Document Title	Explanation of Time Sensitive Network features
Document Category	Explanation (EXP)
Document UID	1091
Document Status	published
Part of AUTOSAR Standard	Foundation
Part of Standard Release	R26-11

Document Change History			
Date	**Release**	**Changed by**	**Description**
2026-11-26	R26-11	AUTOSAR Release Management	• Update of the ATS description
2025-11-27	R25-11	AUTOSAR Release Management	• Update of the ATS description
2024-11-27	R24-11	AUTOSAR Release Management	• No content changes
2023-11-23	R23-11	AUTOSAR Release Management	• Initial release

Disclaimer

This work (specification and/or software implementation) and the material contained in it, as released by AUTOSAR, is for the purpose of information only. AUTOSAR and the companies that have contributed to it shall not be liable for any use of the work.

The material contained in this work is protected by copyright and other types of intellectual property rights. The commercial exploitation of the material contained in this work requires a license to such intellectual property rights.

This work may be utilized or reproduced without any modification, in any form or by any means, for informational purposes only. For any other purpose, no part of the work may be utilized or reproduced, in any form or by any means, without permission in writing from the publisher.

The work has been developed for automotive applications only. It has neither been developed, nor tested for non-automotive applications.

The word AUTOSAR and the AUTOSAR logo are registered trademarks.

Table of Contents

1 Introduction 5
1.1 Objectives . . . 5
1.2 Scope . . . 5
2 Definition of terms and acronyms 6
2.1 Acronyms and abbreviations . . . 6
2.2 Definition of terms . . . 6
2.2.1 Ethernet packet . . . 6
2.2.2 Ethernet frame . . . 6
2.2.3 Stream . . . 7
2.2.4 IEEE1722-based stream . . . 7
2.2.5 Stream data producer . . . 7
2.2.6 Stream data consumer . . . 7
2.2.7 AVTP presentation time . . . 7
2.2.8 Max transit time . . . 7
2.2.9 Media clock . . . 8
2.2.10 Media clock provider . . . 8
2.2.11 Media clock consumer . . . 8
3 Related Documentation 9
3.1 Input documents & related standards and norms . . . 9
4 Time Sensitive Network in AUTOSAR 11
4.1 Constraints and Assumptions . . . 11
4.2 Overview . . . 11
4.2.1 IEEE 802.1AS . . . 11
4.2.2 IEEE 802.1Qav . . . 12
4.2.3 IEEE 802.1Qbv . . . 12
4.2.4 IEEE 802.1Qcc . . . 12
4.3 Usage of Shapers in the Network . . . 13
4.3.1 Fundamental Shaper and Scheduler Concepts . . . 13
4.4 Congestion . . . 13
4.4.1 Congested Traffic . . . 13
4.4.2 Congesting Traffic . . . 14
4.4.3 Summary and Conclusion . . . 14
4.5 Asynchronous Traffic Shaper . . . 15
4.5.1 Overview . . . 15
4.5.2 Description of the ATS algorithm . . . 16
4.5.3 ATS normal shaping operation . . . 17
4.5.4 Alternative functional ATS description . . . 18
4.5.5 Actual ATS transmission bursts . . . 18
4.5.6 Guideline for `MaxResidenceTime` configuration . . . 19
4.5.7 ATS implicit frame length limit . . . 20

4.5.8 ATS implicit rate limit 20
4.5.9 Consequences of the `groupEligibilityTime` 21
4.5.10 Guidance on the configuration 21
4.5.11 ATS Configuration parameters 23
4.6 Credit Based Shaper 24
4.6.1 Overview 24
4.6.2 CBS Configuration parameters 24
4.6.3 CBS credit estimation 24
4.6.4 Guidance on the configuration 25
4.7 Ensure Accuracy of Time Synchronization 25
4.7.1 Recommendations to support accuracy in an Ethernet network . . . 25
4.7.1.1 Parts in Documents which Reflect Accuracy 25
4.7.1.2 System requirements 26
4.8 Measurement of the Time Synchronization Quality (Accuracy) 26
4.8.1 Overview 26
4.8.2 Usage of PPS 26
4.8.2.1 Configuration of PPS within an Ethernet controller 26
4.8.2.2 Configuration of PPS by using OCU in combination with a PHC 27

1 Introduction

1.1 Objectives

Time-Sensitive Networking (TSN) is a set of standards specified by [1, IEEE 802.1 TSN Task Group]. It provides end-to-end data transmission with extremely low latency and high reliability over Ethernet. As a result, TSN is considered as one of the key enabler for deterministic communication on standard Ethernet, which can further be used to accurate timing and guarantee data delivery for automotive Ethernet.

The purpose of this document is to give an overview on the relevant TSN features, and provide examples how to use them together with AUTOSAR platform.

1.2 Scope

This document provides explanatory information and help the automotive engineer to understand the fundamental TSN features, as well as the basic approach to design a deterministic network on automotive Ethernet.

The content of this document will address following topics:

- Constraints and Assumptions
- Overview
- Usage of Shapers
- Accuracy of Time Synchronization
- Measure Quality (Accuracy) of Time Synchronization

2 Definition of terms and acronyms

2.1 Acronyms and abbreviations

Abbreviation / Acronym:	Description:
TSN	Time Sensitive Networking
TSpec	Traffic specification as defined by [2, IEEE Std 802.1Q-2022]
AVTP	Audio/Video Transport Protocol as defined by [3, IEEE Std 1722-2016]
AVTPDU	Audio/Video Transport Protocol Data Unit as defined by [3, IEEE Std 1722-2016]
IIDC	Instrumentation and Industrial Digital Camera as defined by [3, IEEE Std 1722-2016]
61883_IIDC	IEC 61883/IIDC format as defined by [3, IEEE Std 1722-2016]
AAF	AVTP Audio Format as defined by [3, IEEE Std 1722-2016]
CRF	Control Reference Format as defined by [3, IEEE Std 1722-2016]
TSCF	Time Sensitive Control Format as defined by [3, IEEE Std 1722-2016]
NTSCF	None Time Sensitive Control Format as defined by [3, IEEE Std 1722-2016]
ACF	AVTP Control Format as defined by [3, IEEE Std 1722-2016]
ACF_CAN	Controller Area Network (CAN)/CAN with Flexible Data-Rate (CAN FD) message as defined by [3, IEEE Std 1722-2016]
ACF_CAN_BRIEF	Abbreviated CAN/CAN FD message as defined by [3, IEEE Std 1722-2016]
ACF_LIN	LIN® message as defined by [3, IEEE Std 1722-2016]
RVF	Raw Video Format as defined by [3, IEEE Std 1722-2016]
gPTP	generalized Precision Time Protocol [4, IEEE Std 802.1AS]
PPS	Pulse Per Second physical signal
1PPS	1 Pulse Per Second physical line
GPS	Global Positioning System

