



SLA Delivery over RPR

Krishna Pattabhiraman

Kumar Kovvali

Ramji Raghavan





SLAs: Carrier Perspective

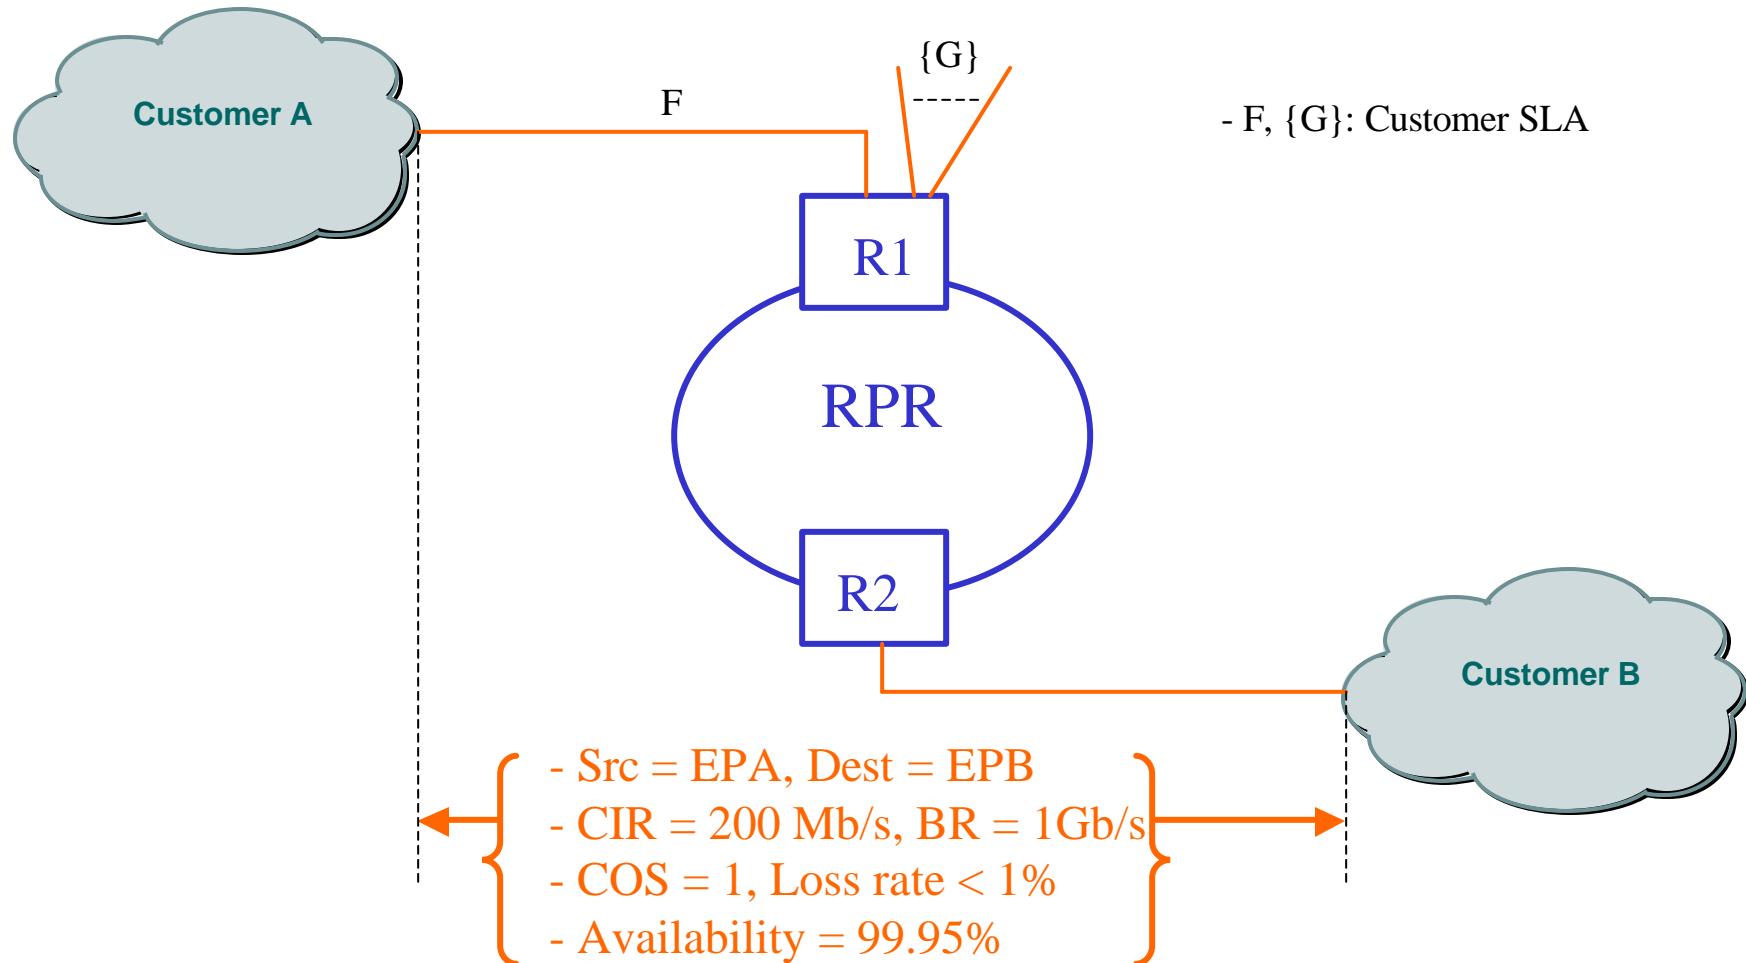
- Billable and enforceable services
- Better control of network resources and usage
- Challenge:
 - Adhere to the SLAs in a stat-muxed network



