ERICSSON

Loop Free Alternates for SPBM Multicast

David Allan, András Császár, János Farkas, Panagiotis Saltsidis, Jeff Tantsura

March 18, 2014

Introduction

- P802.1Qca D0.6 addresses LFA for SPBM unicast
  - subclause 45.3.1 Loop free alternates for unicast data flows
- Comment #77 on P802.1Qca D0.3 requested the support for multicast

- A presentation on LFA for SPBM unicast is available at: http://www.ieee802.org/1/files/public/docs2013/ca-farkas-LFA-SPBM-unicast-0513-v02.pdf

Background

Recap – Destination Rooted Tree

a destination rooted Shortest Path Tree (SPT)

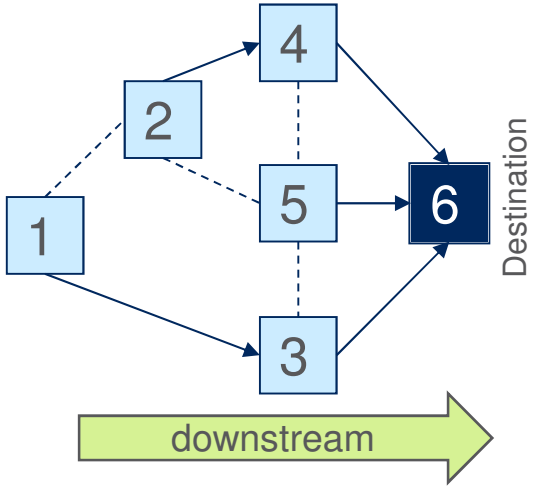

the same SPT as a spanning tree

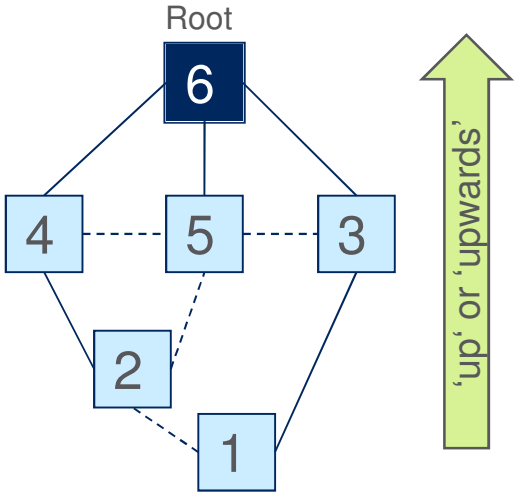

PLR: Point of Local Repair

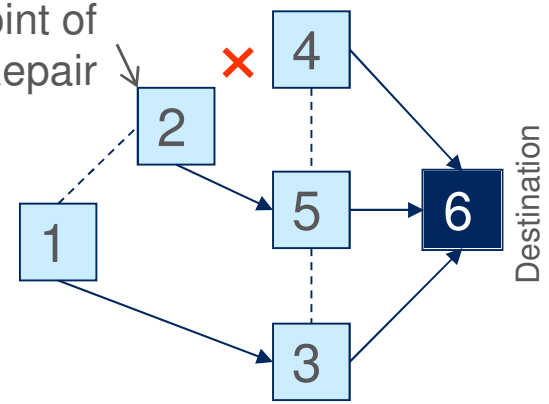

