Automatic Protection Switching for *p*-Cycles in WDM Networks

Dominic A. Schupke

Abstract

p-Cycle recovery relies on a protection switching protocol. We detail several issues for such a protocol taking the evolution from ring networks to *p*-cycles into account. In particular, we propose and evaluate a protocol enhancement to provide means for node failure protection. For the evaluation, we describe an integer linear program, which is applied to network design case studies, and formulate availability models for *p*-cycles. The case studies show that the protocol enhancement improves availability at marginal additional design cost.

Index Terms

Protection, p-Cycles, Node Failures, APS, Network Design.

I. Introduction

The concept of preconfigured protection cycles [1], or "*p*-cycles" for short, has attracted much interest within the research on resilient networks [2], [3], [4], [5], [6], [7]. Besides the applicability of *p*-cycles in many networks, such as WDM, SONET/SDH, or IP/MPLS networks, this concept is attractive because of high capacity-efficiency and—by analogy to automatic protection switching of SONET/SDH rings—fast protection switching times.

Recently, several issues and concerns related to operation and deployment of *p*-cycles in WDM networks have come into discussion. After providing a tutorial-style overview of *p*-cycle operation, we aim to address these issues and concerns in this paper.

Specifically, while WDM *p*-cycles share many properties with SONET/SDH rings (thereby answering many questions), we still undergo a technology change (from SONET/SDH to WDM) and a structure change (from rings to *p*-cycles). In the former one we deal, e.g., with possible WDM sublayers at which *p*-cycles can or should operate. In the latter one we have to ask whether the main network objectives for rings can also be met by *p*-cycles. In this context we reveal in this paper that *p*-cycle node failure recovery cannot be as effective as for rings and we propose a novel concept which enables *p*-cycles to recover from node failures for selectable paths. We also evaluate the capacity and availability performance of networks designed using this concept.

The remainder of this section summarizes related work. In Section II we describe the protection principle of *p*-cycles to recover from link failures. Protocols for *p*-cycles are discussed in Section III. An enhancement of the *p*-cycle APS protocol to protect against node failures is suggested in Section IV. A corresponding design method and and evaluation is presented in Section V. In Section VI we analyse the concept from the availability point-of-view. In Section VII we conclude the paper.

The author is with Siemens, Corporate Technology, Information and Communications, Munich, Germany (dominic.schupke@siemens.com). This work is part of the work while the author was with the Institute of Communication Networks at Technische Universität München, Munich, Germany.

Preprint copy—Please cite as:
Schupke, D.A., "Automatic Protection Switching for p-Cycles in WDM Networks," Optical Switching and Networking (OSN), Elsevier, accepted April 2005.

A. Related Work on Node-Protecting p-Cycles

For node-protection, Stamatelakis and Grover [8], [2] propose node-encircling *p*-cycles. These *p*-cycles are routed through all adjacent neighbor nodes of a node to be protected, but exclude the protected node itself, thus protecting the traffic transiting the node. In certain networks, nodes can only be protected by non-simple cycles. Node-encircling *p*-cycles are most suitable for networks in which the *p*-cycles' protection capacity can be arbitrarily shared (e.g., an IP/MPLS network). If we apply in the same manner a reserved *p*-cycle (as in transport networks), its capacity must be able to carry the traffic transiting the node which can require too excessive capacity values per *p*-cycle. To make this plausible, consider a node with degree d, where the degree of a node is the number of its incident links. Then, the *p*-cycle may have to protect up to $\frac{d(d-1)}{2}$ neighbor traffic relationships through the node, e.g., for $d = 5$, the node-encircling *p*-cycle protects up to 10 traffic relationships; this is much more than for a link-protecting *p*-cycle which protects always only one. As in the best case the transiting traffic is bifurcated on the two segments of the cycle, half of transiting traffic is required as *p*-cycle capacity. Hence, if the transiting traffic is high, the *p*-cycle capacity is still high. Since this capacity would be reserved exclusively per node on all links of its *p*-cycle, we will not consider such node-encircling *p*-cycles.

Shen and Grover suggest the concept of path-segment-protecting *p*-cycles, or flow-*p*-cycles, in [9] and in the extended journal publication [10]. They define a flow as any single contiguous segment of a working end-to-end path. The entities, which the *p*-cycles protect, are then flows constituting the end-to-end working paths. As a link (or "span") is a special case of a flow, flow-*p*-cycles generalize the concept of link-protection *p*-cycles. To deploy flow-*p*-cycles, a notification scheme is proposed, which signals a failure on the path-segment to the switching nodes on the cycle. Another concept is that of providing "optical bypass" at the intermediate nodes of the path-segment. Flow *p*-cycles can be more efficient than link-protecting *p*-cycles and are able to protect against node failures. Path-segments, however, introduce another logical layer, since end-to-end working paths are mapped into path-segments which are in turn mapped into links. Except for the bypass concept, flow *p*-cycles do not have the simplicity of failure detection and protection switching by the nodes adjacent to the failure (as link-protection *p*-cycles), since a failure signaling mechanism between detecting nodes (which can be outside the *p*-cycle) and switching nodes is required.

II. Protection Principle

p-Cycles for link protection are particularly interesting for optical networks, since link failures are the most frequent failure events in these networks. Often it is also feasible to assume single link failures (i.e., only one failure at a time) in the network. While we assume single link failures in this section, the network may fail for node failures (Section III-A) and multiple failures [11], [12], [13].

In this paragraph, assume a network in which connections are switched individually (e.g., a WDM network with protection-switching on the OCh level). Figure 1 depicts the protection principle of *p*-cycles for link protection. The *p*-cycle in Figure 1(a) is preconfigured as a closed connection on the cycle B-C-D-F-E-B. Preconfiguration means that the configuration is done before a failure occurs. The *p*-cycle is able to protect working capacity on its own links, called *on-cycle* links, as shown in Figure 1(b). Upon failure of on-cycle link B-C, the *p*-cycle offers protection by the route on the remainder of the cycle (C-D-F-E-B). The protectable capacity on on-cycle links is thus one capacity unit. The protection of on-cycle links is logically equal to multiplex-section shared protection rings (MS-SPRings) in SDH and bidirectional line-switched rings (BLSRs) in SONET [14]. Unlike these rings, however, *p*-cycles also protect links outside the *p*-cycle path: Each link which has both its end points on the *p*-cycle can also be protected. Figure 1(c) shows the protection of such a link (E-D) which is called *straddling* link. We can provide two protection routes for straddling links, in the example, routes E-B-C-D and E-F-D. In effect, we can protect two working capacity units of straddling links.

We can extend the concept to packet-switched networks (e.g., IP/MPLS networks) or circuit-switched networks in which connections are protection-switched in groups (e.g., a WDM network with protection-switching on the OMS level). Then, the *p*-cycle is a closed path with a given (protection) capacity. Packet-switched networks with virtual connections can share this capacity (a bitrate) arbitrarily (it is only used

during active failure recovery). In circuit-switched networks, the protected working links of a *p*-cycle share this capacity (the number of connections in the *p*-cycle group), which is, however, reserved capacity; out of the scope of this paper is that circuit-switched *p*-cycles mutually share their protection capacity or offer it to pre-emptible traffic (see also Section III-A).

In summary, the *p*-cycle in Figure 1 can protect five on-cycle links (B-C, C-D, D-F, F-E, and E-B) and two straddling links (C-F and E-D). The protectable capacity of a straddling link is double the protectable capacity of an on-cycle link.

