IEEE P802.1Qdd **R**esource **A**llocation **P**rotocol (**RAP**)

Domain Detection and Stream Protection

Feng Chen

Siemens AG

IEEE 802.1 Interim Meeting, Edinburgh, September 2019

Introduction

This deck discusses two topics for Qdd

- Domain detection
  - Recap of domain definition and detection in MSRP
  - RAP domain attribute specifications in Qdd/D0.0
  - Proposal for detection of directional domain boundary in RAP
- Stream protection
  - Priority regeneration at domain boundaries
  - Use of PSFP within a RAP domain

Recap: SRP Domain Definition

An **Stream Reservation Protocol (SRP) Domain**

- identifies a set of SRP stations, their ports and the attached individual LANs, that support **a given SR class** with a uniform configuration and parameterization.
- limits the scope of stream reservations for a given SR class within a region of connected bridges that support both the reservation functions (for stream reservation on control plane) and the QoS functions (for stream transmission on data plane) required by that SR class.
  - ❖ Consistency on management plane is not addressed by a SRP domain.
- prevents streams reserved within the SRP domain from being disrupted by the traffic entering from outside the domain by means of priority regeneration applied at the domain boundaries.

Recap: SRP Domain Boundary Detection

- MSRP detects **per-SR class** SRP domain boundary through the exchange of the MSRP domain attributes between link partners.
  - A consistency check is performed on the **SRclass Domain parameters** carried in the domain attributes to determine the Boolean value of the **per-port** and **per-SR class** *SRPdomainBoundaryPort* parameter. (if all matching -> core port; else boundary port)
  - In MSRP, only a 2-tupe **{SRclassID, SRclassPriority}** is exchanged for consistency check; other SRclass domain parameters such as shaper algorithm (CBSA only) and Class measurement interval (A: 125µs, B: 250µs) use only fixed values and need not to be explicitly exchanged.
- **SRclassVID** is also carried in MSRP domain attributes, but not considered a SRclass Domain parameter and playing no role in domain boundary detection.
  - A VLAN is used by MSRP to control propagation of reservation attributes (MAP context).
  - Multiple VLANs can be used by a SRP domain for a given SR class.
  - A single VLAN can be shared by multiple SRP domains for different SR classes.

  => VLAN is not a per-SR class domain property

Example SRP Domain Detection

- One SRP Domain exists for each SR class (SRclassID) mapped to the same priority.
- SRP Domain boundary exists only outside a bridge; bridge-internal domain boundary, e.g. when two bridge ports use different priorities for the same SRclassID, is not supported.

