

1 Technical Descriptions for
2 Cut-Through Forwarding in Bridges

3 DCN 1-22-0042-09-ICne

4 Author: Johannes Specht

5 November 3, 2022

6	Contents	
7	I. Introduction	8
8	1. Purpose	9
9	2. Relationship to IEEE Standards	10
10	3. Status of this Document	11
11	II. Cut-Through Forwarding in Bridges	12
12	4. Overview and Architecture	13
13	5. Modeling Principles	15
14	5.1. Frame Types	15
15	5.2. Data Resolutions	15
16	5.2.1. Bit-Accurate Modeling	15
17	5.2.2. Parameter-Accurate Modeling	16
18	5.3. Temporal Control	17
19	5.3.1. Processing Stalls	17
20	5.3.2. Late errors	17
21	5.3.3. Fall-backs to S&F	17
22	5.3.4. Instantaneous Operations	18
23	6. Generalized Serial Convergence Operations	19
24	6.1. Overview	19
25	6.2. Service Primitives	21
26	6.2.1. M_DATA.indication and M_DATA.request	21
27	6.2.1.1. DA	21
28	6.2.1.2. SA	21
29	6.2.1.3. SDU	21
30	6.2.1.4. FCS	21
31	6.2.2. M_UNITDATA.indication and M_UNITDATA.request	21
32	6.3. Global Constants	22
33	6.3.1. PREAMBLE	22
34	6.3.2. LEN_OCT	22
35	6.3.3. LEN_ADDR	22
36	6.3.4. LEN_FCS	23

37	6.3.5.	LEN_MIN	23
38	6.3.6.	LEN_MAX	23
39	6.3.7.	LEN_DATA	23
40	6.4.	Global Variables	23
41	6.4.1.	RxBitEnable	23
42	6.4.2.	RxBit	23
43	6.4.3.	RxBitStatus	24
44	6.4.4.	RxDataEnable	24
45	6.4.5.	RxData	24
46	6.4.6.	RxDataStatus	24
47	6.4.7.	TxBitEnable	25
48	6.4.8.	TxBit	25
49	6.4.9.	TxBitStatus	25
50	6.4.10.	TxDataEnable	25
51	6.4.11.	TxData	25
52	6.4.12.	TxDataStatus	25
53	6.5.	Generic Data Receive process	26
54	6.6.	Generic Frame Receive process	26
55	6.6.1.	Description	26
56	6.6.2.	State Machine Diagram	26
57	6.6.3.	Variables	26
58	6.6.3.1.	cnt	26
59	6.6.3.2.	len	26
60	6.6.3.3.	status	26
61	6.6.4.	Functions	28
62	6.6.4.1.	append(parameter,bit)	28
63	6.6.4.2.	FCSValid(FCS)	28
64	6.7.	Receive Convergence process	28
65	6.8.	Generic Data Transmit process	28
66	6.9.	Generic Frame Transmit process	28
67	6.9.1.	Description	28
68	6.9.2.	State Machine Diagram	28
69	6.9.3.	Variables	30
70	6.9.3.1.	cnt	30
71	6.10.	Transmit Convergence process	30
72	7.	Bridge Port Transmit and Receive Operations	31
73	7.1.	Overview	31
74	7.2.	Bridge Port Connectivity	32
75	7.3.	Priority Signaling	32
76	7.3.1.	Receive path operations	32
77	7.3.2.	Transmit path operations	33
78	7.4.	Translations between Internal Sublayer Service (ISS) and Enhanced In-	
79		ternal Sublayer Service (EISS)	33
80	7.4.1.	Receive path operations	33