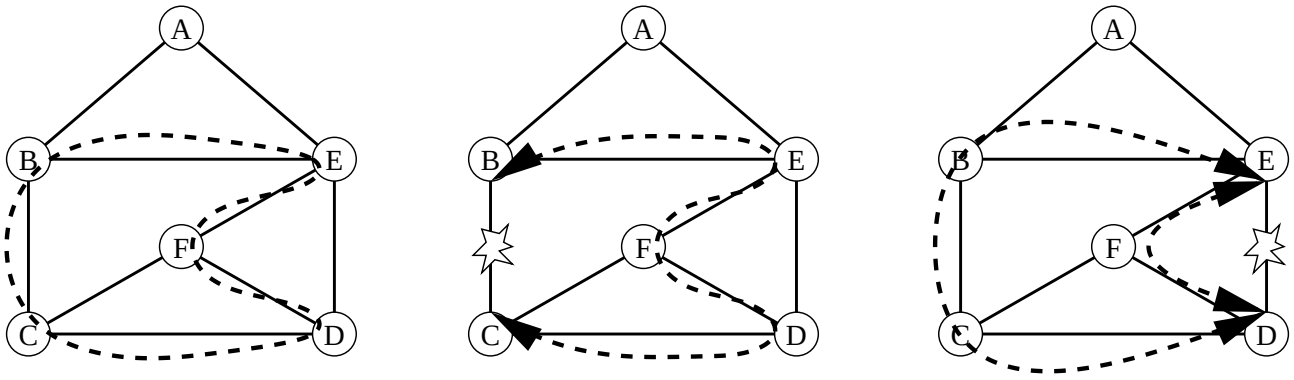

(a) A *p*-cycle (dashed line) in a network. (b) Protection of an on-cycle link. (c) Protection of a straddling link.

Fig. 1. Protection principle of *p*-cycles for link protection.

With Figure 1 we can realize pictorially that the protection routes for straddling links are on average half the *p*-cycle circumference whereas for on-cycle links they are approximately the whole circumference (as in BLSR rings). This is of virtue for cases where the length of a route is restricted (e.g., because of signal degradation).

III. Protection Switching Protocols

We describe how known ring protection switching protocols can be enhanced for *p*-cycles. Reusing functionality of existing protocols can significantly ease implementation, operational migration, and personnel training. Grover and Stamatelakis have already indicated [1] that *p*-cycles can also operate by employing the existing shared protection ring protocols with modifications. Naturally, *p*-cycle protocols designed from scratch are theoretically also possible. It is clear that "oblivious" protection switching is unlikely for application in networks. With this kind of switching, *p*-cycles are present as preconfigured closed connections; only the two nodes adjacent to a link failure are involved in the recovery process and switch to the protection paths obliviously, i.e., without failure signaling. This only works under the assumption of single link failures and it can lead to misconfiguration, if other failure scenarios occur (e.g., node failures). An interesting improvement with respect to multiple link failures is that a *p*-cycle advertises its current state to all its nodes [15]. Then, after protection switching upon a first link failure, a second switch action following a further link failure can be stalled to avoid misconfiguration. After node failures this concept can involve a short period of misconfiguration. Note, however, that it is suffient to assume switching is oblivious when doing research in the common context of all single failures only, and to be aware that the concept used in practice inevitably has to have a signaling protocol.

A. Ring Protocols

We resort to the established standard ITU-T Recommendation G.841 [16], which defines SDH shared protection rings. The key element is the automatic protection switching (APS) protocol. We expect protocol implementations for *p*-cycles in optical networks to be close to this ring APS protocol, since the only

major extension is the capability of straddling link protection. Although the shared protection rings in [16] perform protection switching of a group of channels (multiplex section), shared protection rings working on individual channels have the same functional behavior [14].

Let us summarize the ring APS protocol functionality. Each ring node has an address, a ring map (i.e., the sequence of the nodes on the ring), knowledge about traffic passed through that node, and information about the source and destination addresses of traffic on the ring. APS uses special signaling channels of SDH for request messages, which are coded in the "K1,K2 bytes" and are transmitted by a node in both directions on the ring. Our discussion puts an emphasis on the actions for protection against single physical link and node failures.

Regard the cycle B-C-D-F-E-B in Figure 1(a) as a *ring* of suitable capacity. After failure of link B-C (Figure 1(b)), nodes B and C (denoted switching nodes in [16]) detect a signal failure condition. Upon detection, node C sends over route C-D-F-E-B a request to B to perform protection switching (denoted bridge request in [16]). Likewise, node B sends a request to C over the reverse route. By the request messages, the intermediate nodes D, F, and E enter a pass-through state [16]. Node B, upon reception of the request from node C, switches to the protection path (more precisely [16], node B bridges traffic sent toward C on the disrupted link B-C to the protection path and selects the protection path to receive traffic from C). Node C operates analogously. After that, signaling reaches steady-state and the link failure is recovered.

As a reaction to node failures, the ring APS protocol behaves similarly to link failures and we describe only the differences. We depict the ring substructure of Figure 1(a) in Figure 2(a) including three connections which are protected by the ring. After failure of node C, the detecting ring neighbor nodes B and D send bridging requests to node C in direction away from the failure. Node D can deduce from the received bridging request (with destination C and sourced by node B) and its own detection of signal failure toward node C that node C failed. It then squelches (overwrites with logical ones) all traffic destined for node C (in the example the connection between C and D) and executes a ring bridge and switch for all other traffic. Node B operates analogously. With this operation, traffic passing through the failed node (in the example the connection from B to F) can survive and traffic destined for the failed node is squelched (since node C failed, it cannot source traffic). Squelching is necessary to avoid misconfiguration; in the example, without squelching connections C-E and C-D can be misconfigured to an unwanted connection D-E. Note that in this case, oblivious protection switching results in misconfiguration.

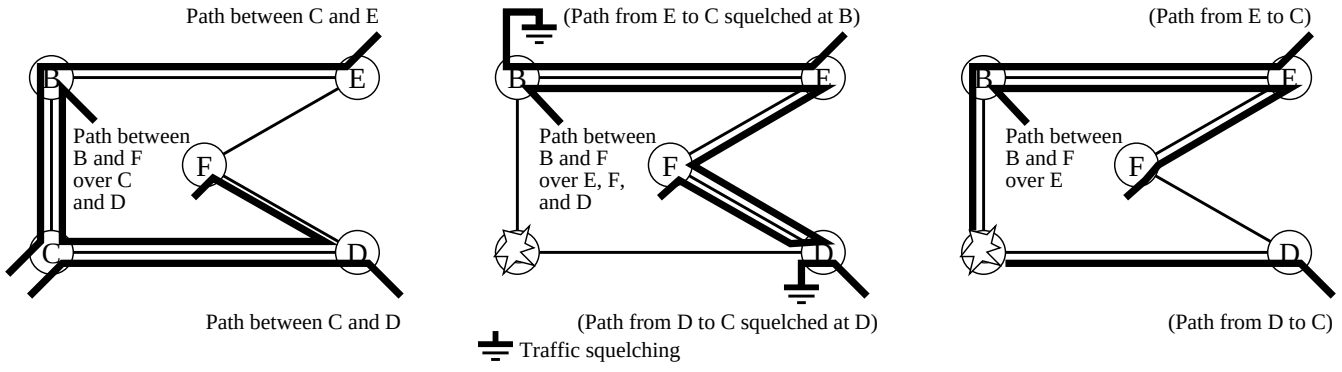

(a) A ring with three protected connections.

