highly likely that the lightpath routing algorithms themselves will not be the subject of standardization. Conversely, this is certainly an area of vendor-value add, and many suppliers will prefer implementing their own proprietary algorithms/policy control as best suited to their individual customer needs. Therefore, what needs to be standardized is more the actual informational framework required to perform proper lightpath RWA computation.

3.2 Protection Signaling

It is safe to assume that operators will demand SONET/SDH-type recovery

timescales for protected optical ring services (i.e., 50 ms ceiling), and meeting this stringent requirement is perhaps the foremost concern when trying to apply the MPL(ambda)S/GMPLS framework to optical ring control [GHANI2]. Now for most optical ring types (excluding 1+1 uni-directional O-UPSR/2 designs), millisecond recovery requires fast "APS-like" signaling capabilities, akin to the SONET/SDH K1/K2-byte APS protocol. Generally speaking, all such schemes can be subsumed under a more encompassing model, namely that of two (or more) switching end-point nodes and intermediate, physically disjoint protection resource(s). (This excludes loop-back switching techniques, which are largely deemed unfavorable for optical networks, Section 2.2.2). For example, for channel protection, the end-point nodes are either the source and destination nodes (O-BPSR/2, edge-to-edge near-side and far-side O-BPSR/4), or the appropriate intermediate nodes (intermediate near-side O-BPSR/4). Likewise, for span protection, the end-point nodes are simply the adjacent optical ring nodes. By developing appropriate switchover signaling capabilities to implement this generic model, conceivably most relevant ring protection schemes can be covered.

For the special case of optical ring networks, two possible options exist for implementing such fast protection switching. One is to develop enhancements to the existing RSVP-TE/CR-LDP LSP protection (survivability) signaling proposals and tailor them for "optical lightpath LSP" protection, termed herein as the direct interworking approach (originally proposed in [GHANI2]). The other would be to develop an altogether new, dedicated protection-switching protocol, namely an optical APS (O-APS) protocol, to complement the overall MPL(ambda)S/GMPLS framework. This new protocol would only perform protection switchover signaling for fault events but not any setup provisioning (relegated to existing setup signaling mechanisms, as detailed previously in Section 3.1). These two cases are presented in a broader context in Figure 9, which shows an example of multi-level recovery protocols. On the packet routing level, service recovery

actions can be performed by existing IGP path re-computation/re-routing schemes (longer convergence times). On the virtual circuit (i.e., packet LSP) level, emerging enhancements to RSVP-TE/CR-LDP signaling can effect improved recovery timescales (sub-second or lower). Finally, on the lightpath circuit level, recovery actions can be implemented either via further RSVP-TE/CR-LDP enhancements or an (above-mentioned) O-APS protocol, Figure 9. Further details are now discussed.

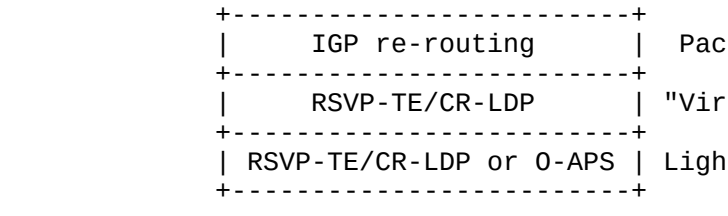

Figure 9: Service recovery protocols (packet, flow, circuit levels)

3.2.1 Direct Interworking

It is instructive to first briefly review MPLS LSP protection concepts. Clearly, simple non-signaled protection is possible by establishing

multiple (LSP) paths between source and destination LSR nodes and streaming data across these paths, essentially like 1+1 ("make-before-break") protection. However, more advanced protection signaling proposals are also beginning to emerge within the extended RSVP-TE and CR-LDP protocols framework [BHANDARI],[HUANG],[KINI1-2],[OWENS1-3]. The basic idea with MPLS LSP protection is to provision back LSP (sub)-paths, and in case of fault discovery, perform a signaled switchover. Generic protection switch LSR (PSL) and protection merge LSR (PML) nodes are defined and these entities define the edges of the protected LSP segments. Specifically, a desired LSP segment, termed working (or active) path [OWENS1], is setup for protection by having the PML/PSL nodes source and sink two distinct (sub)-paths, working and protection, as shown in Figure 10. As a generalization, PSL/PML node pairs can protect multiple LSP segments, termed protected MPLS traffic group (PMTG) [OWENS1], reducing signaling overheads for improved scalability. Downstream nodes detecting a fault event propagate a failure indication signal (FIS) in the upstream direction, containing a list of protected LSP's on the failed PMTG entity. Various timer mechanisms are used to control the inter-FIS packet timing, duration of FIS transmissions, and hold-off time for initial FIS indication, see [OWENS1] for discussions on timer settings. Upon receiving the FIS message, the PML node performs a switchover from the working to protection sub-paths for all affected LSP's specified in the PMTG. Additionally, a failure recovery signal (FRS) is also propagated after the fault has been repaired (along the same route as the FIS message). Similar timer mechanisms as with the FIS message also exist for the FRS message, and neither message type requires reliable transport, e.g., no TCP connection. Note that both the FIS and FRS message types are "protection-related" additions to the MPLS signaling framework (CR-LDP, RSVP). Owing to the generic nature of this specification, the PML and PSL nodes need not be the "end-point" source and destination nodes, respectively, and hence technically speaking, judicious placement thereof allows this framework to incorporate path, sub-path, and hop protection schemes. Although this overall framework seems most applicable to 1:1 or 1:N protection schemes (downstream nodes signal fault switchover requests to upstream nodes), a 1+1 protection type is also mentioned in [OWENS3]. Finally, proposals for sharing protection resources between multiple protection paths (and lower-priority traffic) are also beginning to emerge [BHANDARI],[GHANI1], [KINI2].

