



HAL
open science

IP Fast Reroute in Networks with Shared Risk Links

Yan Li, Mohamed G. Gouda

► **To cite this version:**

Yan Li, Mohamed G. Gouda. IP Fast Reroute in Networks with Shared Risk Links. 9th International IFIP TC 6 Networking Conference (NETWORKING), May 2010, Chennai, India. pp.213-226, 10.1007/978-3-642-12963-6_17. hal-01056311

HAL Id: hal-01056311

<https://inria.hal.science/hal-01056311v1>

Submitted on 18 Aug 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

IP Fast Reroute in Networks with Shared Risk Links

Yan Li¹ and Mohamed G. Gouda²

¹ Department of Computer Science
The University of Texas at Austin
1 University Station (C0500)
Austin, TX 78712-0233

² The National Science Foundation, and
The University of Texas at Austin
1 University Station (C0500)
Austin, TX 78712-0233
{yanli, gouda}@cs.utexas.edu

Abstract. IP fast reroute is a mechanism that is used to reroute packets around a failed link as soon as the link fails. Most of the IP fast reroute mechanisms, that have been proposed so far, focus on single or dual link failures but can not handle Shared Risk Link Group (SRLG) failures when several links fail at the same time because of some common underlying component failure. Furthermore, most of current work is based on the assumption that each node in the network has access to some global topology information of the network. In this paper, we present the first IP fast reroute mechanism for SRLG failures that is not based on the assumption that the nodes in the network have global topology information of the network. In our mechanism, nodes in the network use “relay bits” to identify themselves as “relay nodes” for a reroute link in a fully distributed manner. Through simulation, we show that our mechanism succeeds in rerouting around SRLG failures almost 100% of the time, with average length of a reroute path about 1.5 times the re-converged shortest path.

keywords. IP Fast Reroute, Shared Risk Link Group Failure, Distance Vector, Failure Recovery, Reliability

1 Introduction

Demands on reliability and availability of the Internet are becoming more and more stringent, especially with the development of more real-time applications like VoIP and Video on demand [5]. Unfortunately, failures are very common in daily operations of a network and what makes things worse is that most failures are not predictable. It is reported in [22] that 80% of all failures are unexpected. Among these unexpected failures, besides the most common single link failures, another significant part is Shared Risk Link Group (SRLG) failures. Links that

belong to the same SRLG share some underlying component either in the optical infrastructure like a fiber or at a router like a line card.

The convergence process for failure recovery in traditional routing protocols, link state and distance vector, is time consuming and may result in instability in case of frequent transient link failures. Although much work has been dedicated to reduce the convergence time of routing to even under a second [12], it is still quite far from the 50 milliseconds target for mission critical applications [25]. Recently, IP fast reroute [10, 13, 26] has been proposed to proactively compute backup paths before a failure happens. And as soon as a failure is detected, the backup path can be invoked immediately to reroute around the failure during the convergence period. Thus the routing disruption time can be limited to only the failure detection time. Although several mechanisms have been developed within the IP fast reroute realm, most of them focus on single or dual link failures and can not handle SRLG failures [3, 4, 6, 14, 16, 20, 23, 24]. Also, most of existing work relies on the existence of some global topology information to precompute backup paths [3, 4, 6, 8, 14, 16–18, 23, 24, 27].

When global topology information is not available (for example, in distance vector like routing protocols), to recover from SRLG failures, IP fast reroute faces more challenges: how to get necessary information to compute alternative backup paths to avoid all the links in the same SRLG, without changing the original routing tables? To address this challenge, we design an IP fast reroute mechanism for handling SRLG failures with the following goals:

- ***No Global Topology Information:*** Our mechanism does not assume nodes in the network have access to any global topology information: neither the connectivity information of the network nor any additional IP addresses associated with each node. Each node only has access to its local topology, i.e., links associated with the node itself. Also, each node only knows to which SRLG its associated links belong but does not know other links in the same SRLG.
- ***Distributed Computation of Backup Paths:*** Each backup path for a reroute link is designated by a node called *relay node*. In order to find relay nodes in a fully distributed way, we introduce two *relay bits* for each reroute link. Using the two relay bits, a node in the network can automatically decide if itself can serve as a relay node for a reroute link or not, in a distributed manner.
- ***Reroute Only When You Want:*** Rerouting information is propagated only for links that are currently under protection using IP fast reroute, which are called *reroute links*. Reroute links can be changed at any time. So the cost of our mechanism is dynamically associated with the number of links currently protected in the network.

We propose a tunneling scheme to ensure that loops are never formed. We also propose an algorithm to suppress redundant relay notification messages. Finally we show, through extensive simulations on a variety of networks of different sizes and varying SRLG size, that the coverage of our mechanism is close to 100%. Suppression can effectively cut about 80% notification messages when the

network has at least one hundred nodes. Also, the average length of a reroute path is around 1.5 times the average length of the re-converged shortest path.

