

RPR Topology Discovery Proposal

Jeanne De Jaegher, Alcatel

Jason Fan, Luminous

Henry Hsiaw, NEC

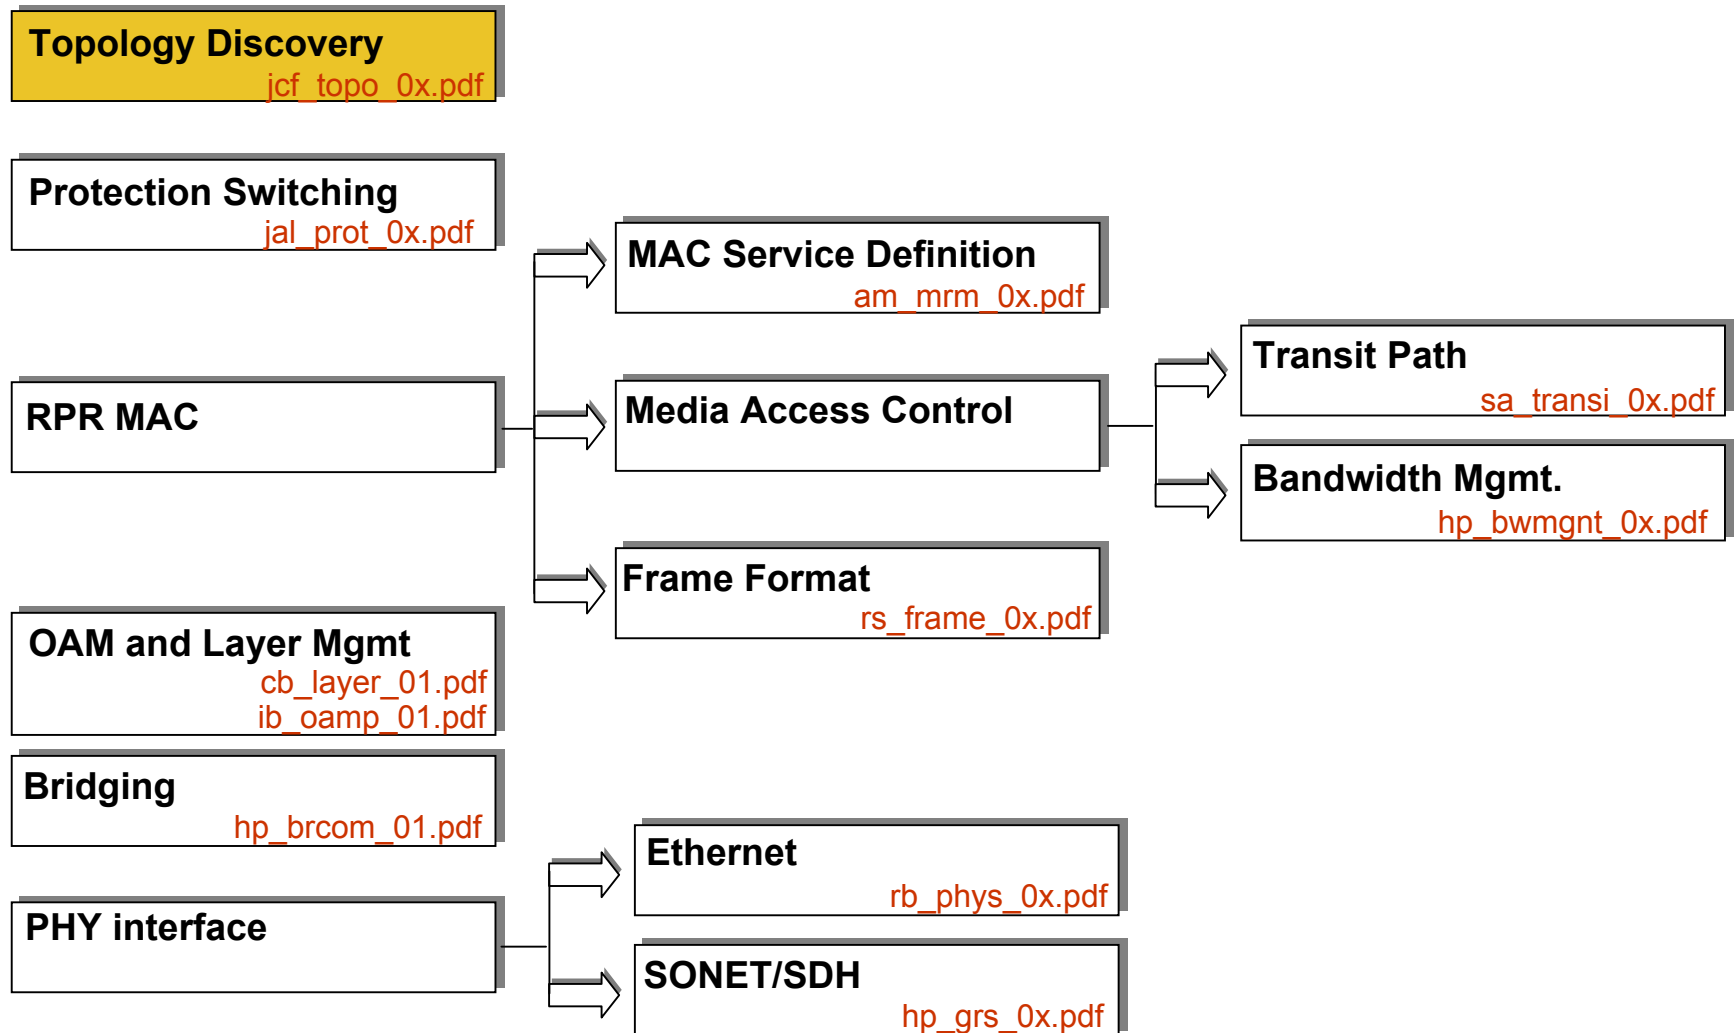
John Lemon, Lantern

Robin Olsson, Vitesse

Harry Peng, Nortel

Frederic Thepot, Dynarc

Components of a complete RPR proposal



Goals

- Scalable from 1 to 100's of stations
- Determine/validate connectivity and ordering of stations on the ring
- Ensure all stations on the ring have a uniform and current image of the topology
- Immediate reaction to changes
- Tolerant of message loss
- Operate without any master station on the ring
- Operate independently of and in the absence of any management systems

Goals, continued

- Usable with all supported topologies: multiple ringlet ring, linear (broken ring), and “star” (single station)