(b) Protection switching and squelching after failure of node C.

(c) Protection switching after failure of node C for the transoceanic application.

Fig. 2. Node protection principle of shared ring protection.

B. Adaptation to p-Cycles

In WDM networks, link-protection *p*-cycles can detect physical link or physical node failures on different levels. Nodes can detect link failures at the termination points of fiber links (WDM system layer) and of

WDM logical links (WDM logical layer).

Nodes detecting a signal failure (or degradation) initiate APS protocol instances depending on the layer. An APS protocol instance can be invoked

1) per failed fiber in the physical link, i.e., on the system WDM layer, *p*-cycle (this is similar to four fiber shared protection rings in SDH [16]).
2) per failed channel group in the fibers of the physical link, i.e., on the logical WDM layer (this is similar to two fiber shared protection rings in SDH [16]). The channel group has to belong to the same cycle.
3) per failed channel in the fibers of the physical link, i.e., on the logical WDM layer (this is similar to optical channel shared protection rings [14]).

Conceptually, the APS protocol processing effort in a node can become larger from Variant 1 to 3 for the same type of failure, since more instances become active. In a similar manner, signaling traffic increases. This becomes evident when relating the last two variants to Variant 1 which works on the fiber level. Variant 2 works on the channel group level and we expect few protection channel groups (at least one) per fiber with protection capacity. Variant 3 works on the channel level and, e.g., for systems with 160 wavelength channels, we deal with few dozens of protection channels per fiber with protection capacity.

On the other hand, the protection granularity becomes finer [17]: (i) the protection selectivity [18] between unprotected and protected traffic can be made on a per-channel basis in Variant 2 and 3 while it is on a per-fiber basis (thus, multiple channels) in Variant 1, (ii) a single channel failure (e.g., failure of channel termination equipment such as lasers) can affect protection switching of (still operating) other channels in Variant 1 and 2.

In this paper we assume Variant 3 (or an equivalent implementation of Variant 2), since the problem of high processing effort can, e.g., be resolved with high performance or parallel processing units per node. The signaling transmission effort can also be accommodated, since the APS can work on signaling information coded in few bytes. If signaling channels are provisioned per wavelength channel (e.g., overhead channels as defined in [19]), the operation will be equivalent to SDH. If signaling messages share a signaling channel, e.g., an optical supervisory channel (OSC), the channel bitrate has to be sufficient; for an estimation, assume 1000 APS instances each requiring 2×64 kbit/s; then a shared signaling channel with 128 Mbit/s is required. Although this signaling traffic is small compared with the bitrates of the user traffic, the system has to be designed for these signaling channel bitrates, which is, e.g., important for the evolution of point-to-point WDM systems whose OSCs have much lower bitrate requirements.

By choosing Variant 3, a *p*-cycle has one protection capacity unit on the links of a cycle, i.e., it is a closed connection loop. This is the way we use *p*-cycles in the investigations.

For adapting ring APS to *p*-cycles, we have a closer look to several *network objectives* in [16] for shared protection rings (with N nodes):

1) The switch completion time for a failure on a single physical link shall be less than 50 ms. This is required under the condition that all ring-nodes are in idle state, the ring does not carry extra traffic, and the ring circumference (physical cycle length) is less than 1200 km. On rings under all other conditions, the switch completion time can exceed 50 ms.
2) Extent of protection:
   a) For a single physical link (node) failure, the ring will restore all traffic that would be passing through the link (node) had no failure occurred.
   b) The ring shall restore all traffic possible.
3) Bidirectional protection switching (i.e., both directions are always switched to protection, even for unidirectional failures) shall be provided.
4) APS protocol:
   a) The protocol shall accommodate $N \leq 16$ nodes.
   b) Squelching shall be done at the switching nodes.

For *p*-cycles, we discuss and refine these network objectives.

It is plausible that Objective 1 (50 ms recovery time) can be met by *p*-cycles if it is met by rings. As

nodes also have to detect failures on straddling links and APS instances also have to process straddling links, *p*-cycle APS can have longer switching times than ring APS. We expect a resulting augmented delay to be marginal, since the extensions can exploit parallelization (especially for detection) and rely on simple rules and structure of the APS protocol. Since the protection switching for single on-cycle link failures is equivalent to shared protection rings (where the signaling path for ring switching passes through $N-1$ links), we can assume same protection switching times for on-cycle links. As the protection signaling path for straddling links, which passes through $\left\lfloor \frac{N-2}{2} \right\rfloor$ links in the worst case and only 2 links in the best case, is always shorter than for on-cycle link protection (except for artificial cases), signaling message transmission delay is not an issue introduced by straddling links.

For *p*-cycles therefore, the on-cycle links are the links for which the physical cycle length limit of 1200 km is essentially required, to achieve the 50 ms recovery time. For reasons as efficiency [20], [3] or geographical extent, it can be desirable to admit longer cycles. The "transoceanic application" of shared protection rings, as defined in Annex A of [16], relaxes this requirement. The transoceanic application extension is suggested if the distances between the nodes on the ring exceed 1500 km in SDH [16] as well as to simplify the analog engineering rules for ring design in optical networks [21]. Effectively, it reduces the path length during protection switching as depicted in Figure 2(c) for the node C failure (link failures are treated analogously). Instead of the nodes adjacent to the failure, the nodes terminating the connections on the ring (E and F) are responsible for protection switching, avoiding loops as on link D-F in Figure 2(b). This also abolishes the squelching mechanism, since source nodes do only protection switching for traffic other than destined for the failed node. Note that the routing in the transoceanic application for rings is similar to the "steering" concept in resilient packet ring (RPR) [22].

The transoceanic application is not a direct method to make switching faster for all connections, because only those recovered connections not terminated by the nodes adjacent to the failure location can profit from shorter signaling paths, and thereby can have less switching time than without the application. For the transoceanic application, however, [16] also relaxes the switch completion time requirement, to a maximum of 300 ms, although a limit of the ring circumference is not specified. If switching times longer than 50 ms but under 300 ms are possible for a network and if *p*-cycles should conform to ring standards [23], APS with transoceanic application can be an alternative for *p*-cycles. If switching times longer than 50 ms are admissible in general, we can use APS with or without transoceanic application for *p*-cycles.

Objective 2a (single failure recovery) can only be met for physical *link* failures. Figure 3 explains a situation where a static *p*-cycle cannot recover all traffic passing through a failed node. The figure shows a unit-capacity *p*-cycle A-B-C-D-E-F-A and three unit-capacity demands B-F, C-D, and C-E which are protected by the *p*-cycle and which pass through node A. The *p*-cycle can protect all single link failures. After failure of node A, however, the protection capacity is only available on the *p*-cycle segment B-C-D-E-F which can carry only *one* of the demands passing through node A. Therefore, if more than one path traverses a node (this is only possible for nodes with incident straddling links), the *p*-cycle can only protect at most one of the paths against node failure. Still, in the sense of Objective 2b (restore all traffic possible), the protocol should aim to restore as much traffic as possible, for which a protocol rule is necessary (e.g., as in Section IV). Similarly, the protocol has to incorporate rules for general multiple failures.

