IP Fast Reroute using “notvia” addresses

<draft-bryant-shand-IPFRR-notvia-addresses-00.txt>
Stewart Bryant (stbryant@cisco.com)
Mike Shand (mshand@cisco.com)

Notvia Properties

- Repairs ALL non-partitioning failures.
- Intuitive, easily predictable repair paths.
- Suitable for IP, MPLS/LDP.
- Looks like a good starting point for Multicast FRR.

Notvia Overview

- When a failure occurs, the repairing router needs to get the packet to its destination **not via** the failure.
- For each protected network component – link, node, LAN, or SRG we calculate the required set of notvia paths.
- To repair, we encapsulate the packet to the notvia address on the far side of the failure.
- Use of early terminated incremental SPF makes the calculation tractable.

Notvia Repairs

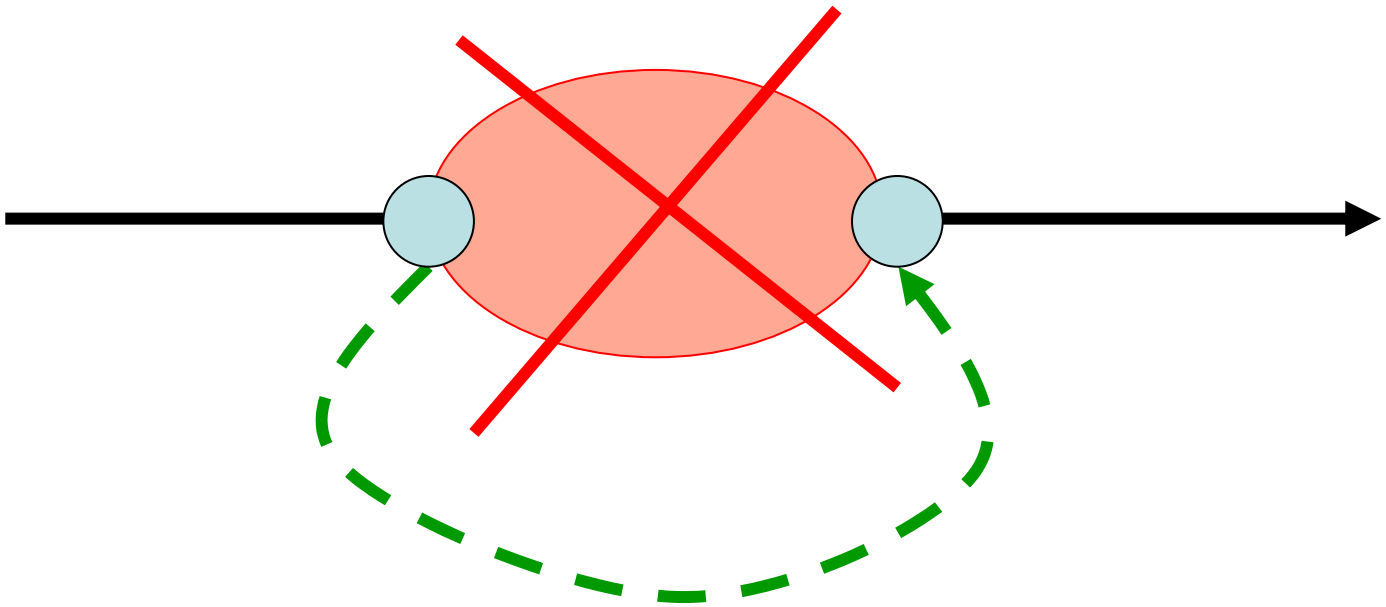

The repairing router delivers the packet downstream of the failure, NOTVIA the failure.