81	7.4.2. Transmit path operations	34
82	7.5. Higher Layer Compatibility	34
83	7.6. CTF Sublayer	34
84	7.6.1. Receive Path Operations	34
85	7.6.2. Transmit Path Operations	35
86	7.6.3. Inconsistent frame handling	35
87	8. Bridge Relay Operations	36
88	8.1. Overview	36
89	8.2. Active Topology Enforcement	38
90	8.2.1. Overview	38
91	8.2.2. Learning	38
92	8.2.3. Initial set of potential transmission Ports	38
93	8.3. Ingress Filtering	39
94	8.4. Frame Filtering	39
95	8.5. Egress Filtering	39
96	8.6. Flow Classification and Metering	40
97	8.6.1. General	40
98	8.6.2. Stream Filtering	41
99	8.6.3. Maximum SDU size filtering	41
100	8.6.4. Stream Gating	41
101	8.6.5. Flow Metering	41
102	8.7. Individual Recovery	42
103	8.8. Sequence Recovery	42
104	8.9. Sequence Encode	42
105	8.10. Active Stream Identification	42
106	8.11. Queuing Frames	43
107	8.12. Queue Management	43
108	8.13. Transmission Selection	44
109	9. Management Parameters	45
110	9.1. Overview	45
111	9.2. Control Parameters	45
112	9.2.1. CTFTransmissionSupported	45
113	9.2.2. CTFTransmissionEnable	45
114	9.2.3. CTFReceptionSupported	46
115	9.2.4. CTFReceptionEnable	46
116	9.3. Timing Parameters	46
117	9.3.1. CTFDelayMin and CTFDelayMax	46
118	9.4. Error Counters	46
119	9.4.1. CTFReceptionDiscoveredErrors	46
120	9.4.2. CTFReceptionUndiscoveredErrors	47
121	III. Cut-Through Forwarding in Bridged Networks	48

122	IV. Appendices	50
123	A. Interaction of the Lower Layer Interface (LLI) with existing Lower Layers	51
124	A.1. PLS Service Interface	51
125	A.1.1. Overview	51
126	A.1.2. Service Primitives	51
127	A.1.3. Global Variables and Constants	52
128	A.1.3.1. BitTick	52
129	A.1.3.2. LEN_FRAMEGAP	52
130	A.1.4. Global Constraints	52
131	A.1.5. Transmit Bit Clock process	52
132	A.1.6. PLS Transmit process	52
133	A.1.6.1. Description	52
134	A.1.6.2. State Machine Diagram	53
135	A.1.6.3. Variables	53
136	A.1.7. PLS Receive process	53
137	A.1.7.1. Description	53
138	A.1.7.2. State Machine Diagram	54
139	A.1.7.3. Variables	54
140	Bibliography	54

141 List of Figures

142	4.1. Architecture of a Cut-Through Forwarding (CTF) Bridge.	13
143	6.1. Overview of the generalized serial convergence operations.	19
144	6.2. State Machine Diagram of the Generic Frame Receive process.	27
145	6.3. State Machine Diagram of the Generic Frame Transmit process.	29
146	7.1. Bridge Port Transmit and Receive (VLAN-unaware).	31
147	7.2. Bridge Port Transmit and Receive (VLAN-aware).	32
148	8.1. Forwarding process of a CTF bridge.	37
149	8.2. Flow classification and metering.	40
150	A.1. Processes and interactions for interfacing between LLI and PLS service	
151	primitives.	51
152	A.2. State machine diagram of the PLS Transmit process.	53
153	A.3. State machine diagram of the PLS Receive process.	54

154 List of Algorithms

155	6.1. Signature of the M_DATA.indication service primitive.	21
156	6.2. Signature of the M_DATA.request service primitive.	21
157	6.3. Signature of the M_UNITDATA.indication service primitive.	22
158	6.4. Signature of the M_UNITDATA.request service primitive.	22
159	6.5. Definition of data type low_data_t.	24
160	8.1. Queuing rules for frames under reception.	43

161

Part I.

162

Introduction

1. Purpose

Purpose of this document is to provide input for technical discussion in pre-PAR activities of IEEE 802 (i.e., Nendica). The contents of this document are technical descriptions for the operations of Cut-Through Forwarding (CTF) in bridges. The intent is to provide more technical clarity, demonstrate technical feasibility, and thereby also address the desire expressed by individuals during the IEEE 802.1 closing plenary meeting in July 2022 to a certain extent.

