

IEEE 802.3 SPEP2P SG: Half-Duplex and Preemption

Peter Jones - Cisco

Background - CFI

Expanding Ethernet in OT

Familiar Topologies

- *Network topology driven by use case*
- *Point to point, multidrop, trunk & spur*
- *Reach, cable type*

Power Delivery

- *Power small device (e.g., sensor, field switch)*
- *Power control system for larger device (e.g., HVAC air handling unit)*

TSN - synchronization

- *Precision Time Protocol (e.g., IEEE 1588 default, 802.1AS, IEC 62439-3, SMPTE 2059)*
- *Tracking events*
- *Coordinating actions*

TSN – Latency

- **Frame Preemption (IET)**
- *Credit Based Shaper*
- **Scheduled Traffic**

TSN – Reliability

- *Frame Replication and Elimination*
- *Path Control and Reservation*
- *Per-Stream Filtering and Policing*

TSN - Resource Management

- *Stream Reservation Protocol*
- *Link-local Registration Protocol*
- *LLDPv2 for MultiframeData Units*
- *Multicast and Local Address Assignment*

Background - CFI

So, What's the Problem?

- Clause 99.1 in IEEE Std 802.3-2018: “specifies an optional MAC Merge sublayer for use with a pair of full-duplex MACs and a single PHY operating at 100 Mb/s or higher on a point-to-point link”
 - This makes perfect sense in that many 10 Mb/s PHY do not support the PCS and thus will not recognize the the SMD which is the Start of Mpacket Delimiter
 - However the newer 10 Mb/s PHY technologies (T1L and T1S) do support the PCS and will work with the MAC Merge sublayer
- Other TSN features (scheduled traffic, FRER, ATS, etc.) are already compatible with these PHY technologies.

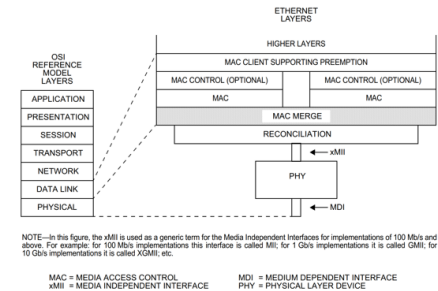


Figure 99-1—Relationship of MAC Merge sublayer to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model


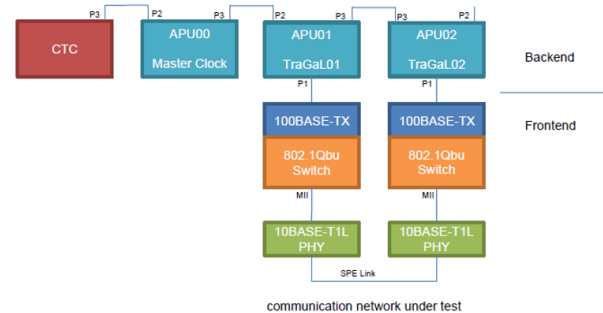


Figure 99-4—mPacket format

Background - CFI

What is left out of TSN for 10BASE-T1L?

- Addition of 10BASE-T1L to MAC MERGE clause should be easy and straightforward
- Should be a simple project
- Study group should first look and make sure nothing else is missing



10BASE-T1L MACMERGE demonstration
Source: Martin Ostertag (private communication)

Background – Preemption and 10BASE-T1L

Content

- Investigation of available PHY hardware (finished)
 - Shows that existing 10BASE-T1L PHYs accept Start mPacket Delimiters (SMDs) according to IEEE 802.3-2018 clause 99.3. Frames with different SMDs are accepted for transmit (TX) and receive (RX). A mixed vendor setup was used to perform and document tests with different combination of SMDs.
- Feasibility Demonstration (reduced latency for high priority traffic - ongoing)
 - The positive impact of implementing MAC Merge on a 10BASE-T1L link on high priority traffic latencies is demonstrated in a mixed-vendor setup with a fictive application.

What about 10BASE-T1S Point to Point?