- Support dynamic addition and removal of stations to/from the ring
- Detect mis-cabling between stations
- Provide means of sharing additional information between stations
- Cause minimal overhead

Information Sharing

- RPR Topology Image used by other algorithms
 - Steering algorithm uses Topology Image to know when steering is needed
 - Congestion avoidance uses Topology Image to know where congestion is being experienced

Topology Discovery Triggers

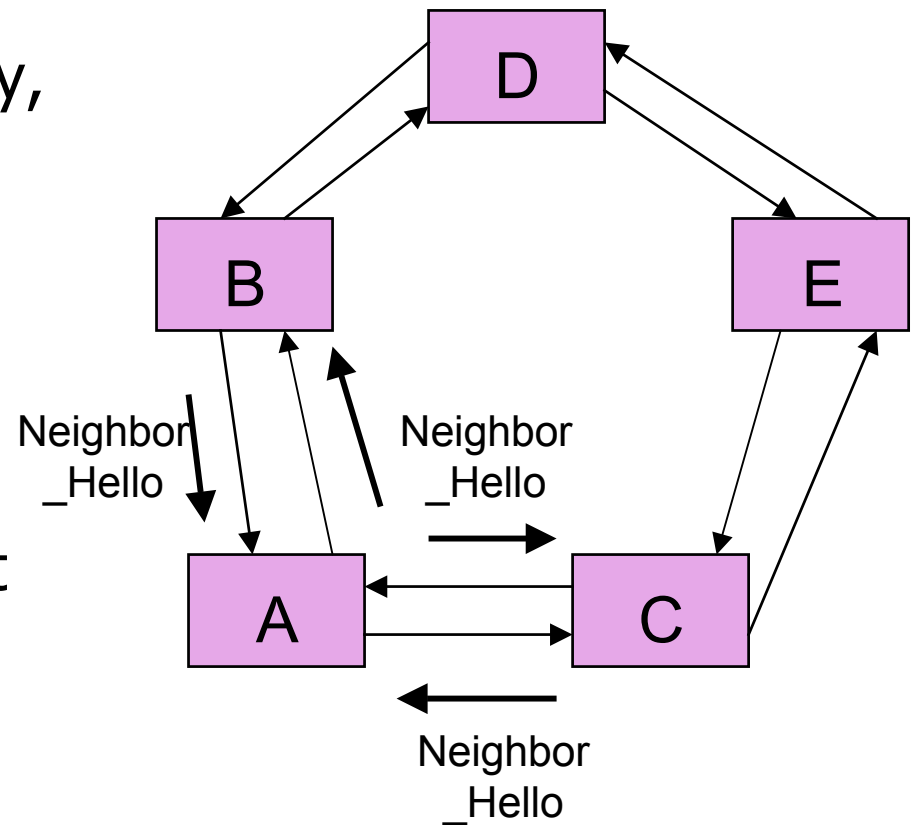
- Neighbor change at any station
 - Addition or deletion of neighbor
 - Change in link status
- Detection of validation failure at any station
 - Station lacking topology image
 - Station with outdated or corrupted topology image