The structure and organization of this paper follows from our technical report [20] which focuses on how to handle single link failures using IP fast reroute. Section 2 presents the concept of reroute links. Then in Section 3, we introduce the concept of a relay node for a reroute link that is a member of a SRLG and how to use the relay node in rerouting. Section 4 presents how a node learns a relay node, without access to the global topology information, for a reroute link which belongs to a SRLG. Section 5 describes the suppression mode to suppress redundant notification messages. We show the efficiency and overhead of our rerouting algorithms in Section 6. Related work is reviewed in Section 7. Finally we conclude in Section 8.

2 Reroute Links in Shared Risk Link Group

We model a network as an undirected graph where each node represents a router and each (undirected) edge between two nodes represents two links. For each link $a \rightarrow b$, node a is called *the source* of link $a \rightarrow b$ and node b is called *the sink* of link $a \rightarrow b$. We assume that packets are routed between different nodes in the network using distance vector routing protocols [2, 21].

A Shared Risk Link Group (SRLG) is a set of links that share the same underlying physical point of failure such as a fiber cut or a line card failure. We assume that some links in the network are partitioned into SRLGs with unique identifiers. Thus, each link belongs to at most one SRLG. We also assume that the source node a of any link $a \rightarrow b$ in the network knows the identifier of the SRLG to which link $a \rightarrow b$ belongs, if any. (But node a does not need to know the other links in the same SRLG to which link $a \rightarrow b$ belongs, if any.)




Fig. 1. An example network N_1

Table 1. Routing table $RT.a$ of node a

dest.	next hop	dist.
a	-	0
b	b	2
c	b	3
d	b	5
e	e	1
f	e	2
g	e	2
h	e	3
i	e	4

As an example, Figure 1 shows a network N_1 in which each edge is labeled with a distance. The routing table $RT.a$ of node a in network N_1 is shown in Table 1. Also, there are three SRLGs in N_1 with SRLG number g_1 , g_2 and g_3 . Link $a \rightarrow b$, $c \rightarrow d$ and $e \rightarrow f$ belong to SRLG g_1 ; link $b \rightarrow f$ and $f \rightarrow g$ belong to SRLG g_2 ; link $g \rightarrow h$ and $h \rightarrow i$ belong to SRLG g_3 .

Now consider the situation where node a has a packet whose ultimate destination is node d . But then node a notices that link $a \rightarrow b$ used to reach destination

d has failed. The question now is “what does node a do with this packet that need to be transmitted over the failed link $a \rightarrow b$ before the entries of the routing table $RT.a$ have converged to their new correct values?”

Node a has two options in this situation. The first option is that node a drops every packet. This option is attractive if link $a \rightarrow b$ is very reliable or if the rate of packets that need to be transmitted over link $a \rightarrow b$ is very small. The second option is that node a anticipates the failure of link $a \rightarrow b$ and maintains alternative routes that can be used to reroute around link $a \rightarrow b$ when this link fails. To do so, node a needs to advertise to every node in the network that link $a \rightarrow b$ has been designated (by node a) to be a *reroute link*. Thus every node in the network can proceed to help node a identify and maintain alternative routes that can be used to reroute around link $a \rightarrow b$ when it fails.

Each node a in the network is provided with a *rerouting table* $RR.a$ that has four columns: (*rlink*, *srlg*, *rbits*, *relay*). The first column, *rlink*, in every rerouting table lists all the links that have been designated, by their source nodes, as reroute links. The second column, *srlg*, lists the id of the shared risk link group that this reroute link belongs to. The other two columns, *rbits* and *relay*, are discussed below in Section 4.

Initially, the rerouting table $RR.a$ of each node a is empty except that node a adds one entry for each link $a \rightarrow s_n$ that a wants to designate as a reroute link. Whenever node a sends a copy of its routing table $RT.a$ to each neighboring node, node a also sends a copy of its rerouting table $RR.a$ (excluding the “relay” column) to the neighboring node. This periodic exchange of routing and rerouting tables between neighboring nodes in the network eventually causes every link that has been designated as a reroute link to have one entry in every rerouting table in the network. At any time, each node a can change the set of links that it has designated as reroute links by adding new links to this set or by removing old links from this set.

3 Relay Nodes for SRLG failures

In this section, we introduce the concept of a relay node for a reroute link that is a member of a SRLG in a network. (Note that the relay node defined in this section also works for reroute links that does not belong to any SRLG in the network.) We then discuss how a relay node for a reroute link can be used in rerouting around its reroute link when all the links that belong to the same SRLG fail.