In this paper we stick to Objective 3 (bidirectional protection switching). Bidirectional protection switching can ease the repair of links if only a unidirectional part failed. Using unidirectional *p*-cycles, bidirectional protection switching necessitates coordination between *p*-cycles after unidirectional failures. Instead, as for rings, we postulate bidirectional *p*-cycles to avoid this coordination. Bidirectional *p*-cycles are not used in all investigations, notably in [4], [5], [6], [3], [24]. Note that bidirectional *p*-cycles can still imply different wavelengths between the two directions on the *p*-cycle [3].

Although Objective 4a ($N \leq 16$) can be met, it is desirable to allow more than 16 nodes for a *p*-cycle, since the longer *p*-cycles are, the more efficient they become [20], [3]. Note that the network and the system practically set length limitations, since larger *p*-cycles require longer protection paths that are more likely to suffer multiple failures and suffer from physical layer constraints.

Objective 4b (squelching at switching nodes), which is inapplicable for the transoceanic application, is

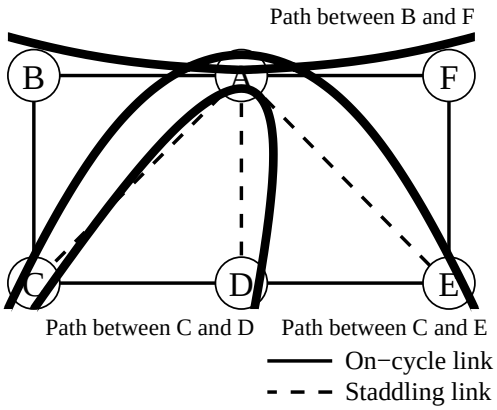

Fig. 3. Example for a *p*-cycle protecting three paths traversing a single node.

an issue for non-digital networks like transparent optical networks. To conform to this objective, means for squelching a general optical signal are necessary at the switching nodes.

Last but not least, the *p*-cycle APS protocol has to be able to protect straddling links. For this, we have to extend the logic of the ring APS protocol. This includes detection of failures on straddling links, issuing of switch commands from straddling links to on-cycle links, and the generation and processing of bridge requests to straddling neighbors. The latter extension requires the nodes' ring maps to be augmented by straddling link relationships. As a *p*-cycle protects two capacity units on a straddling link, an arbitration (which can be automatic) has to be supported at the link's end nodes to decide which unit is routed over which path during protection switching, i.e., an arbitration straddling link working capacity unit to east or west protection capacity unit. Note that Grover and Stamatelakis have previously detailed a way to add straddling links to BLSR switching functionality in [7], [25].

Recommendation [16] defines other features of shared protection rings, including support of "span switching" (only for four fiber rings) and "extra traffic" (traffic carried over the protection capacity and pre-emptible during failure events). It has to be decided if these features can be dropped for *p*-cycle APS to reduce protocol complexity. As *p*-cycles often need only little protection capacity [1], [3], it is arguable if the extra traffic feature (using this capacity) along with the required traffic differentiation scheme offers a significant added value.

It is also proposed that different *p*-cycles share their protection capacity [26]. This is possible at links protected by *p*-cycles which are routed on otherwise link-disjoint cycles, i.e., among the cycles the only common link is the shared link. Protection capacity sharing needs *p*-cycle inter-coordination in WDM networks, since multiple failures can cause contention for the shared protection capacity. The concept makes multiple *p*-cycle networks also more vulnerable to multiple failures. Therefore, the introduced complexity and vulnerability has to be weighted against the efficiency sharing brings in. It is worth to note that these "shared capacity *p*-cycles" are not fully pre-failure cross-connected structures anymore, rather they are equal to link-restoration with an artificial restriction to use only two restoration routes.

Although APS with transoceanic application is an attractive protocol candidate for optical network *p*-cycles, we assume APS without transoceanic application in the following, to be in line with the other published work.

C. Protection Switching Times

We discuss the technological feasibility of protection switching times. By the analogy to SDH rings we can say so far that a goal of 50 ms can be met. Operators can still insist in having 50 ms recovery time, since

service tariffs rely on it, there is too much effort in traffic segregation, or personnel may be averse to dealing with longer protection times [23]. But, relaxed protection switching times seem also to be acceptable [7]. For example, recovery time of large SONET rings can range from 100 to 150 ms [27].

Generalizing from SDH networks to WDM networks comes with a change in switching technology. Emerging optical switches, e.g., MEMS [14], are able to meet the timing requirements. An important point, favorable to speed up switching for *p*-cycles, especially for those working on the logical link level, is that the connection controller interface (CCI) should support parallel switching. But even in presence of longer switching times, it is worth noting that *p*-cycles have been invented to help in this issue, since the research which has led to *p*-cycles aimed at "an average-case speed-up for mesh networks where cross-connect time was the limiting factor" [28]; for *p*-cycles only two ring-like switchings are necessary, under the assumption of no protection capacity sharing with other *p*-cycles.

A general reasoning for *p*-cycles to fulfill the recovery time requirements, however, is that *p*-cycles do not oblige hard switching time requirements compared with other recovery mechanisms. In other words, if switching systems are designed to fulfill general switching time requirements for recovery, they will be usable for *p*-cycles.

IV. A Protocol Enhancement to Protect Against Node Failures

We exploit that an APS *p*-cycle protocol is able to protect some traffic against node failures. Unlike [10], where *p*-cycles protect path-segments and thus protect against node failures on the path-segment, we pursue a simpler approach requiring only an APS *p*-cycle protocol which is designed primarily to protect links. We assume only non-pre-emptible traffic, and as failure scenarios single physical link failures and single physical node failures. *p*-Cycles protect single link failures by default. As traffic originated or terminated at failed nodes cannot be recovered (by any scheme), we can concentrate on intermediate single node failures on a path.

In extension to the APS protocol in Section III, we introduce qualifiers for paths which should be protected against node failures. These qualified paths are routed only through those nodes, which are member of the same *p*-cycle. In other words, the path is protected against link and node failures from the first link after the ingress node of the *p*-cycle to the last link before the egress node of the *p*-cycle. As explained with Figure 3, at most one path per *p*-cycle can survive an intermediate node failure. Therefore, a single *p*-cycle can support multiple qualified paths given that each node on the *p*-cycle sees at most one qualified path traversing though it, i.e., qualified paths have to be node-disjoint per *p*-cycle.

The *p*-cycle nodes have to store the information which paths are qualified. After failure of an intermediate node of a qualified path, the adjacent (on-cycle and straddling) nodes detect the failure and initiate the signaling for protection switching. By deducing that a node failure is present, the adjacent nodes (or the ingress and the egress node in the transoceanic application method) perform protection switching for the qualified paths.

Several methods are viable to provide paths which survive intermediate node failures for a network with multiple link-protection *p*-cycles. We can guarantee for a path, which is routable as qualified path on a single *p*-cycle, node protection. Whenever a path, however, is protected by more than one *p*-cycle, the nodes at which the path changes from one *p*-cycle to another one is not protected by default. To protect against failures of these nodes, reconfiguration or *p*-cycle interconnection with dual homing ([14]-Section 10.2.5) can be deployed, especially for *p*-cycles on the system layer. Another approach is to ensure by design that between nodes demanding node-protected paths enough qualified paths are present, i.e., that the demanding node-pair is member of a single cycle, see the next section.

