b - plus

automotive | mobile automation | embedded systems

Credit Based Shaper with Preemption Support

Christian Boiger
christian.boiger@b-plus.com

IEEE 802.1 Interim
September 2014
Ottawa, Canada

Interspersing Express Traffic

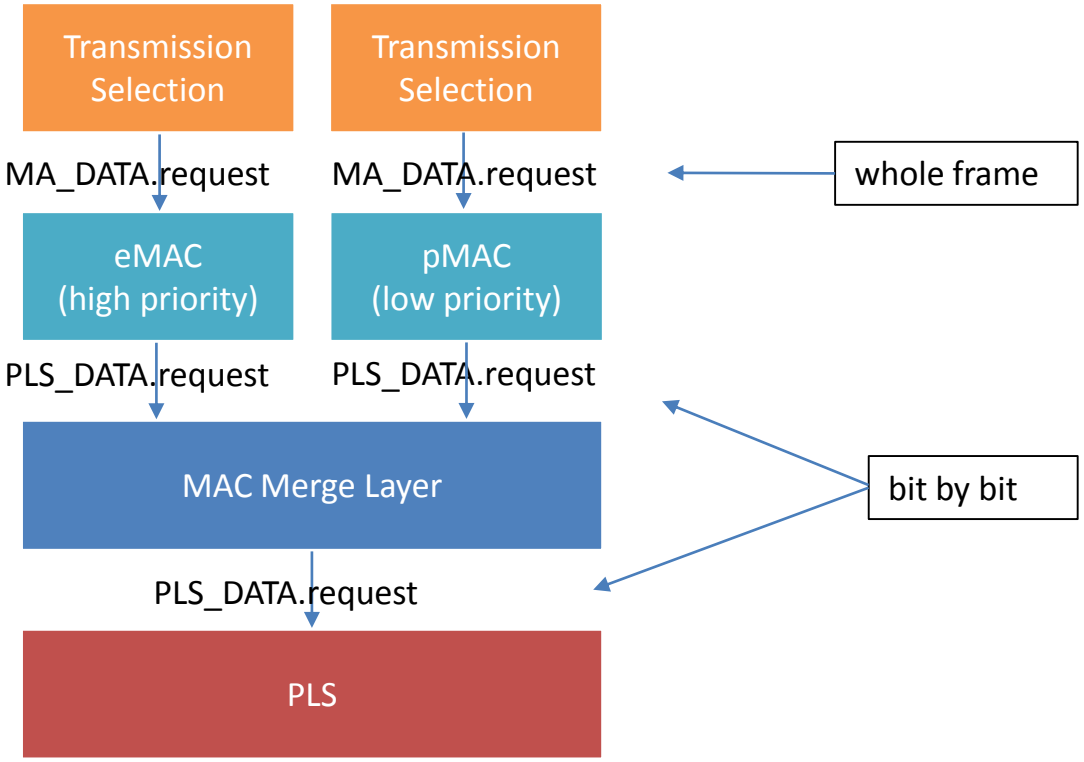

CBS Credit Calculation

- **“idleSlope. The rate of change of credit, in bits per second, when the value of credit is increasing (i.e., while transmit is FALSE).”**
- **“transmit. Takes the value TRUE for the duration of a frame transmission from the queue; FALSE when any frame transmission from the queue has completed.”**
- **“credit. The transmission credit, in bits, that is currently available to the queue. If, at any time, there are no frames in the queue, and the transmit parameter is FALSE, and credit is positive, then credit is set to zero.”**
- **“sendSlope. The rate of change of credit, in bits per second, when the value of credit is decreasing (i.e., while transmit is TRUE).”**

Reserved Traffic Queue is Connected with eMAC

- AVB stream frames can preempt non AVB stream frames
- Use Cases:
  - Improvement of AVB latency
  - Transmission of “too big” non AVB stream frames