****

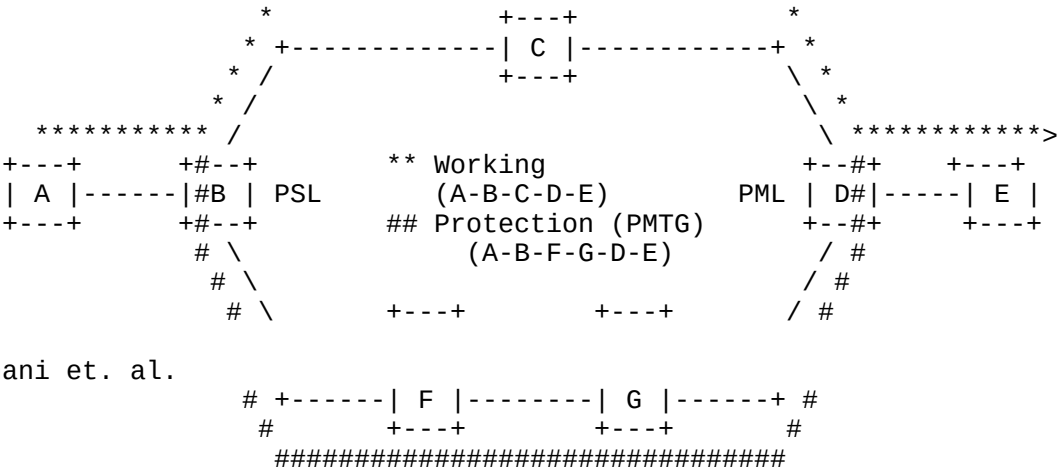

Figure 10: MPLS LSP protection concept (PSL/PML LSR nodes)

A paramount concern for network operators is fast recovery times. The MPLS LSP protection proposals are increasingly aware of this need, especially in comparison with the relatively longer timescales of IP re-routing schemes. Along these lines, the MPLS LSP protection framework includes the concept of a reverse notification tree (RNT) [OWENS1-2] entity that traverses from the PSL to the PML node. This provides an "express" signaling path for protection and recovery messages, significantly more efficient that "flooding-type" recovery schemes. The RNT is basically an "inverse" label-lookup (cross-connect) table that is constructed at the time of working/protection LSP setup and allows for resolving the incoming links on which to forward the backwards-propagating FIS message. As such, this construct implements "near-side" protection signaling. By using the RNT, hop-by-hop routing of FIS messages can be avoided, helping to expedite switchover times. In the latest specification, hop-by-hop routing (layer 3), packet LSP (MPLS), or SONET K1/K2 bytes (layer 2) mechanisms can be used to implement the RNT (see [OWENS1]). Also note that the RNT concept extends to multicast LSP's and is implemented for both working and protected paths. The latter allows it to be used to indicate failures on the protection path (requiring subsequent manual operator intervention, however). The actual setup of protection segment is implemented via extensions to the ER field of CR-LDP and RSVP-TE setup messages, e.g., LABEL_REQ, PATH [HUANG],[OWENS3]. Specifically, attributes are added for identifying the PSL/PML pair, protection type (1:1, 1+1) RNT implementation, timer values, etc. Note that there have also been related proposals for augmenting IGP protocols to support LSP protection (e.g., delineate active/back bandwidths), see [KINI1]. These can be extended to the optical case to specify active/ backup wavelength sets, etc.

Now consider the application of the above MPLS "packet LSP" protection framework within the context of protection switching in optical (ring) networks. For the case of channel (OCh) protection, the optical (O-ADM, OXC) LSR devices can now serve as PML and PSL nodes and "disjoint" protection lightpaths (or hops) can be specified between the two nodes, as per [GHANI2]. The PMTG entity at this level is the lightpath channel. For example, for edge-to-edge channel protection (e.g., O-BPSR/2, O-BPSR/4), the PSL/PML nodes can be the (sub)connection end-points themselves. Alternatively, for intermediate near-side channel protection (O-BLSR/4 case only), the PSL/PML pair can be the appropriate intermediate ring nodes. Given the appropriate information (via requirements specified in Section 3.1.2), RWA algorithms can appropriately setup the ring working/protection routes and switching points by using the ER signaling function. Note, further, that the PMTG concept can be used to group "lambdas" and define appropriate class of service (CoS) for the optical domain (i.e., from the service protection perspective).

The case of line protection, as proposed in O-BLSR/2/4 schemes, is somewhat different, since spans are more static (physical) entities and not dynamically created ones, as are lightpaths. However, using the protection group concept, all wavelengths on a given fiber span can be grouped into a common "span" PMTG and the diverse PMTG "span" route established. This route can either be a single span (O-BLSR/4) or a series of spans (O-BLSR/2 with loop-back), with the two adjacent nodes serving as the PML/PSL pair. Note that depending upon the span-switching implementation, wavelength switching may not be required. More clearly, O-BLSR/4 schemes using simple 2x1 switches for fiber protection do not permit wavelength re-use on protection fibers. In this case, a FIS message (pertaining to a fiber cut) will simply trigger a 2x1 span switch. However, simpler O-BLSR/2 schemes and more elaborate O-BLSR/4 schemes (e.g., without 2x1 span switches) can carry lower-priority traffic on protection wavelengths. In these cases, all individual channels of a PMTG have to be switched.

Although the above high-level interworking seems amenable, there are some concerns regarding recovery timing, particularly with regards to RNT setups and fault signaling. Consider the RNT issue first. During MPLS LSP setup, LSR nodes must keep track of the upstream node, incoming link and interface, and list of LSP(s) (unicast case) in order to construct the RNT. The procedure assumes bi-directional links between intermediate LSR nodes, since FIS messages are subsequently transmitted on the "reverse-table" incoming link interface. This implies an "inband" signaling setup. However, in optical rings (even meshes), especially transparent rings (meshes), there is likely a much higher degree of orthogonality between control and data flows. For example, if control signaling is done on out-band OSC channels and not "embedded" in data wavelengths, even though RNT setups can extract the above-detailed state at channel setup time, the actual FIS (and FRS) messages are not sent on the "reverse-lookup" incoming interface links. Additionally, the current MPLS RNT setup performs near-side protection signaling, since fault messaging traverses the same set of nodes but in the opposite direction. For long-side protection signaling (as required per some O-BLSR/O-BPSR designs, Section 2.2), however, protection signaling is required on the RNT of the protection path. This is slightly different from the existing possibility of MPLS protection-path RNT signaling [OWENS1], since it implies failure of the working and not protection side. All of these intricacies will require further setup signaling considerations.