Notvia Addresses

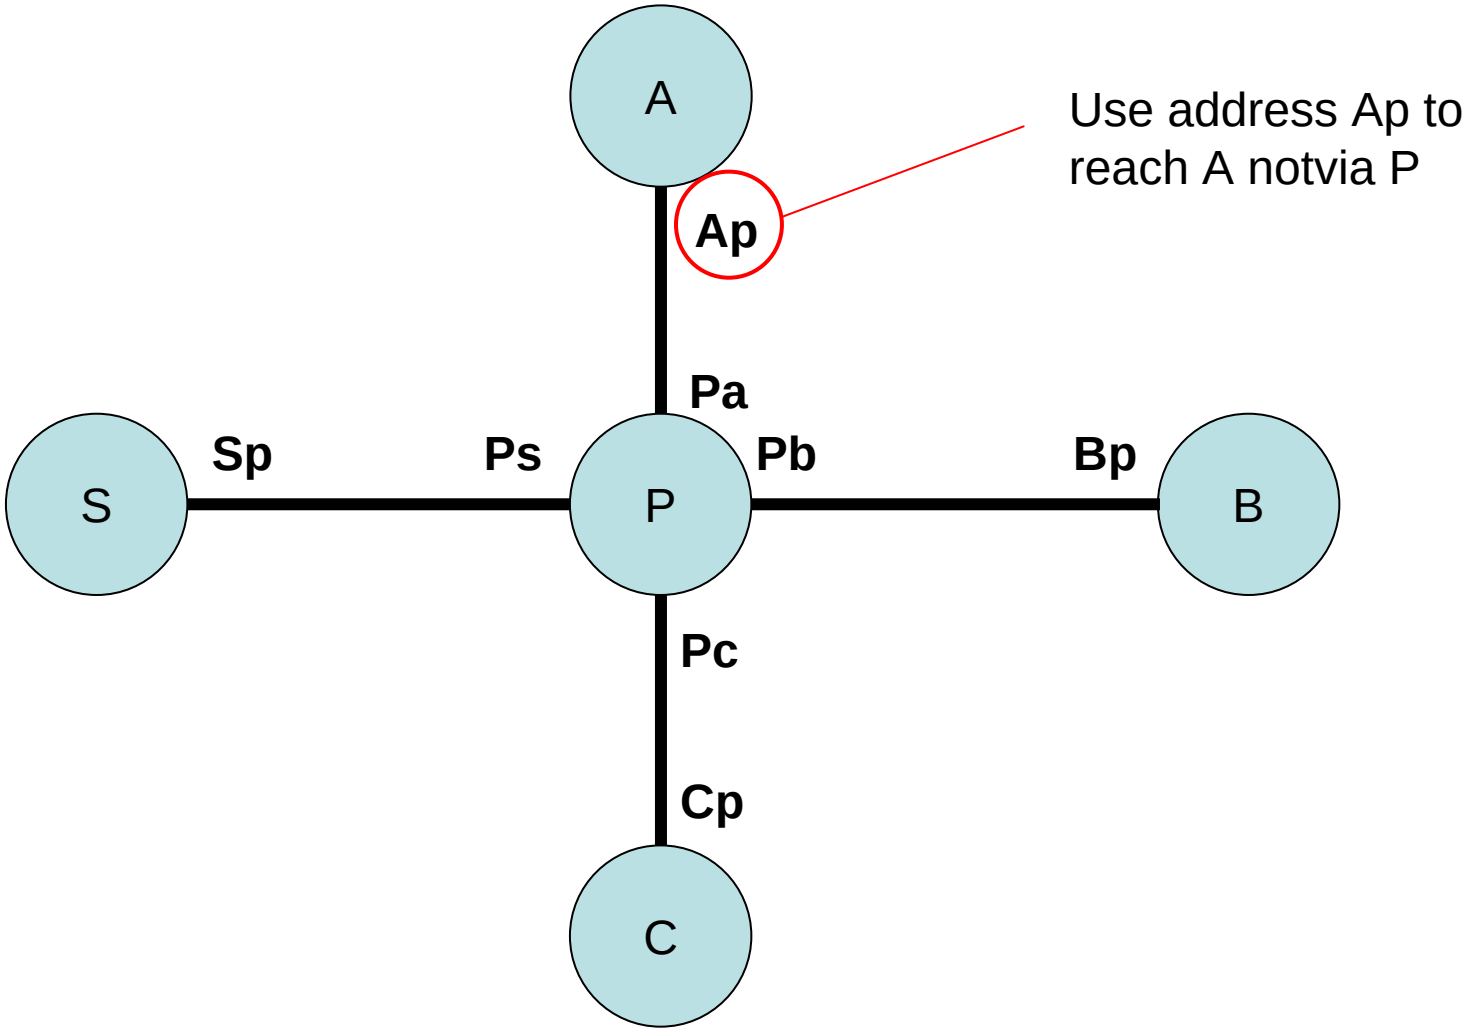

Notvia Repairs

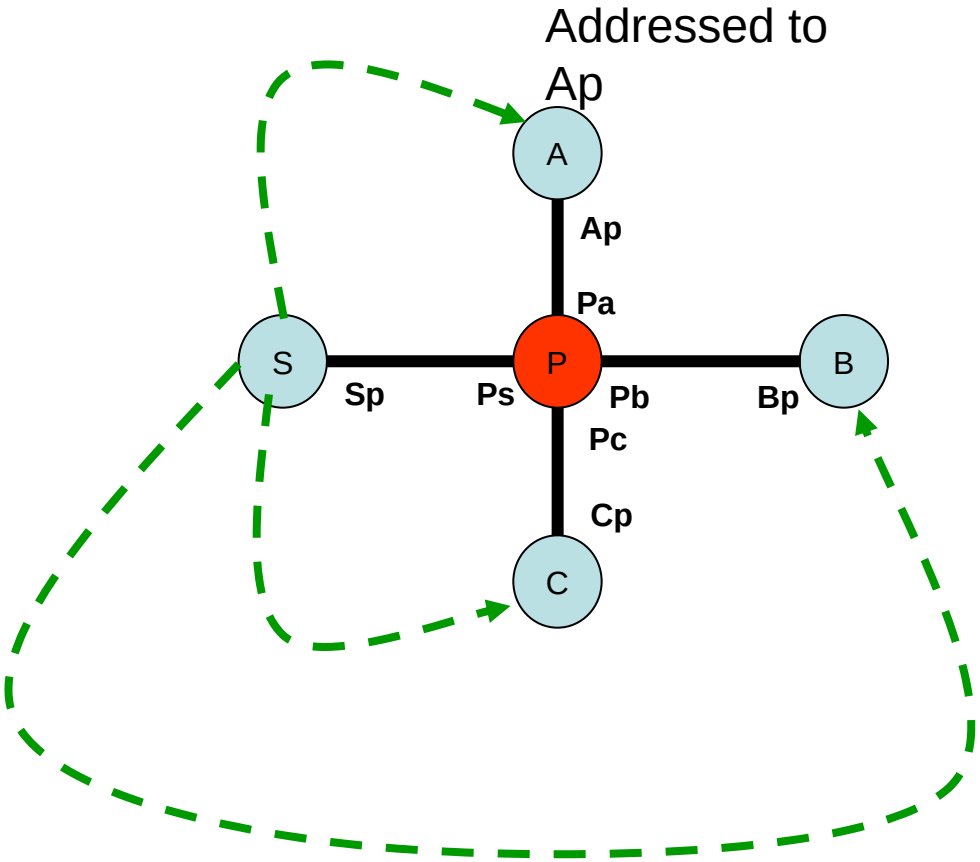

All repairs take the shortest path from S to P's neighbour

Notvia Path Computation

- All routers calculate path to all notvia addresses.
- Cheaper to do the calculation that to find out if calculation is needed.
- For nodes, fail node and calculate path to neighbours of that node.
- Use incremental SPF with early termination.
- For real networks of 40 to 400 nodes, worst case takes between 5 and 13 full SPF compute times.

MPLS-LDP

- Notvia addresses are ordinary IP addresses, and will be distributed by LDP.
- “Just Works” with MPLS TE hardware.