Let s and d be two nodes in a network N , and let $R(s, d)$ denote the shortest route from node s to node d as determined by the routing tables. A node r is called a *relay node* for a reroute link $a \rightarrow b$ in N iff neither the route $R(a, r)$ nor the route $R(r, b)$ contains any link which belongs to the same SRLG as link $a \rightarrow b$ (including $a \rightarrow b$ itself).

As an example, consider network N_1 in Figure 1. Link $a \rightarrow b$, $c \rightarrow d$ and $e \rightarrow f$ belong to the same SRLG g_1 . If link $a \rightarrow b$ in network N_1 is designated, by node a , as a reroute link, then by definition, node g , h and i are relay nodes for link

$a \rightarrow b$. However, if link $b \rightarrow c$ and $c \rightarrow h$ belong to the same SRLG (not shown in N_1), and link $b \rightarrow c$ is designated, by node b , as a reroute link, then no node in network N_1 is a relay node for link $b \rightarrow c$. These two examples simply show that a reroute link in a SRLG can have one or more relay nodes or even no relay nodes. Fortunately, we show by extensive simulations below in Section 6 that reroute links in a SRLG that have no relay nodes are extremely rare.

Next we describe the procedure for rerouting around a reroute link $a \rightarrow b$, when this link fails, assuming that node a knows the identity of a relay node r for link $a \rightarrow b$:

1. Assume that a packet is to be sent from a node s to a node d along the route $R(s, d)$ which contains the reroute link $a \rightarrow b$. In this case, the IP header of the packet can be represented as (*from s , to d*).
2. Assume also that when this packet reaches node a , node a discovers that the reroute link $a \rightarrow b$ has failed and so decides to reroute the packet towards the relay node r for link $a \rightarrow b$. In this case, node a encapsulates the packet in two outer IP headers (*from a , to b*) and then (*from a , to r*) and forwards the encapsulated packet towards r .
3. When the encapsulated packet reaches the relay node r , node r removes the outermost IP header (*from a , to r*) and discovers that the packet has an inner IP header (*from a , to b*). Thus node r forwards the encapsulated packet towards b .
4. When the encapsulated packet reaches node b , node b removes the outer IP header (*from a , to b*) and discovers that the packet has an inner IP header (*from s , to d*). Thus node b forwards the packet towards d . Note that now the packet is not encapsulated any more.
5. Assume that, while the packet is traversing the route $R(b, d)$ (non-encapsulated now like a normal packet), the packet reaches a node x that needs to transmit the packet over a reroute link $x \rightarrow y$ except that it discovers that link $x \rightarrow y$ has failed. Node x can use the same procedure, step 1 through 4 as describe above, to reroute the packet around the failed link $x \rightarrow y$, no matter whether link $x \rightarrow y$ belongs to the same SRLG as link $a \rightarrow b$ or not.
6. However, assume that while the encapsulated packet is traversing the route $R(a, r)$ or the route $R(r, b)$, the packet reaches a link $x \rightarrow y$ that has failed. Recognizing that the packet is being rerouted because it is an encapsulated packet, node x drops the packet and does not attempt to reroute it a second time.

The fact that no routing loops are created due to the repeated rerouting of the same packet is established in the following lemma and theorem. The proofs can be found in [19].

Lemma 1. *Let r be a relay node for a reroute link $a \rightarrow b$ in a network N , and let d be any node in network N such that the route $R(a, d)$ contains the reroute link $a \rightarrow b$. If link $a \rightarrow b$ fails, and node a reroutes a packet, whose ultimate destination is node d , to node r , then this packet will not traverse any loop in the network before reaching node b no matter whether link $a \rightarrow b$ belongs to a SRLG or not.*

Theorem 1. *No routing loops are created due to repeated rerouting of the same packet to its ultimate destination using the rerouting procedure, no matter whether all encountered failed links belong to the same SRLG or not.*

4 Relay Bits to identify Relay Nodes

In the previous section we presented a procedure by which a node a can reroute packets around a reroute link $a \rightarrow b$ when all the links that belong to the same SRLG fail. This procedure is based on the assumption that node a knows a relay node for the reroute link $a \rightarrow b$. So the question now is “How does node a know a relay node for link $a \rightarrow b$ which belongs to a SRLG without access to the global topology information?” In this section we present a fully distributed procedure by which node a learns all the relay nodes for link $a \rightarrow b$ although node a does not know any other links that belong to the same SRLG as link $a \rightarrow b$. This procedure consists of the following three parts.

The First Part: For node a to announce that it has designated link $a \rightarrow b$ as a reroute link, node a adds the entry $(a \rightarrow b, g_1, 00, -)$ to its rerouting table $RR.a$. Recall that each entry in a rerouting table consists of four components $(rlink, srlg, rbits, relay)$, where $rlink$ is a reroute link; $srlg$ lists the id of the SRLG that this link belongs to; $rbits$ are two relay bits (to be discussed shortly) for the reroute link; and $relay$ is the set of all known relay nodes for the reroute link. The initial value of $relay$ is “-” which indicates that node a does not know yet this value.