Now consider the MPLS fault signaling message types, namely FIS and FRS and their usage for optical channel protection. Initially, the various fault detection (isolation) schemes, Section 2.4, are expected to trigger FIS message transmissions within a few milliseconds of an occurring fault (note that associated FIS hold-off timers must set appropriately). Once the FIS messages are generated, the remaining recovery latency is largely controlled by MPLS-layer signaling protocols and ensuing optical switchover times. The latter issue depends upon the actual switching technology used in the ring node’s protection stage, Figure 1, and realistically, millisecond timeframes can be expected via solutions such as MEMS or (O-E based) EXC designs. Meanwhile, this stresses the need for expedient FIS processing in order to match stringent benchmarks set by SONET/SDH APS. Here, the RNT architecture is of particular importance (as detailed above). It is

expected that high-priority MPLS packet LSP’s (routed on the OSC) will be required to expedite fault message transmissions along the reverse path. Specifically, improvements can be achieved using a variety of solutions. One is to use priority queuing for (reverse) FIS messages, and dedicate a fixed minimum amount of bandwidth via some

scheduler mechanisms. A further extension would be to perform FIS message processing (e.g., RNT label lookups and fast switchover) via dedicated hardware, such as FPGA devices. Clearly, both of these schemes entail added system complexity, and demonstrable evidence is required to determine if SONET/SDH recovery times can be effectively matched. Otherwise, the advantage of fast protection switching yielded by ring networks cannot be realized.

Another important issue arises with regards to "operational modes." Specifically, the emerging MPLS protection signaling framework still lacks some of the vital, "externally-initiated" [GR1230] features which SONET operators are well-accustomed to. Namely, the SONET K1/K2 byte protocol enables multiple operating "modes" via a well-defined message priority structure. For example, messages are defined (in decreasing order of priority) for lockout, forced switching, fault events (signal fail, signal degrade), and manual switching, see [GR1230]. Such procedures are vital to operations-related tasks and are used during various phases (i.e., maintenance, diagnostics, and upgrades). Controlling the "operating mode" is instrumental in avoiding excessive service disruptions to live customer traffic. Undoubtedly, similar functions must eventually be provided by MPL(ambda)S/GMPLS-based optical signaling protocols, in both ring and mesh networks, if optical channel services are to be deployed in carrier-class networks. This area has not received much attention to date and significant further work will be required. Provisioning such operating modes will require additional message types to be added to RSVP-TE and CR-LDP messaging, e.g., a forced or manual switching message type, etc. In summary, there are clearly a number of issues that need to be resolved before MPLS LSP protection schemes can be confidently applied to optical ring networks.

3.2.2 O-APS Protocol

As an alternative to generalizing MPLS LSP protection capabilities, a specialized, fast optical APS (O-APS) protocol is possible for optical rings. This entity can be considered as an orthogonal addition to the MPL(ambda)S/GMPLS protocols suite to achieve fast protection signaling, see Figure 9. For some, there are various compelling reasons to develop such an alternative. First of all, given the relatively stringent recovery requirements, many may argue that modifying or specializing MPLS signaling protocols (e.g., added failure-recovery messages, prioritized processing/implementations) may become too complicated and lengthy a process. Instead, a lightweight O-APS protocol can be designed, and this would be functionally equivalent to an "optical" version of the ubiquitous SONET K1/K2 byte protocol. Nevertheless, unlike the SONET K1/K2 byte APS protocol [GR1230] the O-APS protocol should be defined as "fast" packet-based protocol, in order to keep it in-line with the packet-oriented control philosophy of MPLS networks.

How the O-APS protocol’s packet messages are actually transported on control channels, however, can be left open to vendor implementation, but likely, bandwidth guarantees will be necessary in order to meet recovery timing requirements (similar to discussions for FIS message transport, Section 3.2.1). For example, some vendors may choose to explicitly map the message bytes into appropriate inband overhead bytes (e.g., SONET/SDH or digital wrappers bytes). Note that this case is somewhat different from that of standardizing explicit signaling bytes, i.e., "non-packetized" O-APS protocol. However, since protection timing is such a critical issue, some guidelines will likely be required to ensure satisfactory performance across larger networks (consisting of multi-vendor equipment). Such guidelines are for further study and can include guard-band times for message processing, etc.

Some of the key components of an O-APS protocol are briefly highlighted here, although a more detailed specification is clearly beyond the

scope of this discussion and intended for further study. Among other things, message fields must identify the switching nodes, lightpath channels/spans, fault type (channel, span, node), and requested protection actions (channel or span switching, near-side, far-side), etc. Additional parameters must also be specified for alarm messaging, such as durations, spacings, even priorities (e.g., span, channel). A complete state machine definition and related rules are also required, and examples include triggering recovery actions, starting/stopping alarm messaging, alarm squelching for multiple types of alarms (e.g., channel versus span, etc). Another issue is inter-node keepalive messaging. Such "hello" message formats are common in IGP protocols and are directly embedded into the SONET APS protocol, i.e., non-alarm K1/K2 byte fields serve as constant "hello" updates. O-APS peer nodes must also have this capability, and one alternative is to add explicit hello messaging for non-failure time periods. Note that the LMP protocol also has some provisions for "liveness" message updates, but this protocol is currently more geared towards mesh network support, i.e., OXC-to-OXC or router-to-OXC connectivity maintenance with likely longer inter-message periods, see [LANG],[FREDETTE]. (Nevertheless, new WDM-related provisions are being considered for LMP, and their applicability within the O-APS context is discussed later). Hence a fast, dedicated liveness/hello mechanism (and fast detection mechanism) is desirable for optical rings. Finally, since the O-APS protocol will be "new" protocol, it presents a good opportunity to properly define crucial "operator-initiated" functionalities, Section 3.2.1. For example, explicit message types (or fields, as appropriate) and appropriate priorities can be assigned for features such as resource lockout, forced and/or manual protection switching, etc. In fact, this option is one clear advantage of defining an altogether new protection O-APS switching protocol. However, significant further work is required to specify a truly generalized O-APS framework to implement the previously-defined transparent optical ring architectures, Section 2. Overall, an O-APS function will be an orthogonal, complimentary addition to the MPL(ambda)S/GMPLS suite.

Note that from a broader perspective, a dedicated O-APS protocol can also be deployed in a "standalone" manner, an added benefit. This is

important for many vendors who need to provide optical ring solutions, but at the same time, want to gradually transition into a full-blown "MPLS-based" control frameworks. In such cases, the orthogonal nature of the O-APS protocol will allow vendors to either couple its protection switching features with their own (proprietary) NMS-based provisioning solutions, or with their MPL(ambda)S/GMPLS-based control framework. In the former case, an NMS controller(s) will explicitly setup and takedown ring channel lightpaths and "fill in" the required information for the O-APS protocols to operate from, e.g., such as ring maps, etc. However, future migrations towards truly open, distributed provisioning paradigms (i.e., in lieu of proprietary NMS-based provisioning setups) will clearly necessitate added interworkings between the O-APS protocol and the other (orthogonal) MPL(ambda)S/GMPLS components. In particular, proper interfaces have to be identified (and developed) to enable any information exchange. Although the details of such interworkings are for further study, some preliminary possibilities can be highlighted here.