LDP use of Notvia Repair

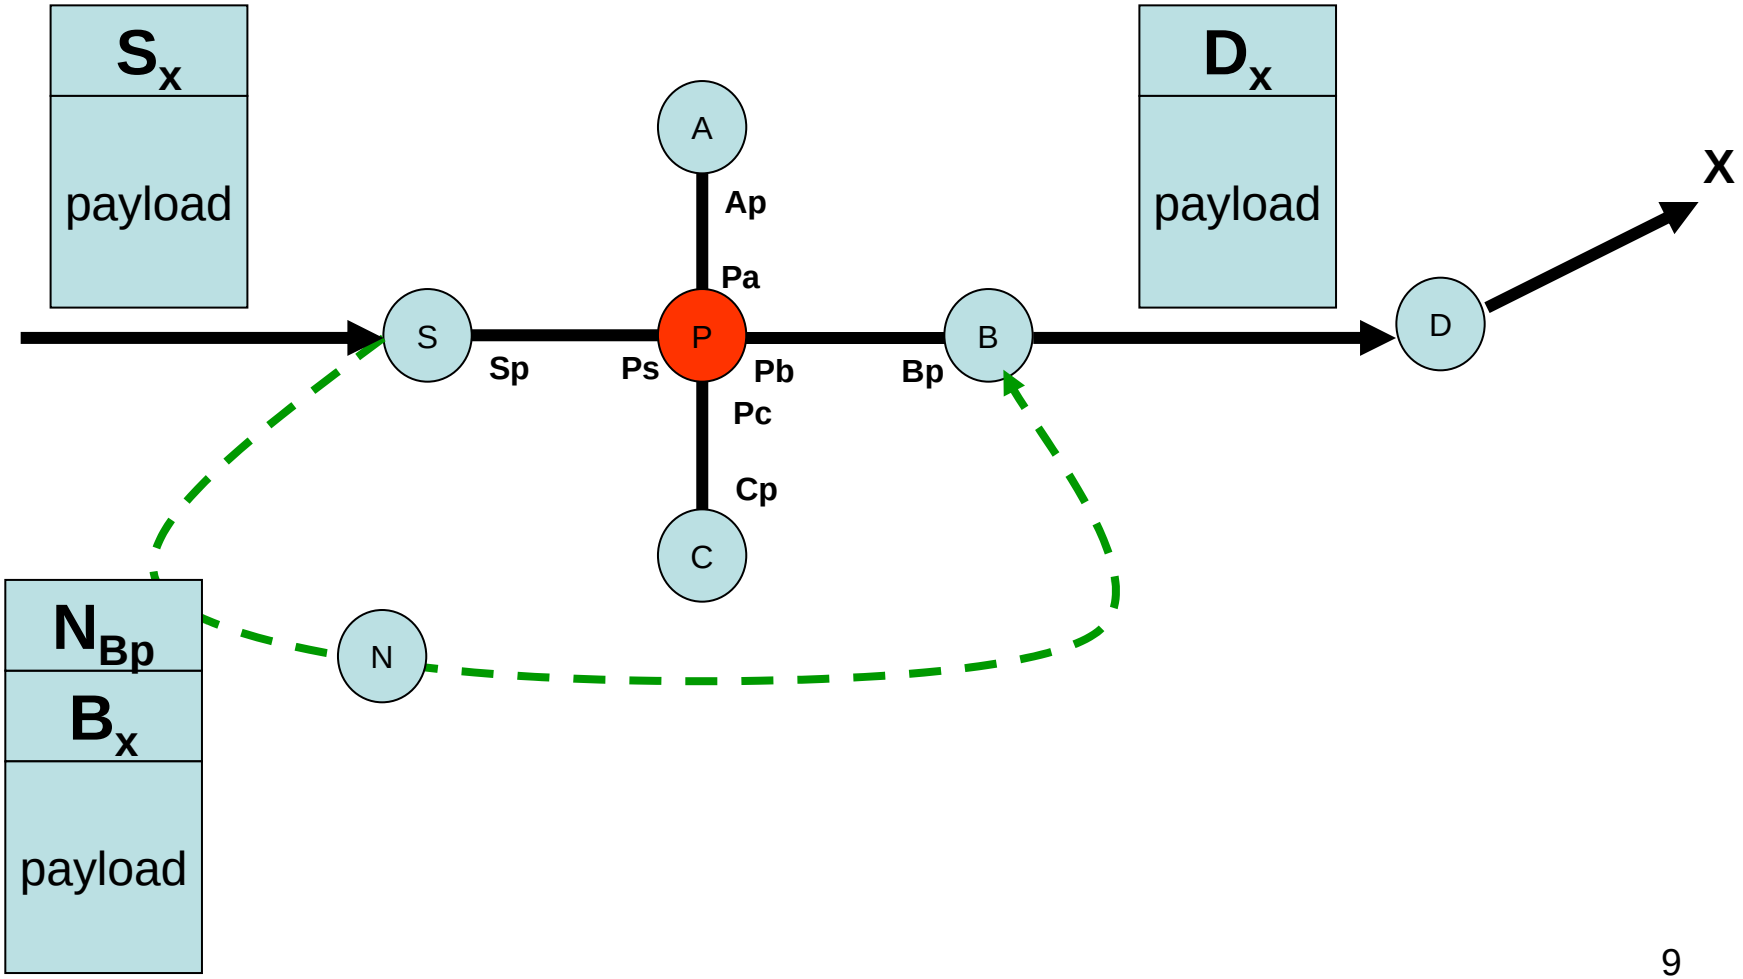

Labels Needed By S

- S already has N’s notvia labels.
- For link repair S already has P’s labels
- For node repair S needs to use directed LDP or <draft-shen-mpls-ldp-nnhop-label-01.txt>
- For SRG S needs directed LDP

Encapsulation

- Any IETF specified IP in IP encapsulation may be used to carry a notvia repair in IP network
  - IP in IP
  - GRE
  - L2TPv3
- Only ONE level of encapsulation needed.

…but it's not just pt-pt unicast nodes, there's

