Time Aware Shaping and Pre-Emption

Christian Boiger
christian.boiger@hdu-deggendorf.de

IEEE 802.1 Interim
January 2012
Munich, Germany

Traffic Classes and Shaper

- Traffic classes:
  - Class A
  - Class B

- Possible AVB Gen2 Shaper:
  - 125µs CBS (AVB Gen1 Class A) – CBS125
  - 250µs CBS (AVB Gen1 Class B) – CBS250
  - 125µs CBS w/pre-emption – CBS125P
  - 250µs CBS w/pre-emption – CBS250P **?**
  - TAS with pre-emption in bridges and end stations – TAS

Credit Based Shaper with Pre-Emption

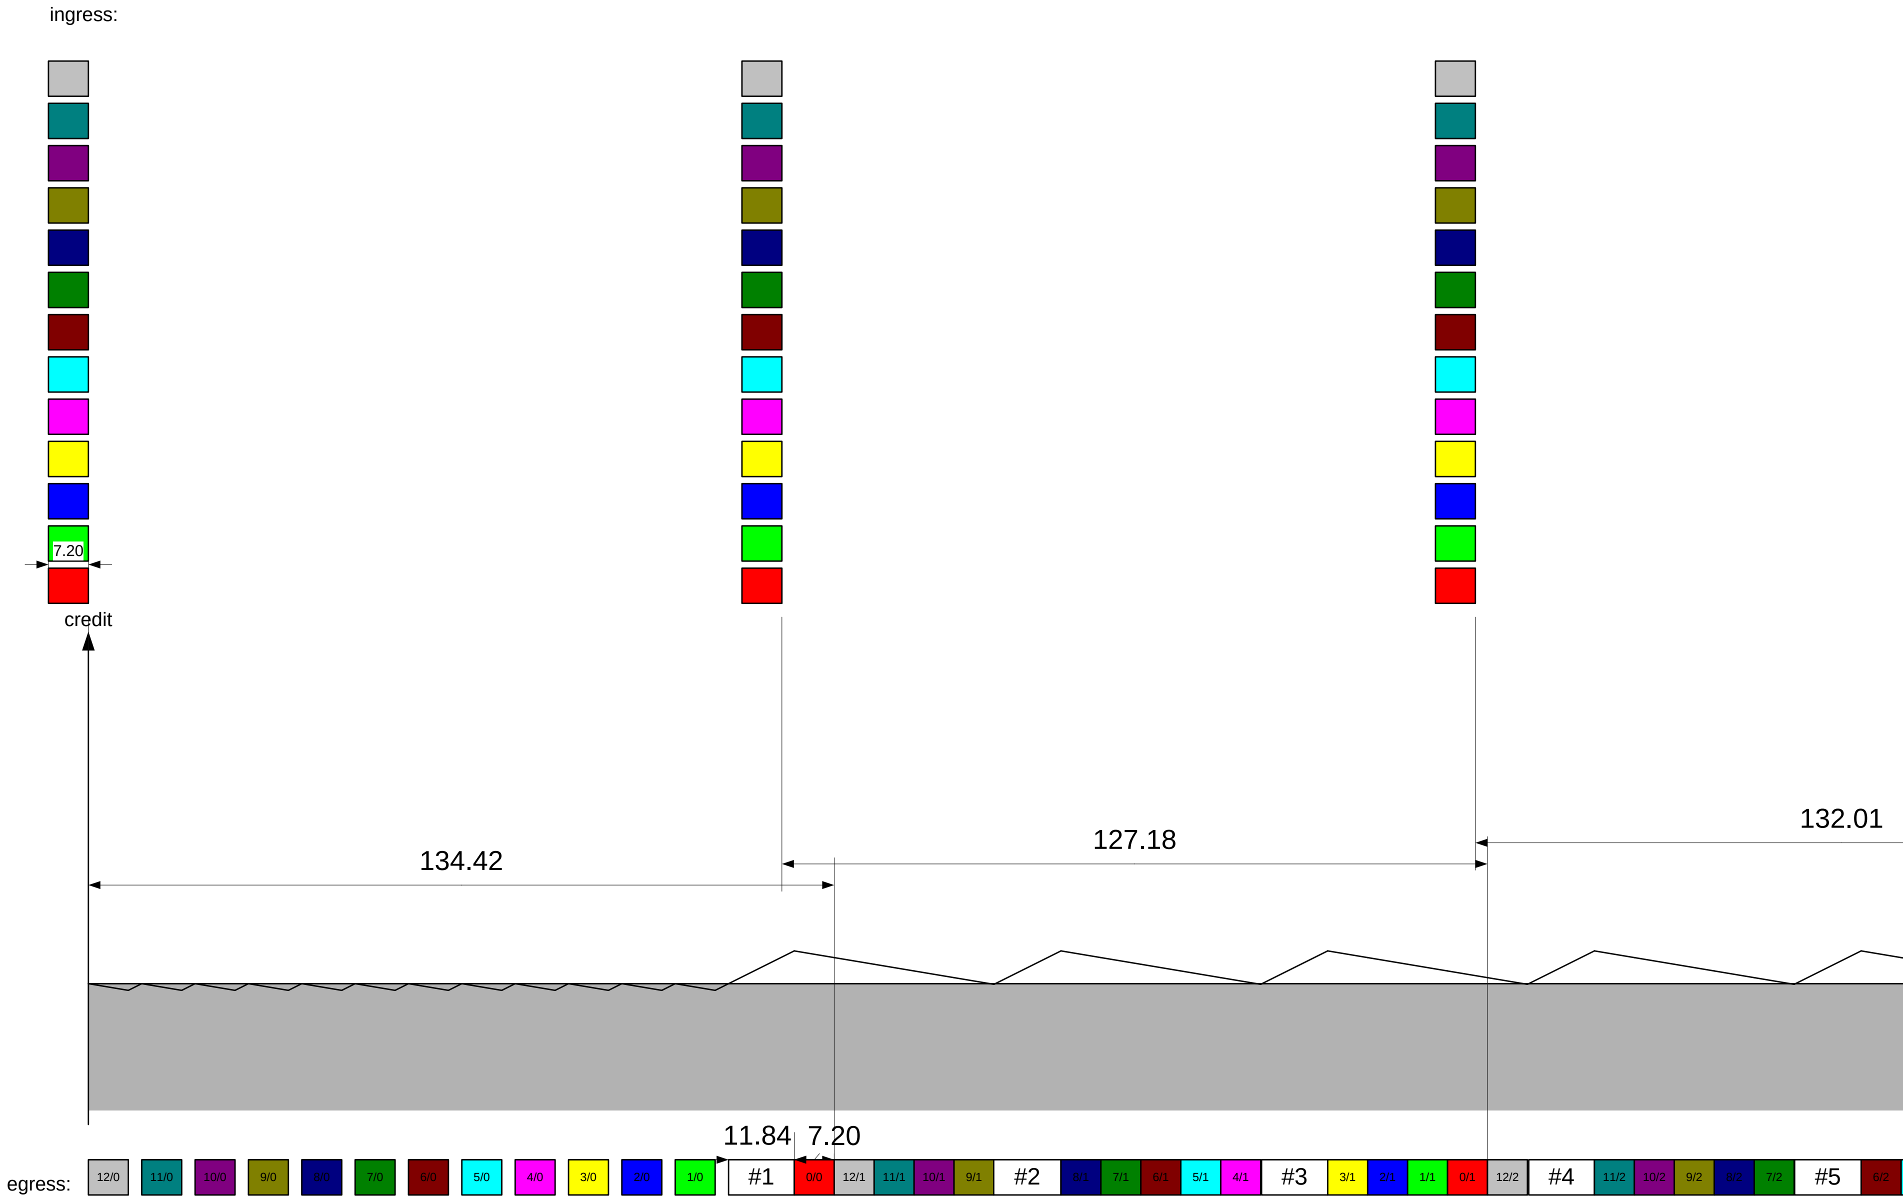

Credit Based Shaper with Pre-Emption

- CBS125P is used as Class A shaper
- Pre-empt if something is in the queue and credit >= 0
- Worst case talker latency (Class A, CBS125P):
  - FE interfering frame CBS125: 123.36 µs
  - FE interfering frame CBS125P: 11.84 µs (128 bytes fr.)
  - ➔ 123.36 µs - 11.84 µs = 111.52 µs
  - GE interfering frame CBS125: 12.336 µs
  - GE interfering frame CBS125P: 1.184 µs
  - ➔ 12.336 µs – 1.184 µs = 11.152 µs

