Int. J. Electron. Commun. (AEÜ) 67 (2013) 470–478

Contents lists available at SciVerse ScienceDirect

International Journal of Electronics and Communications (AEÜ)

journal homepage: www.elsevier.com/locate/aeue

Hybrid FRR/*p*-cycle design for link and node protection in MPLS networks☆

Chang Cao[a,b,*], George N. Rouskas[c], Jianquan Wang[d], Xiongyan Tang[d]

[a] China United Network Communications Co., Ltd., Postdoctoral Workstation, Beijing 100033, China
[b] Beijing University of Posts and Telecommunications, Beijing 100876, China
[c] Department of Computer Science, North Carolina State University, Raleigh, NC 27695-8206, USA
[d] China Unicom Research Institute, Beijing 100032, China

ARTICLE INFO

Article history:
Received 7 July 2012
Accepted 11 November 2012

Keywords:
Multi-protocol label switching
Fast reroute
Pre-configure cycle

ABSTRACT

Survivable MPLS technologies are crucial in ensuring reliable communication services. The fast reroute (FRR) mechanism has been standardized to achieve fast local repair of label switched paths (LSPs) in the event of link or node failures. We present a suite of hybrid protection schemes for MPLS networks that combine the well-known *p*-cycle method with FRR technology. Whereas with pure FRR backup paths are planned by each node individually, the hybrid schemes employ a set of *p*-cycles that may be selected using techniques that take a holistic view of the network so as to share protection bandwidth effectively. The hybrid FRR/*p*-cycle methods are fully RFC 4090-compliant, yet allow network operators to leverage a large existing body of *p*-cycle design techniques. Numerical results on realistic network topologies indicate that the hybrid approach is successful in combining the advantages of *p*-cycle design and FRR.

1. Introduction

Multi-protocol label switching (MPLS) [1], originally developed to enable fast packet forwarding, has also facilitated traffic engineering, quality-of-service (QoS) routing, and differentiated services support in IP-based metro and backbone networks [2]. MPLS technology is widely deployed and is crucial to the operation of the Internet and its ability to support critical communication services efficiently. Consequently, MPLS survivability mechanisms [3] are key to ensuring that the network may continue to provide reliable services even in the presence of failures. In particular, with today's multi-layer network architectures, it may be more economical for IP/MPLS layer operators to restore traffic within their own IP/MPLS logical environment rather than relying on physical layer restorability [4–6].

There are mainly two types of failures that network operators must design the network to withstand. Link failures (e.g., due to fiber cuts or the malfunction of active components such as transponders) are usually handled by the physical layer first. But if such a failure is not restored within a certain period of time (typically, a few tens of milliseconds), the MPLS layer must initiate its own recovery actions. Node failures may be due to router crashes or router restarts after the routine application of software patches or upgrades, and may occur as frequently as, or even more often than link failures [7]. Such failures must be dealt with directly at the MPLS layer.

The IETF has standardized the fast reroute (FRR) mechanism [8] for protecting label switched paths (LSPs) in MPLS networks. FRR calls for local repair actions in the event of a link or node failure. Specifically, affected traffic is re-directed onto pre-configured backup tunnels by two nodes adjacent to the failed link or node. As a result, all affected LSPs are rerouted to backup paths within a few tens of milliseconds.

The pre-configure cycle (*p*-cycle) scheme [9] also employs local repair actions to re-direct traffic from the failed link or node onto a backup path along a pre-configured cycle. This method provides ring-like protection speeds with mesh (span-restorable) capacity efficiency, and it has been studied extensively (for a survey of related work, see [10]). Although originally designed for protection in the optical layer, *p*-cycle technology can be applied to the IP [7] or MPLS [11] layers.

There is an extensive technical literature on network survivability design, optimization, and performance evaluation using *p*-cycles in the WDM and IP layers [10]. For packet-switched networks, in particular, it was shown in [7] that by using integer linear programming (ILP) methods, *p*-cycle design can be as capacity-efficient as optimized span restoration. However, obtaining exact optimal solutions is an NP-hard combinatorial problem that does not scale to large networks. As a result, a number of relaxations and heuristics must be considered for practical application of *p*-cycle selection to realistic network topologies. Another two studies

☆ Chang Cao was supported by China National Science and Technology Major Project (2010ZX03004-002-02).
* Corresponding author at: China United Network Communications Co., Ltd., Postdoctoral Workstation, Beijing 100033, China. Tel.: +86 13810021220.
E-mail address: caochang@chinaunicom.cn (C. Cao).

1434-8411/$ – see front matter © 2012 Elsevier GmbH. All rights reserved.
http://dx.doi.org/10.1016/j.aeue.2012.11.005

(first [11] and then [12]) presented a group of mixed integer program (MIP) formulations for p-cycle design in MPLS networks, and investigated the relationship between protection bandwidth requirements and traffic load distribution. Besides MIP formulations' discussion, study in [13] focused on the issue of p-cycles protection switching protocols. It explained that p-cycles may not be able to recover all traffic transiting through a failed node as rings did, and proposed a protocol enhancement which protected qualified paths against node failures. The concept of path-segment protecting p-cycles was first described in [14], it extended the ability of p-cycles to protect node by restoring all of relevant path segments. Later, this method was developed in [15], which let the cycles act as p-cycles for end-to-end paths between nodes on the cycle, but only allowed each cycle to provide protection relationships to a group of paths whose routes are all mutually disjoint. In order to simplify the design for node protection, study in [16] reported a new strategy that integrated the BLSR-like behavior of ordinary p-cycles under on-cycle node failure conditions with a new straddling-oriented use of the p-cycles for node protection, and employs only one set of p-cycles over the whole network. More recently, a new shared-segment protection to restore node failure using ordinary p-cycles was introduced in [17].

The motivation for our work is based on the observation that both FRR and p-cycle are local repair protection schemes, hence we expect the network operation, delay, and overhead incurred for failure detection, notification, and triggering of restoration action to be similar for the two technologies. However, FRR backup tunnels are typically planned individually by the nodes adjacent to the protected link or node, whereas p-cycle design takes a more holistic view of the network in determining the protection cycles so as to share spare resources effectively. Therefore, we first introduce the signaling and local repair (protection switching) methods of FRR into p-cycle's restoration, which are not employed in p-cycles' operation before, and then propose several novel design technique (i.e., area division p-cycles and FRR-based p-cycles). Together with planning methods of Hamiltonian p-cycle and node encircling p-cycle in previous studies [10], we compare these four types of p-cycles' performance in MPLS networks with different topologies. Another contribution of our work is to introduce network holistic p-cycle designs, which means to use the same set of p-cycles to protect both link and node failures.

The rest of the paper is organized as follows. In Section 2, we briefly review the FRR method for MPLS link and node protection. In Section 3, we describe how to combine the p-cycle and FRR methods, and describe several approaches for selecting the set of p-cycles. In Section 4 we present three performance metrics to compare the pure FRR and hybrid schemes. We present numerical results in Section 5, and conclude the paper in Section 6.