170 2. Relationship to IEEE Standards

171 This document **IS NOT** an IEEE Standard or an IEEE Standards draft, it is an
 172 individual contribution by the author containing technical descriptions. This allows
 173 readers to focus on the technical contents in this document, rather than additional
 174 aspects that are important during standards development. For example:

- 175 1. The structure of this document does not comply with the structural requirements
 176 for such standards (e.g., this document does not contain mandatory clauses for
 177 IEEE Standards [1]).
- 178 2. Usage of normative keywords has no implied semantics beyond technical lan-
 179 guage. For example, usage of the words *shall*, *should* or *may* **DOES NOT**
 180 imply conformance requirements or recommendations of implementations.
- 181 3. This document contains references, but without distinguishing between norma-
 182 tive and informative references.
- 183 4. This document does not contain suggestions for assigning particular contents
 184 to *vehicles* (e.g., IEEE 802 Working Groups, potential amendment projects for
 185 existing standards, or potential new standard projects). As a consequence, the
 186 clause structure of this document is intended for readability, rather than fitting
 187 into the clause structure of a particular Standard (which would especially matter
 188 for potential amendment projects).

189 3. Status of this Document

190 This document is work-in-progress. It contains technical and editorial errors, omis-
191 sions and simplifications. Readers discovering such issues are encouraged for making
192 enhancement proposals, e.g. by proposing textual changes or additions to the author
193 (johannes.specht.standards@gmail.com).

194

Part II.

195

Cut-Through Forwarding in Bridges

196

4. Overview and Architecture

This part of the document comprises technical descriptions for supporting CTF in bridges. While this document is not a standard, there are published IEEE 802.1 Standards describing the operation of bridges without the descriptions herein. For differentiation between bridges with support for CTF and bridges according to the published IEEE 802.1 Standards (e.g., IEEE Std 802.1Q[2]), term *CTF bridge* is used in this document to refer to the former, whereas term *S&F bridge* is used in this document to refer to the latter. Like in IEEE Std 802.1Q, CTF bridges may or may not support Virtual Local Area Networks (VLANs), and therefore terms *VLAN-aware* and *VLAN-unaware* are used to distinguish between bridges with and without support for VLANs.

The architecture of CTF bridges is widely aligned with the bridge architecture in IEEE Std 802.1Q [2, 8.2]. It is shown in Figure 4.1 (itself likewise aligned with the architectural figures in IEEE Std 802.1Q [2, Figure 8-2, 8-3, 8-4, ff.]) in a compact form.




Figure 4.1.: Architecture of a Cut-Through Forwarding (CTF) Bridge.

This architecture comprises the following elements:

1. One or more higher layer entities above the MAC Service (MS) interface.

- 214 2. A bridge relay entity (8) that relays frames between different bridge Ports.
- 215 3. Generalized serial convergence operations (6) that translate between the Internal
216 Sublayer Service (ISS) interface and Lower Layer Interface (LLI) per bridge Port.
- 217 4. Bridge Port transmit and receive operations (7) per Bridge port that transform
218 and transfer service primitive invocations between the bridge relay entity, higher
219 layer entities and the generalized serial convergence operations.
- 220 The operation of CTF bridges is described in this document in the chapters referred
221 to before, typically limiting on describing the additions and potential differences to
222 the operations of S&F bridges.
- 223
- 224 Excluded from this document are several details on higher layer entities¹ above the
225 MAC Service interface and elements of the bridge relay entity other than the forwarding
226 process²:
- 227 – For frames to and from higher layer entities, the bridge port transmit and receive
228 operations of a CTF bridge establish the behavior of S&F bridge at the MAC
229 service interface (7.2), allowing higher layer entities to operate according to the
230 behavior specified in IEEE 802.1 Standards unaltered.
- 231 – The forwarding process of a CTF bridges (re-)establishes the behavior of S&F
232 bridges at interaction points with other elements of the bridge relay entity.
- 233 In general, this part of the document limits the use of Cut-Through to operations
234 standardized in IEEE Stds 802.1Q[2], 802.1AC[3] and 802.1CB[4].

¹ Examples for higher layer entities are Spanning Tree Protocols and Multiple Registration Protocols, supported by LLC entities above the MAC service interface [2, item c) in 8.2 and b) in 8.3].