Table 2.1: Acronyms and abbreviations used in the scope of this Document

2.2 Definition of terms

2.2.1 Ethernet packet

Definition: An "Ethernet packet" is an on wire format defined by [5, IEEE Std 802.3-2022] which includes the following parts: Preamble (7 bytes), SFD (start frame delimiter, 1 byte), Ethernet frame (up to 2000 bytes), IPG (inter package gap, 12 byte times)

2.2.2 Ethernet frame

Definition: An "Ethernet frame" is on wire format defined [5, IEEE Std 802.3-2022] which includes the following parts: MAC destination address field (6 bytes), MAC source address field (6 bytes), Type field (2 bytes), MAC client data field (include op-

tional Q-Tag (4 bytes)) (up to 1982 bytes), optional PAD (padding bytes), FCS (frame check sequence, 4 bytes)

2.2.3 Stream

Definition: A "stream" represent multiple Ethernet frames which are grouped by similar frame attributes (e.g. MAC source address)

2.2.4 IEEE1722-based stream

Definition: A "IEEE1722-based stream" represent multiple Ethernet frames which have EtherType set to AVTP Ethertype (0x22F0) (see [3, IEEE Std 1722-2016])), carry an IEEE1722 specified AVTP header and an AVTP stream data subtype (e.g. AAF (AVTP Audio Format))

2.2.5 Stream data producer

Definition: A "stream data producer" represent an end node in an Ethernet network which (continuously) produces data. The data is transmitted via a stream and received by 1 or multiple end nodes (stream data consumer)

2.2.6 Stream data consumer

Definition: A "stream data consumer" represent an end node in an Ethernet network which (continuously) consumes data. The data is received via a stream.

2.2.7 AVTP presentation time

Definition: The "AVTP presentation time" is defined by [3, IEEE1722-2016]) and represents the globally synchronized time at which designated data within an AVTPDU payload is transferred to a time-sensitive application of an stream data consumer. Please note: AVTP presentation time is calculated as "TavtpPresentationTime" = "TcurrentGlobalTime" + "TmaxTransitTime".

2.2.8 Max transit time

Definition: "Max transit time" is defined by [3, IEEE1722-2016]) and represents the maximal acceptable delay when data of a stream data producer is added to its egress queue, transferred across the network via an IEEE1722-based stream and this data

is forwarded from the ingress queue of the receiving stream data consumer to a time-sensitive application.

2.2.9 Media clock

Definition: "Media clock" is defined by [3, IEEE1722-2016]) and represents an entity which generates a periodic signal with a constant rate (e.g. precise hardware clock with a constant rate (e.g. 48kHz)). The media clock is hosted by the media clock provider.

2.2.10 Media clock provider

Definition: A "media clock provider" is an end node in the network which hosts an media clock. The media clock provider transmit an IEEE1722-based stream to 1 or multiple media clock consumer. The IEEE1722-based stream is of subtype Clock Reference Format (CRF) and contain several presentation timestamps which correlates to the media clock rate.

2.2.11 Media clock consumer

Definition: A "media clock consumer" is an end node in the network which receives a IEEE1722-based stream of subtype Clock Reference Format (CRF) from a media clock provider. The media clock consumer performs a recovery of its media clock (e.g. PLL) based on the received encapsulated data from the media clock provider.

3 Related Documentation

3.1 Input documents & related standards and norms

[1] IEEE 802.1 TSN Task Group
https://1.ieee802.org/tsn/

[2] IEEE 802.1Q-2022 - IEEE Standard for Local and Metropolitan Area Network - Bridges and Bridged Networks
https://ieeexplore.ieee.org/

[3] IEEE Standard 1722-2016 - IEEE Standard for a Transport Protocol for Time-Sensitive Applications in Bridged Local Area Networks

[4] IEEE 802.1AS-2011 - IEEE Standard for Local and metropolitan area networks - Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks, Rev. 2011

[5] IEEE 802.3-2022
https://www.ieee802.org/3/

[6] IEEE 802.1CB-2017 - IEEE Standard for Local and Metropolitan Area Network - Frame Replication and Elimination for Reliability
https://ieeexplore.ieee.org/

[7] IEEE 802.1AS-2020 - IEEE Standard for Local and metropolitan area networks - Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks, Rev. 2020

[8] IEEE 1588-2019: IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems

[9] IEEE Standard for Local and metropolitan area networks – Bridges and Bridged Networks - Amendment 25: Enhancements for Scheduled Traffic
https://ieeexplore.ieee.org/

[10] IEEE Standard for Local and metropolitan area networks-Virtual Bridged Local Area Networks-Amendment 12: Forwarding and Queuing Enhancements for Time-Sensitive Streams
https://ieeexplore.ieee.org/

[11] IEEE Standard for Local and Metropolitan Area Networks–Bridges and Bridged Networks – Amendment 31: Stream Reservation Protocol (SRP) Enhancements and Performance Improvements
https://ieeexplore.ieee.org/

[12] P802.1Qcr - Bridges and Bridged Networks Amendment: Asynchronous Traffic Shaping

[13] Specification of Ethernet Switch Driver
AUTOSAR_CP_SWS_EthernetSwitchDriver

[14] Time Synchronization over Ethernet Protocol Specification
AUTOSAR_FO_PRS_TimeSyncOverEthernetProtocol

[15] Specification of Time Synchronization over Ethernet
AUTOSAR_CP_SWS_TimeSyncOverEthernet

[16] 802.1AS Recovered Clock Quality Testing
https://avnu.org/wp-content/uploads/2014/05/Avnu-Testability-802.1AS-
-Recovered-Clock-Quality-Measurement-1.0_Approved-for-Public-Release.pdf

4 Time Sensitive Network in AUTOSAR

4.1 Constraints and Assumptions

AUTOSAR does not support the following TSN features:

- time aware shapers
- active stream identification
- frame replication

AUTOSAR assumes that all Ethernet switches in the network are VLAN-aware. VLAN-unaware Ethernet switches are not supported.

