Preamble

The subsequent slides (not including this slide) contain **draft** material proposed for inclusion into a planned 802 tutorial on CTF (see https://mentor.ieee.org/802.1/dcn/21/1-21-0015-04-ICne-ctf-study-item-planning-proposal.pdf):

- **At the time this draft slide set is published, an 802 tutorial has not been approved!**
- However, the contents of the following slides are designed to show the final content, including indications for such a tutorial, as it would look like if such a tutorial would be approved.

The current version of this slide set contains the proposed introduction to the topic, intended to be followed by existing (and potential upcoming) use-case presentations, and subsequent material outlining one potential integration of CTF into IEEE 802.1.

The existing use-case presentations are the following ones:

- Industrial Automation
  https://mentor.ieee.org/802.1/dcn/21/1-21-0018-00-ICne-ctf-industrial-use-case.pdf
- Data Center Networks
  https://mentor.ieee.org/802.1/dcn/21/1-21-0019-01-ICne-ctf-for-dcn.pdf

Tutorial: Cut-Through Forwarding (CTF) in Bridges and Bridged Networks

Johannes Specht, Jordon Woods, Paul Congdon/Lily Lv, TBD

Abstract

Cut-Through Forwarding (CTF) is a known method to improve the delay performance in Bridged Networks. In contrast to the store and forward operation of standardized switched Ethernet, CTF allows frame transmission in Bridges before reception is completed. Although not standardized in IEEE 802, CTF is already implemented in existing Bridge implementations. It is therefore technically feasible, but different implementations face interoperability problems that can be resolved by standardizing CTF in IEEE 802.

This tutorial introduces CTF on a technical level, explains application areas, markets and use-cases for CTF, and describes one possible integration of CTF into switched Ethernet.

Disclaimer

This presentation should be considered as the personal views of the presenters not as a formal position, explanation, or interpretation of IEEE.

Per IEEE-SA Standards Board Bylaws, August 2020:

At lectures, symposia, seminars, or educational courses, an individual presenting information on IEEE standards shall make it clear that his or her views should be considered the personal views of that individual rather than the formal position of IEEE.

Table of Contents

Introduction

Johannes Specht

IEEE 802.1 TSN Context, Basic CTF Operation Guaranteed Latency, CTF Performance, Reasons for standardizing CTF