Credit Based Shaper with Pre-Emption

- Worst case talker latency (Class A, CBS125P):
  @FE 138.76 µs (Class A, CBS125: 250.28 µs)
  @GE 126.376 µs (Class A, CBS125: 137.528 µs)
- Smaller worst case latencies
- (Q)BurstSize = f(MaxLatency) => shorter bursts
- ➔ Less problems in bigger “weird” networks

- But: CBS is almost the worst case (bandwidth efficiency) for pre-emption => CBS produces many framelets

Time Aware Shaper with Pre-Emption

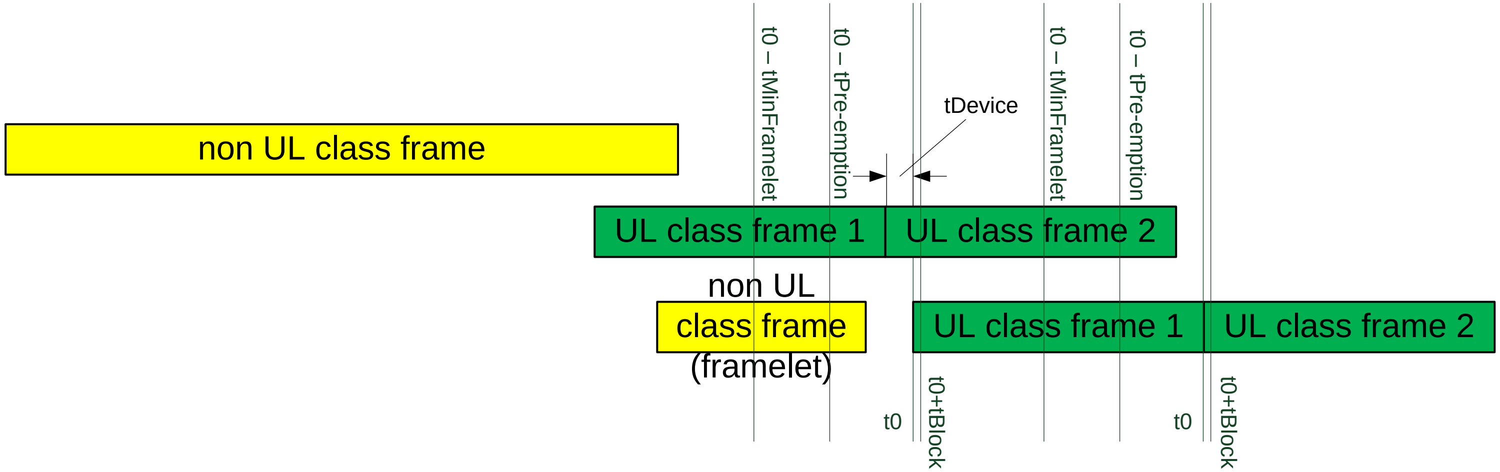

1. Block all non Ultra-Low Latency (UL) class queues at t0-tMinFramelet
2. Block the transmission of framelets at t0-tMinFramelet (part of the pre-emption PAR?)
3. Start pre-emption mechanism at t0-tPre-emption (time depends on how the pre-emption mechanism works)