- LANs
- Multicast
- LFA and ECMP
- Incremental deploy'nt
- Routing Extensions
- Link repair
- Misdiagnosis
- Multi-homed prefixes
- SRLG

These requirements apply to ALL FRR solutions

Local Area Networks - 1

S knows that it is not seeing BFD responses from P.

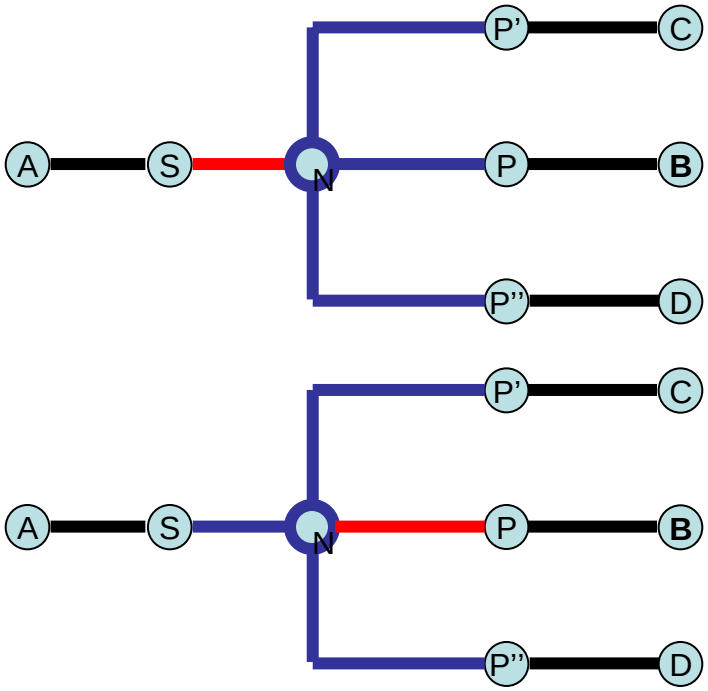

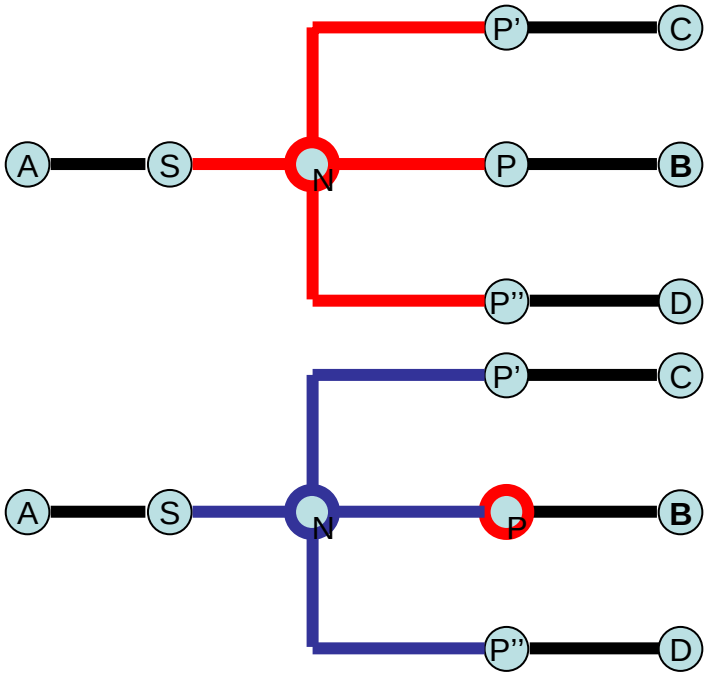

Without further diagnostics, it does not know whether its connection to the LAN, P's connection to the LAN, the LAN, or P has failed.