2. Pure FRR protection design

RFC 4090 [8] defines two schemes for dealing with link and node failures, respectively. Consider first the case of link failure, e.g., the failure of directed link $B \rightarrow C$ in the 8-node MPLS network shown in Fig. 1. The upstream router B is referred to as the "point of local repair" (PLR) with respect to protecting traffic in the event that the link fails, while the downstream router C adjacent to the other end of the link is known as the "merge point" (MP). On the other hand, if a node fails, all links incident to the node (i.e., the three bi-directional links $C \leftrightarrow D$, $D \leftrightarrow H$, and $D \leftrightarrow G$, in the case of failed node D in the network of Fig. 1) are affected. In other words, we can regard a node failure as the simultaneous loss of all link pairs (i.e., 2-hop paths) with the failed node in the middle. In this situation, all neighbors of the failed node act as the PLR or MP, where the specific role of each neighbor is determined by the direction of the specific traffic flow considered.

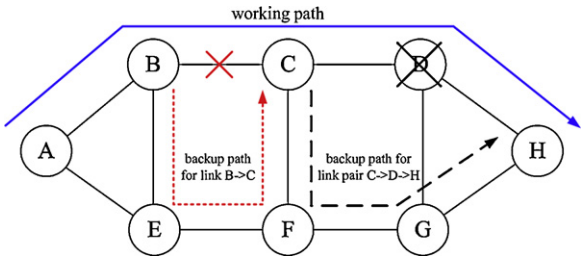

Fig. 1. Link/node protection with pure FRR.

Regardless of whether the failure involves a link or node, implementation of FRR protection consists of three steps [8]:

1. **Planning**. The key idea in FRR is to find, for each protected link or link pair and before any failure takes place, a backup path from the PLR node to the MP node. Consider first the case of link failure, e.g., of directed link $B \rightarrow C$ in Fig. 1. In MPLS FRR, the MP node C is the next-hop (NHop) router of the PLR node B with respect to the link $B \rightarrow C$. Hence, B may select the path $B \rightarrow E \rightarrow F \rightarrow C$ on which to re-direct traffic in the event that link $B \rightarrow C$ fails.

   Let us now consider the case of node failure, e.g., of node D in Fig. 1. Suppose that node C, a neighbor of D, has traffic that passes through D on its way to node H and beyond, i.e., traverses the link pair ($C \rightarrow D$, $D \rightarrow H$). In FRR, from the point of view of protecting this link pair, C is the PLR node and H is the next-next-hop (NNHop) MP node. The NNHop scheme consists of finding a backup path from the PLR to the NNHop router; Fig. 1 shows that the path $C \rightarrow F \rightarrow G \rightarrow H$ has been selected to protect the link pair ($C \rightarrow D$, $D \rightarrow H$). Similar actions are taken by all neighbors of the failed node D to protect all link pairs through this node.
2. **Backup LSP signaling**. Backup LSPs are established along the backup paths using the same signaling mechanisms (e.g., RSVP-TE) as for setting up working LSPs. Although backup LSPs do not carry traffic under normal conditions, they are ready to accept traffic re-directed from failed working LSPs once the backup signaling is finished. This step is implemented identically for both link and node protection.
3. **Local repair**. When a link (e.g., link $B \rightarrow C$ in Fig. 1) or a node (e.g., node D in Fig. 1) fails, the failure will be detected and confirmed after several signaling actions between its PLR and MP nodes. From that instant, and until a global routing update takes effect, any packets that would have been forwarded along the failed link or node, are instead re-directed by the PLR node onto the corresponding backup path, e.g., as shown in Fig. 1. In re-directing traffic affected by the failure, the PLR uses a new label (i.e., FRR, or protection, label) in place of the former working label, so that packets be forwarded along the backup LSP. Once the packets arrive at the MP node over the backup LSP, they are forwarded toward their destination as if they had arrived over the working LSP.

Note that the second (signaling) and third (local repair) steps of this process must conform to the relevant MPLS standards, especially RFC 4090 [8]. However, the first step (planning) is outside the purview of standards, and network operators are free to employ customized algorithms to select a backup path for each protected link. In pure FRR, the backup path for each protected link or link pair (in case of node failure) is selected by the PLR node, typically using a constraint-based shortest path first (CSPF) algorithm [8, Section 6.2]. Since the PLR of a protected link/link pair executes the CSPF algorithm independently of other routers, it makes a locally optimal decision based on its own information.

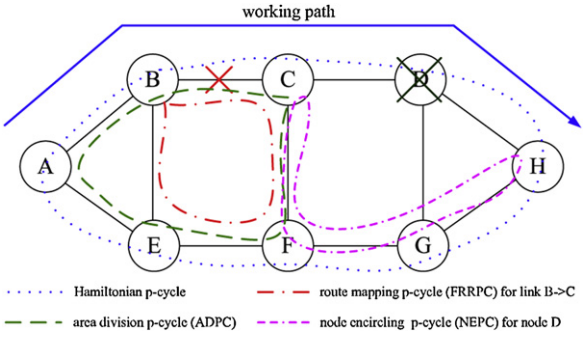

Fig. 2. Hybrid FRR/*p*-cycle protection.

However, these locally optimal backup paths may not constitute a globally optimal solution, i.e., one that optimizes a network-wide objective such as backup resource cost and/or utilization. Since the planning step may take place *offline*, it is possible to employ a more sophisticated design methodology that takes a more holistic view of network performance and cost in selecting backup paths. In the next subsection we will describe how to apply such a backup capacity design based on the *p*-cycle concept.

3. Hybrid FRR/*p*-cycle design

The *p*-cycle design is simply an alternative way of carrying out the planning step of link or node protection. In this step, the whole network may be protected by a single Hamiltonian *p*-cycle or a set of smaller *p*-cycles that may be selected according to various methodologies. As we mentioned in Section 1, there has been extensive research in developing optimization techniques for selecting optimal sets of *p*-cycles. However, general variants of the problem are NP-hard, with the computational complexity increasing quickly with the size and density of the network. As reported in [18], hop-limited *p*-cycle designs may take several hours to solve optimally even for a relatively small network.

Our goal in this work is *not* to present optimal *p*-cycle solutions, but rather, to quantify the benefits of incorporating such designs within the FRR framework. To this end, we consider four schemes for determining *p*-cycle sets that can be used to protect an MPLS network from a single link or node failure. Each method uses a fast technique to yield a "good" *p*-cycle set; taken as a whole these *p*-cycle sets are representative of the performance improvements that are achievable relative to pure FRR. To the degree that optimal *p*-cycle techniques might yield additional performance improvements, they would provide further support to our argument of using a hybrid FRR/*p*-cycle design. Nevertheless, quantifying the benefits of optimal design is outside the scope of this work.

We now discuss four schemes for selecting *p*-cycles, as illustrated in Fig. 2.