Recap –
Source Rooted Tree for Multicast

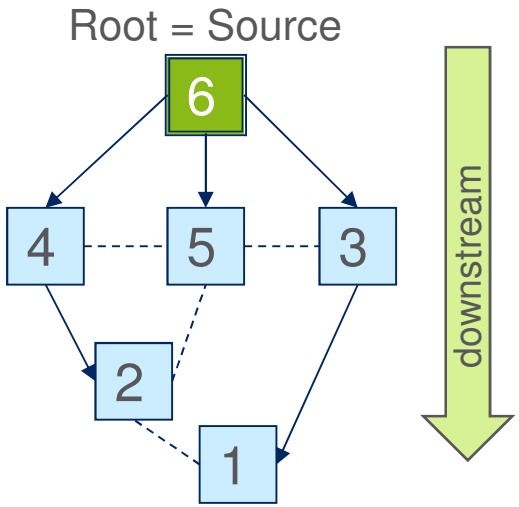

Recap – RSTP/MSTP

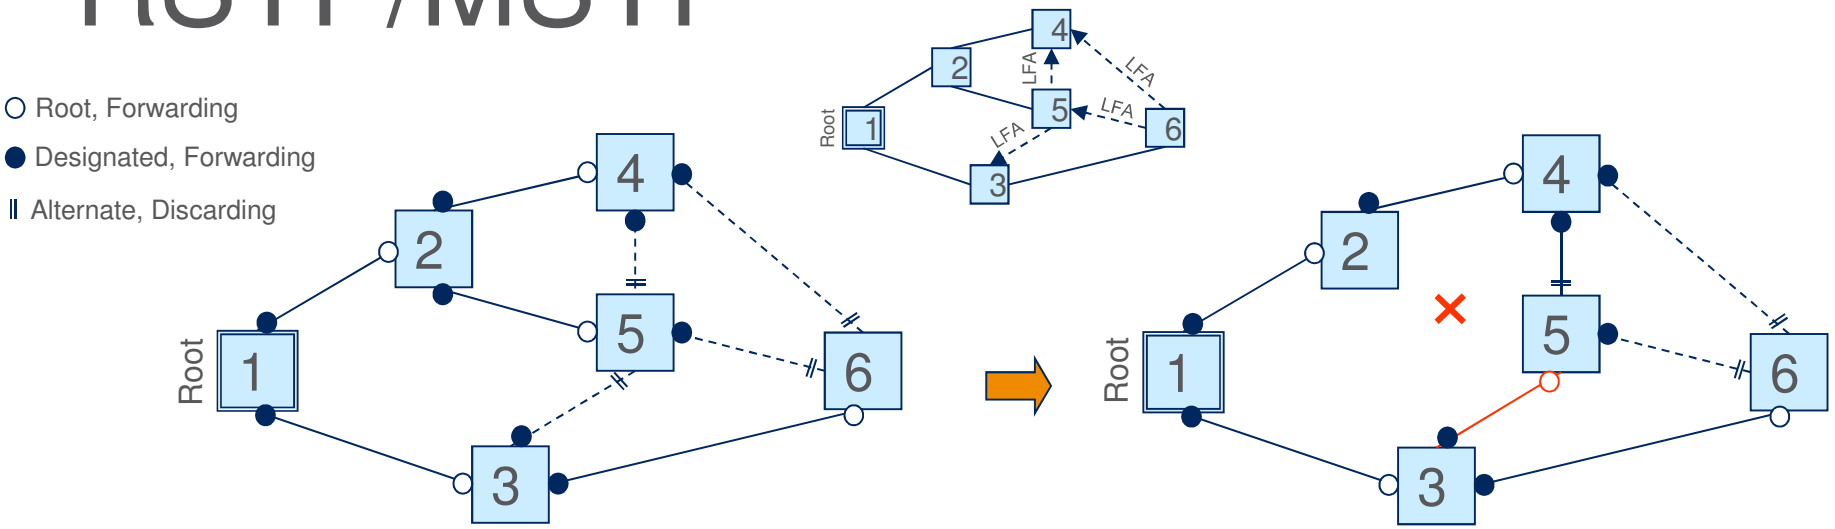

- Forwarding is along a (non-directed) SPT rooted at the Root Bridge
- RSTP implements downstream LFA by means of Alternate Ports
  - If connectivity is lost at a Root Port, then the Alternate Discarding port becomes Root Forwarding (e.g. 2-5 link goes down, then the formerly Alternate Discarding port of 5 towards 4 becomes Root Forwarding)
  - It is just a local repair! (the only action is at bridge 5 in the example)
- Loop prevention
  - Alternate Ports do not cause loop
  - Proposal-Agreement handshake during restoration
- Symmetry is kept as there is only one active path between any node pair