V. Design of *p*-Cycle Networks for Enhanced Node Failure Restorability

In this section we propose a method to design networks, which protect paths against failures of intermediate nodes. This method is based on the qualified path approach of Section IV. Conceptually, we can

embed the approach in a multiple quality of protection (QOP) environment. We differentiate between demands which should be protected against single link failures (QOP_l) and demands which should be protected against both single link failures and single node failures ($\text{QOP}_{n,l}$). The differentiation can be done according to service level agreements, in which the more-protected class $\text{QOP}_{n,l}$ is offered at a higher cost than class QOP_l, or according to availability requirements, e.g., where $\text{QOP}_{n,l}$ is chosen for paths over unreliable nodes.

In the design, a demand unit of class $\text{QOP}_{n,l}$ is served by a qualified path, i.e., a path which is protected by a single *p*-cycle and which is node-disjoint to other $\text{QOP}_{n,l}$ paths protected by that *p*-cycle. Paths for class QOP_l demands can traverse multiple *p*-cycles.

A precondition for this design is that we can interconnect each $\text{QOP}_{n,l}$ demand pair by at least one *p*-cycle, as conveyed by Figure 4. This is possible in two-node-connected networks. If this is impossible because of other reasons (e.g., capacity restrictions), we can modify the design, e.g., to minimizing the number of *p*-cycles these paths traverse.

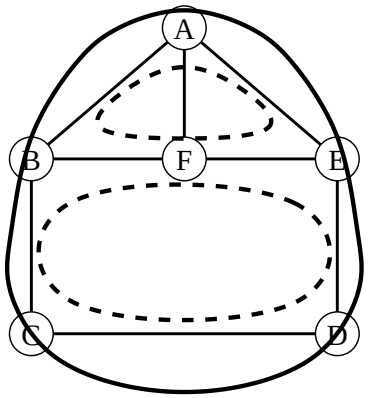

Fig. 4. Example for the design of *p*-cycle networks for enhanced node failure restorability. Using the dashed cycles, link protection of the network is possible, however, another cycle, e.g., the solid cycle, is needed to make the network node-protected by the qualified path approach. Using both the dashed cycles and the solid cycle (or alternatively only an Hamiltonian cycle), every node pair is reachable by a single cycle.

We extend the basic "joint model" [4], [5], [29] by constraints enforcing the required node-survivability. These constraints ensure the solution contains *p*-cycles that fully house the qualified paths. We assume WDM networks with full wavelength conversion or with enough converters deployed (e.g., using the architectures in [30]), such that wavelength constraints do not impose further restrictions on the design.

The network to be protected is modeled by graph $G = (V, E)$. Associated with the graph are edge costs $z_e, \forall e \in E$, working capacities $w_e, \forall e \in E$, and protection capacities $p_e, \forall e \in E$.

Based on (precomputed) cycles with index-set C, the two size-$(|E|, |C|)$ matrices $\mathbf{\Pi}$ and $\mathbf{\Phi}$ are given ($|S|$ is the number of elements in set S). The cycles can be computed by, e.g., a breadth first search algorithm. The cycle set can comprise all possible cycles in the network or a limited subset to decrease computation time (see, e.g., [29], [31]). An entry $\pi_{e,k} \in \{0, 1\}$ of matrix $\mathbf{\Pi}$ indicates if edge $e \in E$ is element of cycle k. An entry $\phi_{e,k} \in \{0, 1, 2\}$ of matrix $\mathbf{\Phi}$ indicates how many working capacity units on edge $e \in E$ are protectable by a *p*-cycle unit on cycle $k \in C$. These latter matrix entries correspond to the useful paths, as denoted in [1], i.e., the number of useful paths a *p*-cycle can offer for protection of the edge.

For the *p*-cycles configuration, we are interested in the number of *p*-cycle units n_k for each cycle $k \in C$ that is needed, i.e., the cycle capacity (or *p*-cycle group capacity) on cycle k.

For the joint model, we additionally have to model working capacity computation. For a path-flow formulation, we are given a set of demand pairs $D \subset V \times V$, and for each demand pair $\delta \in D$ the number of demands d_δ, an index-set of eligible working paths Q_δ, and indicators $\gamma_{e,q} \in \{0, 1\}$ being one if link $e \in E$ is on working path $q \in Q_\delta$, zero otherwise. Additional solution variables are the demand flow variables $f_{\delta,q}$.

We can design the network with the following integer linear program (ILP) formulation. As additional

input, we are given for each demand pair $\delta \in D$ the number of demands d^*_δ, $d^*_\delta \leq d_\delta$, which require QOP$_{n,l}$.

With a separate function we identify those path-cycle-pairs for a demand pair $\delta \in D$, which satisfy that the edges of the path are entirely protectable by the *p*-cycle. We define the function as

$$\eta_{\delta,q,k} = \begin{cases} 1 & \text{if } \gamma_{e,q} \leq \phi_{e,k},\ \forall e \in E \\ 0 & \text{otherwise} \end{cases}, \quad \begin{matrix} \forall \delta \in D, q \in Q_\delta, \\ k \in C \end{matrix}. \tag{1}$$

We state the ILP as follows:

$$\min \quad \sum_{e \in E} z_e (w_e + p_e) \tag{2}$$

$$p_e = \sum_{k \in C} \pi_{e,k} n_k,\ \forall e \in E \tag{3}$$

$$w_e \leq \sum_{k \in C} \phi_{e,k} n_k,\ \forall e \in E \tag{4}$$

$$w_e = \sum_{\delta \in D} \sum_{q \in Q_\delta} \gamma_{e,q} f_{\delta,q},\ \forall e \in E \tag{5}$$

$$d_\delta = \sum_{q \in Q_\delta} f_{\delta,q},\ \forall \delta \in D \tag{6}$$

$$d^*_\delta = \sum_{q \in Q_\delta} \sum_{k \in C} \eta_{\delta,q,k}\, f^*_{\delta,q,k},\ \forall \delta \in D \tag{7}$$

$$n_k \geq \sum_{\substack{(a,b) \in D, q \in Q_{(a,b)}, \\ (n,m) \in E : a \neq m, b \neq m}} \frac{1}{2} \gamma_{(n,m),q}\, \eta_{(a,b),q,k}\, f^*_{\delta,q,k}, \quad \forall m \in V, k \in C \tag{8}$$

$$w_e \in [0, \infty),\ \forall e \in E \tag{9}$$

$$p_e \in [0, \infty),\ \forall e \in E \tag{10}$$

$$f_{\delta,q} \in \{0, 1, 2, \ldots\},\ \forall \delta \in D, q \in Q_\delta \tag{11}$$

$$f^*_{\delta,q,k} \in \{0, 1, 2, \ldots\}, \quad \begin{matrix} \forall \delta \in D, q \in Q_\delta, \\ k \in C : \eta_{\delta,q,k} > 0 \end{matrix} \tag{12}$$

$$n_k \in \{0, 1, 2, \ldots\},\ \forall k \in C \tag{13}$$

The Objective (2) minimizes the cost-weighted total capacity. Constraints (3) determine the protection capacity allocation, and Constraints (4) ensure the working capacity to be protected. Note that the right-hand side of (4) is the protect*able* capacity which can be larger than the working capacity. Constraints (5) determine a link's working capacity by the sub-flows of all demands on the link. For a demand pair, Constraints (6) equate the demand to its individual sub-flows on the given paths. Constraints (7) require that the demand units of QOP$_{n,l}$ are assigned to a path-cycle-pair. Constraints (8) ensure that the capacity of a cycle is at least as high as the QOP$_{n,l}$ traffic passing through each of its nodes, i.e., that each *p*-cycle has at most one intermediate node of a QOP$_{n,l}$ path. Hence, Constraints (8) guarantee that the protocol modification for node protection can be used for these paths.