Preemption and Half-Duplex

- Preemption reduces the delay to transmit a high priority (express) packet when a station is currently transmitting a lower priority (preemptable) packet
- On a full-duplex link, this reduces the worst-case insertion delay for a single express packet from MTU size to 64 bytes
- On a half-duplex link, this reduction applies only when the local station (not the peer) is transmitting

10BASE-T1S – Half-Duplex and Full-Duplex

IEEE Std 802.3cg-2019

IEEE Standard for Ethernet—Amendment 5: Physical Layer Specifications and Management Parameters for 10 Mb/s Operation and Associated Power Delivery over a Single Balanced Pair of Conductors

147. Physical Coding Sublayer (PCS), Physical Medium Attachment (PMA) sublayer and baseband medium, type 10BASE-T1S

147.1 Overview

This clause defines the type 10BASE-T1S Physical Coding Sublayer (PCS) and type 10BASE-T1S Physical Medium Attachment (PMA) sublayer. Together, the PCS and PMA sublayers comprise a 10BASE-T1S Physical Layer (PHY). Provided in this clause are full functional and electrical specifications for the type 10BASE-T1S PCS, PMA, and MDI.

The 10BASE-T1S PHY is specified to be capable of operating at 10 Mb/s in several modes. All 10BASE-T1S PHYs can operate as a half-duplex PHY with a single link partner over a point-to-point link segment defined in 147.7, and, additionally, there are two mutually exclusive optional operating modes: a full-duplex point-to-point mode over the link segment, defined in 147.7, and a half-duplex shared-medium mode, referred to as multidrop mode, capable of operating with multiple stations connected to a mixing segment, defined in 147.8. The medium supporting the operation of the 10BASE-T1S PHY is defined in terms of performance requirements between the attachment points (Medium Dependent Interface (MDI)), allowing implementers to specify their own media to operate the 10BASE-T1S PHY as long as the normative requirements included in this clause are met.

Preemption Details

802.1 Qbu

- Does not include any assumptions about speed or-duplex

802.3 Clause 4

- 4.3.3 Services required from the Physical Layer – “The collisionDetect signal is generated only during transmission and is never true at any other time; in particular, it cannot be used during frame reception to detect collisions between overlapping transmissions from two or more other stations.”

802.3 Clause 22

- 22.2.1.4 Mapping of PLS_SIGNAL.indication - SIGNAL_STATUS assumes the value SIGNAL_ERROR when the MII signal COL is asserted, and assumes the value NO_SIGNAL_ERROR when COL is deasserted

Preemption Details (cont.)

802.3 Clause 99. MAC Merge sublayer

▪ Today

- “This clause specifies an optional MAC Merge sublayer for use with a pair of full-duplex MACs and a single PHY operating at 100 Mb/s or higher on a point-to-point link.”
- 99.4.5 Receive processing – “PLS_SIGNAL.indication is never produced by Receive processing since it does not apply to full-duplex PHYs.”
- 99.4.6 Express filter – “PLS_SIGNAL.indication is never produced by Express filter since it does not apply to full-duplex PHYs.”

▪ Thoughts

- As per “Figure 99–3—MAC Merge sublayer Functional Block Diagram”, RS:PLS_SIGNAL.indication is sent to both “Receive Processing” and “Express Filter”.
- If these functions just pass this signal to both the eMAC and the pMAC, then these MACs can deal with the collision impact. Need to check what happens when a collision is signaled to a MAC that’s not currently transmitting.
- A similar approach is taken with PLS_CARRIER.indication - “If a PLS_CARRIER.indication is received from the PLS, PLS_CARRIER.indications with the same CARRIER_STATUS shall be sent to the pMAC and the eMAC.”