Local Area Networks - 2

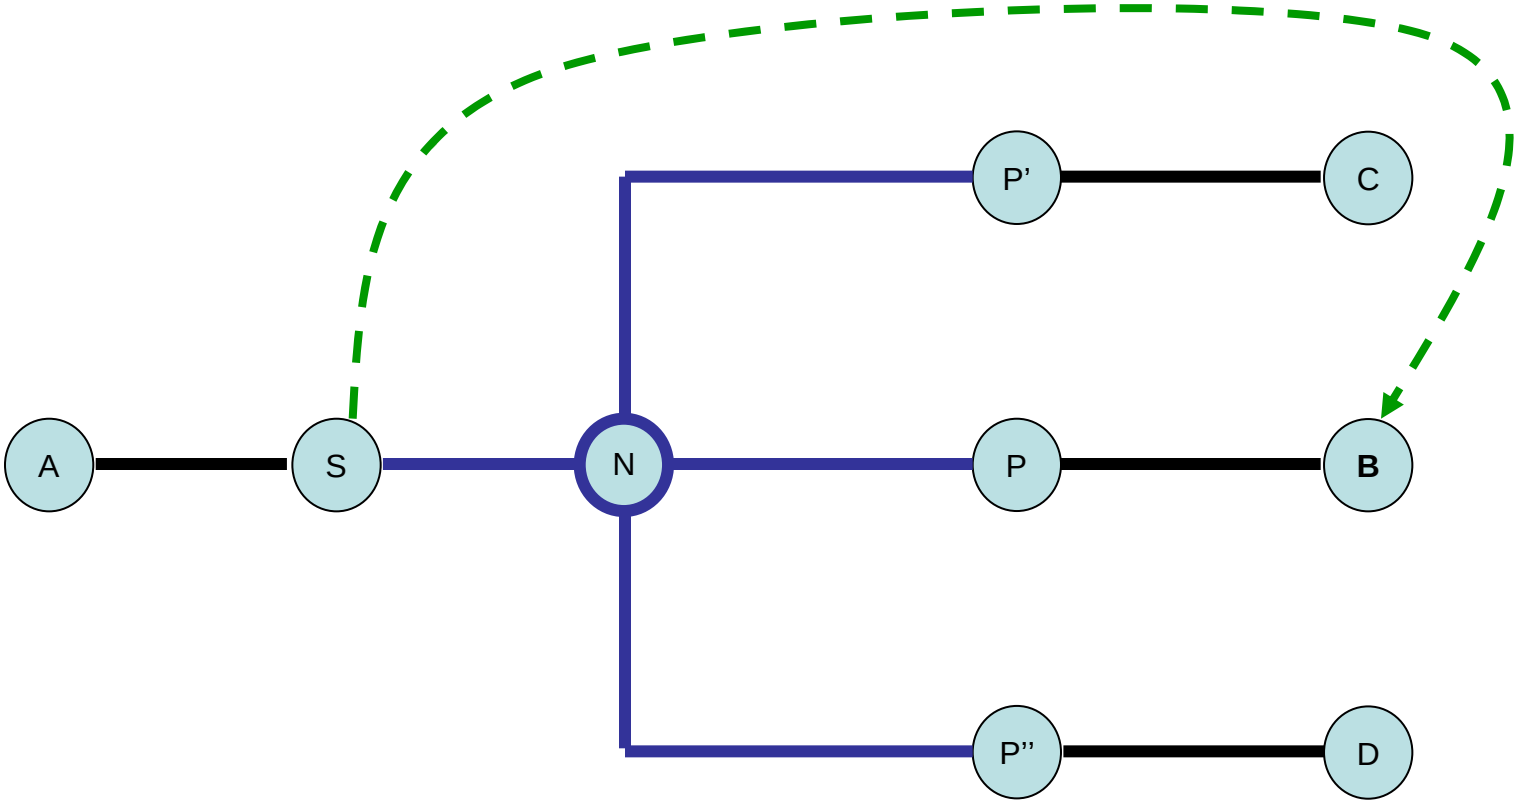

Without further diagnostics S must treat SP adjacency failure as a failure of the P and the WHOLE LAN

Local Area Networks - 3

- S can correlate adjacency checks from P, P', and P'', and diagnose failure of P, or failure of the LAN.