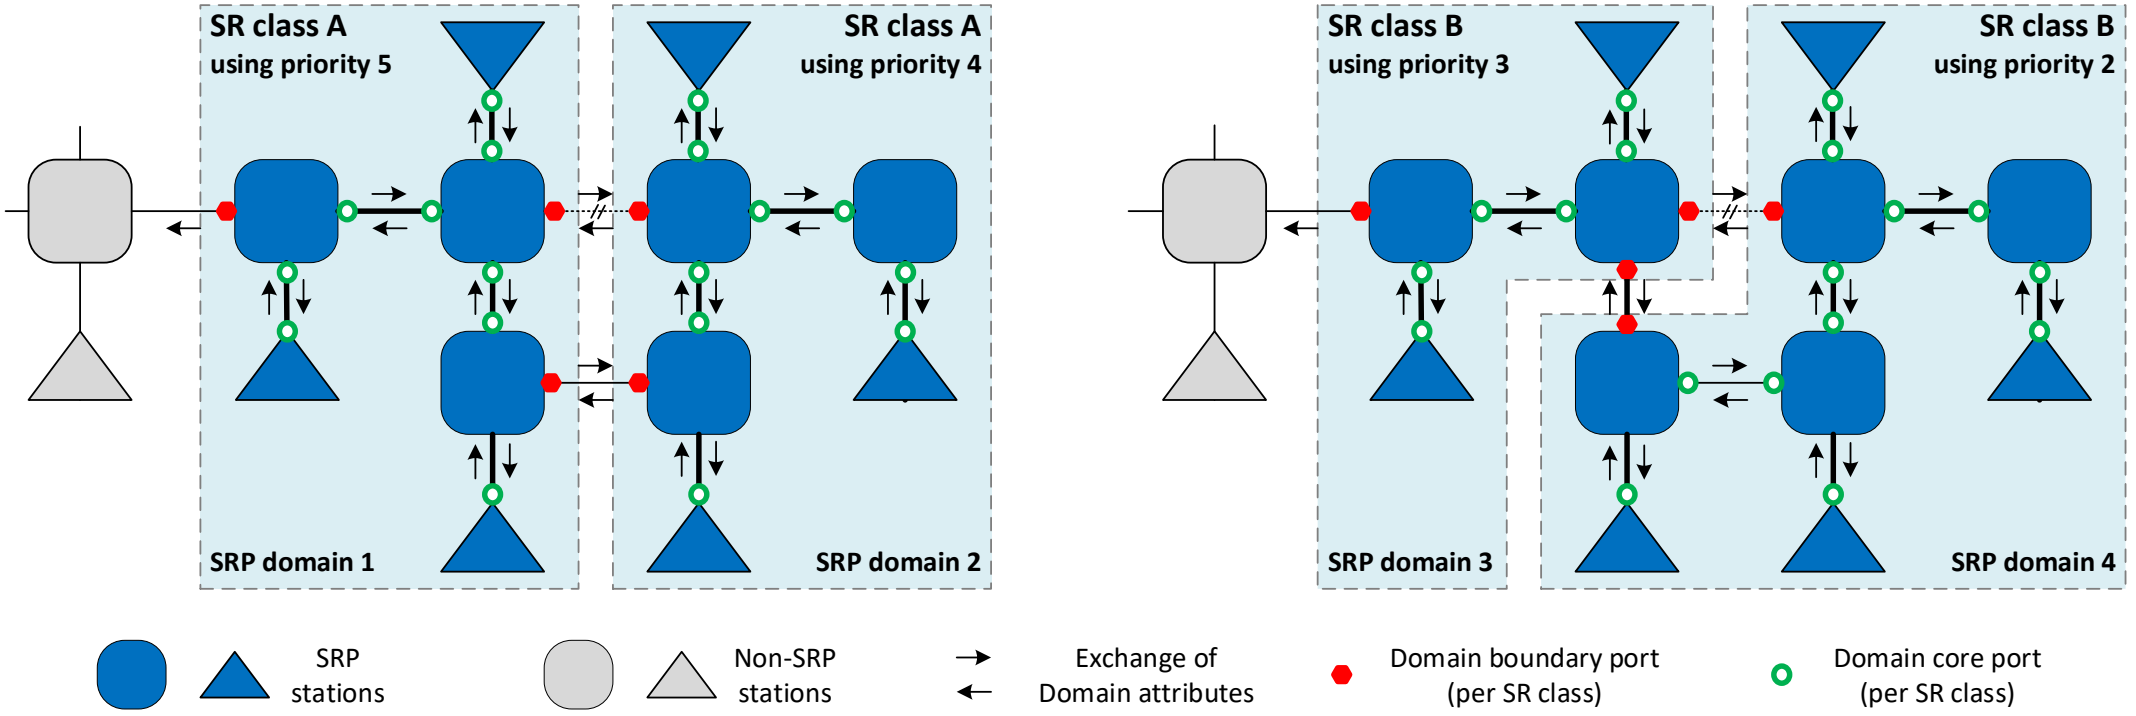

RAP Domain Specifications in Qdd/D0.0

RAP domain inherits the basic principles of SRP domain

- a per-SR class property for SR class based stream reservations
- a region providing common control-and data-plane capabilities for a given SR class
- automatic domain boundary detection through exchange of domain attributes

RAP domain attributes specified in 99.5.3 of Qdd/D0.0

- per-SR class *SRclass Descriptor sub-TLV* contains the settings of SR class domain parameters that are subject to consistency check in detection of domain boundary
  - The *Reservation Mode* (99.5.3.1.3) parameter encodes an organization-specific value (OUI/CID+Index) that describes a suite of tools (QoS function, latency calculation algorithm, TSpec format etc.) used by the associated SR class.
- per-device *Stream VID List sub-TLV* contains the VIDs available for use by streams
  - not associated with a particular SR class, no effects on domain detection
  - notifying end stations of the stream-related VLAN settings in the network

=> RAP domain attributes are fully configurable via managed objects.