1. **Hamiltonian *p*-cycle**. The concept of Hamiltonian *p*-cycle has been known for years, and it generally works well for small size networks. As shown in Fig. 2, a single Hamiltonian *p*-cycle may be configured through all eight nodes of the network. Such a *p*-cycle may restore any single link or node failure by having the PLR reroute the affected traffic toward the MP router, along the part of the cycle that is accessible to the PLR after the failure. While a Hamiltonian represents a straightforward *p*-cycle solution for protection, the restoration time increases with the size of the network due to the long backup paths.
2. **Area division *p*-cycles (ADPC)**. Rather than using a single Hamiltonian *p*-cycle for the whole network, the key idea is to divide the network into several smaller areas and build a Hamiltonian cycle for each area; the resulting set of *p*-cycles is then used for link and/or node protection in the original network. Therefore, the areas must be determined such that collectively, they include all network links (for link protection only), all link pairs (for node protection only), or both (for link and node protection). By carefully selecting the areas, the length of the corresponding Hamiltonian cycles (and, hence, of the backup paths) can be kept well below that of a single Hamiltonian *p*-cycle, without sacrificing backup resources (as the numerical results, to be presented shortly, indicate).
3. **Node encircling *p*-cycles (NEPC)**. Node encircling *p*-cycles (NEPC) have been proposed for node protection at the IP/MPLS layer [7]. Each such cycle is designed to protect a specific network node, and includes all the neighbors of this node (but not the node itself). We observe that by selecting NEPCs appropriately, all links of the network may be covered as well. Therefore a set of NEPCs may also be used for link protection.
4. **FRR-based *p*-cycles (FRRPC)**. Observe that, taken together, the working and protection paths used by pure FRR to protect from a given link or link pair failure form a cycle. Therefore, we may build a set of *p*-cycles based on the working and backup paths constructed by FRR, referred to as FRR-based *p*-cycles (FRRPC). The main difference from pure FRR is that by organizing the paths into *p*-cycles, backup capacity may be shared among backup paths that protect different links and/or nodes. This is simply not possible for pure FRR since each node individually constructs backup tunnels to protect its adjacent links or nodes, hence there is no capacity sharing among the backup paths constructed by different nodes. Since this set of *p*-cycles uses the same routing as pure FRR, it allows us to investigate the potential improvement in backup capacity requirements that is due to simply taking a comprehensive view of protecting the whole network through *p*-cycles, rather than having each node make its own protection arrangements independently of other nodes in the network.

Note that, in the context of MPLS networks, a *p*-cycle is a *logical* entity, and its purpose is simply to define the backup path for each link or node that it protects. Consider the link $B \rightarrow C$ in Fig. 2 which happens to be an on-cycle link for the Hamiltonian and ADPC *p*-cycles shown. The backup path for this link is the path from *B* to *C* along the counter-clockwise direction on either *p*-cycle. Similarly, all on-cycle links are backed up by the (unique) reverse path to their MP node. On the other hand, there are two potential backup paths for each straddling link (e.g., link $B \rightarrow E$ with respect to the Hamiltonian), one along each direction around the *p*-cycle. Operators may use one of these two paths (e.g., the shortest one), or both. Similar observations hold for node protection.

Once the set of *p*-cycles has been selected, the second (signaling of backup LSPs) and third (local repair) steps take place exactly as the standard [8] specifies. We discuss the local repair step in more detail in Section 4.3.

4. Performance metrics and analysis

Now we discuss three performance metrics to evaluate the relative merits of the pure FRR and hybrid FRR/*p*-cycle schemes. Our goal is to protect the network from any one of three scenarios: (1) single link failure only, (2) single node failure only, or (3) either a single link or single node failure. For the pure FRR scheme, we assume that the backup path of each link or node is given; while for the hybrid scheme we assume that the *p*-cycle set is given and that each straddling link (in a link failure case) or link pair (in a node failure case) is protected by sending its traffic along the two backup paths around the *p*-cycle. Unless we explicitly specify otherwise, whenever we refer to a link *l* we assume that the link is directed. Symbols that are used in our metrics are defined in Table 1.

Table 1
Symbols used for performance metrics.

Symbols	Meanings	Symbols	Meanings
B	Total backup capacity	W	Total working capacity
W_l	Working capacity carried by any link l	W_{l_1,l_2}	Working traffic that flows on both links of pair (l_1, l_2)
W_l^s	Working capacity on the short backup path of a cycle	W_l^l	Working capacity on the long backup path of a cycle
p_l	Backup path for link l protection	p_{l_1,l_2}	Backup path for link pair (l_1, l_2) protection
d_l	Length of backup path for link l protection	d_{l_1,l_2}	Length of backup path for link pair (l_1, l_2) protection
$\mathcal{L}_n^{in}/\mathcal{L}_n^{out}$	Set of incoming/outgoing links of node n	$\mathcal{L}_c^{cw}/\mathcal{L}_c^{ccw}$	Set of clockwise/counter-clockwise links of p-cycle c
$\mathcal{L}_c^{str}$	Set of straddling links of p-cycle c	B_{frr}	Total amount of backup capacity for pure FRR
H_n	Weighted backup hop cost for node n	B_{hfrr}	Total amount of backup capacity for hybrid FRR/p-cycle
B_l	Backup capacity for on-cycle links if link l fails	B_l^{on}/B_l^{off}	Backup capacity for on-cycle/straddling (l_1, l_2) protection
d_c	Length of a p-cycle c	d_c^s/d_c^l	Length of the short/long backup path on a p-cycle c
$\overline{H}_{frr}$	Weighted backup hop cost for pure FRR	$\overline{H}_{hfrr}$	Weighted backup hop cost for hybrid FRR/p-cycle
H_l	Weighted backup hop cost for link l	H_{l_1,l_2}	Weighted backup hop cost for link pair (l_1, l_2)

In order to give better analysis, we put working capacity with different units on spans of Fig. 3. Furthermore, here we assume that only following 8 links are needed to be protected, which are in clockwise direction as $A \to B \to C \to D \to H \to G \to F \to E \to A$. Working capacity of these directed links are set as 3, 5, 1, 3, 3, 2, 5, 4 respectively. The backup paths for traffic restoration complies with the design of either FRR or hybrid FRR/p-cycle schemes, and here the planning of ADPC is indicated with three cycles of different colors.

4.1. B/W ratio

By setting up bandwidth-guaranteed backup LSPs, it is possible for the MPLS network to protect all working LSPs. The B/W ratio is an important metric to compare p-cycles' performance, which has been used in lots of literature like [7–15].

Given the routing and amount of traffic carried by each working LSP, it is straightforward to compute the total working traffic W_l carried by any link l (assuming the routing of working LSPs is independent of how backup LSPs are selected). Therefore, W in the network is equal to: $W = \sum_l W_l$.

Assuming that in each of the three scenarios we consider, the network must be protected from all link and/or all node failures, then the amount of working capacity to be protected is equal to W in each case. In the following subsections, we derive expressions for the backup capacity B that is needed for each scenario under the pure FRR and hybrid FRR/p-cycle schemes.