At channel setup time, the O-APS protocol may require various pieces of information from the related setup signaling entities (CR-LDP or RSVP-TE, Section 3.1) in order to perform its functions, i.e., since the O-APS protocol itself does not implement any channel provisioning functionalities. As a particular example, "connection ring map" information must be supplied after the appropriate signaling procedures have setup the associated lightpath channels, identifying the source and destination endpoints of the lighpath connection. Additional information will likely be required from the lighpath routing engine which computes

details of the working/protection routes, e.g., protection types (e.g., channel, span), switching endpoints (source/destination or intermediate node pairs), etc. For example, the source/destination ring nodes are the switching end-points for edge-to-edge long/near-side channel protection (as per O-UPSR and O-BPSR designs), whereas the selected intermediate nodes are the end-points for near-side intermediate channel switching (as per some O-BPSR/4 designs). Alternatively, for span protection, the end-points are the two nodes adjacent to the failure. Another requirement for information exchange (with the O-APS protocol) can also arise during fault event occurrences. Specifically, it was stated earlier that optical rings provide the added benefit of decoupling fault detection mechanisms from the subsequent recovery procedures, Section 2.4. Now in order to develop a more structured, formal mapping between the actual fault detection, notification, and recovery mechanisms, interworking with the emerging LMP protocol [LANG] can be considered. Specifically, LMP provides generic fault correlation/ notification functionalities which are independent of the actual fault detection schemes, a very germane feature. Moreover, recent proposals for new WDM-transport related considerations within the LMP framework [FREDETTE] will undoubtedly help improve its scalability and fault notification timings in optical (ring) networks. As this work matures, mapping LMP notifications to O-APS recovery mechanisms (e.g., via defining switching triggers) can improve overall architectural modularities/orthogonalities and this requires further investigation.

3.2.3 Multi-Layer Escalation Strategies

Assuming that fast optical (ring) lightpath protection schemes will emerge, inter-layer protection "collisions" will be of concern. Since multiple protocols can provide recovery mechanisms operating across multiple domains, the simultaneous interference of such functionalities (e.g., optical lightpath protection, SONET/SDH APS, MPLS LSP protection switching, IP flow re-routing) can lead to serious shortcomings, such as reduced resource utilization and data routing instabilities [DEMEESTER], [MANCHESTER2]. For example, optical lightpath recovery times can overlap with (client) SONET/SDH circuit or MPLS LSP protection timescales. Clearly, a mechanism is required to coordinate recovery actions between the various layers (packet, circuit, wavelength, fiber). This issue is commonly termed as escalation strategy design and has been treated in the broader research literature [GHANI1],[DEMEESTER],[MANCHESTER2]. Specifically, two types of escalation strategies have been proposed, namely bottom-up and top-down approaches, see [DEMEESTER] for full details. The former scheme assumes that "lower-level" recovery schemes (e.g., optical ring protection) are more efficient and expedient, and therefore inhibits higher-layer protection switching (such as IP re-routing, MPLS/ATM LSP protection switching, or SONET/SDH APS). Alternatively, the top-down approach attempts service recovery at the higher layers first before invoking "lower layer" (e.g., optical) recovery. The reasoning here is that higher-layer protection can be more service selective, and therefore efficient. Clearly, these are both advanced mechanisms and require complex signaling and hold-off timer mechanisms [GHANI2] to coordinate the different layer recovery procedures. Overall, the SLA’s between the network operators and their clients will determine the necessary timescales for protection recovery (e.g., 50 ms, 200 ms, 5 minutes, etc) and will also impact escalation strategy design. Note that a broader delineation of escalation strategies is also presented in [MANCHESTER2], i.e., serial and parallel approaches.

As far as the proposed optical (ring) protection framework is concerned, escalation strategies can be implemented using either MPLS/GMPLS or non-MPLS (non-GMPLS) type control-planes. Carefully note that this pertains to how protection capabilities are initiated and not the subsequent switching signaling actions. Consider the former case, in which the "higher layers" (e.g., packet LSP) are also controlled

(provisioned and protected) by the MPLS (GMPLS) framework. Assuming a generalized MPLS LSP restoration framework [XU] at all layers, escalation strategy timing is facilitated by this common control framework. The appropriate LSP protection timer mechanisms can specify hold-off times, alarm message (FIS) spacings, and alarm message durations. Clearly, judicious choices of these parameters at different LSP levels (packet, circuit, wavelength lightpath, fiberpath) can be used to design advanced "inter-layer" escalation strategies. For example, at the wavelength LSP level, small hold-off times and FIS spacings can be used to enact fast (sub-50 ms) recovery. Additionally, the duration of lightpath-level FIS messaging can be restricted to a timescale window, beyond which lightpath FIS notification is terminated. This duration (plus an acceptable guard-time) can be the hold-off time for "higher-layer" packet LSP FIS message generation. Note that this example details a "bottom-up" recovery case, and a complimentary "top-down" case can also be detailed. Specifically, lightpath recovery hold-off times can be set larger than

packet LSP notification durations, thereby permitting more selective "per-CoS" or "per-LSP/G-LSP" re-routing (based upon priorities, policies, etc). However, given the increased complexity and signaling requirements of top-down approaches, many operators may not find them very attractive in practical network settings.

Meanwhile, for the non-MPLS (non-GMPLS) control specific case, escalation strategy design can be more complicated since generalized timing control and signaling mechanisms may not exist at all protocol layers. In particular, this situation will arise if MPLS (GMPLS) protection is not used at the various networking "levels", e.g., O-APS at optical lightpath level, SONET/SDH APS at the circuit tributary level, MPLS LSP protection at flow level, etc. In general, this makes it more difficult to control inter-layer protection recovery timings, since inter-layer synchronization needs to be addressed/defined. For example, optical (WDM)-SONET/SDH protection interworkings may possibly need SONET/SDH hold-off timers, requiring changes to existing standards and deployed equipment [MANCHESTER2], raising even further challenges and complexities. In such cases, simpler "all-or-nothing" interworkings may be more feasible. For example, for the case of traditional "SONET-over-WDM", either optical ring or SONET APS recovery can be disabled. Nevertheless, for higher-layer IP packet traffic, "bottom-up" escalation strategies can usually be implemented safely by simply ensuring small enough FIS message windows, i.e., versus IGP re-routing timescales. In general, escalation strategy design is a complex issue and needs significant investigation.

