802.3br IET Proposal
Architecture and Encapsulation

2014-01-15

Pat Thaler

Contents

- MAC Merge Architecture
- Mframe format

MAC Merge layer encapsulation goals

- Preserve frame integrity
  - No increase in undetected errors
- Indicate which MAC receive frame belongs to
- Minimize impact on throughput
- Transparent to existing non-deprecated PHYs above 10 Mb/s

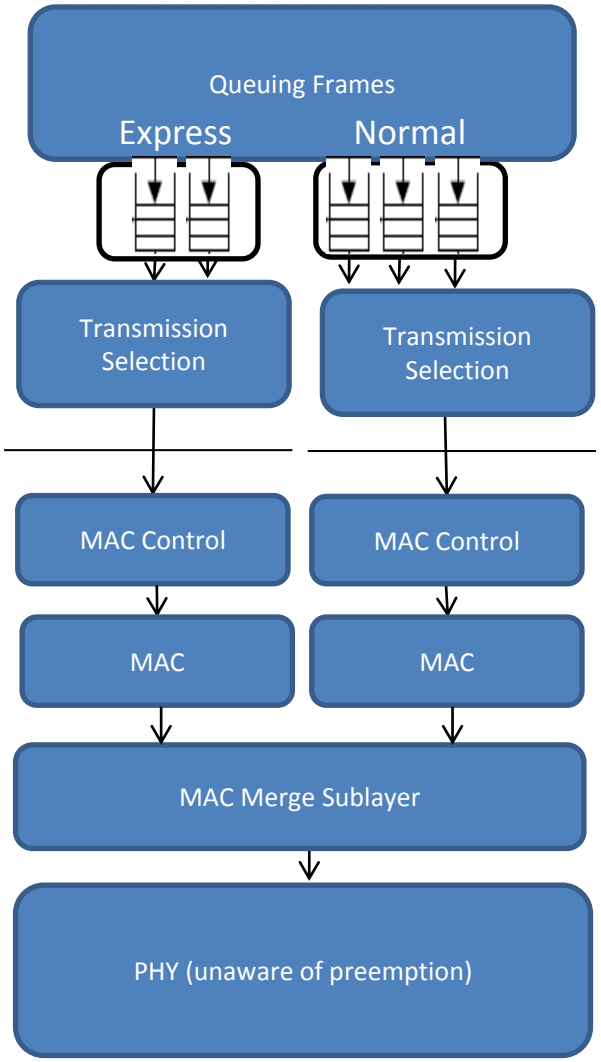

Terminology

- Express frame – frames with the lowest latency
- Normal frame – frames that are not Express frames.
- Mframe -- A transmitted unit from MAC Merge that includes both whole frames and fragments of premptable frames – Mframe
    - Which is stands for MAC Merge frame – a unit that looks like a frame at the PHY layer but may contain a whole frame or a fragment of a MAC layer premptable frame. I'm leaning toward this one.

Hold Primitive need

A hold primitive allows:

- A MAC Client that has a schedule for Express traffic to preempt Normal frames before the scheduled Express traffic arrives When scheduled frame arrives, it can be transmitted immediately
- EPON MAC Control to preempt a frame near the end of the Gate and send a Report before the end of the Gate.

MM.Request from MAC Client

- MM.request: Primitive to carry the indication to Hold (i.e. preempt any frame in transmission and prevent the start of new frames) or release the normal transmission path