Recap – SPBV

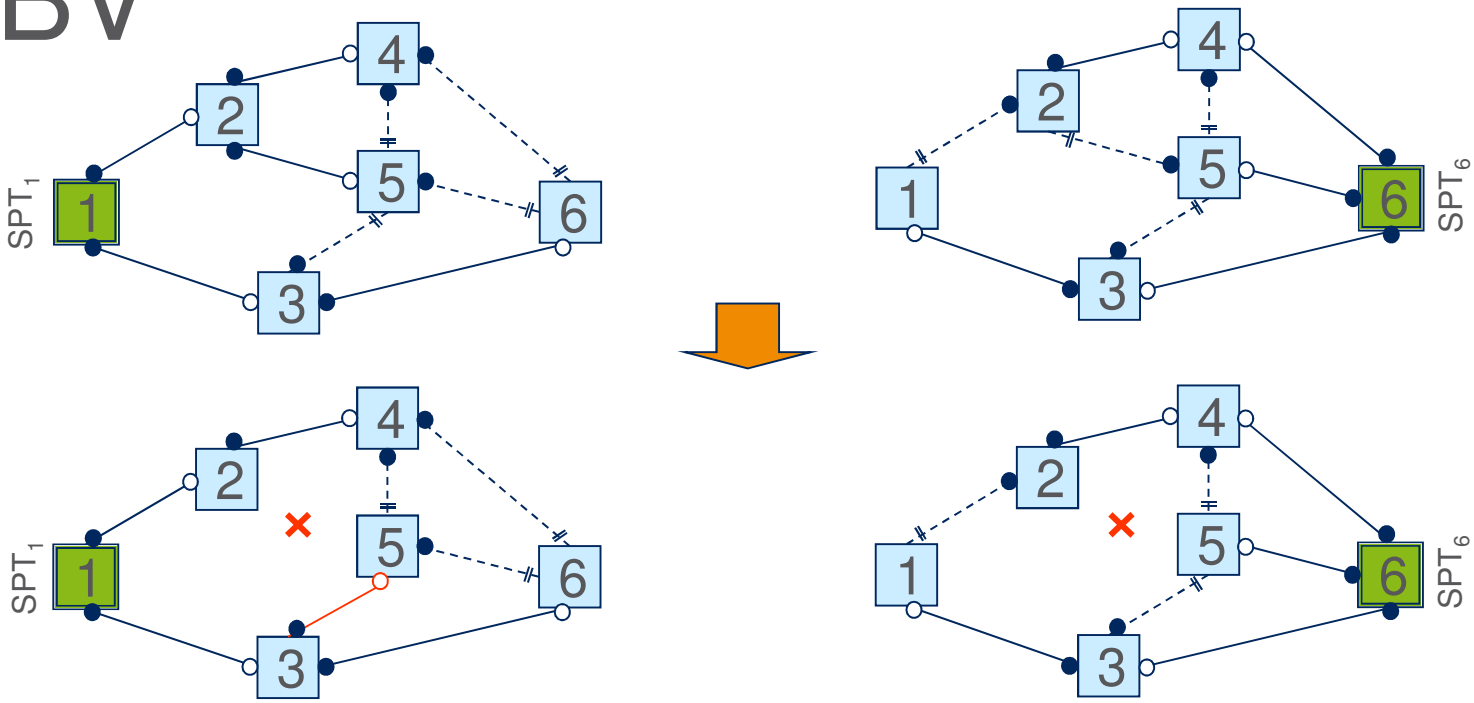

- SPTs can be implemented by the per tree Port State and Port Role variables (13.17)
  → Alternate Port based LFA is just there
- Loop prevention
  - Alternate Ports do not cause loop
  - Agreement Protocol during restoration
- Symmetry is maintained if the same ECT-Algorithm is used for the selection of the Alternate Ports as for the computation of the SPT

Recap – SPBM Forwarding

- Unicast
  - Destination rooted SPTs
  - Symmetry during basic forwarding is ensured by ISIS-SPB
  - Loops are mitigated by ingress checking

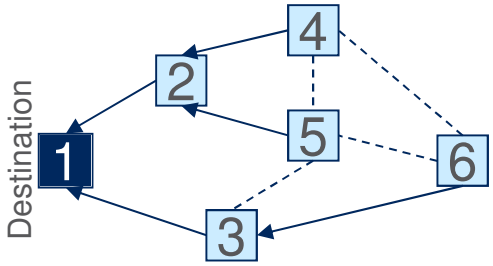

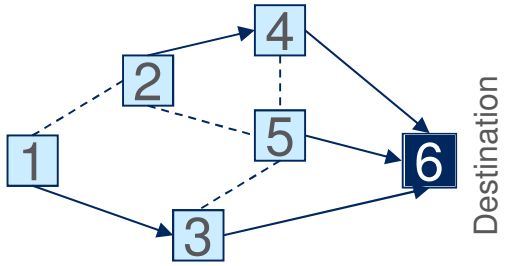

- Multicast
  - Source rooted SPTs as in case of SPBV
  - Symmetry, congruency during basic forwarding is ensured by ISIS-SPB
  - During restoration, loops are prevented by the Agreement Protocol

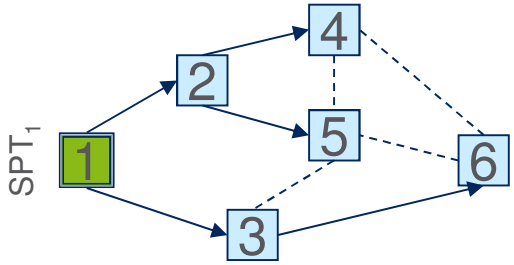

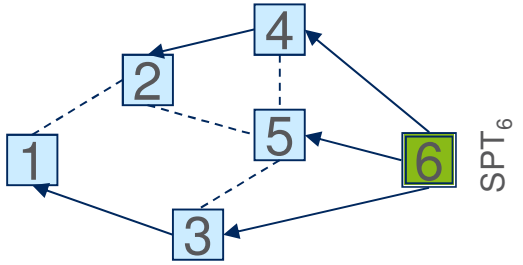

Recap – LFA for SPBM Unicast

- Enhance ingress checking = extend, relax it
  - Frame reception from all upstream neighbours has to be allowed
  - It remains loop-free as the loop-free conditions are still met
- Ensure safe updates after a topology change
  - Change detected
    - Update ingress checking to block unsafe reception
    - Remove unsafe LFAs
  - Re-convergence is over
    - Install states for primary paths
    - Install states for new LFA path
    - Update ingress checking to allow reception from upstream neighbours