3.3 Additional Considerations

Albeit detailed, the above discussions have only focused on basic optical ring definitions and provisioning issues. Clearly, many more advanced concerns relating to optical rings can be tabled, but their detailed treatment is beyond the scope of this document. Nevertheless, a brief synopsis is presented in order to stimulate further work.

3.3.1 Multi-Ring Provisioning

In most current SONET networks, multi-ring architectures are very common. Specifically, smaller rings are used to aggregate traffic from local domains onto larger rings spanning increased distances (metro, regional), and standards exist for so-called dual ring interworking (DRI) interconnection between multiple SONET/SDH rings [GR1230],[G.842]. Likewise, as optical rings emerge (most likely re-using much of the existing SONET/SDH ring fiber infrastructures), there will be a strong requirement for similar optical ring interworkings, namely to route lightpaths in between multiple rings. In addition to applying the conventional DRI concepts, inter-connection can be achieved by simpler, static "back-to-back" O-ADM co-location or via more advanced, dynamic OXC switching devices [ARIJIS]. Now conceivably, different optical

ring types (e.g., O-BLSR, O-BLSR, O-UPSR) can be used for an end-to-end circuit connection, along with their respective "localized" protection mechanisms (i.e., protection zones). Clearly, this may also permit greater latitudes in user SLA definitions.

>From a routing/provisioning point of view, there are various ways to handle such "multi-ring" architectures. In the longer run (and for

larger rings) it may be advantageous to move to a hierarchical routing setup. Specifically, individual rings would be grouped into separate domains, e.g., autonomous systems (AS), and multi-ring provisioning would be performed under the broader context of inter-domain provisioning [GHANI1],[GUO2],[PAPADIMITRIOU2],[RAJAGOPALAN]. Here, added enhancements to emerging (optical) network interface definitions (O-NNI) [PAPADIMITRIOU2] may be required, e.g., for setup signaling, protection switching between multiple rings, etc (further study needed). Alternatively, a more immediate alternative is that of intra-domain provisioning between rings, especially where multiple smaller rings constitute a domain, i.e., single ring represents an area instead. Here, since opaque LSA's only have area scope, further work is required in order to define summary LSA's to provide enough information for inter-ring (i.e., intra-domain) provisioning, yet without flooding the network with message updates.

3.3.2 Hybrid Mesh-Ring Interworking

Similar to the case of multi-ring provisioning (above), the broader evolution towards mesh/mesh-ring network topologies is also an important concern. From an operator migration point of view, both ring and mesh topologies have their respective advantages/disadvantages. Ring topologies allow for fast, well-defined protection switching concepts but have reduced connectivity (degree two). Meanwhile, mesh networks offer better connectivity and improved resource efficiencies but lack well-defined protection switching features. Hence, a likely, cost-effective migration path will be for operators to first migrate their existing SONET/SDH (TDM-based) rings to counterpart optical rings and then move towards mesh or "hybrid" mesh-ring topologies, inline with growth in traffic demand and operational experience. These evolutions can be done either via phased expansions to existing ring topologies (i.e., adding fibers between non-adjacent ring nodes to "break" the ring) or altogether new (i.e., "greenfield-type") deployments. Such cases present two foreseeable interworking requirements, namely for ring emulation and ring-mesh interconnection purposes (and others may also emerge) [GUO2]. Either way, it is clear that provisioning features for hybrid topologies will be a crucial requirement for operators as they move to deploy or expand their optical network offerings.

On a high level, ring emulation basically entails provisioning/operating "virtual" rings on top of mesh (network) topologies, e.g., via the concept of ring covers [PAPADIMITRIOU3]. For operators accustomed to operating ring networks, this capability will still allow them to expand to mesh topologies, and is particularly germane to the case of phased-in ring-to-mesh expansions. For example, different ring types (O-UPSR, O-BLSR, O-BPSR) can be deployed in selective parts of a mesh network topology, thereby exploiting the advantages of fast ring-based protection switching. Additionally, for richly-connected mesh networks, operators can offer virtual private optical ring (VPOR) services to large clients, an attractive proposition. Note that ring emulation will require that specific network nodes (i.e., those sitting on multiple rings) have more advanced spatial switching characteristics, as yielded by OXC/PXC designs and not basic O-ADM designs. Moreover, these nodes must be "ring-enabled", and most notably, be capable of meeting fast protection switching requirements.

Meanwhile, the issue of ring-mesh interconnection arises when provisioning lighpath channels across multiple, separate ring and mesh topologies (e.g., metro-regional rings to long-haul meshes). Here, the mesh network segments themselves may or may not perform ring emulation, and therefore this becomes a generalization of the multi-ring interconnection case (Section 3.3.1). However, many of the concepts discussed for the multi-ring case, such as intra- and inter-domain partitioning, are also applicable here. Of particular concern will be achieving commensurate protection switching timescales in the mesh-network segments. However, here it is expected that as standards evolve, many of the optical ring protection switching concepts/protocols will likely also be leveraged for mesh architectures. For example, mesh span protection or mesh end-to-end channel protection (via diverse routing) can re-use or extend the working/protection channel setup and protection signaling mechanisms developed for optical rings.

Overall, hybrid topology provisioning, for both ring emulation and ring-mesh interconnection, will require additional specifications. Considering the more likely case of intra-domain provisioning, topology/ resource discovery methods will need to perform summarization/aggregation of ring information. Some early work along these lines is emerging, proposing "ring ID" and "ring type" sub-TLV definitions for opaque LSA's, see [GUO2],[PAPADIMITRIOU3]. Provisioning lightpaths across ring-mesh topologies or "virtual ring" mesh networks will require new resource/ constrained-routing algorithms also. A related, key issue here will be implementing protection switching signaling (as per the user SLA requirements), e.g., protecting a fiber cut between multiple "virtual" rings or an optical channel failure across multiple network ring/mesh network segments. Cleary, there is a strong need for more detailed work in the area of hybrid topology provisioning.