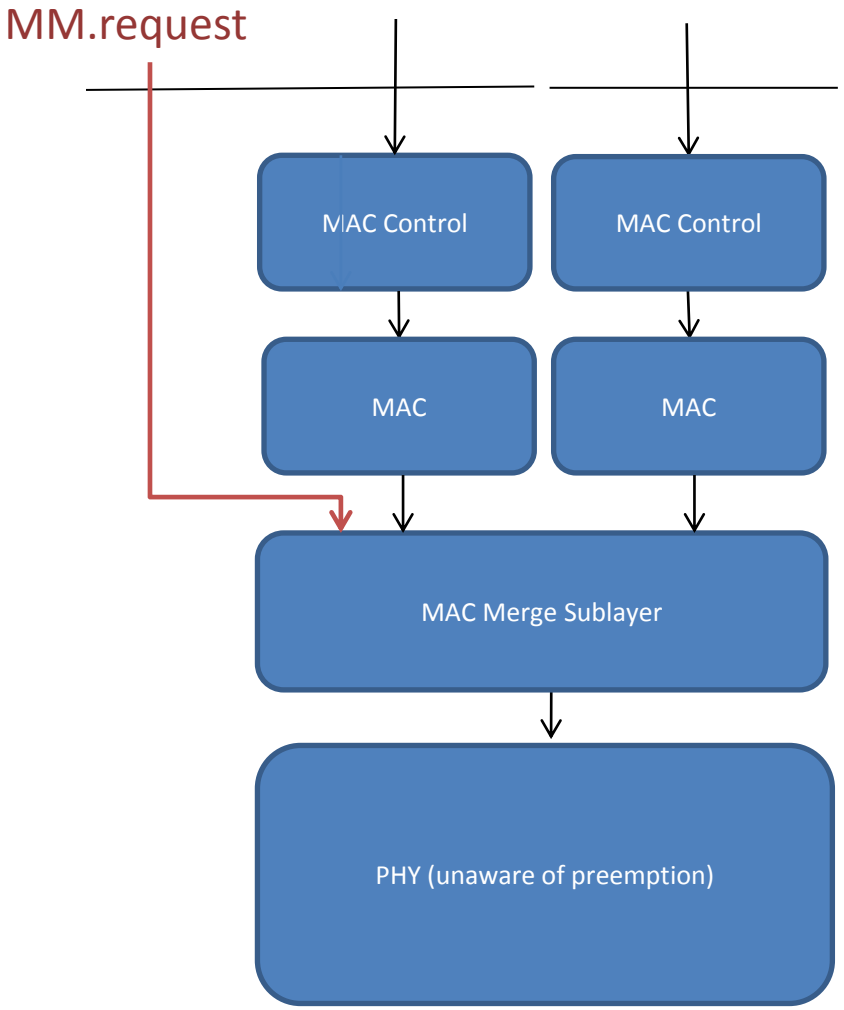

MM.request(hold_req)

- hold_req parameter takes one of two values:
  - hold – asserts hold variable in MAC Merge sublayer
  - release – clears hold value in MAC Merge sublayer
- MAC Merge preempts whenever hold = TRUE or DMLT MAC PLS_DATA.request has a bit to transmit.