² An example element of the bridge relay entity other than the forwarding process is the learning process [2, item c) in 8.2 and b) in 8.3].

235 5. Modeling Principles

236 5.1. Frame Types

237 If necessary, distinct terms are used for frames for describing their current state,
238 as follows:

239 **frame under reception** A frame that is being serially received from a LAN's physical
240 medium for which reception began bit did not finish.

241 **received frame** A frame that was serially received from a LAN's physical medium that
242 finished.

243 **frame under transmission** A frame that is being serially transmitted to a LAN's phys-
244 ical medium for which transmission began bit did not finish.

245 **transmitted frame** A frame that was serially transmitted to a LAN's physical medium
246 for which transmission finished.

247 5.2. Data Resolutions

248 5.2.1. Bit-Accurate Modeling

249 All invocations of service primitives in this document are atomic. That is, each in-
250 vocation is non-dividable (see also 7.2 of IEEE Std 802.1AC[3]). Service primitive
251 invocations are modeled more explicitly in this document, allowing for accurate de-
252 scription of operations within a Bridge, while retaining atomicity. This explicit model
253 comprises the following:

- 254 1. A service primitive provides two attributes¹, *'start* and *'end. These attributes
255 are used in subsequent descriptions to indicate the start and the end of the
256 indication, respectively.*
- 257 2. The parameters of a service primitive are explicitly modeled as bit arrays.
- 258 3. The values of parameters during invocations of a service primitive are passed
259 according to a call-by-reference scheme.

¹The concept of *attributes* is inspired by the *Very High Speed Integrated Circuits Hardware De-
scription Language*, VHDL[5], which provides predefined attributes (e.g., *'transaction*) that allow
modeling over multiple VHDL simulation cycles at the same instant of simulated time.

260 In a series of sequential *processing stages* (e.g., the processes introduced in 6.1 or a
 261 sub-process of the forwarding process in 8), this model allows later processing stages
 262 to access contents in service primitive parameters that are incrementally added by an
 263 earlier processing stage.

264 5.2.2. Parameter-Accurate Modeling

265 At higher levels processing stages, service primitives of frames and processing of these
 266 frames themselves is modeled at parameter level accuracy. The purpose of this model
 267 is to

- 268 1. provide means for compact description of temporal control (5.3) in and across
 269 processing stages,
- 270 2. enable re-use of existing transformation rules from IEEE 802.1 Stds by reference,
 271 and
- 272 3. avoid low level details that would not provide any value to the clarity and un-
 273 ambiguous descriptions.

274 The parameter-accurate operates at the resolution of symbolic and/or numeric param-
 275 eters instead of bit arrays (5.2.1). A parameter is said to be *complete* at the earliest
 276 instant of time at which the *minimal information* is available to *unambiguously* deter-
 277 mine the parameter's value within the specified valid value range of such parameter.
 278 The minimal information may be

- 279 1. a coherent sequence of bits in a frame,
- 280 2. the result of composition and/or computation across bits located at various lo-
 281 cations in a frame,
- 282 3. based on out-of-band information, or
- 283 4. any combination of the aforesaid.

284 As an example, the `vlan_identifier` parameter of `EM_UNITDATA.indication` (7.4)
 285 invocations can be derived from a subset of underlying bits of the associated SDU
 286 parameter of `M_DATA.indication` invocations (6.2.1) that are located in a VLAN Tag
 287 [2, 9.6] according to the specification of the Support for the EISS defined in IEEE Std
 288 802.1Q [2, item e) in 6.9.1] or originate from out-of-band information like a configured
 289 per-Port PVID parameter [2, item d) in 6.9, item f) in 6.9.1 and 12.10.1.2]. If the
 290 VLAN tag is required to unambiguously determine the `vlan_identifier` parameter, the
 291 parameter is complete when all bits of the VID parameter² in the VLAN Tag where
 292 received.
 293

²The bits and potential out-of-band information form the minimal information, and exclude any
 redundant information, most prominently the (in-band) redundant encoding of the VID parameter
 in the frame's FCS parameter.