	Octet	Length
1 or more SRclass Desriptor sub-TLVs (99.5.3.1)	1	variable
Stream VID List sub-TLV (99.5.3.2)	variable	variable

Figure 99-xx—Value of Domain Attribute TLV

		Octet	Length
streamVlanIdList (99.4.1.3)	VID 1	1	2
	VID 2	3	2
	…	5	2

Figure 99-xx—Value of Stream VID List sub-TLV

	Octet	Length
SRclass ID (99.5.3.1.1)	1	1
Priority (99.5.3.1.2)	2	1
Reservation Mode (99.5.3.1.3)	3	4
Class Measurement Interval (99.5.3.1.4)	7	4
Stream Max SDU Size (99.5.3.1.5)	11	2
Stream Min SDU Size (99.5.3.1.6)	13	2

Figure 99-xx—Value of SRclass Descriptor sub-TLV

A Gap in SRP Domain Detection

An SRP Domain attribute sent through a bridge port carries information only about the capability of transmitting streams on the declared SR class(s) through this port.

If this port (saying **X**) does not support transmission on a given SR class (saying **A**), **X** will become a boundary port for **A**. This causes all stream requests (Talker Advertise) on **A** sent from an upstream station to **X** to become a Talker Failed, meaning that streaming on **A** entering from X is always blocked.

However, such behaviors disregard the fact that, if at least one of the other ports on the same bridge is capable of transmitting on **A**, e.g. **Y** is a core port on **A**, streaming on **A** from **X** to **Y** is still possible.

Example: Domain Boundary Detection in MSRP

Reservation procedures for s1 and s2 are failed by MSRP due to propagation of Talker Advertise across domain boundaries, which does not reflect that B2 does support the required SR class at the transmission ports for s1 and s2.