TSN Context

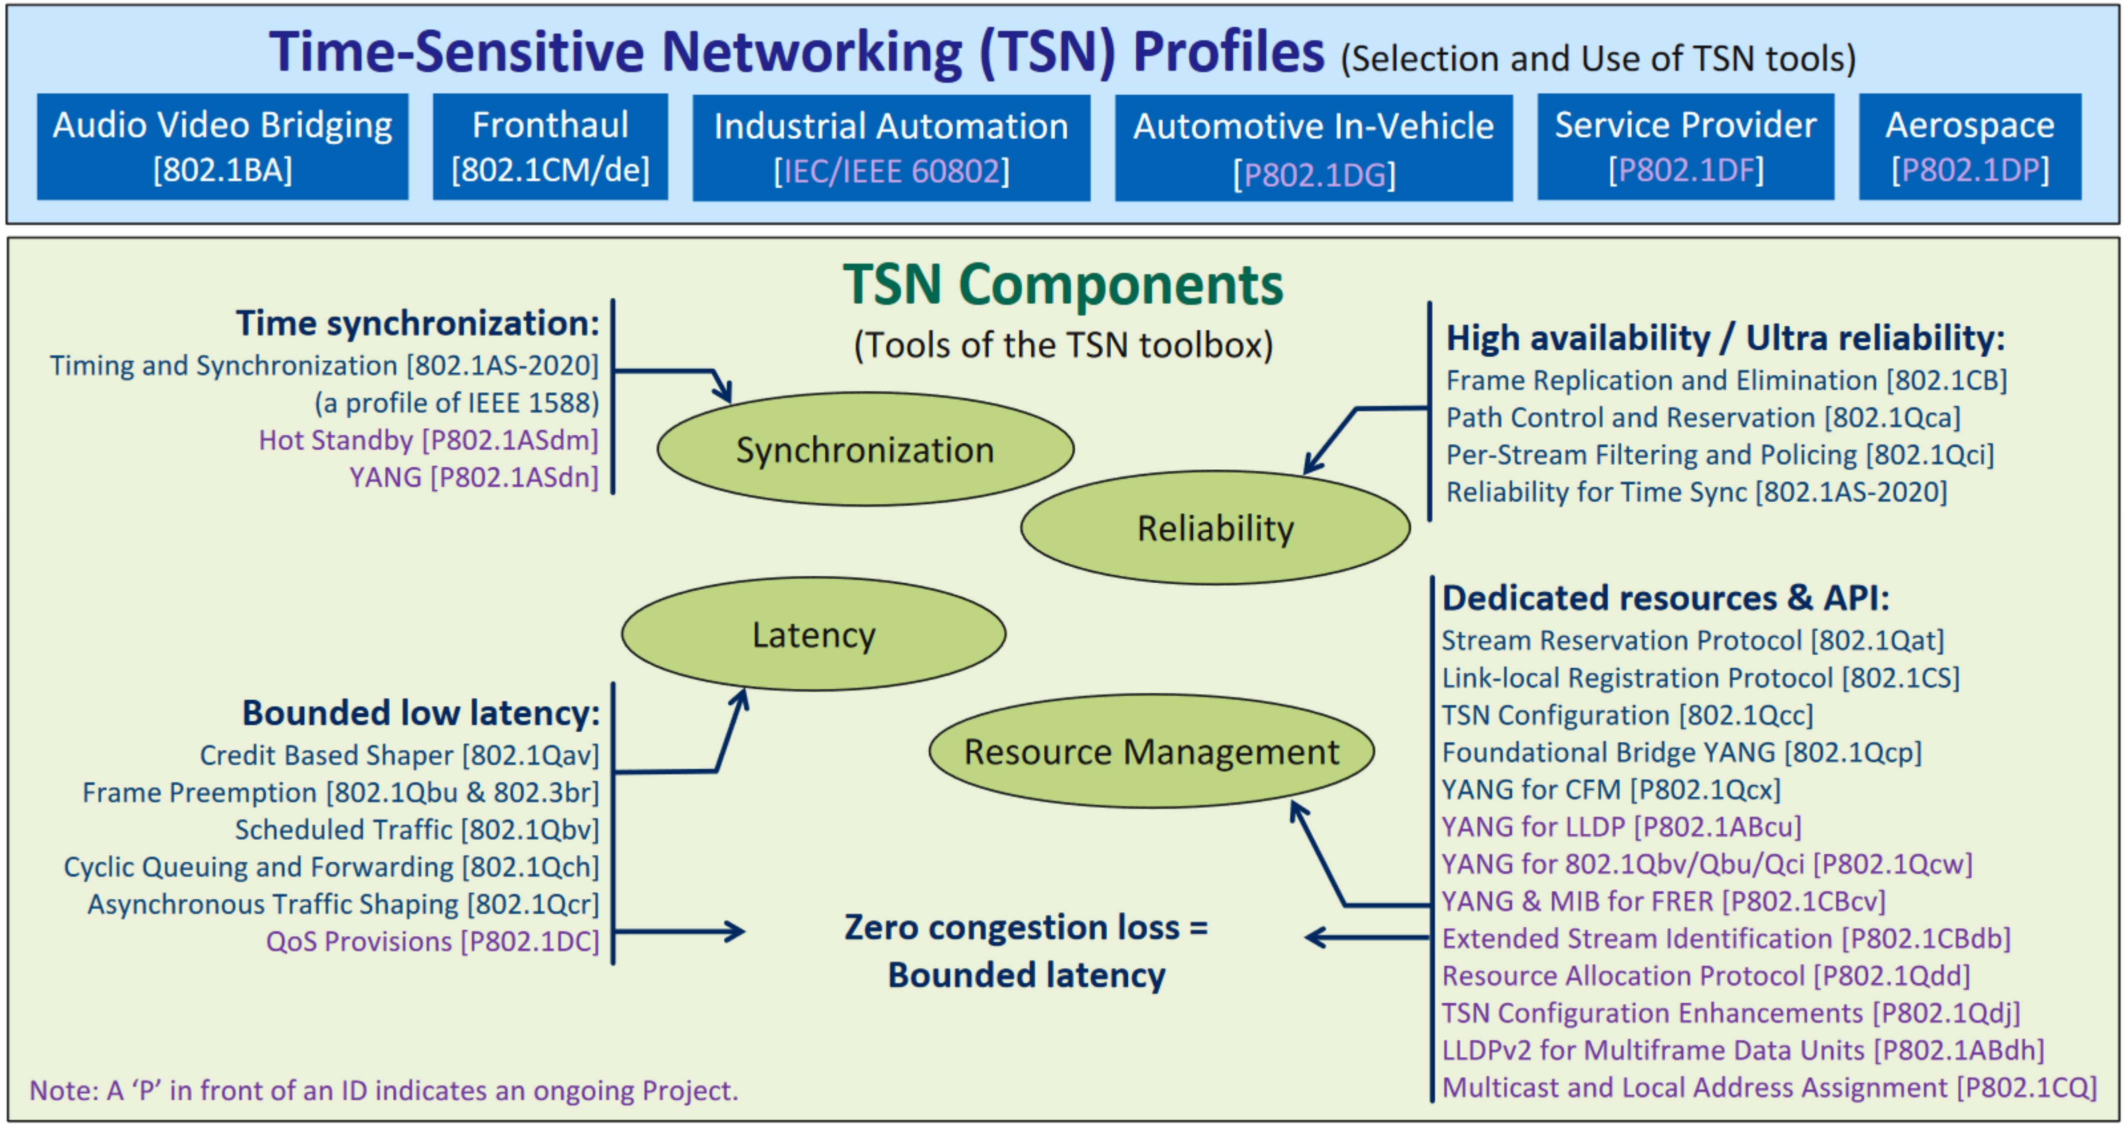

Source: https://www.ieee802.org/1/files/public/docs2021/admin-tsn-summary-0221-v01.pdf

Traditional and Deterministic Services

- Traditional Service
  - Curves have long tail
  - Average latency is good
  - Lowering the latency means losing packets (or overprovisioning)
- Deterministic Service
  - Packet loss is at most due to equipment failure (zero congestion loss)
  - Bounded latency, no tails
  - The right packet at the right time

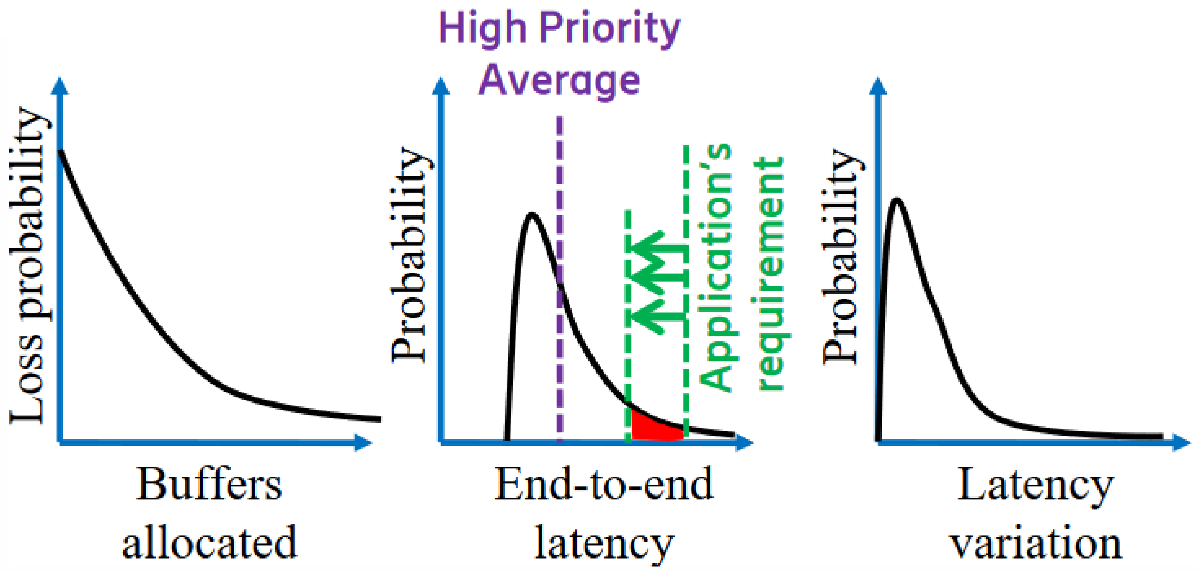

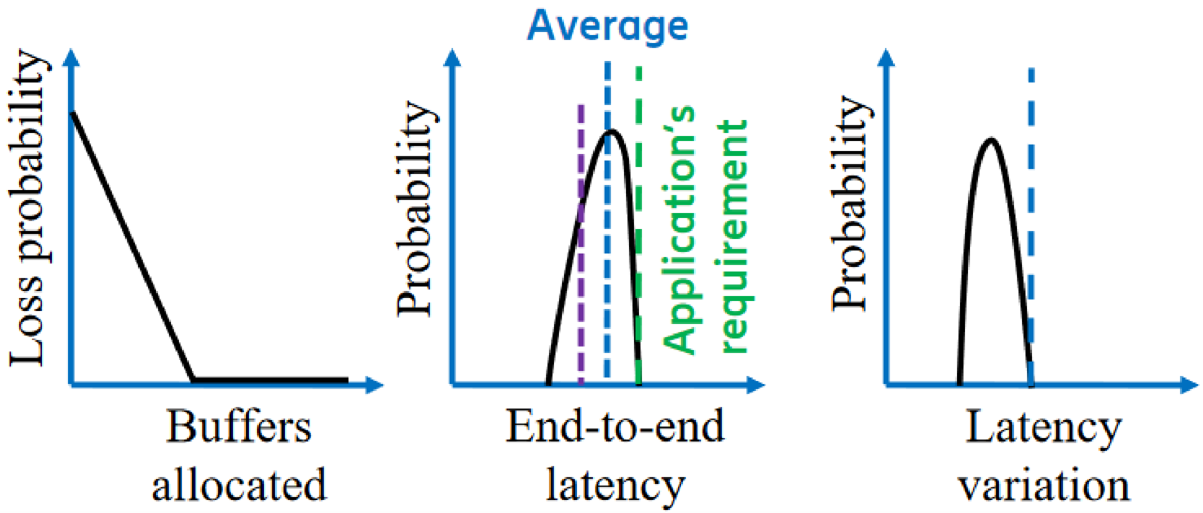

Source: https://www.ieee802.org/1/files/public/docs2018/detnet-tsn-farkas-tsn-basic-concepts-1118-v01.pdf