Time Aware Shaper with Pre-Emption

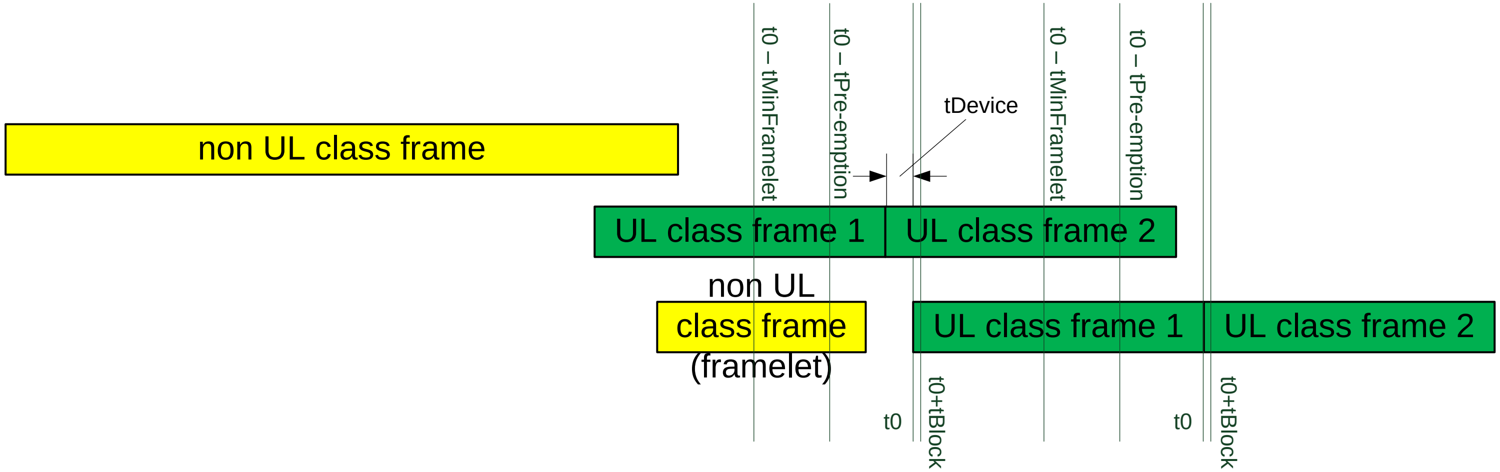

4. De-block the UL class queue at t0
5. Block the UL class queue at t0+tBlock (very short time after the scheduled start of the UL class frame)
6. De-block the framelet transmission at t0+tBlock
7. De-block the non UL class queues

Time Aware Shaper – Per Transmission Period

- One UL class time slot per transmission period

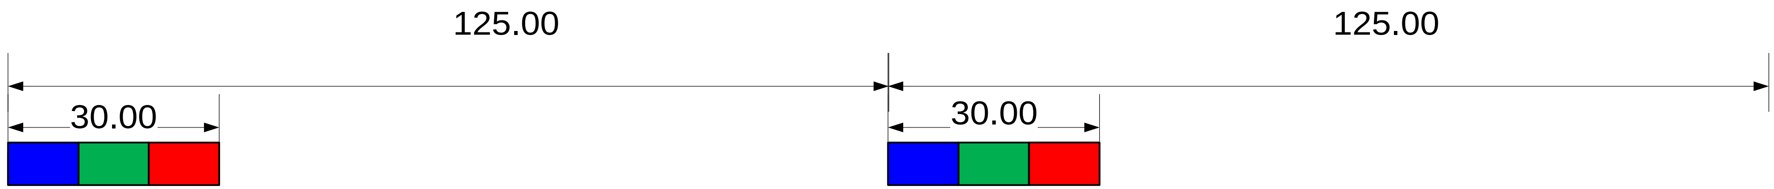

Disadvantages:

- Wasting bandwidth (gaps between the frames)
- One fixed transmission period
- Problems with complex topologies and stream routings
- Higher transmission periods (e.g. 500 µs) can cause a long interference with the Class B streams (e.g. CBS125) ➔ long Class B bursts

Time Aware Shaper – Impacts on Class B

- As the TAS has a very deterministic behavior Class B still has a determinable worst case latency
- But if the UL class streams have a higher transmission period than 125 µs and the streams together are bigger than a maximum size interfering legacy frame, the UL Class has an not negligible effect on the Class B (CBS125) streams
- ➔ Longer bursts
- ➔ Higher Class B CBS125 latency than AVB Gen1 Class A