Most of the data transformations between bits in a frame, frame parameters and potential out-of-band information is already unambiguously specified in the relevant IEEE 802.1 Standards. This document omits repetition of already specified transformations and instead just refers to the relevant data transformations in existing IEEE 802.1 Standards.

5.3. Temporal Control

5.3.1. Processing Stalls

Parameter-accurate modeling allows formulating temporal control in processing stages. A processing stage (5.2.1) may *stall* further processing of a frame under reception, including (but not limited to) passing this frame to a subsequent processing stage, until one or more parameters are complete (5.2.2), subject to the implicit discarding due to late errors (5.3.2). Most processing stalls are given due to the data dependencies already specified in IEEE 802.1 Standards (e.g., Ingress Filtering as part of the forwarding process in IEEE Std 802.1Q[2, 8.6.2] depends on the availability of a frame's VID, which therefore implicitly requires completion of the `vlan_identifier` parameter of `EM_UNITDATA.indication` invocations), however, explicit modeling of processing stalls may be expressed by formulations in natural language.

Example formulations:

1. *“Processing **stalls** pending the **vlan_identifier** parameter.”*
2. *“Further execution in a CTF bridge is **stalled** pending the **destination address** of a frame under reception prior to the filtering database lookup of the destination ports.”*

5.3.2. Late errors

In a sequence of processing stages, an earlier processing stage may discover an error in a frame under reception and then notify all subsequent (not antecedent) processing stages, which may then implement error handling upon this such notification. This is termed as a *late error*, which is raised by the earlier processing stage and associated with a particular frame under reception. If any of the subsequent stage stalls processing pending one or more parameters of the associated frame under reception when the error is raised, the frame is discarded in the subsequent stage and thereby neither further processed nor passed to any other following processing stage.

5.3.3. Fall-backs to S&F

The descriptions of the processing stages use *fall back to S&F* as a modeling shortcut to summarize the following sequence:

1. Processing of a frame under reception stalls pending the frame's end of reception, which is a shortcut by itself for stalling processing pending all parameters of a frame under reception, including the FCS.

- 331 2. Dependent on whether or not a late error was indicated by an earlier processing
332 stage for that frame:
- 333 a) Late error indicated:
334 The frame is discarded prior to any further processing by any stage.
- 335 b) No Late error indicated:
336 The frame continues subsequent processing through subsequent processing
337 stages according to the standardized behavior of an S&F bridge.

338 5.3.4. Instantaneous Operations

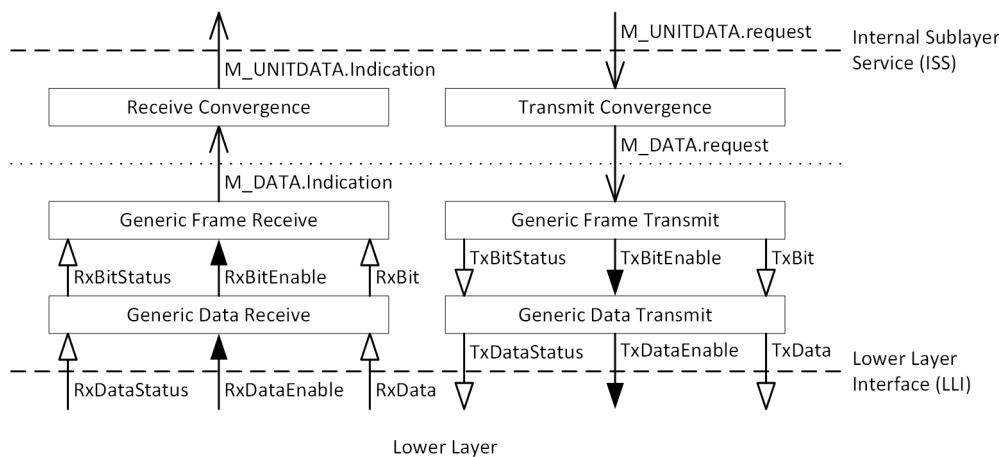
339 In absence of processing stalls, processing stages in this document perform their oper-
340 ations instantaneously. It is clear that instantaneous operations, in terms of 0-delay at
341 an infinite high resolution³, are not possible in real world implementations. Physical
342 constraints and design decisions introduce additional delays in such implementations.
343 The model is not intended to upper limit such delays. It is there for describing data
344 dependencies, late error handling and the resulting externally visible behavior. Ad-
345 ditional delays (e.g., real world implementations starting transmissions on a physical
346 medium later than the model) are not described by the model, but could be determined
347 by observation/measurement and are available as management parameters (9.3).