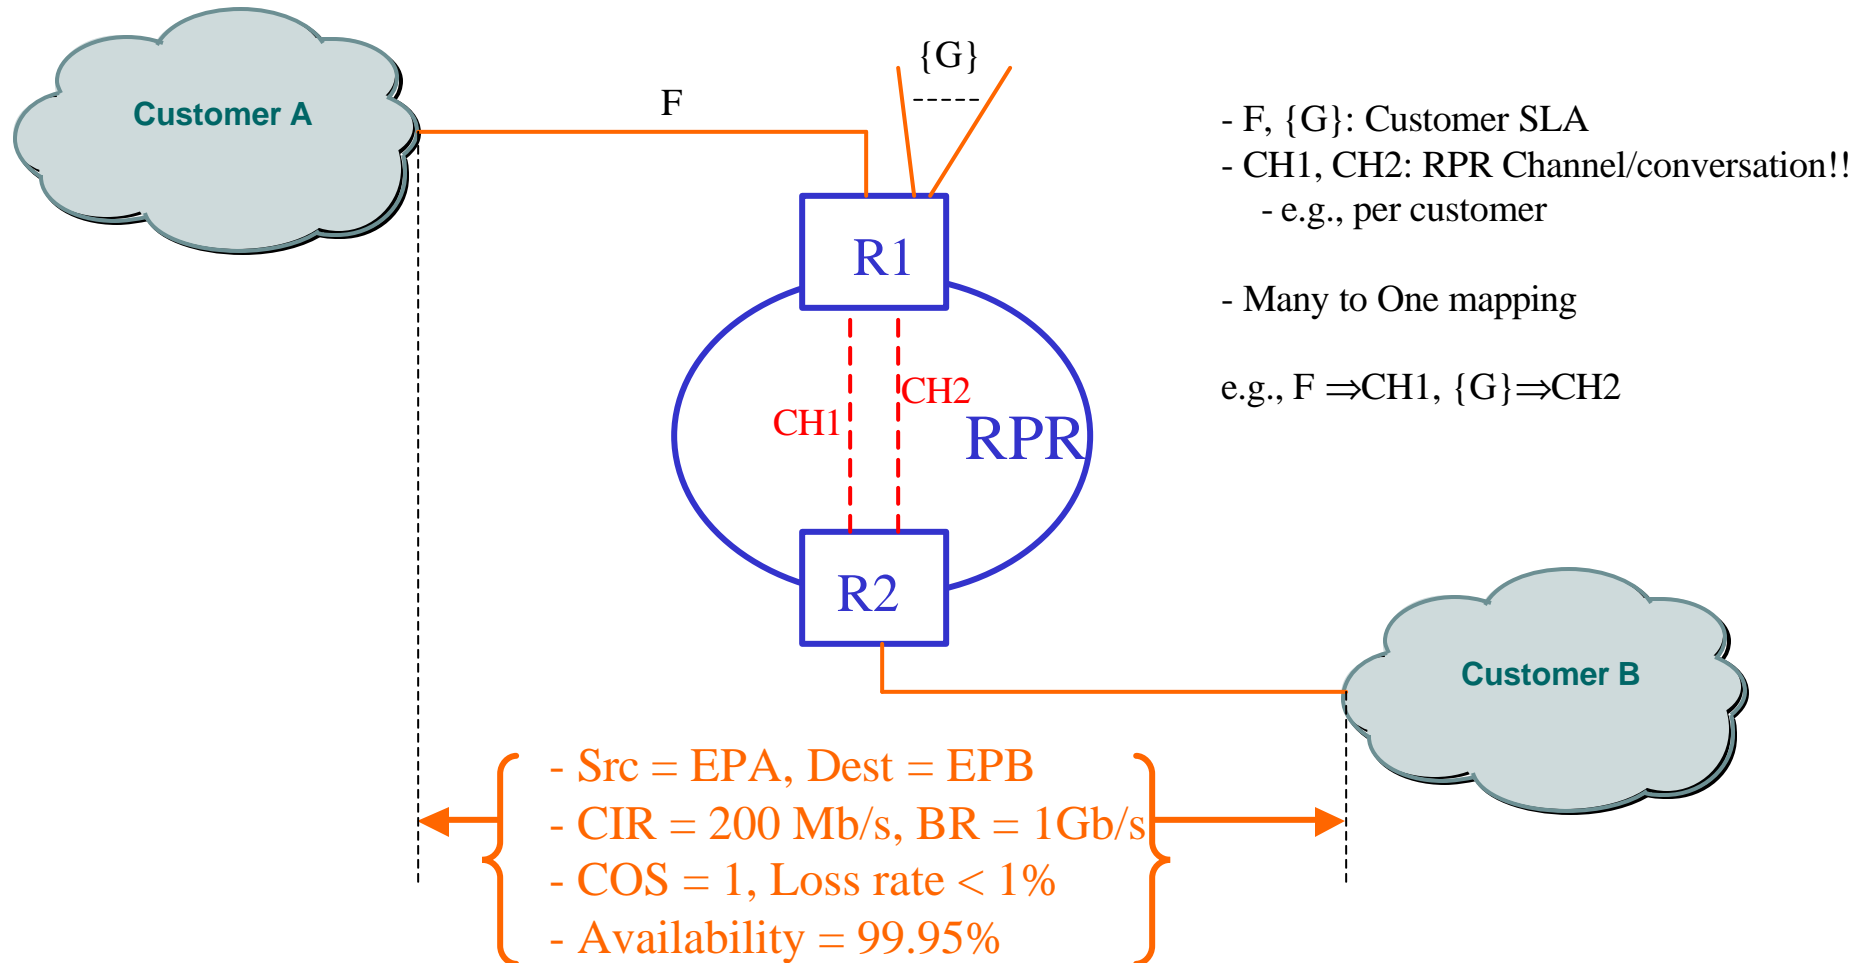
What are SLAs?

