802.1aq: link-state protocol and loop prevention

János Farkas

ERICSSON
TAKING YOU FORWARD

Outline

- Background
- Transient loops may appear
- Loop elimination possibilities
- Summary

Background

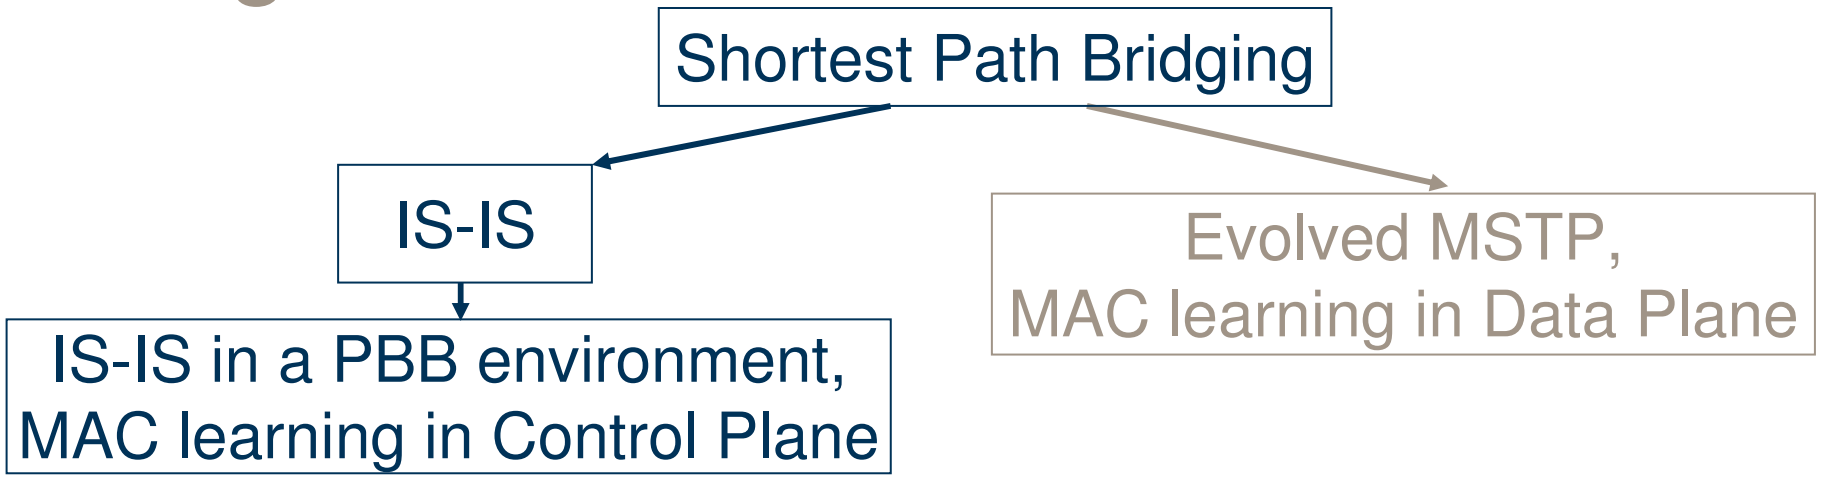

Ongoing discussions:

- Source tree identification
  - VLAN ID
  - Destination MAC address
  - Source MAC address
- Number of trees
  - One tree per source, Root Bridge: BEB
  - Sub-Optimal trees, Root Bridge: BCB
  - ECMT
- Loop prevention mechanism
  - Ingress checking
  - Ingress checking + TTL

First version described in the draft

Loop mitigation

- Ingress checking (e.g. RPFC)
  - Frames not arriving on the shortest path from the Source Bridge are discarded
  - Makes the tree directed
  - Good for loop prevention in most cases
  - Transient loops may appear
  - Ingress filtering has to be modified
- TTL (hop-count)
  - Looped frames are discarded after a while
  - Spreading of multicast frames lasts shorter but not eliminated
  - New field in header