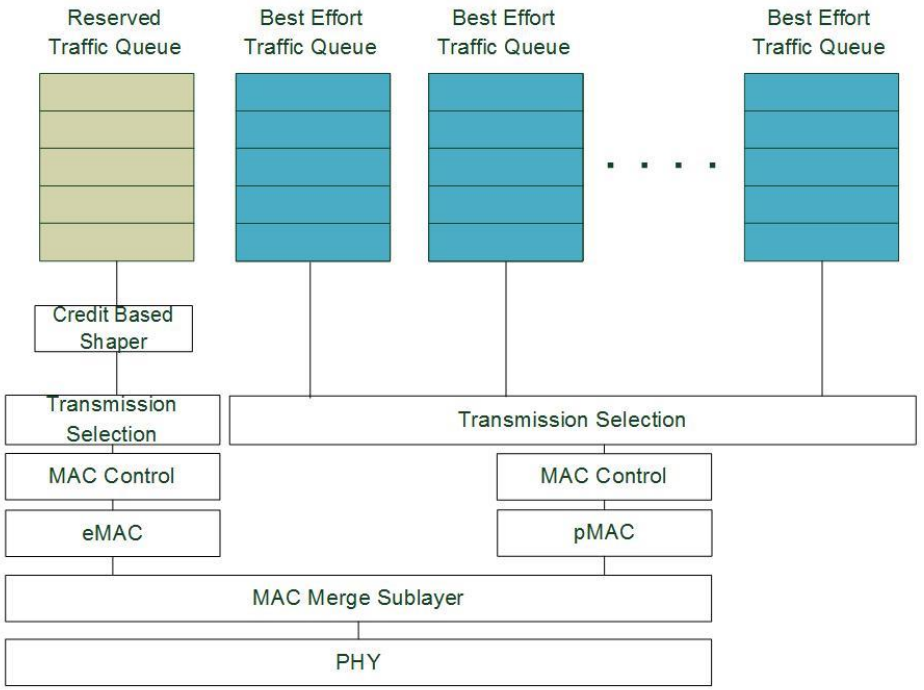

Reserved Traffic Queue is Connected with eMAC

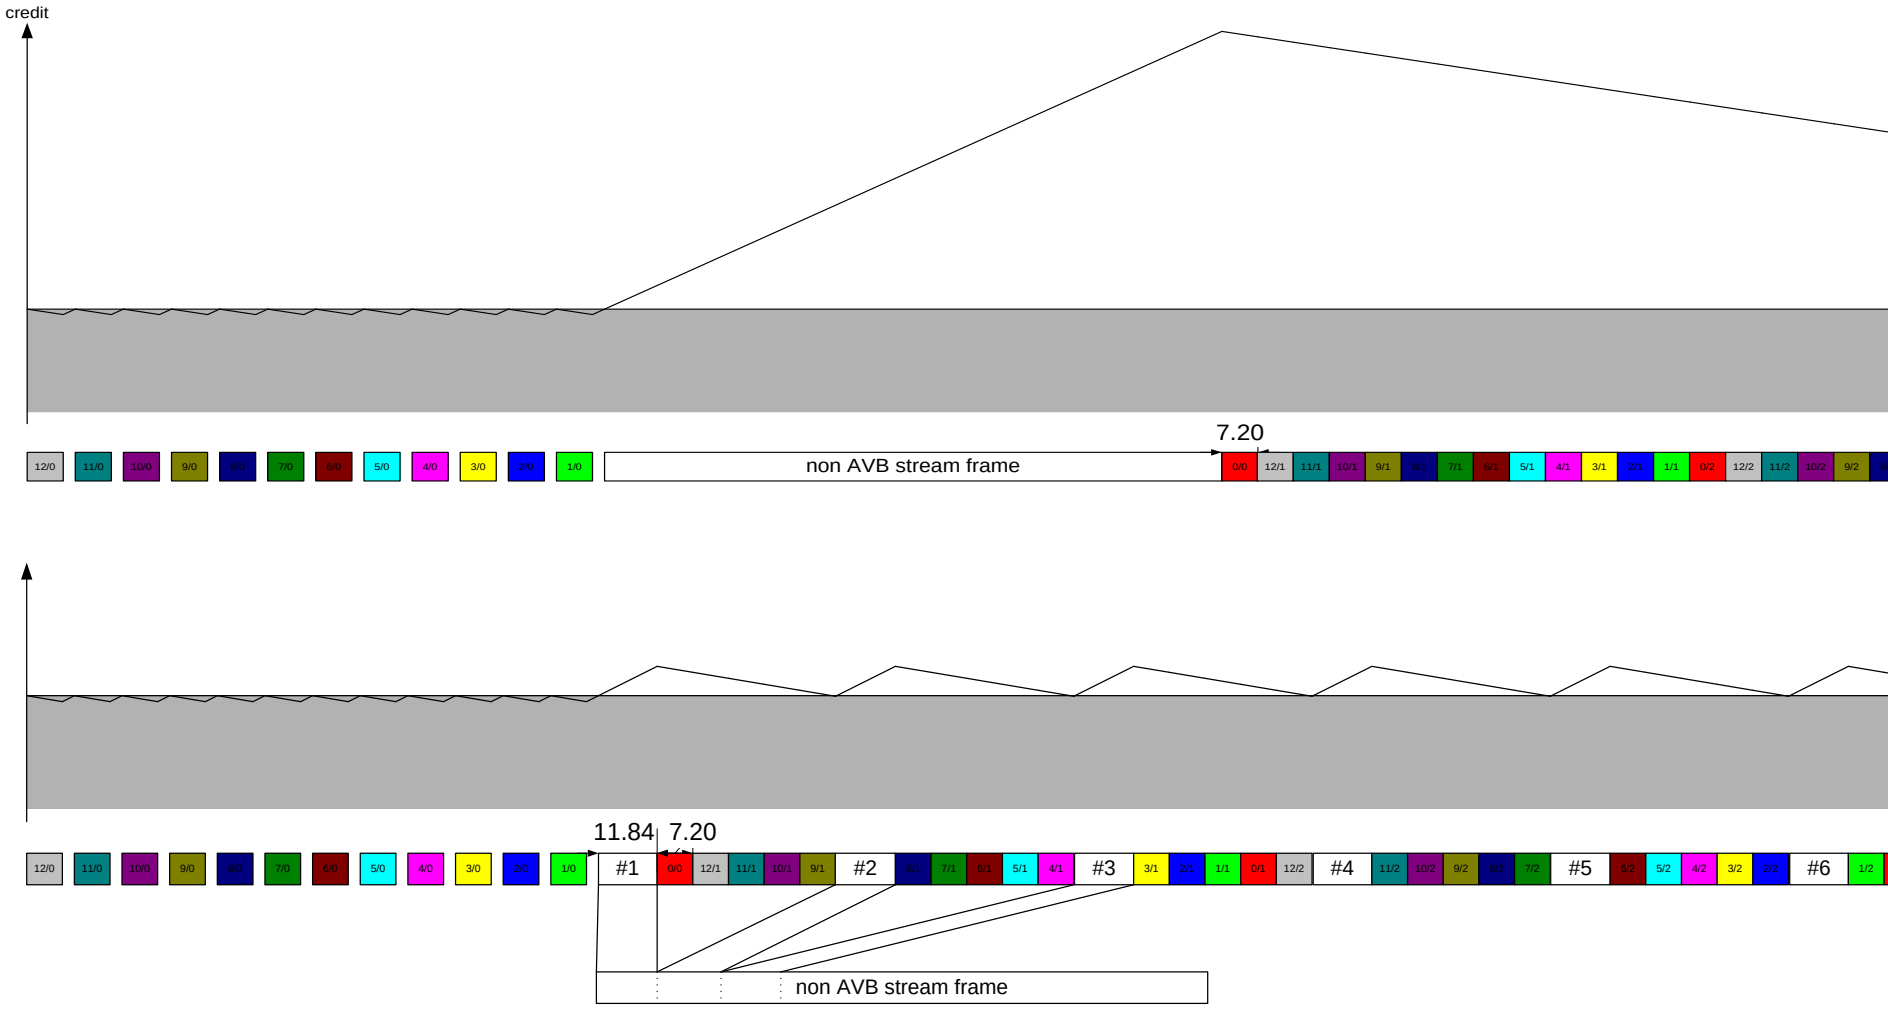

When is the Credit Increasing?

- **“idleSlope. The rate of change of credit, in bits per second, when the value of credit is increasing (i.e., while transmit is FALSE).”**
- **“transmit. Takes the value TRUE for the duration of a frame transmission from the queue; FALSE when any frame transmission from the queue has completed.”**

- The credit calculation depends on the transmit parameter.
- It is not stated in 802.1Q, how the duration of the frame transmission is exactly determined (or I haven't found it).
- As 802.1Q magically knows when the next MA_DATA.request can be issued, it somehow seems to be possible to know the duration of a frame transmission in 802.1Q (in the absence of preemption).