3.3.3 Resilient Packet Ring (RPR) Synergies

So far, only "full-granularity" lightpaths have been considered, i.e., ring channels utilizing full wavelength capacity, such as OC-48/STM-16 and OC-192/STM-64. However, new generations of advanced integrated edge devices (IED's) are beginning to appear, integrating (sub-rate) packet, circuit, and wavelength/fiber switching capabilities onto a common platform. In a timely manner, the broader GMPLS framework is also being extended to provision related sub-rate tributary channels [ASHWOOD1], [XU] (both packet and circuit). Therefore, for improved generality, GMPLS ring provisioning mechanisms/concepts can also be considered for various "sub-rate ring" architectures. Sub-rate rings can improve wavelength utilization and provide finer-granularity connections between smaller users. A very good example of a sub-rate (packet) ring technology is the emerging resilient packet ring (RPR) architecture (IEEE 802.17). Recently there has been significant interest and development in this space, as noted by work on IP over RPR (IPoRPR) architectures, see [HERRERA].

Since packet rings operate on an "electronic" level and require visibility into the packet stream, from a technology point of view, they are quite different from the optical rings proposed herein. Nevertheless, there are many generic architectural aspects of optical

rings which can also apply to packet rings, and these can be further investigated. Examples include setup signaling procedures, protection signaling, constrained routing, etc. Furthermore, RPR's can be mapped onto optical ring wavelength channels, thereby permitting potential traffic/resource engineering synergies, i.e., advanced "multi-ring" architectures with "embedded" sub-rate packet ring channels being carried in larger granularity optical ring channels. Additionally, given the fast fault detection/recovery mechanisms being proposed for RPR designs (in comparison with mesh IP packet re-routing), the likelihood of protection "collisions" between optical ring and RPR

recovery mechanisms increases. Hence protection escalation strategies (Section 3.2.3) are of importance in such interworkings and should be designed to properly arbitrate between the respective protection protocol(s) or different levels thereof (packet, circuit). All of these topics need further, more detailed investigation.

4. APPENDIX A: Review of SONET Ring Architectures

SONET (SDH) ring architectures have emerged to dominate the transport landscape. Termed also as self-healing rings (SHR), perhaps their defining characteristic is stringent recovery timescales. Namely, SONET (and SDH) standards stipulate a service recovery time of 50 ms after the fault condition (i.e., including detection, guard time, switching time, ring propagation delays, and re-synchronization). These values are derived from frame synchronization at the lowest frame speed, namely DS1 (1.5 Mb/s). A very brief outline of the SONET/SDH ring framework is given here. However, this summary is only intended to serve as a background reference, and interested readers are referred to the specifications for complete details [ANSI],[ITU], [GR1230]. As will be seen, these existing architectures will form the basis for much of the counterpart optical ring frameworks.

4.1 Uni-directional Path-Switched Ring (UPSR)

The UPSR concept is designed for channel level protection in two-fiber rings. Although two-fiber BLSR architectures also exist, termed BLSR/2, the UPSR architecture is significantly less complex. UPSR rings dedicate one fiber for working TDM channels (timeslots) and the other for corresponding protection channels (counter-propagating directions). Traffic is permanently bridged at the head-end and sent along both fibers, namely 1+1 protection. UPSR working traffic travels in the clockwise direction and protection traffic travels in the counter-clockwise direction. This implies that bi-directional connections will consume resources on all working and protection fibers, restricting ring throughput to that of a single fiber. Clearly, UPSR rings represent simpler designs and do not require any notification or switchover signaling mechanisms between ring nodes, i.e., receiver nodes perform channel switchovers. As such, they are resource inefficient since they do not re-use fiber capacity (both spatially and between working/protection paths). Moreover, span (i.e., fiber) protection is undefined for UPSR rings, and such rings are typically most efficient in access rings where traffic patterns are concentrated around collector hubs.

4.2 Bi-Directional Line Switched Ring (BLSR)

BLSR rings are designed to protect at the line (i.e., fiber) level, and there are two possible variants, namely two-fiber (BLSR/2) and four-fiber (BLSR/4) rings. The BLSR/2 concept is designed to overcome the spatial reuse limitations associated with two-fiber UPSR rings and only provides path (i.e., line) protection. Specifically, the BLSR/2 scheme divides the capacity timeslots within each fiber evenly between same-direction working and protection channels (with working channels on a given fiber being protected by protection channels on the other fiber). Therefore bi-directional connections between nodes will now traverse the same intermediate nodes but on differing fibers. This allows for sharing loads away from saturated spans and increases the level of spatial re-use (sharing), a major advantage over two-fiber UPSR rings. Protection slots for working channels are pre-assigned based upon a fixed odd/even numbering scheme, and in case of a fiber cut, all affected timeslots are looped back in the opposite direction of the ring. This is commonly termed "loop-back" line/span protection and avoids any per-channel processing. However, loop-back protection increases the distance and transmission delay of the restored channels (nearly doubling path lengths in the worst case). More importantly,

since BLSR rings perform line switching at the switching nodes (i.e., adjacent to the fault), more complex active signaling functionality is required. Further bandwidth utilization improvements can also be made here by allowing lower-priority traffic to traverse on idle protection spans.

Four-fiber BLSR rings extend upon the BLSR/2 concepts by providing added span switching capabilities. In BLSR/4 rings, two fibers are used for working traffic and two for protection traffic (counter-propagating pairs, one in each direction). Again, working traffic can be carried in both directions (clockwise, counter-clockwise), and this minimizes spatial resource utilization for bi-directional connection setups. Line protection is used when both working and protection fibers are cut, looping traffic around the long-side path. If, however, only the working fiber is cut, less disruptive switching can be performed at the fiber level. Here, all failed channels are switched to the corresponding protection fiber going in the same direction (and lower-priority channels pre-empted). Overall, the BLSR/4 ring capacity is twice that of the BLSR/2 ring, and the four-fiber variant can handle more failures. Also, it should be noted that both two- and four-fiber rings provide node failure recovery for pass-through traffic. Essentially, all channels on all fibers traversing the failed node are line switched away from the node.

As mentioned above, BLSR rings (unlike UPSR rings) require a protection signaling mechanism. Since protection channels can be shared, each node must have global state, and this requires state signaling over both spans (directions) of the ring. This is achieved via an automatic protection switching (APS) protocol running on the "embedded" K1/K2 bytes in the SONET/SDH frame overhead, also commonly termed as the SONET APS or BLSR K1/K2 byte protocol [GR1230], [T1.105.01],[G.841]. This protocol uses a 4-bit node identifier (in the K1 byte) and hence only allows up to 16 nodes per ring. Additional

bits are designated to identify the type of function requested (e.g., bi-directional or unidirectional switching) and the fault condition (i.e., channel state). Control nodes performing the switchover functions utilize frame-persistency checks to avoid premature actions and discard any invalid message codes. Further details can be found in the associated references.