4.2 Overview

TSN is a set of IEEE 802 standards based on Ethernet. It specifies functions to ensure the transmission with bounded low latency and high availability for applications in automotive or industrial control facilities. The deterministic transmission of time-critical frames may be achieved by using TSN features of high reliability, low delay, zero packet loss through traffic shaping, seamless redundant transmission, filtering, priority-based scheduling, ingress policing and rate limitation.

Some of these IEEE documents are stand-alone, like [6, IEEE Std 802.1CB-2017] which are each further developed through annexes and revisions, or they are initially annexes of [2, IEEE Std 802.1Q] which get rolled up into the main document and may be updated, as the main document is developed further through revisions.

While it is customary to often refer back to the original Q-annex a feature was introduced in, the reader is advised to always check the latest revision of the main document (e.g. currently [2, IEEE Std 802.1Q-2022]) for updates and changes to the functionality.

The list of documents considered part of the TSN family is constantly evolving. The reader is advised to refer to the IEEE's TSN task group website (`https://1.ieee802.org/tsn/`) for a current list.

The following subsections only give highlights and are not to be assumed current nor complete.

4.2.1 IEEE 802.1AS

IEEE 802.1AS ([7, IEEE Std 802.1AS-2020]) was derived as a profile of IEEE Std 1588 ([8, Precision Time Protocol]) operating on the data link layer (ISO/OSI layer 2). It may be surprising to some, that only very few TSN features require time-synchronization as a basis. The most well known are the Enhancements for scheduled traffic in section

8.6.8.4 of [7, IEEE Std 802.1AS-2020], originally developed in [9, IEEE Std 802.1Qbv-2015] (see 4.2.3).

Time synchronization is used mostly for application level communication and not as an infrastructure element to manage the network.

4.2.2 IEEE 802.1Qav

IEEE 802.1Qav [10, IEEE Std 802.Qav] can be used to ensure that traditional asynchronous Ethernet data traffic does not interfere with AVB's real-time audio and video streams. To avoid competition for network resources between ordinary data traffic and AVB traffic, the AVB switch differentiates time-sensitive audio and video streams from ordinary data streams, queues real-time frames and asynchronous frames separately, and assigns the highest priority to real-time frames. Therefore a so-called Credit Based Shaper (CBS) could be used for bandwidth limitation to convergence time-sensitive and time-insensitive traffic.

CBS is described in more detail in section 4.6.

4.2.3 IEEE 802.1Qbv

IEEE 802.1Qbv [9, IEEE Std 802.Qbv] specifies the feature of time-aware shaping, which allows the scheduling of the transmission of frames with different priorities in time-triggered windows. Time-aware shaping divides time into fixed windows, and schedules the transmission of frames within these windows based on their priority. Time-critical frames are normally given with higher priority and transmitted in a separate window, while lower priority frames are transmitted in the remaining time.

4.2.4 IEEE 802.1Qcc

Most networks need to be configured while the network is down, which is not suitable for applications such as in industrial control and automotive network communication. IEEE 802.1Qcc [11, IEEE Std 802.Qcc] can be used to introduce a Centralized Network Configuration (CNC) and a Centralized User Configuration (CUC) to implement dynamic network configuration, and flexibly configure new devices and data flows when the network is running.

4.3 Usage of Shapers in the Network

4.3.1 Fundamental Shaper and Scheduler Concepts

The following definitions are used in this section:

1. Scheduling: Is the action of assigning resources
2. Traffic Shaping: Is a bandwidth management technique
3. Policing: Is the rule (based on header fields, ingress or egress ports, ...) or bandwidth (data per time) based discarding of frames

If the resource assigned by a scheduler is (access to) bandwidth, then a scheduler will implicitly perform Traffic Shaping. E.g.: The Enhancements for Scheduled Traffic (of [2]) are based on scheduling access to a transmission port based on a timetable.

Strict Priority Queuing (SPQ) assigns access to a transmission port based on priority. But SPQ does not manage the bandwidth, as the highest priority can get 100% access to the transmission port.

An ingress rate limiter discards frames (polices them) in case they arrive too frequently. Less bandwidth is required for the left over egress traffic, but one would usually not call this shaping.

Schedulers and shapers can ever only delay transmission. This will increase the worst case latency of the frames being shaped. On the other hand, if the shaped traffic is in a high priority, the shaper will create gaps for lower priority frames to be transmitted, thereby reducing the latency of these lower priority flows.

Shapers can be viewed as `non-work-conserving`, i.e. they will not transmit frames of a high priority flow, even if no frames of a lower priority flow are available to transmit, if the high priority flow would exceeded its configured bandwidth limit.

The de-bouncing features in Autosar are not `bandwidth-conserving` as they require a minimum time between two transmissions. If a transmission is delayed, this delay will carry forward and there is no way to make up for the lost bandwidth.

4.4 Congestion

4.4.1 Congested Traffic

Congested Traffic is under-served on egress, i.e. frames accumulate in the egress buffer of the relay as they do not get selected for transmission.

Traffic can become congested, if:

1. The shaper on egress is configured with too low a rate or too a small burst size.
2. Excessively large lower priority Frames block the transmission.

3. A large higher priority burst blocks the transmission.

Therefore in order to prevent traffic from becoming congested:

1. Bandwidth needs for all traffic must be well known.
2. All traffic goes through a max. SDU Size Filter ([2, IEEE Std 802.1Q-2022]:8.6.5.3.1).
3. All higher priority traffic (and therefore all traffic, except the lowest TC) goes through a Flow Meter ([2, IEEE Std 802.1Q-2022]:8.6.5.5).