MAC Merge Block Diagram

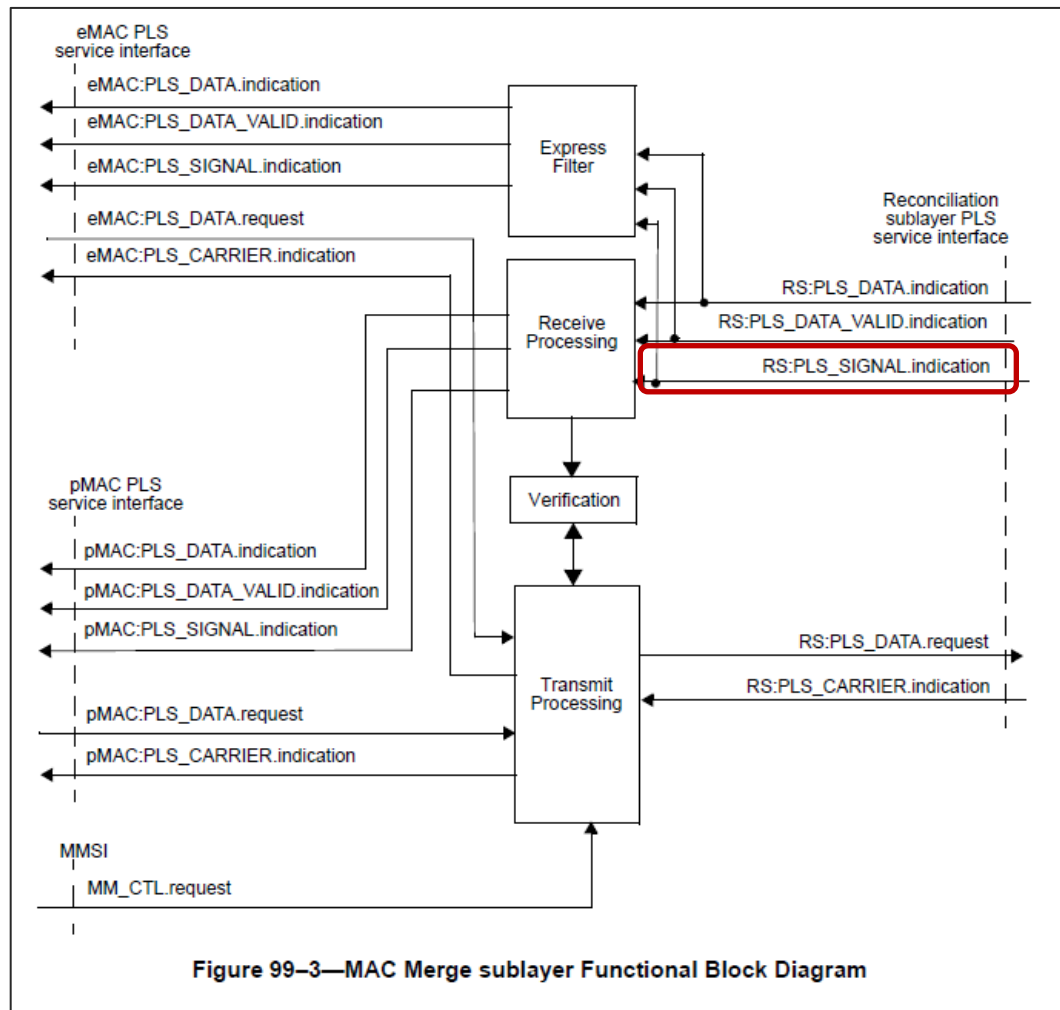


Figure 99-3—MAC Merge sublayer Functional Block Diagram

Preemption and Half-duplex

Based on my reading, it seems that if MAC Merge passes PLS_SIGNAL.indication to eMAC and pMAC (like it does with PLS_CARRIER.indication), preemption and half-duplex should work correctly. The Clause 4 MAC spec only pays attention to collisionDetect in WatchForCollision() and it's guarded by transmitSucceeding

In half duplex mode, TransmitLinkMgmt monitors the medium for contention by repeatedly calling WatchForCollision, once frame transmission has been initiated:

```
procedure WatchForCollision;  
begin  
  if transmitSucceeding and collisionDetect then  
    begin  
      if currentTransmitBit > (slotTime – headerSize) then lateCollisionError := true;  
      newCollision := true;  
      transmitSucceeding := false;  
      if burstMode then  
        begin  
          bursting := false;  
          if not burstStart then  
            lateCollisionError := true {Every collision is late, unless it hits the first frame in a burst}  
          end  
        end  
      end  
    end;  
end; {WatchForCollision}
```

Proposed Text Changes

99.1 Introduction

This clause specifies an optional MAC Merge sublayer for use ~~with a pair of full-duplex MACs and a single PHY operating at 100 Mb/s or higher~~ on a point-to-point link. The two MACs are:

- a preemptable MAC (pMAC), which carries the preemptable traffic, and
- an express MAC (eMAC), which carries the express traffic.