CTF in the TSN Context

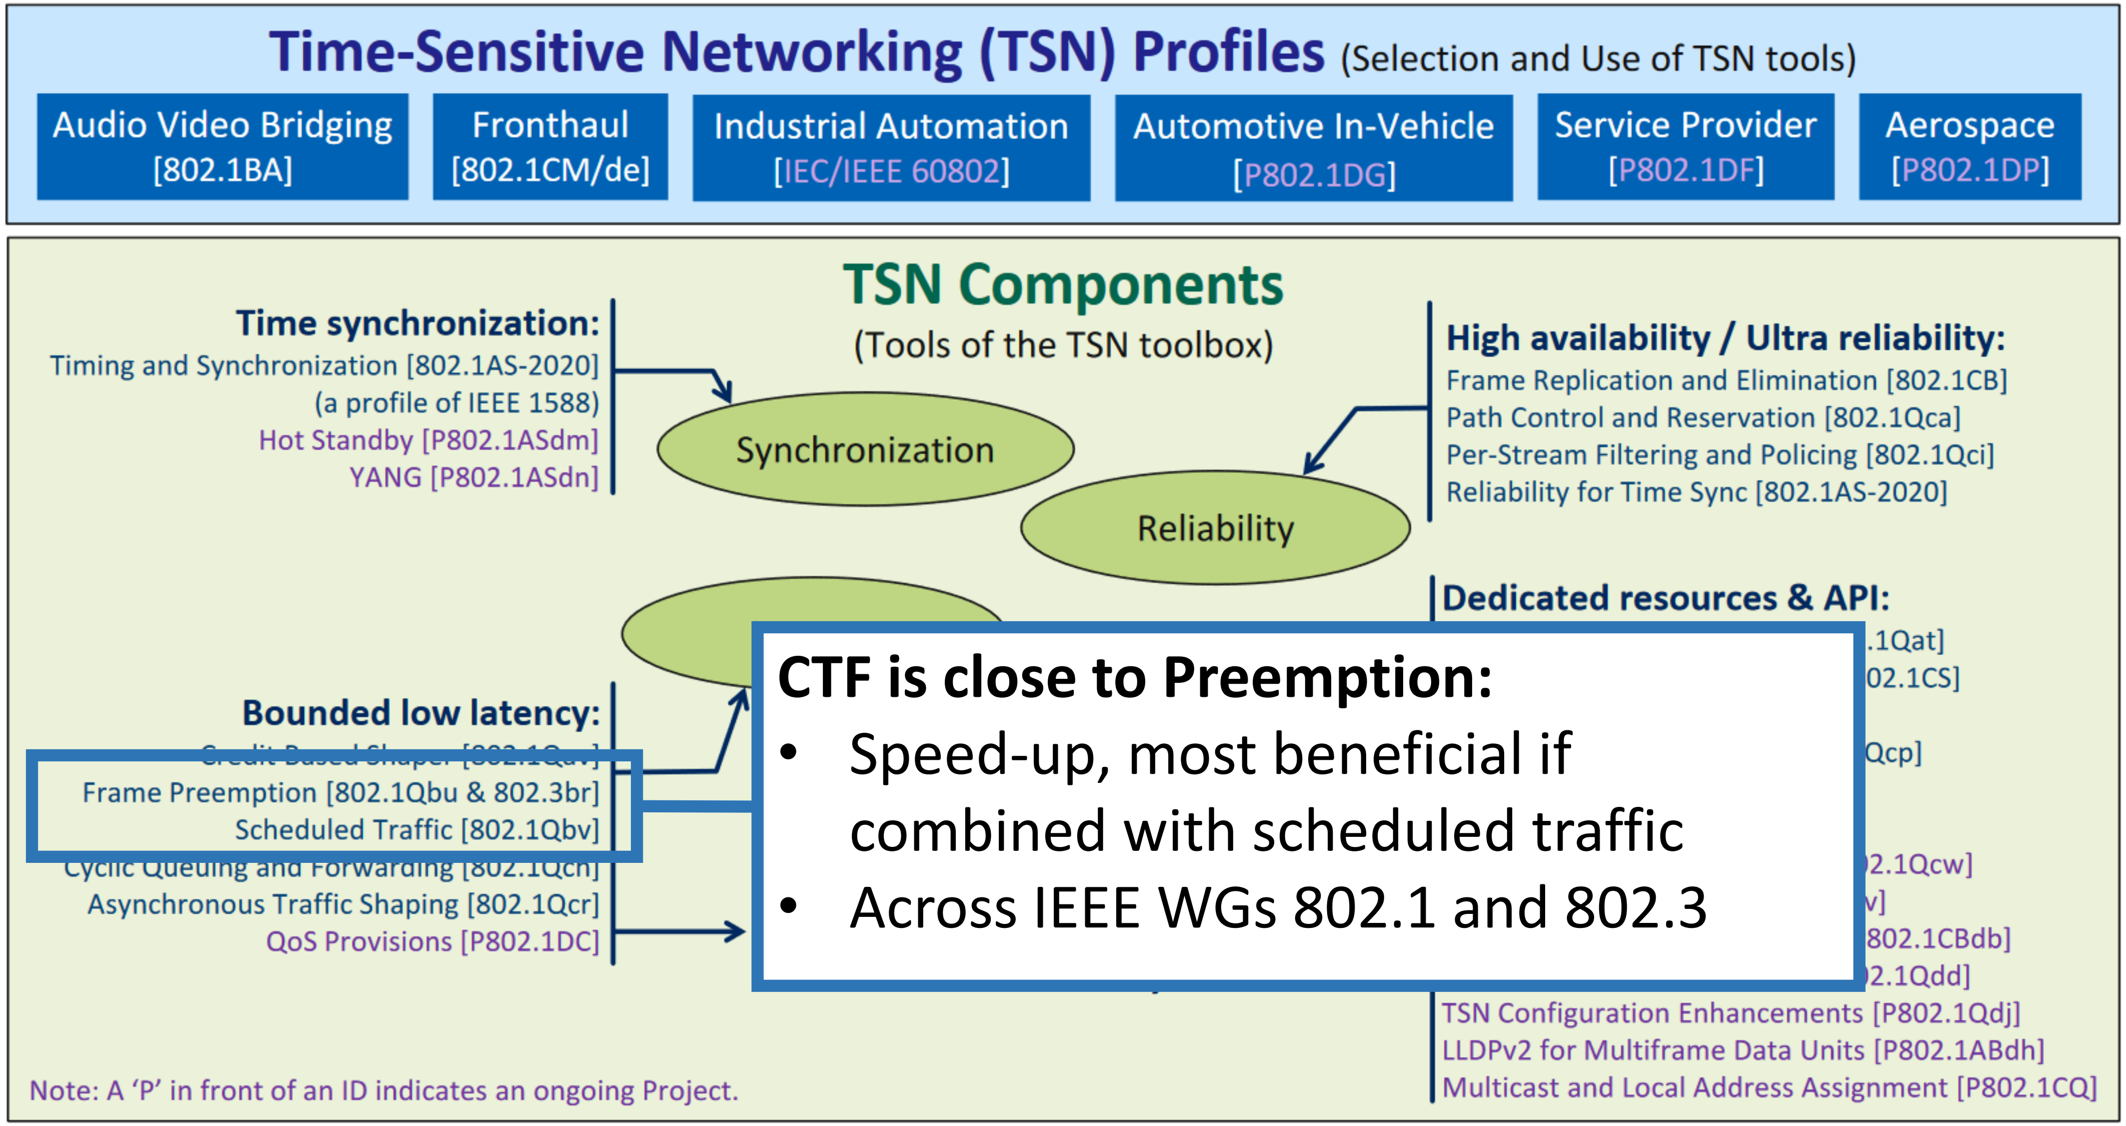

CTF is close to Preemption:

- Speed-up, most beneficial if combined with scheduled traffic
- Across IEEE WGs 802.1 and 802.3

Source: https://www.ieee802.org/1/files/public/docs2021/admin-tsn-summary-0221-v01.pdf