Because the rerouting table of every node is sent periodically to every neighbor of this node, the fact that link $a \rightarrow b$ has been designated a reroute link, as well as the SRLG g_1 it belongs to, is eventually recorded in every rerouting table in the network according to the following rule. If a node x receives a rerouting table $RR.y$ from a neighbor y , and the next hop for reaching node a in the routing table $RT.x$ of node x is node y , and if $RR.y$ has an $a \rightarrow b$ entry but $RR.x$ does not have $a \rightarrow b$ entry, then node x adds an entry $(a \rightarrow b, g_1, 00, -)$ to its rerouting table $RR.x$. Conversely, if $RR.y$ has no $a \rightarrow b$ entry but $RR.x$ has an $a \rightarrow b$ entry, then node x removes the $a \rightarrow b$ entry from its rerouting table $RR.x$.

The Second Part: For each reroute link $a \rightarrow b$, node x maintains two bits, named the *relay bits* of link $a \rightarrow b$, in its rerouting table $RR.x$. These two bits are denoted $rbits.x[a \rightarrow b]$ and each of the two bits has anyone of two values. The value “0” in the first bit indicates two cases: either node x does not know yet the correct value of the bit (i.e., initial value of the bit), or node x has checked that some link that belongs to the same SRLG g_1 as link $a \rightarrow b$ occurs in the route $R(x, a)$. The value “1” in the first bit indicates that x has checked that no link that belongs to the same SRLG g_1 as link $a \rightarrow b$ occurs in the route $R(x, a)$. Similarly, the value in the second bit indicates the same meaning except that node x checks whether there is any link that belongs to the same SRLG g_1 as link $a \rightarrow b$ occurs in the route $R(x, b)$. Only when the two bits are both “1”s, i.e., no link that belongs to the same SRLG g_1 as link $a \rightarrow b$ occurs in route $R(x, a)$ and route $R(x, b)$, node x is a relay node for reroute link $a \rightarrow b$.

Next, we describe how to set up the two relay bits for a reroute link. Initially, the value of $rbits.x[a \rightarrow b]$ is “00” in the rerouting table $RR.x$ in every node x in the network, meaning that every node does not know the correct value of the bits yet. The source node a of link $a \rightarrow b$ assigns the relay bits $rbits.a[a \rightarrow b]$ in its rerouting table $RR.a$ the value 10. The first bit “1” means that no link that belongs to the same SRLG g_1 as link $a \rightarrow b$ occurs in the route $R(a, a)$, and the second bit “0” means that link $a \rightarrow b$ occurs in the route $R(a, b)$. Then the sink node b of link $a \rightarrow b$ assigns the bits $rbits.b[a \rightarrow b]$ in its rerouting table $RR.b$ the value 01. The first bit “0” means that link $b \rightarrow a$ that belongs to the same SRLG g_1 as link $a \rightarrow b$ occurs in the route $R(b, a)$, and the second bit “1” means that no link that belongs to the same SRLG g_1 as link $a \rightarrow b$ occurs in the route $R(b, b)$.

Then every other node x in the network assigns each of the two relay bits $rbits.x[a \rightarrow b]$ in its rerouting table $RR.x$ the value val , where val is either 0 or 1, according to the following rule: If x receives $RR.y$ from neighbor y , and if the next hop for reaching node a in the routing table $RT.x$ of node x is node y , then node x checks whether the SRLG of link $x \rightarrow y$ is the same as the SRLG of link $a \rightarrow b$, if yes, then node x assigns the first bit $rbits.x[a \rightarrow b][0]$ in its $RR.x$ the value 0. Otherwise, node x assigns the first bit $rbits.x[a \rightarrow b][0]$ in its $RR.x$ the value of the first bit in $rbits.y[a \rightarrow b][0]$. Similarly, node x assigns the second relay bit.

The first and second parts outlined above are part of updating the rerouting table $RR.x$ after node x receives the rerouting table $RR.y$ from the neighboring node y shown in Algorithm 1.

