802.1Qav + P802.1Qbv Time-gated Shapers

Norman Finn
Cisco Systems

V01 November 5, 2014

Priority queuing and AVB queuing

- 802.1Q: Priority (including weighted round robin)

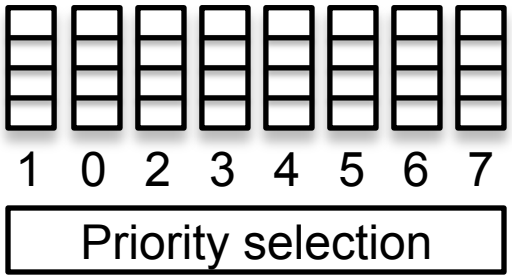

- 802.1Q-2012 (802.1Qat) adds shapers. ∇ Shaped queues have higher priority than unshaped queues, because they can still guarantee bandwidth to the highest unshaped priority.

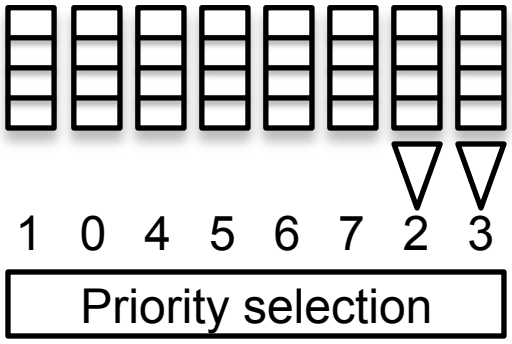

Time-gated queues

- 802.1Qbv: A circular schedule controls a gate between each queue and the priority selection function.

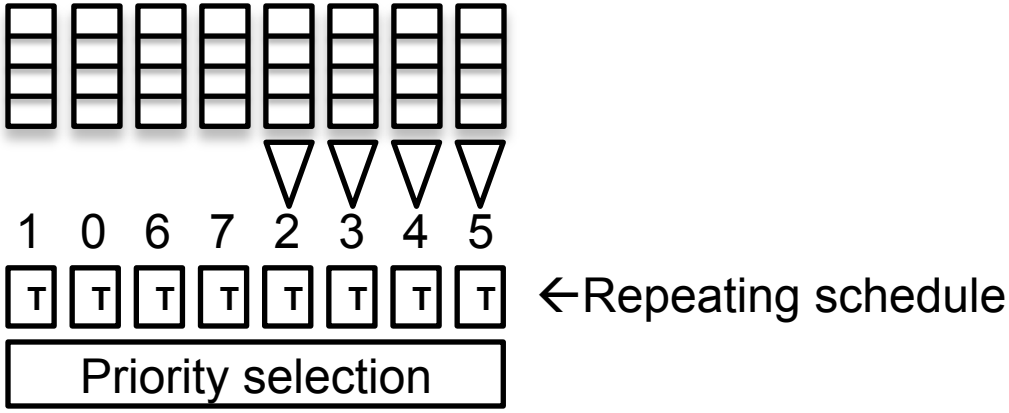

- (Don’t worry about how many queues have shapers. Maybe none. Maybe most. Read on.)

N. Finn's comment on P802.1Qbv D2.1

- My comment on P802.1Qbv D2.1 was related to what happens to the AVB shaper when the time gate is turned on or off. You could:
  1. Continue to build up credit when the gate is off.
  2. Freeze credit when the gate is off, except that negative credit builds up to, but not past, 0.
  3. Freeze credit when the gate is off, period.

Long-cycle CQF: One use case

- I may want CQF not to obtain low-latency, but to obtain zero congestion loss. Finite guaranteed latency is a by-product.
- I may want to have lots of flexibility in the ratio of lowest to highest bandwidth allowed for a reservation.
- I may want to support lots of zero-congestion loss streams through one port.
- Either of these may result in having a large CQF cycle time (and consequently, large queues/buffers).
- By the way – I have no interest in getting rid of AVB queues. I like AVB queues. They have nice properties, even if they don't provide guarantees.