Transient loop in case of a link-state protocol with ingress checking

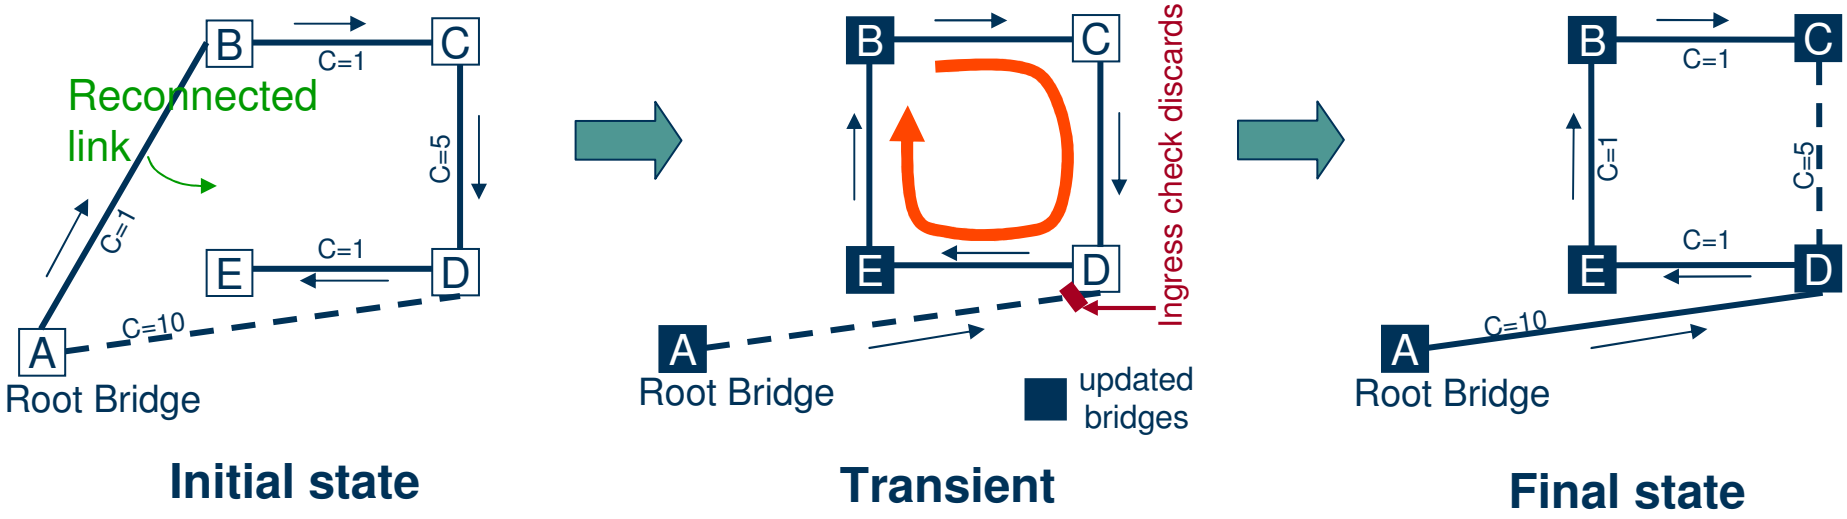

- Unpredictable order of active topology updates
- Transient loop may appear
- Multiplication of multicast frames: Catherine wheel

Possibilities to eliminate transient loops

- Transient loop is not a problem for unicast traffic
  - Little effect on the network
  - Use ingress checking only
- Multicast traffic getting more dominant
  - Severe problem for multicast
  - A chance of network melt-down remains if one does not care

- Elimination methods already discussed (besides ingress checking)
  - TTL: Spreading of multicast frames lasts shorter but not eliminated
  - Neighbor checking
    - Handshake mechanism applied e.g. in RSTP/MSTP
    - Eliminates the looping possibility
    - Implementation possibilities
      - As a new functionality within IS-IS
      - Run MSTP only for neighbor checking parallel to IS-IS
  - Ordered FIB update
    - Proper order of the update of FIBs eliminates loops
    - Order depends on the type of topology change → Dynamically calculated
    - Hold-down and network specific timers
      - Hard to find optimal values
      - Slow operation → Acceleration uses handshake

- **New proposal: Strict Sequential Update**
  - Fully controlled update of the active topology
  - Details in next slide