5. Security Considerations

Security considerations are for future study, in particular with regards to signaling extensions and a possibly new O-APS protocol. The overall optical ring provisioning framework, however, poses the same security requirements as those present in existing MPL(ambda)S or GMPLS provisioning architectures.

6. References

[ABOULMAGD] O. Aboul-Magd, et al, "Signaling Requirements of the Optical UNI," Internet Draft, draft-bala-mpls-optical-uni-signaling-01.txt, November 2000.

[ANSI] "Synchronous Optical Network (SONET): Basic Description Including Multiplex Structure, Rates, and Formats," ANSI T1.105-1995, 1995.

[ARIJS] P. Arijs, et al, "Design of Ring and Mesh Based WDM Transport Networks," Optical Networks, July 2000.

[ARVIND] K. Arvind, et al, "Optical Domain Service Interconnect (ODSI) Signaling Control Specification," Version 1.4.5, ODSI Coalition, November 2000.

[ASHWOOD1] P. Ashwood-Smith, et al, "Generalized MPLS-Signaling Functional Description," Internet Draft, draft-ietf-mpls-generalized-signaling-01.txt, November 2000.

[ASHWOOD2] P. Ashwood-Smith, et al, "Generalized MPLS Signaling-RSVP-TE Extensions," Internet Draft, draft-ietf-mpls-generalized-rsvp-te-00.txt, December 2000.

[ASHWOOD3] P. Ashwood-Smith, et al, "Generalized MPLS Signaling-CR-LDP Extensions," Internet Draft, draft-ietf-mpls-generalized-cr-ldp-00.txt, November 2000.

[AWDUCHE] D. Awduche, Y. Rekhter, J. Drake, R. Coltun, "Multi-Protocol Lambda Switching: Combining MPLS Traffic Engineering Control with Optical Crossconnects," Internet Draft, draft-awduche-mpls-te-optical-00.trt, October 1999.

[BERNSTEIN] G. Bernstein, V. Sharma, "Some Comments on GMPLS and Optical Technologies," Internet Draft, draft-bernstein-gmpls-optical-00.txt, November 2000.

[BHANDARI] R. Bhandari, S. Sankaranarayanan, E. Varma, "High-Level Requirements for Optical Shared Mesh Restoration," Internet Draft, draft-bhandari-optical-restoration-00.txt, May 2001.

[CEUPPENS] L. Ceuppens, et al, "Performance Monitoring in Photonic Networks in Support of MPL(ambda)S," Internet Draft, draft-ceuppens-mpls-optical-00.txt, March 2000.

[CHEN] J. Chen, T. Shiragaki, "Routing of OCh Shared Protection Ring," T1X1 Forum, T1X1.5/99-256R1, October 1999.

[CHIU] A. Chiu, J. Strand, R. Tkach, "Unique Features and Requirements for the Optical Layer Control Plane," Internet Draft, draft-chiu-strand-unique-olcp-01.txt, December 2000.

[CVIJETIC1] M. Cvijetic, T. Shiragaki, "Standardization of OCh Shared Protection Ring and Its Open Issue List," T1X1 Forum, T1X1.5/99-255R1, October 1999.

[CVIJETIC2] M. Cvijetic, T. Shiragaki, A. Weissberger, "OCh Shared Protection Ring," T1X1 Forum, T1X1.5/99-178, July 1999.

[DEMMESTER] P. Demmester, et al, "Resiliency in Multilayer Networks," IEEE Communications Magazine, March 2000.

[DOSHI] B. Doshi, et al, "Optical Network Design and Restoration," Bell Labs Technical Journal, January-March 1999.

[DOVOLSKY] D. Dovolsky, I. Bryskin, "Calculation of Protection Paths and Proxy Interfaces in Optical Networks Using OSPF," Internet Draft, draft-dovolsky-bryskin-ospf-pathprotect-proxy-00.txt, June 2000.

[DUTTA] R. Dutta, G. Rouskas, "A Survey of Virtual Topology Design Algorithms for Wavelength Routed Optical Networks," Optical Networks, January 2000.

[FREDETTE] A. Fredette, et al, "Link Management Protocol (LMP) for WDM Transmission Systems," Internet Draft, draft-fredette-lmp-wdm-00.txt, December 2000.

[GHANI1] N. Ghani, "Lambda-Labeling: A Framework for IP over WDM Using MPLS," Optical Networks, April 2000.

[GHANI2] N. Ghani, "Survivability Provisioning in Optical MPLS Networks," 5th European Conference on Networks and Optical Communications, June 2000.

[GHANI3] N. Ghani, et al, "COPS Usage for ODSI," Version 2, ODSI Coalition, August 2000.

[G.709] D. Brungard, "TD Draft ITU-T G.709 for February 2001 SG15 Meeting," T1X1 Forum, T1X1.5/2001-043, March 2001.

[G.841] "Types and Characteristics of SDH Network Protection Architectures," ITU-T Recommendation G.841, 2000.

[GFP] E. Hernandez-Valencia, "GFP Draft Specification-Revision 2 (Synchronous Optical Network (SONET)-Generic Framing Protocol," January 2001.

[GR1230] GR-1230-CORE, SONET Bi-directional Line-Switched Ring Equipment Generic Criteria, Issue 4, December 1998.

[GR3009] GR-3009-CORE, Optical Cross-Connect Generic Requirements, Issue 1, January 1999.

[GUO1] D. Guo, et al, "Extensions to RSVP-TE for Bi-directional Optical Path Setup," Internet Draft, draft-sorrento-rsvp-bi-osp-00.txt, July 2000.

[GUO2] D. Guo, et al, "Hybrid Mesh-Ring Optical Networks and Their Routing Information Distribution Using Opaque LSA," Internet Draft, draft-guo-optical-mesh-ring-00.txt, December 2000.

[HERRERA] A. Herrera, et al, "A Framework for IP Over Resilient Packet Rings," work in progress, January 2001.

[HUANG] C. Huang, V. Sharma, S. Makam, K. Owens, "Extensions to RSVP-TE for MPLS Path Protection", Internet Draft, draft-chang-rsvpte-path-protection-ext-00.txt, June 2000.

[ITU] "Network Node Interface for the Synchronous Digital Hierarchy (SDH)," International Telecommunication Union, G.707, March 1996.