What I want, in effect, for long-cycle CQF

- The two highest-priority queues are the CQF queues. They have to be highest priority, because priority 7 has no bandwidth limit. But, it is essential, since my CQF queues hold a lot of data, that they have an AVB shaper. Otherwise, I will have a horrible latency for priority 7 BPDUs. (For that matter, I may have an AVB-shaped queue either above or below CQF.)

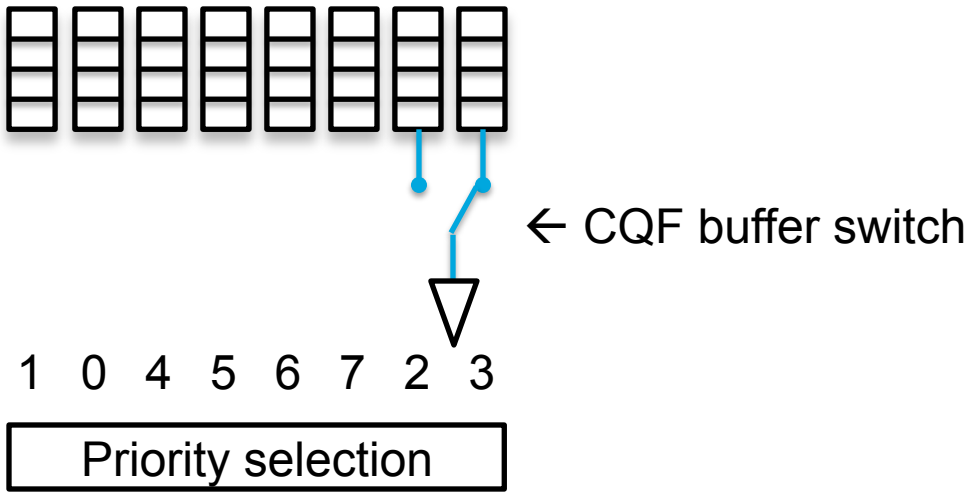

(Thank you, Vahid Tabatabaee!)

What I want, in effect

- Of course, the way I implement the CQF buffer switch is (or can be) the 802.1Qbv time gates.

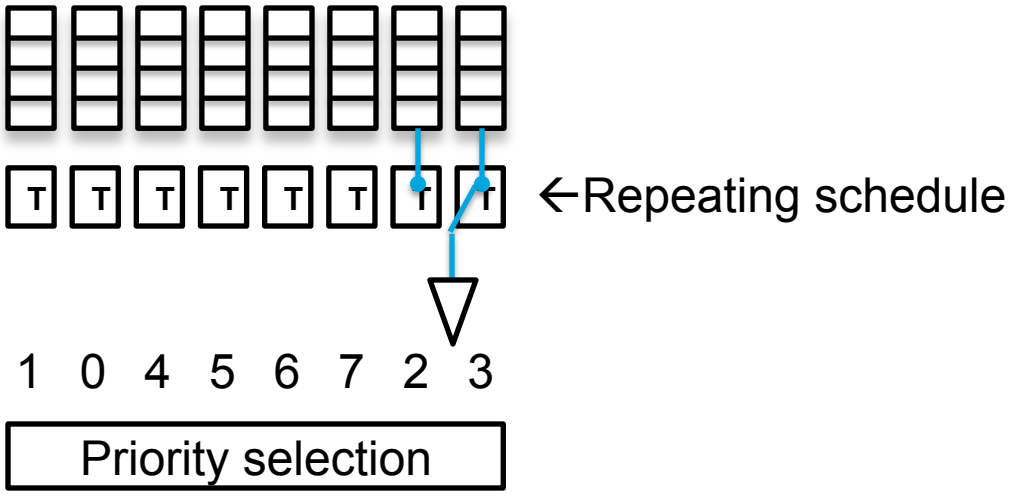

Oops!