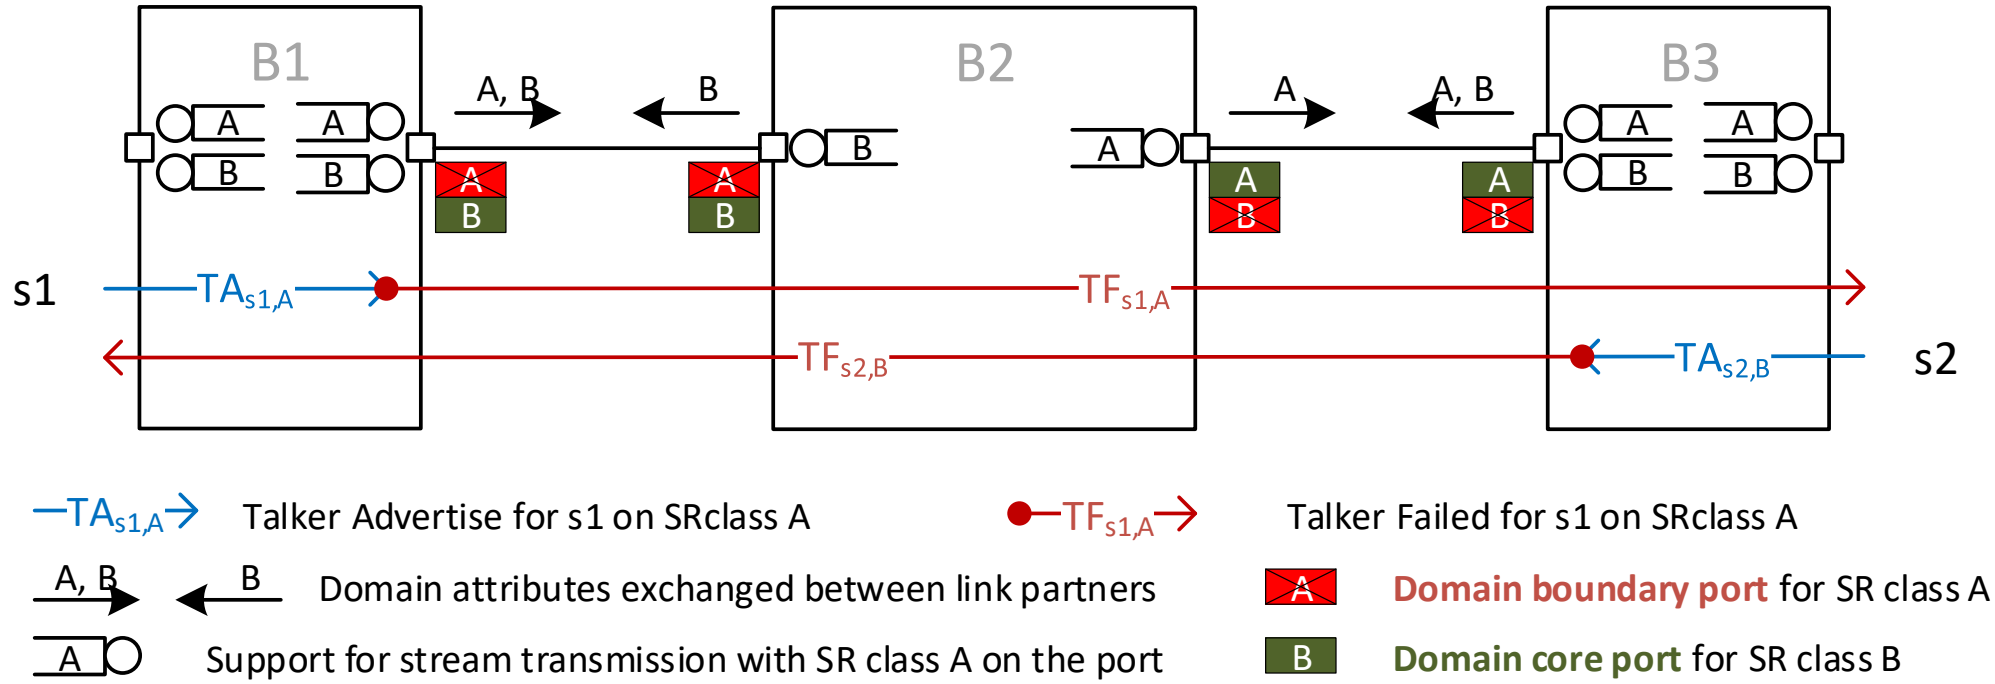

Proposal for Qdd:
Directional Domain Boundary Detection in RAP

RAP domain attributes sent through a port carry information about both transmission and reception capabilities of this port for a given SR class,

- **Transmission-capable**: this port supports that SR class
- **Reception-capable**: at least one of the other ports on the same bridge supports that SR class

which are used to determine the values of two (*per-port, per-SR class)* variables, *RAPdomainBoundaryTxPort and RAPdomainBoundaryRxPort.*

