Fast Reroute for IP and LDP based Networks

- By Albert J. Tian and Naming Shen

MPLS 2004

Redback NETWORKS

Need for IP/LDP Fast Reroute

- IP/LDP network is no longer best effort
  - Mission Critical Application
  - Delay sensitive services such as VoIP, Video over Broadband, Video on Demand, Pseudo Wire etc
  - L3VPN, L2VPN, and L1VPN mainly use IP/LDP in the provider backbone
- Fast Reroute can quickly detect and reroute traffic around failure
- Goal is to achieve under 50ms repair time when a network element fails

Existing Solutions insufficient for IP/LDP

- Fast reroute for RSVP-TE LSP
  - Protect RSVP-TE LSP traffic
  - 50ms repair time
  - Difficult to protect IP traffic: need full mesh edge to edge TE LSPs
- IGP/LDP fast convergence
  - Minimize packet loss for IP/LDP traffic
  - Sub-second repair time, hard to achieve under 50ms

Desirable Attributes for IP/LDP Fast Reroute

- Repair time within 50ms
- Protect variety of traffic types: IP unicast, IP multicast, LDP, RSVP-TE, etc
- 100% repair coverage regardless of topology
- Work across area/level/domain boundary
- Guaranteed repair time regardless of network topology and size
- Solution complexity do not increase with network size

Three Questions on Fast Reroute

- Repair Path is key concept for fast reroute
- Q1: Where to terminate the repair paths?
- Q2: How to calculate the repair paths?
- Q3: How to implement the repair paths?

Answers: Nexthop Fast Reroute (NHFRR) and Alternative Shortest Path (ASP)

- Termination: Next-Hop Fast Reroute:
  - terminate at nexthop or next-nexthop
    <draft-shen-nhop-fastreroute-01.txt>
- Calculation:
  - Alternative Shortest Path: exclude the link/node being protected and re-calculate SPF
    <draft-tian-frr-alt-shortest-path-01.txt>
  - Full CSPF with QoS constraints
- Implementation:
  - RSVP-TE
  - (or, a stack of LDP labels)

Next-Hop Fast Reroute
Link Protection

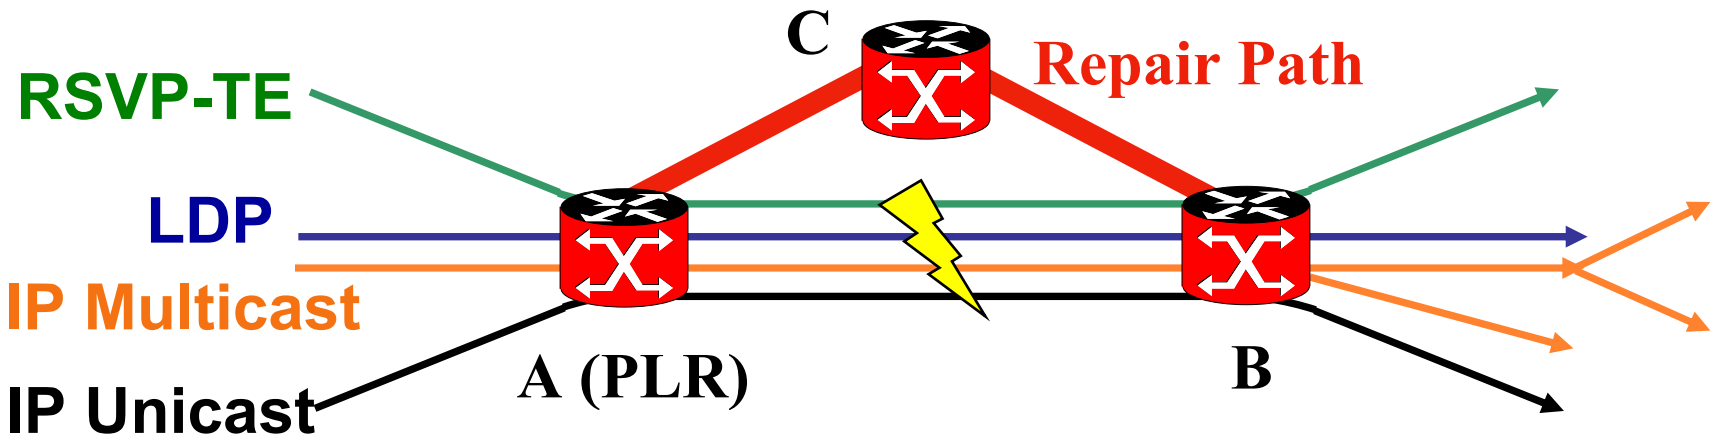

- repair path terminate on nexthop B
- can protect many types of traffic on link A-B