Basic CTF Operation

CTF is an alternative forwarding method to Store & Forward (S&F) in Bridges

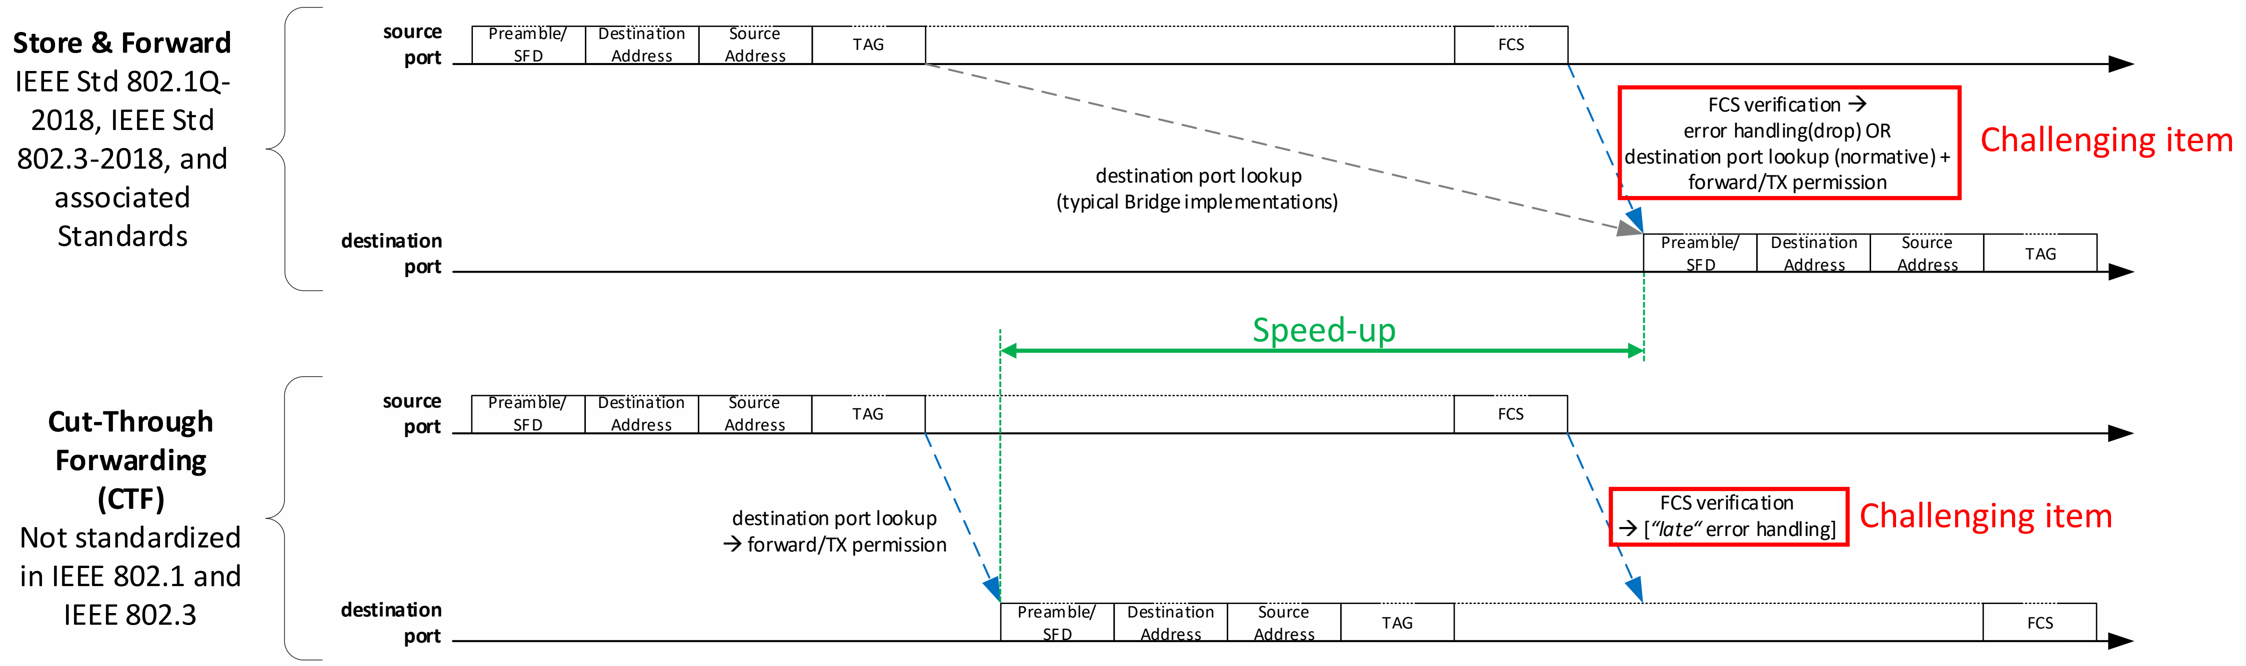

Delay performance enhancements

- Reduced residence times of frames in Bridges ("speed-up")
- Reduced frame length dependent jitter/delay variation

(Main) Challenges

- Transmission of frames with errors discovered by FCS verification, and the associated consequences
- S&F operation "deeply" manifested in IEEE 802.1 and 802.3 Standards

CTF Speed-up Analysis: Assumptions (1)

Purpose

- The following assumptions assemble a simplified model to focus on a simple speed-up analysis:
  - Some assumptions can be valid for some real systems, while being invalid for others.
  - The assumptions here are not intended as requirements or limitations for real systems with CTF.

Topology/Network

- Chain Network/Network segment
- Identical Link Speeds, Full-Duplex, negligible propagation delays
- CTF possible on all interconnections *except* from/to end stations (i.e., S&F at first and last hops)
- Strict Priority Transmission Selection Algorithm, optional with Enhancements for Scheduled Traffic

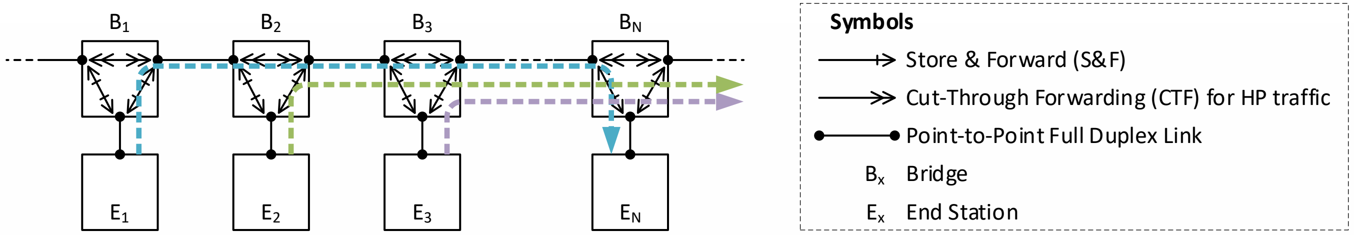

Errors

- Error free environment → no data corruption in frames
- However, errors, including late error handling, is addressed later in this tutorial

CTF Speed-up Analysis: Assumptions (2)

Traffic – Focus on Bounded Latency

- High Priority (HP): Focus of the Analysis
  - At most one stream sent by each end station, and each end station receives HP streams from at most one direction of the chain
  - Constant frame length[1]
  - Periodic (same period for all streams)
  - Period < maximum end-to-end latency
  - Nominal transmission times at sending end stations
- Low Priority (LP): Background
  - Always Store & Forward
  - Interferes with CTF traffic
    - Without preemption: 1542 octets (max. LP frame[1,2])
    - With preemption: 155 octets (max. non-preemptible LP frame[1,3])

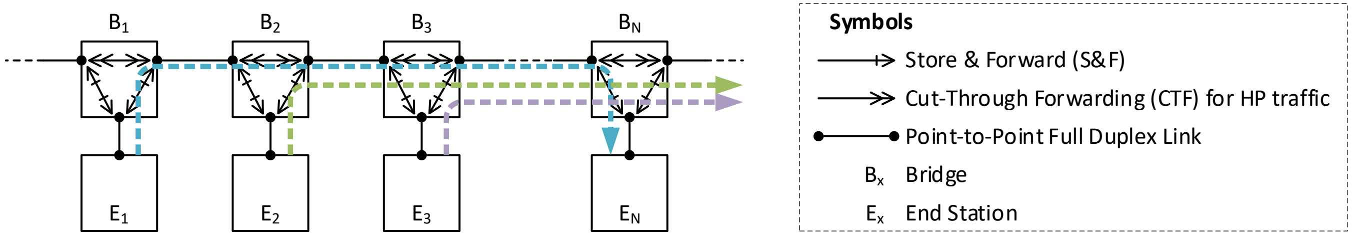

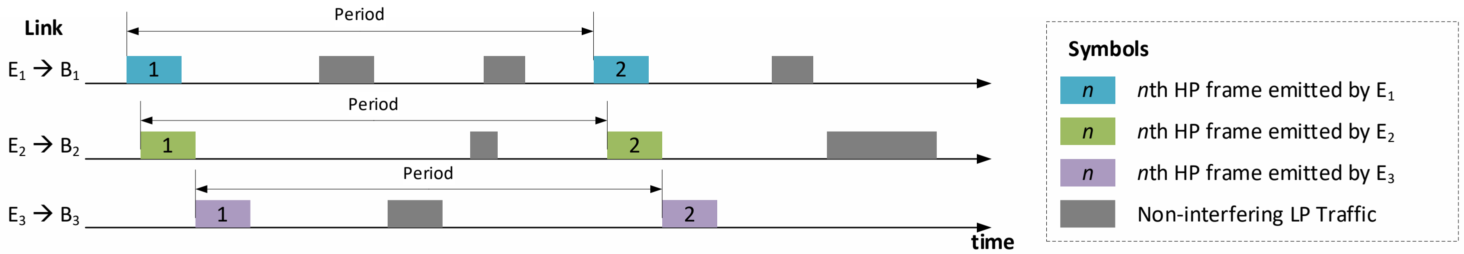

1) Includes all media-dependent overhead for IEEE 802.3 point-to-point full duplex media (Preamble, SFD, minimal Interpacket Gap).
2) Upper limit of 1500 octets payload in a tagged frame.
3) Defined upper limit for addFragSize=0 (cmp. 99.4.8 of IEEE Std 802.3br-2016).