4.1.1. Backup capacity for pure FRR

Let us first consider protection from single link failure only (first scenario). In this case, an amount of backup capacity equal to W_l must be provisioned along each link of the backup path for link l. Let p_l be the backup path for link l, and let d_l be the length (in hops) of path p_l. Then, the total amount of backup capacity for FRR under scenario 1 is:

$$B_{frr}^1 = \sum_{l=1}^{L} d_l \times W_l, \tag{1}$$

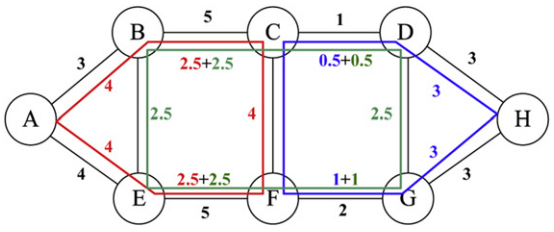

Fig. 3. Demonstration of spans' working capacity and ADPC planning.

where L is the number of (directed) links in the network. For example, in order to protect 8 selected (directed) links in Fig. 3, the total allocated backup capacity will be calculated as: $5 \times 3 + 1 \times 3 + 3 \times 2 + 3 \times 2 + 4 \times 2 + 5 \times 3 + 2 \times 3 + 3 \times 2 = 65$

Consider now scenario 2, i.e., protection from single node failures only. Whenever a node n fails, all the traffic on the links adjacent to n (in either direction) must be protected, *except* traffic that originates or terminates at node n. Let $\mathcal{L}_n^{in}$ and $\mathcal{L}_n^{out}$ denote the set of incoming and outgoing links, respectively, of node n. Let (l_1, l_2) be a pair of (directed) links passing through n, i.e., $l_1 \in \mathcal{L}_n^{in}$ and $l_2 \in \mathcal{L}_n^{out}$. Let p_{l_1,l_2} be the backup path that is used to bypass this link pair in the event that n fails, and d_{l_1,l_2} be the length of this path. Let W_{l_1,l_2} denotes the amount of working traffic that flows on both links of pair (l_1, l_2), and must be protected if n fails; in other words, this is the traffic that travels from link l_1 to link l_2 through node n, not including any traffic that terminates at, or originates from, n. The total amount of backup capacity required for scenario 2 can then be obtained as:

$$B_{frr}^2 = \sum_{n=1}^{N} \left(\sum_{l_1 \in \mathcal{L}_n^{in}, l_2 \in \mathcal{L}_n^{out}} d_{l_1,l_2} \times W_{l_1,l_2} \right) \tag{2}$$

where N is the total number of nodes in the network. For instance, if the working capacity of link pair $A \to B \to C$ is 3, and $E \to B \to C$ is 2, totally $3 \times 3 + 2 \times 2 = 13$ unit will be cost to protect node B. As for other nodes, the principle is similar.

For scenario 3 (i.e., protection from either a single link or a single node failure), both the NHop and NNHop schemes must be activated independently under pure FRR. Hence, the total amount of protection bandwidth required in this case is:

$$B_{frr}^3 = B_{frr}^1 + B_{frr}^2. \tag{3}$$

4.1.2. Backup capacity for hybrid FRR/p-cycle

For link protection, let us assume that C, $C \geq 1$, p-cycles have been configured for protecting the network links. Note that, if a link l is a straddling link in some p-cycle c, then an amount of backup capacity equal to $W_l/2$ on the on-cycle links (in both directions) is sufficient to protect this link. On the other hand, if a link l is an on-cycle link, then W_l units of backup capacity need to be provisioned on all other links of the p-cycle *in the opposite direction*. However, if a link is an on-cycle link of k different p-cycles, each p-cycle needs to provision only W_l/k units of backup capacity. Let $\mathcal{L}_c^{cw}$ (respectively, $\mathcal{L}_c^{ccw}$) denotes the set of clockwise (respectively, counter-clockwise) links of p-cycle c, and $\mathcal{L}_c^{str}$ denote the set of straddling links of p-cycle c. Based on these observations, for scenario 1, protection from

a single link failure only, the backup capacity for the clockwise on-cycle links of p-cycle c is given by:

$$B_l^1 = \max\left\{\max_{l'\in\mathcal{L}_c^{ccw}}\left\{\frac{W_{l'}}{k_{l'}}\right\}, \max_{l'\in\mathcal{L}_c^{str}}\left\{\frac{W_{l'}}{2}\right\}\right\}, \quad l\in\mathcal{L}_c^{cw}. \tag{4}$$

A similar expression can be written for the backup capacity of counter-clockwise links, while straddling links need no backup capacity, i.e., $B_l^1 = 0$, $l\in\mathcal{L}_{str}$. We also note that if a link belongs to multiple p-cycles, the backup capacity that is reserved on this link is the maximum capacity assigned by any p-cycle c from the corresponding expression (4). According to the Fig. 3, since the biggest value of each link's working capacity is 5, for a Hamiltonian p-cycle, it only costs $8\times5=40$ unit to protect any single link failure. If we use ADPC or other p-cycle design, by putting area division like Fig. 3 and calculating the protection capacity in each smaller "*Hamiltonian*" p-cycle, $4\times5+3\times5+2.5\times6=50$ will be used for the whole link protection.

For scenario 2 (protection from single node failure only), we distinguish between two cases. If a node n on a p-cycle fails, the p-cycle can be used to to protect a link pair (l_1, l_2) through node n if (1) both links are on-cycle links, (2) one link is an on-cycle link and the other one is a straddling link on the p-cycle, or (3) both links are straddling links of the p-cycle. In each of these cases, there is only one backup path around the p-cycle to restore affected traffic, and the amount of capacity on each link of the backup path is equal to: $B_l^{on} = \max\{W_{l_1,l_2}\}$, where the maximum operation is over all link pairs (l_1, l_2) for which on-cycle link l is on their backup path. For example, if node B failure happens in Fig. 3, link pair $A\to B\to C$ will be protected by the other part of a Hamiltonian p-cycle going from A through C, assuming traffic on $A\to B\to C$ is 3 unit, it will cost $6\times3=18$ unit for this link pair protection. However, as for ADPC, it only uses $3\times3=9$ unit to restore the failure.

If a node n not on the p-cycle fails, the p-cycle can be used to protect a link pair (l_1, l_2) if the path (l_1, l_2) is a straddling path of the p-cycle (e.g., as in node-encircling p-cycles). In this case, an amount of capacity equal to $B_l^{off} = \max\{W_{l_1,l_2}/2\}$ must be reserved on each on-cycle link l, where the maximum operation is taken over all link pairs (l_1, l_2) that are straddling paths of the p-cycle. Hence, under scenario 2, the backup capacity for the on-cycle links of p-cycle c is given by:

$$B_l^2 = \max\left\{B_l^{on}, B_l^{off}\right\}, \quad l\in\mathcal{L}_c^{cw}\cup\mathcal{L}_c^{ccw}. \tag{5}$$

As in scenario 1, straddling links need no backup capacity.

For scenario 3, i.e., protection from either single link or single node failure, the backup capacity on each link must be equal to the maximum of the capacities required for scenarios 1 and 2, i.e.,

$$B_l^3 = \max\{B_l^1, B_l^2\}. \tag{6}$$

Consequently, the total backup capacity for hybrid FRR/p-cycle under scenario i, $i=1,2,3$, can be computed as:

$$B_{hfrr}^i = \sum_{l=1}^{L} B_l^i, \quad i=1,2,3. \tag{7}$$

4.2. Traffic weighted backup hop cost

When a link or a node fails, all traffic on the affected link(s) is re-directed along the backup path(s), incurring additional delay that depends on the length of backup paths. Here traffic weighted backup hop cost is used for measuring that in order to protect one link or node of the whole network, how many additional resources are needed in average. This is a new metric that first introduced by this paper.

