INTEC

MPLS Recovery

Didier COLLE

Pim VAN HEUVEN

Adelbert GROEBBENS

Chris DEVELDER

Mario PICKAVET

Piet DEMEESTER

MPLS recovery: single layer

- Introduction to:
  - MPLS and MPλS technologies
  - MPLS Recovery techniques:
    - Study of IETF proposals
    - Development of FTCR scheme
- Porting MPLS recovery to MPλS
- Spare resource dimensioning

MPLS and MPλS

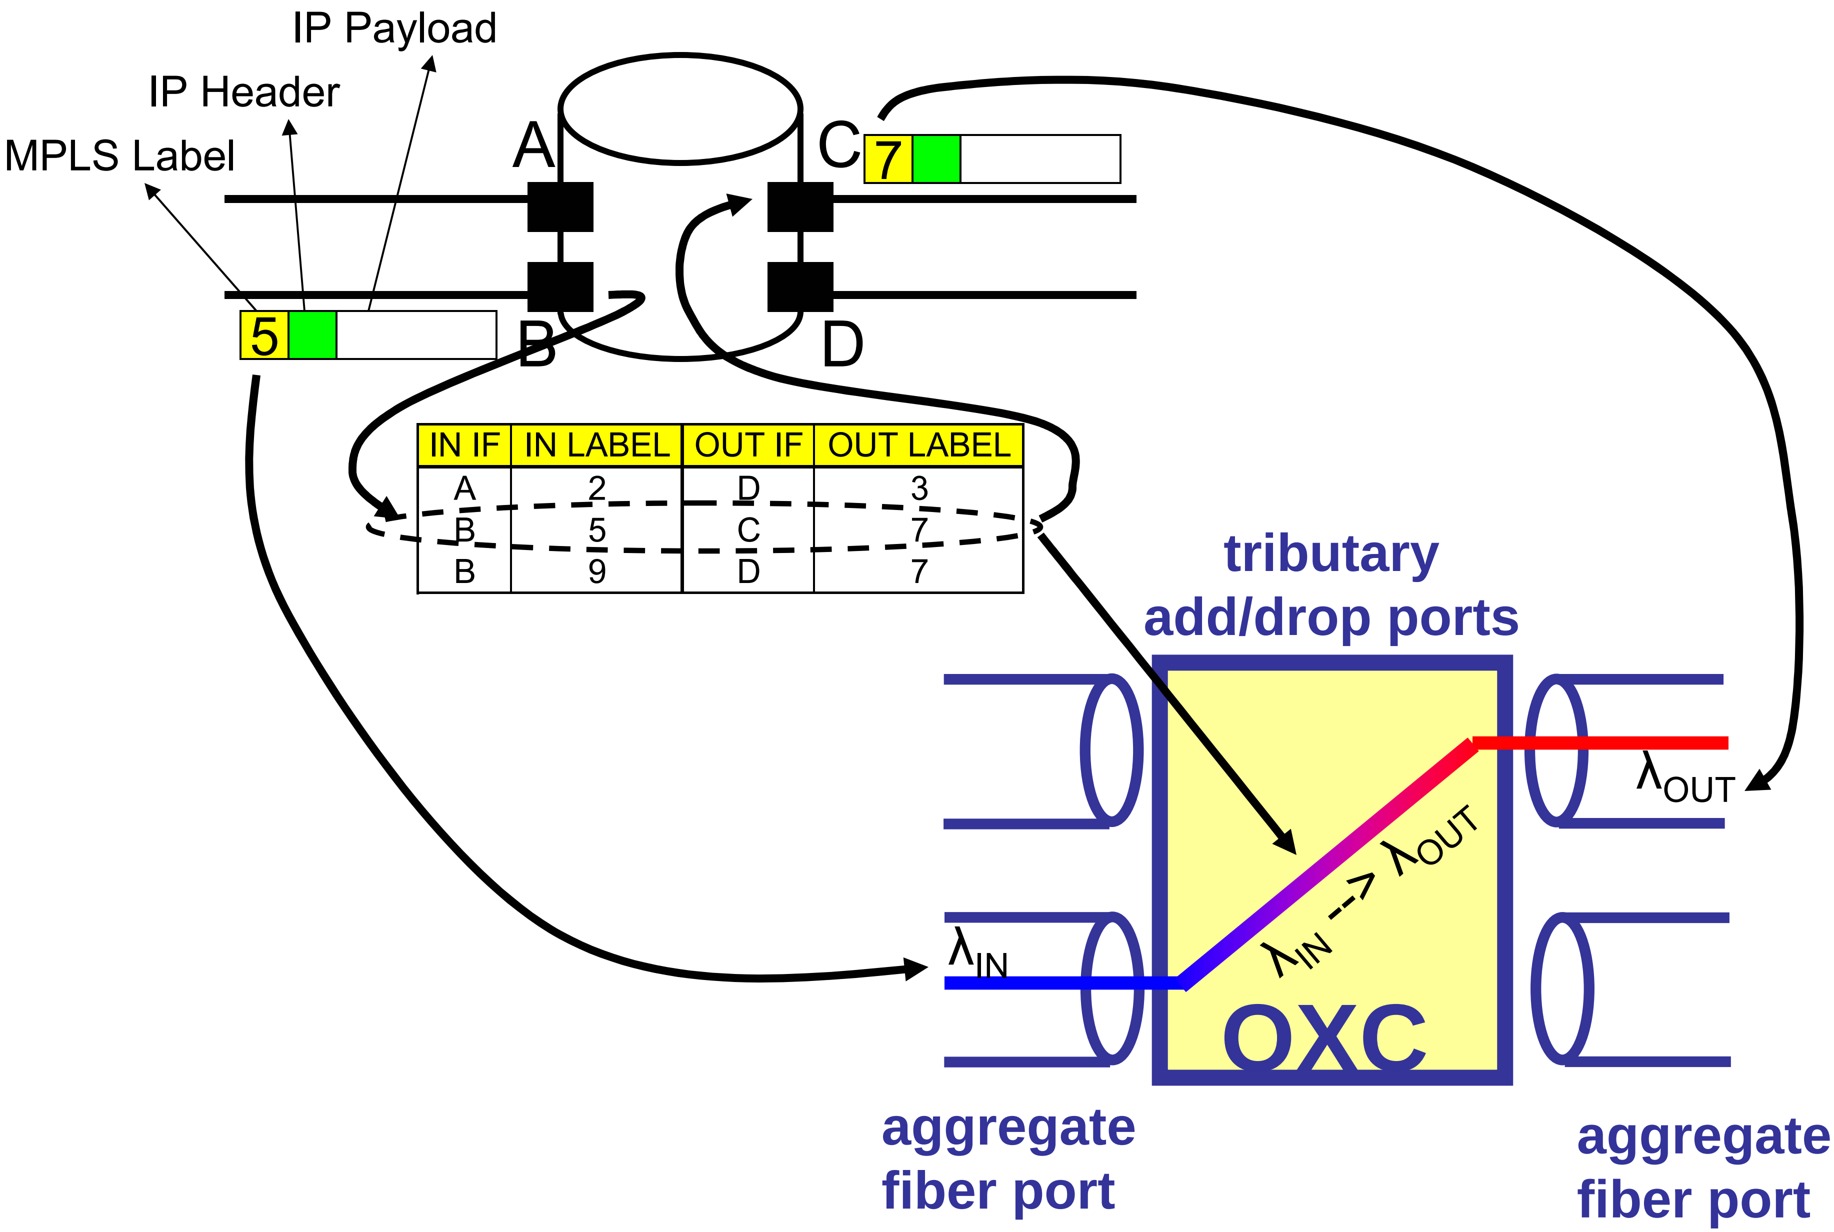

IN IF	IN LABEL	OUT IF	OUT LABEL
A	2	D	3
B	5	C	7
B	9	D	7

MPLS protection

- Pre-establish backup LSP
  - Protected segment:
    - local (link or node)
    - subnetwork
    - end-to-end
  - Upstream: Protection Switch LSR (PSL)
    - protection switching
  - Downstream: Protection Merge LSR (PML)
    - no protection switching, but merging