CTF Speed-up Analysis: Math

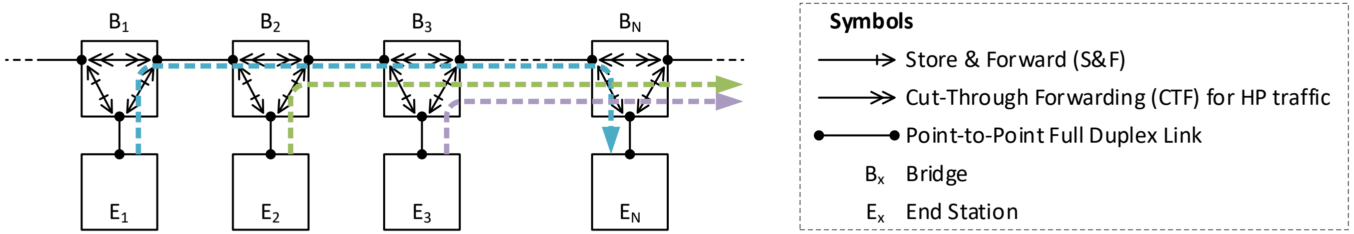

Delay until forwarding to destination ports happens. Assumed that the lookup starts after l_{Hdr} octets and finishes after d_{LU} µs. Note that the lookup can finish after frame completion during reception.

$$d_{SFF}^{max} = (H+2)\left(\max\{l_{HP}d_{Oct}, l_{Hdr}d_{Oct} + d_{LU}\} + d_Q\right) + \left((H+1)l_{LP} + Hl_{HP}\right)d_{Oct}$$

Maximum interference by crossing high priority traffic (l_{HP}) and crossing low priority traffic (l_{LP}). Dependent on the subsequently introduced communication schemes, either one or both types of interference exist or not (e.g., full TDM avoids both).

$$d_{CTF}^{max} = 2\left(\max\{l_{HP}d_{Oct}, l_{Hdr}d_{Oct} + d_{LU}\} + d_Q\right) + H\left(l_{Hdr}d_{Oct} + d_{LU} + d_Q\right) + \left((H+1)l_{LP} + Hl_{HP}\right)d_{Oct}$$

Separates the H interconnections (CTF) from the first and last ones (S&F). Note that, if the lookup finishes after frame completion during reception, then CTF provides no lower delay than S&F. The other way around, if the lookup is "fast enough", then CTF provides lower delays than S&F.

Symbol	Description
d_{SFF}^{max}	Maximum end-to-end delay without CTF of HP frames, in μs.
d_{CTF}^{max}	Maximum end-to-end delay with CTF of HP frames, in μs.
H	Number of possible CTF interconnections (e.g., N-2 for the stream of E_1).
l_{HP}	Frame size of high priority traffic (i.e., the traffic that can be subject to CTF), including all media dependent overhead, in octets.
l_{LP}	Frame size of low priority traffic (always S&F), including all media dependent overhead, in octets. Assumption: 1542 octets without preemption, 155 octets with preemption.
l_{Hdr}	Header length required for destination port lookup in Bridges, in octets. Assumption: 24 octets (preamble, start of frame delimiter, DA, SA, VLAN-Tag).
d_{Oct}	Nominal duration of an octet reflecting the link speed, in μs.
d_{LU}	Destination port lookup duration after l_{Hdr} octets were received, in μs. Assumption: 0.16 μs (e.g., 20 clock cycles @ 125 MHz).
d_Q	Interference-independent queuing delay (MAC delay, PHY delay, etc.), in μs. Assumption: 0.32 μs.

CTF Speed-up Analysis: Both Extremes

Uncoordinated

Interference by low priority and other high priority (CTF) traffic

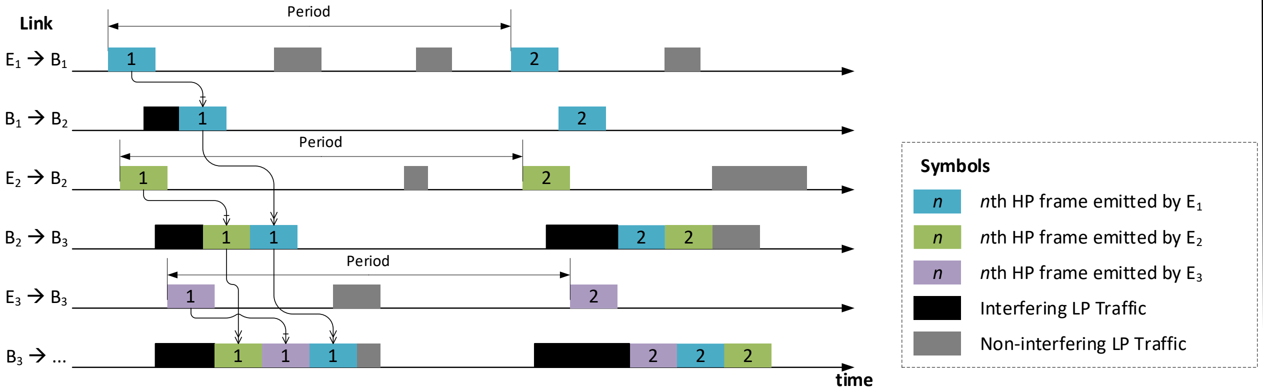

		SFF-to-CTF ratio									
		Preemption unsupported					Preemption supported				
H	Link \ l_{HP}	128	256	512	1024	1542	128	256	512	1024	1542
2	100 Mbps	96%	93%	88%	83%	80%	85%	80%	76%	74%	73%
4	100 Mbps	96%	91%	85%	79%	75%	82%	75%	70%	67%	66%
16	100 Mbps	95%	90%	82%	74%	70%	77%	68%	62%	59%	57%
64	100 Mbps	94%	89%	81%	73%	68%	76%	66%	60%	56%	54%
2	1 Gbps	97%	94%	89%	84%	81%	89%	82%	78%	75%	74%
4	1 Gbps	96%	92%	86%	80%	76%	86%	78%	72%	68%	67%
16	1 Gbps	96%	91%	83%	75%	70%	83%	72%	65%	60%	58%
64	1 Gbps	96%	90%	82%	74%	69%	82%	71%	62%	57%	55%
2	2,5 Gbps	98%	95%	90%	84%	81%	94%	86%	80%	76%	75%
4	2,5 Gbps	98%	93%	87%	81%	77%	92%	83%	75%	70%	68%
16	2,5 Gbps	97%	92%	85%	76%	71%	90%	78%	69%	62%	60%
64	2,5 Gbps	97%	92%	84%	75%	70%	90%	77%	67%	60%	57%

Full Time Division Multiplexing

No Interference

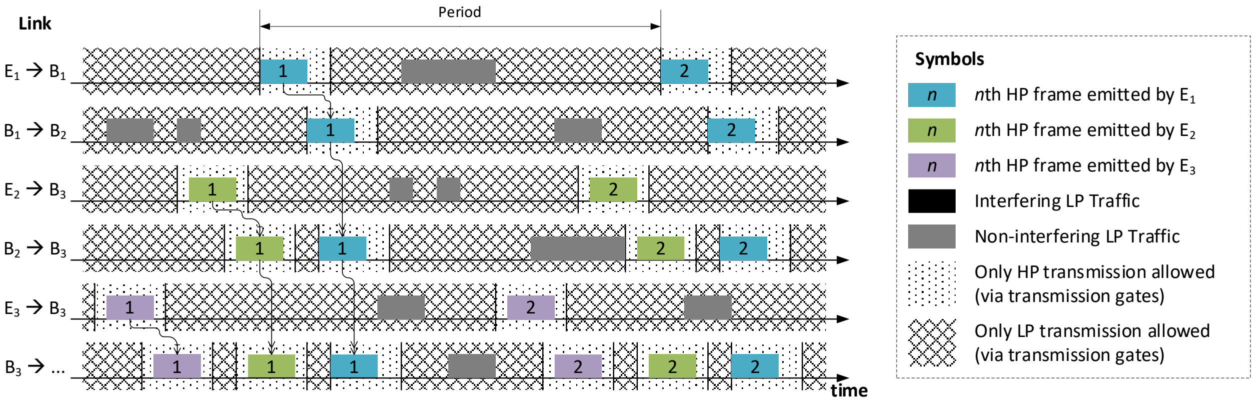

		SFF-to-CTF ratio				
		Preemption supported or not				
H	Link \ l_{HP}	128	256	512	1024	1542
2	100 Mbps	61%	56%	53%	51%	51%
4	100 Mbps	48%	41%	37%	35%	35%
16	100 Mbps	31%	21%	16%	14%	13%
64	100 Mbps	25%	14%	9%	6%	5%
2	1 Gbps	75%	64%	58%	54%	53%
4	1 Gbps	67%	52%	43%	39%	37%
16	1 Gbps	56%	36%	25%	18%	16%
64	1 Gbps	52%	31%	18%	11%	8%
2	2,5 Gbps	88%	74%	64%	58%	55%
4	2,5 Gbps	84%	66%	52%	44%	40%
16	2,5 Gbps	79%	55%	36%	25%	21%
64	2,5 Gbps	77%	50%	31%	18%	13%

Lower percent values indicate higher end to end delay performance improvements of CTF over S&F.