Algorithm 1: Update rerouting table $RR.x$ after x receives rerouting table $RR.y$ from neighbor y

```

1: for  $(s_r \rightarrow s_n \in \text{rlink}.x)$  and  $(s_r \rightarrow s_n \notin \text{rlink}.y)$  do
2:   if  $\text{nextthop}.x[s_r] == y$  then
3:     remove  $s_r \rightarrow s_n$  entry from  $RR.x$ ;
4: for  $(s_r \rightarrow s_n \notin \text{rlink}.x)$  and  $(s_r \rightarrow s_n \in \text{rlink}.y)$  do
5:   if  $\text{nextthop}.x[s_r] == y$  then
6:     add  $s_r \rightarrow s_n$  entry to  $RR.x$ ;
7: for  $s_r \rightarrow s_n \in \text{rlink}.x$  do
8:   if  $x == s_r$  then
9:      $\text{rbits}.x[s_r \rightarrow s_n] := 10$ ;
10:  else if  $x == s_n$  then
11:     $\text{rbits}.x[s_r \rightarrow s_n] := 01$ ;
12:  else
13:    if  $\text{nextthop}.x[s_r] == y$  then
14:      if  $\text{srlg}.x[x \rightarrow y] == \text{srlg}.x[s_r \rightarrow s_n]$  then
15:         $\text{rbits}.x[s_r \rightarrow s_n][0] := 0$ ;
16:      else
17:         $\text{rbits}.x[s_r \rightarrow s_n][0] := \text{rbits}.y[s_r \rightarrow s_n][0]$ ;
18:    if  $\text{nextthop}.x[s_n] == y$  then
19:      if  $\text{srlg}.x[x \rightarrow y] == \text{srlg}.x[s_r \rightarrow s_n]$  then
20:         $\text{rbits}.x[s_r \rightarrow s_n][1] := 0$ ;
21:      else
22:         $\text{rbits}.x[s_r \rightarrow s_n][1] := \text{rbits}.y[s_r \rightarrow s_n][1]$ ;

```

Table 2. Rerouting table $RR.a$ of node a in network N_1 with the relay nodes for the reroute links $a \rightarrow b$ whose source node is a

rlink	srlg	rbits	relay
a→b	g_1	10	g, h, i
a→e	-	10	b, c, d
c→d	g_1	00	-
e→f	g_1	10	-
b→f	g_2	11	-
f→g	g_2	11	-
g→h	g_3	10	-
h→i	g_3	10	-
d→i	-	11	-

The Third Part: When node a receives a $notify(x, a \rightarrow b)$ message, Node a then adds x to the set of $relay.a[a \rightarrow b]$ in the rerouting table $RR.a$ of node a .

In Figure 1, assume that all the links in three SRLGs g_1 , g_2 and g_3 have been designated by their respective source nodes to be reroute links. Also, assume that link $a \rightarrow e$ and link $d \rightarrow i$ have been designated as reroute links too. But these two links do not belong to any SRLG. Then the rerouting table $RR.a$ of node a , after these links have been designated as reroute links, is shown in Table 2.

Correctness of the procedure for updating the relay bits follows from the next theorem. The proof can be found in [19].

Theorem 2. *For any node x in a network N , if the relay bits in node x for a reroute link $a \rightarrow b$ are both ones, i.e., $rbits.x[a \rightarrow b]=11$, then neither route $R(a,x)$ nor route $R(x,b)$ contains any link that belongs to the same SRLG as reroute link $a \rightarrow b$.*

5 Suppression Mode

There is one problem concerning the second and third part of the procedure discussed in the previous section: for some reroute links many nodes in the network qualify to be relay nodes and so these many nodes start to send notify messages to the source node of the link, and the source node has to process all the notify messages even though only one relay node is enough to reroute packets around the link when it fails.

In order to minimize the notification messages sent in the network, we introduce a *suppression mode*. In the suppression mode, when the relay bits $rbits.x[a \rightarrow b]$ in the rerouting table $RR.x$ have the value 11, node x recognizes that it is a relay node for link $a \rightarrow b$ and so it sends a $notify(x, a \rightarrow b)$ message to its next hop for reaching node a , which either drops the message (as explained below) or forwards the message to its next hop for reaching node a . Thus, if the $notify(x, a \rightarrow b)$ message reaches, along the route $R(x, a)$, a node y where the relay bits $rbits.y[a \rightarrow b]$ in the rerouting table $RR.y$ have the value 11, then node y drops the $notify(x, a \rightarrow b)$ message knowing that its own $notify(y, a \rightarrow b)$ message is sufficient for node a to have one relay node for link $a \rightarrow b$.

If the suppression mode is used in network N_1 in Figure 1, then for reroute link $a \rightarrow b$, the notify message from node i is dropped by node h , and the notify messages from h is dropped by g .

The actions of a node x concerning the sending and receiving of notify messages are shown in Algorithm 2.

