

Multiple Cyclic Queuing and Forwarding

(slides to accompany df-finn-multiple-CQF-0919-v01)

Norman Finn
Huawei Technologies Co. Ltd
nfinn@nfinnconsulting.com
df-finn-multiple-CQF-slides-0919-v01

Why this paper?

Cyclic Queuing and Forwarding is perceived as having major flaws, which prevent its adoption:

- It is difficult to pick a cycle time:
 - A long cycle time produces long end-to-end latency.
 - A short cycle time severely restricts the number of flows handled.
 - No one cycle time serves more than a narrow range of requirements.
- Having 2-buffers per hop precludes long links.


With a little imagination, these issues can be solved without adding anything to IEEE 802.1Q, so the solutions are suitable for P802.1DF.

Step 1: You can run with 3 buffers, not 2


- With 2 buffers, you change input buffers at the same moment that you change output buffers.
- Link delay and forwarding delay have to be subtracted from the useful bandwidth, to prevent filling and emptying a buffer at the same time.
- If you have 3 buffers, then you can change to the next input buffer at a different phase of the cycle from when you change to the next output buffer.
- Each buffer spends part of a cycle idle; one cycle per hop is added to the latency.
- But, the link delay and forwarding delay are no subtracted from the bandwidth..

Two buffers


Two buffers per port. Input and output buffers swap at the same moment, once every cycle, period T_C . Small guard band to allow for link delay and forwarding delay. All bridges are synchronized and swap buffers at the same moment. Cycle time $T_C > \text{transit time} + \text{forwarding time} + \text{clock inaccuracy} + \text{max data transmit time}$.



Two buffers



Two buffers




Three buffers


Input buffer swap is out-of-phase with output buffer swap to allow for arbitrary link delay.

No guard band needed for link/forwarding delay.