- But now, I have:
  1. Inverted the order of the time gates and the shapers.
  2. Introduced the concept of 2 queues feeding the same shaper.
  3. Begged the question, “What happens if both queues feeding one shaper are turned on at the same time? Do you need another level of priority resolution?”

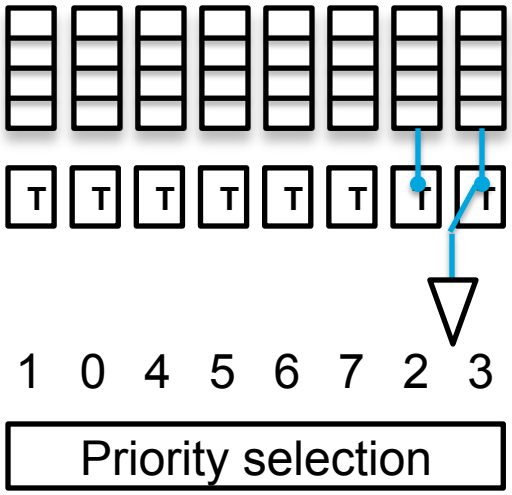

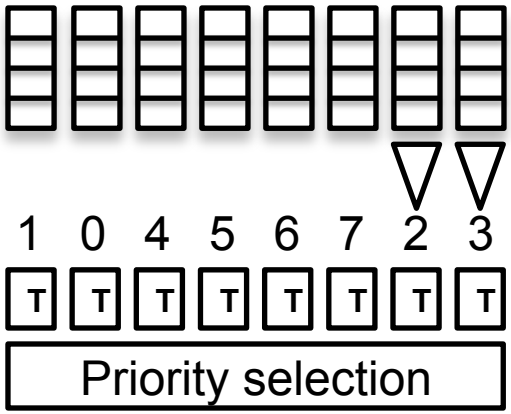

Hence, my ballot comment on Qbv D2.1

- Given the current state of P802.1Qbv, when CQF switches between queues 2 and 3, the shaper always has enough credit to burst the entire contents of the buffer as if there were no shaper.
- If the credit-based shapers are “frozen” while the time gate is closed, then the current architecture gives me what I want for long-cycle CQF.

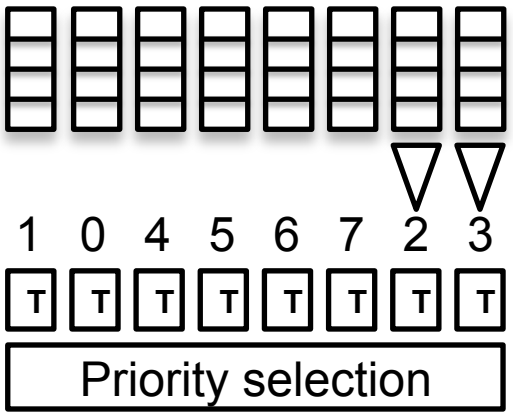

This is the same as the marker model

- Let me also add that, if I choose to implement CQF using a single queue, with a “pause until the next cycle” marker In the queue, and if I then put a shaper on that queue, I will also get exactly what I want, assuming that the “pause” marker looks like an empty queue to the shaper and resets its credit to 0, which seems appropriate.
- I think it’s a good idea that the 2-queue time-gated model for CQF gives the same results as the 1-queue marker model, when we use it with a shaper.

Recommendation

- My comment on P802.1Qbv D2.1 was related to what happens to the AVB shaper when the time gate is turned on or off. You could:
  1. Continue to build up credit when the gate is off.
  2. Freeze credit when the gate is off, except that negative credit builds up to, but not past, 0.
  3. **Freeze credit when the gate is off, period.**
- I think option 3 gives the behavior that I'm asking for.
- **If you have a use case that option 3 would break, please speak up on the next ballot.**

Thank you.