

RPR Topology Discovery Proposal

Italo Busi, Alcatel

Jeanne De Jaegher, Alcatel

Jason Fan, Luminous

Chi-Ping Fu, NEC

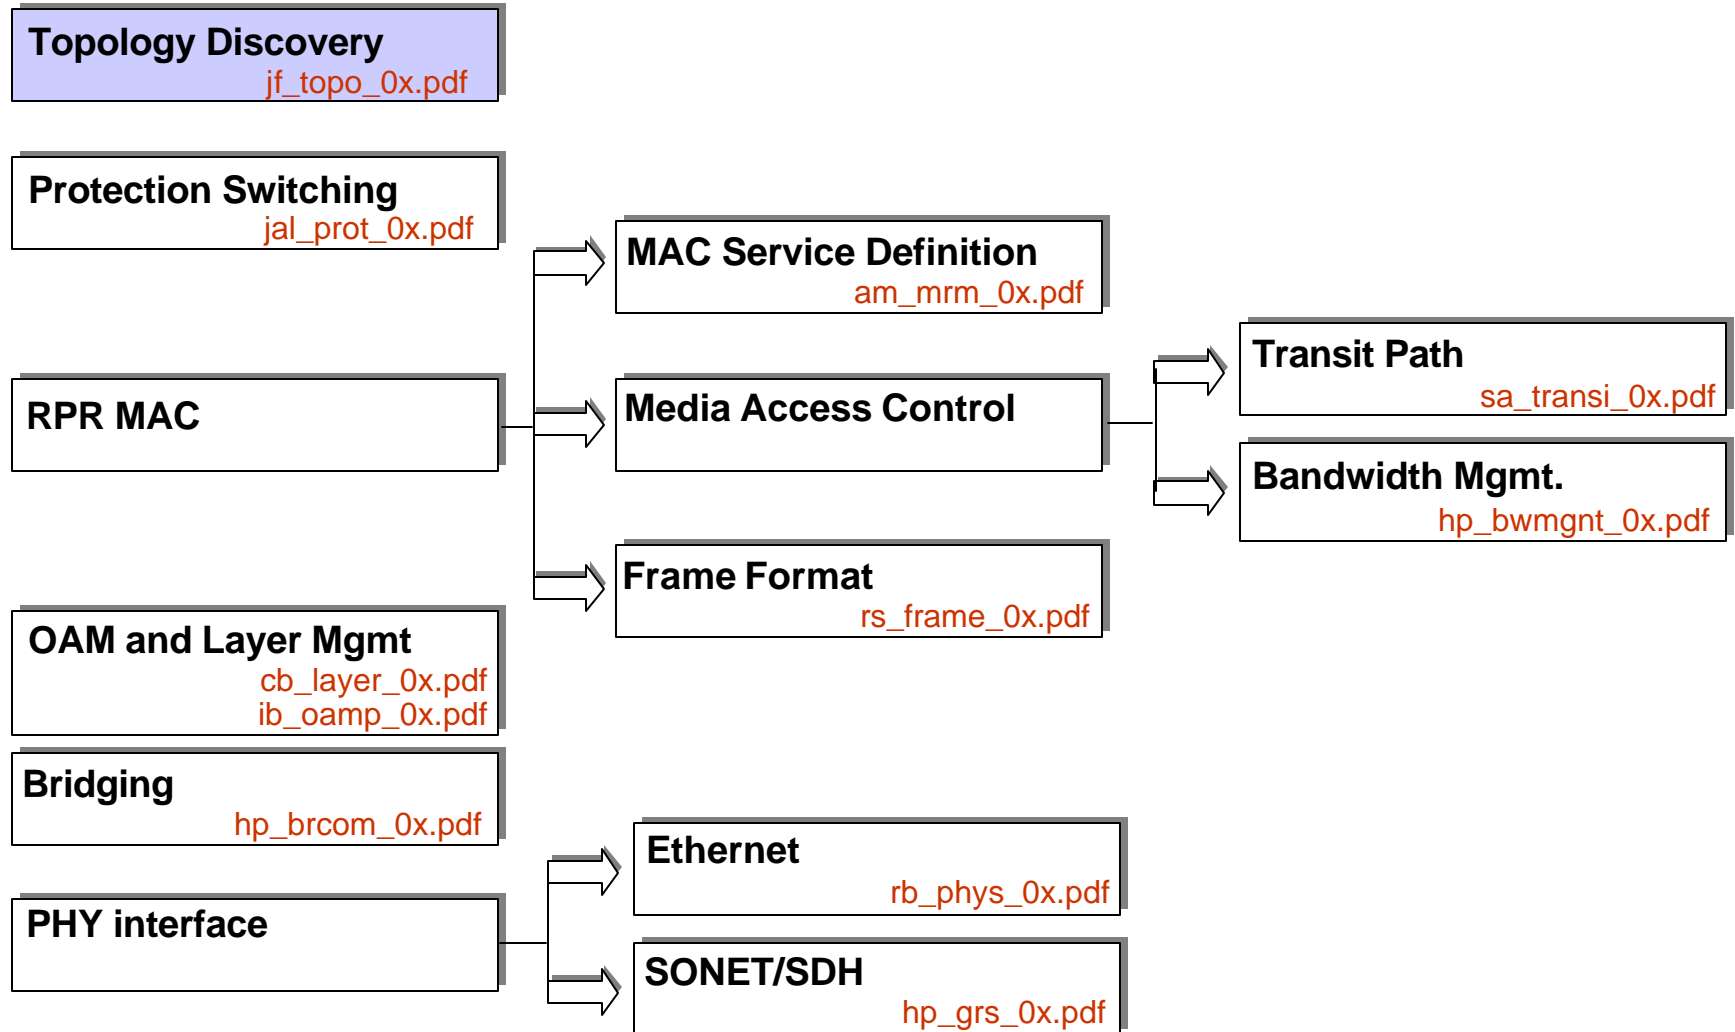
John Lemon, Lantern

Robin Olsson, Vitesse

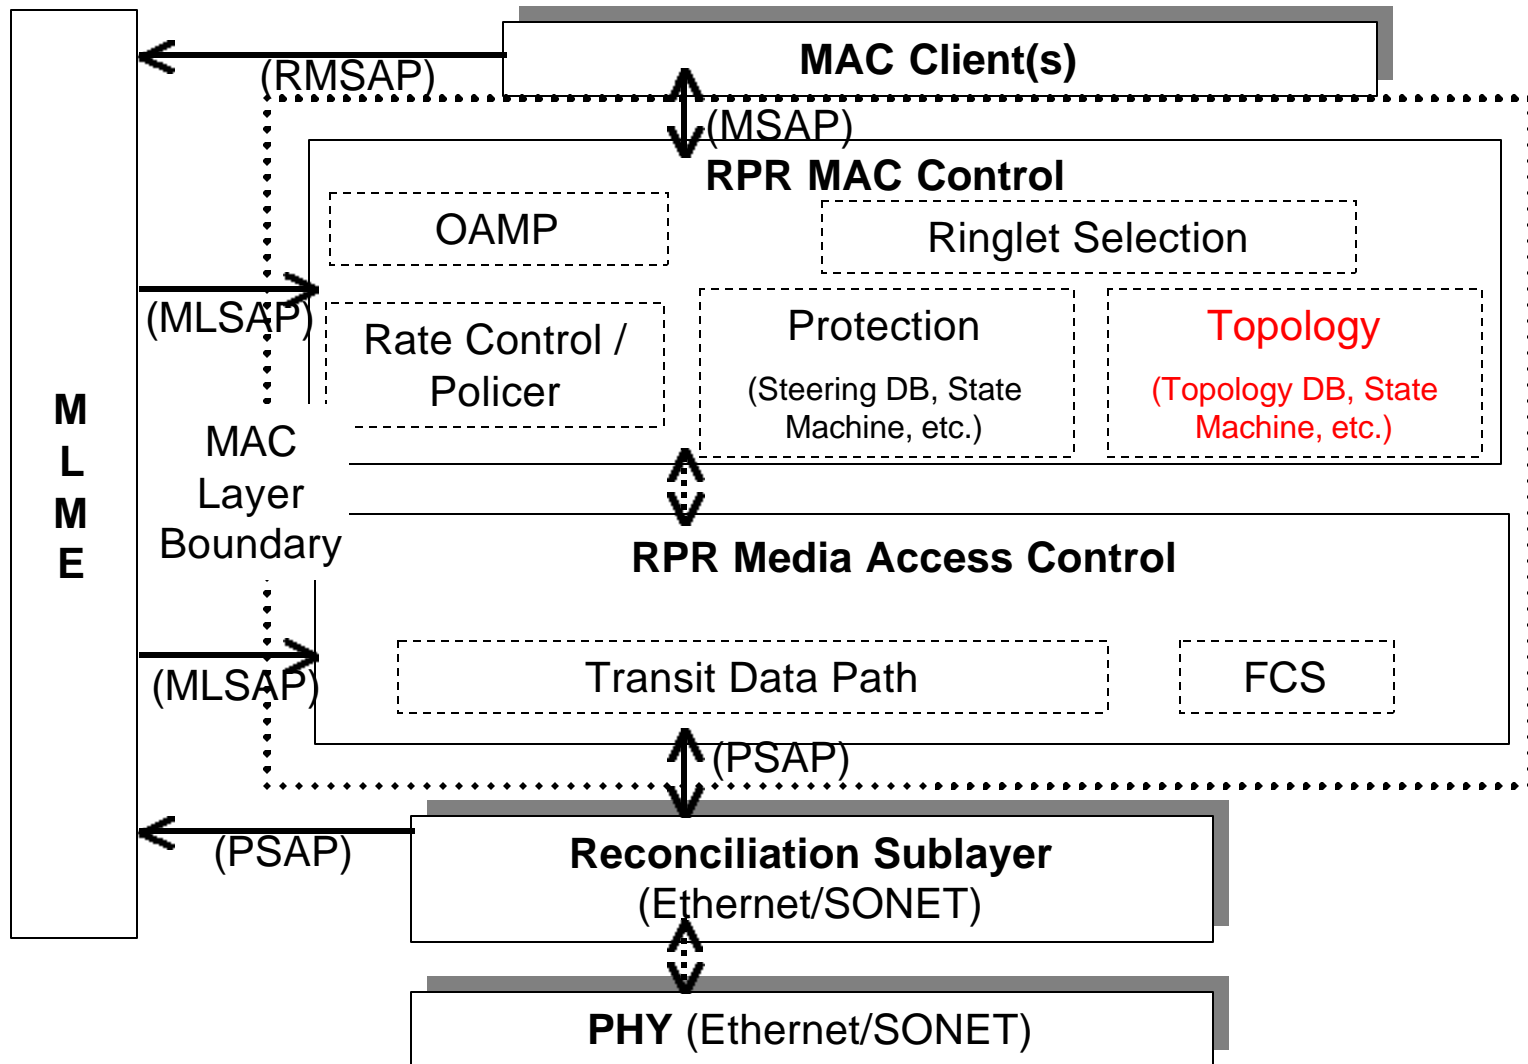
Harry Peng, Nortel

Yiming Yao, Luminous

Components of a complete RPR proposal



802.17 MAC Components



Information Sharing

- RPR Topology Image consists of station addresses, link connectivity between stations, and status for each link
- RPR Topology Image used by other algorithms
 - Steering algorithm uses Topology Image to know for which destinations steering is needed
 - Congestion avoidance uses Topology Image to know where congestion is being experienced
 - Ringlet selection uses Topology Image to know which ringlets are available to reach each destination

Goals of this Presentation

- In September the details of the topology discovery proposal were presented
 - Minor modifications made in the written proposal based on discussions with a variety of companies
- This presentation will cover
 - Discussion of methods for topology discovery
 - Additional simulation results

Method 1: Neighbor Discovery and Distribution

- Each station:
 - Discovers neighbors using periodic messaging
 - Broadcasts neighbor information on change
 - Collects neighbor information from stations detecting change
 - Verifies with adjacent stations that discovered topologies match
 - Updates topology database if verification successful
- On bring-up a station:
 - Follows the steps above with one modification
 - Broadcasts neighbor information with version indication set to 0 to cause other stations to send neighbor information to it