Time Aware Shaper – Per Stream

- Each stream has its own transmission time slot

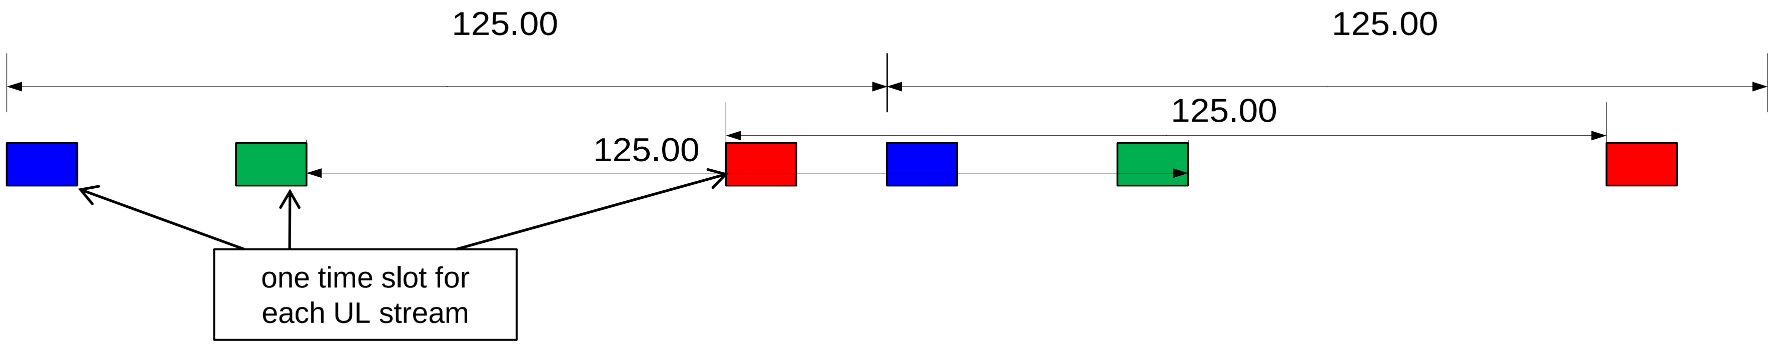

- Bandwidth efficient (but problems with very small gaps between the streams)
- Highly configurable
  - Variable transmission periods
  - Complex network topologies
  - Support of many talker and listener

Time Aware Shaper – Time Synchronization Error

- Two possible ways to deal with the Time Synchronization Error in each bridge:
- Latency optimized version:

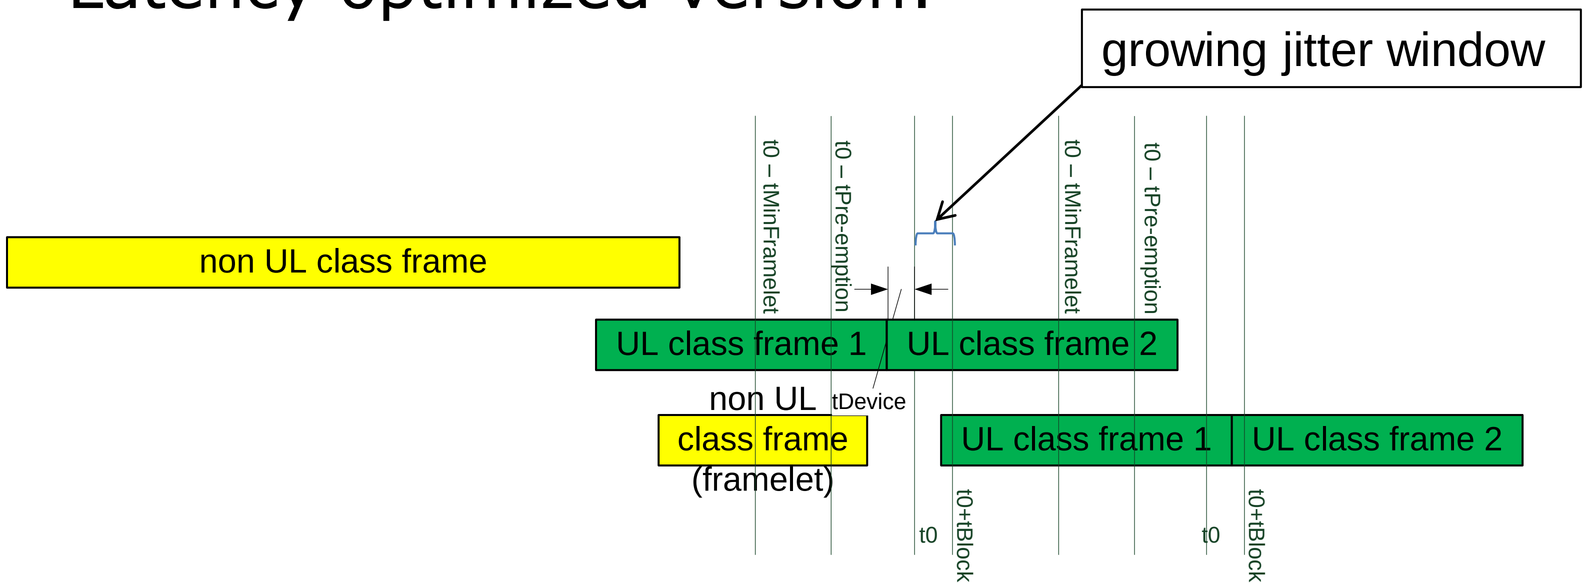

- The time slot is growing at every hop by the MTIE
- Problems: Wasting bandwidth, the topology is limited

Time Aware Shaper – Bandwidth Optimized

- Bandwidth optimized:

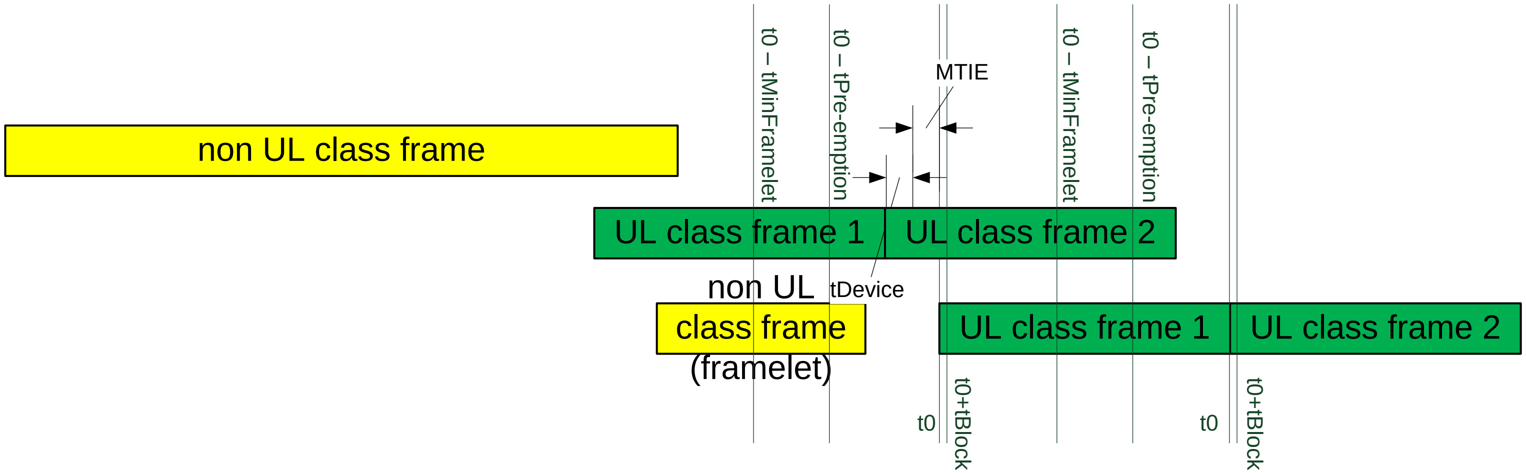

- The TAS de-blocks the queue when the UL class frame is already in the queue
- Higher latency but the time slot has the same size in every bridge

One Additional “Shaper”?