Mframe format

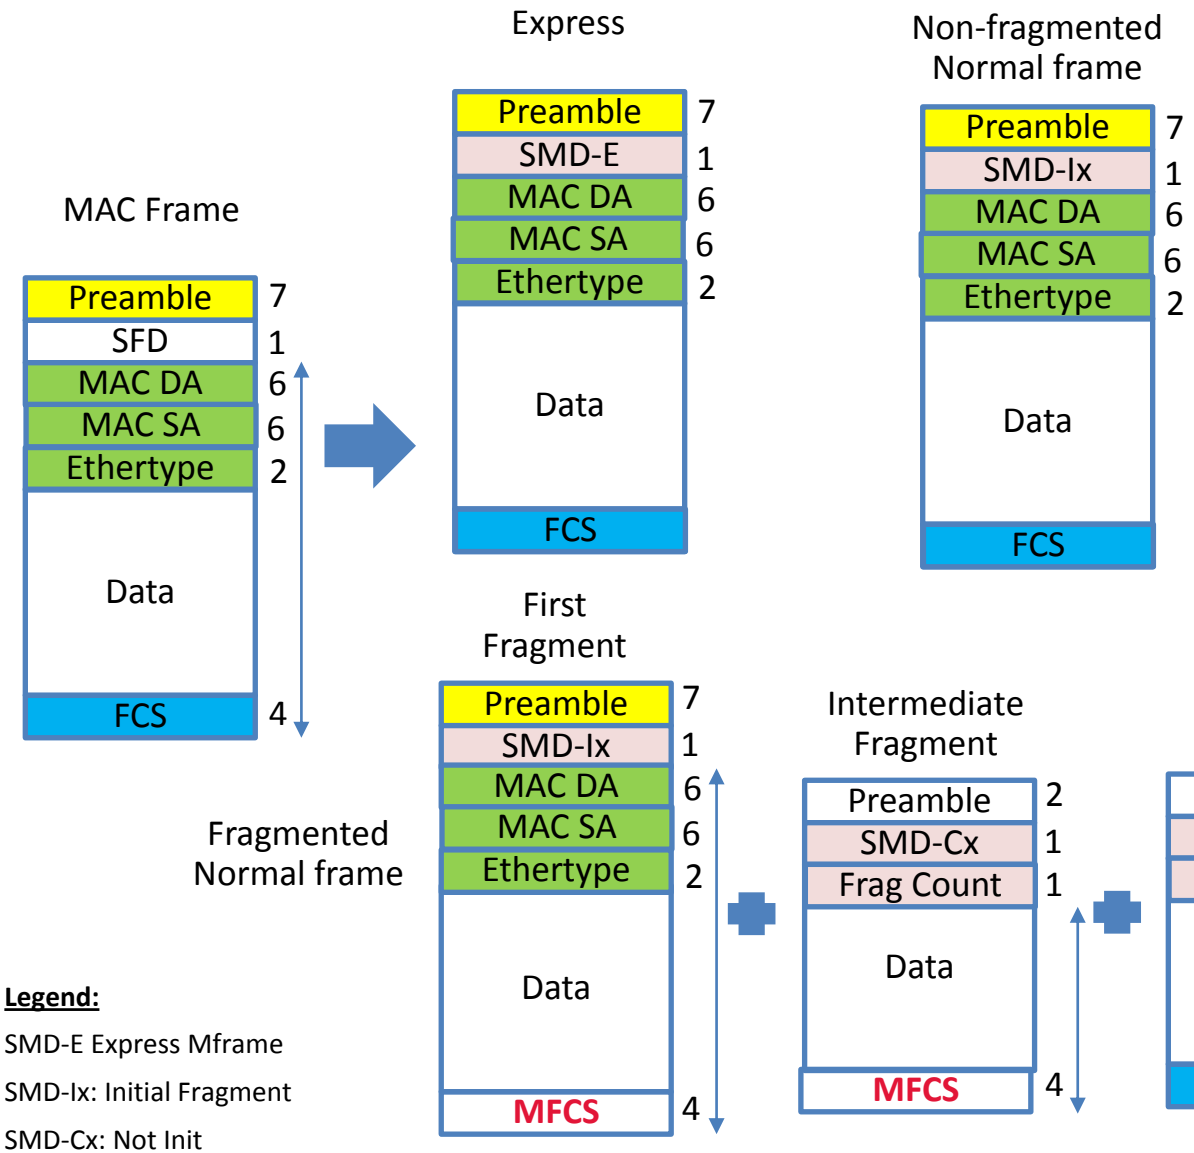

MFCS is the CRC of a non-final fragment.

Value is the same as the FCS of the frame bytes transmitted XOR FF00

Alternatively, could use a one-byte fragmentation trailer on each non Express Mframe

Legend:

SMD-E Express Mframe

SMD-Ix: Initial Fragment

SMD-Cx: Not Init (Continue) Fragment

Mframe format elements

- Preamble
  - Minimize use of preamble bytes for non-initial fragments
  - Provide at least 2 bytes of preamble for fragments
    - The 1000BASE-X PHY can drop up to 2 bytes of preamble and insert the SFD over another byte. In practice, implementations only drop 1 and many delay the start of preamble rather than drop and only overwrite 1 byte (similar to 10GBASE-X alignment)
    - Open issue: use full or shortened preamble for 10 Mbps PHYs?
- Identify Mframe as Express or start of a Normal frame or a later fragment of a Normal frame
- Protection for reassembly errors when an Mframe is lost
  - Frame number – circular count from 0 to 3
  - Fragment number – circular count from 0 to 3
- Identify last Mframe of a Frame
  - Mark end of Normal Mframe

Fragment size constraints

- Preempted fragment size will be no smaller than 64 bytes
  - Therefore a packet less than 127 bytes will not be preempted.
- To simplify implementation, non-final fragments will have an alignment constraint
  - That is, their payload will be a multiple of a set number of octets
  - 4-octet alignment or 8-octet alignment?

Mframe start

- For start of non-initial fragments
  - Insert 2 bytes of preamble followed by
  - SMD byte (Start Mframe Delimiter)
  - Frag byte (Fragment count)
- Normal Frame start and Express frame
  - Replace SFD with SMD
- SMD values have Hamming distance 4 from each other
- Frag values have Hamming distance 4 from each other

SMD and Count byte encodings

Mframe type	Frame #	SMD
SMD-E	NA	0x33
SMD-Ix Premptable frame start	0	0x66
	1	0xCC
	2	0xFF
	3	0xAA
SMD-Cx Non-initial fragment	0	0xE1
	1	0xD2
	2	0x1E
	3	0x2D

Frag Count	Frag
0	0x66
1	0xCC
2	0xFF
3	0xAA

Frame and Mframe CRC

- Frame CRC is generated by the MAC over the same frame bits as always. It is not altered by the MAC Merge sublayer

- MAC Merge layer calculates the Mframe CRC which it adds to any non-final fragment Mframes. Mframe CRC on a fragment is calculated over the bytes of the frame transmitted up to the end of the current Mframe. It is the same as the MAC CRC calculation XOR'ed with FF00

Fun with CRCs

- IEEE 802.3 inverts the calculated CRC and appends the result to the frame.
- Transmitting MAC Merge sublayer could invert the second two bytes of the intermediate CRC result (the CRC computed over the bytes of the MAC frame that have been transmitted so far) at the end of non-final fragments.
- Receiving MAC Merge sublayer runs a CRC calculation as the frame is received. When an Mframe ends, it compares the calculated value with the
  - If the difference between that and the last 4 bytes of the Mframe is 0xFFFF, it's the end of a MAC frame
  - If the difference between that and the last 4 bytes of the Mframe is 0x00FF, it's a non-final fragment.

Thank You.

Questions or Comments?