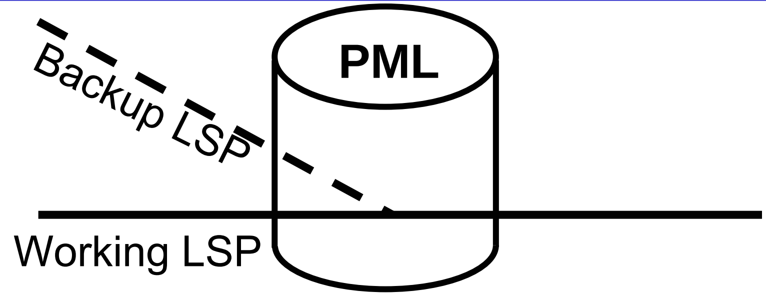

IN IF	IN LABEL	OUT IF	OUT LABEL
A	1	C	3
B	2	C	3

– **End-to-end Prot**

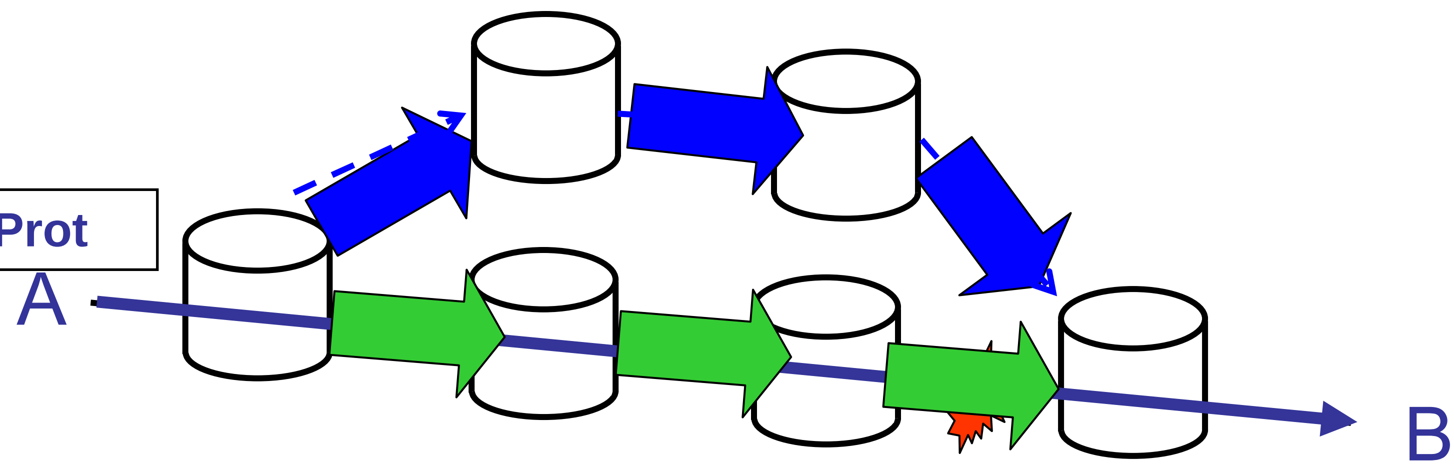

MPLS protection: local loop-back

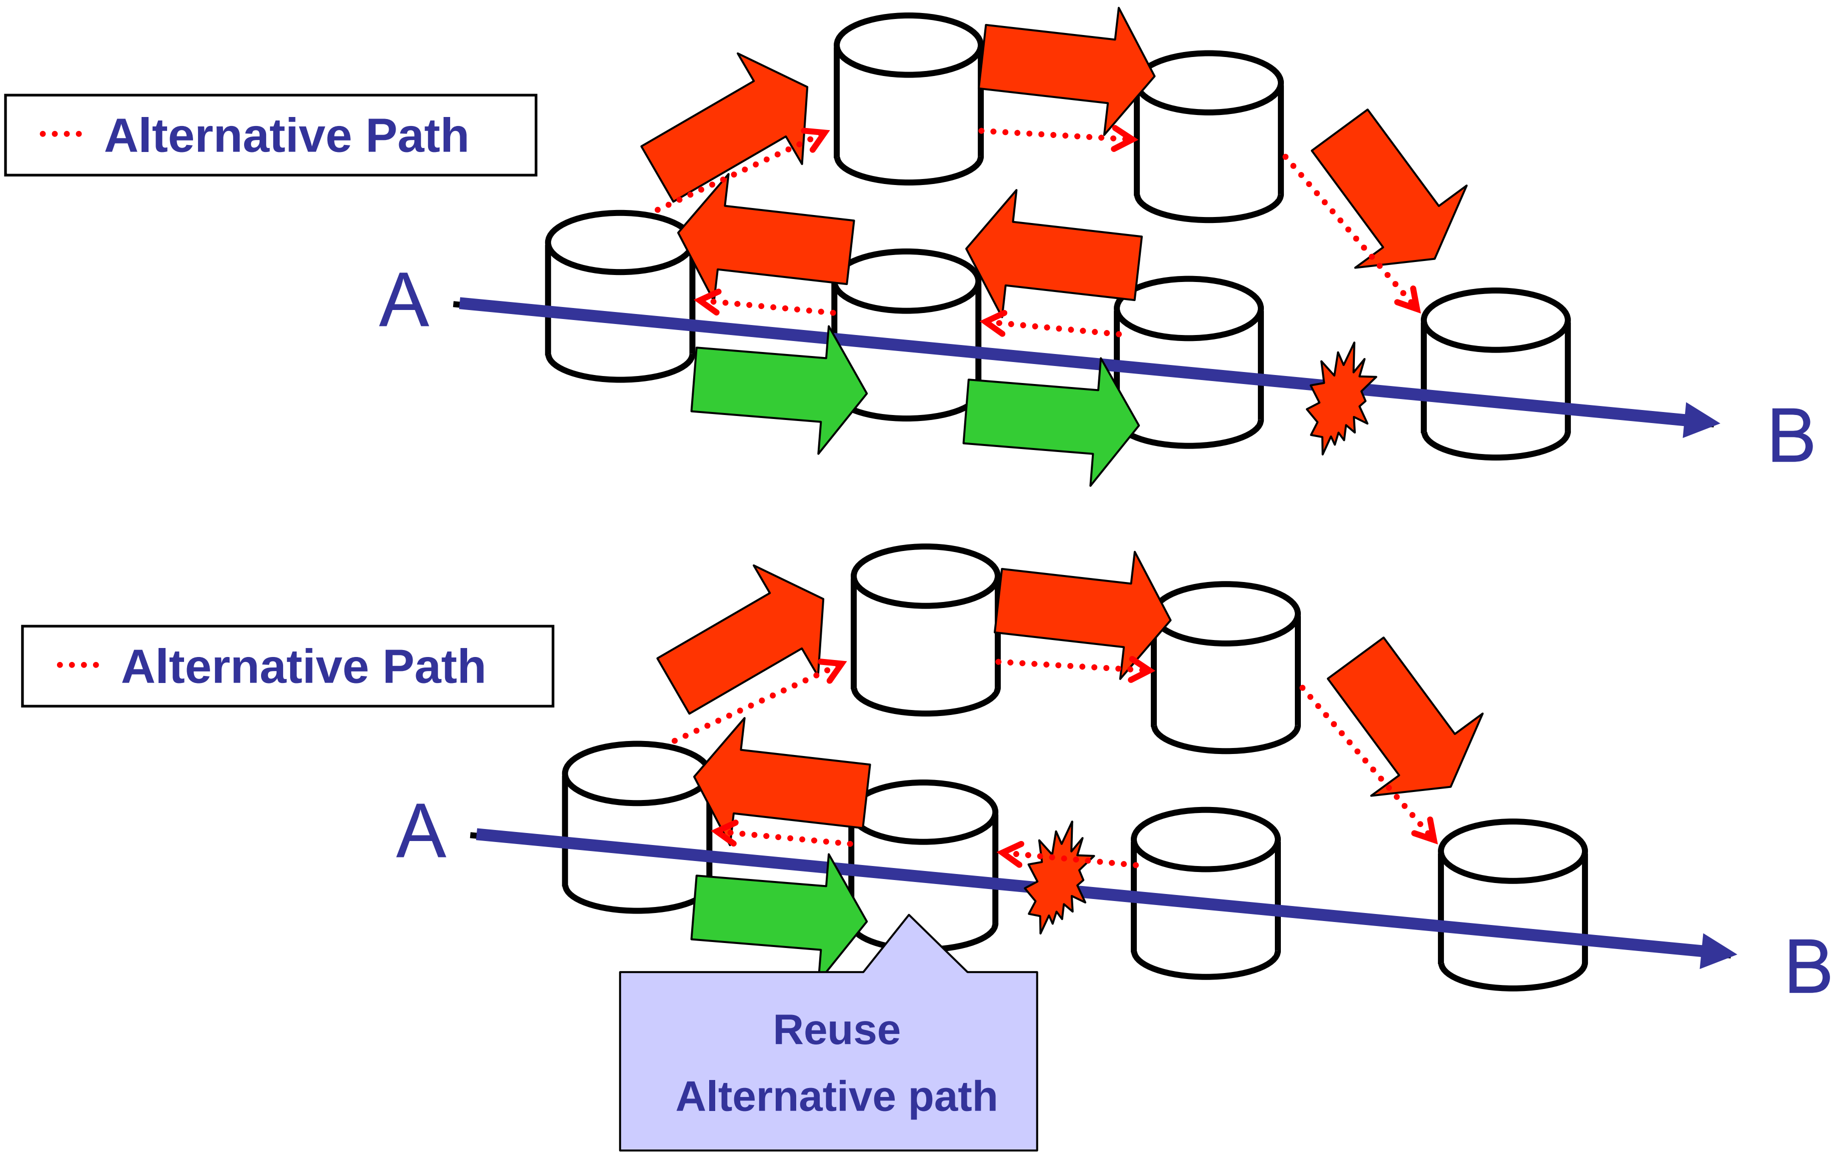

