



RPR Topology Discovery Proposal

Italo Busi, Alcatel

Jeanne De Jaegher, Alcatel

Jason Fan, Luminous

Yiming Yao, Luminous

Chi-Ping Fu, NEC

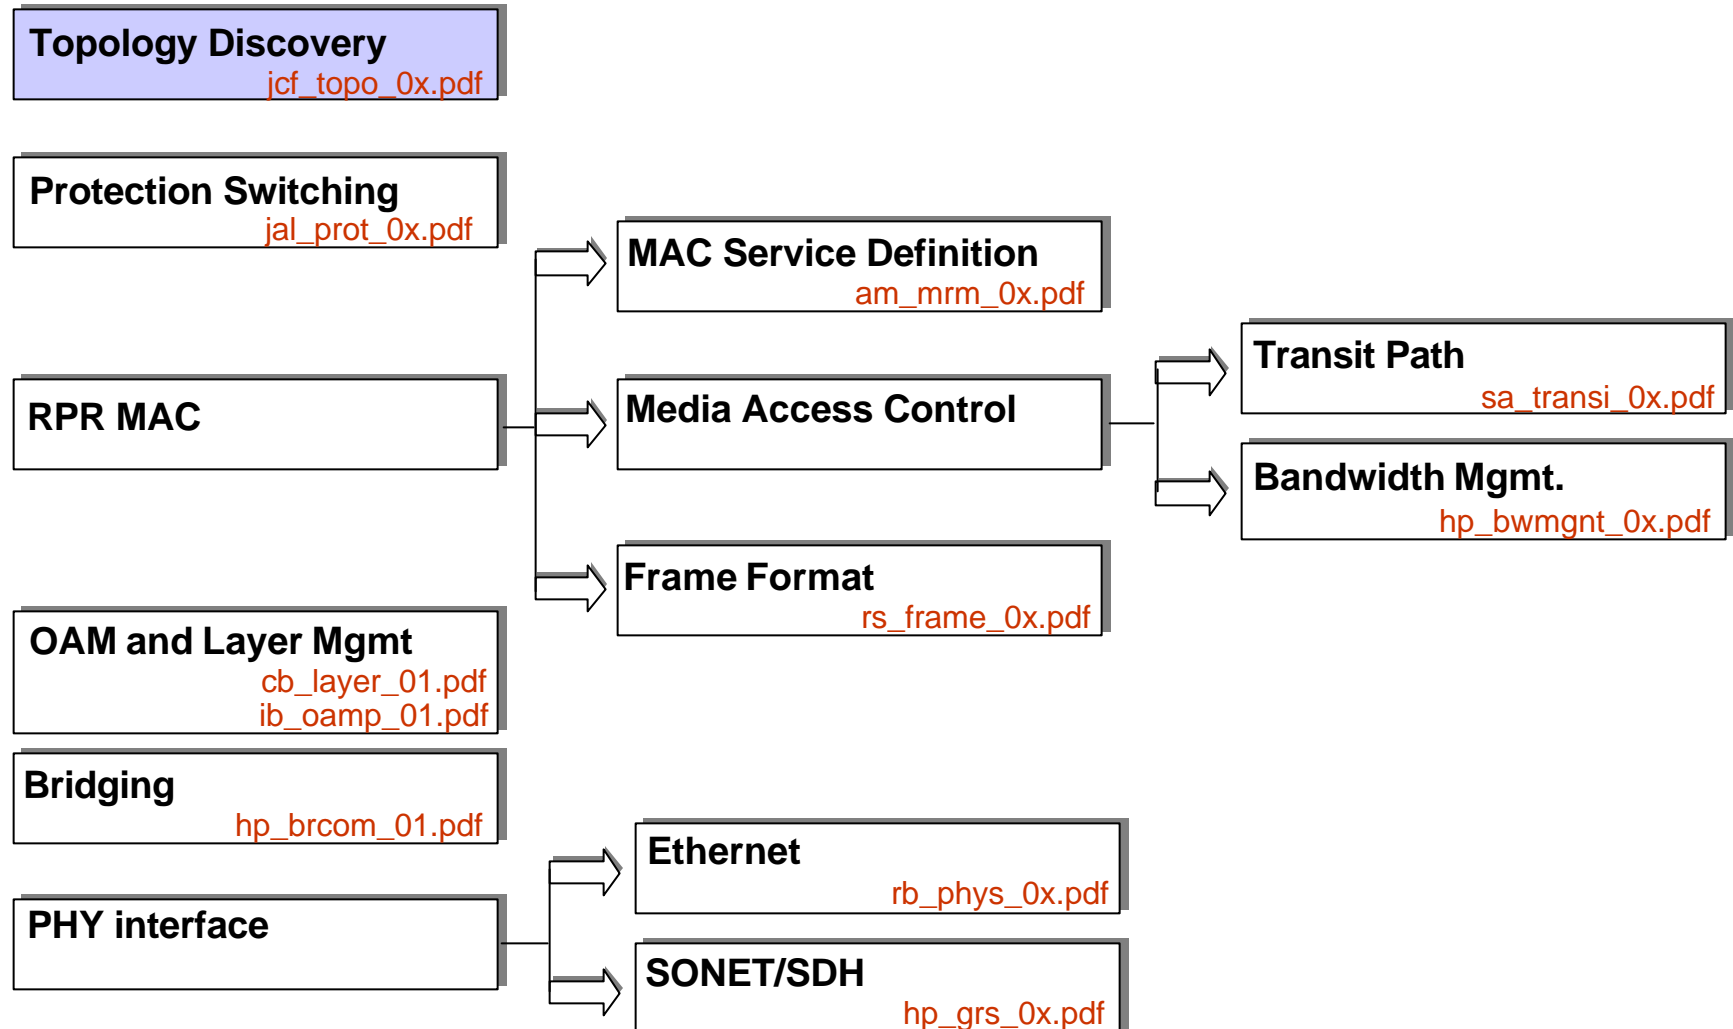
John Lemon, Lantern

Robin Olsson, Vitesse

Harry Peng, Nortel



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 Topology_StatusOperate 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
- 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 - Station with outdated or corrupted topology image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
- Broadcast on each ringlet with $TTL = Max_Ring_Size$
 - Removed by source

