

PRIORITY SUPPORT FOR PLCA

WHAT MINIMALLY NEEDS TO BE DONE IN 802.3
SO THAT THE REST CAN BE DONE IN 802.1

DON PANNELL,
DONALD.PANNELL@NXP.COM

SEP 2018 – SPOKANE



SECURE CONNECTIONS
FOR A SMARTER WORLD



Agenda

- Overview
- The Problem
- Proposal to Support Priorities in PLCA
- Estimated Overhead & Logic Costs
- Things Still to Do
- Summary
- Backup Analysis & Proof of the Need



Supporters (alphabetical by last name):

- Gordon Bechtel, 802.1, Bechtel
- Tobias Belitz, 802.3, Renesas
- Piergiorgio Beruto, 802.3, Canovatech
- Rodney Cummings, 802.1, National Instrument
- Norm Finn, 802.1, Huawei
- Claude Gauthier, 802.3, NXP
- Craig Gunther, 802.1, Craig Gunther Consulting
- Mark Hantel, 802.1, Rockwell
- Don Pannell, 802.1, NXP
- Jordon Woods, 802.1, Analog Devices
- Wang Xinyan, 802.3, Huawei



PART I

Overview

First, When did Priority become an 802.3 Issue?

- Or, “Why is this the 1st time Priority support in an 802.3 PHY has come up?”
 - Because it wasn’t relevant or needed until now:
 - Priorities were added to Ethernet Frames (802.3ac) & Bridges (802.1p) in Sep 1998
 - And the 802.3cg project is the 1st PHY to use the 802.3 CSMA/CD MAC since then!
 - All other past 802.3 PHY projects supported full-duplex point-to-point links (or other MACs) so Priority considerations in a PHY was never an issue (EPON is point-to-multipoint w/TDM)
- Why isn’t this an 802.1 problem to solve?
 - It is...
 - But 802.1 can’t completely solve the problem by itself in the case where shared media is used
 - “802.1 only” solutions do not lower the worst case latency & can’t support TSN effectively
- What is needed from 802.3?
 - Optional Priority information exchange so that 802.1 can do all the policy decisions

Why is Priority support needed?

- To be consistent with what 802.1 has defined for TSN (Time Sensitive Networking)
 - The TSN architecture is a proven mechanism to support converged data on a single wire
 - Traditional Ethernet traffic (e.g., IP) shares the same link as Reserved &/or Time Critical traffic
 - The Ethernet frame's Priority Code Point bits (the “user priority” bits defined in 802.3ac-1998, now moved to 802.1Q Annex G as referenced by 802.3-2015 1.4.347) are used to distinguish between these various traffic types so that differentiated services can be applied to them
 - The mechanism used in TSN to differentiate the service applied to various traffic types is priority
 - For TSN to work, some frames must be treated better (unfairly) compared to other frames
 - This is needed to support the lowest possible latencies, to guarantee the reserved bandwidth, etc.
 - The success of TSN is evident by the increasing number of 802.1 use case TSN Profiles:
 - Audio Video Bridging (AVB) Systems – IEEE 802.1BA-2011
 - Time-Sensitive Networking for Fronthaul – IEEE 802.1CM-2018
 - TSN Profile for Industrial Automation – IEC/IEEE 60802 (PAR approved May 2018)
 - TSN Profile for Automotive – IEEE 802.1?? (PAR to be voted on in Nov 2018)

802.3 & 802.1 have a Long History of Jointly Solving Issues

- VLAN Tag (802.3ac) – where the Priority bits reside
- Frame Expansion (802.3as) – initially for MACSec's extra fields
- Energy-efficient Ethernet (802.3az) – need to know frames need to be transmitted
- MAC Control Frame for Priority-based Flow Control (802.3bd) – for data centers
- Ethernet Support for the IEEE P802.1AS Time Synchronization Protocol (802.3bf)
- Interspersing Express Traffic (802.3br) & Preemption (802.1Qbu)

- We need to work together on this issue too!
 - 802.1 can't solve the problem by itself without additional data from 802.3
(<http://www.ieee802.org/1/files/public/docs2018/new-TSN-pannell-QoS-for-PLCA-0518-v02.pdf>)

- This proposal has 802.3 transferring priority information only to/from 802.1 and then all the priority policy decisions are made in 802.1
 - This minimizes the impact to 802.3 & maximizes the flexibility for 802.1

Status of TSN's Fast-Path Standards on Today's PLCA Networks

- 802.1AS, Precise Timing Protocol, works (even w/out PLCA – demonstrated)
- 802.1Qbv, Time Aware Shaper, works & could be a replacement for PLCA
 - At a large loss of bandwidth & increased latency due to fixed schedules (demonstrated)
- **802.1Qav, Credit Based Shaper, BROKEN – see Backup 1**
 - Fixed with the proposed Priority Advertise in PLCA
- **802.1Qcr, Asynchronous Traffic Shaper, BROKEN – see Backup 2**
 - Fixed with the proposed Extended Priority Advertise in PLCA (with 4-bit priority values)
- 802.1Qch, Cyclic Queuing & Forwarding, **probably BROKEN**
 - These frames need to egress in certain windows and Backup 1 & 2 can insure that
 - 802.1Qch is really an improved version of the Credit Based Shaper – at a bit more cost
- Future new Shapers, **likely BROKEN**
 - Future shapers are going to follow the 802.1 TSN Transmit Architecture – explained next
 - So these would likely work with the proposed Priority Advertise as well!

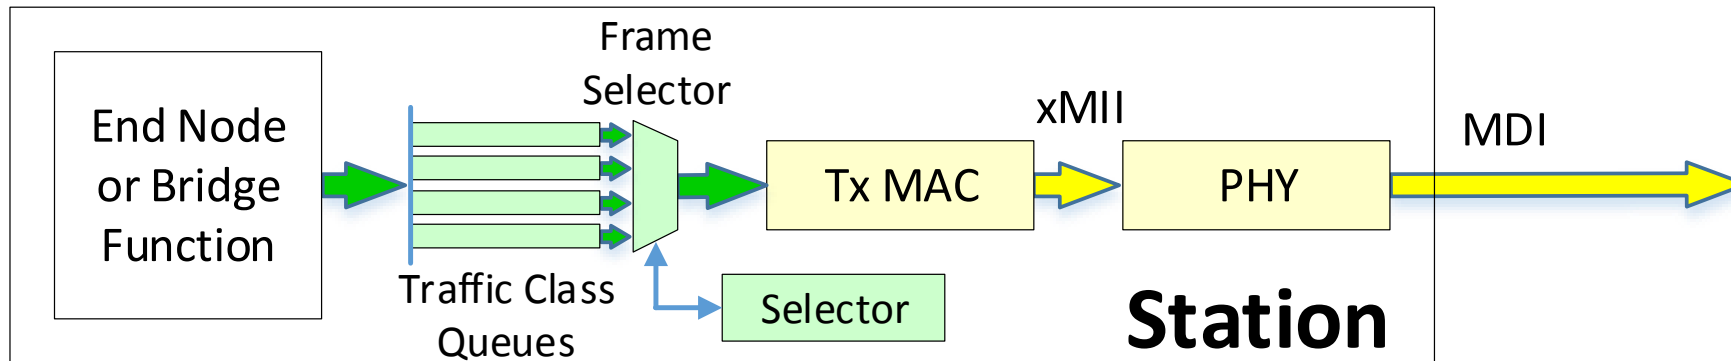


PART 2

The Problem

The TSN Transmit Architecture

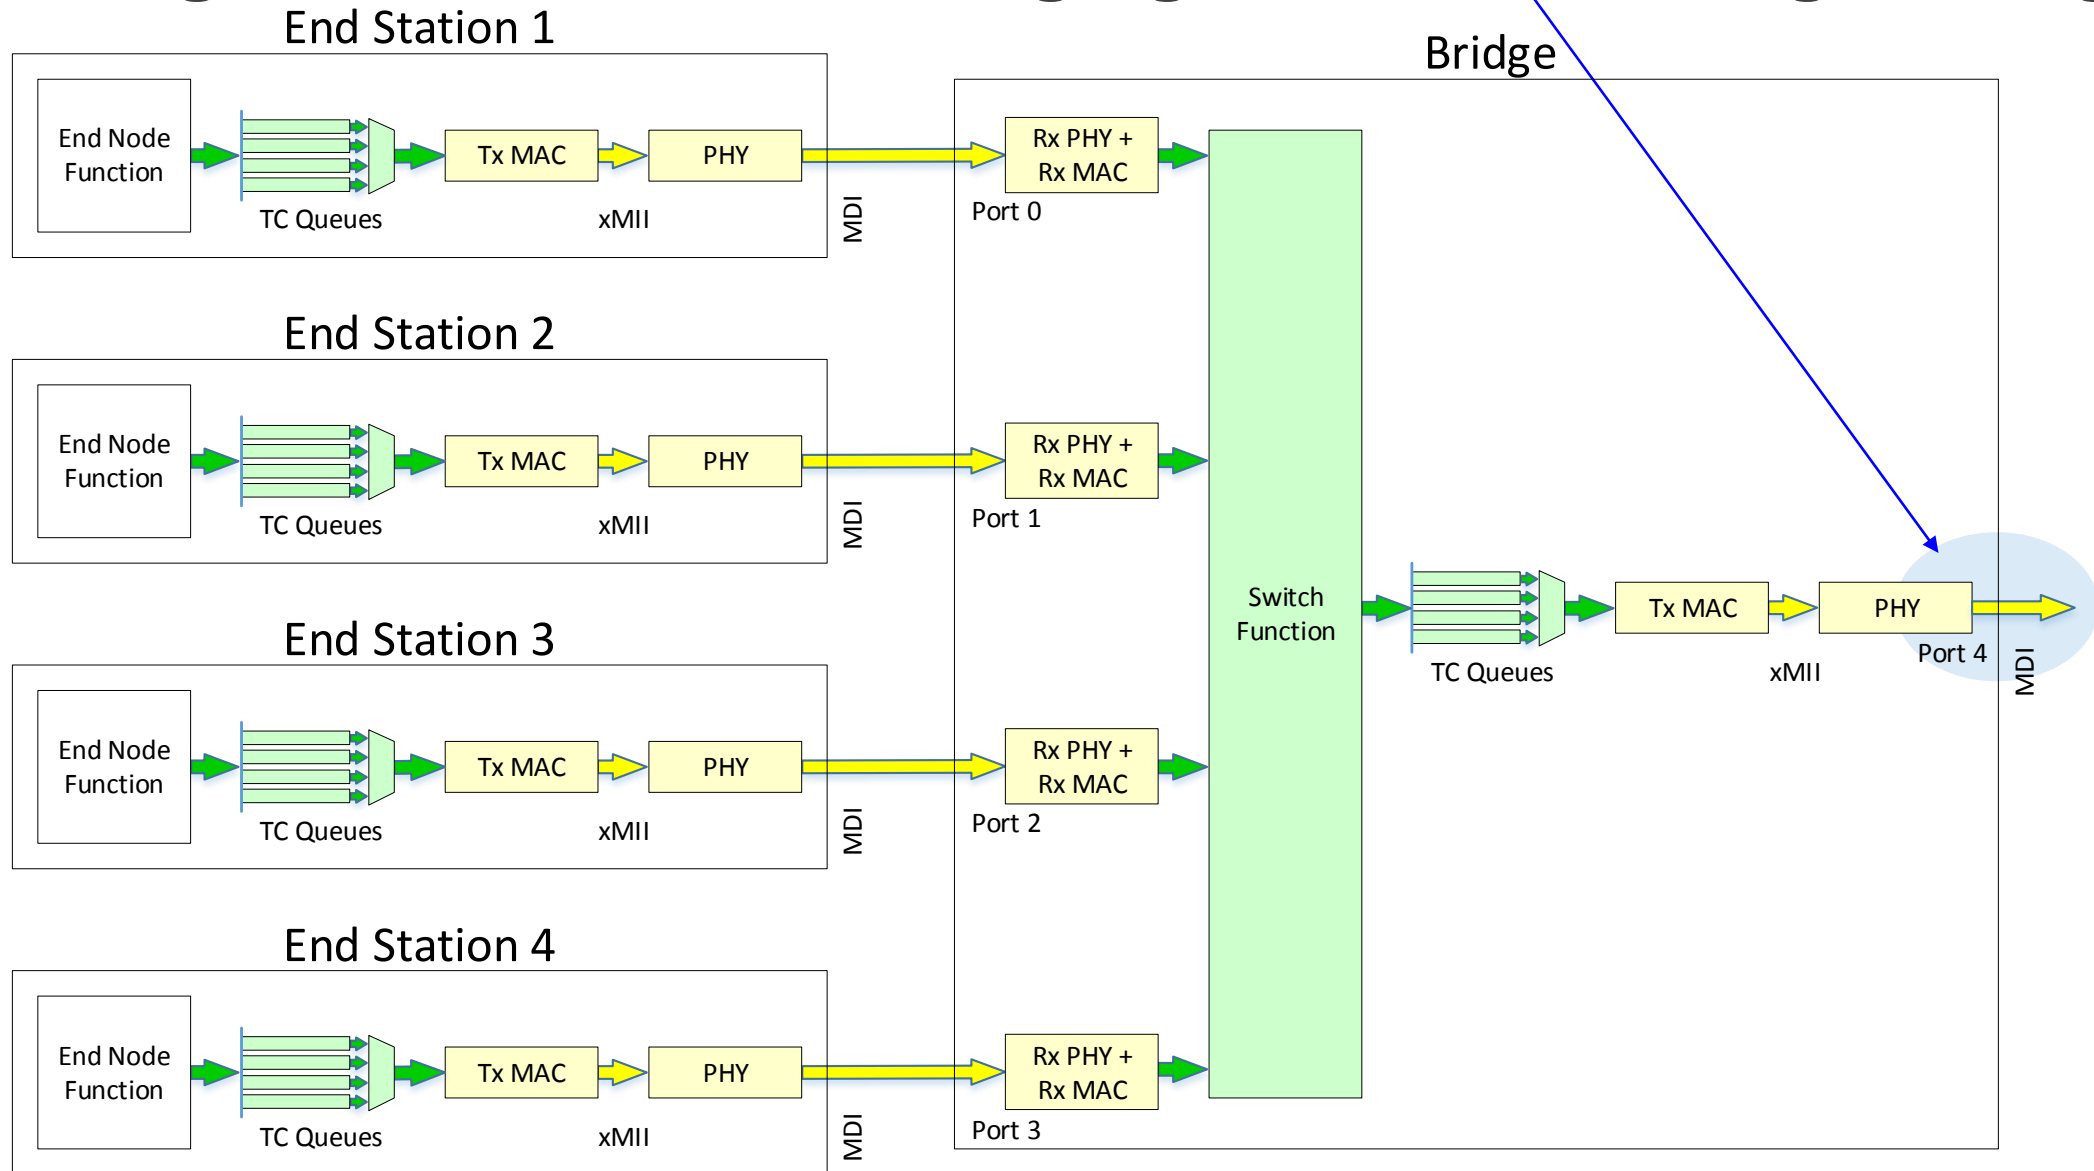
- The 802.1 portion:
 - Stations have 1 or more transmit Traffic Class Queues
 - In TSN, these Queues are used to separate the various traffic types in order to support the differentiated services that each traffic type needs
 - The Selector knows what Queues have frames to transmit & the “Queue’s configured traffic type”, and based on this, selects the next frame to transmit
 - Various Selector algorithms are supported in 802.1 to support the various traffic types
 - Normally this next frame is released to the Tx MAC right away