Re-routing in MPLS

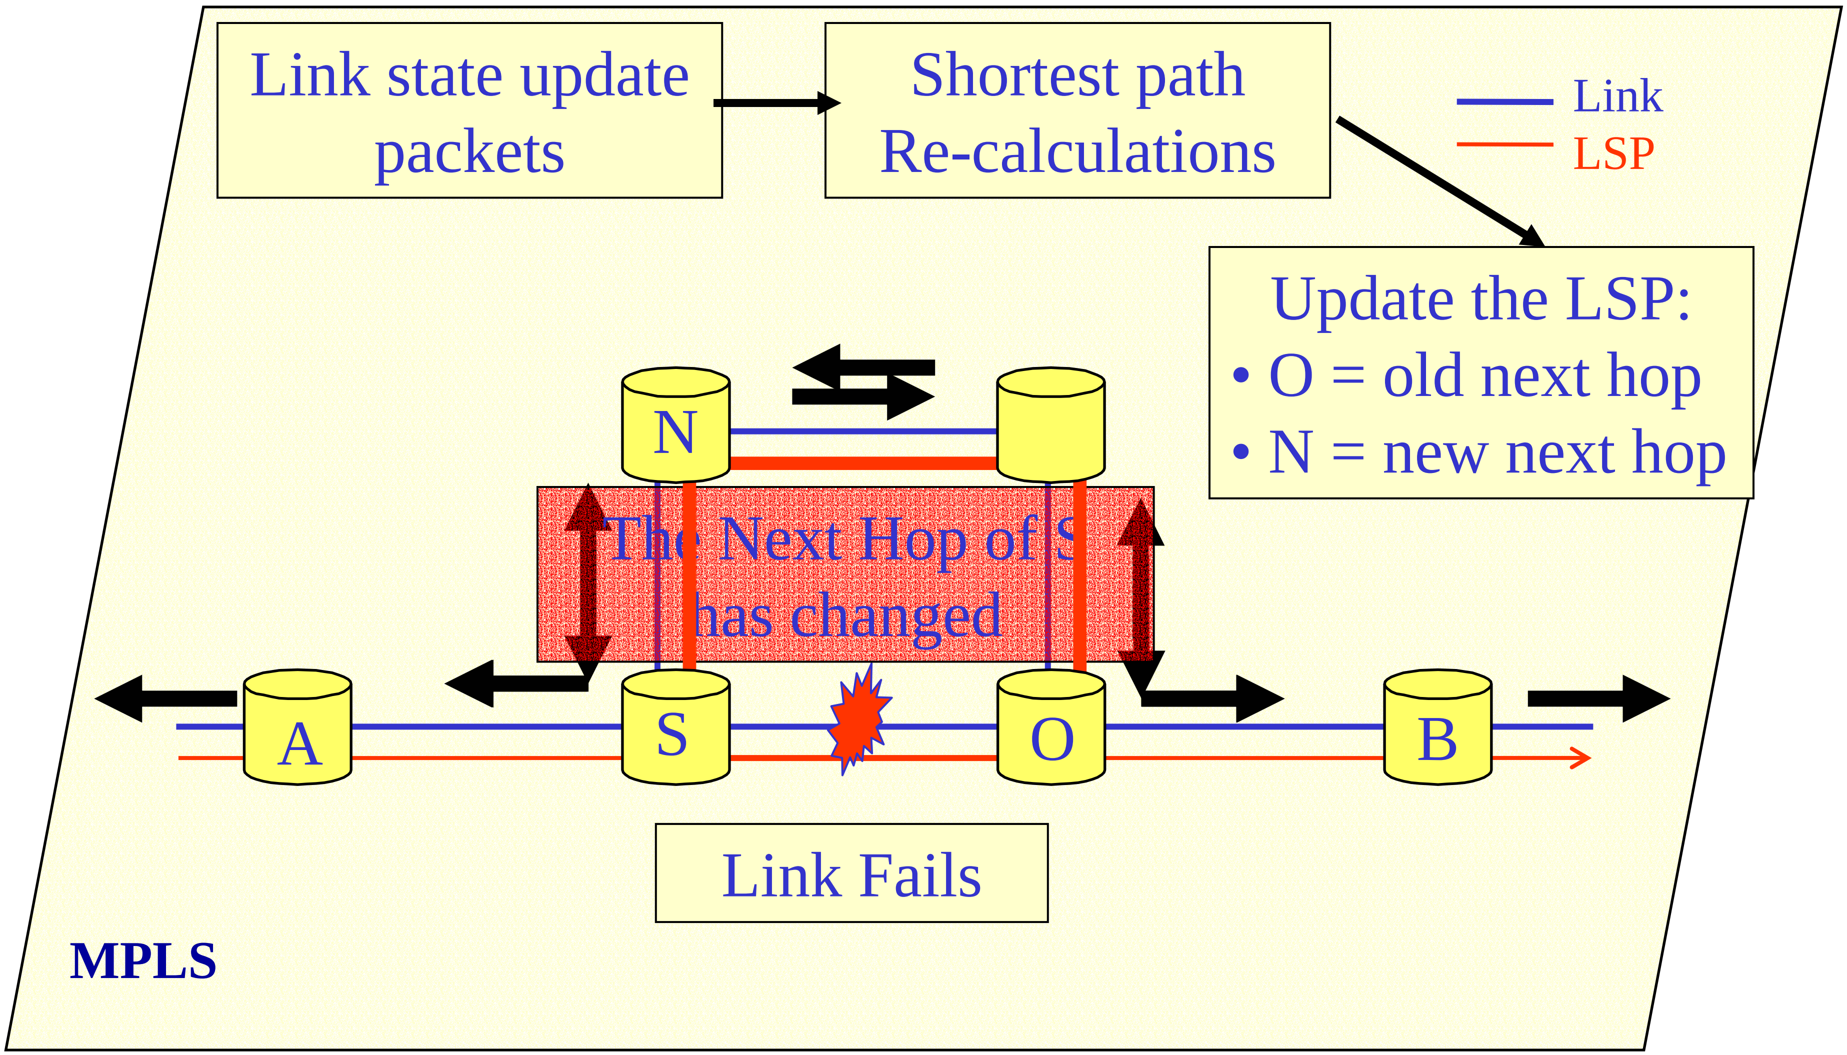

FTCR: Fast Topology-driven Constraint-based Rerouting

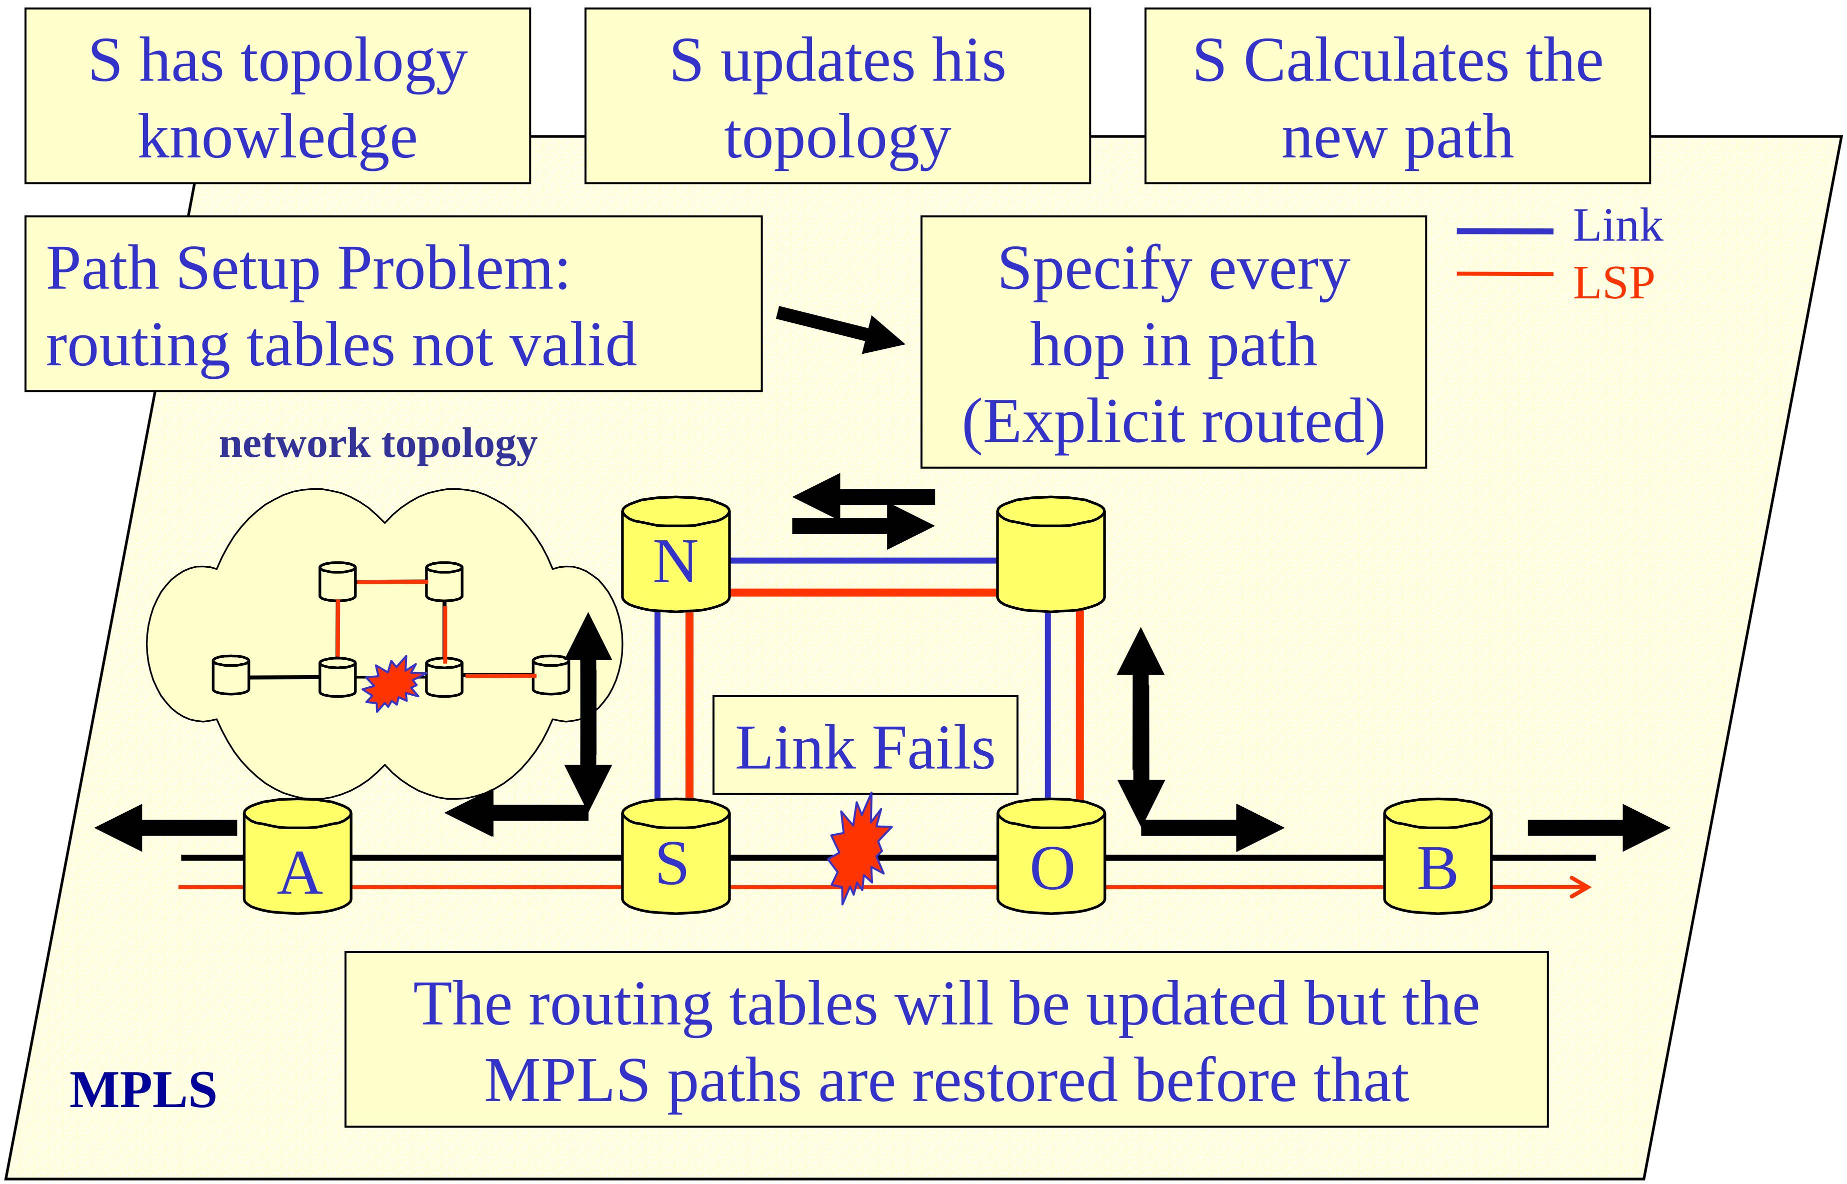

Porting MPLS protection to MPλS