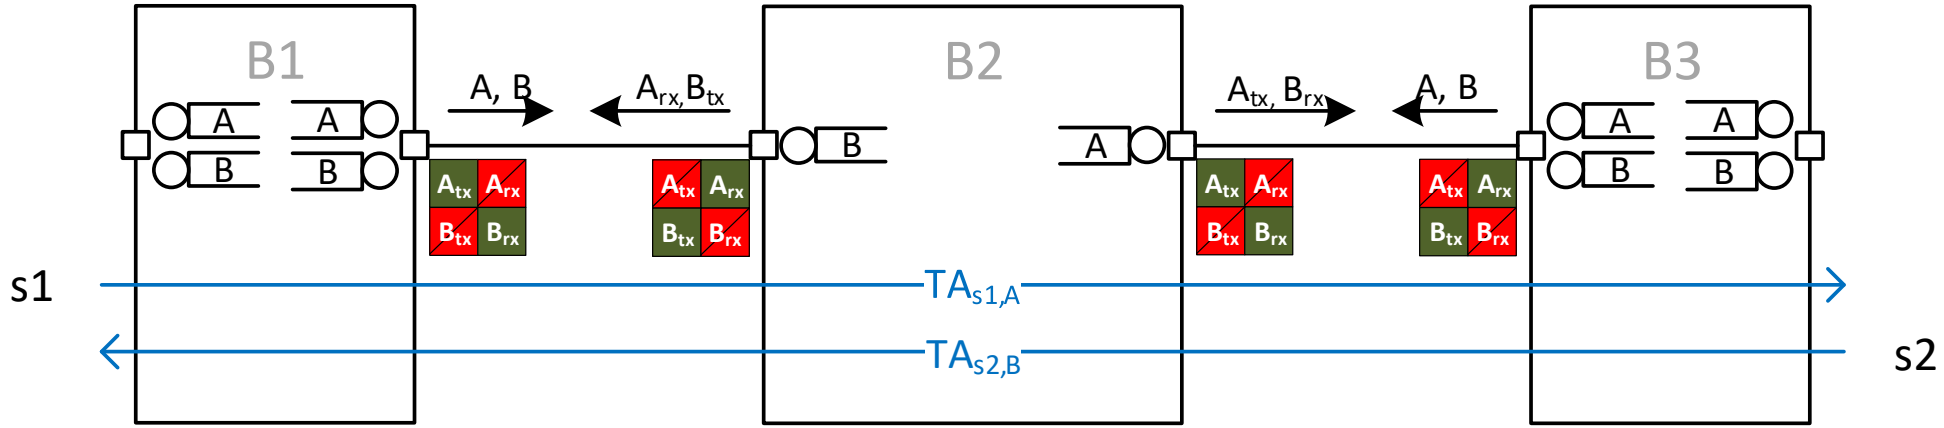

A_{tx} A_{rx}
B_{tx} B_{rx}

Domain core Tx port and **domain boundary Rx port** for SR class A

Domain boundary Tx port and **domain core Rx port** for SR class B

A_{tx}, B_{rx} Domain attributes carrying Tx/Rx capabilities for each supported SR class

Stream Protection in MSRP

MSRP protects against unauthorized access to the stream queues by applying priority regeneration to the traffic that enters from **outside a SRP domain** and carries a priority associated with a SR class of that domain at the domain boundaries (on ingress).

- remapped to a priority associated with a numerically lower traffic class than all SR classes on that bridge (adjustable via management).