Reasons for standardizing CTF in IEEE 802

Interoperable and deterministic data plane (examples)

- Distinguish CTF Traffic from S&F Traffic
  - TAGs, Addresses, Ports?
- "Late" error handling
  - Shorten/truncate erroneous frames?
  - Mark erroneous frames?
  - Do nothing?
- Behavior of existing 802.1 Bridge mechanisms for CTF traffic
  - Flow Metering (e.g. Max. SDU size filters, MEF 10.3)?
  - Transmission selection algorithms?
  - Transmission gates?
  - Link speed transitions?[1]

Unified Management

- Elements
  - Configuration Parameters (e.g., enable/disable CTF)
  - Device properties (e.g., timing)
  - Status Variables (e.g., erroneous CTF frame counters)
- Required, for example, for automated, efficient and consistent TDM configuration (e.g., centralized network controller [802.1Qcc-2018])

Application and limitations of CTF in Networks

- Quality of Service[1,2]

Limit circulating erroneous frames in topological loops; limit bandwidth loss by erroneous frames

- Security[1]

Prevent exposure of frame contents (CTF <u>and</u> S&F) to untrusted network segments

1) See also https://ieee802.org/1/files/public/docs2017/new-tsn-thaler-cut-through-issues-0117-v01.pdf
2) See also https://www.ieee802.org/1/files/public/docs2019/new-seaman-cut-through-scissors-0119-v01.pdf

Possible integration into IEEE 802.1: General

Johannes Specht

Location in IEEE 802.1

Dedicated IEEE 802.1 Standard for CTF

- Not one or more amendment[s] to existing IEEE 802.1 Standards.

Reference Usage

- Select/import and adjust existing protocols and protocol procedures from other IEEE 802.1 Standards:
  1. IEEE Std 802.1Q-20xx
  2. IEEE Std 802.1CB-20xx
  3. IEEE Std 802.1AC-20xx

Some Implications

- At least some of the implications:
  1. No distribution of CTF across multiple IEEE 802.1 Standards documents
  2. Existing protocols and protocol procedures not addressed are basically “beyond specification”
  3. A simple way for inclusion without adjustment is basically *“as specified in x.y.z of IEEE Std 802.1A.B.C”*

Main Contents

Requirements for CTF in Bridges

CTF in Networks

- Structure and elements (e.g., “CTF Bridge”)
- QoS Maintenance/Requirements[1]
- Usage/Performance aspects[2]

CTF in Bridges

- Bridge data plane behavior and managed objects (YANG)
  - MAC Relay Entity/Forwarding Process
  - Bridge Port Transmit and Receive[3]

“Features” for QoS Maintenance and usage

1) Issues introduced by CTF (cmp. 6.5 of IEEE Std 802.1Q-20xx)
2) See earlier slides in this slide set
3) To the extent possible in IEEE 802.1

Possible integration into IEEE 802.1: CTF in Bridges

Johannes Specht

CTF in Bridges: Feature Set

- Required:
  1. IEEE Std 802.1Q-20xx: "Basic" VLAN/MAC Bridge Operations
  2. New for CTF: Fallbacks from CTF to S&F (i.e., to behavior from existing IEEE 802.1 Stds)
  3. New for CTF: Late error handling
- Options/within specification:
  1. IEEE Std 802.1Q-20xx: Per-Stream Filtering and Policing (PSFP)
  2. IEEE Std 802.1Q-20xx: Enhancements for Scheduled Traffic (EST)
  3. IEEE Std 802.1Q-20xx: Preemption
  4. IEEE Std 802.1Q-20xx: Frame Replication and Elimination for Reliability (FRER)
- For later discussion:
  1. New for CTF: Header check sequences[1]

1) Not necessarily required - header check sequences imply several challenges (interoperability with non-CTF Bridges, loose definition of headers, etc.). This topic can be considered thoroughly during a IEEE 802.1 standards development project.

CTF in Bridges: Traffic Identification, Separation and Transmission

Cut-Through Forwarding (CTF)
Not standardized in IEEE 802.1 and IEEE 802.3

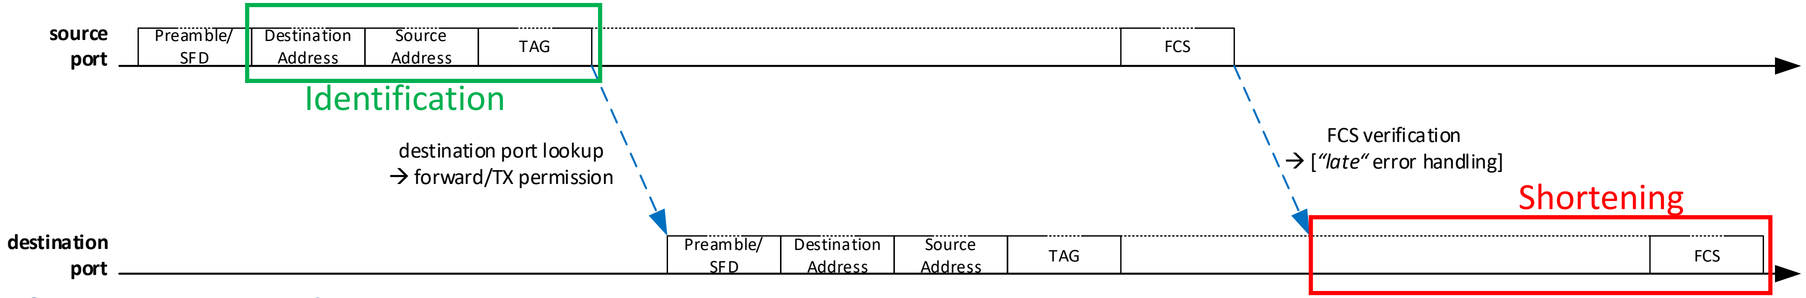

1. Identification by Port and Priority

Reception on a Port for which CTF has been enabled

AND (

Priority decoded from VLAN-TAG (6.9 and 6.20 of IEEE Std 802.1Q-20xx)

OR

FRER Stream Identification (IEEE Std 802.1CB-20xx),
used by PSFP (IEEE Std 802.1Q-20xx) for Internal Priority Value (IPV) assignments[1]

)

New Management Parameter(s)
- CTFReceiveEnable (Boolean, RW, default False)
- Per-Port

2. Separation by traffic classes

Queuing in traffic classes (8.6.8 of IEEE Std 802.1Q-20xx) for which CTF is supported and has been enabled

New Management Parameter(s)
- CTFTransmitEnable (Boolean, RW, default False)
- CTFTransmitSupported (Boolean, RO)
- Per-Port per traffic class

3. Transmission of CTF frames

- Strict priority transmission selection algorithm plus EST transmission gates (if supported)
- Abort transmission/shorten frames if FCS verification fails

1) The Mask-and-Match stream identification, as currently under development in IEEE P802.1CBdb, effectively enables a priority to be determined by at least the Destination Address. As one result, there are different (potentially co-existing) perceptions of a "header".

CTF in Bridges: Fallbacks to S&F

1. Implicit
   1. Interferences by other frames
2. Explicit (interference-independent operation)
   1. CTF reception is disabled on a Bridge Port
   2. CTF is disabled/unsupported by a traffic class on a Bridge Port
   3. No matching filtering entry in the FDB (i.e., flooding)
   4. Association of a frame under reception with a FRER recovery function
   5. Transmission Port link speed differs from that of the associated reception Port
   6. Content changes - TAG removal, insertion or replacement
3. Frames (copies) leaving the main path
   1. To Higher Layer Entities
   2. To the FDB for learning