Impact of Preemption

- Without Preemption 802.3 did not have a queue
- Preemption adds a queue (one frame) to 802.3. It is possible that at the same point in time two frames are in the MAC Layer (one in the eMAC and one in the pMAC)
- This queue reduces the knowledge about the current transmission status
- After an eMAC MA_DATA.request a frame is not immediately transmitted if another frame from the pMAC is currently in transmission.
    - The time interval between two MA_DATA.requests on the eMAC can not be used to measure the “duration of the frame transmission”.

When is “transmit” true?

- **“transmit. Takes the value TRUE for the duration of a frame transmission from the queue; FALSE when any frame transmission from the queue has completed.”**

- Transmit is true from the beginning of a MA_DATA.request until the next MA_DATA.request could be issued.

  → This would lead to a wrong credit calculation as it would include the time that is needed to preempt a pMAC frame

When is “transmit” true?

Another way to look at it:

- The CBS magically knows the exact time when the frame is transmitted by the eMAC (might be a little bit too much magic) and the transmit parameter is set accordingly.

- **“If, at any time, there are no frames in the queue, and the transmit parameter is FALSE, and credit is positive, then credit is set to zero.”**

  → If the frame is in the eMAC and waits for transmission and no other frame is in the AVB stream queue, the credit is set to zero, as transmit would be FALSE

  → This would result in a wrong credit calculation, as the credit would be set to zero during the time an AVB frame waits for its transmission.

Reserved Traffic Queue is Connected with pMAC

- AVB stream frames can be preempted by eMAC frames
- Use Cases:
  - Converged network with low latency control streams (e.g. Scheduled Traffic)

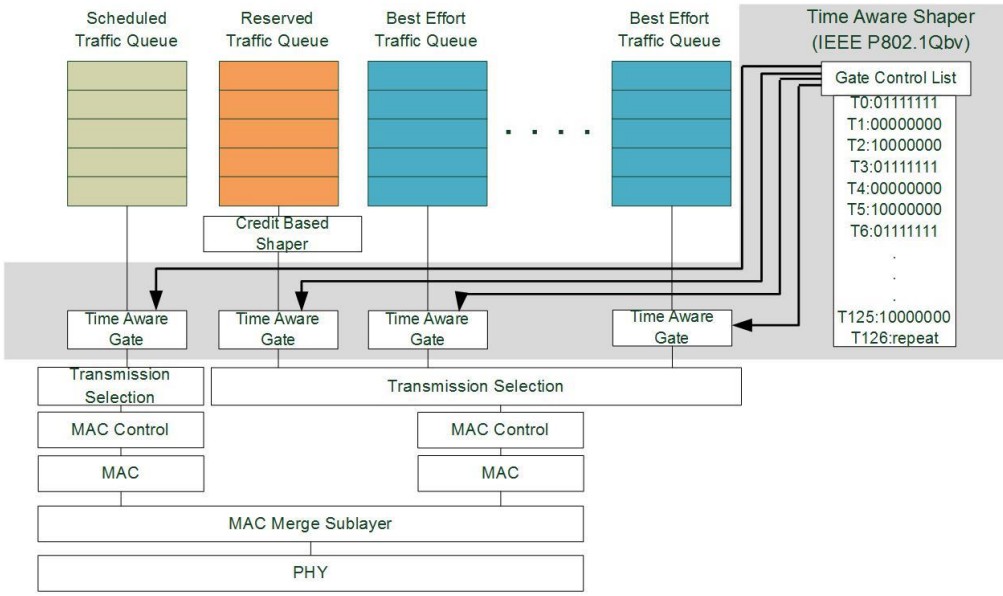

Additional Issues

- The transmission of pMAC frames can be interrupted by eMAC frames and the start of transmission can be delayed due to an eMAC frame in transmission.
  - → The time interval between two MA_DATA.requests on the pMAC can not be used to measure the “duration of the frame transmission”.

- The length of pMAC frames on the wire is unknown to the higher layer as it is the sum of the length of all fragments.

Possible Solutions

- The Credit Based Shaper magically knows everything (the duration of the frame transmission of each frame as well as the exact point in time when a frame is transmitted)
  - → This would need only some changes to the current definition -> simplest solution; however, this seems to be a little bit too much magic.
- The Credit Based Shaper uses the available information like frame length, time intervals between MA_DATA.requests, etc. to “simulate” the credit
  - → Might be possible but complicated (e.g. in order to determine if a “set credit to zero” event happened or not; or if a frame got preempted or not)
- 802.3br could provide information about the transmission status (e.g. the transmit parameter, or number of bits of current frame transmitted for each MAC)
  - → Would require to pass additional parameters to 802.1Q through the MAC Merge Service Interface
- Other possibilities?