Controlled transient: Strict Sequential Update

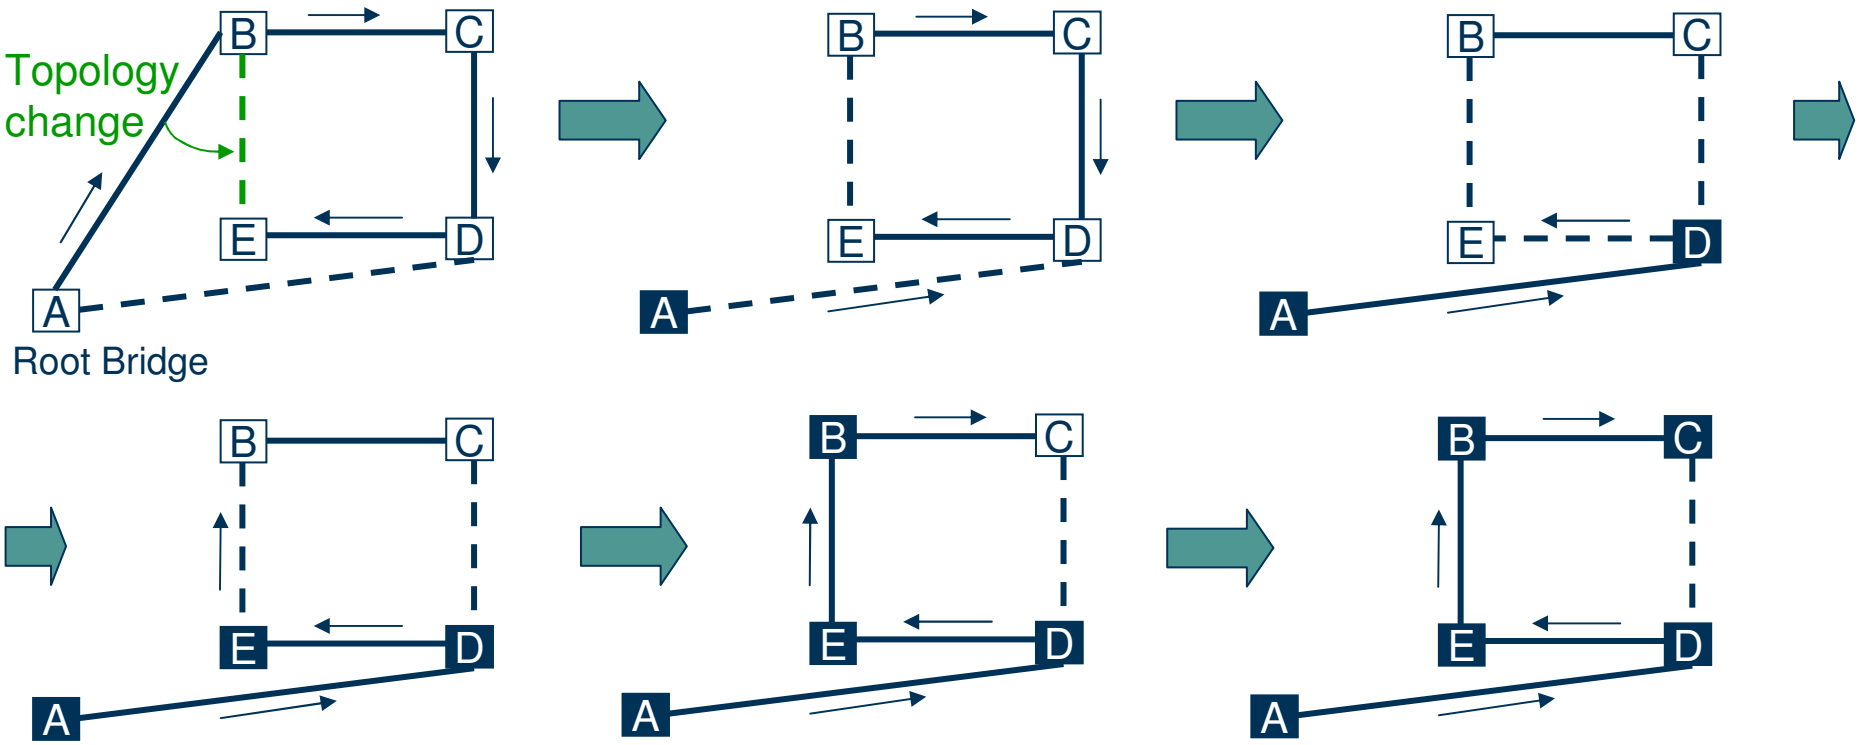

- Root Bridge controls the update
- The control of the update travels along the new tree
- Intermediate bridge forwards the control message after its own update
- Implement e.g. in a new sub-TLV (RFC 3784)

Summary

- Transient loops may appear if ingress checking is the only loop prevention mechanism in IS-IS based SPB
- Additional method needed besides ingress checking in order to eliminate Catherine-wheels
- Apply Strict Sequential Update
  - Full control on topology change
  - Simple extension to IS-IS
  - Eliminates any loop without ingress checking even if MAC learning is in data plane
- Next step
  - Work out detailed operation of Strict Sequential Update