Image Versions

- Station_Image_Version
 - Starts at 0 (indicating no valid image)
 - Incremented upon each change in local status
 - Independent value for each station
- Ring_Image_Version
 - Checksum of all Station_Image_Versions for all known stations (including self)
 - Common value for each station

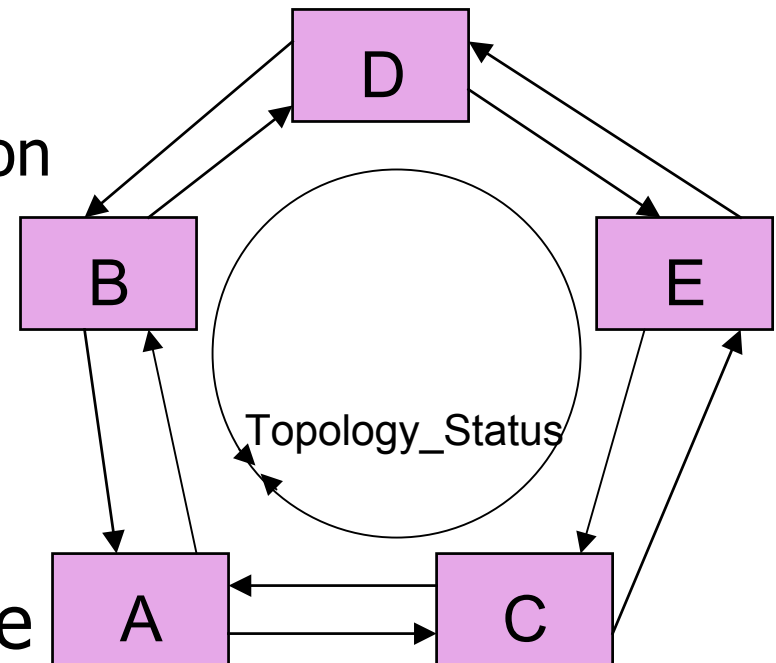
Neighbor_Hello Control Message

- Reports presence, identity, and topology version of neighbor station
- Key fields
 - Ring image version
 - Ringlet ID
- Broadcast on each ringlet with TTL = 1
 - Removed by neighbor