Problem

SPBM Multicast Example – Source Rooted SPTs

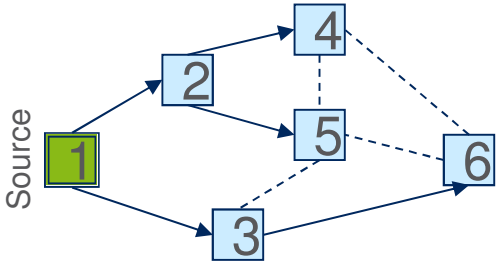

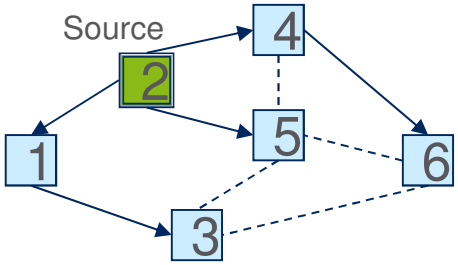

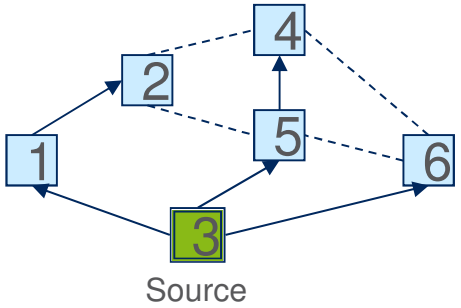

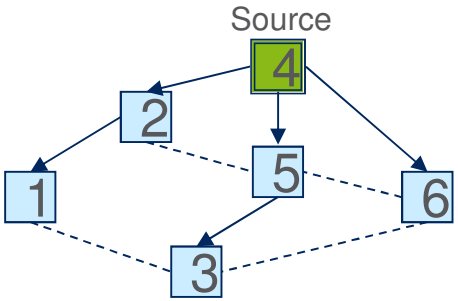

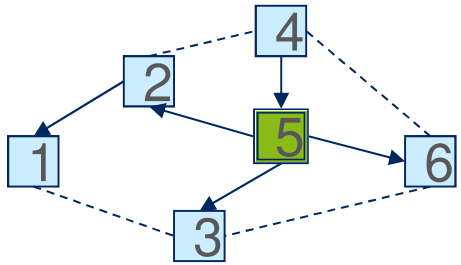

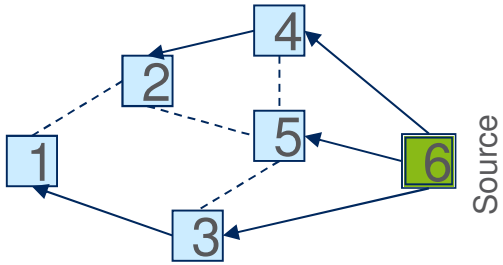

SPBM Multicast Example – Breakdown of 2-4 Link

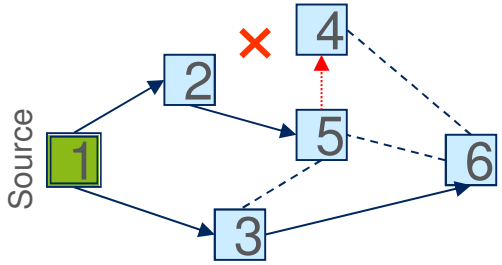

- 4 can admit 1→4 frames
- 5 should install state

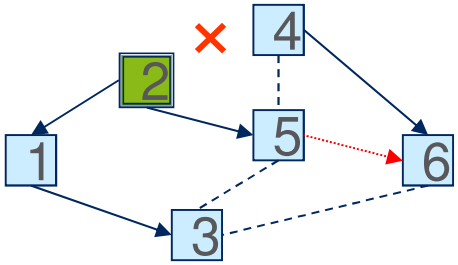

- 6 could admit 2→6 frames
- 5 should install state

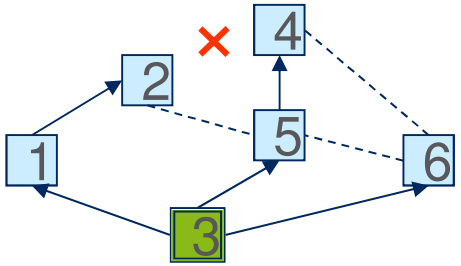

- nothing to do

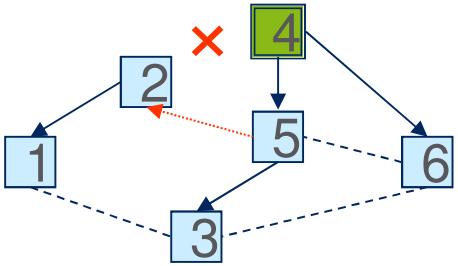

- 2 can admit 4→1 frames
- 5 should install state

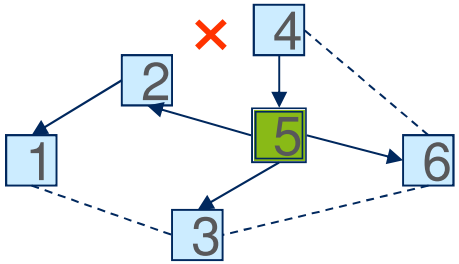

- nothing to do

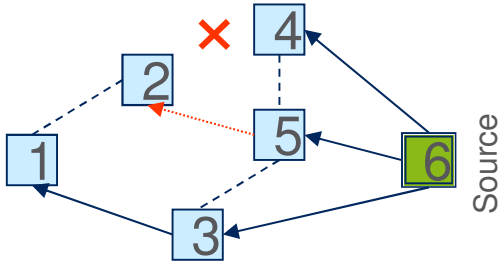

- 2 can admit 6→2 frames
- 5 should install state

› Bridges not being aware of the failure should install LFA state

SPBM Multicast Example – Breakdown of 1-2 Link

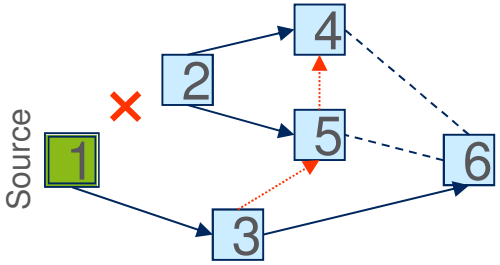

- 5 could admit 1→4 frames
- 4 could admit 1→4 frames
- 3 and 5 should install state