The working and protection capacity variables, which are auxiliary variables here, are defined in (9) and (10), respectively. The integer flows are contained in (11). Variables (12) indicate that a portion of the QOP$_{n,l}$ demand is assigned both to a single path and to a single cycle. In a more detailed but also more complex approach, we can perform this association with individual demand units instead of demand portions. Finally, the integer *p*-cycle units are defined in (13).

This approach adds several variables and constraints to the basic joint model. In the worst case, the number of additional integer variables from (12) is $(\sum_{\delta \in D} |Q_\delta|)|C|$ and the number of additional constraints from (8)-(7) is $|V||C| + |D|$.

A. Comparative Results

In this section we compare the *p*-cycles QOP design with dedicated path protection. We study two optical transport network reference scenarios as suggested in [32]. The `ger_net` network (see Figure 5) is based on a hypothetical Germany backbone network as used and the `nsf_net` network (see Figure 6) is based on the National Science Foundation Network (NSFNET). We scale the demands of [32] by 10. The link

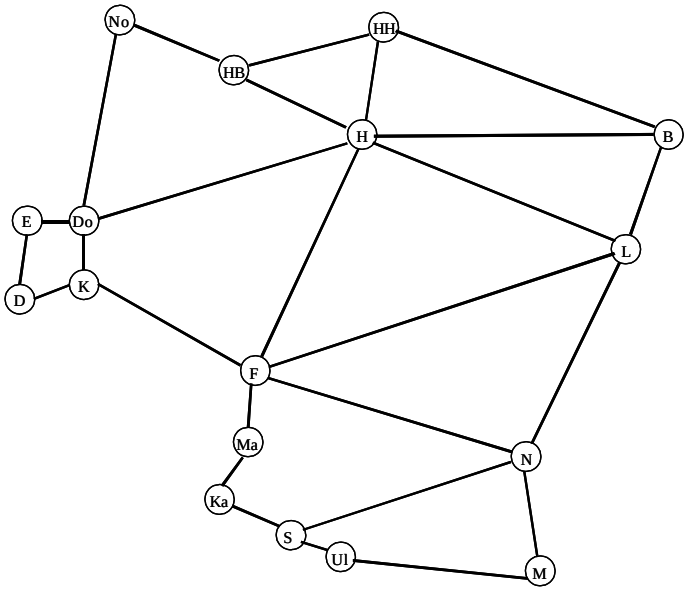

Fig. 5. The `ger_net` topology.

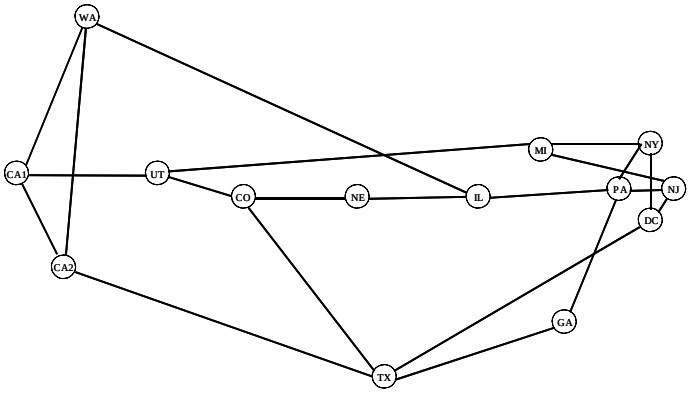

Fig. 6. The `nsf_net` topology.

capacity cost is based on length-weighted utilization, i.e., the cost z_e of a link $e \in E$ is set to its physical length. Note especially for `nsf_net` that we are predominately interested in the topological structure here, without regarding limitations through the geographical extent. The problems are solved using CPLEX [33].

The design for dedicated path protection (e.g., 1+1 path protection) is obtained as follows. Per demand unit, two link-disjoint paths are found which minimize the sum of their physical lengths. The overall cost is the sum of these lenghts-sums over all demand units, which is also the minimum overall cost solution. Since node-disjoint routing includes link-disjoint routing, the cost values are a lower bound for dedicated path protection that guarantees protection of (intermediate) nodes. Hence, the cost values for node-protection can be even higher.

For both networks, all possible cycles are precomputed and used as input to the *p*-cycles ILP. For the QOP designs we assume each demand pair to have a fixed portion ϵ of its demand for QOP$_{n,l}$. Table I shows the cost results for *p*-cycles and dedicated path protection relative to the cost for an unprotected network.

TABLE I

DESIGN COSTS RELATIVE TO THE COST FOR AN UNPROTECTED NETWORK FOR TWO STUDY NETWORKS OPERATED WITH *p*-CYCLE DESIGN UNDER QOP DIFFERENTIATION AND DEDICATED PATH PROTECTION. FOR *p*-CYCLES, ϵ DENOTES THE PORTION OF THE NODE PROTECTED DEMAND.

	ger_net	nsf_net
Dedicated path protection	2.55	2.80
p-Cycles with $\epsilon = 0$	1.72	1.81
p-Cycles with $\epsilon = 0.5$	1.72	1.81
p-Cycles with $\epsilon = 1$	1.78	1.83

The cost in excess to an unprotected design is significantly higher with dedicated path protection than with *p*-cycles. It is also apparent that *p*-cycles can provide full node-failure survivability (i.e., $\epsilon = 1$) at marginal additional cost (compare with $\epsilon = 0$). Without node protection ($\epsilon = 0$), the optimal *p*-cycle solutions for the networks contain long cycles (also Hamiltonian cycles), which eases dedication of paths to single *p*-cycles (as in the proposed manner). Therefore, we can expect these smaller cost differences when increasing ϵ, but we expect larger ones when cycle length restrictions are introduced.

VI. AVAILABILITY MODELS FOR *p*-CYCLES

In this section we derive availability models for *p*-cycles, to consider the influence of straddling links (which is an important parameter, e.g., in migration considerations from rings to *p*-cycles) and the improvement in availability for the (node-protected) QOP$_{n,l}$ traffic as in Section V.

In accordance with [34], we define the instantaneous availability as the probability that an entity is in the up-state at a given instant of time, and the mean availability as the normalized integral of the instantaneous availability in a given time interval. The entity can be a node or a link, or it can be a connection. We deploy the term "availability" in the sense of mean availability.

We calculate the all-terminal availability and the two-terminal availability [35]:

- The *all-terminal availability* is the probability that all nodes in the network are operating and can communicate with each other. The all-terminal availability reflects the network operator's point-of-view.
- The *two-terminal availability* is the probability that two given nodes in the network can communicate with each other, independently of the states of other parts in the network. The two-terminal availability reflects the user's point-of-view.

For an availability assessment of the *p*-cycle concept, we consider a *p*-cycle represented by its nodes and links on the physical layer. Nodes and links can have two states, either operating or failed. In addition, we assume that failures of nodes and links occur independently and consider only those states in which, at any time, at most either one node failure or one link failure is present, i.e., we neglect higher order failures. A model including more detail is also possible [36], if the exact protocol functionality and the configuration details are given. The model, however, provides sufficient accuracy, if node and link failures are the most probable failure scenarios, which is often the case. Furthermore, we assume all nodes to have the same availability A_n and all links to have the same availability A_l. Since links can differ in their lengths and thus in their availabilities, we can obtain worst-case approximations by using the lowest availability value of the links.