99.4.5 Receive processing

~~PLS_SIGNAL.indication is never produced by Receive processing since it does not apply to full-duplex PHYs."~~

If a PLS_SIGNAL.indication is received from the PLS, PLS_SIGNAL.indication with the same SIGNAL_STATUS shall be sent to the pMAC.

99.4.6 Express filter

~~PLS_SIGNAL.indication is never produced by Express filter since it does not apply to full-duplex PHYs.~~

If a PLS_SIGNAL.indication is received from the PLS, PLS_SIGNAL.indication with the same SIGNAL_STATUS shall be sent to the eMAC.

Conclusions

- 10BASE-T1S point to point is an important member of the 10BASE-T1S family
- Half-duplex is mandatory for 10BASE-T1S point to point
- TSN already supports half-duplex media (e.g., 802.11)
- We should support preemption on 10BASE-T1S half-duplex in this project
- Add preemption with half-duplex to objectives

Enable use of Clause 146 and Clause 147 point-to-point full-duplex *and half-duplex* 10 Mb/s PHYs with Clause 99 MACMERGE

Consensus

WE BUILD IT.

Connect with us on:



Facebook: <https://www.facebook.com/ieeesa>



Twitter: @ieeesa



LinkedIn: <http://www.linkedin.com/groups/IEEESA-Official-IEEE-Standards-Association-1791118>