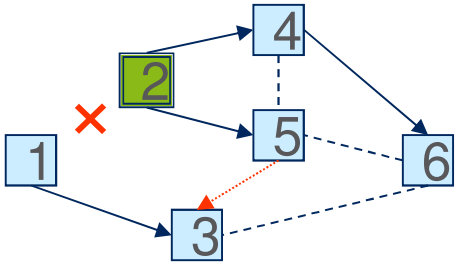

- 3 could admit 2→3 frames
- 5 should install state

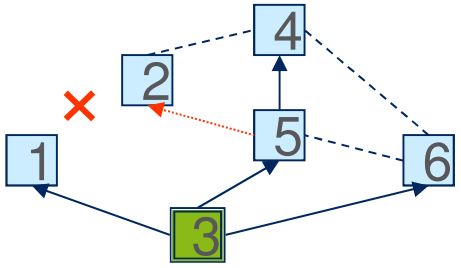

- 2 can admit 3→2 frames
- 5 should install state

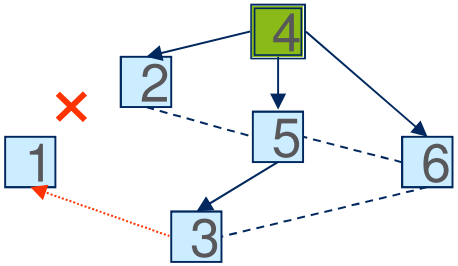

- 1 can admit 4→1 frames
- 3 should install state

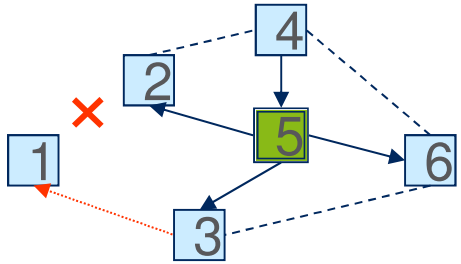

- 1 can admit 5→1 frames
- 3 should install state

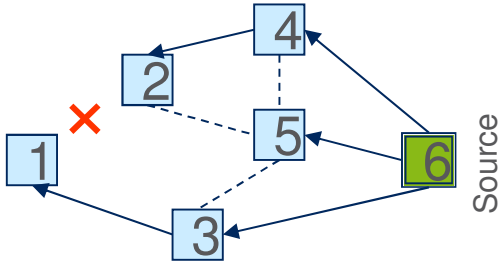

- nothing to do

Bridges not being aware of the failure should install LFA state

SPBM Multicast Example – Breakdown of Bridge 2

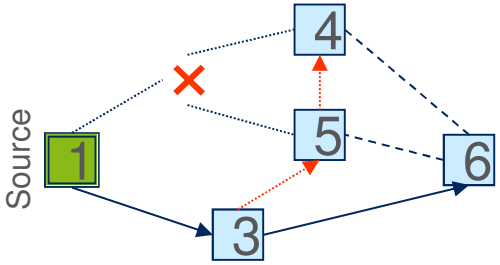

- 5 can admit 1→4 frames
- 5 can install state
- 4 can admit x→4 frames
- 3 should install state

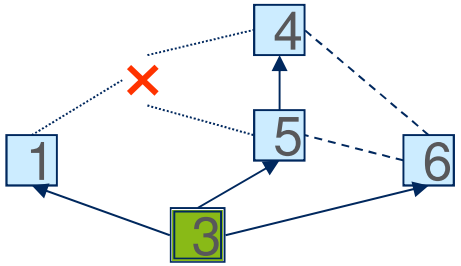

- nothing to do

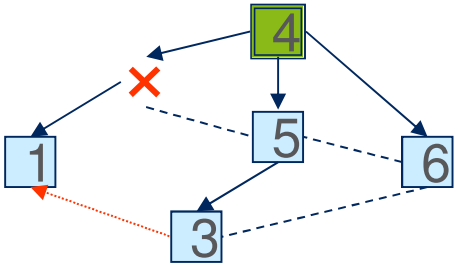

- 1 can admit x→1 frames
- 3 should install state

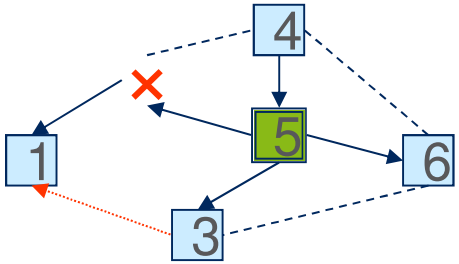

- 1 can admit x→1 frames
- 3 should install state

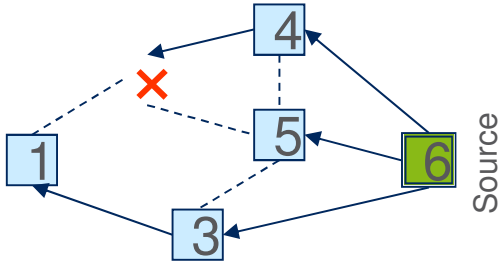

- nothing to do

Bridges not being aware of the failure should install LFA state

The Difficulty

- The bridges that need to react to the failure are not aware of it.
- Therefore, local repair is not enough.