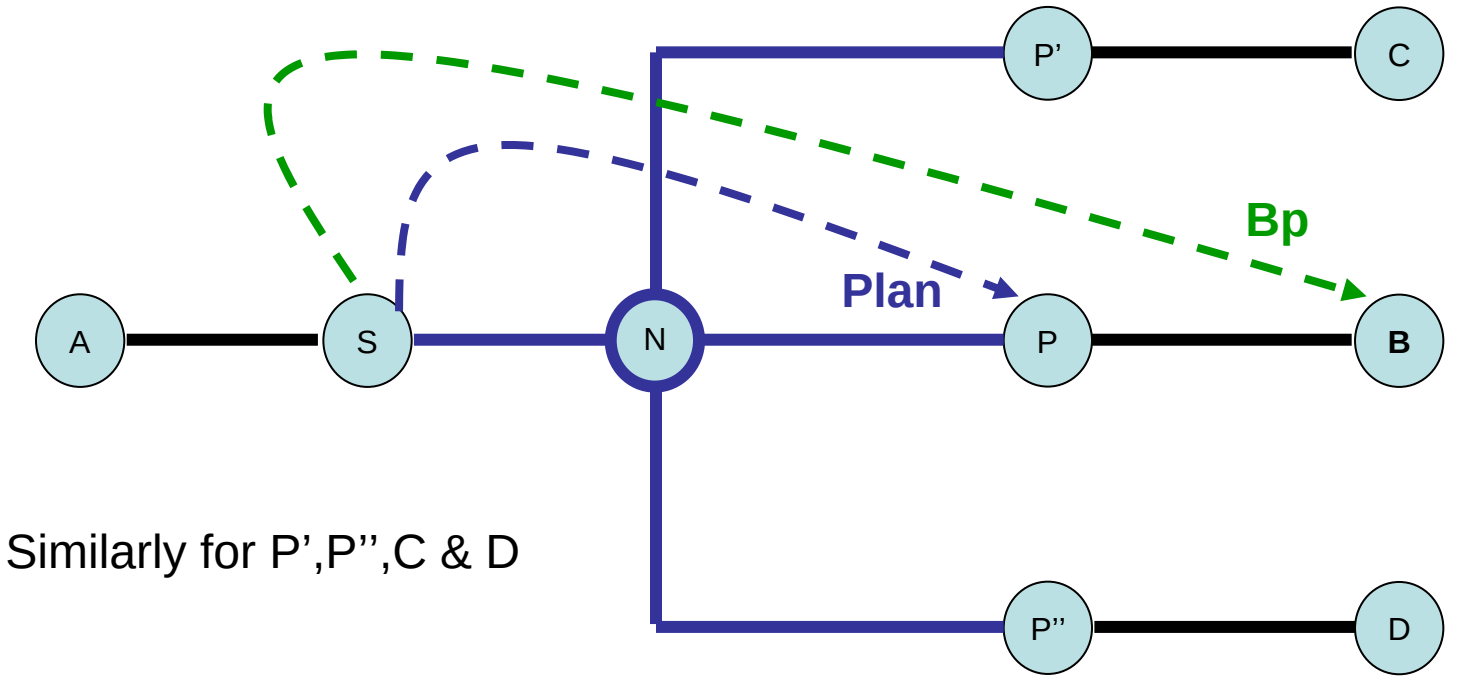

Similarly for P',P'',C & D

Multicast

- Although this is a hard problem, a large proportion of high value, interruption sensitive, traffic is multicast traffic. (e.g. broadcast television and financial trading information).
- To repair multicast it is necessary to get the packet back into the delivery tree appearing to have come from the original input interface.
- This condition can be satisfied using notvia addressing.
- A lot more work is required on this important topic.

Loop Free Alternatives

- Where an LFA exists, S may use this in place of the notvia repair mechanism for unicast packets.
- Multicast traffic seems to require the use of a repair encapsulation.

Equal Cost Multi-Path

- A router can use an equal cost multi-path (ECMP) repair in place of a notvia repair for unicast packets.
- A router computing a notvia repair path MAY subject the repair to ECMP.

Incremental Deployment

- Exclude routers that are not calculating notvia routes from the base repair topology.
- Repairs may be steered around island of routers that are not IPFRR capable.
- Routers that are **protecting** a network component need to have the capability to encapsulate and decapsulate packets.
- Routers that are **on the repair path** only need to be capable of calculating notvia paths and including the notvia addresses in their FIB (i.e. no h/w changes).

Routing Extensions

- IPFRR router directly connected to a protected network component must advertise a notvia address for that component.
  - i.e. one notvia address per neighbor
- Advertisement must associate protected component (router or SRG) and the notvia address.
- Notvia capable routers advertise in the IGP that they will calculate notvia routes.
- It is necessary for routers to advertise the type of encapsulation that they support (LDP, GRE [RFC1701], L2TPv3 etc).

Link Failure

- Normal mode of operation is node failure (most conservative approach)
- We could require no single points of node failure.
- Alternatively accept that some destinations only reachable via neighbor.
- For these destinations only, assume link failure.
- Repair to Ps.
  - Path to Ps already calculated for a node repair of S.

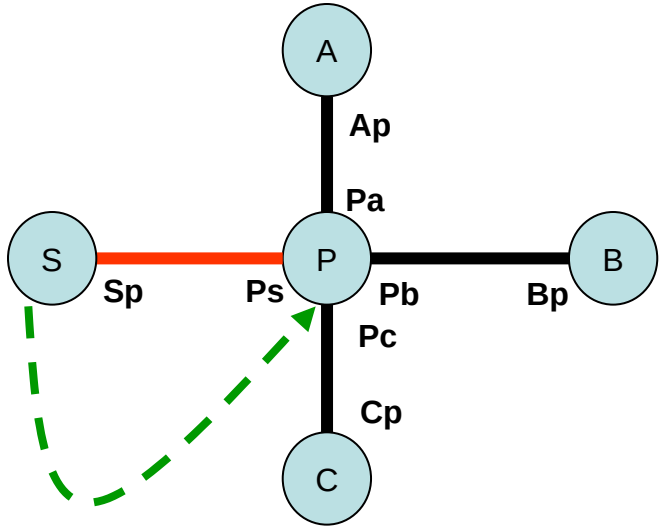

Misdiagnosis

- All solutions must have a strategy for dealing with misdiagnosis.
- For example, attempting link repair in the presence of node failure could give rise to looping.
- We do not provide repair paths for notvia addresses.
- This avoids the problem.