Green = IEEE 802.1

Yellow = IEEE 802.3

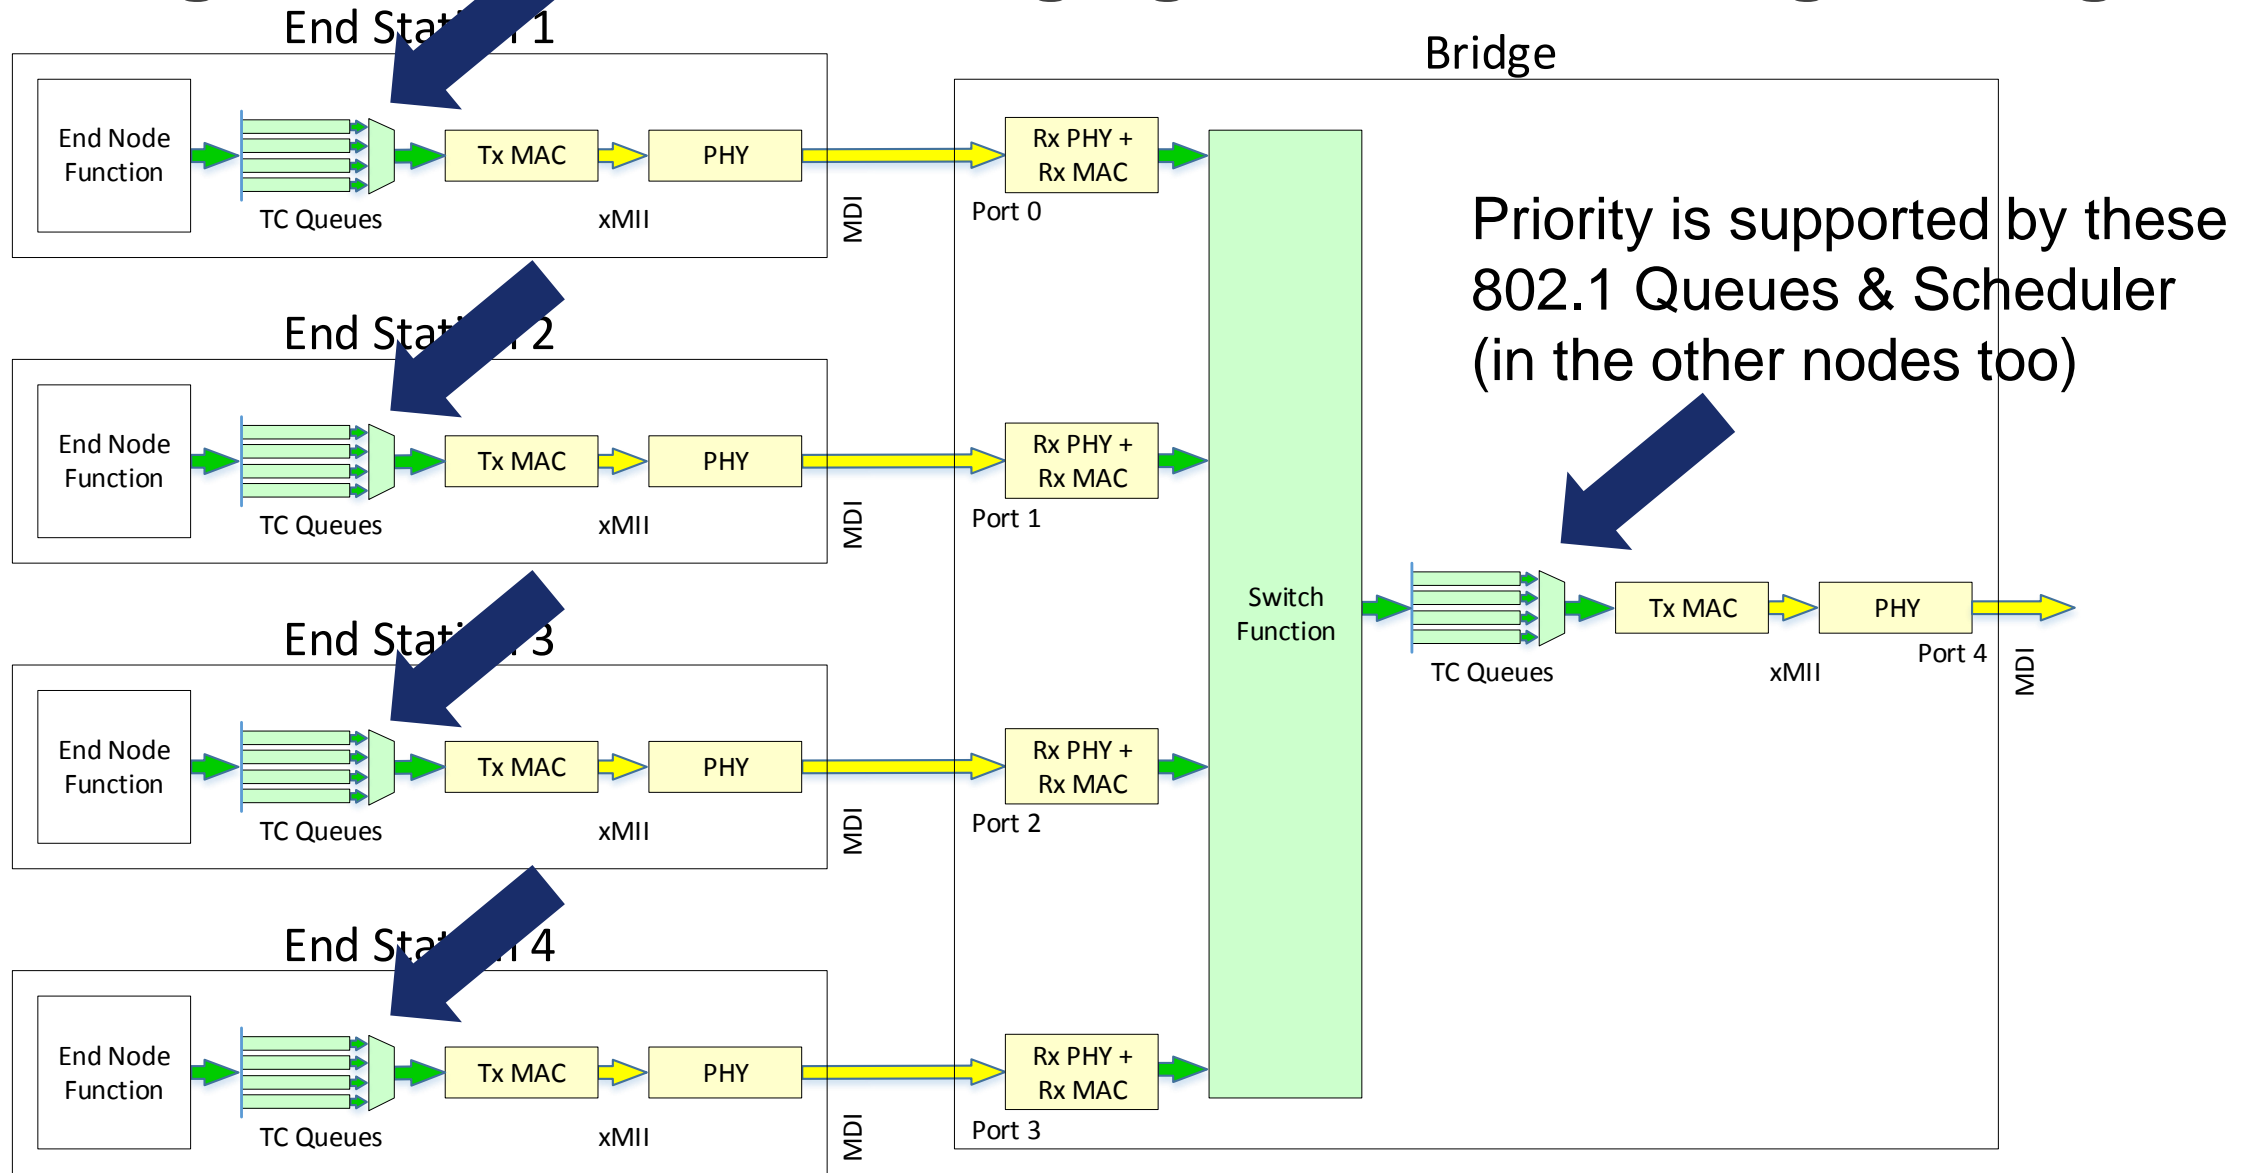
Extending this to 4 Stations merging to 1 Link – Using a Bridge



Green = IEEE 802.1

Yellow = IEEE 802.3

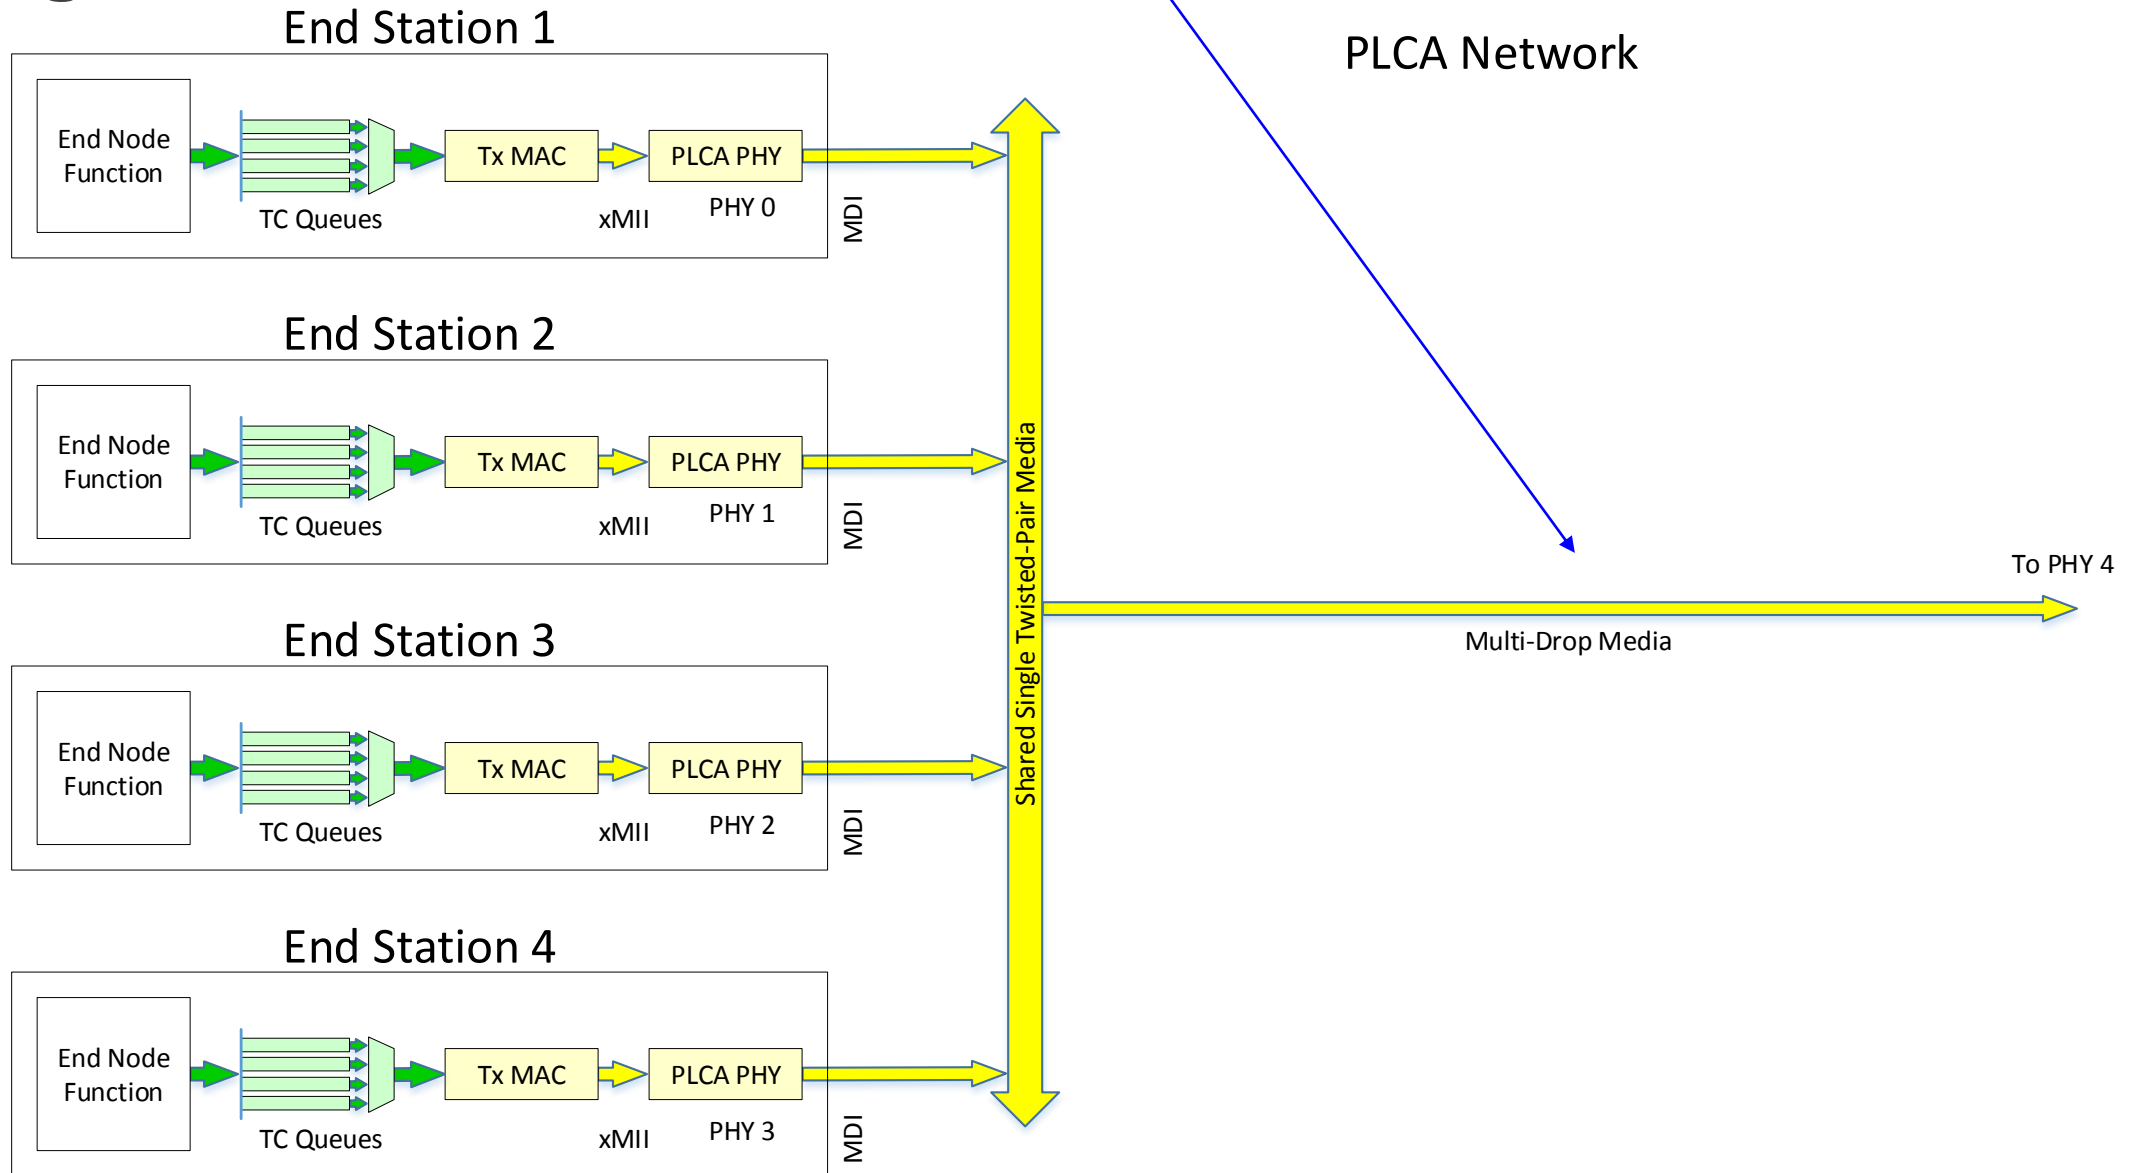
Extending this to 4 Stations merging to 1 Link – Using a Bridge