ERICSSON

Solution Possibilities

Basic Idea

- Use notification(s) to enable LFA for multicast
  - Simplest notification: LSP
    - Remote LFA action is keyed/triggered by LSP/LSDB update
    - NOTE that it is NOT the restoration triggered by an LSP!
    - See details later on
  - Dedicated notification messages:
    - Leaf Notification
    - Root Notification as an optional add-on
    - See details later on
- Safe alternate paths are determined a-priori for each source rooted SPT

To Be Determined A Priori

- Safe alternate paths
- Reaction to the failure
  - Depends on what type of notification is used and on the data frame forwarding discipline applied
  - It can be
    - Blocking primary path and opening a safe alternate path
    - Start sending data frames on alternate path

LSP-based Notification

- LSPs are used to notify remote nodes
- Remote nodes react as determined a-priori
  - This is NOT the link state restoration!
  - The reaction to the failure is keyed on the received LSP

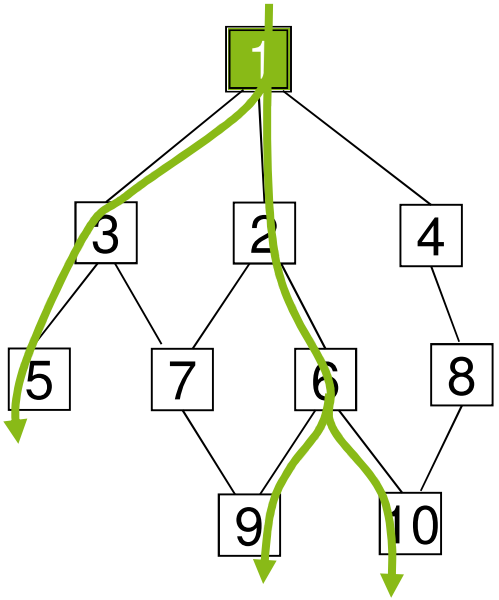

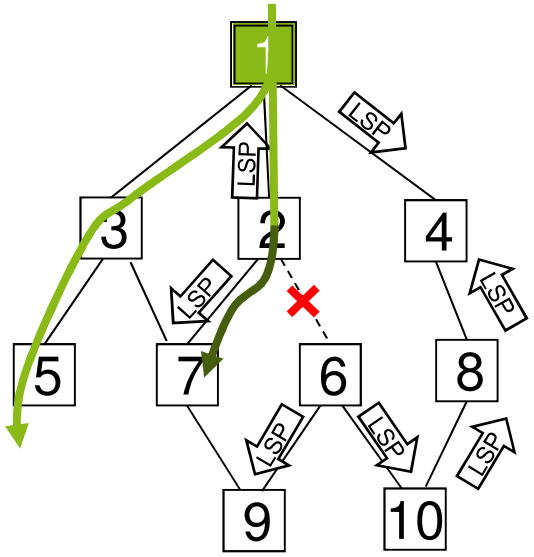

- Bridges detecting the failure send LSPs, which are flooded
- Bridge 2 performs local repair, i.e. redirects towards bridge 7

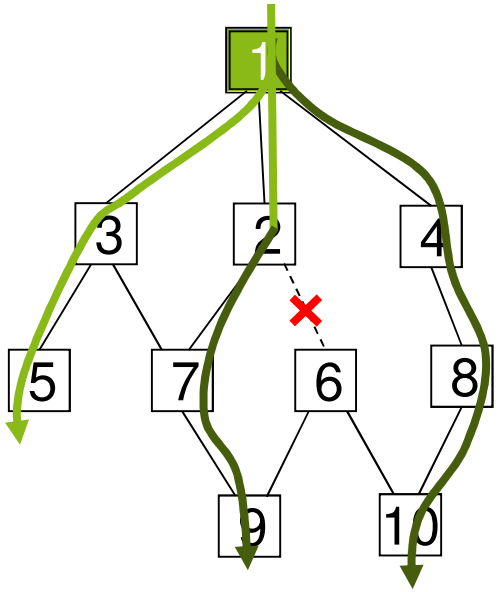

Bridges react keyed on the LSP encoding the failure

- Bridges 1, 4, 8 forward to 10,
- Bridge 7 forwards to 9
- Bridges 9 and 10 admit on alternate

Leaf Notification (LN)

- Leaf Notification (LN) message is sent by the detector of a failure towards the leafs on the SPT
- Data frames: multiple copies are needed if only LN is used
  - A copy of the data frame is sent on each alternate path aside the primary path
  - Blocking is needed on the alternate paths where they merge the primary path, thus ensuring that a single copy of the frame is sent onwards
- Data frame reception is unblocked on one of the alternates upon reception of LN

