

Differentiated Services for 10BASE-T1S/M - Half Duplex Transmission Selection

Peter Jones – Cisco Systems

Overview

The problem statement presented in [Consensus Building for 10BASE-T1S Differentiated Services: PLCA Multidrop Latency, Jitter, and scalability enhancements](#) is

10BASE-T1-S multidrop are being widely adopted in automotive Ethernet applications across e.g., body, chassis, and powertrain control applications, whose functions span a broad range of latency and jitter requirements, from tightly bounded control loops to more timing-tolerant comfort and convenience functions.

The Physical Layer Collision Avoidance (PLCA) layer developed as a part of 802.3cg and refined in 802.3da provides a fair and equal access to stations on the mixing segment while providing bounded maximum access latency. The maximum access latency is not sufficiently bounded for many of the tightly bounded control loops. Likewise, the variable length of the PLCA cycle results in high jitter (lack of determinism) that such applications cannot tolerate. PLCA as defined today is cannot treat traffic differently based on latency and jitter requirements. Additionally, adding traffic to multidrop networks engineered to meet latency and jitter requirements can change the latency and jitter of existing traffic in ways that are difficult to predict resulting in scalability issues and the need to frequently reengineer networks.

I restate this as:

The current set of 802.1 and 802.3 tools do not provide sufficient control over the selection of the next frame to be transmitted on the shared media.

802.1Q provides extensive tools for a station to choose which frame to transmit next (8.6.8 Transmission selection), but the selection happens independent of the station gaining access to the media. Several of these tools have been designed for use with full duplex media.

802.3 Media Access Control is class/priority unaware (802.1AC-2016 - Ethernet convergence function and 802-3-2022 Detailed service specification). Using 802.3 half duplex media with:

- CSMA/CD: provides no guarantee of when a given station will get access to the media
- PLCA: provides bounded access time for a station with worst case delay in the order of (station count * MTU transmission time).

Modifying PLCA can change when a specific station gets access to the media, but not which frame is sent by that station. Which frame is sent is controlled by 802.1Q Transmission selection. 802.1Q Transmission selection operates independently of media access. A frame is selected for transmission regardless of potential media access delay. This is an issue for half duplex media where access may be granted significantly later than the transmission selection decision. A similar issue can be seen on full duplex media using 802.2 PAUSE function (802-3-2022 Annex 31B).

I think addressing the variable delay between transmission selection and media access is key to addressing the problem statement above, and the solution space includes 802.1, 802.3 half duplex CSMA/DC and 802.2 full duplex PAUSE.

Transmission Selection for Half Duplex Networks

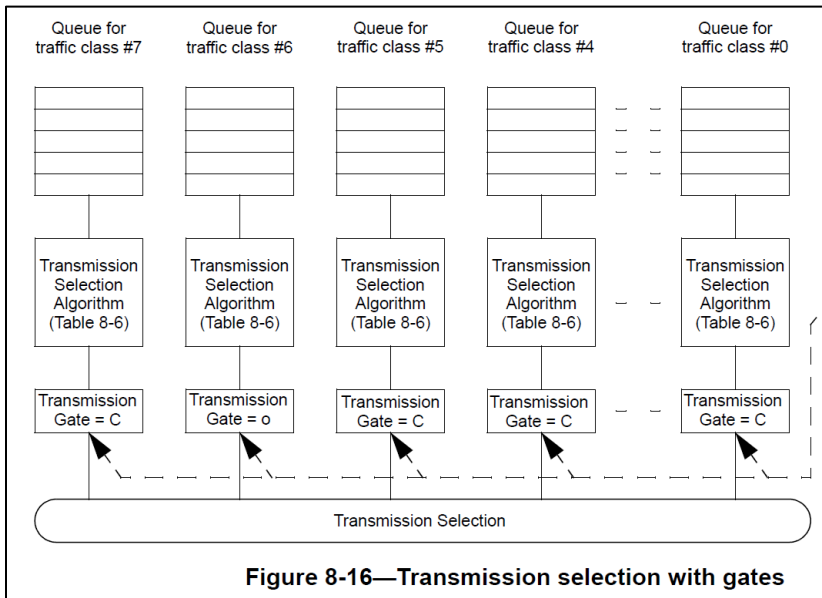
Background

Half duplex networks create issues with 802.1Q Transmission Selection(8.6.8). Transmission Selection supports several scheduling algorithms as shown in Table 8-6.