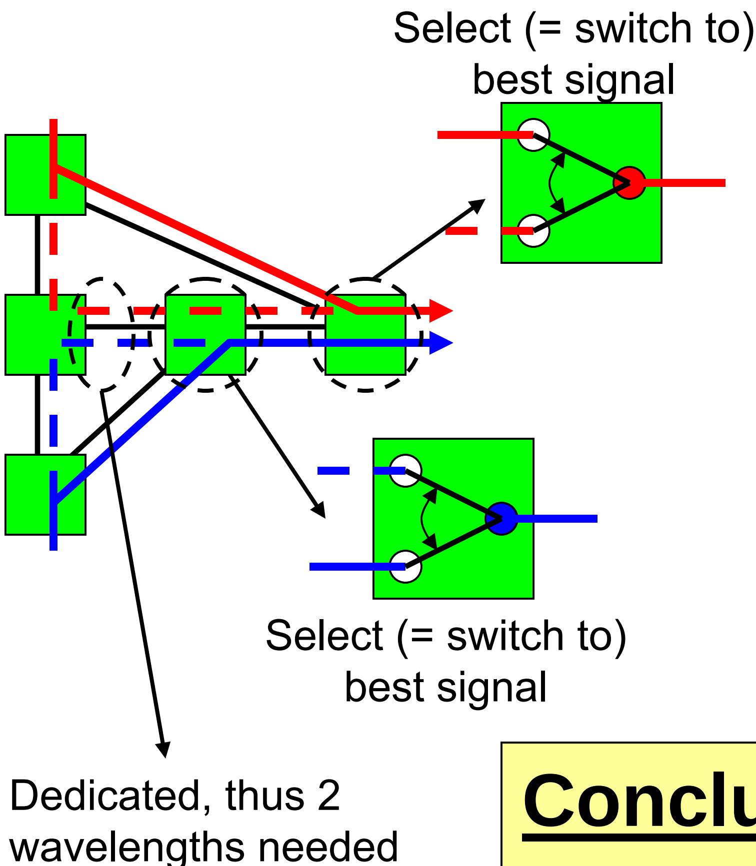

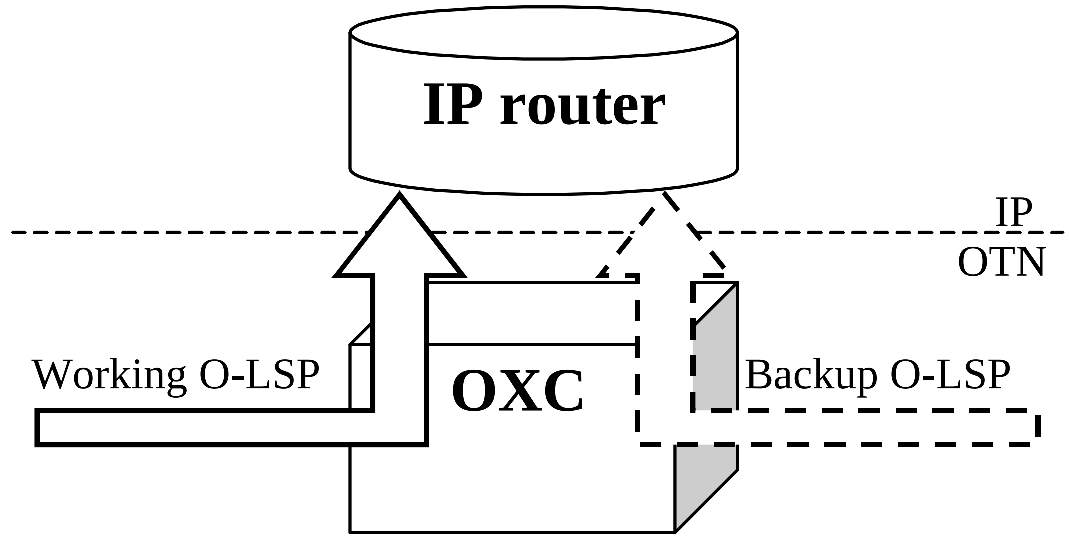

Conclusions:

- Dedicated protection
- Merging problem
  - solve by simulating with passive selector/switch
  - shift merging to client (i.e., IP layer).