Green = IEEE 802.1

Yellow = IEEE 802.3

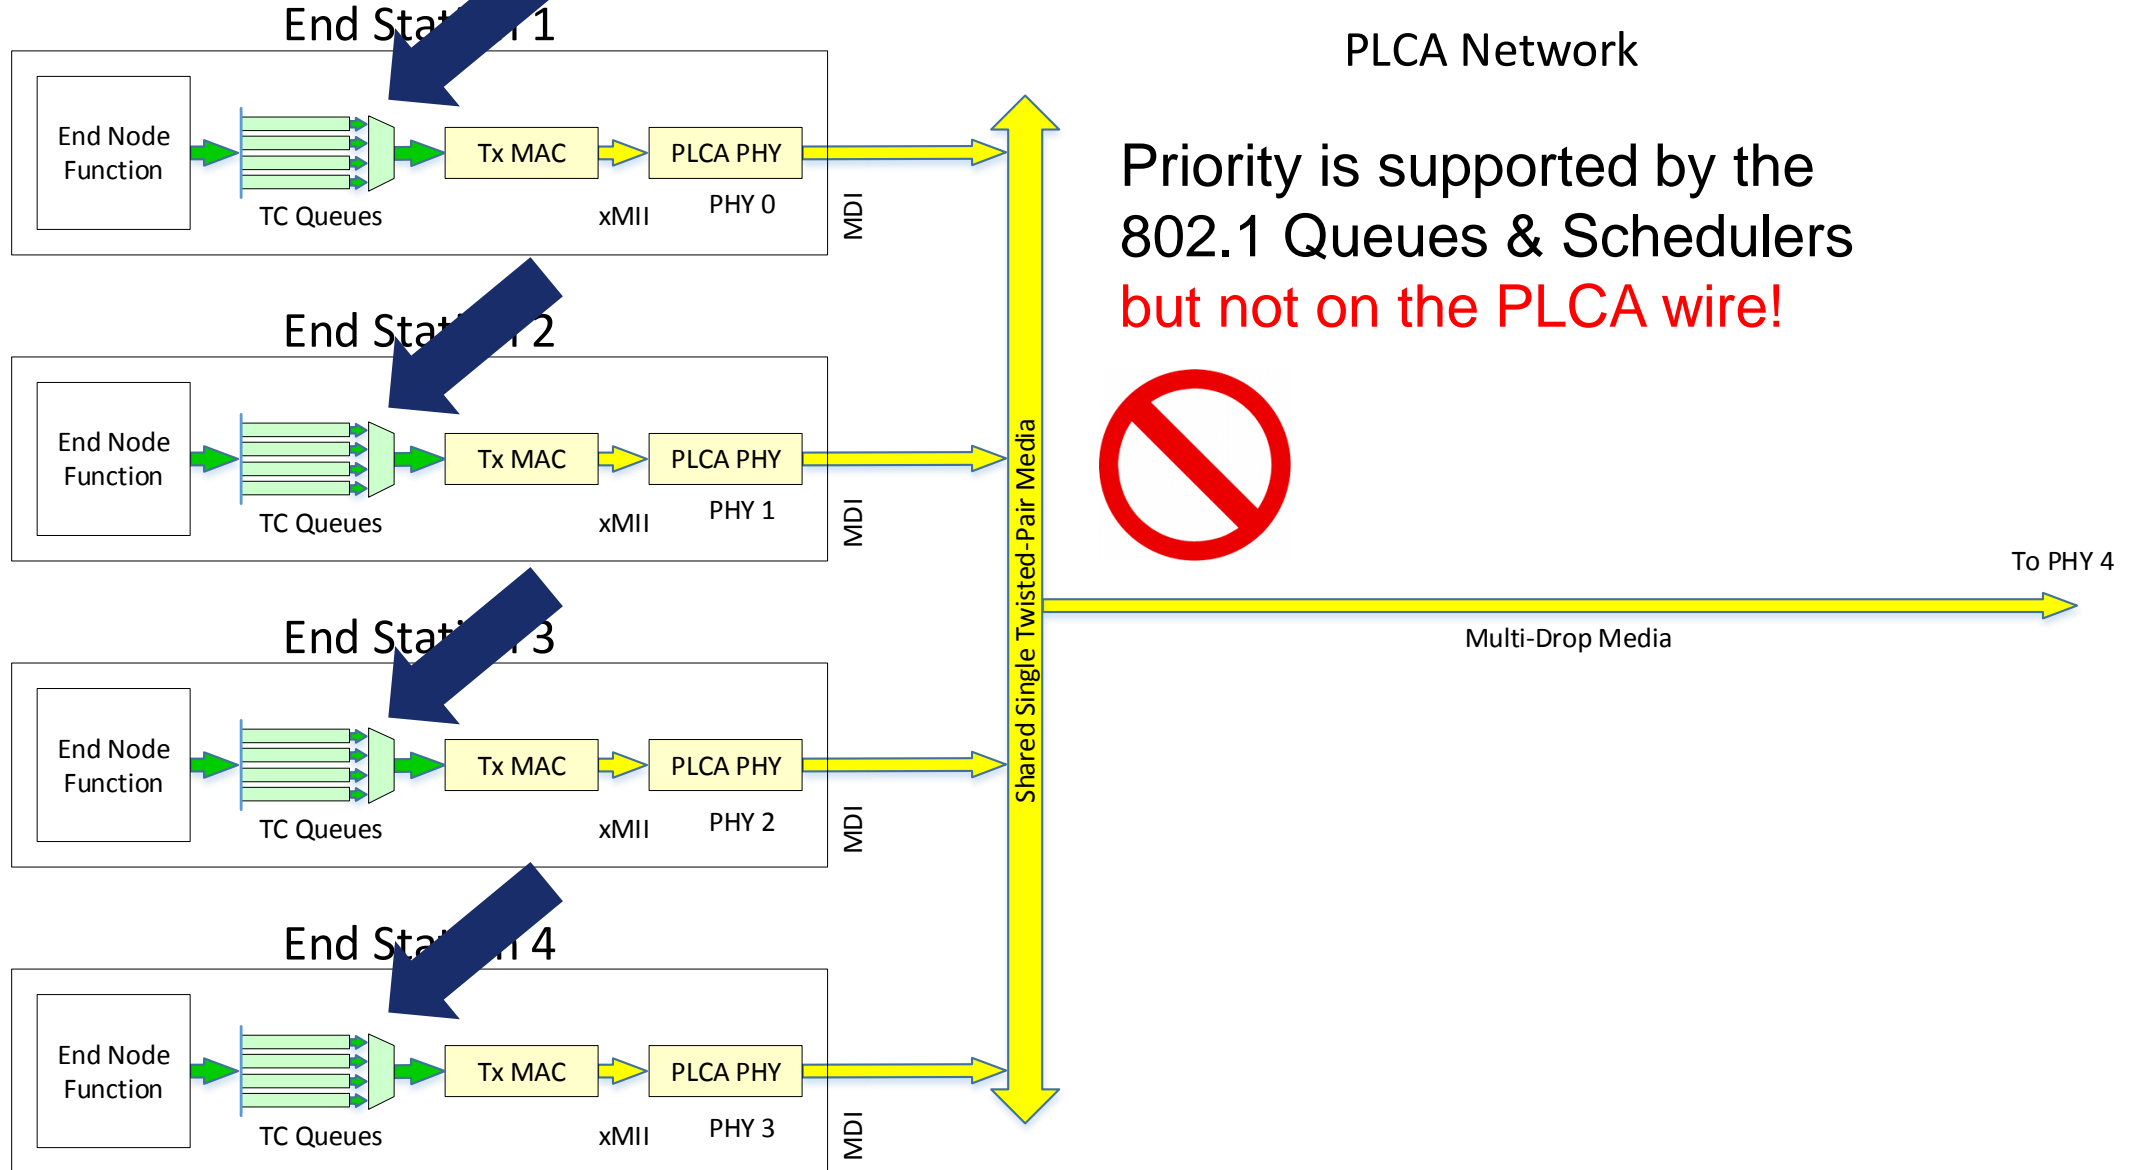
Change the 4 End Stations to 1 Link to Now Use PLCA



Green = IEEE 802.1

Yellow = IEEE 802.3

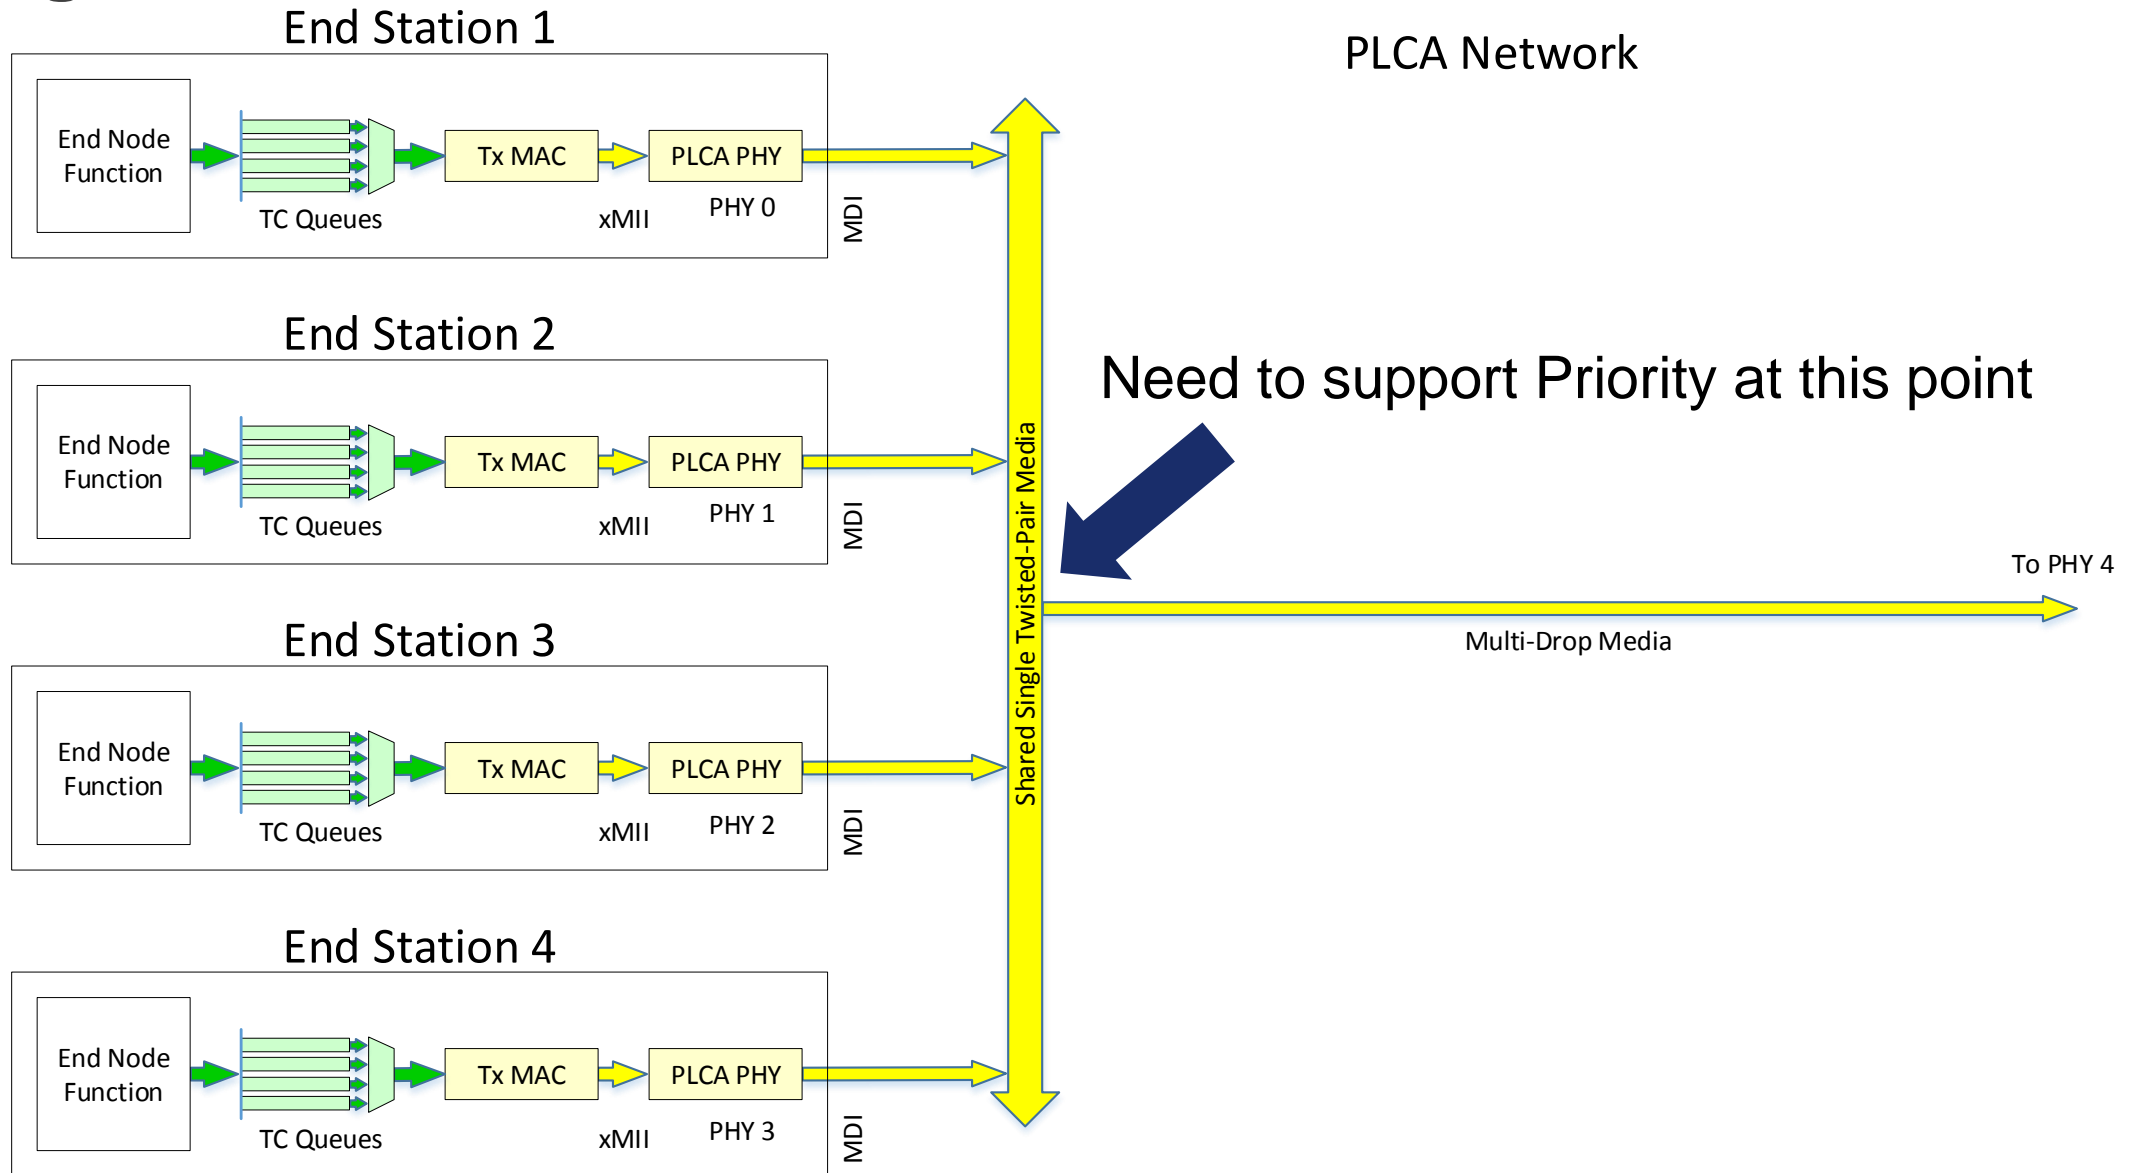
Change the 4 End Stations to 1 Link to Now Use PLCA



Green = IEEE 802.1

Yellow = IEEE 802.3

Change the 4 End Stations to 1 Link to Now Use PLCA



Why is Priority support needed?