The models provide a method to analyze the influence of straddling links and the availability contribution for QOP$_{n,l}$ traffic compared with QOP$_l$ traffic. We can calculate the availabilities for a *p*-cycle with ν on-cycle links and μ straddling links as follows.

Two mutually exclusive events contribute to the all-terminal availability. In the first event, all parts of the *p*-cycle operate

$$A'_{all} = A_n^{\nu} A_l^{\nu+\mu} \tag{14}$$

and, in the second event, some link failed (this can happen by failure of one of the $\nu + \mu$ links)

$$A''_{all} = (\nu + \mu) A_n^{\nu} (1 - A_l) A_l^{\nu+\mu-1}. \tag{15}$$

In general, the all-terminal availability sums to

$$A_{all} = A'_{all} + A''_{all}. \tag{16}$$

Three mutually exclusive events contribute to the two-terminal availability. The first two yield (14) and (15). For the third event, we distinguish between $A_{st,1}$, the availability of paths which are protected against any single failures of their intermediate nodes, and $A_{st,2}$, the availability of paths which are not recovered if intermediate straddling nodes fail (a node on a *p*-cycle is called straddling node if it terminates at least one straddling link of the *p*-cycle). We associate $A_{st,1}$ and $A_{st,2}$ to services QOP$_{n,l}$ and QOP$_l$, respectively. The availability for QOP$_l$ can be slightly higher, since the *p*-cycle can still aim to recover a QOP$_l$ path from a straddling node failure if no QOP$_{n,l}$ path passes through the node.

After failure of a node without straddling links, *p*-cycles can recover any traversing path, be it of class QOP$_{n,l}$ or of class QOP$_l$. The difference between the two classes appears at straddling nodes. Denote χ as the number of intermediate straddling nodes on the path between the two terminal nodes.

For QOP$_{n,l}$, if some node except for the terminal nodes has failed (this can happen by failure of one of the $\nu - 2$ non-terminal nodes), the terminals can still communicate with each other

$$A'_{st,1} = (\nu - 2)(1 - A_n) A_n^{\nu-1} A_l^{\nu+\mu}. \tag{17}$$

For QOP$_l$, if some node except for the terminal nodes and except for nodes with incident straddling has failed (this can happen by failure of one of the $\nu - 2 - \chi$ non-terminal non-straddling nodes), the terminals can still communicate with each other

$$A'_{st,2} = (\nu - 2 - \chi)(1 - A_n) A_n^{\nu-1} A_l^{\nu+\mu}. \tag{18}$$

Note that this availability is zero in the worst case, i.e., if a path for terminals, which are neighbors by an on-cycle link, is routed on-cycle on the long route and every node is straddling node (then, $\chi = \nu - 2$). We compute the two-terminal availabilities as

$$A_{st,1} = A_{all} + A'_{st,1} \text{ and } A_{st,2} = A_{all} + A'_{st,2}. \tag{19}$$

Hence, the worst case difference in two-terminal availability between QOP$_{n,l}$ and QOP$_l$ is $A'_{st,1}$ in (17). In general for QOP$_l$, the availability is path dependent, since by (18) and if both A_n and A_l are close to one, the availability for QOP$_l$ will be diminished by a node's unavailability $(1 - A_n)$ every time a straddling node is visited. Therefore, routing should aim to minimize the number of straddling nodes for QOP$_l$ paths.

Now we compare the availability performance of *p*-cycles and rings. We assume that the availability of *p*-cycle nodes (links) is equal to the availability of ring nodes (links). It is easy to show that A_{all}, $A_{st,1}$, and $A_{st,2}$ decrease monotonically with μ, if the number of nodes stays the same and, for $A_{st,2}$, if χ increases monotonically with μ, e.g., if the paths on the *p*-cycle do not change. Hence, whenever we introduce straddling links to a ring (or generally to a *p*-cycle), both the all-terminal availability and the two-terminal availability decrease (or stay for $A_l = 1$). This is expected, since adding straddling links does not contribute to more protection path opportunities. A fairer comparison, however, considers the number of links which can be protected. Say, we are able to protect τ links by either a ring or a *p*-cycle. Then, for a ring we set $\nu_r = \tau$ ($\mu_r = 0$) and for a *p*-cycle we set $\nu_p + \mu_p = \tau$, to obtain

$$\frac{A_{all}(\text{Ring})}{A_{all}(p\text{-Cycle})} = A_n^{\mu_p}. \tag{20}$$

If $A_n < 1$ holds, the all-terminal availability of a *p*-cycle realization with at least one straddling link is higher than the all-terminal availability of the ring realization, since less nodes are involved for *p*-cycles in the protection of the links.

We cannot make the comparison for the two-terminal availability, since the node-failure event enumeration in (17) and (18) is different. As *p*-cycles realize protection with fewer nodes than rings, the availability contribution because of single node failure events is less in the calculation (17) and (18).

VII. Conclusions

This paper addressed several *p*-cycle protection switching protocol issues. We envisaged an extended ring Automatic Protection Switching (APS) protocol as candidate for *p*-cycles and discussed the involved protocol changes and network objectives for *p*-cycles.

Because of the ability to protect more links (i.e., straddling links), we explained why *p*-cycles may not be able to recover all traffic transiting through a failed node as rings do. However, we proposed a protocol enhancement which protects qualified paths against node failures. In a case study, we have shown that *p*-cycle network design which is resilient against node failures (by restricting paths to qualified paths) is still much more cost efficient than a dedicated path protection design.

We formulated availability models which are useful for the *p*-cycle network design, to provide service guarantees. Using the availability models, we have shown that the protocol modification (for node protection) enhances the availability for paths passing through straddling nodes. We also compared the availability of rings with the availability of *p*-cycles.

References

[1] W. D. Grover and D. Stamatelakis, "Cycle-oriented Distributed Preconfiguration: Ring-like Speed with Mesh-like Capacity for Self-planning Network Restoration," in *Proceedings of the IEEE International Conference on Communications (ICC)*, (Atlanta, GA, USA), June 1998.

[2] W. D. Grover and D. Stamatelakis, "Bridging the ring-mesh dichotomy with *p*-cycles," in *Proceedings of the International Workshop on Design of Reliable Communication Networks (DRCN)*, (Munich, Germany), Apr. 2000, Invited Talk.

[3] D. A. Schupke, C. G. Gruber, and A. Autenrieth, "Optimal Configuration of p-Cycles in WDM Networks," in *Proceedings of the IEEE International Conference on Communications (ICC)*, (New York City, NY, USA), April–May 2002.

[4] C. G. Gruber, "Resilient Networks With Non-Simple p-Cycles," in *Proceedings of the International Conference on Telecommunications (ICT)*, (Papeete, Tahiti, French Polynesia), Feb. 2003.

[5] C. Mauz, "p-cycle Protection in Wavelength Routed Networks," in *Proceedings of the Working Conference on Optical Network Design and Modelling (ONDM)*, (Budapest, Hungary), Feb. 2003.