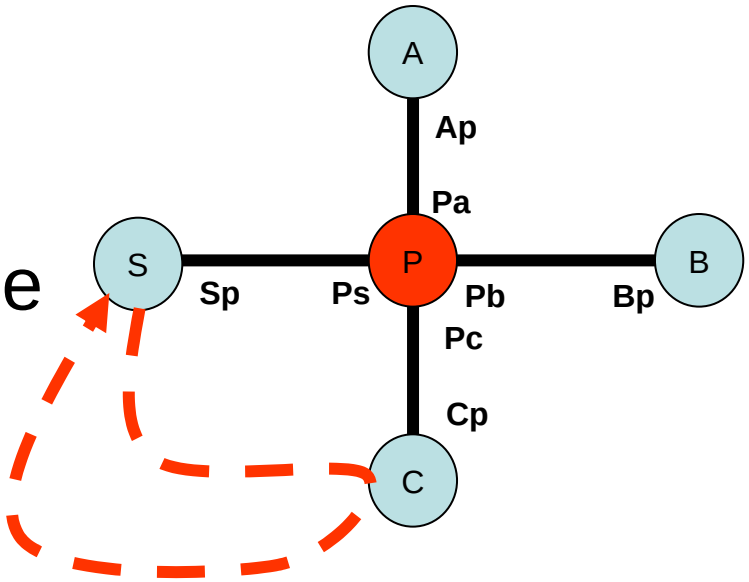

Multi-homed Prefixes

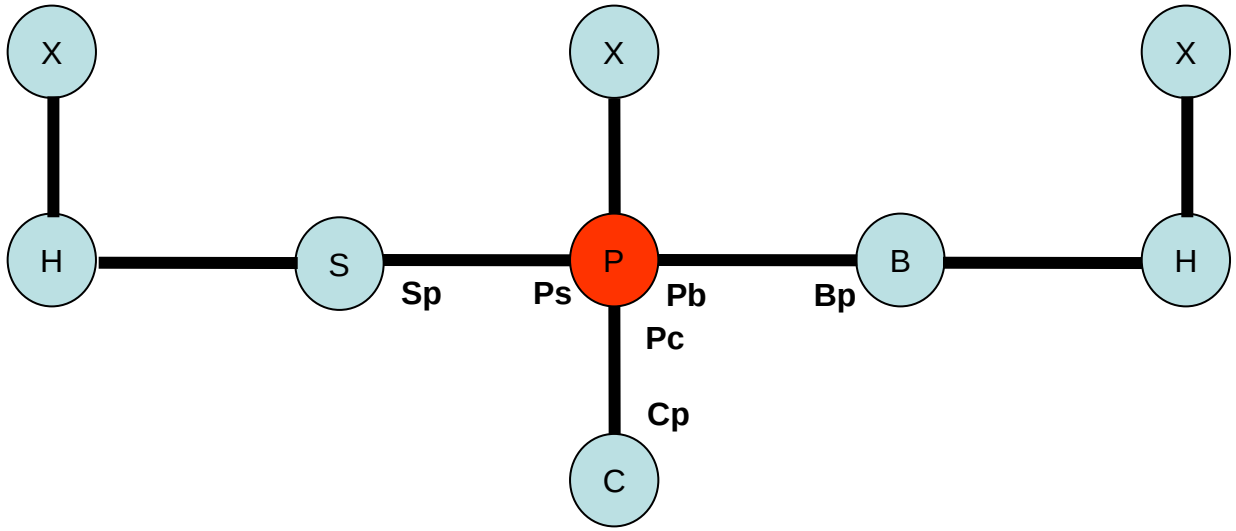

When P fails, X becomes unreachable through P

To calculate the repair strategy the neighbours of P

1. **Remove X from P, and run incremental SPF until X is reattached.**
2. **Look at the next hop to the new home (H) of X**
3. **If it is not P, encap to H.**
4. **If it is P, put on the repair path to H.**