- PLCA supports equal transmit opportunities between all the nodes on the wire
- It also supports unequal transmit opportunities per node such that you could define:
 - A Node transmit opportunity order of 0, 1, 0, 2, 0, 3, 4, 5 or some other order
 - This example gives Node 0 three times the opportunities compared to the other nodes
- While this approach gives higher opportunities to some nodes, ...
- The goal of Priorities is to support different (non-equal) treatment of frames
 - Node 0 may be transmitting only low priority data but it still gets its 3 Tx opportunities!
- The PLCA transmit opportunity schedule is static
 - And this approach must be “engineered” to work right, requiring more ‘user’ effort
- Priorities support differentiated services dynamically self adjusting as expected!
 - The transmit opportunities of each Node needs to change based on the frame it is Tx’ing
 - Priorities are an integral part of TSN (as long as they follow the 802.1 model)
 - Therefore, there is no extra ‘user’ effort needed in order to use Priorities

Why is Priority support needed?

- Other ‘solutions’ have been presented that solve the problem of a specific use case by limiting the size of the lower priority frames (somehow)
- This is an “engineered” solution to a problem that would automatically be solved by using 802.1 TSN methods (i.e., Priorities)
- The TSN target guarantees can only be met if the TSN architecture/methods are used throughout the network
 - Assuming TSN is used on the higher speed links in a car, it would be only natural to want to extend TSN to the 10SPE links, but without Priorities, this is hard, if not impossible
 - Making exceptions in certain parts of a network make it really hard, & takes a lot of extra engineering effort to be convinced, that the TSN portion will work as desired
- That being said, any solution must be cost effective, not just in the purchase of components, but also in the cost of using those components
 - More on this later



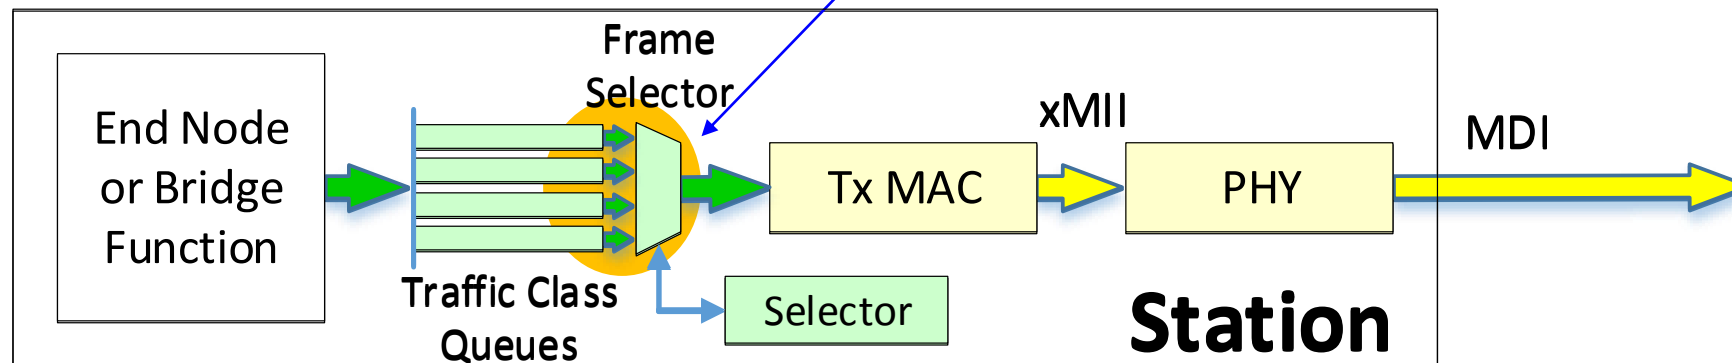
PART 3

Proposal to Support Priorities in PLCA

Just the Priority Information Exchange only!

What is Minimally Needed so 802.1 Can Do the Rest (Policy)?

- Ideally PLCA should mimic as close as possible what is done inside a Bridge
- The Selector in Bridges minimally need to know which Queues (Priorities) have frames ready to be transmitted, and what Queue (Priority) they are in
- That it!
- Based on this information the Selector determines the next frame to transmit
 - This decision (the Policy of which frame goes next) can be quite complex which is why 802.1 has support for many different types of Queue Shapers ahead of the Selector
 - And 802.3 does not need to worry about that part!



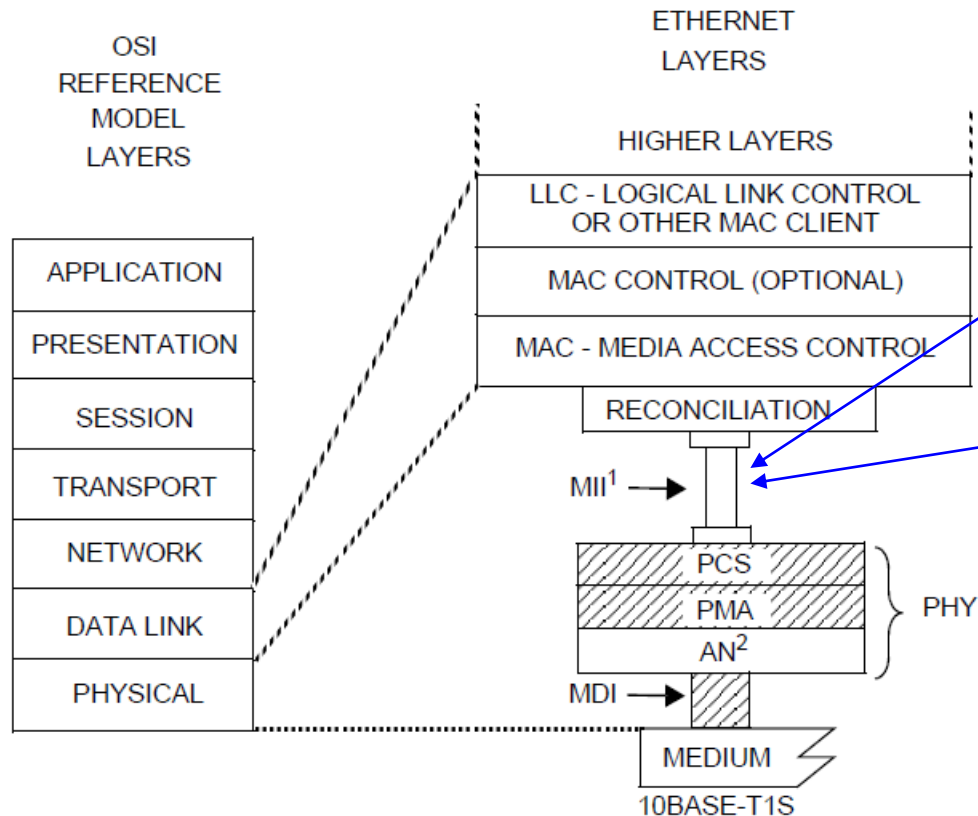
Green = IEEE 802.1

Yellow = IEEE 802.3

Communicating the Priority of the Next Frames to Tx with PLCA

- Before each regular (data) Beacon, an **optional** Advertise Beacon is issued
 - PHY 0 issues this Advertise Beacon and then indicates the Priority of its 'frame to Tx'
 - The Priority to advertise of the 'frame to Tx' is given to the PLCA function by 802.1
 - The PLCA function sends this Priority down the MII using 1 new code followed by the Priority
 - The PHY then transmits this information similar to how a Beacon is indicated on the wire
 - All the other PHYs see this new code and send the observed Priority up their MII's
 - The PLCA function sends to 802.1 the received Advertise Beacon and all observed Priorities
 - Each PHY in the normal sequence issues its 'frame to Tx' Priority in the same way
 - After all PHYs have advertised their Node's 'frame to Tx' Priority, PHY 0 issues the normal (data) Beacon and the data phase begins
 - At this point in time all the 802.1 Selectors know the current Priority for all the frames that wish to be transmitted on this Beacon cycle
 - Each local Selector then release to the Tx MAC its 'frame to Tx' only if it meets the currently configured 802.1 Policy

Communicating the Priority of the Next Frames to Tx over the MII



MDI = MEDIUM DEPENDENT INTERFACE
MII = MEDIA INDEPENDENT INTERFACE

NOTE 1—MII is optional
NOTE 2—Auto-Negotiation is optional

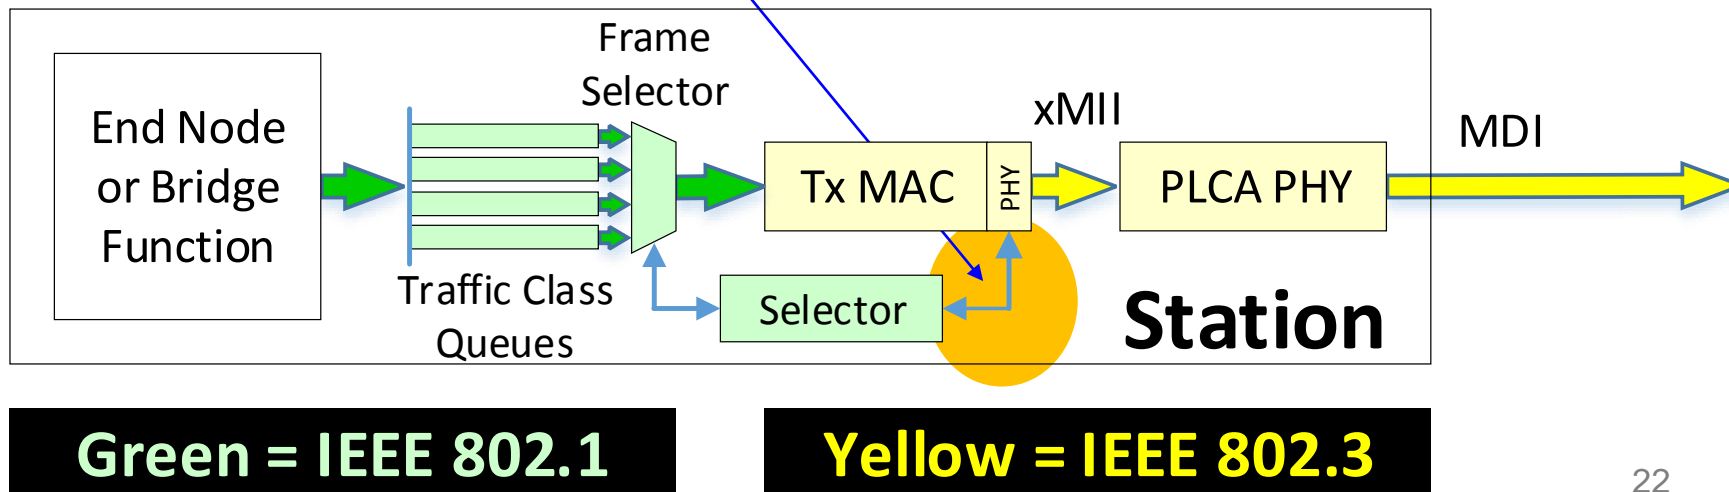
PCS = PHYSICAL CODING SUBLAYER
PMA = PHYSICAL MEDIUM ATTACHMENT
PHY = PHYSICAL LAYER DEVICE
AN = AUTO-NEGOTIATION

- The PLCA function (in the Reconciliation sublayer) issues the locally selected frame's Priority to advertise down the MII
- And all the neighbor's priorities that they advertise gets sent up the MII
- Then 802.1 decides if the local frame gets to be transmitted during this Beacon cycle or not
- Creating a new "Advertise Beacon" allows re-use of existing state machines, saving logic

Figure 147-1—Relationship of 10BASE-T1S PHY to the ISO/IEC OSI reference model and the IEEE 802.3 Ethernet Model

Communicating the Priority of the Next Frames to Tx to 802.1

- The 802.1 Selector knows what frame it wants to transmit next & gives its Priority to the PLCA function (going around the MAC)
- The 802.1 Selector will 'see' the Advertise Beacon and all the advertised Priorities from itself and all the other Nodes on the multi-drop network segment
- The 802.1 Selector will then send its selected frame to its Tx MAC if the frame matches the currently configured 802.1 rules
- A basic rule would be to only release frames that match the highest advertised Priority (this is called Strict Priority)



Communicating the Priority to/from 802.1 Using EEE's Model

78.1.1.1 Reconciliation sublayer service interfaces

Figure 78–1 depicts the LPI Client and the RS interlayer service interfaces.