Simulations: assumptions

- Single layer planning
  - MPLS recovery techniques
  - MPλS recovery techniques
- Routing:
  - shortest path
  - each LSP independent
- Capacity/cost model
  - linear capacity model: line capacity = used capacity
  - cost model: cost to carry unit of capacity proportional with link weight (roughly estimated on distance).
- Traffic matrices: asymmetric
- Random generation (e.g., traffic):
  - set of 10 instances
- MPλS: wavelength conversion assumed

Results: Optical versus Electrical Recovery

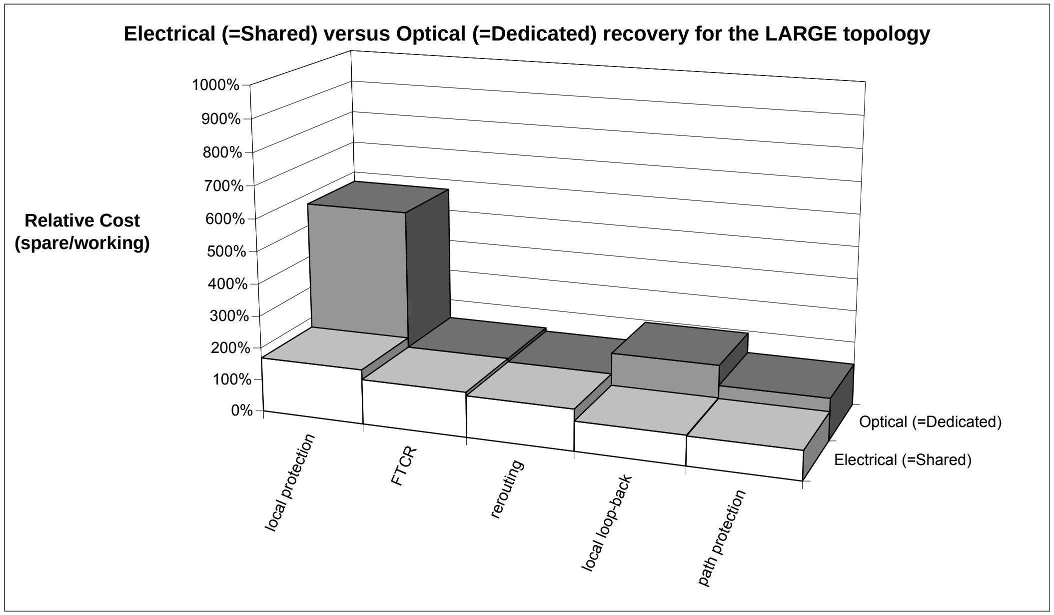

Failure scenarios:
- single link failures (interpreted as a node failure by adjacent LSRs, except for rerouting)
- single node failures

Traffic:
- Uniform pattern
- Randomly generated (integer values)

Last link (of an LSP):
- Protected
- Not reverted (for local loop-back)

Topologies
- Large: 57 links and 44 nodes
- Small: 36 links and 30 nodes

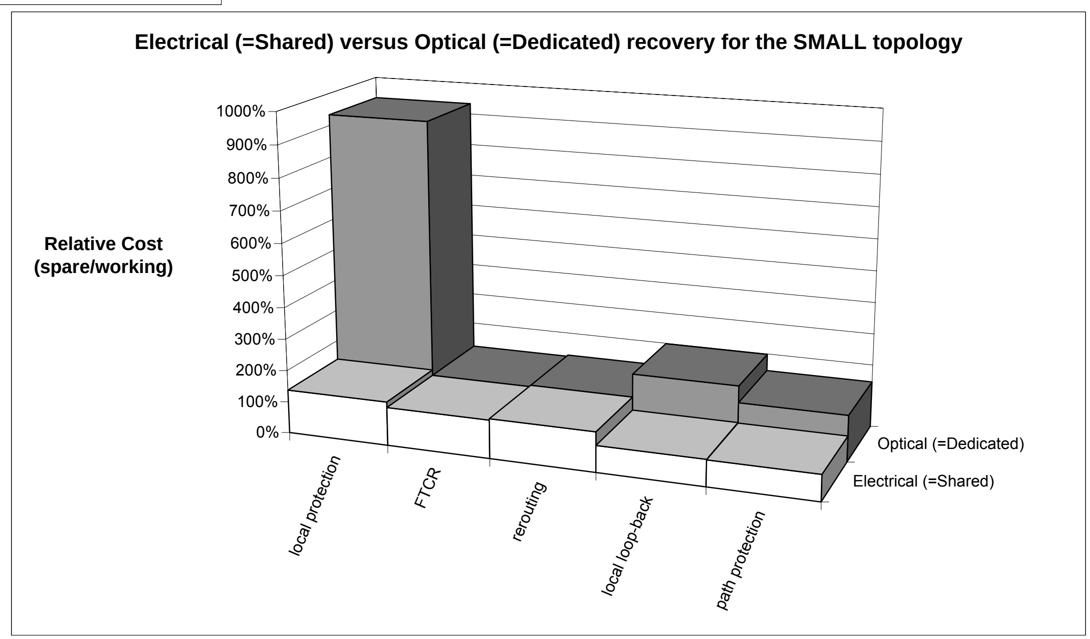

Results: Optical versus Electrical Recovery

- Rerouting and FTCR: no difference
  - When tearing down part of primary LSP downstream of the failure
- Worst case: dedicated versus shared protection
  - No merging possible (eventually simulating merging via switching)
  - Label is scarce product in MPλS, instead of bandwidth in MPLS
  - How to improve this worst case --> see next slides
- Dedicated effect:
  - significant for end-to-end protection or local loop-back
  - does not allow sharing between both direction for local loop-back
  - catastrophe for local protection

Results: Electrical MPLS Recovery

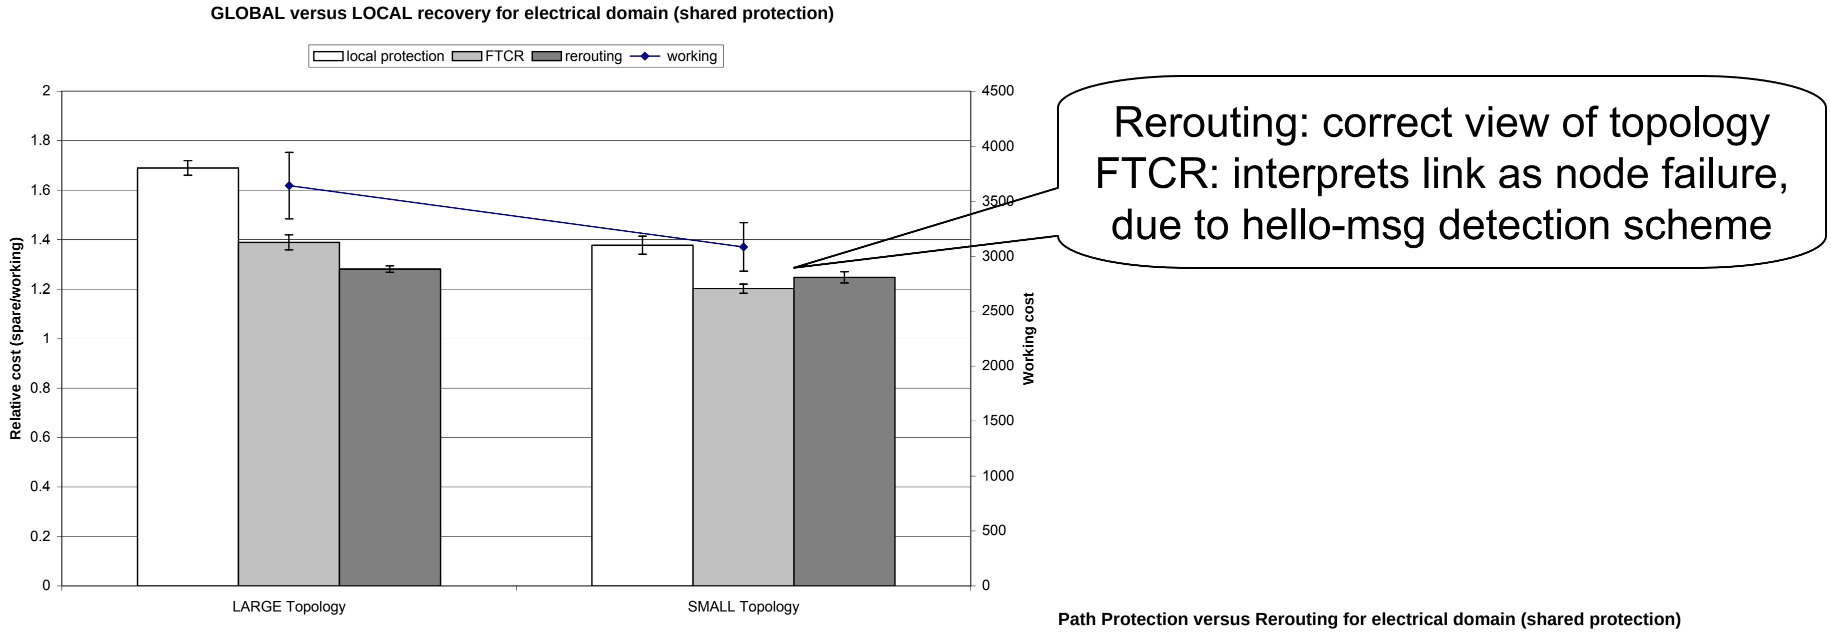

Failure scenarios:
- single link failures (interpreted as a node failure by adjacent LSRs, except for rerouting)
- single node failures