Method 2: “Chaining”

- Each station:
 - Sends a topology message on each ringlet periodically
 - Each station adds its MAC address to each topology message
 - Collects topology information from each ringlet
 - Needs to verify that information is received from each ringlet
 - Updates topology database if verification successful
 - Verifies with adjacent stations that discovered topologies match
- On bring-up a station:
 - Follows the steps above

Goals Met by Both Methods

- Scalable from 1 to 100's of stations
- Determine/validate connectivity and ordering of stations on ring
- Topology database updated based on dynamic addition and removal of stations to/from the ring
- Ensure all stations on the ring have a uniform and current image of the topology
- Immediate reaction to changes
- Operate without any master station on the ring
- Provide means of sharing additional information between stations
- Operate independently of and in the absence of any management systems

Additional Goals of Topology Discovery


- Robustness
- Minimize messaging overhead
- Discovery of spans on which there is a single failed link
- Usable with all supported topologies: multiple ringlet ring, linear (broken ring), and “star” (single station)

Robustness (I)

- Methods for tolerance of message loss
 - Neighbor Discovery
 - Send broadcasts on both directions on the ring
 - Validate discovered topology prior to updating stored topology through version comparison with neighbors
 - Chaining
 - Repeat topology discovery periodically
 - Validate discovered topology prior to updating stored topology (?) or update stored topology after identical topology discovered multiple times consecutively
- Neighbor Discovery is not restricted to its current validation mechanism

Robustness (II)

- Neighbor discovery uses broadcast, which has minimum interaction with the MAC as it is sent on the ring
- Chaining requires the MAC of every station to be able to update topology packets correctly for any station to discover the topology



Minimize Messaging Overhead

- M stations on ring
- Neighbor Discovery is very efficient
 - 1 broadcast per ringlet only on topology change, sent only by stations detecting a neighbor change
 - 1 unicast per ringlet only on topology change, sent by each station directly to a new station
 - Total of 3 broadcasts and M-1 unicasts per ringlet, only on topology change
- “Chaining” is less efficient
 - 1 “broadcast” per ringlet per station, sent every topology discovery period
 - Total of M “broadcasts” per ringlet, every topology period
- Simplification comes at the cost of efficiency

Discovery of Spans with Single Failed Link

- Neighbor discovery enables discovery of single failed link
 - Bandwidth efficiency consideration for steering protection: rings with steering can utilize spans with single available link
- “Chaining” does not enable discovery of single failed link
 - Bandwidth efficiency cost for steering protection
 - No bandwidth efficiency cost for wrapping protection

Multiple Ringlet Support

- Neighbor Discovery enables discovery of more than two ringlets, and does not preclude RPR from supporting these scenarios

Simulation Results (I)

- 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
 - Neighbor_Hello_Timer: 0.5 seconds
 - Topology_Stabilization_Timer: 2.5 seconds

Simulation Results (II)

- Scenarios
 - (0) Bring up all 256 stations at once
 - (1) Starting with link 0->1 removed, then add it back
 - (2) Two failures: links 0->1 and 127->128 are removed, then add 0->1 back
 - (3) Three Failures: span 64-65 plus links in (2) are removed, then add span 64-65 back
- In all the scenarios above, the new topology (for adding link or span back) is discovered within 2.5 seconds.

Simulation Results (III)

- Scenarios with lost Topology_Status messages
 - (4) Two messages lost in (1)
 - (5) Two messages lost in (2)
 - (6) Four messages lost in (3)
- In scenarios (4)-(6), the new topology is discovered within 5 seconds.

Topology Discovery Summary

- Many topology discovery objectives are met by both neighbor discovery and “chaining” methods
- Neighbor discovery method:
 - Is more efficient than “chaining” method
 - Can be designed to have lower failure probability than “chaining” method, at cost of efficiency
 - Increases bandwidth availability in rings using steering protection through usage of all available links
 - Has been simulated to show robustness to loss of multiple messages
 - Has been simulated to show topology discovery in rings with multiple link/span failures