4.4.2 Congesting Traffic

Congesting Traffic is providing excess ingress bandwidth, i.e. more buffer is (or would be) required to store incoming frames than was expected/configured.

Traffic causes congestion, if: a) Bursts of random size occur, as no shaping is implemented. b) A Shaper (CBS, ATS, TAS) with a mis-configured rate or burst size is active. c) The Traffic contains excessively large Frames.

Therefore in order to prevent Traffic from causing congestion: d) All Traffic is shaped on egress at every Talker. e) Bandwidth needs for all Traffic must be well known. f) All traffic goes through a max. SDU Size Filter ([2, IEEE Std 802.1Q-2022]:8.6.5.3.1).

4.4.3 Summary and Conclusion

Congestion separation is to be performed on each Bridge locally. a) A Bridge is configurable to segregate the Buffers for different Traffic, so Buffer overruns in one segregated block can lead to discarding of Frames within that one Traffic aggregate, but not for the other Traffic aggregates. b) Frames are not forwarded on egress, if the bandwidth exceeds a configured maximum. c) Frames are discarded if the bandwidth exceeds a configured maximum on ingress.

4.5 Asynchronous Traffic Shaper

4.5.1 Overview

Asynchronous Traffic Shaping (ATS) was originally specified by [12, IEEE Std 802.1Qcr-2020] and provides a traffic shaping concept that can provide bounded latencies for data streams in a network. It allows for bursts and provides a fair usage of the egress port for streams arriving on different ingress ports. This is a main difference between the Credit Based Shaper (CBS) and ATS, as CBS only considers a particular traffic class of an egress port. The ATS shaping algorithm is loosely based on the token bucket algorithm and uses similar terminology for the configuration parameters.

ATS is divided into two major functional blocks, one, the ”ATS Eligibility Time Assignment” (see 4.5.2) situated on the ingress side of a bridge, operating on streams identified by the stream filters using the `stream_handle` and `priority`, includes the computational part of the shaping function. This means it calculates an eligibility time for every frame received based on the ATS algorithm. The eligibility time is the earliest time when a frame is eligible for transmission on the egress port. The other functional block, the ”ATS transmission selection algorithm” is situated on the egress side and selects the traffic according to the results calculated by the ”ATS Eligibility Time Assignment”. As the eligibility times of frames in streams from different receive ports can be in non-decreasing order, ATS queues need to support non-FIFO queues.

The ATS assumes a time base to be available, this does however not need to be a system wide synchronized timebase. In [2, IEEE Std 802.Q-2022:Annex V] some details are given on what effects different deviations in the time bases of bridges implementing ATS instances within a network can have.

IEEE Terminology	Mathematical Symbol
`CommittedBurstSize`	C
`CommittedInformationRate`	r
`MaxResidenceTime`	Δ_{max}
`eligibilityTime`	$\tau[n]$
`schedulerEligibilityTime`	$\tau_0[n]$
`GroupEligibilityTime`	$\tau_G[n]$
`lengthRecoveryDuration`	$L[n]/r$
`length(frame)`	$L[n]$
`arrivalTime(frame)`	$t_A[n] =$ NOW!
`BucketEmptyTime`	$E[n]$
`bucketFullTime`	$F[n]$
`emptyToFullDuration`	C/r
Virtual Credit (NOT used by IEEE)	$c[t]$
Virtual frame Number (NOT used by IEEE)	n

Table 4.1: Mathematical symbols to describe the ATS algorithm

4.5.2 Description of the ATS algorithm

The ATS algorithm is defined by section 8.6.11.3 of [2, IEEE Std 802.1Q-2022]. It assigns an `eligibilityTime` $\tau[n]$ to a frame n received at `arrivalTime` $t_A[n]$.

The IEEE's ATS algorithm uses `BucketEmptyTime` E to represent the current virtual credit. If `BucketEmptyTime` E is in the past of `arrivalTime` ($t_A[n] > E$), positive virtual credit has been accumulated. If `BucketEmptyTime` E is in the future, virtual credit is negative. Virtual credit is constantly accumulated at `CommittedInformationRate` r.

Therefore, while not actually used in the IEEE's description, ATS keeps track of a virtual credit c, which can usually be calculated by:

$$c[t] = (t - E) \times r \tag{4.1}$$

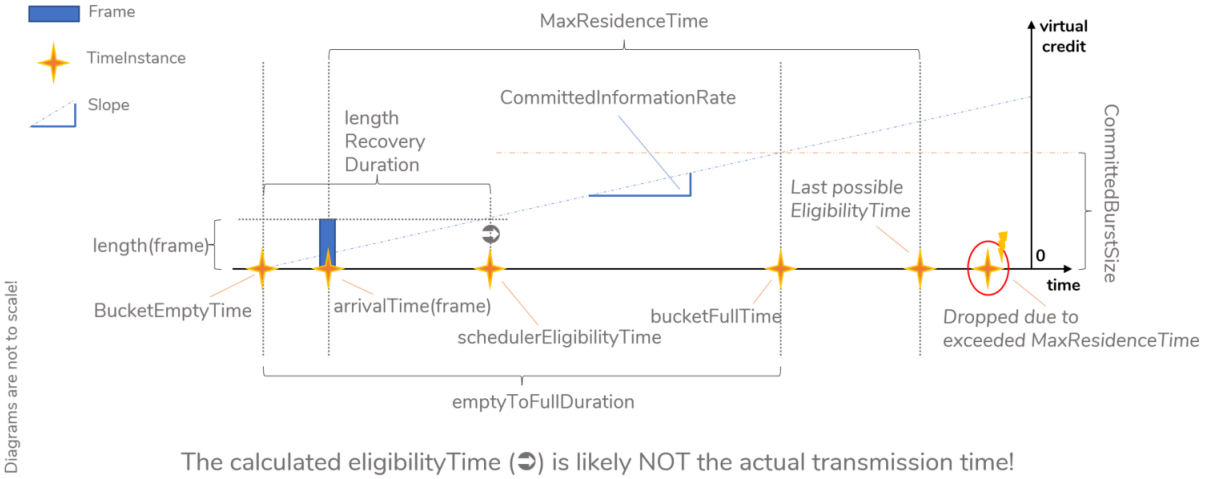

Figure 4.1: ATS Terminology Introduction

At `arrivalTime` the IEEE ATS algorithm initially calculates a `schedulerEligibilityTime`:

$$\tau_0[n] = E[n] + L[n]/r \tag{4.2}$$

and sets the `eligibilityTime` to:

$$\tau[n] = \max\left(t_A[n], \tau_G, \tau_0[n]\right) \tag{4.3}$$

Ignoring τ_G (see 4.5.9 for details) initially in this explanation the frame becomes eligible for transmission immediately (i.e., at t_A) if sufficient virtual credit $c[t_A]$ for the frame of length L is available, or at τ_0 in the future of t_A as soon as sufficient virtual credit has accumulated. There is no check here whether $\tau_0 > E[n] + C/r$, i.e. it is possible for $L > C$ (see 4.5.5 for details).

It then checks for:

$$\tau \leq t_A[n] + \Delta_{max} \tag{4.4}$$

The frame is discarded if τ lies too far in the future, as determined by Δ_{max}. The `BucketEmptyTime` is not adjusted if the frame is discarded. There is no required relation in [2, IEEE Std 802.1Q] between L, C, and Δ_{max} (see 4.5.6 for details).

If the frame is not discarded, the `GroupEligibilityTime` $\tau_G[n+1]$ for the next frame is set to $\tau[n]$, which means the next frame within this Group will not be eligible for transmission before $\tau[n] \leq \tau[n+1]$ according to (4.3), see 4.5.9 for details.

Then the `BucketEmptyTime` is adjusted to account for the reduction in virtual credit according to the frame's length $L[n]$ at its `eligibilityTime` $\tau[n]$. Two cases need to be considered:

1. If the virtual credit would be above the `CommittedBurstSize` at the `eligibilityTime`, i.e.,
$$\tau[n] \geq E[n] + C/r \tag{4.5}$$
Then `BucketEmptyTime` $E[n+1]$ is adjusted, so virtual credit (as per equation (4.1)) reaches `CommittedBurstSize` C after the frame of length $L[n]$ was subtracted at `eligibilityTime` τ:
$$C = ((\tau[n] + L[n]/r) - E[n+1]) \times r \tag{4.6}$$
$$E[n+1] = L[n]/r + \tau[n] - C/r \tag{4.7}$$
This is expressed in [2, IEEE Std 802.1Q-2022] as:
$$E[n+1] = (E[n] + L[n]/r) + \tau[n] - (E[n] + C/r) = \tau_0[n] + \tau[n] - F[n] \tag{4.8}$$
See 4.5.5 how this enables bursting in ATS.
2. Otherwise the new `BucketEmptyTime` is set equal to the `schedulerEligibilityTime`:
$$E[n+1] = \tau_0[n] \tag{4.9}$$
This can be derived from the condition:
$$L[n] = ((\tau[n] + L[n]/r) - E[n+1]) \times r \tag{4.10}$$
See 4.5.3 on how this represents the expected normal shaping operation of the ATS.

The frame along with its `EligibilityTime` is then passed to the next forwarding stages, e.g. in a bridge. At the egress queue the frame will become available for transmission no earlier than its `EligibilityTime`, i.e., it is transmitted at or after its `EligibilityTime` (see 4.5.5 for the effects on egress bursts).

4.5.3 ATS normal shaping operation

In a correctly configured network (see 4.5.10), without overload conditions, data is expected to arrive roughly at the rate r configured in the shaper for the egress under

consideration. In other words frames of identical length arriving within the window between `BucketEmptyTime` and `BucketFullTime` become eligible at the configured `CommittedInformationRate` in equal intervals.

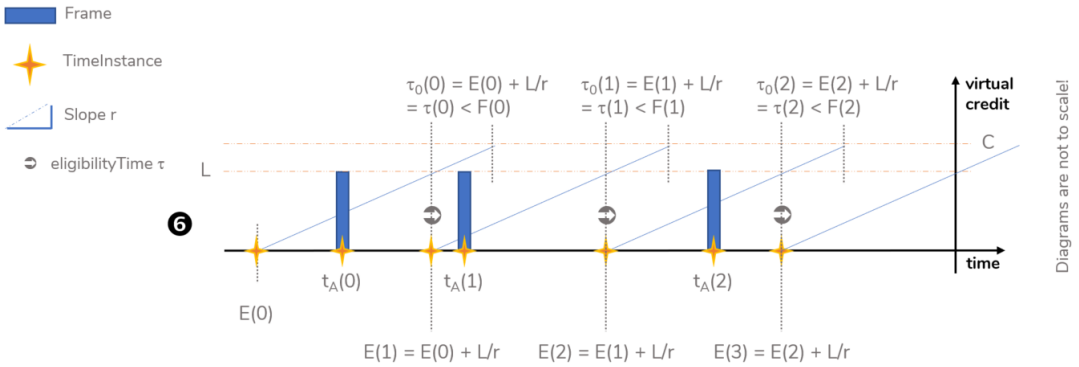

Figure 4.2: ATS normal shaping.

4.5.4 Alternative functional ATS description

The following steps represent a fundamentally equivalent shaping methodology to the algorithm given in IEEE Std 802.1Q-2022 according to Boyer [1]. This does however not handle some special cases, like the discard of a frame due to exceeded `MaxResidenceTime`.