Big Picture for Link Protection

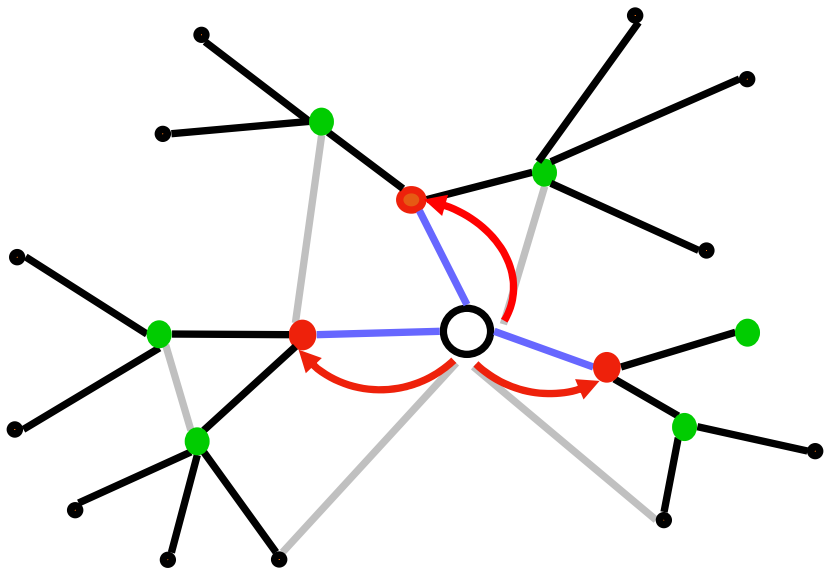

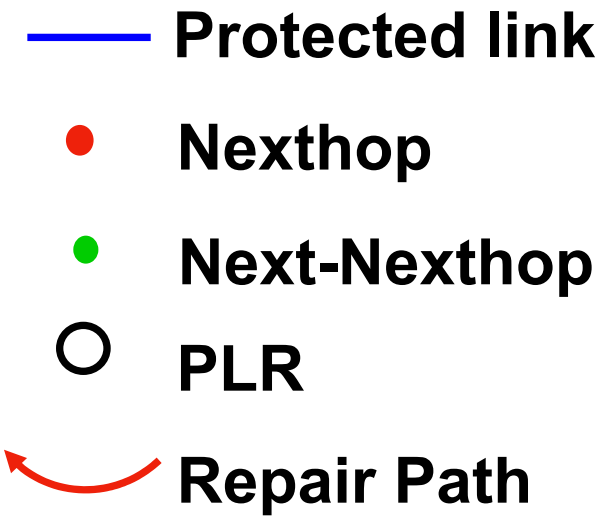

- Black is shortest path tree from SPF
- Minimum one repair path for each protected link
- Need 3 repair paths to cover 3 links