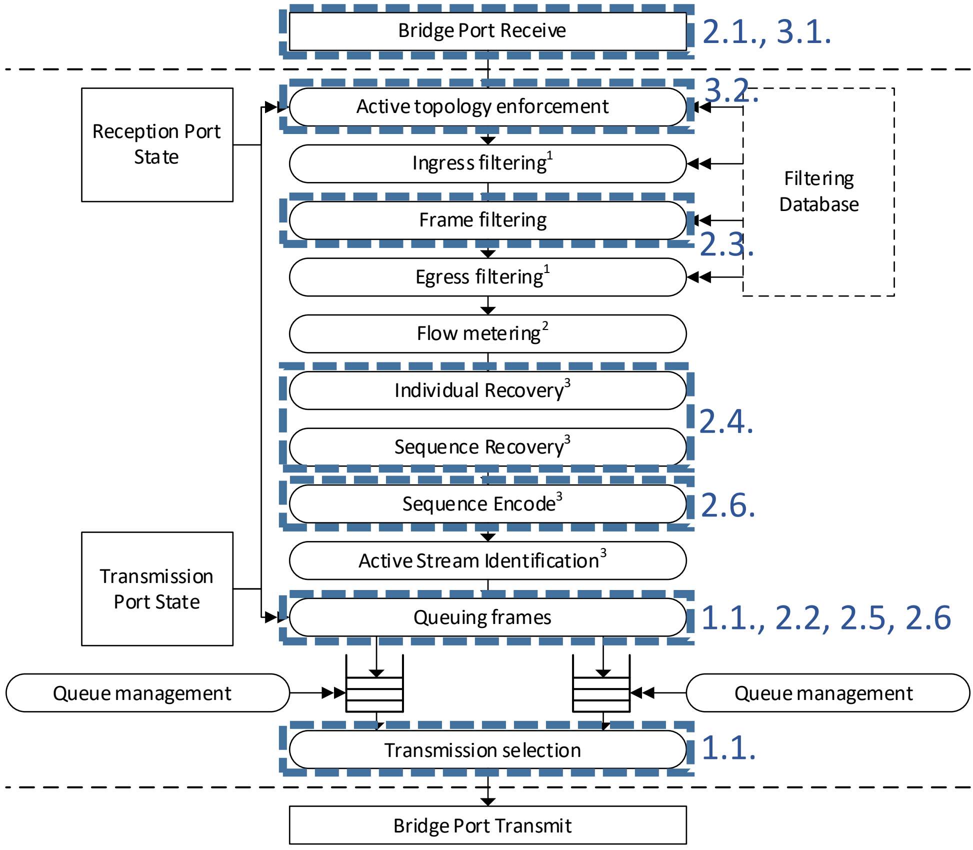

Cmp. 8.6 of IEEE Std 802.1Q-20xx and clause 8 of IEEE Std 802.1CB-20xx.
1) Not present in MAC Bridges
2) Not present if PSFP is unsupported
3) Not present if FRER is unsupported

CTF in Bridges: Late Errors

1. Causes

1. Errors discovered by FCS verification
2. PSFP's Maximum SDU size filtering limit reached during reception
3. PSFP stream gates transition to closed state[1]
4. Color of PSFP flow meters (MEF 10.3) transitions to red
5. The per traffic class maximum SDU size of EST is exceeded

2. Handling

1. Treat the frame end by PSFP's maximum SDU size filtering, stream gates and flow meters (MEF 10.3)
2. Remove the frame from all queues
3. Shorten the end of frame by an implementation-specific amount
4. Mark the end of frame by a special FCS

New Management Parameter(s)
- CTFTransmitShorteningMin (Integer, RO, nanoseconds)
- Per-Port

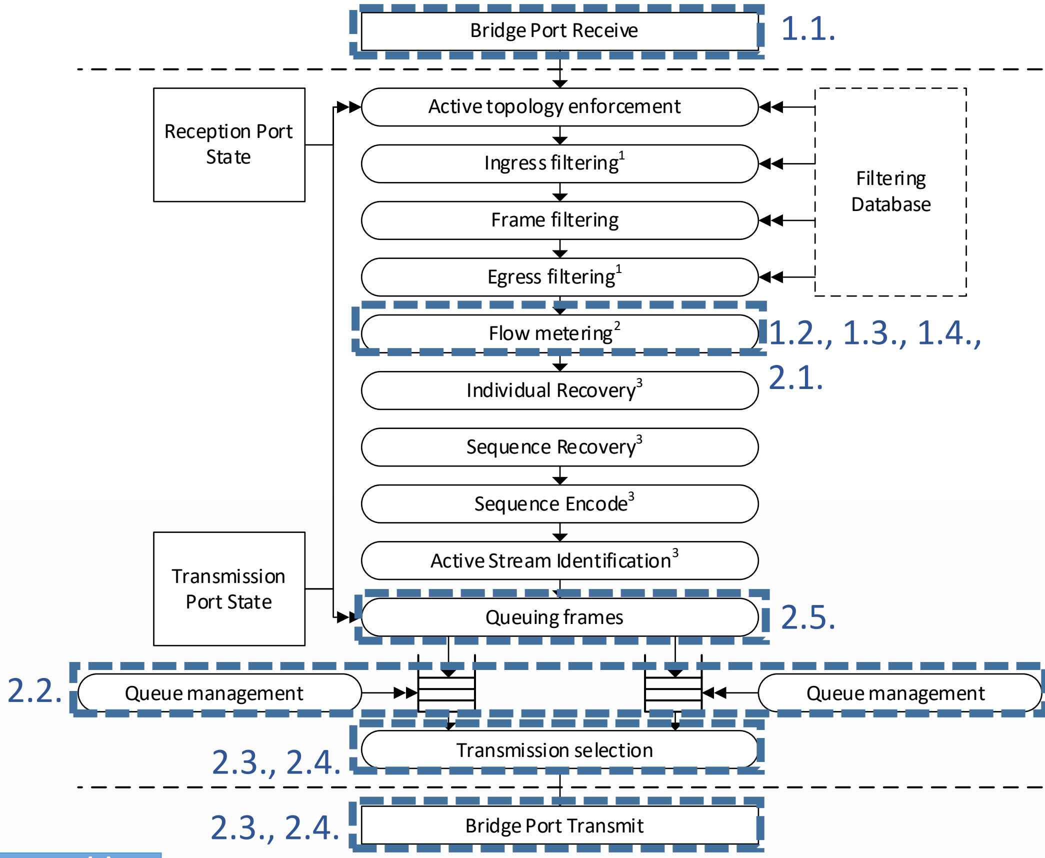

Cmp. 8.6 of IEEE Std 802.1Q-20xx and clause 8 of IEEE Std 802.1CB-20xx.
1) Not present in MAC Bridges
2) Not present if PSFP is unsupported
3) Not present if FRER is unsupported

New Management Parameter(s)
- CTFReceivedErroneousMarked (Counter, RW)
- CTFReceivedErroneousUnmarked (Counter, RW)
- Per-Port

1) In contrast to stream gates, it is not intended to involve late error handling if EST transmission gates transition to a closed state during transmission for compatibility (see 8.6.8.4 of IEEE Std 802.1Q-20xx)

Possible integration into IEEE 802.1: CTF in Networks

Johannes Specht

CTF in Networks: Circulating frames (1)

Problem Description

- Erroneous frame under reception by CTF Bridge are classified for CTF, and are transmitted by unintended Bridge Ports before FCS verification.
- The issue affects networks/network segments with topological loops, in which such frames can circulate for “a while”.

Observation

- It does not matter whether erroneous frames were intended for CTF or S&F
  1. Frames intended for S&F can be misclassified by the receiving CTF Bridge as CTF frames.
  2. Frames intended for CTF can remain classified for CTF, but match a wrong FDB entry (i.e., wrong port map).
  3. Frames misclassified as S&F frames are no issue (i.e., FCS verification prior to transmission).

Goal Definition

Frame removal after at most one round, if FCS verification can discover the error.