[KINI1] S. Kini, M. Kodialam, T. V. Lakshman, "Open Shortest Path First (OSPF) Protocol Extensions for Label Switched Path Restoration," Internet Draft, draft-kini-ospf-lsp-restoration-00.txt, October 2000.

[KINI2] S. Kini, M. Kodialam, T. V. Lakshman, "Shared Backup Label Switched Path Restoration," Internet Draft, draft-kini-restoration-shared-backup-00.txt, October 2000.

[KOMPELLA1] K. Kompella, et al, "OSPF Extensions in Support of GMPLS," Internet Draft, draft-kompella-ospf-gmpls-extensions-00.txt, extensions-00.txt, July 2000.

[KOMPELLA2] K. Kompella, et al, "IS-IS Extensions in Support of MPL(ambda)S," Internet Draft, draft-kompella-isis-ompls-extensions-00.txt, July 2000.

[KOMPELLA3] K. Kompella, et al, "Extensions to IS-IS/OSPF and RSVP in Support of MPL(ambda)S," Internet Draft, draft-kompella-mpls-optical-00.txt, March 2000.

[LANG] J. Lang, et al, "Link Management Protocol (LMP)," Internet Draft, draft-lang-mpls-lmp-01.txt, July 2000.

[MARCENAC] D. Marcenac, "Benefits of Wavelength Conversion in Optical

Ring-Based Networks," Optical Networks, April 2000.

[MANCHESTER1] J. Manchester, P. Bonenfant, C. Newton, "The Evolution of Transport Network Survivability," IEEE Communications, August 1999.

[MANCHESTER2] J. Manchester, P. Bonenfant, "Fiber Optic Network Survivability: SONET/Optical Protection Layer Interworkign," Proceedings of the National Fiber Optics Engineering Conference 1996 (NFOEC'96), Denver, CO, 1996.

[MCADAMS] L. McAdams, J. Yates, K. Bala, "User to Network Interface (UNI) Service Definition and Lightpath Attributes," OIF Forum, OIF2000.061, September 2000.

[OWENS1] K. Owens, et al, "A Path Protection/Restoration Mechanism for MPLS Networks", Internet Draft, draft-chang-mpls-path-protection-02.txt, November 2000.

[OWENS2] K. Owens, V. Sharma, M. Oommen, "Network Survivability Considerations for Traffic Engineered IP Networks", Internet Draft, draft-owens-te-network-survivability-00.txt, March 2000.

[OWENS3] K. Owens, et al, "Extensions to CR-LDP for MPLS Path Protection," Internet Draft, draft-owens-crldp-path-protection-ext-00.txt, December 2001.

[PAPADIMITRIOU1] D. Papadimitriou, et al, "Inference of Shared Risk Link Groups," Optical Internetworking Forum, OIF2001.066, January 2001.

[PAPADIMITRIOU2] D. Papadimitriou, et al, "Optical Network-to-Network Interface Framework and Signaling Requirements," Internet Draft, draft-papadimitriou-onni-frame-01.txt, November 2000.

[PAPADIMITRIOU3] D. Papadimitriou, "Optical Rings and Hybrid Mesh-Ring Optical Networks," Internet Draft, draft-papadimitriou-optical-Rings-00.txt, February 2001.

[RAJAGOPALAN] B. Rajagopalan, et al, "IP Over Optical Networks-A Framework," Internet draft, draft-many-optical-framework-02.txt, November 2000.

[SOULLIERE] M. Soulliere, "Proposed ITU-T Contribution on Transparent OCh SPRings," T1X1 Forum, T1X1.5/2001-027, January 2001.

[SZERENYI] L. Szerenyi, "Approach to OSC Standardization," T1X1 Forum, T1X1.5/2001-002, January 2001.

[T1.105.01] S. Gorshe, "Synchronous Optical Network (SONET) Automatic Protection Switching, ANSI T105.01 (T1X1.5/99-065R3), November 1999.

[XU] Y. Xu, et al, "Generalized MPLS Control Plane Architecture for Automatic Switched Transport Network," Internet Draft, draft-xu-mpls-ipo-gmpls-arch-00.txt, November 2000.

[XUE] Y. Xue, et al, "Carrier Optical Services Framework and Associated UNI Requirements," Internet Draft, draft-many-carrier-framework-uni-00.txt, November 2000.

[YU] J. Yu, et al, "RSVP Extensions in Support for OIF Optical UNI Signaling," Internet Draft, draft-yu-mpls-rsvp-oif-uni-00.txt, July 2000.

[ZANG] H. Zang, J. Jue, B. Mukherjee, "A Review of Routing and

Wavelength Assignment Approaches for Wavelength-Routed Optical WDM Networks," Optical Networks, January 2000.

7. Authors Information

Nasir Ghani
Sorrento Networks Inc.
9990 Mesa Rim Blvd,
San Diego, CA 92121, USA
nghani@sorrentonet.com

James Fu
Sorrento Networks Inc.
9990 Mesa Rim Blvd,
San Diego, CA 92121, USA
jfu@sorrentonet.com

Dan Guo
Sorrento Networks Inc.
San Diego, CA 92121, USA
dguo@sorrentonet.com

Xinyi Liu
Sorrento Networks Inc.
San Diego, CA 92121, USA
xliu@sorrentonet.com

Zhensheng Zhang
Sorrento Networks Inc.
9990 Mesa Rim Blvd
San Diego, CA 92121, USA
zzhang@sorrentonet.com

Paul Bonenfant
Photuris
20 Corporate Place South
Piscataway, NJ 08854, USA
pbonenfant@photuris.com

Leah Zhang
Photuris Inc.
20 Corporate Place South
Piscataway, NJ 08854, USA
lzhang@photuris.com

Antonio Rodriguez Moral
Photuris Inc.
20 Corporate Place South
Piscataway, NJ 08854, USA
arodmor@photuris.com

Murali Krishnaswamy
Photuris Inc.
20 Corporate Place South
Piscataway, NJ 08854, USA
murali@photuris.com

Dimitri Papadimitriou
Alcatel IPO-NSG
B-2018 Antwerpen, Belgium
Phone: +32 3 240-8491
dimitri.papadimitriou@alcatel.be

Sudheer Dharanikota
Nayna Networks Inc.
157 Topaz Street
Milpitas, CA 95035, USA
sudheer@nayna.com

Raj Jain
Nayna Networks Inc.
157 Topaz Street
Milpitas, CA 95035, USA
raj@nayna.com

8. Full Copyright Notice

BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.