1. Virtual credit increases continuously with `committedInformationRate`
2. Virtual credit can never be larger than `committedBurstSize`
3. A frame becomes eligible for transmission as soon as the virtual credit is equal or larger than the frame length
4. The virtual credit of the shaper instance is reduced by the frame's length at it's `eligibilityTime`, i.e. not at it's actual transmission time

4.5.5 Actual ATS transmission bursts

Assuming 2 frames n_1 and n_2 arrive at roughly the same time $t_A[n_1] \approx t_A[n_2]$, the virtual credit is at `CommittedBurstSize` $c[t_A] = C$, and the `CommittedBurstSize` is at least the sum of the frame lengths $C \geq L[n_1] + L[n_2]$. Then the `EligibilityTime` of the first frame is immediately $\tau[n_1] = t_A[n_1]$ as sufficient credit is available (see (4.2)). Because of $c\,[t_A[n_1]] = C$, (4.5) is true and $E[n_2] = L[n_1]/r + t_A[n_1] - C/r$ according to (4.7). As $C > L[n_1]$ this means $E[n_2] < t_A[n_1]$ and $c\,[t_A[n_2]] \geq L[n_2]$. Therefore the `EligibilityTime` of the second frame is also immediately $\tau[n_2] = t_A[n_2]$ as again sufficient credit is available (see (4.2)). In other words the 2 frames carry basically the

[1] Marc Boyer: "Equivalence between the theoretical model and the standard algorithm of Asynchronous Traffic Shaping.", 2022, hal-03788302

same `EligibilityTime` forward and can be burst from the egress queue controlled by the ATS transmission selection algorithm.

The above represents the case of the intended bursting of the ATS, the following two cases must be considered unintended bursts:

As described in 4.5.7, a single Frame shaped by an ATS instance can be larger than `CommittedBurstSize` and it is therefore possible for a queue controlled by the ATS transmission selection algorithm to transmit a frame which is larger than the configured `CommittedBurstSize`. In 4.5.10 it is pointed out, why this could be considered a misconfiguration or an overload condition.

Furthermore, if transmission from the queue controlled by the ATS transmission selection algorithm is blocked through interfering traffic (of either higher or lower priority) a burst of undefined length is possible, as all frames with `EligibilityTime`s in the past will be transmitted as soon as the queue in question is selected for transmission for a sufficient amount of time. This is a necessity for the ATS to be `bandwidth-preserving` while being `non-work-conserving` (see 4.3).

4.5.6 Guideline for `MaxResidenceTime` configuration

As already stated in 4.5.2, IEEE Std 802.1Q does not require a certain relation between L, C, and Δ_{max}. Both (4.2) and (4.4) allow for $L > C$ (see 4.5.7 for details). On the other hand if $L > \Delta_{max} \times r$ the frame will be discarded, even if $L < C$. The later must obviously be viewed as a misconfiguration.

To find a guideline on how to configure `MaxResidenceTime` the `GroupEligibilityTime` can be ignored. For simplicity all frames are assumed to be of equal length $L \leq C$.

In order for (4.4) to be true, `schedulerEligibilityTime` must be greater than `arrivalTime` $t_A < \tau_0$ according to equation (4.3). Therefore $\tau = \tau_0$ and the frame is discarded if:

$$t_A[n] + \Delta_{max} \leq E[n] + L/r \tag{4.11}$$

$$t_A[n] - E[n] \leq L/r - \Delta_{max} \tag{4.12}$$

$t_A[n] - E[n]$ is positive, if frame n was received after $E[n] = E[n-1] + L/r$, if the previous frame $[n-1]$ was scheduled regularly. In any sensible configuration $L/r < \Delta_{max}$ must be true, although this is not mandated by IEEE (see above). Then $L/r - \Delta_{max} < 0$ and (4.12) trivially becomes false, i.e., the frame is forwarded and not discarded.

If $t_A[n] - E[n] \leq 0$, frame n has arrived at or before the `EligibilityTime` $\tau[n-1]$ of the previous frame $[n-1]$.

$$(E[n-1] + L/r) - t_A[n] \geq \Delta_{max} - L/r \tag{4.13}$$

Obviously $t_A[n-1] < t_A[n]$ and therefore $t_A[n] - E[n-1] \geq 0$.

$$t_A[n] - E[n-1] + \Delta_{max} \leq 2 \times L/r \tag{4.14}$$

For the most extreme case $E[n-1] = t_A[n-1]$ and $t_A[n] \approx t_A[n-1]$ yields $t_A[n] - E[n-1] \approx 0$ and therefore frames are discarded if:

$$\Delta_{max}/2 < L/r \tag{4.15}$$

Taking into account the `CommittedBurstSize` C as the upper limit for the frame's length, the conclusion is that frames are being forwarded while:

$$\Delta_{max} \geq 2 \times C/r \tag{4.16}$$

Giving a lower configuration constraint for `MaxResidenceTime`.

4.5.7 ATS implicit frame length limit

Assuming $t_A[n] - E[n] = \epsilon \leq 0$, what is the maximum frame length L_{max} that can still be scheduled without the frame being discarded due to (4.12)?

$$\epsilon \geq L/r - \Delta_{max} \tag{4.17}$$

Since $L/r < \Delta_{max}$ must be true, the maximum frame length results for $\epsilon = 0$:

$$L_{max} = \Delta_{max} \times r \tag{4.18}$$

Therefore it is possible for $L > C$ and the ATS can create a single frame burst which is larger than the configured `CommittedBurstSize` C. Using (4.16) this yields $L_{max} \geq 2 \times C$.

Since this frame consumes more virtual credit than `committedBurstSize`, the `BucketEmptyTime` of the following frame $E[n+1]$ is delayed by a time equivalent to the excessively consumed credit $E[n+1] = F[n] + (L_{max} - C)/r$. If the larger frame length is accounted for by the configured shaping rate this does not cause any issues. If the larger frame length represents excess data beyond the configured `Committed-InformationRate` the ATS will no longer be `bandwidth-preserving`, i.e., it can not catch up and data loss will occur at some later time.

In 4.5.10 the maximum SDU size filter is therefore required to be applied to all streams going through an ATS instance.