Next-Hop Fast Reroute
Node Protection

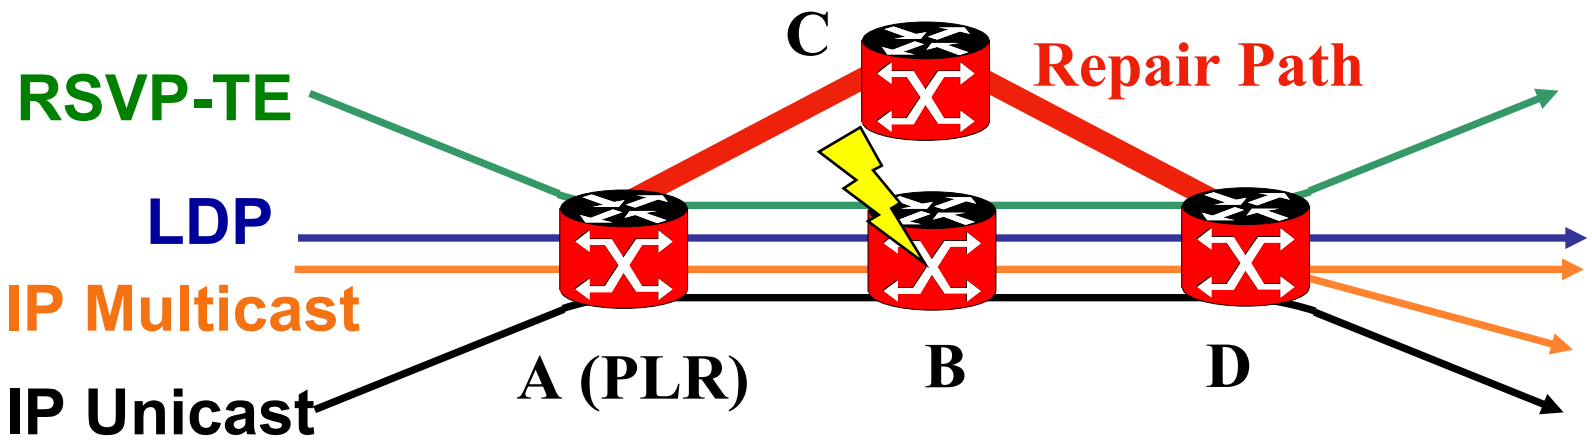

- repair paths terminating on next-nexthops
- Since each nexthop may have multiple next-nexthops, may need multiple repair paths to cover all traffic going through one nexthop

Big Picture for Node Protection

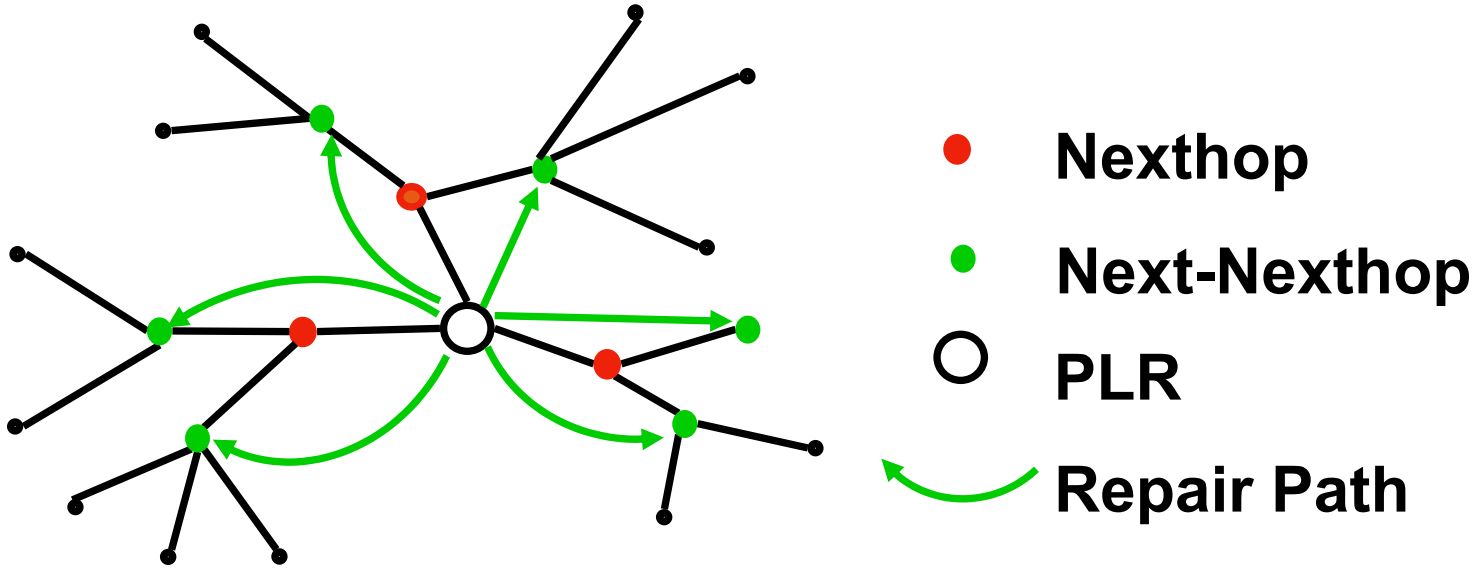

- Black is the shortest path tree as result of SPF
- Minimum one repair path for each next-nexthop
- Need 6 repair paths to cover 3 nexthops