4.2.1. Pure FRR

Let d_l denotes the length (in hops) of path p_l that serves as the backup path of link l. If link l, carrying an amount W_l of working traffic fails, the traffic weighted backup hop cost incurred by link l LSPs is given by: $H_l^1 = W_l\times d_l$. Assuming that all L links are equally likely to fail, the average traffic weighted backup hop cost for pure FRR under scenario 1 can be written as:

$$\overline{H}_{frr}^1 = \frac{\sum_{l=1}^{L} H_l^1}{L} \tag{8}$$

Since we have worked out that it will add totally 65 unit to protect all of single link failure in Section 4.1.1, for every link, the weighted backup hop cost is $65/8=8.125$. In scenario 2, whenever a node n fails, traffic on all link pairs through this node is re-directed on the corresponding backup paths. Using our earlier notation, the traffic weighted backup hop cost for node n is:

$$H_n^2 = \sum_{l_1\in\mathcal{L}_n^{in}, l_2\in\mathcal{L}_n^{out}} W_{l_1,l_2}\times d_{l_1,l_2}. \tag{9}$$

Consequently, the average cost over all node failures is:

$$\overline{H}_{frr}^2 = \frac{\sum_{n=1}^{N} H_n^2}{N} \tag{10}$$

As a result, for pure FRR case, we can get the average cost over all node failures by first adding the amount of every node failure cost together, and then divide this value by the number of nodes.

For scenario 3, if we assume that all link or node failures are equally likely, then the traffic weighted backup cost can be given as:

$$\overline{H}_{frr}^3 = \frac{\sum_{l=1}^{L} H_l^1 + \sum_{n=1}^{N} H_n^2}{L+N} \tag{11}$$

If node and/or link failures have different probability of occurring, the above expression can be modified in a straightforward manner. Since our objective is to investigate the *relative* costs of the pure and hybrid FRR schemes, we will use the above expression for simplicity.

4.2.2. Hybrid FRR/p-cycle

For the FRR/p-cycle hybrid scheme, again assume that C p-cycles have been configured, and let $d_c\geq3$ denote the length (i.e., number of directed on-cycle links) of p-cycle c, $c=1,\ldots,C$. Consider, first, scenario 1. If link l is an on-cycle link for k_l p-cycles, the traffic weighted backup hop cost for this link is:

$$H_l^1 = \sum_{j=1}^{k_l} \frac{W_l}{k_l}\times(d_c-1). \tag{12}$$

For a link l that is a straddling link on p-cycle c, let d_c^s and d_c^l denote the length of the short and long backup paths, respectively, for the link along the p-cycle, i.e., such that $d_c^s\leq d_c^l$ and $d_c^s+d_c^l=d_c$. We send as much working traffic W_l^s as possible on the short backup path, i.e., $W_l^s=\min\{W_l, B_c\}$, where B_c is the spare capacity on the on-cycle links of the p-cycle, and the remaining traffic, $W_l^l=W_l-W_l^s$, if any, on the long backup path. Hence, the weighted cost is:

$$H_l^1 = W_l^s\times d_c^s + W_l^l\times d_c^l. \tag{13}$$

The average traffic weighted backup cost, $\overline{H}_{hfrr}^1$, can be obtained by an expression similar to (8). According to the Fig. 3, since 3 unit is needed to go through all other spans to protect link $A\to B$, 5 unit to protect link $B\to C$, similarly, totally $(5+5+1+3+3+3+4+1+2)\times(8-1)=189$ unit will be reserved for every link failure. Consequently, the weighted backup cost is $189/8=23.625$.

As for ADPC scheme, it will first calculate three cycles as: $(3+4+2.5+2.5)\times(5-1)+(0.5+1+3+3)\times(5-1)+(2.5+2.5+0.5+1)\times(6-1)=110.4$.

So the weighted backup cost is 110.4/8 =13.8.

For scenario 2, if the link pair (l_1, l_2) failed due to the failure of a node on a p-cycle, then

$$H_{l_1,l_2} = W_{l_1,l_2} \times d_{l_1,l_2}, \tag{14}$$

where d_{l_1,l_2} is the length of the backup path for this link pair on the p-cycle. If the link pair (l_1, l_2) is a straddling path in a p-cycle, then

$$H_{l_1,l_2} = W^s_{l_1,l_2} \times d^s_c + W^l_{l_1,l_2} \times d^l_c. \tag{15}$$

The calculation is quite like we have demonstrated for node B with pure FRR scheme. The average traffic weighted backup cost, $\overline{H}^2_{hfrr}$, can be obtained by expressions similar to (9) and (10).

Finally, the average traffic weighted backup hop cost for scenario 3, $\overline{H}^3_{hfrr}$, can also be obtained by using an expression similar to (11).

4.3. Label entry overhead

The number of labels required to establish backup paths is an important metric for MPLS networks, as it determines the size of the forwarding tables at the LSRs. The metric of label entry overhead has been mentioned in some references like [1,2], and here it is first employed to compare the performance of MPLS protection. We assume that the *one-to-one backup* method [8] is used to implement the local repair technique. This method requires the allocation of a different set of protection labels for every traffic component (i.e., LSP). For example, if link l fails, all traffic on the link is sent along the backup path by having the PLR node *switch* on each affected packet the protection label associated with the LSP of the packet. Each intermediate node on the backup path forwards packets based on their protection label, and replaces it with a new protection label, as per the normal MPLS packet forwarding operation. When the MP node receives a packet with a protection label, it replaces it with a new working label and forwards the packet along its original (working) path. In case of node failure, the operation for every node along the backup path (including the PLR and MP nodes) is identical. Therefore, the number of additional labels required to protect a link l or a link pair (l_1, l_2) with the one-to-one backup method is equal to the number of hops along the corresponding backup path *times* the number of traffic components (LSPs) that traverse this link or link pair, respectively.

Since the backup paths in the hybrid FRR/p-cycle scheme are embedded into p-cycles that are determined in advance, it is possible to use a smaller number of protection labels. Consider first the case of link failure, and observe that the backup path always traverses all the links of the p-cycle (in the opposite direction of the working link that failed). Hence, we assign only *two sets* of protection labels for each p-cycle, one set in each direction (clockwise or counter-clockwise). If a link fails, for each affected packet, the PLR node (1) switches the incoming working label to the appropriate outgoing working label (as in normal operation), (2) *pushes* the same outer protection label onto all packets, and (3) forwards all such traffic along the backup path on the appropriate p-cycle. Intermediate nodes on the backup path simply switch the appropriate protection label assigned for the p-cycle and direction of the backup path. When the MP node at the other end of the backup path receives a packet with a protection label, it *pops* this label and continues to forward the packet based on the inner working label assigned by the PLR node. Therefore, for link protection, the number of labels required for all the backup paths on a p-cycle is simply twice the number of links in the p-cycle; i.e., one set of labels for each direction along the p-cycle. This arrangement is possible because each node on the p-cycle may *reuse* the same set of labels to accommodate *any* (on-cycle or straddling) link failure without ambiguity: under any failure, only the MP node of the failed link is aware of the failure and is the one to remove traffic redirected due to the failure on the bypass tunnel from the p-cycle.