6 Simulation Results

We now evaluate the performance of our IP fast reroute mechanism for various size of SRLGs (i.e., the number of links that are members of a SRLG) using simulations. Through simulation, we intend to answer the following questions: 1) What is the repair coverage for various size of SRLGs? 2) what is the efficiency of suppression under different size of SRLGs? 3) What is the chance that a node can have multiple relay nodes to choose for various size of SRLGs? Will the

Algorithm 2: Actions of node x on sending and receiving relay notify messages

```

/* -----sending action----- */
1 for  $s_r \rightarrow s_n \in rlink.x$  do
2   if ( $rbits.x[s_r \rightarrow s_n] == 11$ ) then
3     send  $notify(x, s_r \rightarrow s_n)$  to  $nexthop.x[s_r]$ ;
/* -----receiving action----- */
4 rcv  $notify(z, s_r \rightarrow s_n)$  from a neighbor  $y$  do:
5   if  $x == s_r$  then
6     add  $z$  to  $relay.x[s_r \rightarrow s_n]$ 
7   else if ( $rbits.x[s_r \rightarrow s_n] == 11$ ) then
8     suppress  $notify(z, s_r \rightarrow s_n)$ 
9   else
10    forward  $notify(z, s_r \rightarrow s_n)$  to  $nexthop.x[s_r]$ 

```




Fig. 2. Repair coverage of SRLG failures for Barabasi-Albert networks

suppression affect this? 4) What is the overhead of using a relay path, which may include several relays for links in the same SRLG, instead of using the re-converged shortest path? How does the suppression affect this?

We conduct our simulations using two general networks, generated using the BRITE tool [1]. The first network satisfies the power law distribution based on the Barabasi-Albert (BA) model. The second is a random network based on the Waxman model. For each topology with E edges, we randomly select S edges, $1 \leq S \leq 5$ that are close to each other to form a SRLG (the number of hops between the first selected edge and any other selected edge is no larger than 0.6 the maximum number of hops in the network). For $S = 1$, we count every single link failure. For $S > 1$, we generate up to 1000 different SRLG failures and make sure each SRLG is not a cut of the topology graph.

Let *repair coverage* be the percentage of source-destination pairs in which, when the link $source \rightarrow sink$ of a SRLG used to traverse packets from the source to the destination fails, the source can reroute around any failed link in the same SRLG which appears along the path to reach the destination. Figure 2 and 3 show the repair coverage for SRLG failures for BA and Waxman network respectively. For both BA and Waxman network, no matter what is the size of the network, the repair coverage for smaller SRLG size is greater than the repair coverage for larger SRLG size. However, when the network size is at least 100 nodes, the SRLG size does not have much effect on the repair coverage and our IP fast reroute mechanism can achieve close to 100% repair coverage in these cases.

We measure the efficiency of suppression using *suppress ratio*, which is defined as the percentage of suppressed relay notify messages. As shown in Figure 4 and 5, in both BA and Waxman networks, the size of the SRLG does not affect the suppress ratio much. If the network size is larger than one hundred nodes, then the suppress ratio is about 80%. This demonstrates that suppression will effectively save the processing overhead for the source node and the bandwidth in the network.

In both BA and Waxman networks, no matter what's the size for the SRLG, when the network size is over one hundred nodes and there is no suppression,




Fig. 3. Repair coverage of SRLG failures for Waxman networks




Fig. 4. Suppress ratio of relay notify messages for BA networks




Fig. 5. Suppress ratio of relay notify messages for Waxman networks




Fig. 6. Probability of more than one relay node for SRLG failures in BA networks

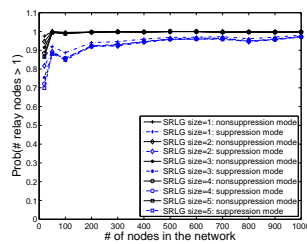


Fig. 7. Probability of more than one relay node for SRLG failures in Waxman networks




Fig. 8. The average path stretch when choosing different relay nodes for SRLG failures in non-suppression mode for BA networks

the chance that a source node can find multiple relay nodes to choose from instead of only one relay node is over 97%, shown in Figure 6 and 7. While in the suppression mode, since some relay notify messages are suppressed, the chance that a source node can find multiple relay nodes drops to over 88% in Waxman networks and to about 80% in BA networks. However, we will show that the suppression mode will not affect the best relay node in terms of reroute path length and it also gives a source node better choices in terms of reroute path length.

For a reroute link, the pre-computed alternative path through a relay node is not necessarily the shortest path. This is because only the source node of the reroute link is aware of the failure and no other nodes are. So compared to the globally re-converged shortest path, IP fast reroute gains the lossless forwarding with a possible longer path penalty. However, we show that the penalty is not significant. Let *path stretch* be the ratio of the length of the pre-computed alternative path going through the relay node(s) divided by the length of the shortest path after re-convergence. When a source node finds that there are multiple relay nodes for a reroute link, which relay node should the source choose? We examine three choices in terms of path stretch: the closest relay node to the source, the farthest relay node to the source and a random relay node.