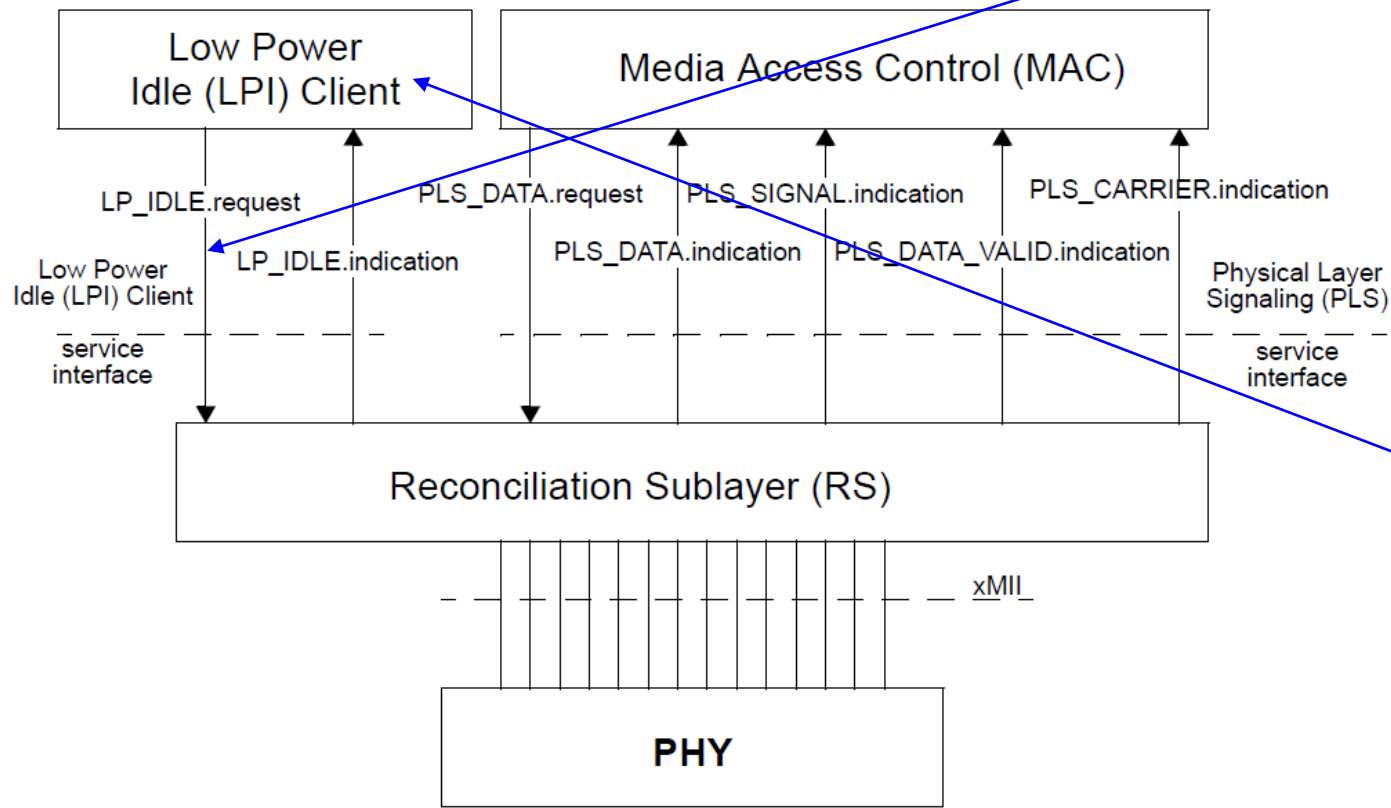
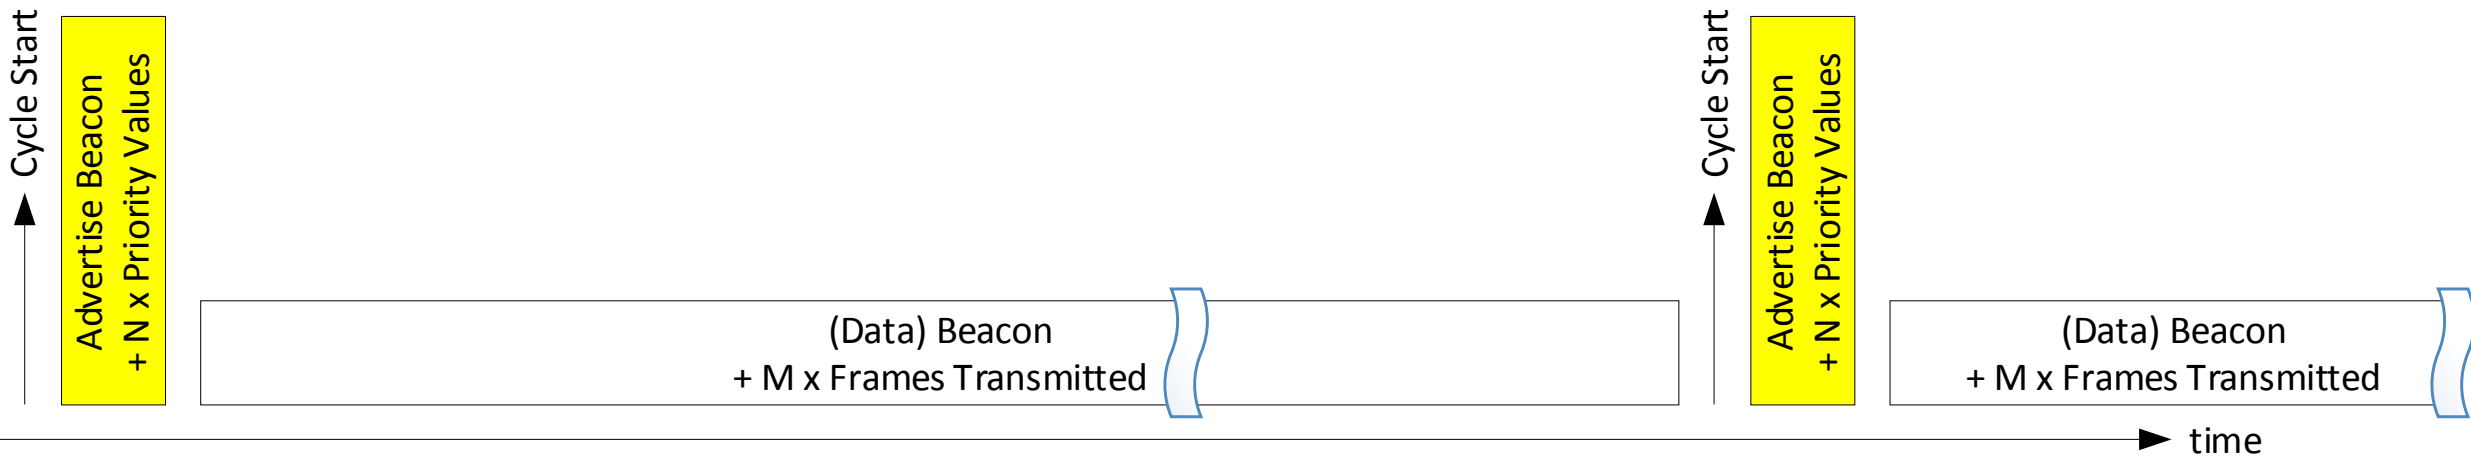


Figure 78–1—LPI Client and RS interlayer service interfaces

- Energy Efficient Ethernet (EEE) needs to know when 802.1's queues are empty so it can power down the PHY
- And it needs to know when 802.1's queues have a frame so it can power up the PHY
- This connection to/from 802.1 is via sideband signals to a Low Power Idle Client so that the MAC is not changed in any way
- The Priorities & Advertise Beacon can use the same method and connect to a **Priority Client** – going around the MAC!

Timing of the Optional PLCA Priority Advertise Mechanism

- The below figures shows where the Priority Advertise takes place
 - It's at the beginning of each cycle
 - You can think of it as an Advertise Phase before the normal Data Phase of each cycle
- This is NOT drawn to scale! That will be shown next
- The location of the Priority Advertise phase, gets 802.1 the information it needs in time for it to support TSN's differentiated services



Where $N = \# \text{ Nodes}$, $M \leq N$



PART 4

Estimated Overhead & Logic Costs

Scale Timing of the Optional PLCA Priority Advertise Mechanism

- This figure shows 8 Nodes transmitting maximum size frames + overhead



- This figure shows the added overhead of 8 Nodes of Priority Advertise (to scale)



- The 1st figure is $8 \times (\sim 1522 + 20) = 12,336$ bytes
 - The 2nd figure is ~ 32 bytes for Priority Advertise (per discussion w/PLCA Editor)
 - The added latency for Priority Advertise with Max size frames is $\sim 0.03\%$
-
- This figure shows 8 Nodes transmitting minimum size frames + overhead



- This figure is $8 \times (64 + 20) = 672$ bytes
- The added latency for Priority Advertise to Min size frames is $\sim 4.5\%$
- In all cases the added latency for Priority Advertise is $\#Nodes \times 4$ bytes

Gate Estimate of the Optional PLCA Priority Advertise Mechanism

- The reason to create an Advertise Beacon along with the original Beacon is to re-use as much logic & state machines as possible (per discussion w/ PLCA Editor)
- It is anticipated that this Advertise Beacon logic added to the current PLCA function will add ~ 1000 gates or less (per discussion w/ PLCA Editor)



PART 5

Things Still to Do

Things Still to Do in the 802.3cg Draft

- Propose the actual text changes for the next draft which was listed in the comment resolution (#573 http://www.ieee802.org/3/cg/comments/802.3cg_draft2p0_Received_By%20Comment%20ID.pdf):
 - 1) Add a new PRIORITY encoding to Tables 22-1 & 22-2 (the MII interface - p25 & p26). A single encoding is all that is needed as the Priority value indication can follow the PRIORITY code.
 - 2) Add PRIORITY 4B/5B encoding to Table 147-1 (p151) or some other mechanism.
 - 3) Update figure 148-3 (p176) to add connections to a "Priority Client" as was done for Energy Efficient Ethernet's Fig 78-1 (p33 of part 6 of 802.3-2015).
 - 4) Update Fig 148-4 (p181) PLCA Control state diagram and associated text to add in the optional Priority communication phase at the start of each BEACON. The goal here is to reuse as much as possible to minimize gate costs.
 - 5) A register bit will be needed to enable this optional feature, a few PICS added, etc.



PART 6

Summary

Summary

- Adding Priority Advertise support to PLCA is:
 - An optional feature for multi-drop 10BASE-T1S only – this does not apply to T1L
 - It's expected gate cost (~1000 gates) is not measurable in \$'s using today's geometries
 - It's latency overhead is tiny when enabled (~ 4 byte times/Node)
 - It's latency overhead is zero when disabled
- It is needed in order to fully support TSN's differentiated services per traffic type
 - It's as close to the mechanism used in Bridges as it can be
 - And it supports the lowest possible latencies, etc. for PLCA's multi-drop mode
- 802.3 only needs to add the mechanism to communicate the Priorities
- 802.1 will do all the policy mechanisms to select which frames are transmitted
 - Initially all that is needed is a Strict Selector which has been defined since 1998
- The MAC and its interfaces are not modified!

Summary – part 2

- Some have stated that this is too complex to add to 802.3cg at that this time
 - It was too complex before, but it is now quite simple
 - All that is needed in 802.3 is the communication of Priority
- Some are concerned that there is no 802.1 project to support Priority Advertise
 - Here we have the famous “chicken and egg – which came first” issue
 - 802.3’s EEE did not wait for an 802.1 project to define what connects to the Low Power Idle Client – the work needed to get done and it got done
 - 802.1’s gPTP did not wait for an 802.3 project to define how to timestamp frames in a PHY – the 802.1 work completed and then an 802.3 project defined the specifics
 - 802.3cg’s addition of Priority Advertise can complete without contingencies on 802.1
 - And the list of Supporters shows that getting a project started in 802.1 will happen
 - The analysis shows that all 802.1 needs to define is the 20 year old Strict Priority anyway
- The time to add Priority Advertise is now – as nobody likes to redefine a PHY

Recommendations for 802.3cg

- Support a 4-bit wide Priority Advertise on the media to support 16 priorities
 - This adds zero overhead from the minimum requirement of 3-bits as the interfaces (the MII & the symbols on the media) already support nibble (4-bit wide) data (see Backup 2)
- Two, very different, TSN use cases are shown next that need Priority Advertise
 - Both of these use cases came from use cases that 802.1 TSN has been asked to support and where standards are done or in process
- The support of multiple use cases using the same Priority Advertise mechanism in very different ways, shows this is the correct location for this mechanism
 - This is no surprise since Priority Advertise's goal was to mimic what 802.1 already did
- 802.1 only solutions to solve these use cases (without using Priority Advertise) have been developed by many people
 - But they don't work as they are much more complex, do not get as low of a latency, & waste a lot of link bandwidth, as many consecutive idle Beacon cycles are required (the only signal 802.1 can 'see' up the MII today)

Thanks

Questions?



Backup 1


Worst Case Latency Improvement

This analysis also shows that one of the key functions of the 802.1Qav Credit Based Shaper (used in AVB) cannot be solved without the proposed per Beacon dynamic Priority information


Worst Case Latency Improvement Example

- 802.1 examines use cases to help determine general needs for a standard improvement
- Then 802.1 likes to look at worst case situations because, as experience has shown, these worst cases ALWAYS show up in networks
- The following example shows 4 End Stations talking to an off-page 5th station
 - The 5th station is not trying to transmit anything so the math is a 4 PHY, 4-to-1 example
 - Makes calculation extrapolations for 8 & 16 node networks easier
 - The 1st three stations have two low priority maximum size frames to transmit
 - And the last station initially has one AVB high priority minimum size frame to transmit
- Again, this is intentionally a worst case setup requiring NO “engineering of the network” as even engineers make mistakes – some which are really hard to find!


End Station 1




End Station 2



End Station 3



End Station 4



PLCA Network

Latency without Priorities

Worst case: low priority max size (red) frames ahead of high priority min size (blue) frames

Shared Single Twisted-Pair Media

Multi-Drop Media

To PHY 4

State of the Queues at T_0

At t_0 , all frames are ready for transmit when PHY 0 issues the Beacon (not shown)