Backup slides

ERICSSON
TAKING YOU FORWARD

Indirect proof of loop elimination

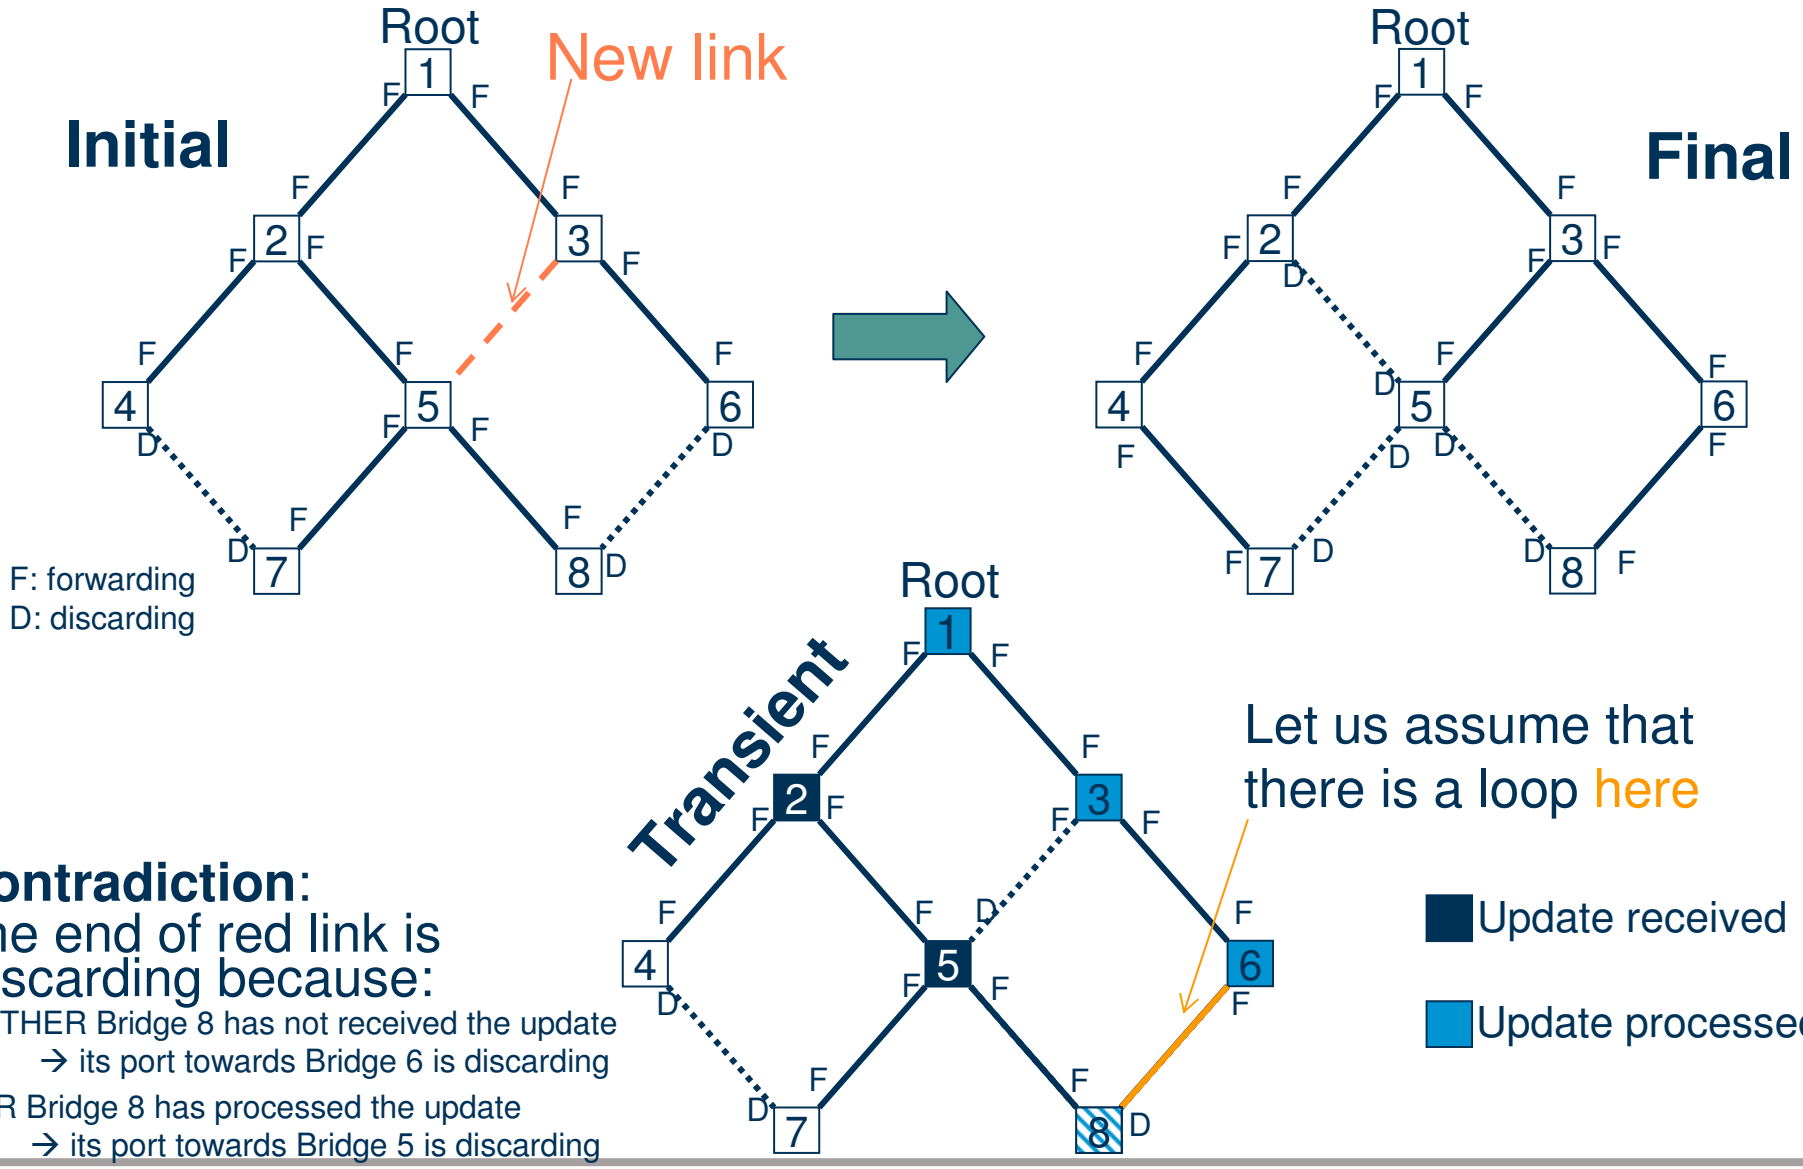

Contradiction:
one end of red link is Discarding because:

- EITHER Bridge 8 has not received the update
  → its port towards Bridge 6 is discarding
- OR Bridge 8 has processed the update
  → its port towards Bridge 5 is discarding

RSTP: Worst-case convergence