- Endpoint specification
- Traffic descriptor
 - CIR, BR, etc
- QOS/COS
 - Latency, loss ratio, etc
- Availability
 - Protection class, etc

SLA Example



SLA Example – RPR Context






SLA Enablers

- CAC or N/W Engg - NOT a MAC function
- Traffic policing – NOT a MAC function
 - E.g., Policer with a marker
- SLA-aware bandwidth management
 - Intelligent discard
 - SLA-aware media access
 - Intelligent congestion notification

Example



Stat-Mux example:

$$\text{SUM (CIR)} \leq \text{RBW}$$

$$\text{SUM (BR)} = K * \text{RBW}, K \geq 1$$

$$\text{RBW} = 2.48 \text{ Gb/s}$$

$$\text{F1: } \{0.38 \text{ Gb/s}, 1.0 \text{ Gb/s}\}$$


$$\text{F2: } \{0.2 \text{ Gb/s}, 1.0 \text{ Gb/s}\}$$

$$\text{F3: } \{1.0 \text{ Gb/s}, 2.48 \text{ Gb/s}\}$$

$$K = 4.48 / 2.48 = 1.8$$

Remaining capacity for burst = 0.9Gb/s

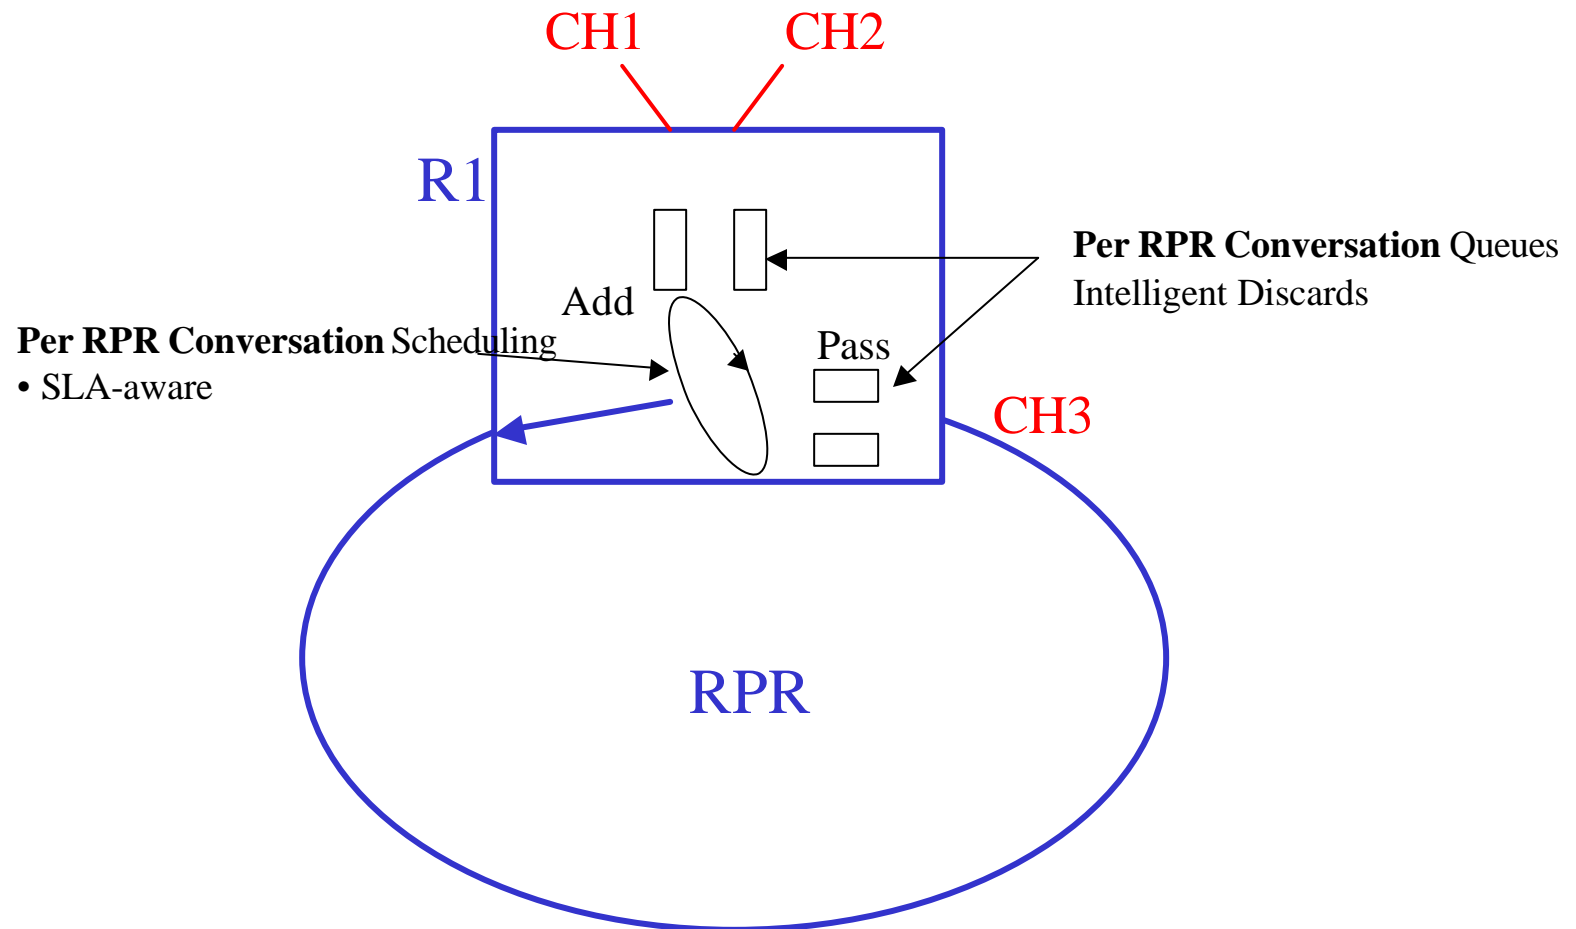
Example



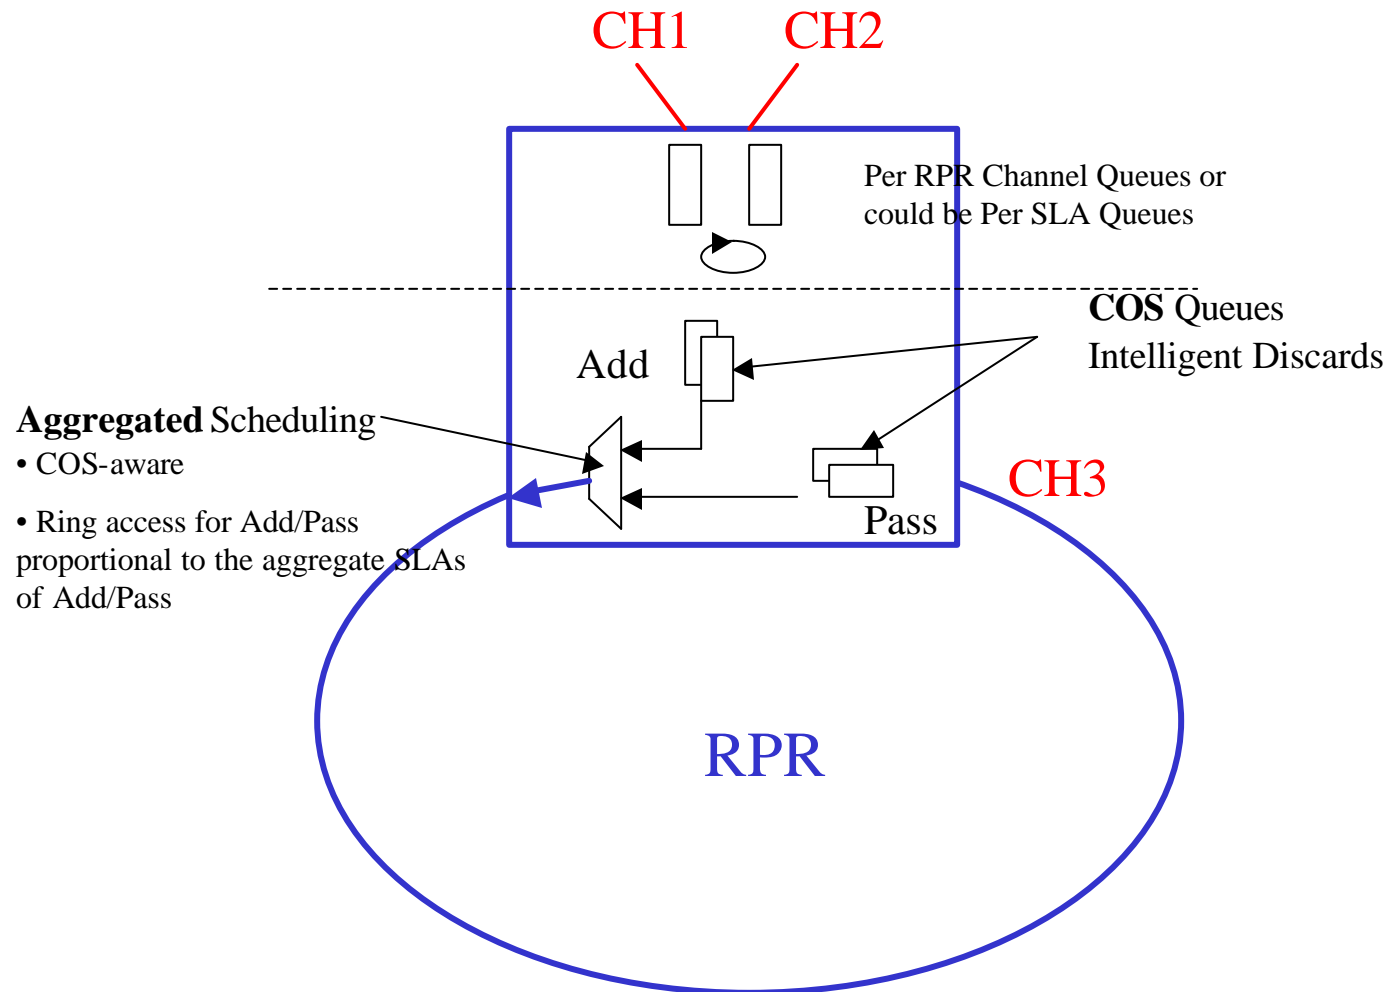
F1, F2, F3: SLAs

CH1, CH2, CH3: RPR Channels/Conversations

Bandwidth Management at each Node during Media Access (Ideal)



Bandwidth Management at each Node during Media Access(Approximate)



Example

Remaining capacity for burst = 0.9Gb/s

Example “fairness” policy under saturation:
Each node distributes the burst capacity equally.

Locally optimal decision @ R1 will yield:


CH1 = 0.68 Gb/s,

CH2 = 0.5 Gb/s,

CH3 = 1.3 Gb/s

This causes congestion at node R3!!!

-Backward congestion notification
(BCN)