End Station 1

PLCA Network

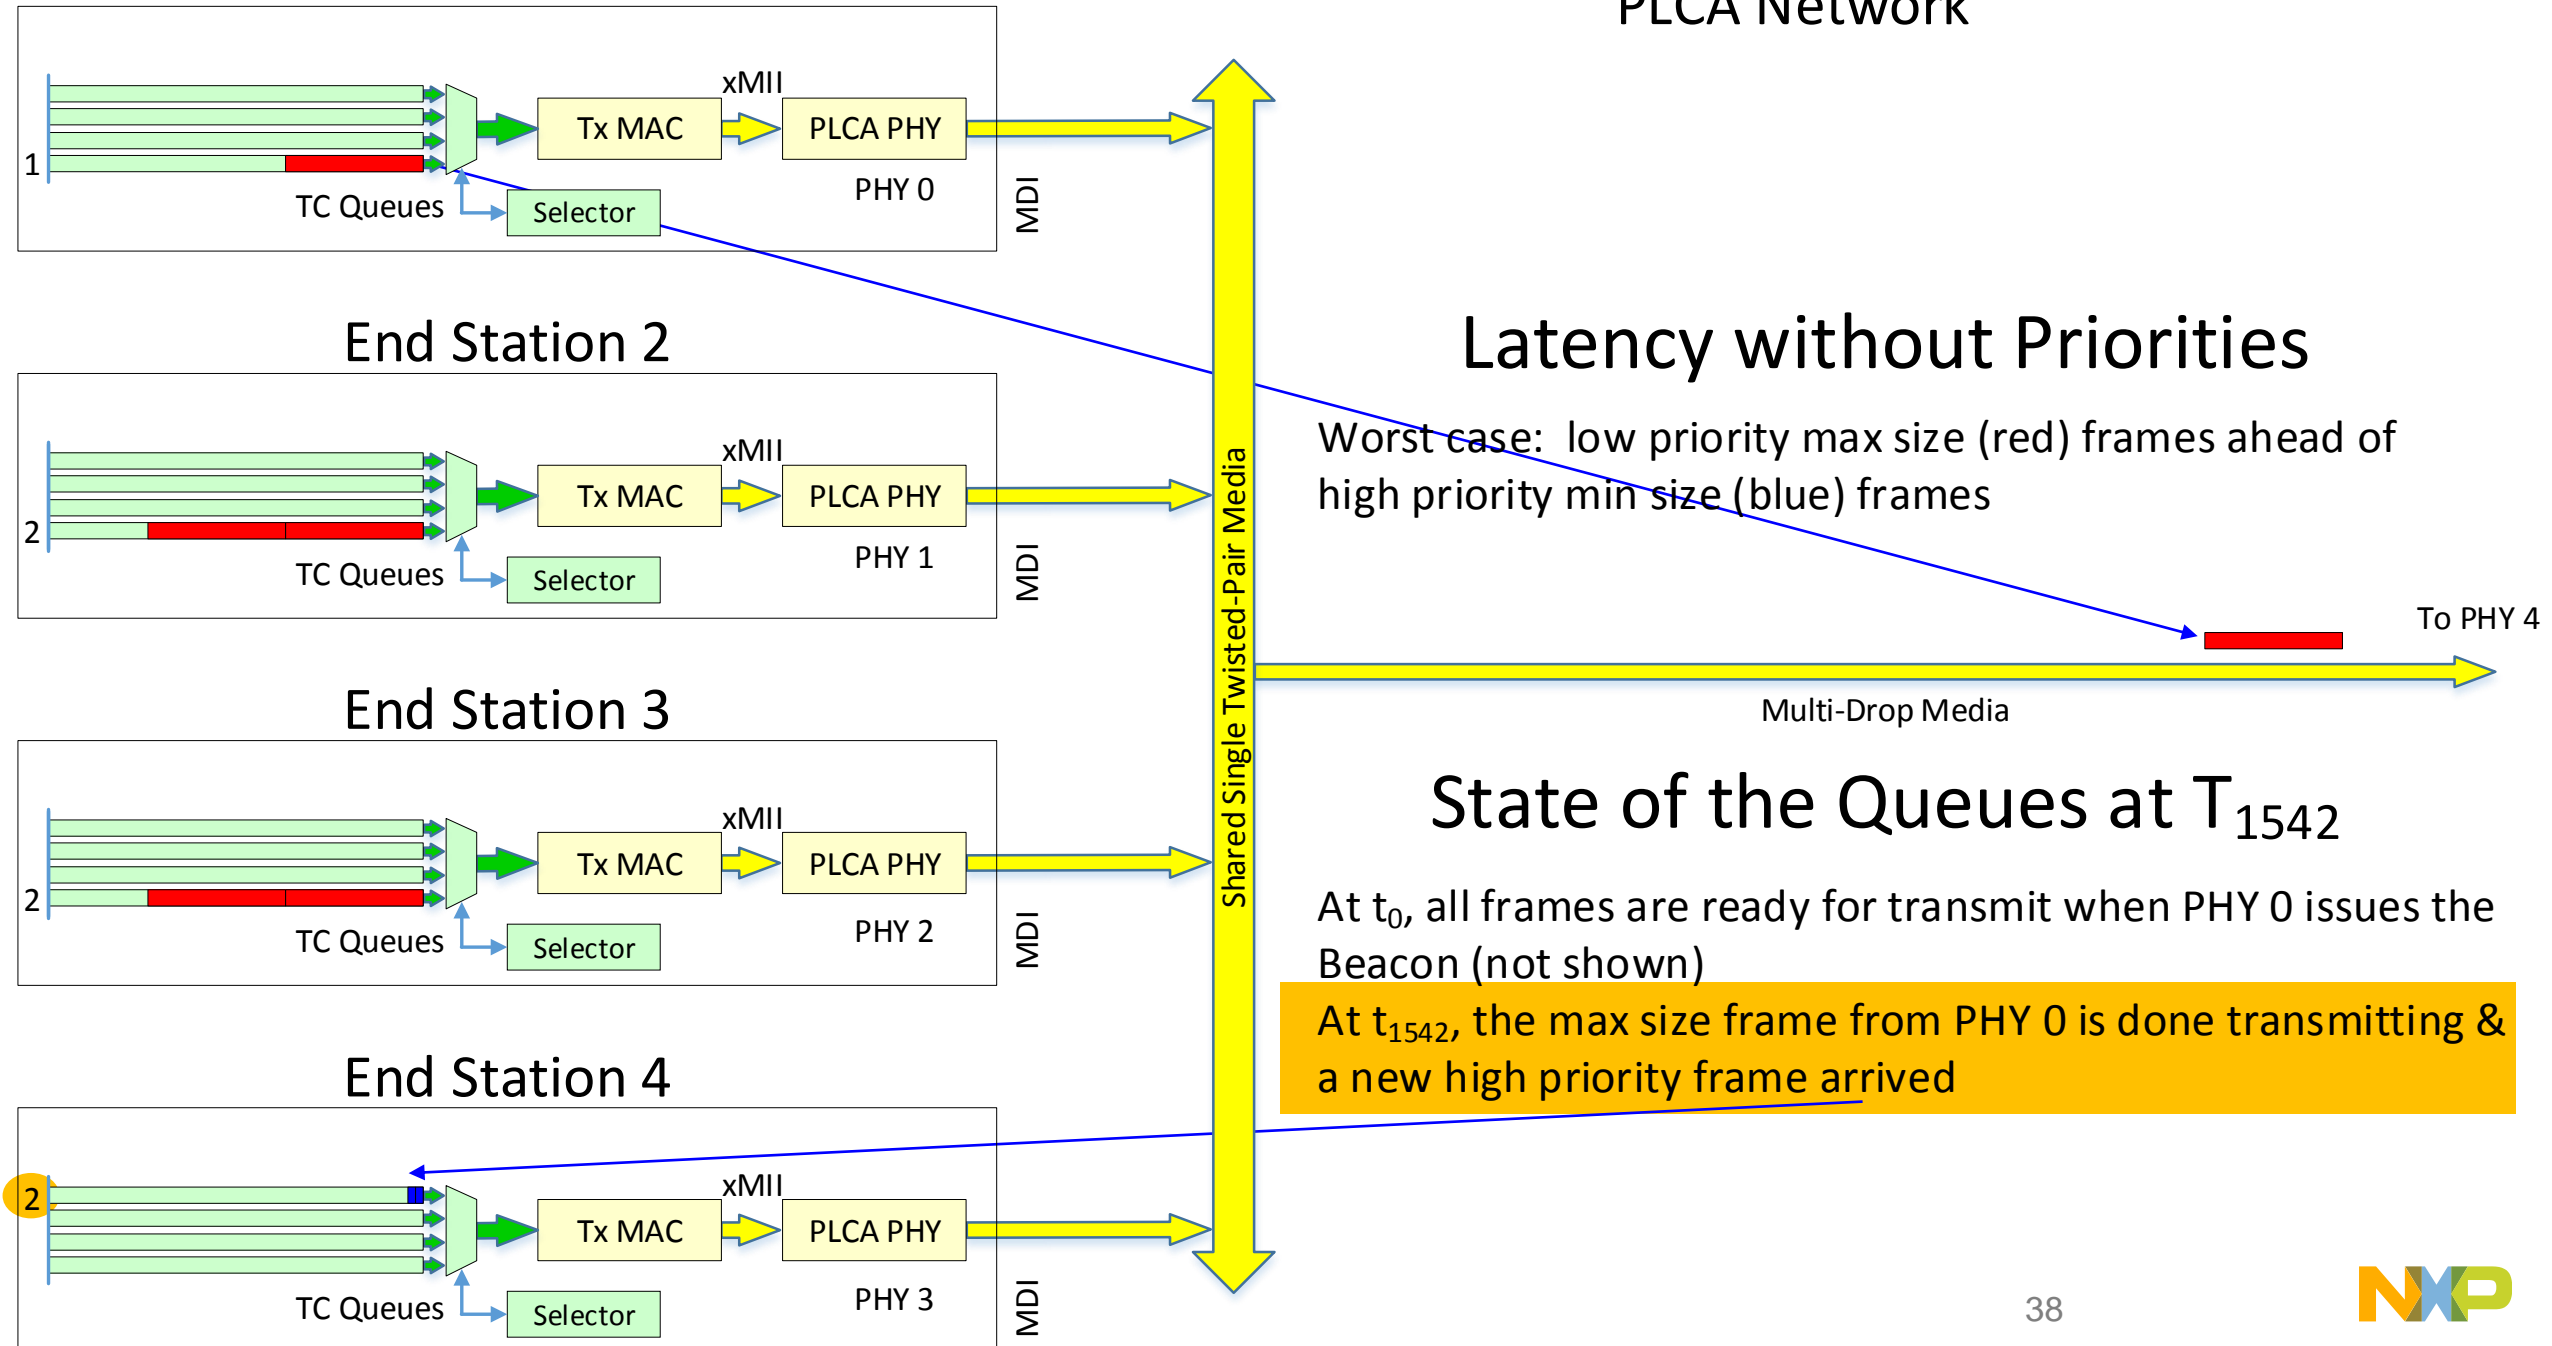
Latency without Priorities

Worst case: low priority max size (red) frames ahead of high priority min size (blue) frames

State of the Queues at T_{1542}

At t_0 , all frames are ready for transmit when PHY 0 issues the Beacon (not shown)

At t_{1542} , the max size frame from PHY 0 is done transmitting & a new high priority frame arrived



End Station 1

PLCA Network

Latency without Priorities

Worst case: low priority max size (red) frames ahead of high priority min size (blue) frames

End Station 2

End Station 3

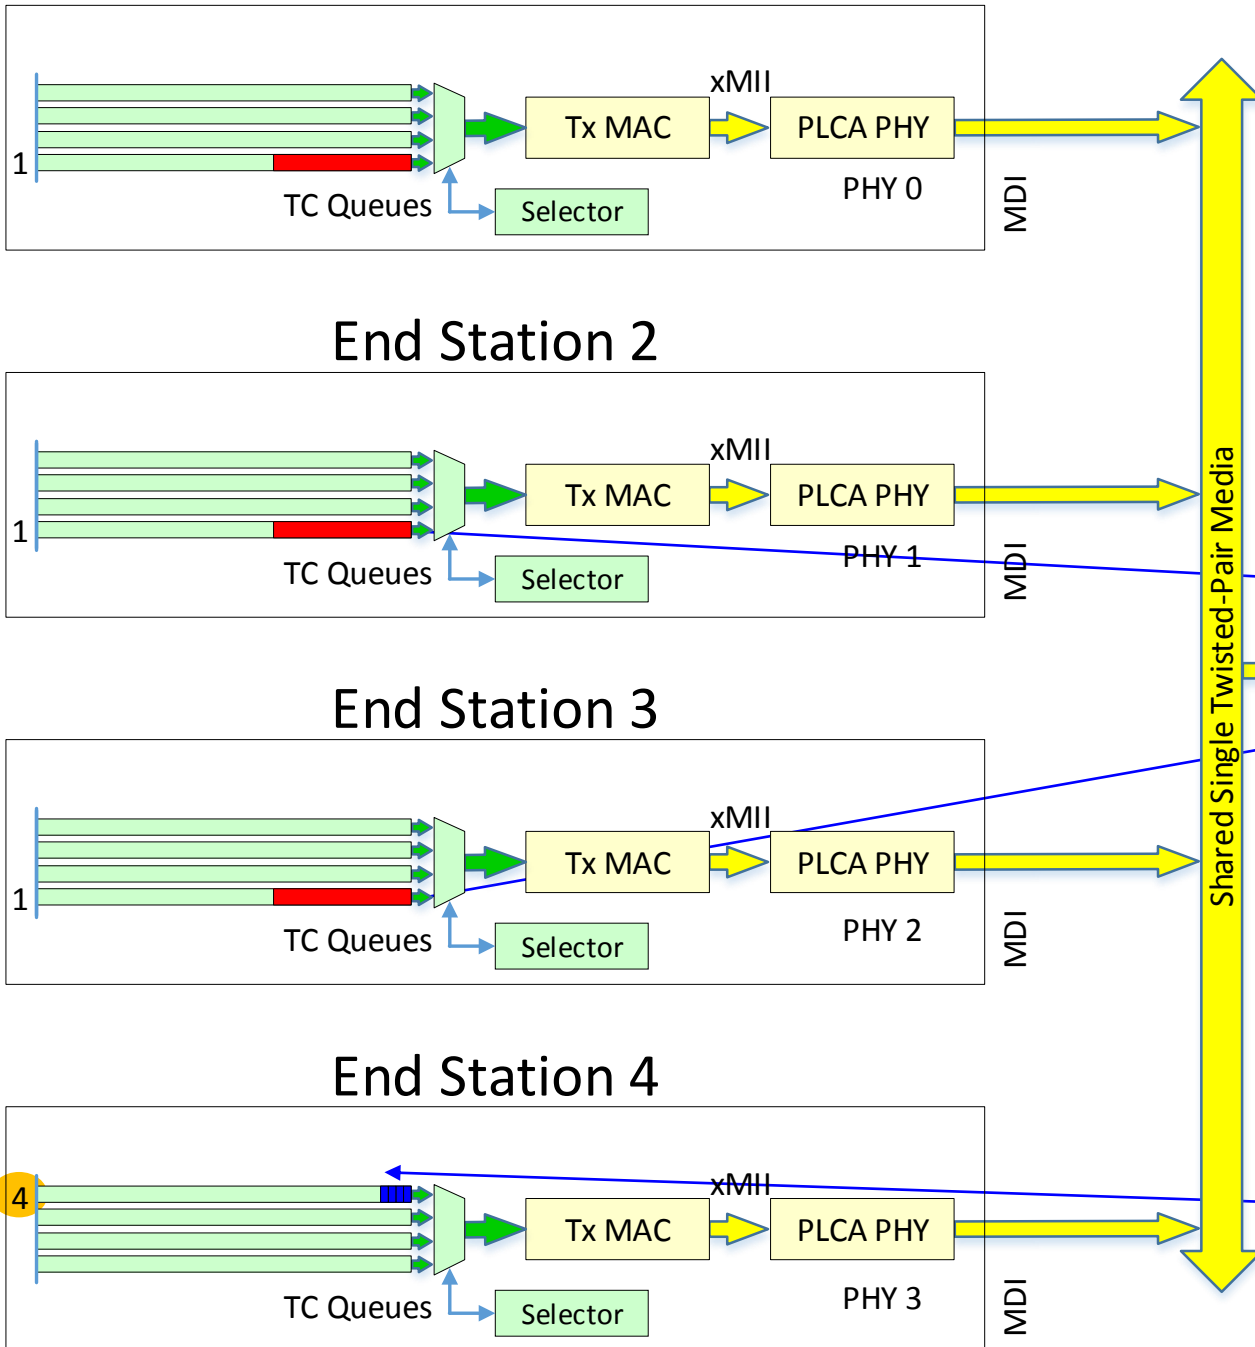
End Station 4

Shared Single Twisted-Pair Media


Multi-Drop Media

State of the Queues at T_{4626}


- At t_0 , all frames are ready for transmit when PHY 0 issues the Beacon (not shown)
- At t_{1542} , the max size frame from PHY 0 is done transmitting & a new high priority frame arrived
- At t_{4626} , the max size frame from PHY 1 & 2 are done & 2 new high priority frames arrived




End Station 1




End Station 2



End Station 3



End Station 4




PLCA Network

Latency without Priorities

Worst case: low priority max size (red) frames ahead of high priority min size (blue) frames


Shared Single Twisted-Pair Media




State of the Queues at T_{4690}

- At t_0 , all frames are ready for transmit when PHY 0 issues the Beacon (not shown)
- At t_{1542} , the max size frame from PHY 0 is done transmitting & a new high priority frame arrived
- At t_{4626} , the max size frame from PHY 1 & 2 are done & 2 new high priority frames arrived
- At t_{4690} , the min size frame from PHY 3 finally gets finished


End Station 1




End Station 2



End Station 3



End Station 4




PLCA Network

Latency without Priorities

Worst case: low priority max size (red) frames ahead of high priority min size (blue) frames

Shared Single Twisted-Pair Media



At the End of Beacon 2's Cycle

The pattern repeats for the next Beacon cycle as there are no Priorities between stations


-- End Station 4's high priority data grows to 6 frames!
Assuming the End Stations 1 to 3 don't have any data to send, End Station 4 is finally able to drain its high priority data

-- But this cannot be guaranteed as low priority frames are generally not Engineered to get out of the way!