- 125µs CBS (AVB Gen1 Class A) – CBS125
- 250µs CBS (AVB Gen1 Class B) – CBS250
- 125µs CBS w/pre-emption – CBS125P
- 250µs CBS w/pre-emption – CBS250P **?**
- TAS with pre-emption in bridges and end stations – TAS

Not covered by a PAR:

- Time Aware Shaper (TAS) in the end stations with pre-emption in the end stations and bridges – TAS-Ready

Why TAS-Ready?

- Pre-emption improves the CBS latency but does not guarantee Ultra-Low Latency
- TAS requires the configuration of the whole network
- There is no other shaper which could be used in combination with pre-emption
- For small networks with applications which do not need the smallest possible latency it might be useful to have TAS-Ready

This would guarantee a scalable system

How Can a TAS-Ready „Shaper“ Look Like

- Different possibilities
- The simplest version would consist of:
  - Small queue size
  - A mechanism to prevent endless bursts

- Problems:
  - Less deterministic
  - Impact on Class B (e.g. CBS125) can not be predetermined
  - Higher jitter
  - Not suitable for big networks

Possible Extension to the TAS PAR Scope

This standard allows bridges and end stations to provide guarantees for ultra-low latency (i.e. minimum latency and delivery variation), loss-sensitive real-time data transmission. It specifies timing aware queue draining mechanisms **for bridges and end stations and bandwidth aware queue draining mechanisms for bridges** for frames that occur on pre-defined regular intervals. This standard uses the timing derived from 802.1AS to achieve the theoretical minimum latency for engineered network topologies. Virtual Local Area Network (VLAN) tag encoded priority values are allocated, in aggregate, to segregate frames among controlled and noncontrolled queues, allowing simultaneous support of ultra low latency, AV traffic and other bridged traffic over wired Local Area Networks (LANs).

Changes to the last version (Nov. 2011) are red and bold.

Combination Matrix (Class A – Class B)

Class A / Class B	CBS125	CBS250	CBS125P	CBS250P	TAS-Ready	TAS
CBS125	N[1]	N[4]	N[5] ?	N[4]	Y	Y
CBS250	Y	N[1]	Y	N[5] ?	Y	Y
CBS125P	N[3]	N[3]	N[1, 2]	N[2]	N[2]	N[2]
CBS250P	N[3]	N[3]	N[2]	N[1, 2]	N[2]	N[2]
TAS-Ready	N[3]	N[3]	N[2]	N[2]	N[1, 2]	N[2]
TAS	N[3]	N[3]	N[2]	N[2]	N[2]	N[1, 2]

[1] Same shaper
[2] Only one shaper w/pre-emption
[3] Shaper w/pre-emption has to have the highest priority
[4] CBS125 guarantees lower latencies than CBS 250
[5] CBSxxx and CBSxxxP shouldn’t be used as separate classes ???

Compatible Shaper

	CBS125	CBS250	CBS125P	CBS250P	TAS-Ready	TAS
TAS	Y/N	Y/N	Y/N	Y/N	Y/N	Y
TAS-Ready	N	N	N	N	Y	
CBS250P	N	Y	N	Y		
CBS125P	Y	N	Y			
CBS250	N	Y				
CBS125	Y					

Only the CBS w/ and w/o pre-emption with the same measurement interval are fully compatible with each other, i.e. they can form one AVB domain.

In theory a TAS shaped stream can be forwarded in a CBS or TAS-Ready domain if certain parameters are guaranteed (jitter (no burst), latency). But this mapping is perhaps not desirable.

The two different TAS approaches can’t be used within one AVB domain, as TAS-Ready does not guarantee a small enough packet jitter.

Other Possible PAR Additions – Copied from the Assumptions

- Ingress policing/monitoring (Yong)
  - Someone's talking when they shouldn't be
    - Talking without a reservation
      - The stream's DA is not known in the filtering database
      - The frame's PCP is AVB to a unicast
  - Talking too much for the amount reserved
    - Exceeding the reservation
  - Is this perfect policing or best effort?
  - Must it stop a Denial of Service attack?

Thank You