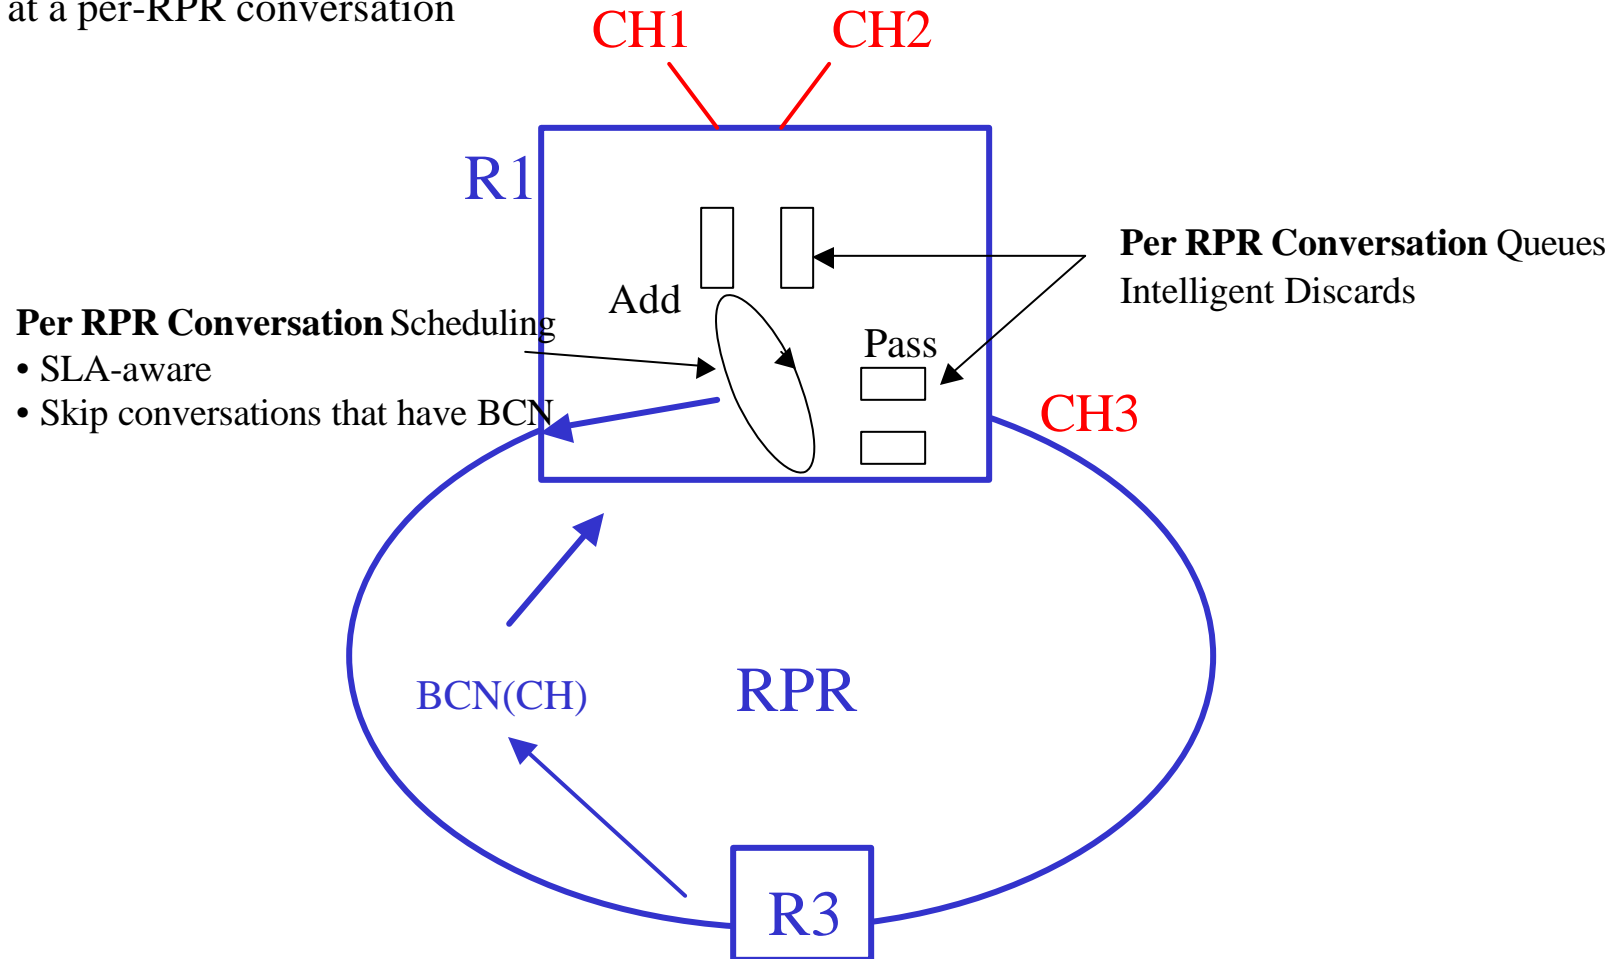


What is required?

- Knowledge of congestion at connection endpoints important
- Aggregated scheduling with aggregated BCN (hop-by-hop) can potentially violate individual SLAs
 - HOL Blocking

Ideal Distributed Solution

- Per-RPR conversation scheduling at each node's MAC
- BCN at a per-RPR conversation level





- Aggregated scheduling (based on aggregated SLAs) at each node’s MAC

The diagram illustrates the RPR (Resilient Packet Ring) architecture. It shows a ring topology with three nodes: R1, R2, and R3. The internal structure of node R1 is detailed, showing the flow of traffic through various queues and scheduling mechanisms.

Node R1 Internal Structure:

- Per RPR Channel Queues:** Two queues at the top, labeled CH1 and CH2, receive traffic from the ring.
- COS Queues:** Traffic from the channel queues is processed by COS (Class of Service) queues, which include intelligent discards.
- Aggregated Scheduling:** Traffic from the COS queues is sent to an aggregated scheduler, which is COS-aware.
- Add/Pass Mechanism:** The scheduler directs traffic to either the 'Add' or 'Pass' path. The 'Add' path is used for traffic going to a destination node that has BCN(R1) set.
- BCN(R1):** A bit control node (BCN) for R1, which is used to manage traffic flow.

Ring Topology:

- The ring consists of three nodes: R1, R2, and R3.
- Traffic flows clockwise from R1 to R2, then to R3, and back to R1.
- Node R3 is shown as a simple box, while R1 and R2 are more complex, indicating they have internal processing capabilities.

- ## Aggregated Scheduling
- COS-aware
 - Ring access for Add/Pass proportional to the aggregate SLAs of Add/Pass
 - Skip or drop packets of *ADD* channels going to a destination node that has BCN(R1) set.



Hooks in the MAC

- Input hooks to the MAC that indicates the proportion of the ring that corresponds to ADD traffic
 - $PASS_BW = RBW - ADD_BW$
 - Proportion determined such that the SLAs are adhered to, and fairness is achieved for the remainder of BW.
- Intelligent Congestion notification from the destination MAC
 - At least at a per source node level
 - Just affects the ADD traffic at the source node for this destination
 - Doesn't affect the PASS traffic at the intermediate nodes



Conclusions

- Key ingredient for ease of SLA delivery
 - SLA-aware media access
 - SLA-aware scheduling at each source node
 - Intelligent congestion notification from destination nodes