Worst Case Latency Improvement Example – Part 2

- The exact same example is repeated, but this time the optional Priority Advertise is used
- A Strict Priority Selector is assumed to be used in the 802.1 portion in this example
 - Strict Priority was the 1st selector defined by 802.1 in Sep 1998
 - It is the required selector mechanism for the highest priority queues when using AVB (802.1BA-2011)
 - Strict Priority selects only the highest priority frames until there are no more frames of that priority, and then the next lowest priority is chosen, and so on, until there are no more frames in the queues, or a frame appears in a higher priority queue


End Station 1




End Station 2



End Station 3



End Station 4



PLCA Network

Latency **with** Priorities

Worst case: low priority max size (red) frames ahead of high priority min size (blue) frames

Shared Single Twisted-Pair Media


Multi-Drop Media

To PHY 4


State of the Queues at T_0

At t_0 , all frames are ready for transmit – i.e., the identical start as before – but this time PHY 0 issues the optional Advertise Beacon with an overhead of ~ 4 bytes/Station


End Station 1




End Station 2



End Station 3



End Station 4



PLCA Network

Latency with Priorities

Worst case: low priority max size (red) frames ahead of high priority min size (blue) frames

Shared Single Twisted-Pair Media

Multi-Drop Media


To PHY 4

State of the Queues at T_{84}


At t_0 , all frames are ready for transmit – i.e., the identical start as before – but this time PHY 0 issues the optional Advertise Beacon with an overhead of ~ 4 bytes/Station

At t_{20} , the normal Beacon (not shown) is issued by PHY 0
At t_{84} , the min size high priority frame from PHY 3 gets transmitted – as all the other stations are held back by 802.1 Strict Priority in this example


End Station 1




End Station 2



End Station 3



End Station 4



PLCA Network

Latency with Priorities

Worst case: low priority max size (red) frames ahead of high priority min size (blue) frames

Shared Single Twisted-Pair Media


Multi-Drop Media

To PHY 4


State of the Queues at Cycle 2_{end}

Cycle 1 ended with only PHY 3 transmitting
Cycle 2 starts with PHY 0 issuing the optional Advertise Beacon with an overhead of ~4 bytes/Station
-- PHY 3 has no data to send, so PHY 0 to 2 get to send
-- During this time PHY 3 accumulates 3 more frames


End Station 1




End Station 2



End Station 3




End Station 4



PLCA Network

Latency with Priorities

Worst case: low priority max size (red) frames ahead of high priority min size (blue) frames



State of the Queues at Cycle 6_{end}

Cycle 3, 4 & 5 drain the 3 pending high priority frames from PHY 3 – like AVB's credit based shaper bandwidth 'catch-up'

Cycle 6 is a repeat of Cycle 2, allowing the low priority maximum size frames out while PHY 3 accumulates 3 more high priority frames

-- and the pattern could repeat like this

Latency Improvement of the PLCA Priority Advertise Mechanism

- Without the PLCA Priority Advertise (4 PHY, 4-to-1 example):
 - The last byte of the high priority AVB frame was received after: 4690 byte times
 - If the interfering frames continue for subsequent Beacon cycles, this 4690 byte times of delay can repeat for each AVB frame, never allowing, or at least delaying the needed TSN bandwidth 'catch-up' supported by the AVB 802.1Qav Credit Based Shaper
 - Bandwidth guarantees, bounded latencies & buffer requirements can't be determined!
- With the PLCA Priority Advertise (4 PHY, 4-to-1 example):
 - The last byte of the high priority AVB frame was received after: 84 byte times
 - An improvement of 4606 byte times! (subject to the Advertise codes chosen)
 - Just as important this approach supports the very important TSN bandwidth 'catch-up' required by the AVB 802.1Qav Credit Based Shaper
 - This approach supports guaranteed bandwidth allocations with defined bounded latencies, two critical aspects of TSN!



Backup 2

Emission Sequence Control

This analysis also shows that the 802.1Qcr Asynchronous Traffic Shaper (aka, the Urgency Based Scheduler in TSN) cannot be solved without the proposed per Beacon dynamic Priority information – and shows the benefit of 4-bit Priority values

802.1 Priority Mapping Ideas for the PLCA Priority Advertise

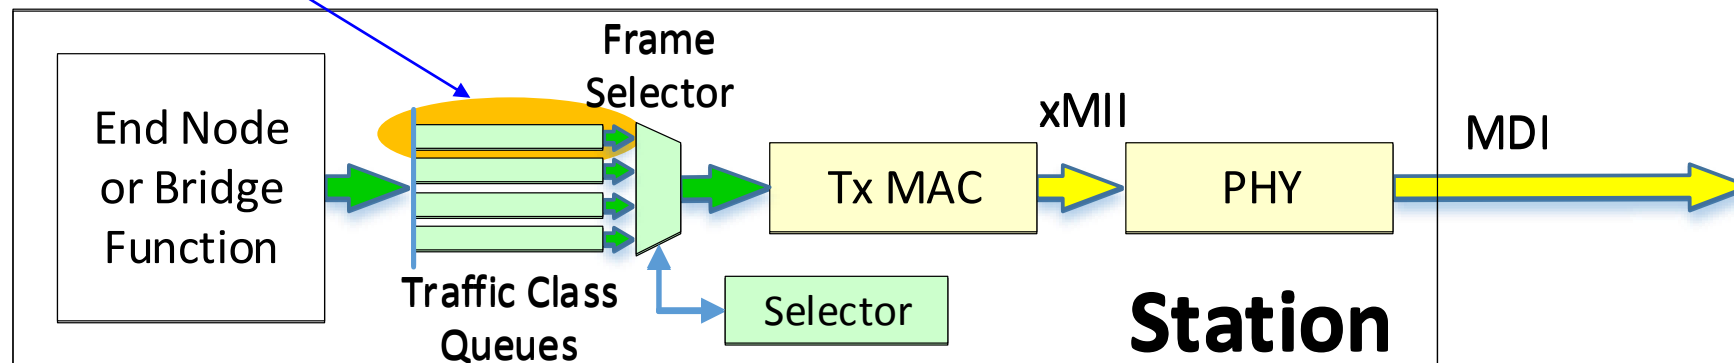
- This section's intention is to show the power and flexibility of the separation between the Priority Advertise function being done in 802.3, and the Priority determination and processing being done in 802.1
- Its showing of a 2nd, very different use case, that efficiently uses & needs the Priority Advertise function, helps show that the architecture and location of Priority Advertise is correct
- This is NOT a proposal for any new 802.1 standards at this time, as there is no need for the PLCA Priority Advertise to have any dependency on 802.1 if the EEE model of 'connecting' to a 'Priority Client' is used in 802.3cg
 - This worked for Energy Efficient Ethernet (EEE) where the Low Power Idle Client 'connects' to 802.1's queue states
 - It can work here too

802.1's Various 'Different' Priorities – Each has a Name & Purpose

- PCP – Priority Code Point: 802.1's name for the 3-bit priority field in 802.3ac "VLAN TAG" frames (since Sep 1998) – an "outside the box" value
 - Indicates the 'priority' (differentiated service class) of the frame externally
- Traffic Class: 802.1's name for a differentiated service queue, where each service needs its own queue in the egress ports (since Sep 1998)
 - Multiple streams of the same TC (Traffic Class) can reside in the same queue
 - But any given queue only contains one kind of differentiated service
- IPV – Internal Priority Value: As its name implies, it is an "inside the box" value
 - It is a recent addition to 802.1Q that supports more than eight Traffic Class queues
- Service Access Priority: 802.1's name for a 3-bit access priority that is used for certain media access methods (like 802.3cg!)

802.1's Mapping of PCP to Traffic Class (Simplified)


- 802.1Q supports devices (end stations and bridges) that contain less than the full compliment of eight Traffic Class queues to support the frame's 3-bit PCP field
 - This is accomplished with an internal 'down-sized' mapping of PCP to IPV
- With the advent of AVB/TSN a PCP value no longer has a linear mapping to Traffic Class queues (it never really did, but now it's more obvious)
 - This allows Class A AVB streams (with a PCP of 0x3) to be mapped to the device's highest Traffic Class queue so the AVB guarantees can be met
 - This would be TC 3 in 4 queue devices (as shown below) or TC 7 in 8 queue devices




Green = IEEE 802.1

Yellow = IEEE 802.3


End Station 1




End Station 2



End Station 3



End Station 4



Devices with Different Queue #'s


All sizes are valid in 802.1. Assume the Blue frames all have the same PCP value of 0x3 they could end up in different TC Queue numbers.

Shared Single Twisted-Pair Media


Multi-Drop Media

To PHY 4


End Station 1




End Station 2



End Station 3



End Station 4



Devices with Different Queue #'s

All sizes are valid in 802.1. Assume the Blue frames all have the save PCP value of 0x3 they could end up in different TC Queue numbers.

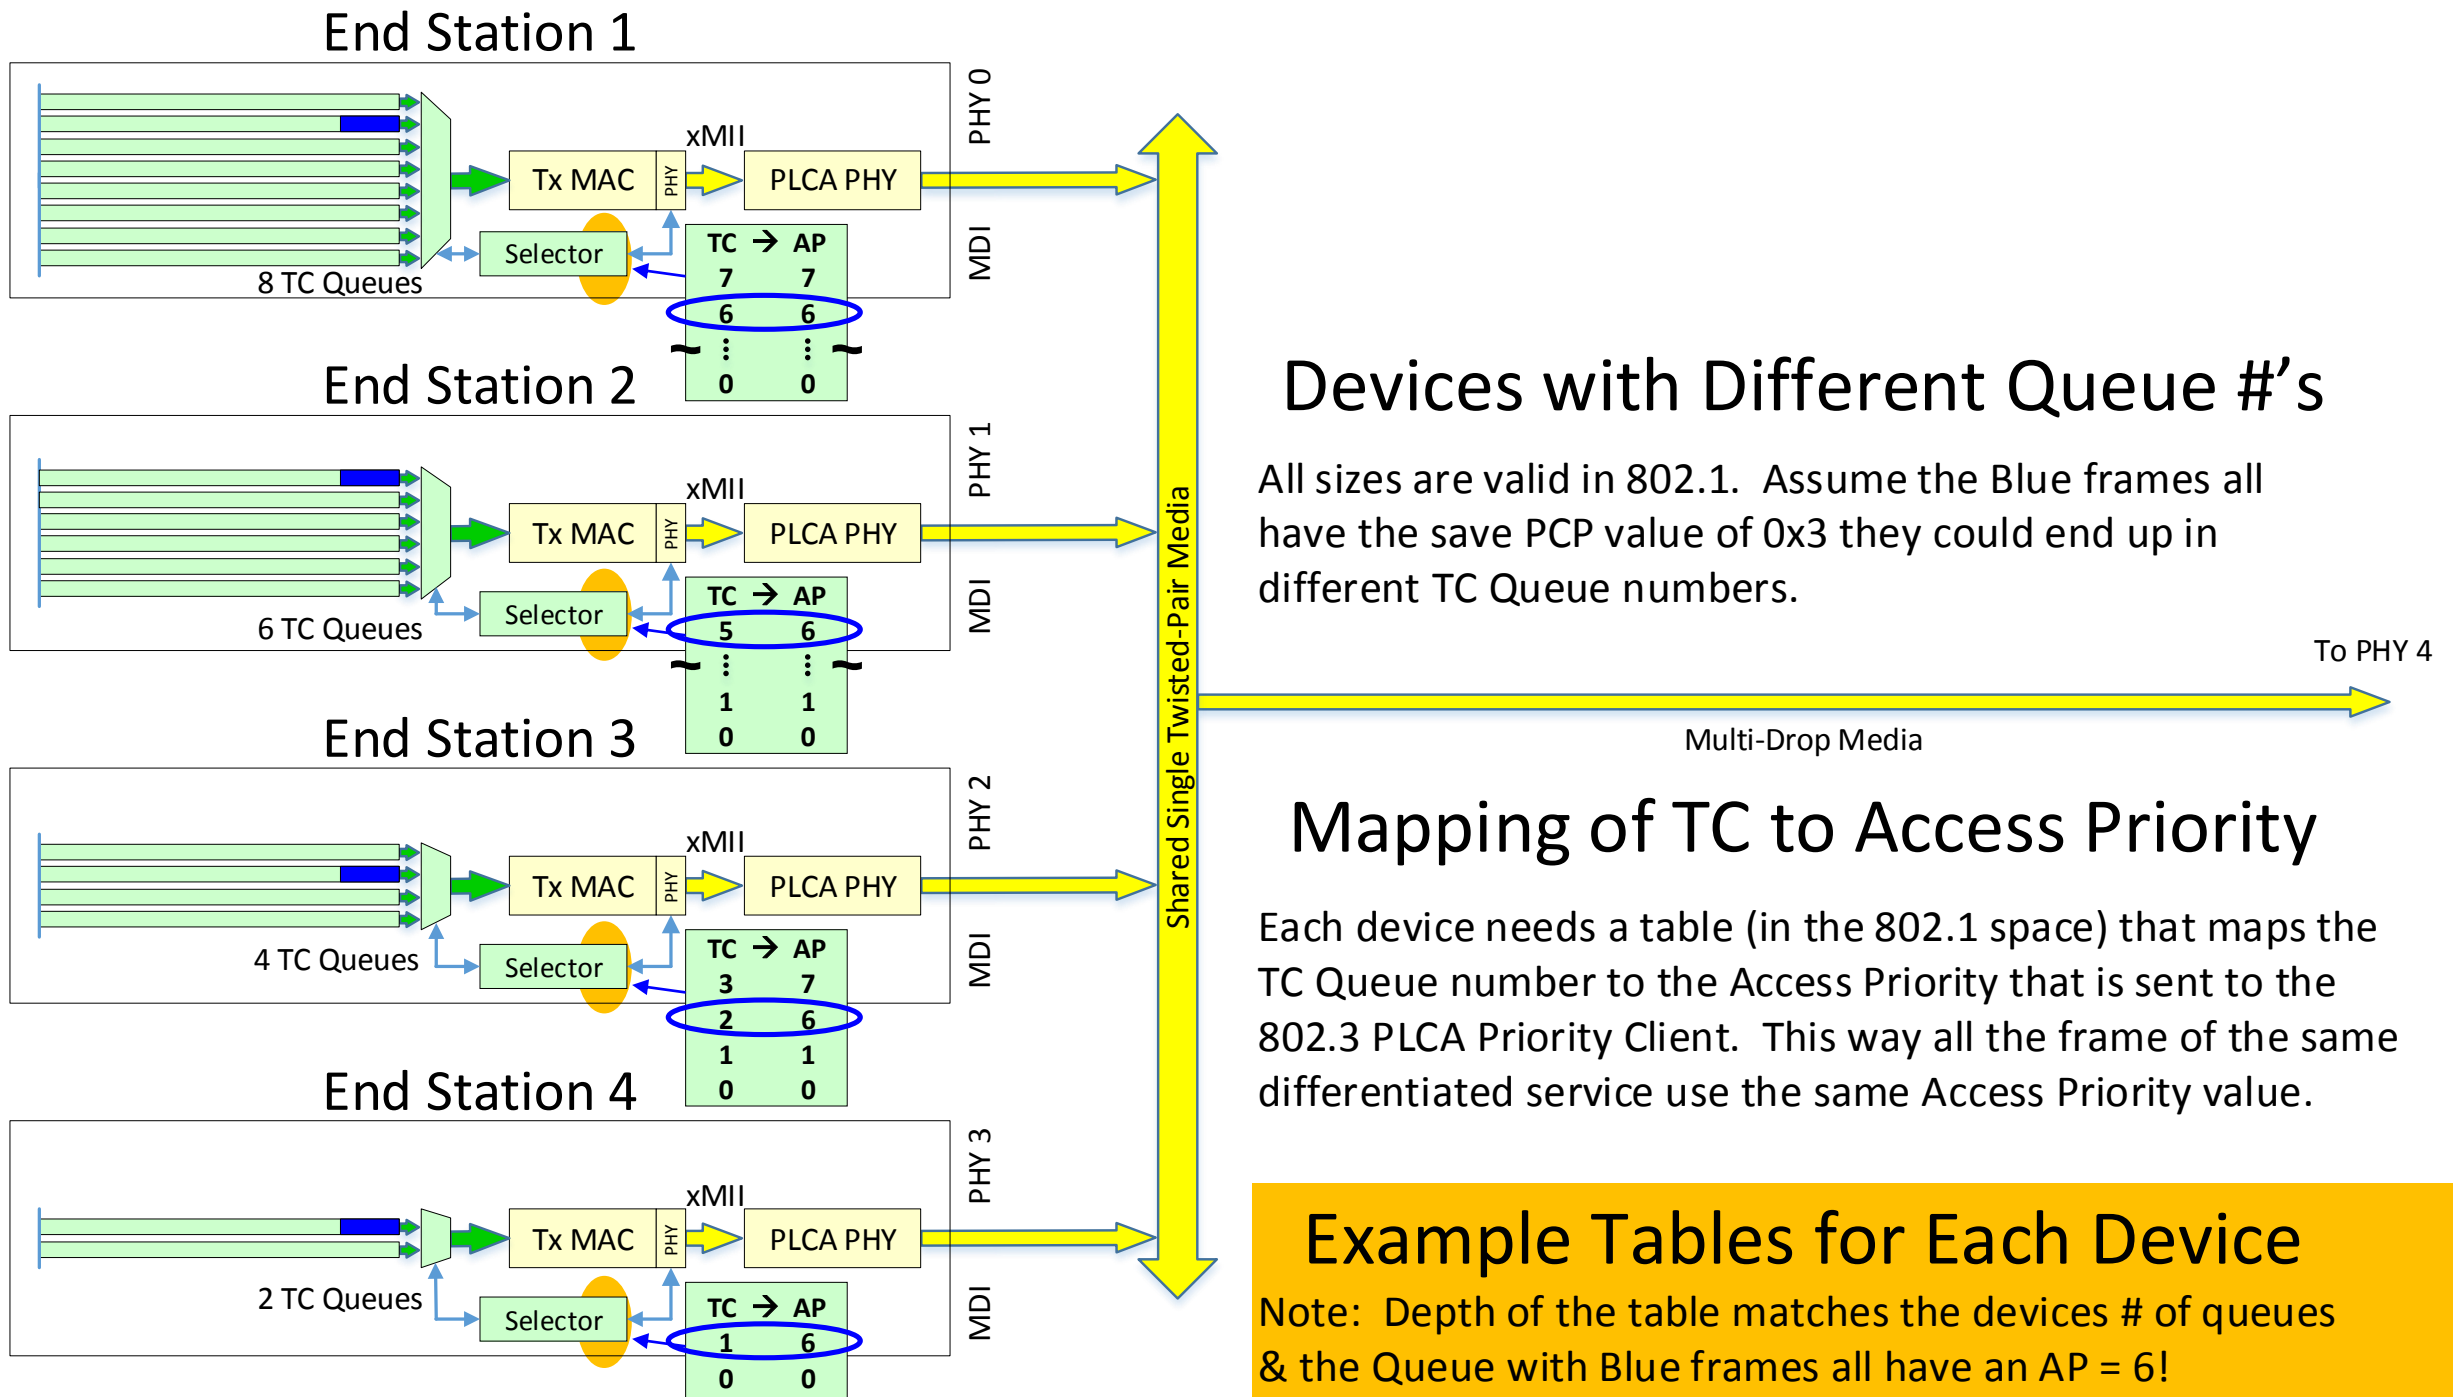
Shared Single Twisted-Pair Media

Multi-Drop Media

To PHY 4

Mapping of TC to Access Priority


Each device needs a table (in the 802.1 space) that maps the TC Queue number to the Access Priority that is sent to the 802.3 PLCA Priority Client. This way all the frame of the same differentiated service use the same Access Priority value.