Each buffer cycles through four states: filling, full, draining, empty




Three buffers




Three buffers




Three buffers




Three buffers



Three buffers




Three buffers



N buffers

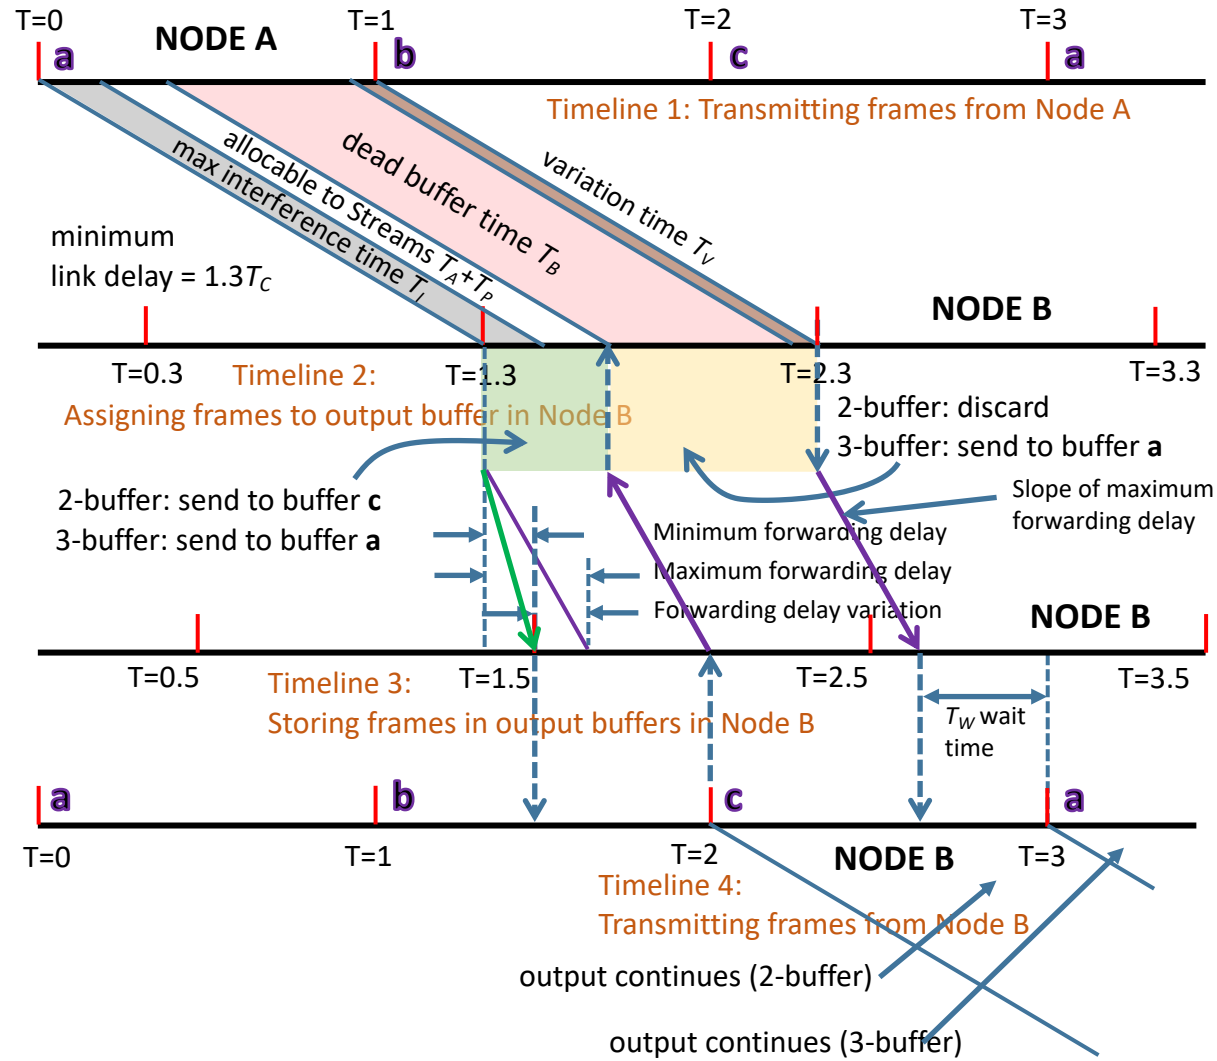
Any number of buffers per port. Useful for delay matching on different paths.



The number of buffers used by a stream


- Different output ports on one system can have different numbers of buffers.
- Different input ports feeding the same output port can each use a different number of the buffers on that output port.
- One Stream can use a different number of buffers at each output port along the path.

df-finn-multiple-CQF-0919-v01 Figure 1 Reference timelines for time-based CQF



Transmission order

- The total bandwidth allocated over all CQF priorities cannot exceed the link bandwidth.
- Draining of CQF buffers are by strict priority, fastest cycles first
- As long as every cycle is an integer multiple of the next-faster cycle, every buffer will empty before the end of its cycle.



More ways to improve the options available to CQF.

- Preemption:
 - Lowers the interference from lower-priority queues, especially best-effort
 - **CQF queues can be preemptable**, because the preemption penalty is bounded.
- Packing varied frame sizes into cycles can require overprovision:
 - “Sausage making,” packing the customers’ packets into fixed-sized encapsulations, requires much less overprovision, and can be done at the edges of a Service Provider network.
- In multiple CQF, deliberate overprovision of a Stream (assigning it to a too-fast cycle) reduces its latency.

The multiple-CQF value proposition

- CQF, as described here, requires network time synchronization.
- CQF requires time-aware state machines on input/output ports.

but,

- Except, perhaps, for the ingress bridge, CQF requires no per-Stream configuration, no per-Stream state machines, no per-Stream knowledge at any node in the network.
- CQF allows a network controller to perform admission control with no per-flow conversations with any (except perhaps the ingress) nodes in the network.

To come ...

There are many more improvements that can be made to TSN, many of which require no new standards. I hope to supply more papers and presentations in this vein.

Thank you