Traffic:
- Uniform pattern
- Randomly generated (integer values)

Last link (of an LSP):
- Protected
- Not reverted (for local loop-back)

Topologies
- Large: 57 links and 44 nodes
- Small: 36 links and 30 nodes

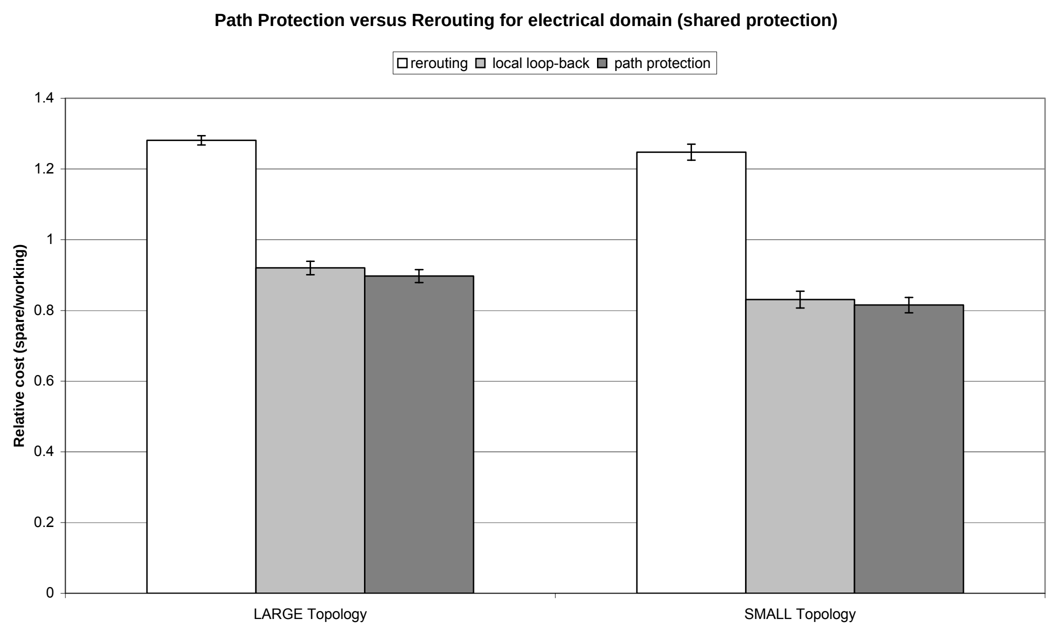

Results: Electrical MPLS Recovery

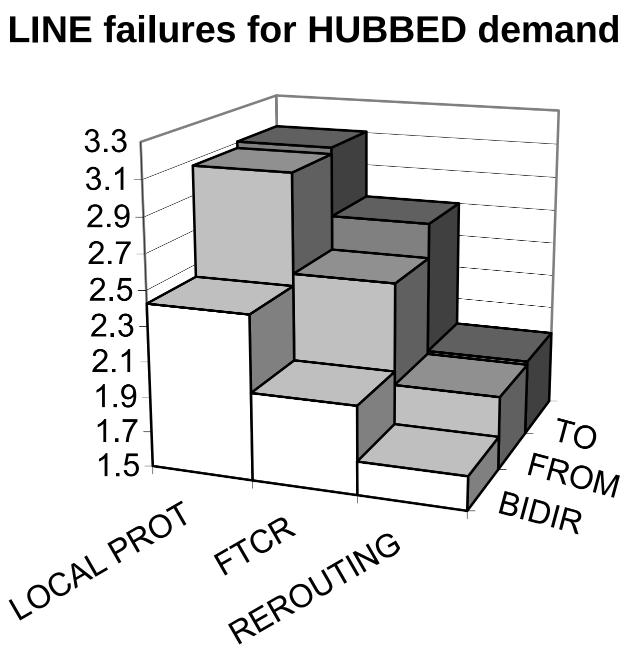

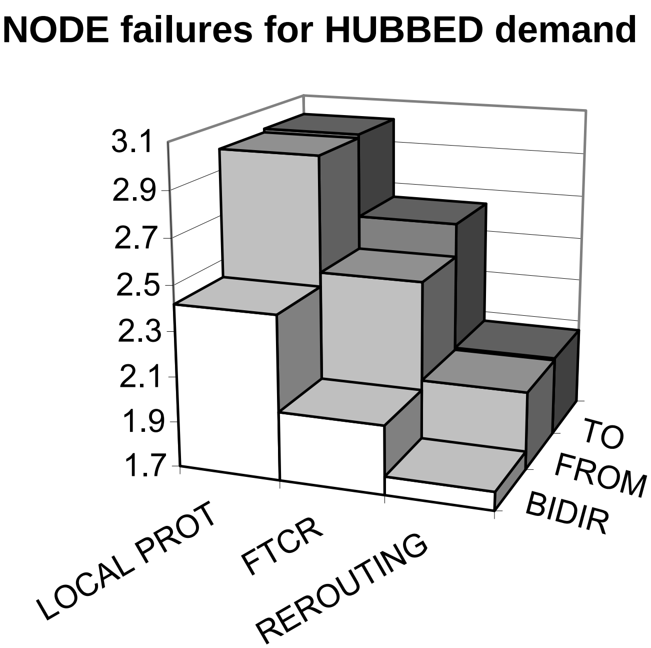

Failure scenarios:
- single link (left) OR node (right) failures
- --> link failures always interpreted as link failures

Traffic:
- pattern:
  - uniform: “bidir”
  - hubbed: “from” or “to” single node
- Randomly generated (integer values)

Topologies
- Large: 57 links and 44 nodes

Why hubbed/star traffic pattern?
- European backbone: gateway to USA
- Residential ISPs
- Traffic to/from a server farm
- Etc.

Results: Electrical MPLS Recovery

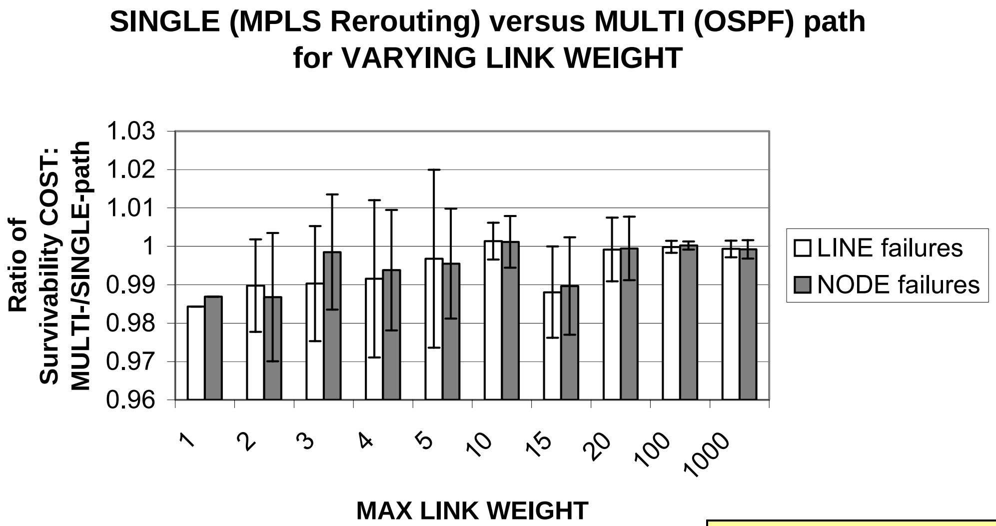

Failure scenarios:
• single link (white) OR node (gray) failures
• --> link failures always interpreted as link failure
Traffic:
• pattern: single, uniform traffic matrix
Topologies
• Large: 57 links and 44 nodes
• Link weights: randomly generated

<u>Single Path</u>
• ***MPLS Rerouting***: single LSP between two nodes, restored by another single LSP
<u>Multi Path</u>
• ***OSPF***: forward packets evenly over all interfaces which have same distance to destination
• ***MPLS rerouting***: consider multiple equal cost LSPs (each to be rerouted!) --> scalability problem!