Fig. 9. The average path stretch when choosing different relay nodes for SRLG failures in nonsuppression mode for Waxman networks




Fig. 10. The average path stretch when choosing different relay nodes for SRLG failures in suppression mode for BA networks

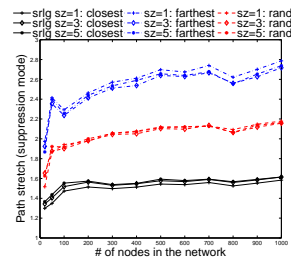


Fig. 11. The average path stretch when choosing different relay nodes for SRLG failures in suppression mode for Waxman networks

In nonsuppression mode (i.e., suppression is not applied), the average path stretch when choosing difference relay nodes for different size of SRLG failures, in BA and Waxman networks is shown in Figure 8 and 9 respectively. In both networks, no matter what's the size of the SRLG failures, choosing the closest relay node gives the smallest path stretch, less than 1.6 compared to the re-converged shortest path length, while choosing the farthest relay node gives the largest path stretch. A random relay has the stretch in between the above two.

Figure 10 and 11 show the corresponding path stretch under the suppression mode. It is clear that suppression will not affect the path stretch for closest relay nodes. However, since suppression filters some farther relay nodes which tend to have larger stretch, the average path stretch for both farthest relay nodes and the random relay nodes is reduced under suppression mode. So a source can also randomly choose a relay node with stretch lower than or about 2 in both types of networks.

7 Related Work

Recently, IP Fast Reroute (IPFRR) has been proposed to recover from failures as soon as a failure is detected using IP-based schemes [26]. However, existing proposals, except [27] which requires substantial number of additional IP addresses, mainly focus on how to handle a single link failure or dual-link failures [3, 4, 6, 14, 16, 20, 24]. Also, most of existing proposals assume each node has the knowledge of some global topology information [3, 4, 6, 8, 14, 16–18, 23, 24, 27]. Instead, our work focuses on shared risk link group failures and assumes that each node has neither global connectivity information of the network nor additional global IP addresses information associated with each node. The idea of precomputing backup paths is also explored for BGP [15, 23, 25, 28, 29].

An IPFRR scheme should be able to avoid micro-loops [7, 9, 11]. Francois et al. [10] and Gjoka et al. [13] evaluate the coverage of several IPFRR mechanisms.

Both Loop-free Alternates [4] and U-turn Alternates [3] pre-computes an alternate next hop before a single link failure. Since these two mechanisms find

alternates only among next hops, the coverage is not high even for single link failures. Tunnels [6] is more generalized than the above two mechanisms in the sense that it is not limited to only use next hops as tunnel endpoints. But again it can only handle single link failures and is only designed for link state protocols. Also it requires a significant number of computations of shortest paths since it computes a reroute path for each of the neighbors of the sink node.

In Multiple Routing Configurations (MRC) [16], each router pre-computes a number of topology configurations by removing rerouted links. Failure Insensitive Routing (FIR) [24] exploits interface-specific forwarding. Both MRC and FIR focus on single link failures. Failure-Carrying Packets (FCP) [17] uses the packet header to carry the list of failed links and requires potentially expensive dynamic computation to route that packet. Path splicing [23] creates multiple routing trees and allows packets to switch paths by inserting a new packet header. In [20], Li et al. explored the idea of using relay nodes to achieve IP fast reroute around single link failures based only on local information. Kini et. al. [14] proposed an approach to handle two simultaneous link failures by assigning three additional addresses to each node.

Each node in Not-via [27] needs d additional Not-via addresses for all the links for which it is a source node, where d is the degree of that node. These additional IP addresses have to be globally known, even when a link is currently not intended to be a reroute link. This significantly increases the size of the routing table and consequently lower the efficiency of forwarding even when there is no failures. Recent work from Li et al. [18] try to improve the efficiency of Not-via by aggregation, but it requires special allocation schemes of Not-via addresses. Enyedi et al. [8] try to reduce the number of Not-via addresses but they also assume the knowledge of global connectivity information.

8 Concluding Remarks

We have presented an IP fast reroute mechanism for Shared Risk Link Group failures in routing protocols without global topology information. In our mechanism, any node x can advertise that it needs to be able to reroute around a link $x \rightarrow y$ when this link fails. Then we leverage a set of relay nodes, computed in advance of any link failures, to tunnel the reroute packets around each failed link right after the detection of a failure. Each node uses a fully distributed algorithm to decide automatically whether it can serve as a relay node for a reroute link or not to avoid all link failures in the same SRLG. Notify messages are sent to the source of a reroute link from relay nodes. We proposed a suppression mode to greatly reduce the number of notify messages. Moreover, our tunneling scheme ensures that loops are never formed even when any number of links fail.

Through simulations on different topologies, we confirmed that our mechanism can achieve close to 100% repair coverage in different types and various size of networks for different SRLG size. The average length of a reroute path is around 1.5 the re-converged optimal paths. As expected, the suppression is

quite effective and cut 80% of notify messages in a network of reasonable size (≥ 100).

Our future work includes migrating our IP fast reroute mechanism to inter-domain routing protocols. Using our mechanism, each AS can potentially leverage the existing Internet topology to achieve fast reroute around Shared Risk Link Group Failures, without changing the BGP advertising and decision process.