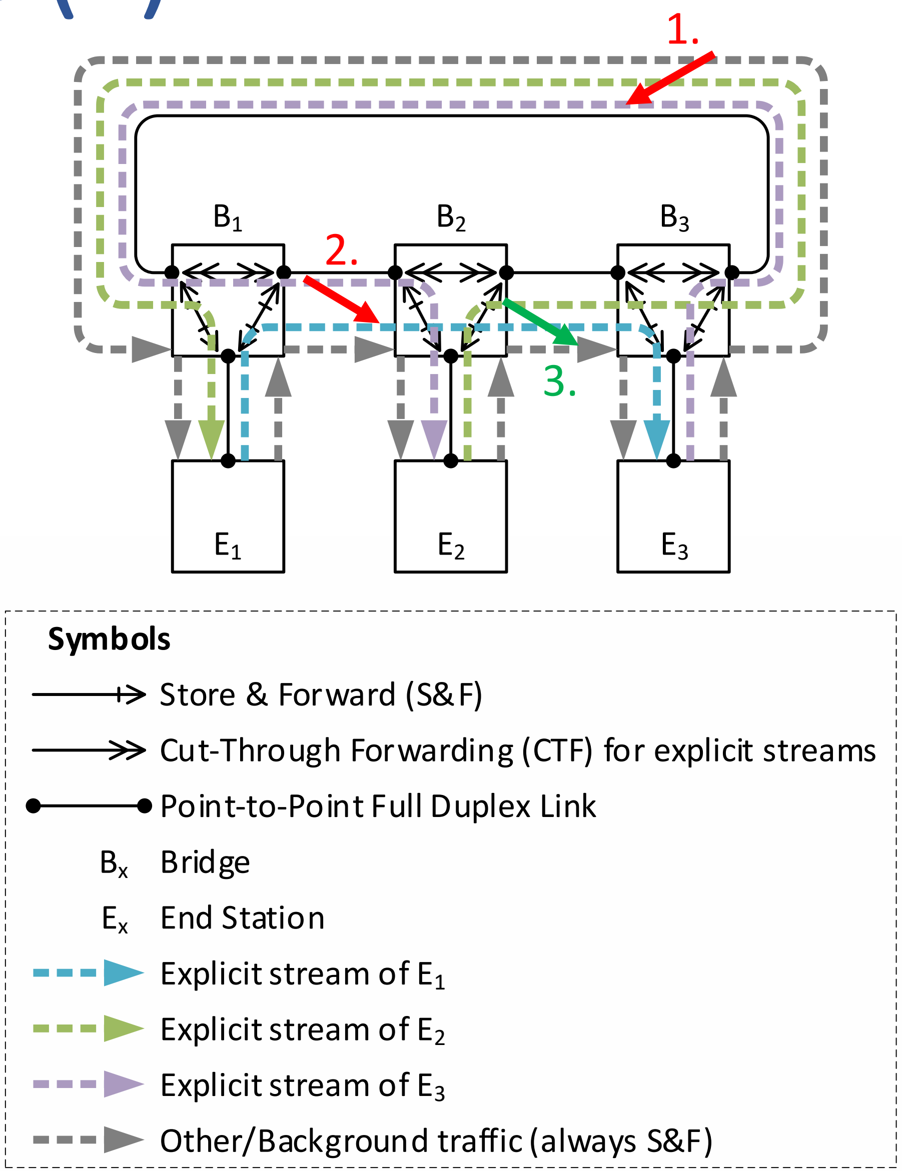

CTF in Networks: Circulating frames (2)

Network Requirements

- **Default**
  At least one S&F-only hop in each topological loop.
- **Potential Alternative**
  Only explicit FDB filtering entries for CTF traffic in all CTF Bridges in a loop
  AND
  the probability of errors affecting the same frame on two or more different links is negligible low.
- **Potential Alternative**
  The topological loop contains sufficient links(hops),
  AND
  all Bridges in the loop limit frame lengths of CTF traffic,
  AND
  the sum of the minimum frame shorting in all Bridges in the loop is greater than the frame length limit.

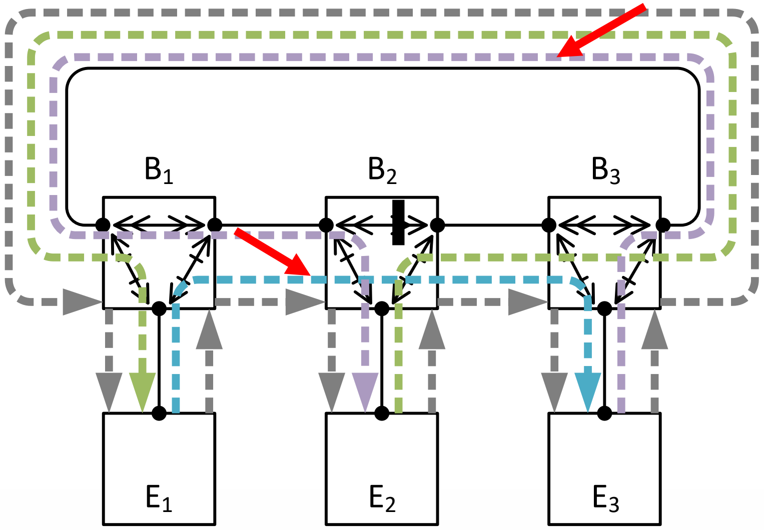

Symbols

- Store & Forward (S&F)
- Cut-Through Forwarding (CTF) for explicit streams
- Point-to-Point Full Duplex Link
- B_x Bridge
- E_x End Station
- Explicit stream of E_1
- Explicit stream of E_2
- Explicit stream of E_3
- Other/Background traffic (always S&F)

CTF in Networks: Bandwidth loss (1)

Problem Description

- Erroneous frame under reception by CTF Bridge are classified for CTF, and transmitted before FCS verification by
  - unintended Bridge Ports
    AND/OR
  - in the wrong traffic class.
- Such frames in the affected traffic class in Bridge transmission Ports can cause unplanned interferences in this traffic class or any higher priority traffic classes (oversized frames) and reduce the bandwidth available for lower priority traffic classes.
- The issue affects every traffic class in Bridge transmission Ports if CTF reception in at least one other Bridge Port is enabled.

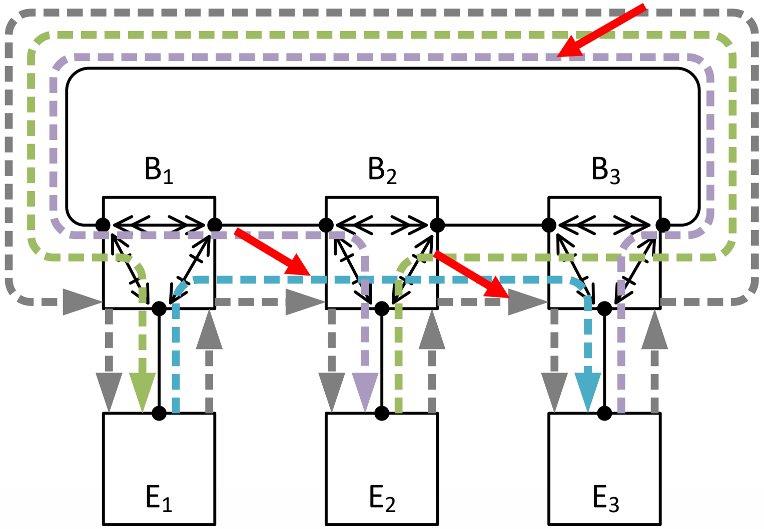

Symbols

- Store & Forward (S&F)
- Cut-Through Forwarding (CTF) for explicit streams
- Point-to-Point Full Duplex Link
- B_x Bridge
- E_x End Station
- Explicit stream of E_1
- Explicit stream of E_2
- Explicit stream of E_3
- Other/Background traffic (always S&F)

1) The planning required to properly configure PSFP can be inacceptable for some systems.
2) See the introduction of this slide set.

CTF in Networks: Bandwidth loss (2)

Network Recommendations

- Plan for additional interference/bandwidth usage
- If applicable[1], use disjoint redundant paths via FRER
- If applicable[2], use PSFP
  - Max. SDU size filtering can limit the effect of oversized frames
  - Proper usage of flow meters and/or stream gates depends on the traffic characteristics - for example[3]:
    - Flow meters (MEF 10.3) can limit the bandwidth of uncoordinated traffic
    - Stream gates can be used for TDM traffic

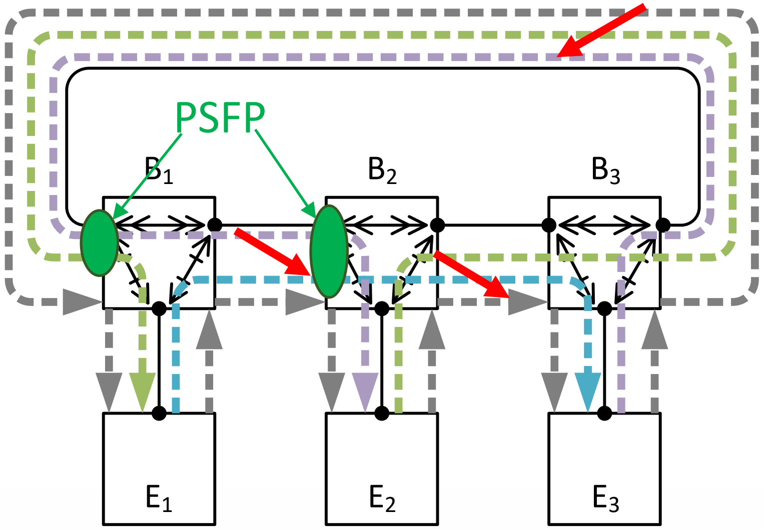

Symbols

- Store & Forward (S&F)
- Cut-Through Forwarding (CTF) for explicit streams
- Point-to-Point Full Duplex Link
- B_x Bridge
- E_x End Station
- Explicit stream of E_1
- Explicit stream of E_2
- Explicit stream of E_3
- Other/Background traffic (always S&F)

1) Disjoint paths are inacceptable for some systems (e.g., due to cost reasons).
2) The planning required to properly configure PSFP can be inacceptable for some systems.
3) See the introduction of this slide set.

Thank you for your Attention!

Questions, Opinions, Ideas?