Node Protection for LDP Traffic

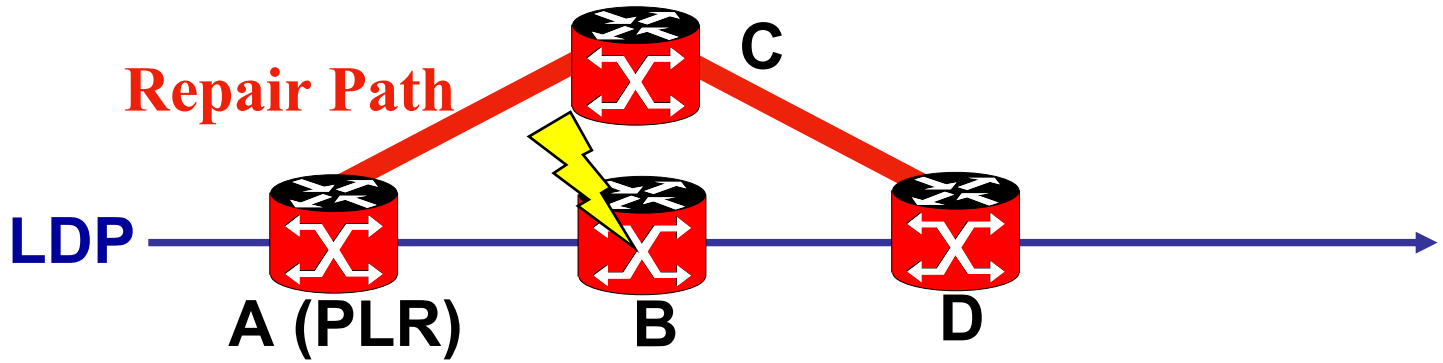

- Rerouted traffic is tunneled to next-nexthop D
- PLR A needs to know next-nexthop D's label
- LDP extension to learn next-nexthop label
  <draft-shen-mpls-ldp-nnhop-label-01.txt>
- LDP targeted neighbor to all next-nexthops

Node Protection for Multicast

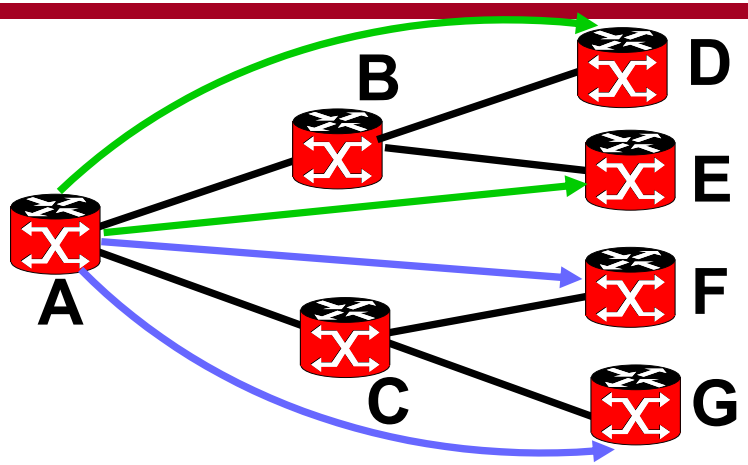

- Repair paths to all downstream next-nexthops
- Replicate onto all repair paths
- Rerouted traffic has to be tunneled (encapsulated) all the way along the repair paths

Finding Next-Nexthops for Multicast

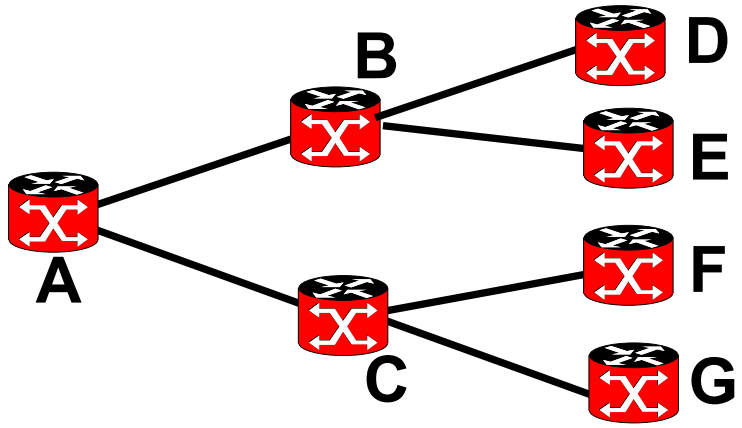

- Normally multicast routing protocols only know immediate downstream neighbors that are interested in a group
- Extensions are needed to learn downstream next-nexthop neighbors that are interested in a group: <draft-shen-pim-nnhop-nodes-01.txt>

Inter-Area Node Protection

- when a border router leaks routes into another area/level, it can optionally attach the nexthop information
- the nexthop information from other areas can be used to setup repair paths across areas/levels
- <draft-shen-isis-interarea-route-info-00.txt>

Repair Path Calculation

- A repair path is an explicit path with a constraint that it can not go through the link or node that is being protected.
- A simple solution - Alternative Shortest Path: exclude the link or node that is being protected and re-calculate SPF (can deal with SRLG by excluding all elements in the same SRLG and re-calculate SPF, so no additional complexity)
- Full CSPF
  - Can take QoS parameters and other policies into account
  - Can produce multiple repair paths to do load sharing