References

1. BRITE: Boston univeristy Representative Internet Topology generator. <http://www.cs.bu.edu/BRITE/>.
2. Enhanced interior gateway routing protocol. http://www.cisco.com/en/US/tech/tk365/technologies_white_paper09186a0080094cb7.shtml.
3. A. Atlas. U-turn alternates for IP/LDP fast-reroute. In *Internet Draft, draft-atlas-ip-local-protect-urn-03.txt*, February 2006.
4. A. Atlas and A. Zinin. Basic specification for IP fast reroute: Loop-free alternates. In *RFC 5286*, September 2008.
5. C. Boutremans, G. Iannaccone, and C. Diot. Impact of link failures on VoIP performance. In *NOSSDAV'02*, May 2002.
6. S. Bryant, C. Filsfil, S. Previdi, and M. Shand. IP fast reroute using tunnels. In *Internet Draft, draft-bryant-ipfrr-tunnels-03*, November 2007.
7. S. Bryant and M. Shand. A framework for loop-free convergence. In *Internet Draft, draft-ietf-rtgwg-lf-conv-frmwk-03*, October 2008.
8. G. Enyedi, P. Szilágyi, G. Rétvári, and A. Császár. IP fast reroute: Lightweight not-via without additional addresses. In *IEEE infocom mini-conference*, 2009.
9. P. Francois and O. Bonaventure. Avoiding transient loops during IGP convergence in IP networks. In *IEEE Infocom*, 2005.
10. P. Francois and O. Bonaventure. An evaluations of IP-based fast reroute techniques. In *CoNext*, 2005.
11. P. Francois, O. Bonaventure, M. Shand, S. Bryant, and S. Previdi. Loop-free convergence using oFIB. In *Internet draft, draft-ietf-rtgwg-ordered-fib-02*, February 2008.
12. P. Francois, C. Filsfil, J. Evans, and O. Bonaventure. Achieving subsecond IGP convergence in large IP networks. In *ACM Sigcomm*, 2005.
13. M. Gjoka, V. Ram, and X. Yang. Evaluations of IP fast reroute proposals. In *IEEE Comsware*, 2007.
14. S. Kini, S. Ramasubramanian, A. Kvalbein, and A. Hansen. Fast recovery from dual link failures in IP networks. In *IEEE Infocom*, 2009.
15. N. Kushman, S. Kandula, D. Katabi, and B. Maggs. R-BGP: Staying connected in a connected world. In *Usenix NSDI*, 2007.
16. A. Kvalbein, A. F. Hansen, T. Cicic, S. Gjessing, and O. Lysne. Fast IP network recovery using multiple routing configurations. In *IEEE infocom*, 2006.
17. K. Lakshminarayanan, M. Caesar, M. Rangan, T. Anderson, S. Shenker, and I. Stoica. Achieving convergence-free routing using failure-carrying packets. In *ACM Sigcomm*, 2007.
18. A. Li, P. Francois, and X. Yang. On improving the efficiency and manageability of NotVia. In *CoNext*, 2007.
19. Y. Li and M. G. Gouda. IP fast reroute in networks with shared risk links. UTCS Technical Report TR-09-38, The University of Texas at Austin, 2009.

20. Y. Li and M. G. Gouda. IP fast reroute without global topology information. UTCS Technical Report TR-09-34, The University of Texas at Austin, 2009.
21. G. Malkin. RFC 2453 - RIP Version 2, November 1998.
22. A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C.-N. Chuah, and C. Dio. Characterization of failures in an IP backbone. In *IEEE Infocom'04*.
23. M. Motiwala, M. Elmore, N. Feamster, and S. Vempala. Path splicing. In *ACM Sigcomm*, 2008.
24. S. Nelakuditi, S. Lee, Y. Yu, Z.-L. Zhang, and C.-N. Chuah. Fast local rerouting for handling transient link failures. *IEEE/ACM Transaction on Networking*, 15(2), April 2007.
25. C. F. Olivier Bonaventure and P. Francois. Achieving sub-50 milliseconds recovery upon BGP peering link failures. *IEEE/ACM Transaction on Networking*, 15(5), October 2007.
26. M. Shand and S. Bryant. IP fast reroute framework. In *Internet Draft, draft-ietf-rtgwg-ipfrr-framework-13.txt*, October 2009.
27. M. Shand, S. Bryant, and S. Previdi. IP fast reroute using Not-via addresses. In *draft-ietf-rtgwg-ipfrr-notvia-addresses-03*, October 2008.
28. F. Wang and L. Gao. A backup route aware routing protocol - fast recovery from transient routing failures. In *IEEE Infocom mini-conference*, 2008.
29. F. Wang and L. Gao. Path diversity aware interdomain routing. In *IEEE Infocom*, 2009.