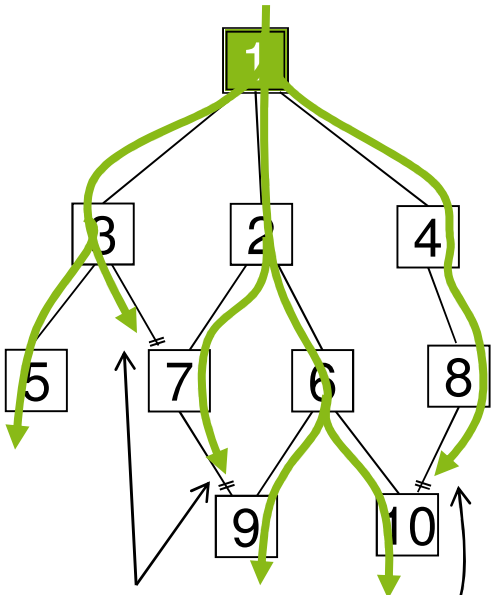

dropping is only
by reception port
at a merger bridge

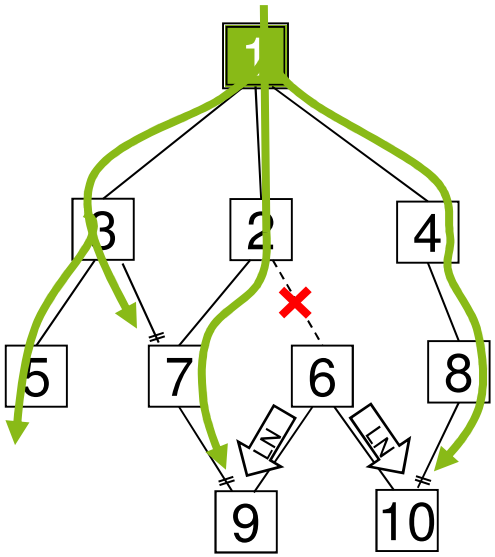

Bridge detecting the failure (e.g. Bridge 6) sends Leaf Notification (LN) meaning:
"Lost connection to Root"

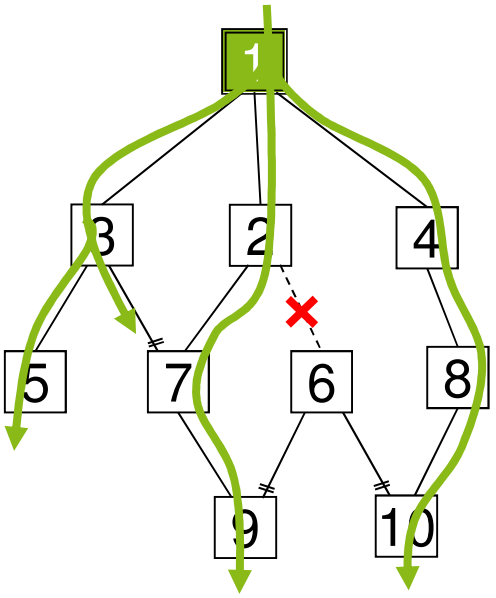

Bridges receiving LN (e.g. 9 & 10)

- Block Primary Port

- Unblock Alternate Port

Leaf Notification – cont'd

- LN message is sent towards the leafs on the SPT → received on the primary path
- LN message means: Root is not reachable on the shortest path
- Reaction to LN message is simple: use alternate
- FDB entries for LN message are to be installed a-priori
- Encoding of LN
  - LN message has to be distinguished from data frames, but
  - It has to convey the information on the SPT it refers to, i.e. B-VID, I-SID and Source
  - Being an LN message can be encoded in the SPSourceID, e.g. have an LN-SPSourceID allocated to the Root=source in addition to the SPSourceID used for data frames

Note: LSP Notification with Multiple Copies of Data Frames

- LSP-based notification can be also used when multiple copies of the data frames are sent
  - A copy of the data frame is sent on each alternate path aside the primary path
  - Alternate paths are blocked where they merge the primary path, thus ensuring that a single copy of the frame is sent onwards
- Data frame reception is unblocked on one of the alternates upon reception of LSP indicating failure

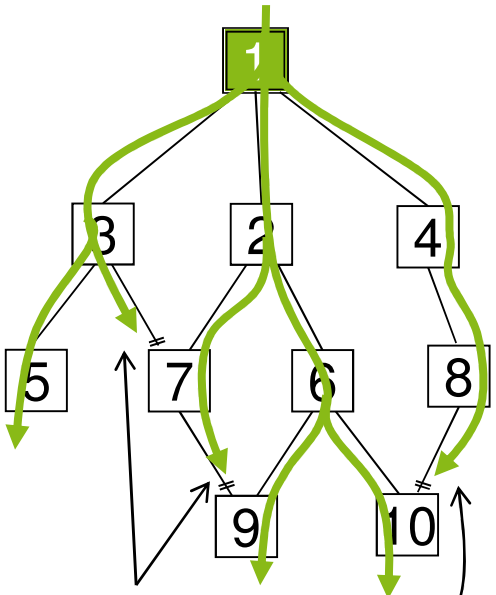

dropping is only by reception port at a merger bridge

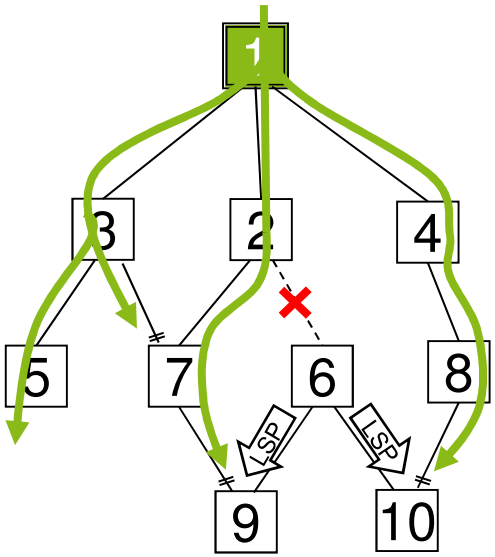

Downstream LSPs (e.g. by Bridge 6) are only important if data frames are sent on the alternate paths too