Results: Electrical MPLS

- Local Protection > FTCR > End-to-end:
  - FTCR is a combination of Local Protection and End-to-end
- End-to-end:
  - Rerouting > end-to-end protection or local-loop back:
    - protection --> less alternative routes --> potentially less spare resources
  - End-to-end protection = +/- Local loop-back:
    - downstream no traffic anymore --> place for local loop-back of opposite direction
- Hubbed Traffic pattern:
  - FTCR performs significantly better for traffic from the hub than for traffic to the hub.
- Single (MPLS Rerouting) versus multipath (e.g., OSPF)
  - Working cost identical
  - Decreasing maximum link weights
    - Multipath seems to perform slightly better
    - But also higher variance on multi/single path ratio.

Sharing in MPλS: local protection

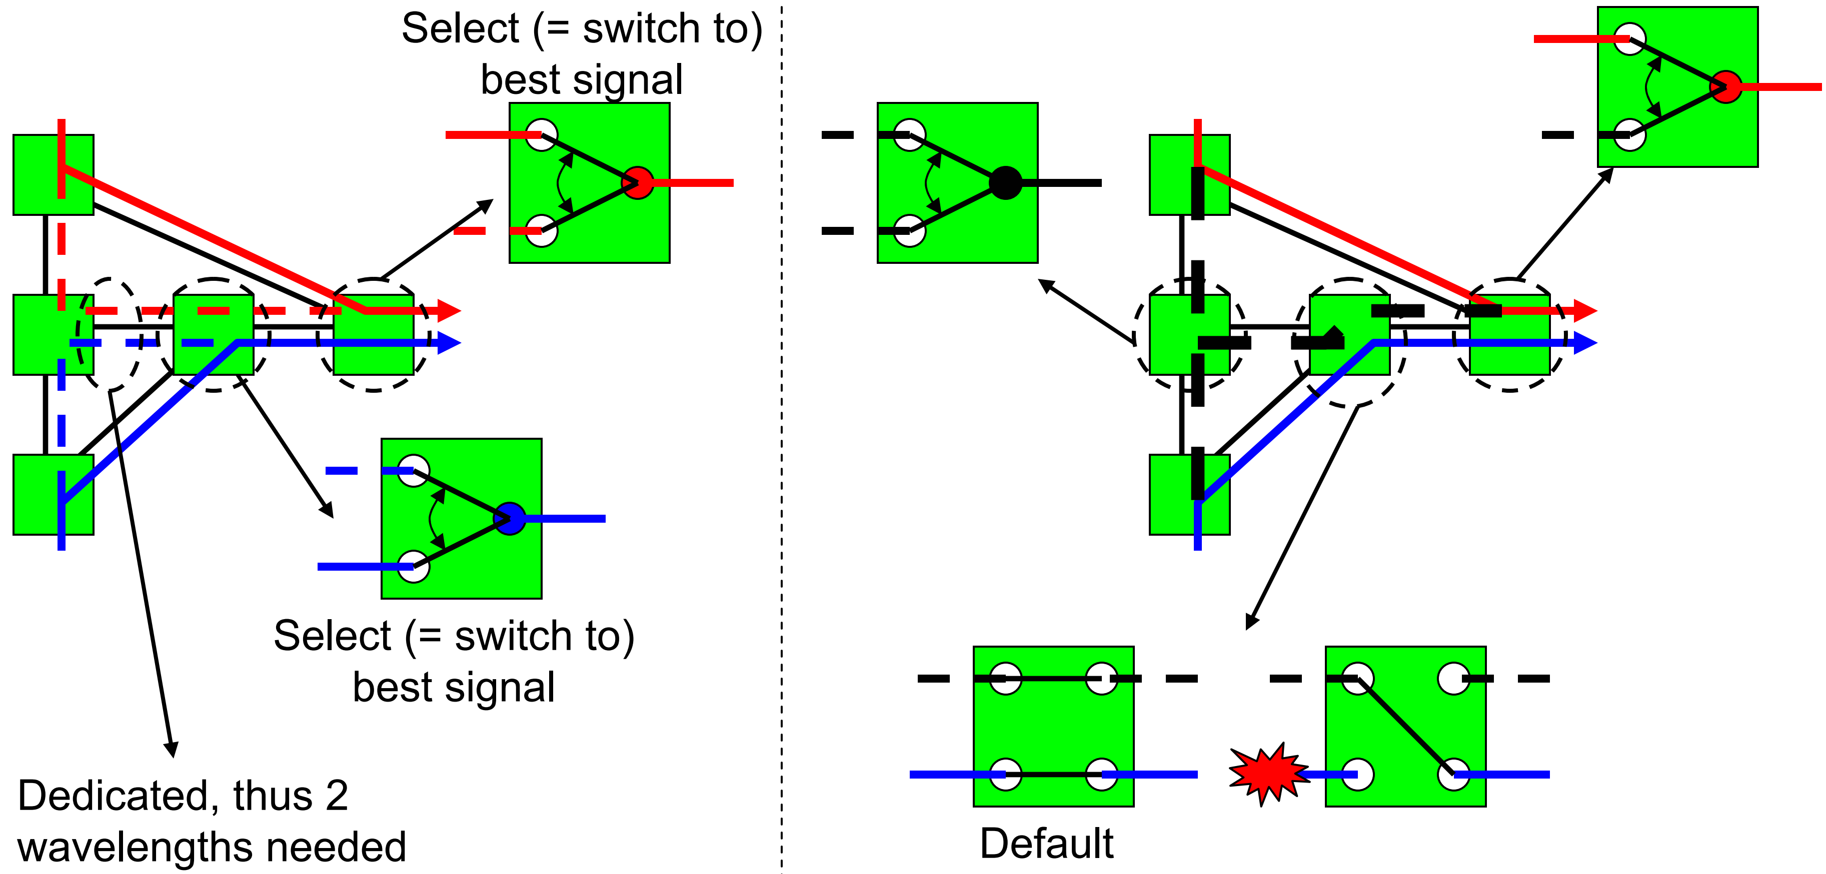

Converging backup Tree:
AT MOST single output
wavelength!!!

Sharing in MPλS: path protection

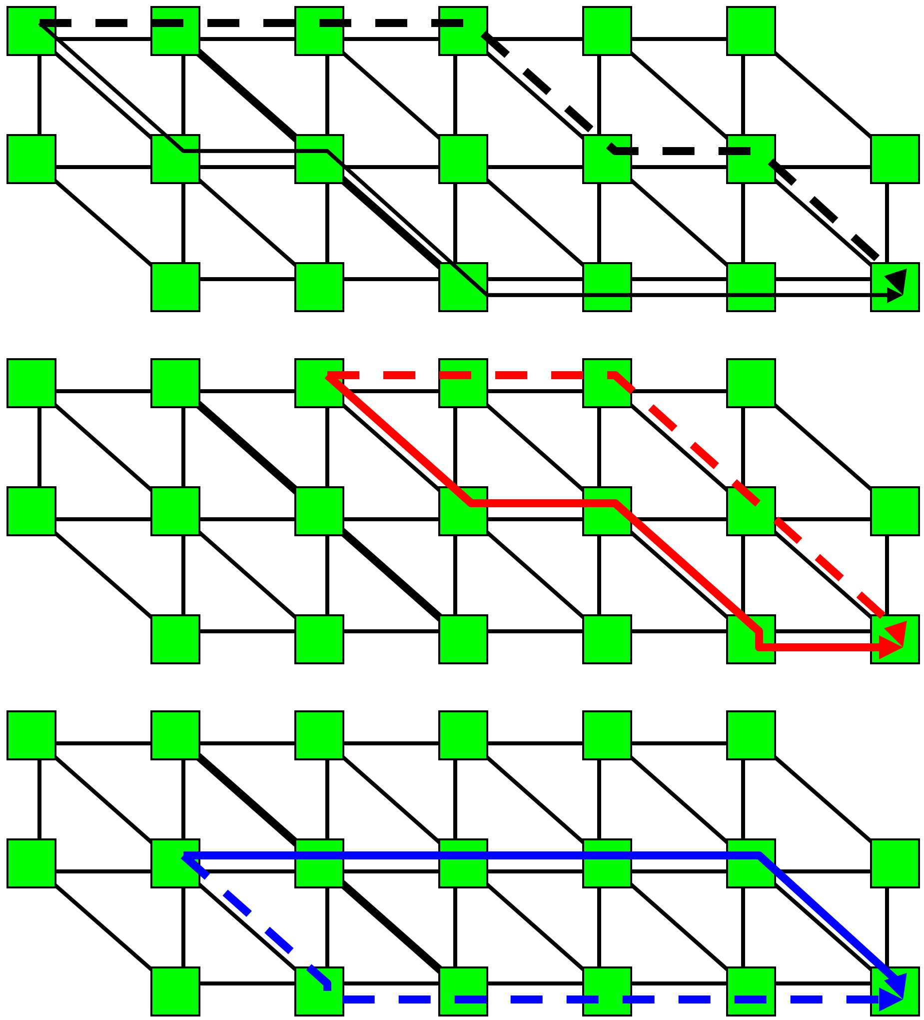

Independent routing!!!