Topology_Status Control Message

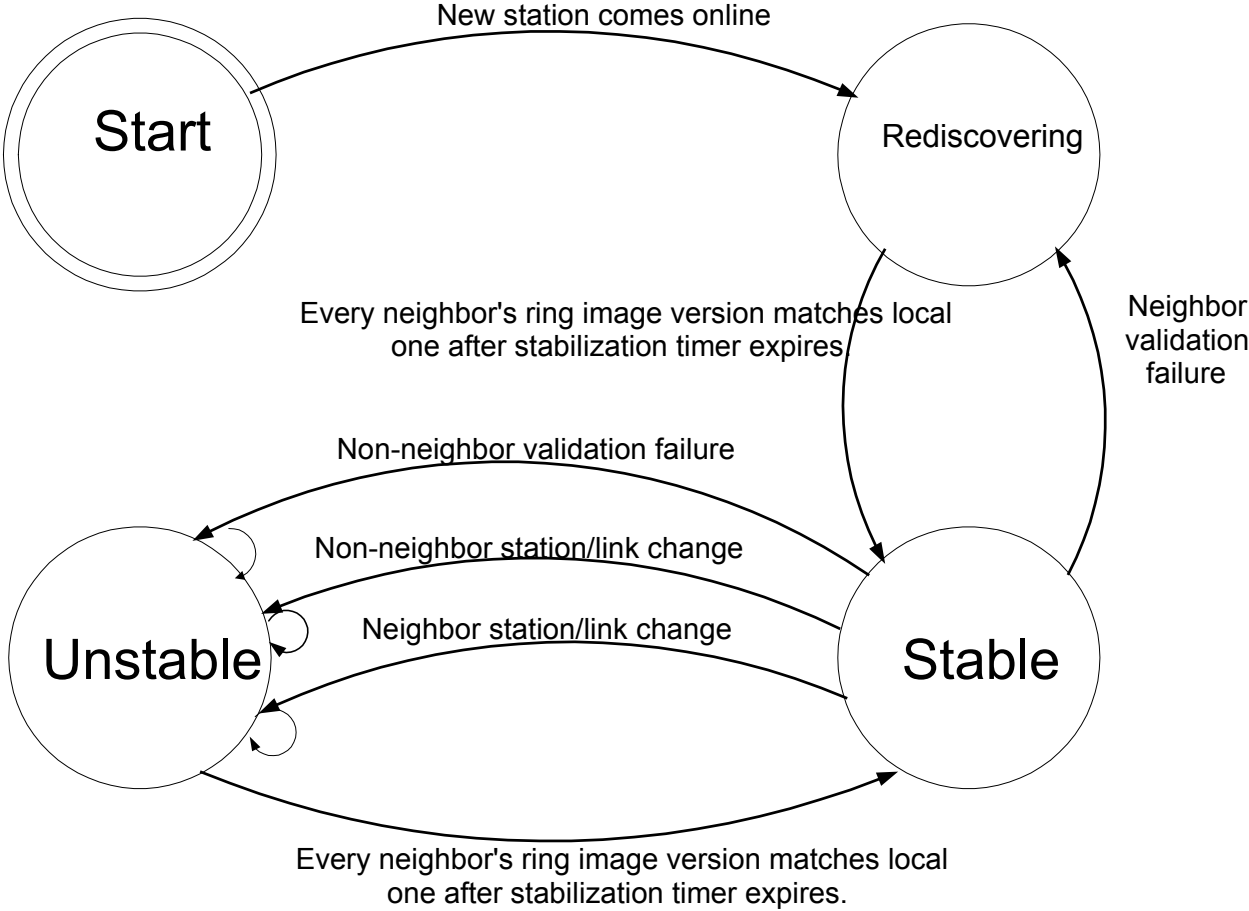
- Reports changes in neighbor identity or link status
- Key fields
 - Source station image version
 - Per ringlet
 - Neighbor MAC addresses
 - Neighbor link statuses
 - Ringlet ID
- Broadcast on each ringlet with $TTL = Max_Ring_Size$
 - Removed by source



Configurable Parameters

- Neighbor_Hello_Timer
- Topology_Stabilization_Timer

State Diagram



State Diagram Details, 1

- Neighbor station/link change
 - Trigger
 - No Neighbor_Hello messages in 3 Neighbor_Hello Periods (NHPs) or
 - Two successive Neighbor_Hellos from a new neighbor in 3 NHPs.
 - Action
 - Increment the local Station_Image_Version
 - Broadcast a Status_Change message
 - Replace the station information in the local topology image
 - Update the local Ring_Image_Version

State Diagram Details, 2

- Non-neighbor station/link change
 - Trigger
 - A higher Station_Image_Version is received in a Status_Change message
 - Action
 - Replace the remote station information in the local topology image
 - Update the remote Station_Image_Version
 - Update the local Ring_Image_Version

State Diagram Details, 3

- Neighbor validation failure
 - Trigger
 - A Ring_Image_Version in a Neighbor_Hello doesn't match the local one, or
 - the local Ring_Image_Version is 0 (a new station)
 - Action
 - Set the local and all the remote Station_Image_Versions = 0
 - Send a Status_Change message

State Diagram Details, 4

- Non-neighbor validation failure
 - Trigger
 - A Status_Change message with Station_Image_Version = 0
 - Action
 - Update the remote Station_Image_Version to 0
 - Broadcast a Status_Change message
 - Update the local Ring_Image_Version

Topology_Stabilization Timer

- Once in any of the above conditions, start the Topology_Stabilization_Timer.
- While the Topology_Stabilization timer is running, do not compare the Ring_Image_Versions.

Simulation Results

- Set up
 - 256 stations
 - 200 km circumference
 - Dual ringlet ring
 - 1 Gbps ring rate
 - Processing times for messages set to exponentially distributed times around mean of 200 usec for Neighbor_Hello and 500 usec for Topology_Status
- Scenario
 - Bring up all 256 stations at once (worst case)
- Results
 - For Neighbor_Hello_Timer of 500 msec, Topology Image complete in 1.65 seconds
 - For Neighbor_Hello_Timer of 1 sec, Topology Image complete in 3.15 seconds