In the case of node failure, backup paths for the various protected link pairs also follow a p-cycle. However, backup paths for different link pairs affected by the failure may terminate at different MP nodes; more importantly, even if backup paths terminate at the same MP node, the traffic on these backup paths may have to take different routes after reaching the MP node, depending on the specific link pair that each backup path protects. Therefore, in addition to the two sets of protection labels that are associated with each direction of the p-cycle (as in the link failure case), we introduce one additional protection label for each link pair protected by the p-cycle. When a node fails, for each affected packet, the PLR node: (1) switches the incoming working label to a protection label associated with the protected link pair that the packet would have traversed under normal operation, (2) *pushes* the same outer protection label onto all packets, and (3) forwards all such traffic along the backup path along the p-cycle. Intermediate nodes simply switch the appropriate outer protection label, forwarding the packet along the p-cycle toward the MP node. When the MP node receives a packet with a protection label, it *pops* this label and examines the inner label. If the inner label is also a protection label (which necessarily corresponds to a specific link pair), it forwards the packet onto the appropriate working path after first switching this inner label with the corresponding working label. If the inner label is not a protection label, then it must be a working label corresponding to a link failure, and the MP node proceeds as we described above. With this arrangement, the number of labels required for each p-cycle is twice the length of the cycle (as in the link failure case), plus the number, say, K of link pairs protected by the p-cycle.

5. Numerical results

We compare the pure FRR to the hybrid FRR/p-cycle schemes on a simulation testbed implemented using the OPNET modeler. For this performance study, we consider the three real network topologies shown in Fig. 4 that have been widely used in survivability research [18,19]. The Cost-239 (N = 11 nodes, L = 52 directed links) topology illustrates a relatively dense network connecting 11 main cities in Europe, with an average node in-/out-degree $\overline{D} = 4.73$, while the Havana topology (N = 17, L = 52) demonstrates a relatively sparse network deployed in Germany, with $\overline{D} = 3.06$); The USA-20 topology (N = 20 nodes, L = 92 directed links), shown in the right part of the figure, and has an average node in-/out-degree $\overline{D} = 4.60$.

Note that here we only evaluate the performance of planning by employing different virtual network design schemes, which can be simulated with a certain number of unit as traffic demands or span bandwidth. As for the physical networks, other two steps like backup LSP signaling and local repair are also very important. They are suggested to comply with protocols of RSVP-TE [20] and RFC 4090 [8], and each of MPLS router in the network must implement these protocols. It will usually cost hundreds of millisecond to set up a new protection path.

Traffic demands are set up between every pair of nodes in each network, and working traffic is routed along shortest paths computed using Dijkstra's algorithm. Let t_{sd} denotes the amount of working traffic carried by the LSP from s to d. To investigate the sensitivity of the relative performance of the pure FRR and the four hybrid schemes, we generated working traffic demands that follow four different patterns:

- *Equal (EQ)*: $t_{sd} = t$ = constant, $\forall(s, d)$.

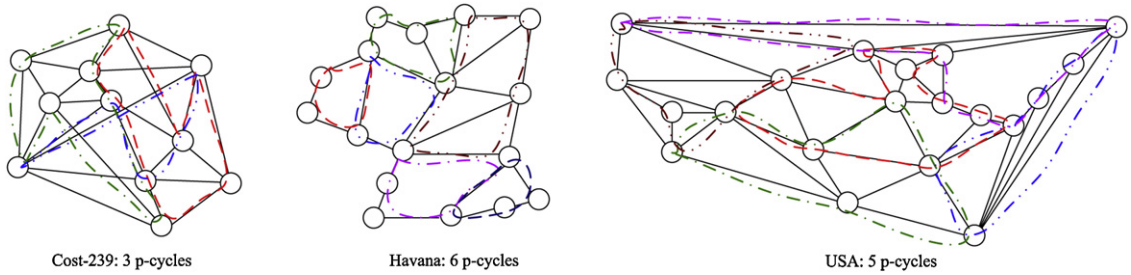

Fig. 4. Network topologies used in the performance study.

- *Uniform (UF)*: t_{sd} is uniformly distributed in the interval [0, 20], $\forall(s, d)$.
- *Locality (LC)*: t_{sd} is uniformly distributed in the interval $[4(h - h_{sd}), 4(h - h_{sd} + 1) - 1]$, where h_{sd} is the length (in hops) of the shortest path between s and d and h is the length of the longest shortest path in the network; in this pattern, the traffic demand between each node pair decreases with the distance between the two nodes, and models the traffic locality observed in some networks.
- *Reverse locality (RL)*: t_{sd} is uniformly distributed in the interval $[4(h_{sd} - 1), 4h_{sd} - 1]$, where h_{sd} is the length (in hops) of the shortest path between s and d, hence, it increases with the length of the shortest path h_{sd}.

For pure FRR, Dijkstra's algorithm was also used to find the shortest backup path for each link l or link pair (l_1, l_2). For the hybrid FRR/p-cycle scheme, we use the four types of p-cycles that discussed in Section 3 for planning the backup paths. In particular, recall that the ADPC scheme is based on subdividing each network into smaller areas, and selecting a Hamiltonian cycle in each area. The various dotted lines in Fig. 4 indicate the smaller areas we used in each topology to select p-cycles for the ADPC scheme.

5.1. Results and discussion

The results of the simulation for the Cost-239, Havana and USA topologies are presented in the two sub-figures of Figs. 5 and 6, respectively. Each sub-figure corresponds to one of the three scenarios we consider, i.e., single link protection only, single node protection only, or both single link and node protection. The sub-figures plot the B/W ratio or the backup hop cost as a function of the traffic pattern (shown on the x axis). In order to make meaningful comparisons, although working traffic demands were generated according to the four traffic patterns described above, the total working traffic in each case was set to 1000 units. In addition, these traffic patterns are plotted with different suffixes as -C, -H and -U, which means the results are given for topologies of Cost-239, Havana and USA respectively. Each sub-figure contains five curves: one for the pure FRR scheme, and four for the hybrid FRR/p-cycle scheme corresponding to the four techniques for selecting p-cycles (as discussed in Section 3). Since Havana is quite a sparse topology, few simple cycles could be found around each node to form NEPCs. However, NEPCs with non-simple cycles are not recommended for widely use, because it will bring a lot of operation confusion and signaling complexity [16]. As a result, here we only employ other three hybrid schemes for protecting the network of Havana.