IEEE-SA Standards Insight blog: <http://standardsinsight.com>



YouTube: IEEE-SA Channel

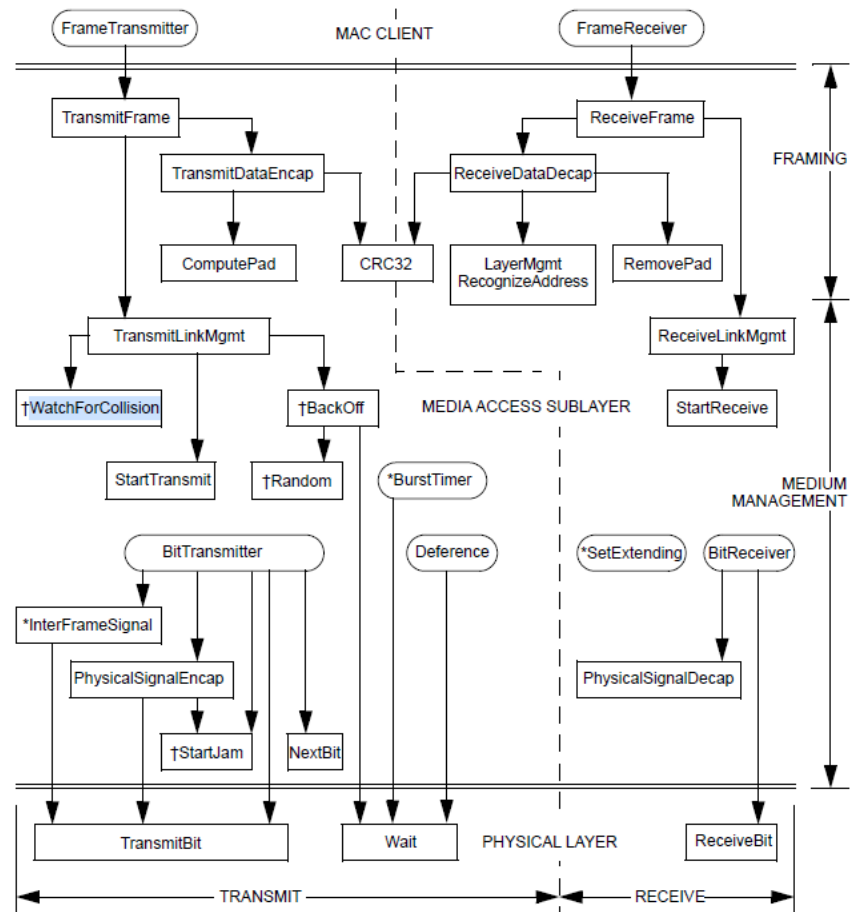
IEEE

standards.ieee.org

Phone: +1 732 981 0060 Fax: +1 732 562 1571

© IEEE

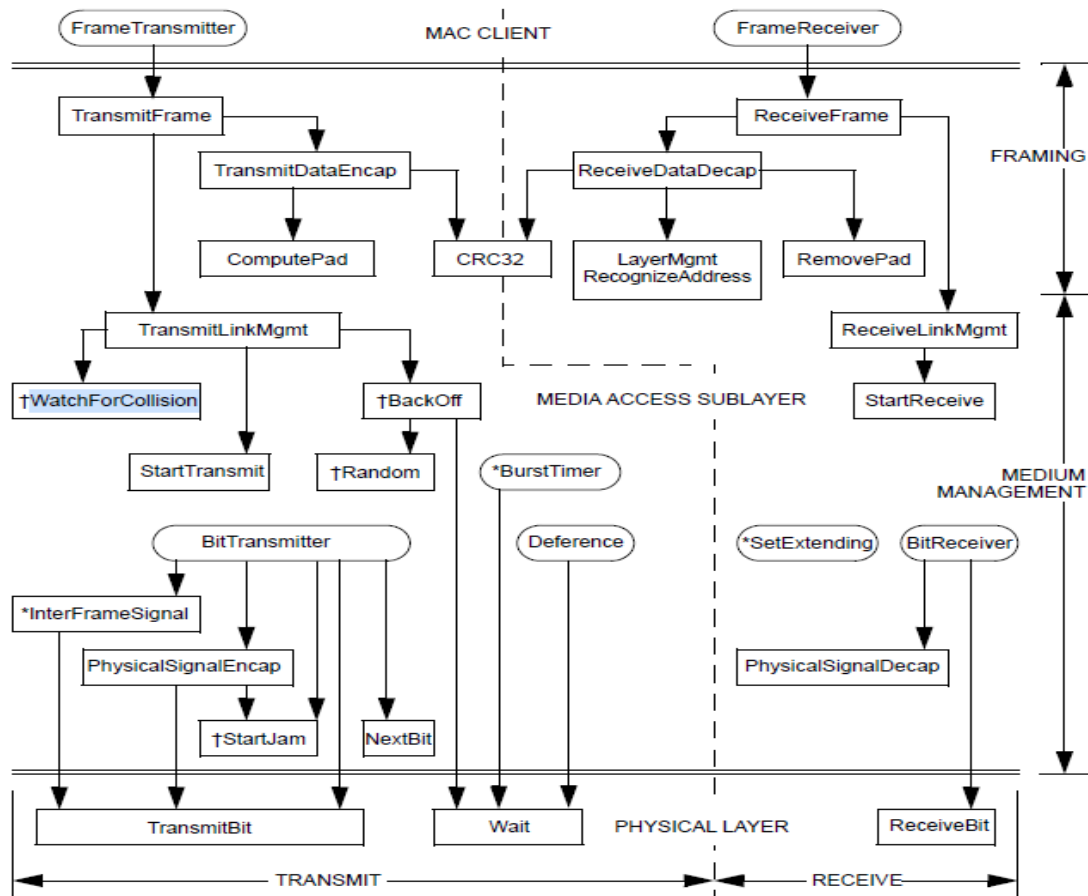
Backup



† Not applicable to full duplex operation.

* Applicable only to half duplex operation at 1000 Mb/s.

Figure 4-1—Relationship among CSMA/CD procedures



† Not applicable to full duplex operation.

* Applicable only to half duplex operation at 1000 Mb/s.