[6] D. Rajan and A. Atamtürk, *Telecommunications Network Design and Management (Editors: G. Anandalingam and S. Raghavan)*, chapter Survivable Network Design: Routing of Flows and Slacks, pp. 65–81, Kluwer Academic Publishers, Boston, Dordrecht, London, 2003.

[7] W. D. Grover, *Mesh-based Survivable Networks: Options and Strategies for Optical, MPLS, SONET and ATM Networking*, Prentice Hall, Upper Saddle River, 2003.

[8] D. Stamatelakis and W. D. Grover, "IP Layer Restoration and Network Planning Based on Virtual Protection Cycles," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 10, pp. 1938–1949, Oct. 2000.

[9] W. D. Grover and G. Shen, "Extending the p-Cycle Concept to Path-Segment Protection," in *Proceedings of the IEEE International Conference on Communications (ICC)*, (Anchorage, AK, USA), May 2003.

[10] G. Shen and W. D. Grover, "Extending the p-Cycle Concept to Path Segment Protection for Span and Node Failure Recovery," *IEEE Journal on Selected Areas in Communications*, vol. 21, no. 8, pp. 1306–1319, Oct. 2003.

[11] D. A. Schupke, "Multiple Failure Survivability in WDM Networks with p-Cycles," in *Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS)*, (Bangkok, Thailand), May 2003, Invited Paper.

[12] D. A. Schupke, "The Tradeoff Between the Number of Deployed p-Cycles and the Survivability to Dual Fiber Duct Failures," in *Proceedings of the IEEE International Conference on Communications (ICC)*, (Anchorage, AK, USA), May 2003.

[13] D. A. Schupke, W. D. Grover, and M. Clouqueur, "Strategies for Enhanced Dual Failure Restorability with Static or Reconfigurable p-Cycle Networks," in *Proceedings of the IEEE International Conference on Communications (ICC)*, (Paris, France), June 2004, Accepted contribution.

[14] R. Ramaswami and K. Sivarajan, *Optical Networks. A Practical Perspective*, Morgan Kaufmann, San Francisco, 2nd edition, 2002.

[15] D. Stamatelakis and W. D. Grover, "OPNET Simulation of Self-organizing Restorable SONET Mesh Transport Networks," in *CD-ROM Proceedings of OPNETWORKS Conference*, (Washington, DC, USA), Apr. 1998.

[16] "ITU-T Recommendation G.841 - Types and characteristics of SDH network protection architectures," International Telecommunication Union, Telecommunication Standardization Sector (ITU-T), Oct. 1998.

[17] P. Demeester, M. Gryseels, A. Autenrieth, C. Brianza, L. Castagna, G. Signorelli, R. Clemenfe, M. Ravera, A. Lajszczyk, D. Janukowicz, K. Van Doorselaere, and Y. Harada, "Resilience in multilayer networks," *IEEE Communications Magazine*, vol. 37, no. 8, pp. 70–76, Aug. 1999.

[18] T1 A1.2, Working Group on Network Survivability Performance, "Technical Report on Enhanced Network Survivability Performance," Tech. Rep. T1.TR.68-2001, Alliance for Telecommunications Industry Solutions, Feb. 2001.

[19] "ITU-T Recommendation G.709/Y.1331 - Interfaces for the optical transport network (OTN)," International Telecommunication Union, Telecommunication Standardization Sector (ITU-T), Feb. 2001.
[20] D. Stamatelakis and W. D. Grover, "Theoretical Underpinnings for the Efficiency of Restorable Networks Using Preconfigured Cycles ("p-cycles")," *IEEE Transactions on Communications*, vol. 48, no. 8, pp. 1262–1265, Aug. 2000.
[21] J. Manchester, P. Bonenfant, and C. Newton, "The evolution of transport network survivability," *IEEE Communications Magazine*, vol. 37, no. 8, pp. 44–51, Aug. 1999.
[22] D. A. Schupke, "A Throughput Study of RPR Networks Subject to Single Failures," in *Proceedings of the European Conference on Networks and Optical Communications (NOC)*, (Darmstadt, Germany), June 2002.
[23] D. E. Smith, E. E. Basch, K. A. DeMartino, R. V. Egorov, S. Gringeri, R. S. Kalbag, S. Liu, and V. Shukla, "Protection Switching Performance in Next Generation Optical Transport Networks," in *Proceedings of the International Workshop on Design of Reliable Communication Networks (DRCN)*, (Banff, AB, Canada), Oct. 2003.
[24] Z.-R. Zhang and W.-D. Zhong, "Design of Survivable WDM Network Using P-Cycles," in *Proceedings of the Conference on the Optical Internet/Australian Conference on Optical Fibre Technology*, (Melbourne, Australia), July 2003.
[25] W. D. Grover and D. Stamatelakis, "Scalable Network Restoration Device," US Patent 6,434,704, Jun 2002.
[26] H. Huang and J. A. Copeland, "A Series of Hamiltonian Cycle-Based Solutions to Provide Simple and Scalable Mesh Optical Network Resilience," *IEEE Communications Magazine*, vol. 40, no. 11, pp. 46–51, Nov. 2002.
[27] L. Lipes, "Understanding the trade-offs associated with sharing protection," in *Proceedings of the IEEE/OSA Optical Fiber Communication Conference (OFC)*, (Anaheim, CA, USA), Mar. 2002.
[28] W. Grover, J. Doucette, M. Clouqueur, D. Leung, and D. Stamatelakis, "New Options and Insights for Survivable Transport Networks," *IEEE Communications Magazine*, vol. 40, no. 1, pp. 34–41, Jan. 2002.
[29] W. D. Grover and J. E. Doucette, "Advances in Optical Network Design with p-Cycles: Joint optimization and pre-selection of candidate p-cycles," in *Proceedings of the IEEE/LEOS Summer Topical Meeting on All-Optical Networking*, (Mont Tremblant, QC, Canada), July 2002.
[30] D. A. Schupke, M. C. Scheffel, and W. D. Grover, "Configuration of p-Cycles in WDM Networks with Partial Wavelength Conversion," *Photonic Network Communications, Journal, Kluwer Academic Publishers*, vol. 6, no. 3, pp. 239–252, Nov. 2003.
[31] J. Doucette, D. He, W. D. Grover, and O. Yang, "Algorithmic Approaches for Efficient Enumeration of Candidate p-Cycles and Capacitated p-Cycle Network Design," in *Proceedings of the International Workshop on Design of Reliable Communication Networks (DRCN)*, (Banff, AB, Canada), Oct. 2003.
[32] R. Hülsermann, S. Bodamer, M. Barry, A. Betker, C. Gauger, M. Jäger, M. Köhn, and J. Späth, "A Set of Typical Transport Network Scenarios for Network Modelling," in *Proceedings of ITG-Fachtagung Photonische Netze*, (Leipzig, Germany), May 2004.
[33] ILOG Inc., "CPLEX," `http://www.ilog.com/`, 2003.
[34] "ITU-T Recommendation G.911 - Parameters and calculation methodologies for reliability and availability of fibre optic systems," International Telecommunication Union, Telecommunication Standardization Sector (ITU-T), Apr. 1997.
[35] D. R. Shier, *Network Reliability and Algebraic Structures*, Clarendon Press, Oxford, 1991.
[36] D. A. Schupke, "Reliability Models of WDM Self-Healing Rings," in *Proceedings of the International Workshop on Design of Reliable Communication Networks (DRCN)*, (Munich, Germany), Apr. 2000.