Complexity of Alternative Shortest Path

- Maximum N SPF computations for PLR, where N is the number of nexthops
- Link Protection:
  - Each SPF is calculated excluding a link being protected (and excluding all links in the same SRLG)
- Node Protection:
  - Each SPF is calculated excluding a nexthop being protected

Repair Path Implementation – RSVP-TE

- PLR signal the repair paths using RSVP-TE
- Assuming there are N nexthops, H next-nexthops
- For Link Protection, Maximum N repair paths need to be signaled
- For Node Protection, Maximum H repair paths need to be signaled
- In reality, N and H are small
- Simple implementations in the data plane as repair paths are associated with nexthops or next-nexthops, not associated with each prefix

Repair Path Implementation – Stack of LDP labels

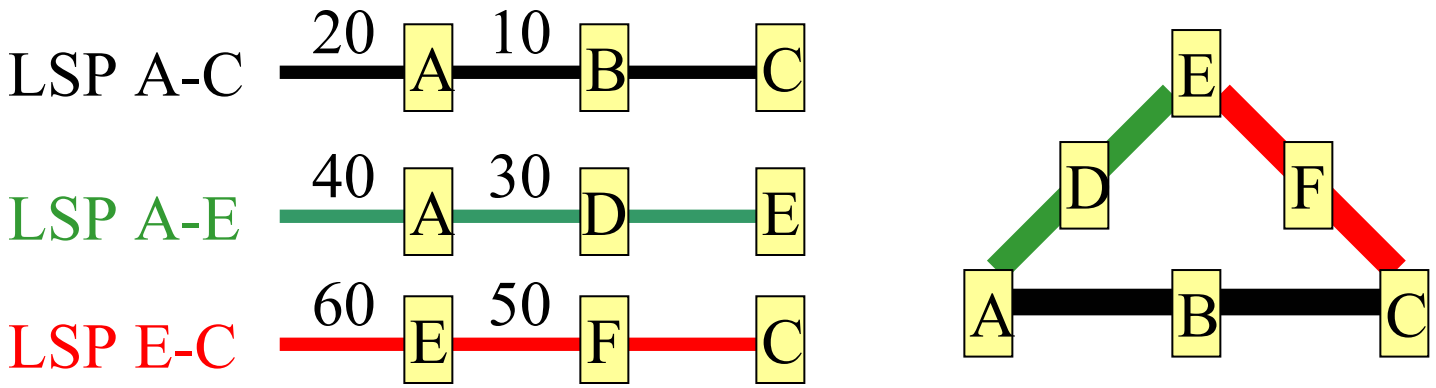

- A stack of LDP labels 40/60 can be used to force an incoming packet on A to go A-E-C path
- LSR A needs to know E's label (60) for LSP E-C
- Need targeted LDP sessions to learn remote labels
- In reality, max 2 additional labels to protect unicast traffic, max 3 additional labels to protect multicast traffic

Summary

- Next-Hop Fast Reroute and Alternative Shortest Path are very simple and intuitive
- 100% repair coverage regardless of topology (*)
- Maintain small number of repair paths, regardless of network size, great scalability
- Repair paths associated with nexthops or next-nexthops
- The only solution that covers multicast so far
- Using RSVP-TE, which is mature technology
- The only solution that can take QoS and other policies into consideration, great flexibility
- Can easily handle SRLG
- Uniform solution protect all traffic types

Extremely Appealing for Link Protection

- With non-stop forwarding only link protection is needed
- Protect all types of traffic
- No extensions needed for LDP
- No extensions needed for PIM
- No inter-area/inter-domain issues
- Simplicity in data plane with per nexthop repair path

Micro Loop Prevention

- Could happen when routers converge in the wrong order
- Ordered convergence would solve the problem
- Orthogonal to Fast Reroute solutions

Compare to Loop-Free Alternative

- <draft-ietf-rtgwg-ipfrr-spec-base-00.txt>
- Computation wise the same as Alternative Shortest Path (without loose segment optimization)
- Loop-free Alternative has limited coverage for unicast
- Loop-free does not cover multicast
- Loop-free uses per prefix (or per egress) repair path in the data plane, where as nexthop FRR uses per nexthop or per next-nexthop repair path
- Loop-free does not need tunnels

Compare to U-Turn

- U-Turn does not have full coverage for unicast
- U-Turn can not protect multicast traffic at the moment
- U-Turn has higher computational complexity
- U-Turn uses per prefix (or per egress) repair path in the data plane, where as nexthop FRR uses per nexthop or per next-nexthop repair path
- U-Turn does not need tunnels