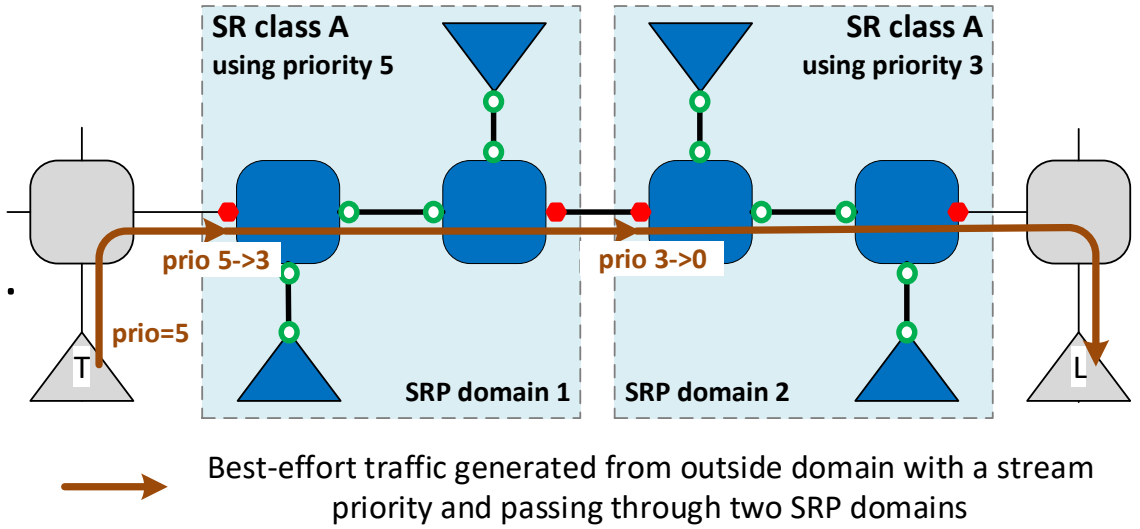

[Gap]: MSRP does not protect against misuse of stream queues caused by traffic generated from **inside a SRP domain** by a misbehaved SRP-capable station, which may

- start transmitting a stream before or even without running the MSRP procedures
  - The *Dynamic Reservation Entries* (8.8.7) in the FDB can only block a stream that already used the MSRP to start the reservation process but has not been successfully reserved.
- transmit the non-stream traffic (e.g. best-effort) using a priority associated with a SR class.

Inside-Domain Stream Protection in RAP

In addition to continuing support for priority regeneration on domain boundaries, RAP should also provide a means for inside-domain stream protection.

Goal: prevent the traffic generated inside the RAP domain and carrying a stream priority associated with a SR class but not having a valid reservation from entering a stream queue.

Solution space:

- Need a data-plane solution, because problems may be caused by a RAP-capable talker intentionally not running the control-plane reservation procedures.
- Priority regeneration at domain core ports doesn't help because it's a per-port per-priority operation done at EISS of a reception port and does not distinguish frames of the same priority in terms of their reservation status.

Example Use of PSFP for Inside-Domain Stream Protection

Per-Stream Filtering and Policing (PSFP) in 802.1Q can be used in combination with RAP to provide inside-domain stream protection. The intention is to **apply certain filtering actions to any frames trying to entering a stream queue but not found in the *Dynamic Reservation Entries***.

- At network start-up, create a PSFP stream filter instance using a wild-card *stream_handle* and a *priority* value for each stream priority on each RAP domain core port. There are two alternative consequent actions:
  - **IPV**: use a permanently open stream gate and an IPV associated with a numerically lower traffic class than all SR classes on that bridge.
    - Pros: connectivity preserved; Cons: IPV effective only locally, needing such a stream filter on every bridge port
  - **Discard**: use a permanently closed stream gate.
    - Pros: stream filter needed only at bridge ports adjacent to talker end-stations; Cons: connectivity disrupted

The author prefers the “Discard” option, because the frames captured by this filter are generated by a RAP-enabled device which has violated the rules of generating traffic within a RAP domain.

Example Use of PSFP for Inside-Domain Stream Protection (cont.)

- Each time when RAP creates a *Dynamic Reservation Entry* for a stream, create a new PSFP stream filter instance, which excludes that stream from the filtering actions done by the wild-card stream filter instance created initially.
  - For the protection purpose discussed here, a simple “pass-through” filter is sufficient. For more comprehensive stream protection, other PSFP functions such as Maximum SDU size filter and flow meter can be used.

- Similarly, each time when RAP deletes an existing *Dynamic Reservation Entry* for a stream, delete the corresponding PSFP stream filter instance created previously for that stream.

[Note]: If the group considers this solution valid and useful, the editor can add it to Qdd drafts, e.g. in terms of a note or an informative Annex.

Next Step

- The first Qdd draft D0.0 was uploaded as 802-1Qdd-d0-0 and introduced in a presentation dd-chen-D0-0-introduction-0719-v01 at the July plenary 2019.

- The editor is preparing a new draft D0.1 and plans to start the first Task Group ballot prior to the November Plenary 2019. D0.1 will
  - incorporate the items shown in this deck upon agreement of the group
  - fill up the remaining items in 99.4 and 99.5 unfinished by D0.0
  - provide an introduction to RAP in 99.1
  - use 802.1Q FrameMaker template (if made available to the editor in time)