Table 8-6—Transmission selection algorithm identifiers

Transmission selection algorithm	Identifier
Strict priority (8.6.8.1)	0
Credit-based shaper (8.6.8.2)	1
Enhanced Transmission Selection (ETS) (8.6.8.3)	2
ATS Transmission Selection (8.6.8.5)	3
Reserved for future standardization	4–254
Vendor-specific Transmission Selection algorithm value for use with DCBX (D.2.8.8)	255
Vendor-specific	A four-octet integer, where the most significant 3 octets hold an OUI or CID value, and the least significant octet holds an integer value in the range 0–255 assigned by the owner of the OUI or CID.

Figure 8-16 shows the structure used by Transmission Selection including transmission gates as specified in Enhancements for scheduled traffic(8.6.8.4):



Frame transmission on half duplex media can produce unexpected results because the transmission selection decision is made before the station has access to the media. Consider a half-duplex point-to-point

100Mb/s Ethernet link. If the peer is transmitting when a frame is selected, the frame will not be sent to the media until the peer is done transmitting. This can delay transmission by up to an MTU worth of delay.

Processing of PAUSE frames (802.3-2022 - Annex 31B MAC Control PAUSE operation) creates a similar situation. Annex 31B specifies that MAC Control generates a MA_CONTROL.indication signaling the pause status to the MAC Client. 802.1AC-2016 clause 13.1 Ethernet convergence function specifies that these indications from the Ethernet MAC are passed to the MAC Client as M_CONTROL.indications. 802.1Q uses M_CONTROL.request and M_CONTROL.indication to implement Priority-Based Flow Control (PFC) but does not process any other M_CONTROL.indications. The function of PAUSE is fully specified in 802.3-2022 - Annex 31.

This situation is exacerbated in multidrop 802.3 segments (e.g., 10BASE-T1S/M) using CSMD/CD with multiple peers. Multiple nodes beginning transmission around the same time can cause collisions for an extended period. Each transmitting node has already performed transmission selection, and the system is controlled by the MAC back-off/retransmit function.

Using PLCA with 10BASE-T1S/M bounds the worst case media access delay to $(\text{nodes} - 1) + (\text{MTU transmission delay})$. PLCA is intended to prevent collisions by delaying access to the media.

Enhanced Half Duplex Transmission Selection

Transmission selection conducted when a node has access to the media produces results much closer to the desired behavior. This can be achieved by passing a “Media Available” signal to transmit selection to act as an additional constraint on the selection process.

This is analogous to the behavior of Priority-based Flow Control(PFC) as specified in 802.1Q Clause 36. PFC controls a vector of `Priority_Paused[n]` bits which act as an additional qualification in Transmission selection (8.6.8):

In a port of a Bridge or station that supports PFC, a frame of priority `n` is not available for transmission if that priority is paused, i.e., if `Priority_Paused[n]` is TRUE (see 36.1.3.2) on that port.

PFC is passed between the MAC Client and the MAC using M_CONTROL PFC primitives. These primitives are specified in 802.1Q 36.1.2 “PFC primitives” and 802.3 Annex 31D “MAC Control PFC operation”.

To support “Media Available” signaling:

1. Define a new MA_CONTROL.indication `MEDIA_AVAILABLE` opcode which notifies the MAC Client of half duplex media availability using a boolean `media_available` parameter
2. Define a new M_CONTROL `MEDIA_AVAILABLE` client in 802.1Q which sets the `Priority_Paused` vector directly from the `media_available request_operand`

Although 802.1AC-2016 mentions the use of M_CONTROL.request and an M_CONTROL.indication in the ISS to implement Priority-Based Flow Control, this text is only in a note. 802.1AC-2016 does not need to be updated when defining the `MEDIA_AVAILABLE` function .

To address:

- PLCA - generate the MEDIA_AVAILABLE indication from the PLCA Data State Machine CARRIER_STATUS variable(148.4.5.2).
- Half-duplex more generally - a new Reconciliation Sublayer extension analogous to clause 90 TSSI could be specified using the generic PLS_CARRIER.indication (6.3.2.1) to generate MEDIA_AVAILABLE indications.
- PAUSE – specify the 802.1Q M_CONTROL MEDIA_AVAILABLE client to also process M_CONTROL PAUSE indications

These options are all shown in the figure below.