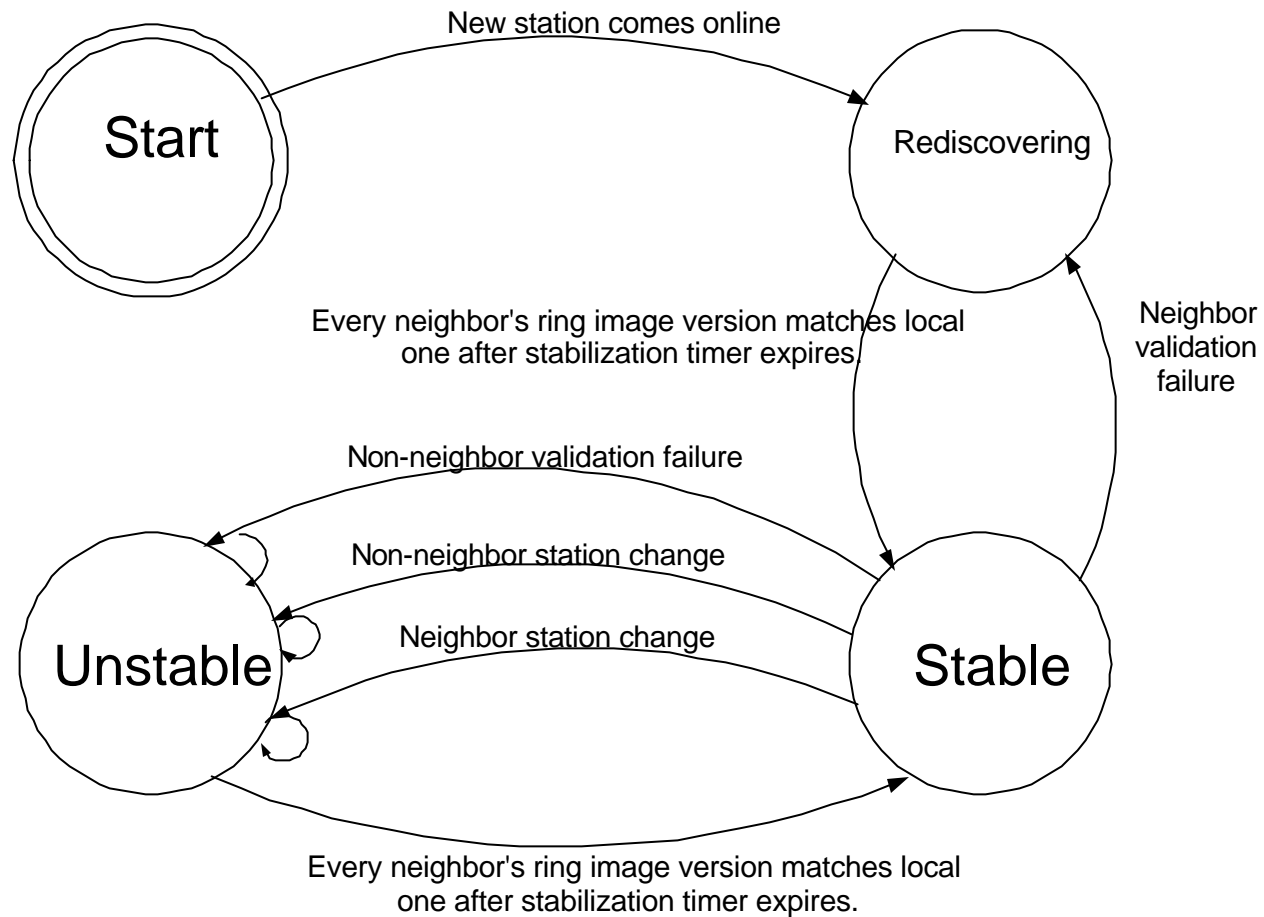




Configurable Parameters

- Neighbor_Hello_Timer
- Topology_Stabilization_Timer

State Diagram





State Diagram Details, 1

- Neighbor station change
 - Trigger
 - No Neighbor_Hello messages in 3 Neighbor_Hello Periods
 - A Neighbor_Hello from a new neighbor
 - Action
 - Increment the local Station_Image_Version
 - Broadcast a Topology_Status message
 - Replace the station information in the local topology image
 - Update the local Ring_Image_Version
 - Start Topology_Stabilization_Timer



State Diagram Details, 2

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



State Diagram Details, 3

- Neighbor validation failure or new station
 - Trigger
 - Neighbor_Hello's Ring_Image_Version != local one AND sending and local stations mutually reachable AND local Topology_Stabilization_Timer not running, 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 Topology_Status message
 - Send a Neighbor_Hello message
 - Start Topology_Stabilization_Timer



State Diagram Details, 4

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




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.
- While the Topology_Stabilization timer is not running, compare Ring_Image_Versions only if the two stations are mutually reachable.

Mutually Reachable Stations

- A and B are mutually reachable, and have the same Ring_Image_Version (=v1).
- C, D, and E are mutually reachable, and have the same Ring_Image_Version (= v2 != v1).
- A [B] will not compare Ring_Image_Versions upon receiving a Neighbor_Hello message from E [C] since A and E [B and C] are not mutually reachable.
- A[B] must check mutual reachability with respect to its neighbors.
- The MAC can report mutual reachability with respect to its neighbors to layers above.





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