We first observe that under scenario 1 (link protection only), the costs associated with protection, i.e., the B/W ratio and backup hop cost, are higher than under scenario 2 (node protection only). This result can be explained by the fact that, when a link fails, all the traffic on the link must be re-directed onto backup paths. On the other hand, when a node fails, traffic originating or terminating at the node cannot be protected. Hence, over all possible link failures, the amount of backup resources required is higher than over all

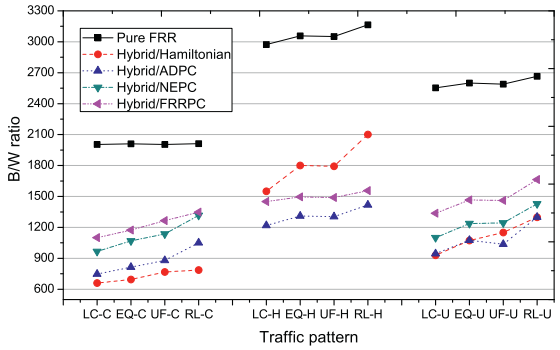

(a) Scenario 1: link protection

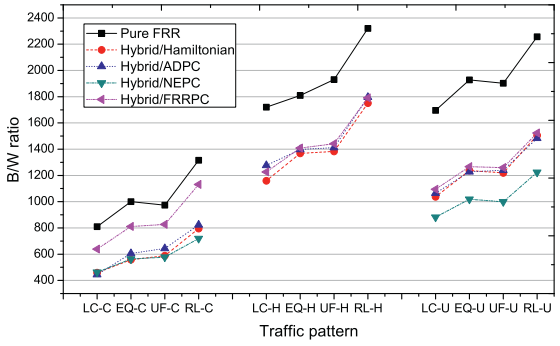

(b) Scenario 2: node protection

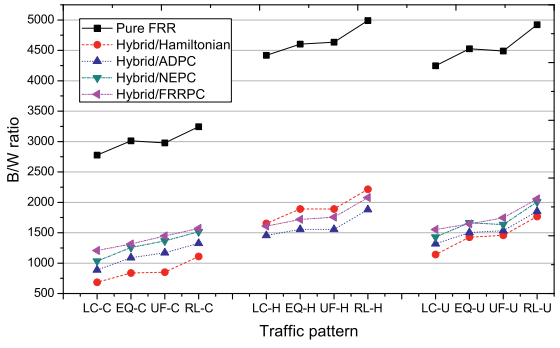

(c) Scenario 3: link and node protection

Fig. 5. Simulation results for the B/W ratio (a): scenario 1: link protection; (b) scenario 2: node protection; (c) scenario 3: link and node protection.

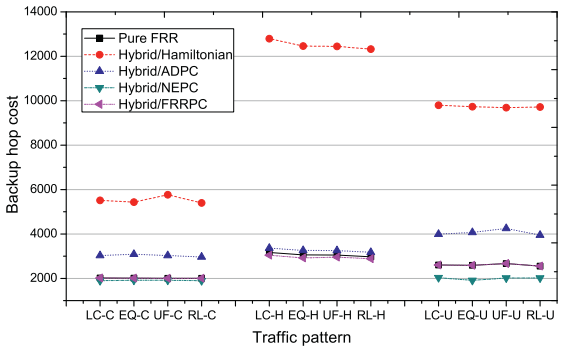

(a) Scenario 1: link protection

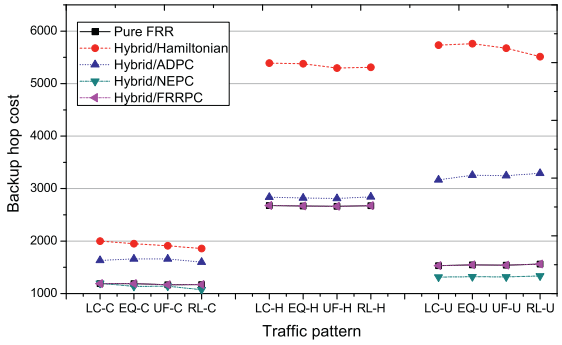

(b) Scenario 2: node protection

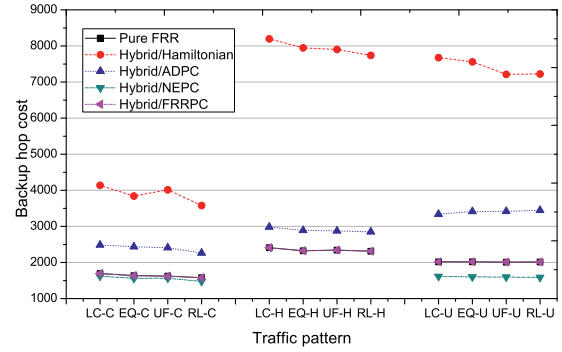

(c) Scenario 3: link and node protection

Fig. 6. Simulation results for the backup hop cost: (a) scenario 1: link protection; (b) scenario 2: node protection; (c) scenario 3: link and node protection.

possible node failures. A similar observation holds for the traffic-weighted backup hop cost. Of course, under scenario 3 (protection of either single link or single node failures), the B/W ratio associated with protection is the highest among all scenarios.

5.2. B/W ratio

From the two figures, it is evident that, under all three protection scenarios and all four traffic patterns, the B/W ratio for the pure FRR scheme is higher than that of the hybrid schemes. With pure FRR, there is no sharing of backup resources as its node plans its own backup paths independently of other nodes. The hybrid designs, on the other hand, are able to share protection bandwidth along the p-cycles. Moreover, in relatively dense topologies such as Cost-239 and USA, there are ample opportunities for links or link pairs to be straddling spans on some p-cycles, hence increasing the protection efficiency (since straddling spans do not need to have any spare capacity). These results illustrate the benefits of taking a comprehensive view of the network and traffic demands in designing backup paths using p-cycles.

Let us now turn our attention to the relative performance of the four hybrid schemes. The results indicate that a Hamiltonian p-cycle is not always the best option in terms of bandwidth efficiency. In fact, the ADPC, FRRPC and NEPC schemes tend to have better performance (even the lowest B/W ratio)across all topologies and the traffic patterns we considered in our study. In particular, the NEPC scheme is always the best for scenario 2, as it is designed specifically for protecting each node. Also note that FRRPC uses the same paths as FRR (but arranged in p-cycles such that backup bandwidth is shared). The fact that FRRPC has generally a much lower B/W ratio than pure FRR is an indication of arranging backup paths so as to share protection resources.

We also note that on one hand, the traffic pattern does affect the B/W ratio, especially for hybrid schemes, but the relative performance among the various schemes is similar. Specifically, the locality (LC) pattern results in the lowest amount of protection capacity: since the majority of traffic is between nodes close to each other in distance, the corresponding backup paths are relatively short, resulting in low overall spare capacity. Similar arguments can be used to explain why the reverse locality (RL) pattern requires the highest amount of backup capacity among the four patterns considered here, while the equal (EQ) and uniform (UF) patterns fall between the other two in terms of this metric. On the other hand, for the topology of Havana, the bandwidth resource for protection is even higher than the topology with bigger scale but more dense, i.e., the USA topology. That is because it is a sparse topology, less spans can be used to protect multiple links/nodes, and more exclusive bandwidth are reserved for each span to protect different links/nodes.