Sharing in MPλS: path protection

Even if red and black working paths do not overlap,
the wavelength cannot be shared on this link,
because they are routed differently downstream.

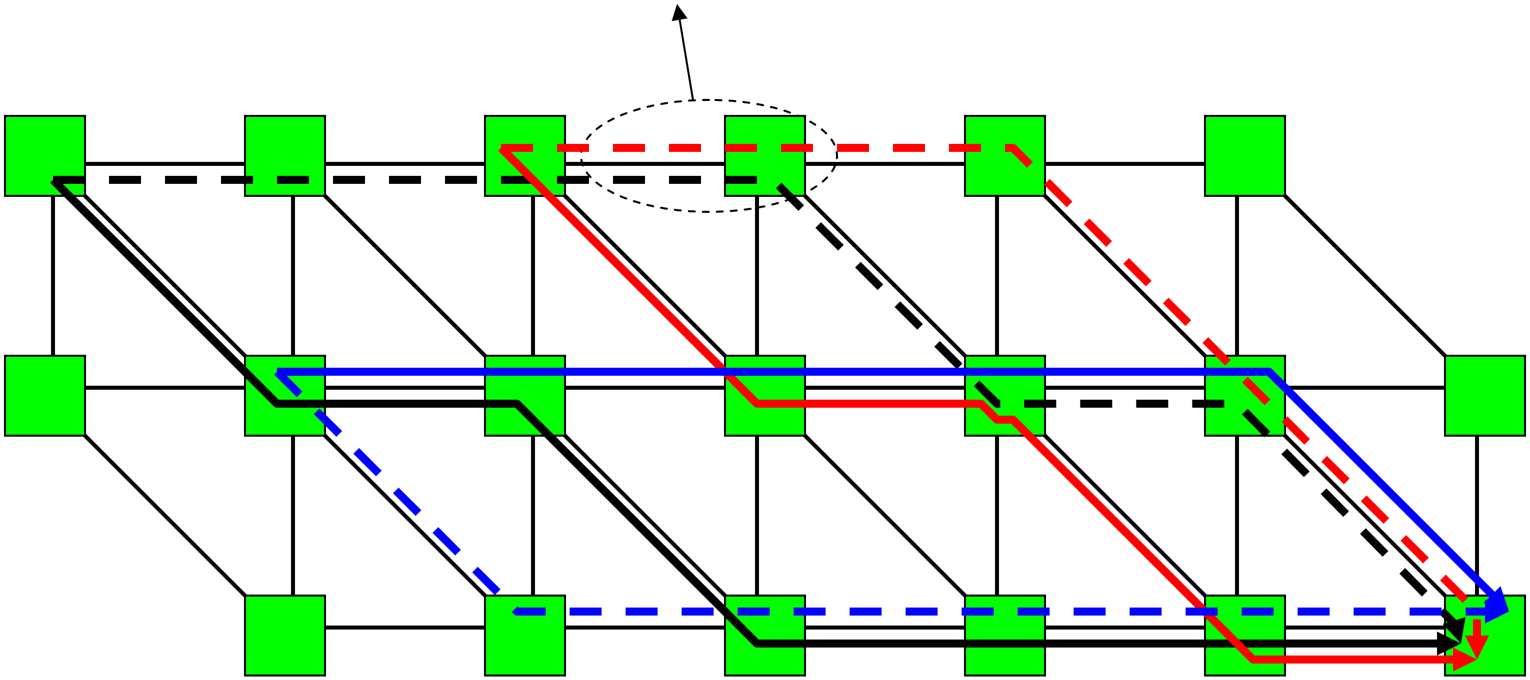

- At most 2 working paths through each piece of equipment.
  Thus at most 2 backup wavelengths needed on each link
- Cost backup wavelengths = 10+5sqrt(2)
  (unit = cost for 1 wavelength per length of horizontal link)

Sharing in MPλS: path protection

- How to force to share backup resources?
  - Limit routing of backup paths to a predefined/predistributed tree
  - Why?
    - Avoid situation that backup paths divert after overlapping
    - Forcing routing so that as much of the backup route is shared with other routes (even if this results in slightly longer backup routes --> to be compensated by the sharing).

- 3 backup routes can share 2 wavelengths
- Cost reduced from 10+5sqrt(2) to 12+2sqrt(2)

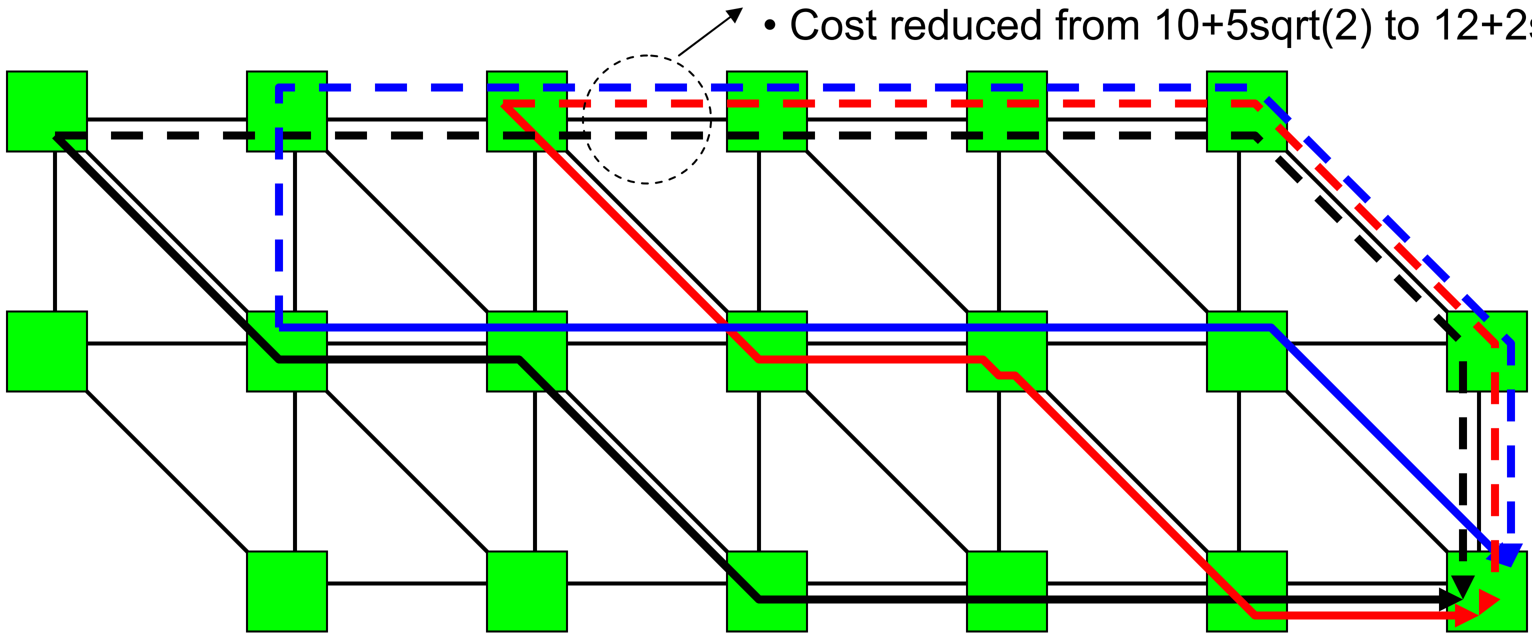

Sharing in MPλS: path protection

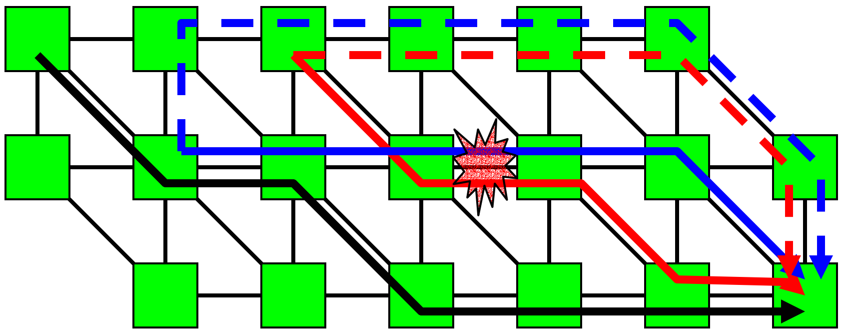

Red and blue should be protected at the same time.
To which color has the backup of the black path to be tuned, in order to share the backup wavelengt

Conclusion: ingress of black path cannot swap to THE backup OLSP, in combination with simple merging downstream.

Red

Blue