Figure 4-1—Relationship among CSMA/CD procedures

PLS_Service Primitives

Reconciliation sublayer

MII Signals

PLS_DATA.request

PLS_SIGNAL.indication

PLS_DATA.indication

PLS_CARRIER.indication

PLS_DATA_VALID.indication

Station Management

TX_ER

TXD <3:0>

TX_EN

TX_CLK

COL

RXD <3:0>

RX_ER

RX_CLK

CRS

RX_DV

MDC

MDIO

**Figure 22-3—Reconciliation Sublayer (RS) inputs and outputs,
and STA connections to MII**

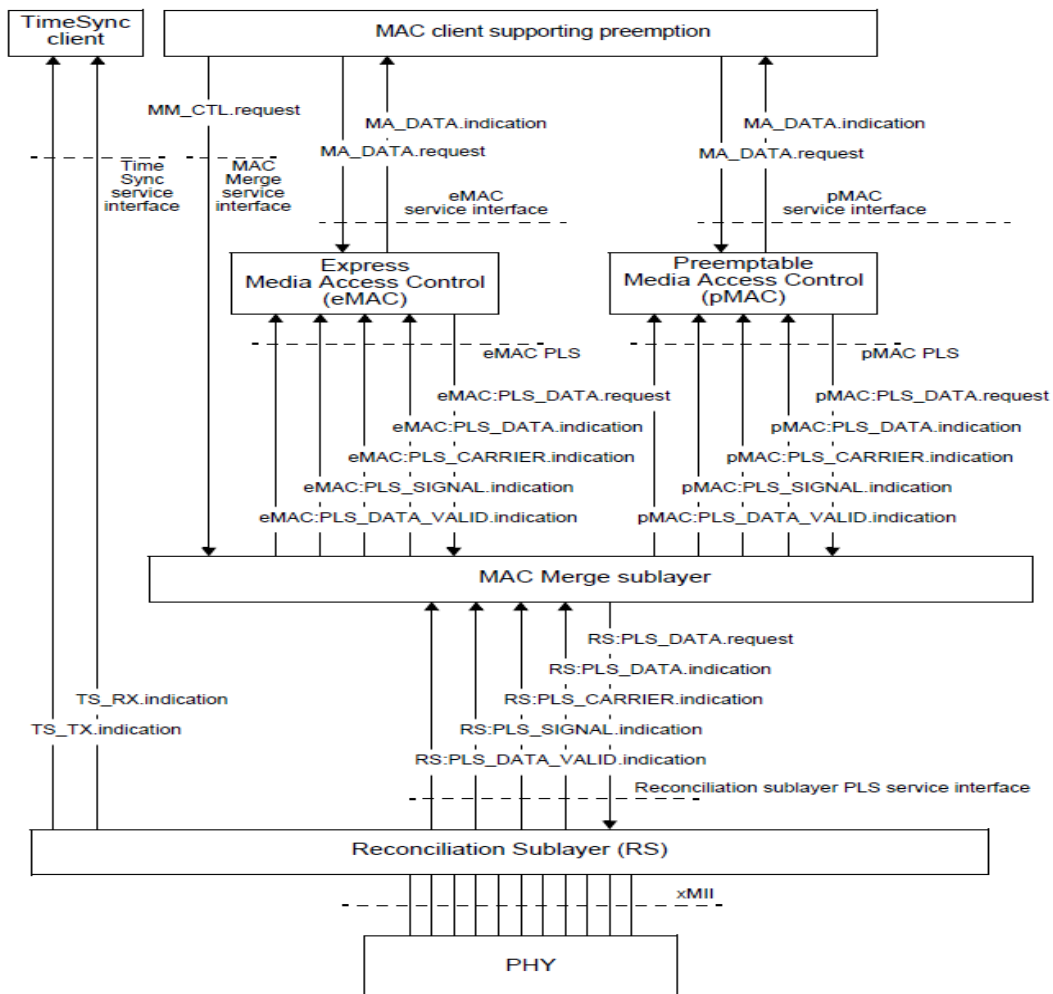


Figure 99–2—MAC Merge sublayer service interfaces diagram

Thanks!