Additional MAC Merge Service Interface Primitives

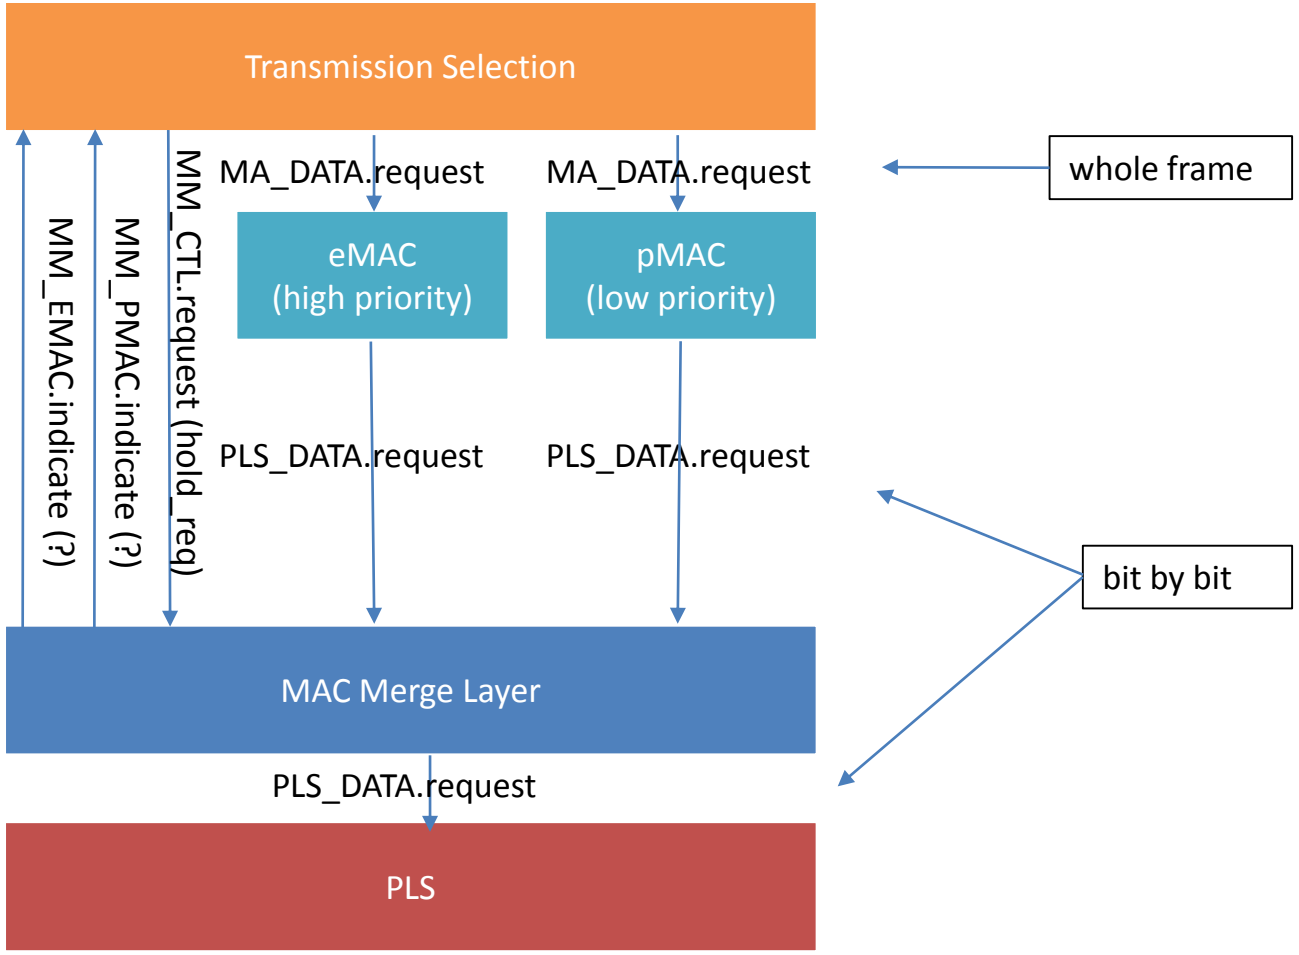

Thank You