We further note that the various p-cycle sets we consider here were not optimized for any specific objective. Hence, the results of B/W ratio are only an upper bound on what can be achieved using p-cycle design; using sophisticated optimization techniques to select the p-cycle sets, additional improvements in capacity efficiency would be possible.

5.3. Traffic weighted backup hop cost

The results show that using a single Hamiltonian path incurs high backup hop cost, due to the long backup paths involved. The ADPC scheme reduces this cost significantly by employing a set of smaller p-cycles. The FRRPC and NEPC schemes use even smaller cycles, reducing this cost further to the level of pure FRR. We also observe that the effect of the traffic pattern on the results is relatively small for the FRRPC and NEPC that use short cycles. However, the backup hop cost more directly depends on the traffic pattern for the Hamiltonian and ADPC schemes that employ longer cycles, especially for scenarios 2 and 3. The less connected network like Havana also has the higher hop cost than Cost-239 and is even higher than USA in scenarios of 1 and 3, because it needs to put more bandwidth for each span, and the backup paths are often more than 3 hops.

5.4. Label entry overhead

Table 2 compares the protection schemes in terms of the number of additional labels needed under the three protection scenarios we consider. For pure FRR, each link or node is protected independently of others by establishing a separate protection LSP per traffic component. Hence, the number of labels is proportional to the number of traffic components and the total length of all backup paths

Table 2
Label entry overhead comparison.

Topology	FRR	Hybrid FRR/*p*-cycle			
		Hamiltonian	ADPC	NEPC	FRRPC
Scenario 1: link protection					
Cost-239	572	22	42	94	152
Havana	896	34	56	Null	198
USA-20	3625	40	80	240	312
Scenario 2: node protection					
Cost-239	270	106	154	178	144
Havana	452	150	206	Null	178
USA-20	2568	252	438	452	426
Scenario 3: link and node protection					
Cost-239	842	106	154	178	296
Havana	1348	150	206	Null	386
USA-20	6193	252	438	452	738

in the network. As a result, the number of protection label entries for the USA topology is much higher than in Hanava topology, and Havana is also higher than Cost-239 topology. In addition, in order to protect both link and node, pure FRR needs to deploy two different mechanisms. Therefore, the total number of label entries for scenario 3 is the sum of the label entries required under scenarios 1 and 2.

For the hybrid FRR/*p*-cycle schemes, all affected traffic is forwarded along on-cycle backup tunnels built for each *p*-cycle that can be realized with only a few labels. Furthermore, under scenario 3, the Hamiltonian, ADPC and NEPC schemes may (re-)use protection labels for both link and node protection. As a result, the label overhead is significantly lower in the hybrid scheme. This is further demonstration of the fact that, by taking a global design approach in protecting the network links or nodes, the *p*-cycle scheme is more efficient in its use of network resources.

6. Conclusions

We have proposed several hybrid FRR/*p*-cycle schemes to implement the local repair method defined for MPLS networks. These schemes all use backup paths along a set of pre-configured *p*-cycles that may be selected using design methodologies that consider the overall network performance, but otherwise are RFC 4090-compliant. Numerical results indicate that using a set of relatively short *p*-cycles outperforms pure FRR in terms of backup capacity and label overhead, and is comparable to pure FRR in terms of backup hop cost. These benefits can be realized for both link and node protection, and become more significant as the network size grows. Our main conclusion is that *p*-cycle designs are an attractive alternative for MPLS network operators.

References

[1] Jorge L, Gomes T. The coming of age of MPLS. IEEE Commun Mag 2011;49(April (4)):78–81.
[2] Alshaer H, Elmirghani J. Multilayer dynamic traffic grooming with constrained differentiated resilience in IP/MPLS-over-WDM networks. IEEE Trans Netw Serv Manage 2012;9(March (1)):60–72.
[3] Cao C, Rouskas G. Hybrid FRR/p-cycle MPLS link protection design. In: Proceedings of GLOBECOM 2011. December 2011. p. 1–6.
[4] Bigos W, Cousin B, Gosselin S, Le Foll M, nakajima H. Survivable MPLS over optical transport networks: cost and resource usage analysis. IEEE Select Areas Commun 2007;25(June (6)):949–62.
[5] Ruiz M, Pedrola O, Velasco L, Careglio D, Fernadez-Palacios J, Junyent G. Survivable IP/MPLS-over-WSON multilayer network optimization. IEEE/OSA J Opt Commun Netw 2011;3(August (8)):629–40.
[6] Liu M, Tornatore M, Mukherjee B. New strategies for connection protection in mixed-line-rate optical WDM networks. IEEE/OSA J Opt Commun Netw 2012;3(March (9)):641–50.
[7] Stamatelakis D, Grover W. IP layer restoration and network planning based on virtual protection cycles. IEEE Select Areas Commun 2000;18(October (10)):1938–49.
[8] Pan P, Swallow G, Atlas A. Fast reroute extensions to RSVP-TE for LSP tunnels. RFC 4090; May 2005.
[9] Yadav R, Yadav R, Singh H. Enhanced intercycle switching in p-cycle survivability for WDM networks. IEEE/OSA J Opt Commun Netw 2010;2(November (11)):961–6.
[10] Kiaei M, Assi C, Jaumard B. A survey of the *p*-cycle protection method). IEEE Commun Surv Tutor 2009;11(3rd Quarter (3)):53–70.
[11] Keng J, Reed M. Bandwidth protection in MPLS networks using *p*-cycle structure. In: Proceedings of DRCN 2003. 2003 October. p. 356–62.
[12] Keng J, Reed M. Network protection for mesh networks: network coding-based protection using *p*-cycles. IEEE/ACM Trans Netw 2010;18(February (1)):67–80.
[13] Schupke DA. Automatic protection switching for p-cycles in WDM networks. Opt Switch Netw (OSN) 2005;2(May (1)):35–48.
[14] Gangxiang S, Grover WD. Extending the *p*-cycle concept to path segment protection for span and node failure recovery. IEEE Select Areas Commun 2003;21(October (8)):1306–19.
[15] Kodian A, Grover WD. Failure-independent path-protecting p-cycles: efficient and simple fully preconnected optical-path protection. J Lightwave Technol 2005;23(Octorber (10)):3241–59.
[16] Onguetou D, Grover WD. A new insight and approach to node failure protection with ordinary p-cycles. In: Proceedings of IEEE international conference on communications. 2008. p. 5145–9.
[17] Jaumard B, Li H. Minimum CAPEX design of segment p-cycles with full node protection. In: Proceedings of 2012 16th international conference on Optical Network Design and Modeling (ONDM). April 2012. p. 1–6.
[18] Jaumard B, Li H. Segment p-cycle design with full node protection in WDM mesh networks. In: Proceedings of 2011 18th IEEE workshop on Local Metropolitan Area Networks. October 2011. p. 1–6.
[19] Li L, Buddhikot M, Chekur C. Routing bandwidth guaranteed paths with local restoration in label switched networks. IEEE Select Areas Commun 2005;23(February (2)):437–49.
[20] Farrel A, Ayyangar A, Vasseur JP. Inter-domain MPLS and GMPLS traffic engineering – resource reservation protocol-traffic engineering (RSVP-TE) extensions. RFC 5151; February 2008.