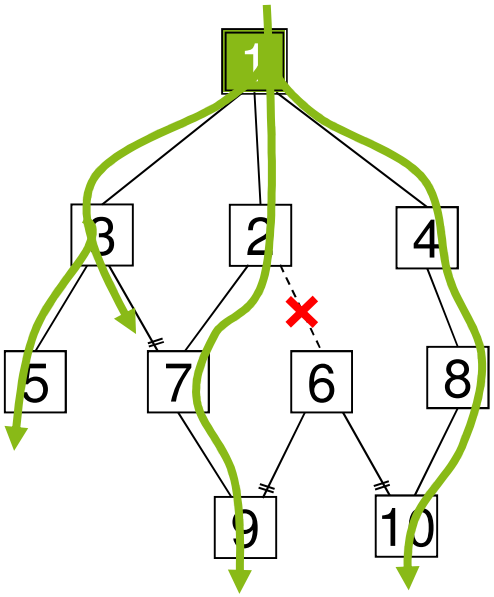

Bridges receiving LSP (e.g. 9 & 10)

- Block Primary Port
- Unblock Alternate Port

Root Notification (RN)

- It is an add-on to Leaf Notification
- An RN is sent to the Root on the alternate path got activated by a bridge receiving an LN
- RN indicates that data frames are to be sent on the alternate path
- Single copy of data frame is enough

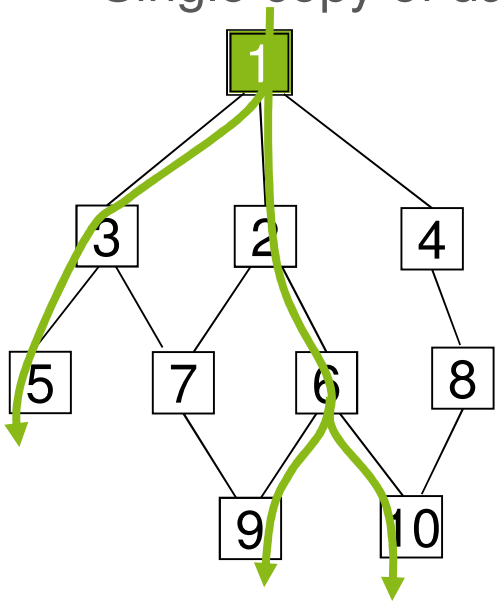

shortest paths are only used

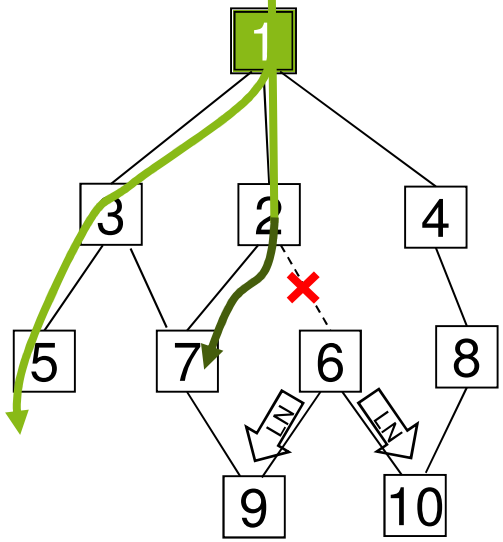

Bridge detecting the failure (e.g. Bridge 6) sends Leaf Notification (LN) meaning: **"Lost connection to Root"**

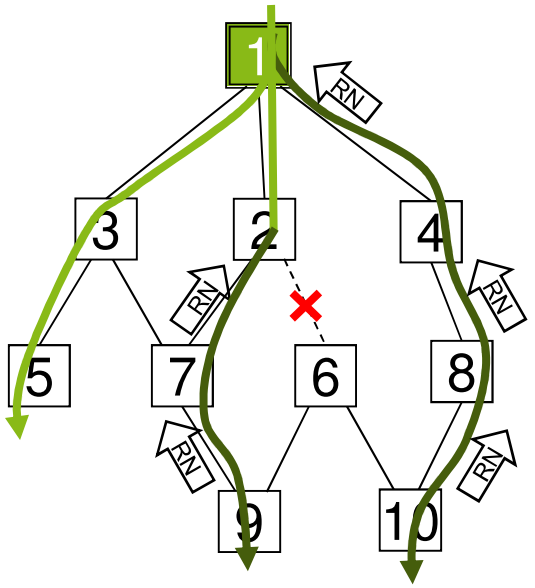

Bridges receiving RN (e.g. 9 & 10) send RN on alternate path
Bridges receiving data + RN (e.g. 1 & 2) start sending on RN reception port

Root Notification – cont'd

- RN message is sent towards the Root on the alternate path
- LN message means: this alternate path is to be used
- Reaction to RN message: start sending data frame on the alternate path on which RN was received if data frames are being received from the Root
- FDB entries for LN message are to be installed a-priori
- Encoding of LN
  - RN message has to be distinguished from data frames, but
  - It has to convey the information on the SPT it refers to, i.e. B-VID, I-SID and Source
  - Being an RN message can be encoded in the SPSourceID, e.g. have an RN-SPSourceID

Summary

- Notification is required for SPBM multicast LFA
- LSP notification: single copy of data frames
- Leaf Notification with multiple copies of data frames
- Leaf Notification + Root Notification: single copy of data frames