4.5.8 ATS implicit rate limit

According to (4.13) a frame n of length L_{max} from (4.18) is discarded if:

$$E[n-1] + L[n-1]/r \geq t_A[n] \tag{4.19}$$

I.e., the maximum length frame must arrive before the `schedulerEligibilityTime` $\tau_0[n-1]$ of the previous frame. Assuming the previous frame was of `CommittedBurstSize` $L[n-1] = C$, an ingress burst of $C + L_{max}$ occurs within `emptyToFullDuration`. The momentary ingress rate leading to the discard of the maximum length frame from (4.18) is therefore:

$$r_{max} = \frac{C + L_{max}}{C/r} = r \times \left(1 + \frac{\Delta_{max} \times r}{C}\right) \tag{4.20}$$

Using (4.16) this yields:

$$r_{max} \geq 3 \times r \tag{4.21}$$

For such an overload condition the ATS will no longer be `bandwidth-preserving`, i.e., it can not catch up and data loss will occur at some later time. In 4.5.10 the ingress rate to the ATS instance is therefore required to be limited for all streams.

4.5.9 Consequences of the `groupEligibilityTime`

If two or more ATS instances are assigned to the same `SchedulerGroupInstanceID`, they share a common `groupEligibilityTime`. According to (4.3) the `EligibilityTime` of a frame being processed can never be earlier than the `groupEligibilityTime`, which is set to the `EligibilityTime` of the frame most recently processed by any one of the ATS instances. In effect frames which are processed by any ATS instance within the same `SchedulerGroupInstanceID` will be assigned consecutive `EligibilityTime`s in the same order in which the frames have been received.

If the `CommittedInformationRate`s of two ATS instances (x and y) within the same `SchedulerGroupInstanceID` are significantly different ($r[x] >> r[y]$) the `groupEligibilityTime` prevents a frame in ATS instance x to be transmitted before the previously processed frame in ATS instance y, even if the higher `CommittedInformationRate` $r[x]$ would require this.

In order to prevent such a dysfunctional network configuration, ATS must be seen as a network wide concept (see 4.5.10) and can not be applied just locally to any one switch.

This is why in Q there is no Groups in End Stations, in order to align the packets correctly at the start!

4.5.10 Guidance on the configuration

In the ATS' ground state the virtual credit is at the `CommittedBurstSize`. Therefore more than one frame can be set eligible for immediate transmission (see 4.5.5). This is different from the Credit Based Shaper (CBS see 4.6), where at its ground state the credit is zero and only a single frame can be transmitted immediately.

Even more importantly the ATS concept can shape different streams in individual ATS instances while still transmitting all frames from a single egress queue, i.e., with the same priority. Looking at IEEE Std 802.1BA-2021 the latency of frames shaped by CBS stems from the fact that all streams (S_i) of the same priority (TC in IEEE Std 802.1Q-2022) arriving on different ingress ports (P_i) with individual rates (R_i) are shaped in a single CBS instance with $r_{CBS} = \sum_i R_i$. Thus a frame from a high rate stream may still be scheduled last. In order to prevent this Specht [2] requires the following ATS instance allocation rule [3]:

SAR1: Two streams may not be assigned to the same ATS instance, if they have been received on two different ingress ports.

The strict priority transmission selection of IEEE Std 802.1Q-2022 requires in 6.5.3 for frames to only change sequence - with respect to the ingress sequence - on egress, if they are in different priorities. While technically this can be achieved through the `groupEligibilityTime` across multiple ATS instances it leads to unacceptable delays as outlined in 4.5.9, leading to another ATS instance allocation rule:

SAR2: Two streams may not be assigned to the same ATS instance, if they have been received on the same ingress port with different priorities.

As is pointed out in 4.5.7 and 4.5.8 the ATS is not `bandwidth-preserving` if the ingress does not match the configured egress shaping. Large frames and bursts of frames at the ingress can lead to the loss of correctly sized and scheduled frames. In order to discard the frames which violate the shaper configuration and not some random frames later, it is important to:

QoS1: Any stream which is to be processed by an ATS instance must go through a max. SDU size filter. Any uncertainty about the SDU size and media-dependent overhead must be accounted for in the configuration of the ATS instance. The maximum feasible values for `CommittedInformationRate` and `CommittedBurstSize` should be chosen to prevent frame loss.

QoS2: Any stream which is to be processed by an ATS instance must go through a flow meter. Any uncertainty about the ingress rate must be accounted for in the configuration of the ATS instance. The maximum feasible values for `CommittedInformationRate` and `CommittedBurstSize` should be chosen to prevent frame loss.

The shaping operation assigns `EligibilityTime`s to the frames according to the configured `CommittedInformationRate`. If the order of frames which have gone through the same ATS instance would be changed before transmission, this shaping would potentially become ineffective. Therefore:

[2] J. Specht and S. Samii, "Urgency-Based Scheduler for Time-Sensitive Switched Ethernet Networks," in 2016 28th Euromicro Conference on Real-Time Systems (ECRTS), Toulouse, France, Jul. 2016, pp. 75 to 85. doi: 10.1109/ECRTS.2016.27.

[3] The rules have been reworded to adjust for terminology and are referred as Shaper instance Allocation Rules (SAR) here.

SAR3: Two streams may not be assigned to the same ATS instance, if they will be transmitted on the same egress port with different priorities.

4.5.11 ATS Configuration parameters

Configuration parameters of the ATS Eligibility Time Assignment:

IEEE Terminology	AR ECU Parameter	Mathematical Symbol
Scheduler Instance Table	EthSwtAtsInstanceTable	
	EthSwtAtsInstanceEntry	
`CommittedBurstSize`	EthSwtPortATSCommitedBurstSize	C
`CommittedInformationRate`	EthSwtPortATSCommitedInformationRate	r
`SchedulerGroupInstanceID`	EthSwtPortAtsSchedulerGroupRef	
Scheduler Group Instance Table	EthSwtAtsGroupInstanceTable	
	EthSwtAtsGroupInstanceEntry	
`MaxResidenceTime`	EthSwtAtsGroupMaximumResidenceTime	Δ_{max}

Table 4.2: Terms, configuration paramteres, and mathematical symbols to describe and configure the ATS

Configuration parameters of the ATS transmission selection algorithm: There are no specific parameters to configure the ATS transmission selection algorithm, but the ATS transmission selection algorithm needs to be configured for the traffic class used to transmit the frames intended to be shaped using ATS.