Multi-homed Prefixes

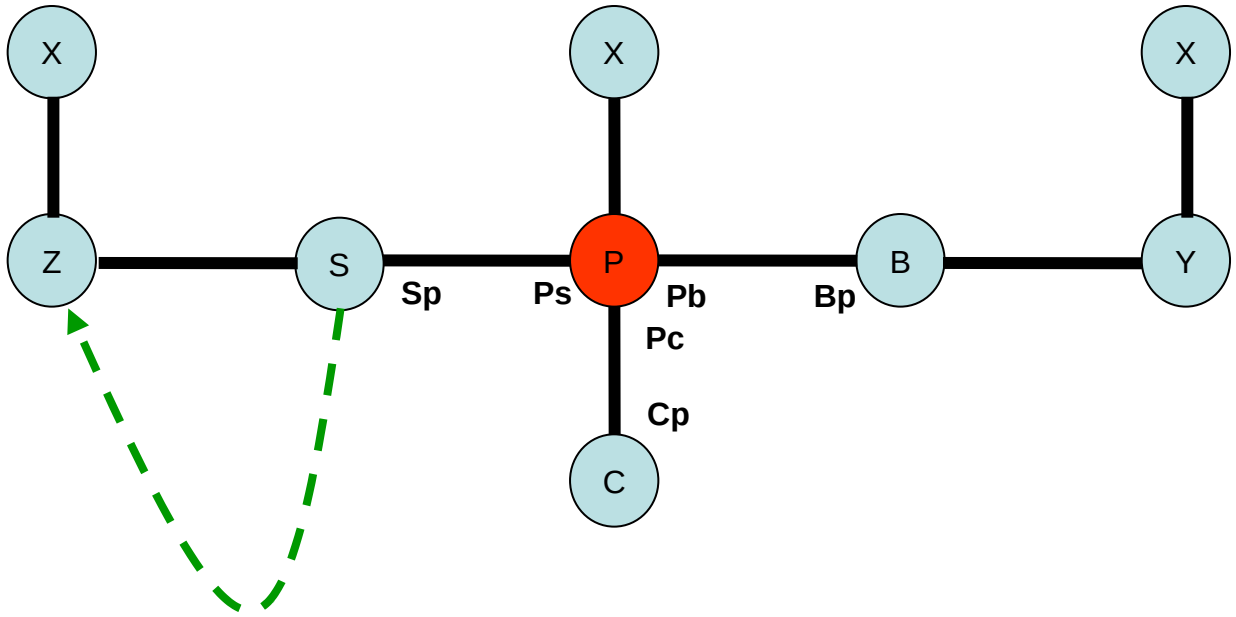

Encap to Z
Z *must* forward to X

Case when Z is the closest reachable attachment point of X after the failure

Multi-homed Prefixes

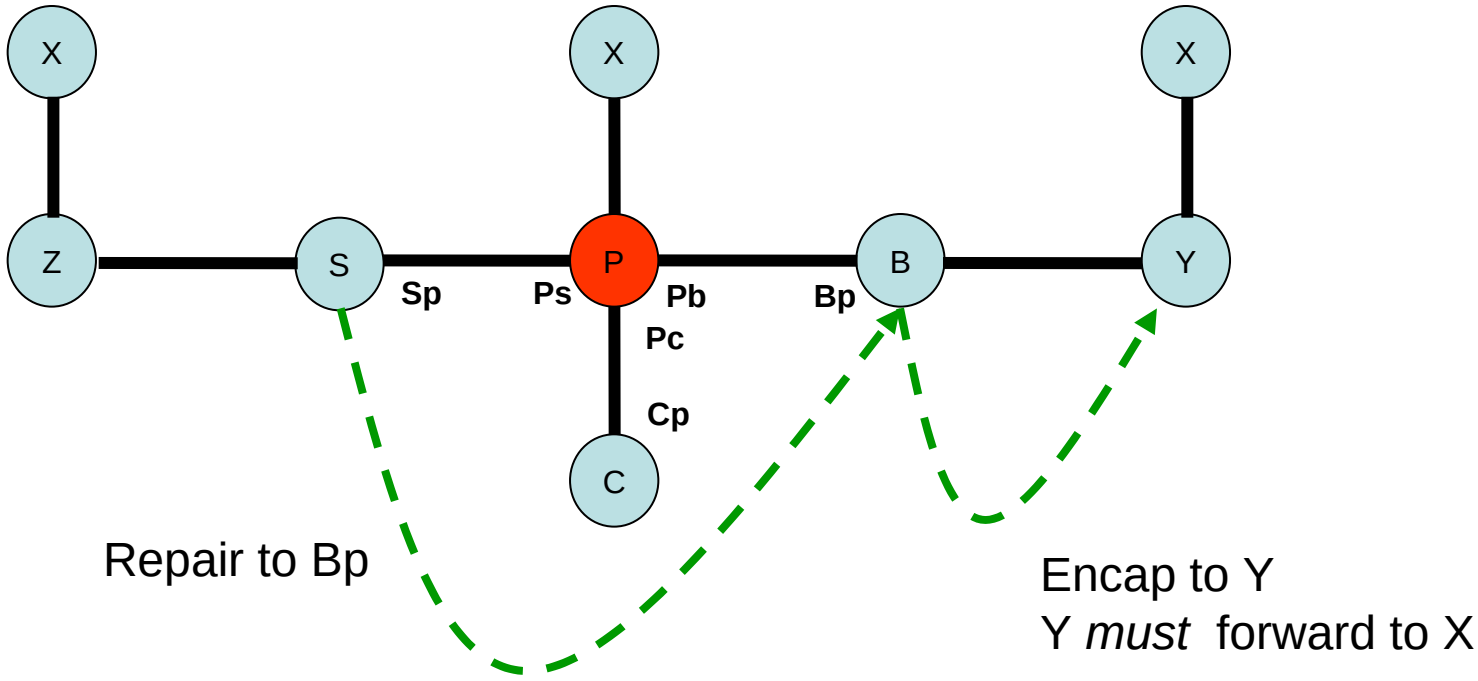

Case when Y is the closest reachable attachment point of X after the failure.

Note – Only one level of encapsulation is needed.

Summary

- A intuitive approach that has 100% coverage of non-partitioning faults.
- Uses existing MPLS FRR hardware, or single level IP encapsulation.
- Repair path compute time is bounded, and comparable with other proposed solutions.