Determination of the Access Priority Size (its # of bits)

- Based on this information, what is the data width needed for the Priority Value between the Priority Client and the Reconciliation Sublayer?
 - And subsequently down and up the MII and on the media?
- The obvious answer is 3 bits to match the PCP & Service Access Priority widths
- The non-obvious answer is 4 or more bits – called Extended Priority Advertise
 - 4 bits won't cost any more than 3 since the MII is nibble wide
 - It takes the same amount of time (clock cycles) to transfer the 3 bits as it does 4 bits
- Why more than 3 bits? – to support Emission Sequence Control, ...
 - 1) 802.1's now has a IPV that is more than 3 bits to support more than 8 queues, and
 - 2) 802.1 devices with less than 8 queues can use the larger Priority Value to mimic CAN like message priorities and to mimic 802.1Qcr's Asynchronous Traffic Shaping!


End Station 1




End Station 2



End Station 3



End Station 4



Emission Sequence Control


The setup: High critical frames (blue) are ready to be transmitted by each station. Lower critical frame (green) are also ready to be transmitted by each station. The blue frames need to egress ahead of the green but in a different station order (as noted by a letter in the frame). The lowest priority frames (red) are transmitted last.

Shared Single Twisted-Pair Media


Multi-Drop Media

To PHY 4


End Station 1




End Station 2



End Station 3



End Station 4



Emission Sequence Control

The setup: High critical frames (blue) are ready to be transmitted by each station. Lower critical frame (green) are also ready to be transmitted by each station. The blue frames need to egress ahead of the green but in a different station order (as noted by a letter in the frame). The lowest priority frames (red) are transmitted last.

Shared Single Twisted-Pair Media


Multi-Drop Media

To PHY 4


Use of Extended Priority Advertise

After the 1st Advertise Beacon, blue A from Station 4 is sent.


End Station 1




End Station 2



End Station 3



End Station 4



Emission Sequence Control

The setup: High critical frames (blue) are ready to be transmitted by each station. Lower critical frame (green) are also ready to be transmitted by each station. The blue frames need to egress ahead of the green but in a different station order (as noted by a letter in the frame). The lowest priority frames (red) are transmitted last.

Shared Single Twisted-Pair Media


Multi-Drop Media

To PHY 4


Use of Extended Priority Advertise

After the 1st Advertise Beacon, blue A from Station 4 is sent. Beacon 2 egresses blue B from 2 & blue C from 3.


End Station 1




End Station 2



End Station 3



End Station 4



Emission Sequence Control

The setup: High critical frames (blue) are ready to be transmitted by each station. Lower critical frame (green) are also ready to be transmitted by each station. The blue frames need to egress ahead of the green but in a different station order (as noted by a letter in the frame). The lowest priority frames (red) are transmitted last.

Shared Single Twisted-Pair Media

Multi-Drop Media

To PHY 4


D C B A

Use of Extended Priority Advertise


After the 1st Advertise Beacon, blue A from Station 4 is sent. Beacon 2 egresses blue B from 2 & blue C from 3.

Beacon 3 lets blue D out from Station 1 – last of the blues.


End Station 1




End Station 2



End Station 3



End Station 4



Emission Sequence Control

The setup: High critical frames (blue) are ready to be transmitted by each station. Lower critical frame (green) are also ready to be transmitted by each station. The blue frames need to egress ahead of the green but in a different station order (as noted by a letter in the frame). The lowest priority frames (red) are transmitted last.

Shared Single Twisted-Pair Media


Multi-Drop Media

To PHY 4


Use of Extended Priority Advertise

After the 1st Advertise Beacon, blue A from Station 4 is sent. Beacon 2 egresses blue B from 2 & blue C from 3. Beacon 3 lets blue D out from Station 1 – last of the blues. Beacon 4 egresses green A, B & C from 1, 2 & 3, respectively.


End Station 1




End Station 2



End Station 3



End Station 4



Emission Sequence Control

The setup: High critical frames (blue) are ready to be transmitted by each station. Lower critical frame (green) are also ready to be transmitted by each station. The blue frames need to egress ahead of the green but in a different station order (as noted by a letter in the frame). The lowest priority frames (red) are transmitted last.

Shared Single Twisted-Pair Media


Multi-Drop Media

To PHY 4


Use of Extended Priority Advertise

After the 1st Advertise Beacon, blue A from Station 4 is sent. Beacon 2 egresses blue B from 2 & blue C from 3. Beacon 3 lets blue D out from Station 1 – last of the blues. Beacon 4 egresses green A, B & C from 1, 2 & 3, respectively. Beacon 5 lets green D out from Station 2 – last of the greens.


End Station 1




End Station 2



End Station 3

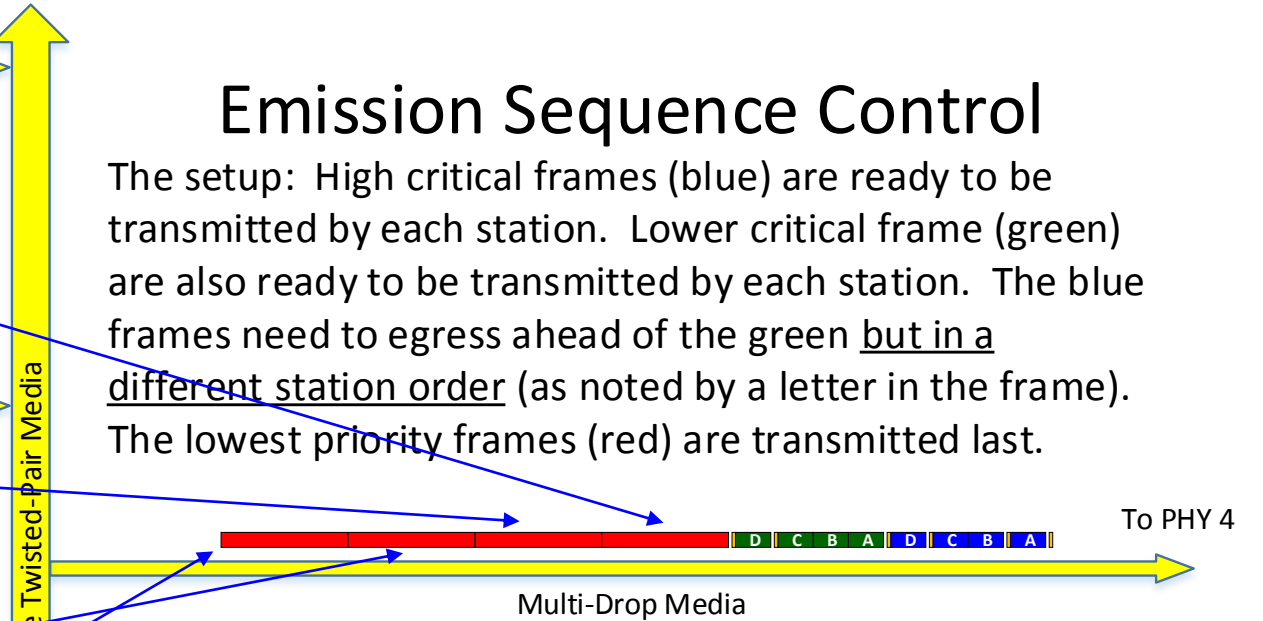


End Station 4



Emission Sequence Control

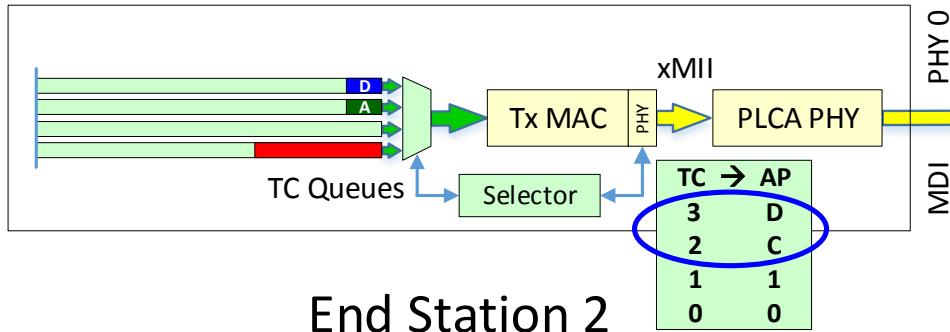
The setup: High critical frames (blue) are ready to be transmitted by each station. Lower critical frame (green) are also ready to be transmitted by each station. The blue frames need to egress ahead of the green but in a different station order (as noted by a letter in the frame). The lowest priority frames (red) are transmitted last.



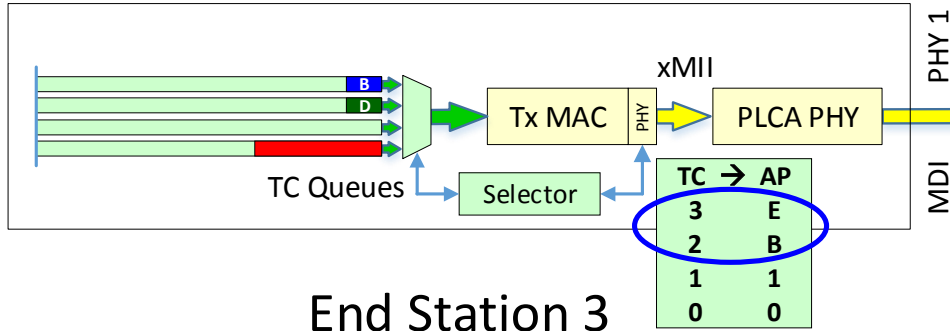
Use of Extended Priority Advertise

After the 1st Advertise Beacon, blue A from Station 4 is sent. Beacon 2 egresses blue B from 2 & blue C from 3. Beacon 3 lets blue D out from Station 1 – last of the blues. Beacon 4 egresses green A, B & C from 1, 2 & 3, respectively. Beacon 5 lets green D out from Station 2 – last of the greens. Beacon 6 egresses the red frame out in Station # order.

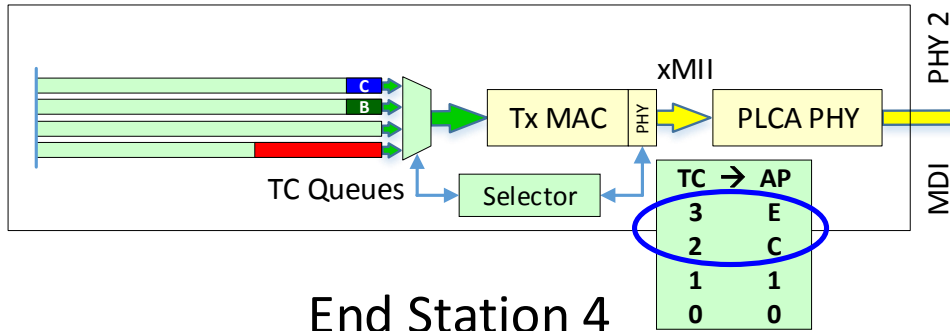
End Station 1



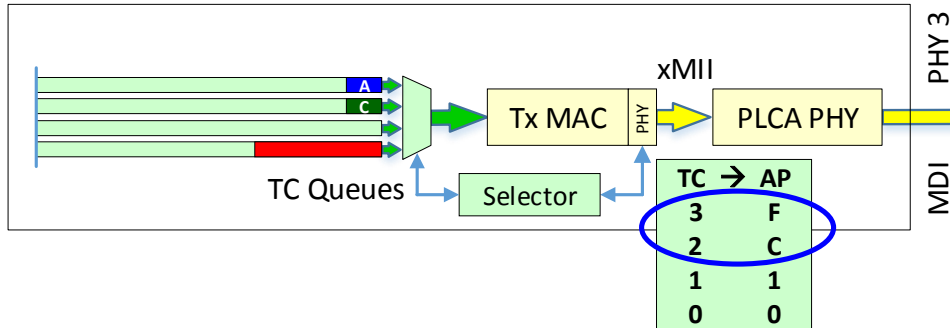
End Station 2



End Station 3




End Station 4



Shared Single Twisted-Pair Media

How This Works

How does Extended Priority Advertise support different emission sequences per Traffic Class? Is all done in the TC to Access Priority mapping table in each device – an N x 3 table that is already needed (where N is the # Queues in the device's port). This example shows that a 4-bit wide Priority Advertise on the media is useful even if the End Stations only have 4 TC Queues!



Each device's Traffic Class (TC) to Access Priority (AP) mapping table values are shown that produce the emission sequence in this example.

Note: Sequence segments that follow the PHY's number assignments don't require a distinct AP value. See E, C & 0

Since the number of 'critical' streams on any PLCA network are expected to be low & since linear sequence segments can use the same value, a 4-bit wide Priority Advertise value on the media is seen to be enough. An 802.1 recommended practice could define AP's 0 to 7 for the 8 (typical TC's) and use AP's 8 to F for Emission Sequencing.

Recommendations for 802.3cg

- Support a 4-bit wide Priority Advertise on the media to support 16 priorities
 - This adds zero overhead from the minimum requirement of 3-bits as the interfaces (the MII & the symbols on the media) already support nibble (4-bit wide) data
- Two, very different, use cases have been shown needing Priority Advertise
 - Both of these use cases came from use cases that 802.1 TSN has been asked to support and where standards are done or in process
- The support of multiple use cases using the same Priority Advertise mechanism in very different ways, shows this is the correct location for this mechanism
 - This is no surprise since Priority Advertise's goal was to mimic what 802.1 already did
- 802.1 only solutions to solve these use cases (without using Priority Advertise) have been developed – but they are much more complex, do not get as low of a latency, and waste a lot of link bandwidth as many consecutive idle Beacon cycles are required
 - So much bandwidth is wasted, Emission Sequence Control is no longer viable