4.6 Credit Based Shaper

4.6.1 Overview

The Credit Based Shaper (CBS)

is

was orginally specified by [10, IEEE Std 802.Qav] and provides a traffic shaping concept for high priority, high data rate streams. It was specified as part of the Audio Video Bridging (AVB) standards. The behavior is similar to ATS, but the CBS is operating only on Traffic Classes on the egress side of a bridge. Therefore, there is no differentiation between individual streams within the traffic class as well as different ingress ports. As a consequence, the Credit Based Shaper is strictly operating on the order frames are received on all ports, there is no reordering of frames received on different ports. The Credit Based Shaper minimizes bursts and therefore protects lower priority traffic from the high priority CBS shaped traffic.

The CBS only exists as one functional block (per traffic class) that is located on the egress side of a bridge port and makes frames available for transmission selection based on a credit value. If there is zero or positive credit, frames are available for transmission, in case there is negative credit, frames of the associated traffic class are not available for transmission selection. The `idleSlope` is the

only

parameter that defines the rates with which credit is accumulated and reduced, and thereby it directly controls the data rate of the assigned traffic class on a port.

4.6.2 CBS Configuration parameters

Configuration parameters of the Credit Based Shaper:

- idleSlope (EthSwtPortEgressQueueCreditBasedShaperIdleSlope) [ECUC_EthSwt_00188] in [13]

4.6.3 CBS credit estimation

In the IEEE mode, the credit can theoretically increase without limit. Obviously this is not possible in any practical implementation. While not being configurable

- `highCredit` and
- `lowCredit`

give useful indications of the credit levels to be expected.

- `highCredit` is calculated based on the maximum amount of interference by traffic from higher priority traffic classes and the `idleSlope` configuration parameter.
- `lowCredit` can be derived from the maximum frame size expected in the traffic class and the `idleSlope` configuration parameter.

Since accurate numbers for these calculations may not always be known, these values give a rough lower bound for the credit levels only.

4.6.4 Guidance on the configuration

1. Frames associated with the CBS are identified by the priority. In case there are ports of the bridge that use the priority intended to be used for CBS traffic for other traffic, the priorities on those ports need to be remapped by configuring the priority regeneration table.
2. The `idleSlope` of the CBS is set according to the average bit rate of all streams sent with this traffic class on that specific port. A little bit of oversubscription is necessary to account for frequency differences of the oscillators of the different bridges in the network. Credit increases with the idleSlope and is reduced with sendSlope (sendSlope = idleSlope - portTransmitRate).
3. To use the CBS algorithm the traffic class(es) used for CBS need to be configured to use the CBS transmission selection. Therefore, `EthSwtPortEgressQueueTransmissionSelectionAlgorithm` needs to be set to `ETHSWT_TRANSMISSION_SELECTION_ALGORITHM_CBS`

   [SWS_EthSwt_00542] in [13]

   .

4.7 Ensure Accuracy of Time Synchronization

4.7.1 Recommendations to support accuracy in an Ethernet network

4.7.1.1 Parts in Documents which Reflect Accuracy

In Time Synchronization Protocol Specification [14, PRS Time Synchronization over Ethernet]:

Time Master and Time Slave shall work with a Time Base reference clock accuracy as defined in [4, IEEE 802.1 AS], ANNEX B.1.2 "Time measurement granularity".

In Specification of Time Synchronization over Ethernet [15]:

Automotive systems requiring a common Time Base for ECUs regardless of which bus system the ECUs are connected to.

4.7.1.2 System requirements

TSN uses gPTP protocol (IEEE 802.1AS) to get precise timestamps for the nodes in the Network. IEEE 802.1AS also targets at accuracy of +/-500ns between any two nodes over 7 hops for Audio Video Bridging using 100 Mb/s links. That means, the hardware has to support the global time on Ethernet with an end to end accuracy of +/- 500ns, where end to end describes a network path from a data producer to a data consumer (e.g. from audio source to an audio sink).

4.8 Measurement of the Time Synchronization Quality (Accuracy)

4.8.1 Overview

The quality of time synchronization needs to be measured in order to assess if devices meet requirements for gPTP recovered clock quality.

The traditional approach to measure the accuracy of the clock synchronization is the 1 Pulse Per Second (1PPS) method. The PPS signal must be driven by a hardware clock in all participating devices so that it generates a pulse every full second. If this is done for multiple hardware clocks in a distributed system their synchronization accuracy can be checked on a measurement device like a scope by measuring the distance of PPS pulses originating from different devices. After their clocks have been correctly synchronized, the PPS pulses triggered every second should only have a marginal time distance in a successful scenario.

However, the 1PPS method is not feasible for automotive environment, where using a separate output pin for generating the 1PPS signal on the ECU is considered costly. Furthermore, static offset in the internal clock of a slave may exist and can be canceled out by an opposite static offset in the logic that drives the 1PPS output. This will result in an error condition that cannot be observable by using 1PPS measurement method.

In order to mitigate the flaws from 1PPS measurement, Advance methods like Ingress Reporting Method and Reverse Sync Method may be considered for alternative measurement methods. More details for Ingress Reporting Method and Reverse Sync Method, please refer to [16, 802.1AS Recovered Clock Quality Testing].

4.8.2 Usage of PPS

4.8.2.1 Configuration of PPS within an Ethernet controller

AUTOSAR support to configure the generation of a PPS signal, if the hardware supports this feature. AUTOSAR support the following configuration options:

- support to configure the rate of a PPS (e.g. 1Hz, 2 Hz, 4 Hz ...)
- support to configure the duty cycle of a PPS (e.g. 50% == 500ms high an 500ms low if rate is configured for 1Hz)

- support to configure the starting point (either rising edge and falling edge) of a PPS (Please consider this point for the SysT)

4.8.2.2 Configuration of PPS by using OCU in combination with a PHC

If OCU is hardware driven, then a PPS could be generated as alternative to an existing PPS hardware unit in the Ethernet controller. Please note: It is recommended to use a PPS hardware unit, since the OCU approach could introduce inaccuracy.