³The semantics of “instantaneous” can vary dependent on the resolution [6, p.11].

6. Generalized Serial Convergence Operations

6.1. Overview

The generalized serial convergence operations are described by a stack of processes that interact via global variables (see 6.4) and service primitive invocations (see 6.2). These processes provide the translation between the Internal Sublayer Service (ISS) and a broad range of lower layers, including (but not limited to) physical layers. Figure 6.1 provides an overview of these processes and their interaction¹. The processes can



NOTATION

- ▷ : A global variable set solely by the originating process.
- ▶ : A global variable set the originating process and reset by the receiving process.
- ➡ : A service primitive.

Figure 6.1.: Overview of the generalized serial convergence operations.

be summarized as follows:

1. A Receive Convergence process (6.7) that translates each invocation of the `M_DATA`.-

¹This interaction model is inspired by clause 6 and 8.6.9 of IEEE Std 802.1Q[2].

- 358 indication service primitive (6.2.1) into a corresponding invocation of the M_UNIT-
359 DATA.indication service primitive (6.2.2).
- 360 2. A Generic Frame Receive process (6.6) that generates M_DATA.indication in-
361 vocations for bit sequences originating from the Generic Data Receive process of
362 at least LEN_MIN (6.3.5) bits.
- 363 3. A Generic Data Receive process (6.5) that translates a lower layer-dependent²
364 serial data stream into delineated homogeneous bit sequences of variable length,
365 each typically representing a frame.
- 366 4. A Transmit Convergence process (6.10) that translates each invocation of the
367 M_UNITDATA.request service primitive into a corresponding invocation of the
368 M_DATA.request service primitive.
- 369 5. A Generic Frame Transmit process (6.9) that translates M_DATA.request invo-
370 cations into bit sequences for the Generic Data Transmit process.
- 371 6. A Generic Data Transmit process (6.8) that translates bit sequences from the
372 Generic Frame Transmit process into a lower layer-dependent serial data stream.
- 373 The generalized serial convergence operations are heavily inspired by the concepts de-
374 scribed in slides by Roger Marks [7, slide 15], but follow a different modeling approach
375 with more formalized description of these functions and incorporate some of the follow-
376 ing concepts, as suggested by the author of this document during the Nendica meetings
377 on and after August 18, 2022. The differences can be summarized as follows:
- 378 – Alignment with state machine diagram conventions of IEEE Std 802.1Q[2, Annex
379 E].
 - 380 – Support for serial data streams from lower layers with arbitrary data word
381 length³.
 - 382 – Explicit temporal modeling of atomic ISS service primitive invocations.
- 383 By keeping ISS service primitive invocations atomic, the approach in this document
384 is intended to provide a higher level of compatibility with existing IEEE 802.1 Stds,
385 similar to the modeling approach via frame look-ahead of service primitive invocation-
386 s/prescient functions[8, slides 7ff.].

²Such a lower layer may be an entity on the physical layer (PHY), but the generalized receive operations are not limited to this.

³This generalization is intended to allow a wide range of lower layers. In addition, the support for word sizes (e.g., 8 bits, 32 bits or 64 bits) may be close to realities found in hardware implementation. It is subject to discussion whether this and other generalizations over [7] introduced by the author are considered to be helpful.

Algorithm 6.1 Signature of the M_DATA.indication service primitive.

M_DATA.indication(DA, SA, SDU, FCS)

Algorithm 6.2 Signature of the M_DATA.request service primitive.

M_DATA.request(DA, SA, SDU, FCS)

387 6.2. Service Primitives

388 6.2.1. M_DATA.indication and M_DATA.request

389 The M_DATA.indication service primitive passes the contents of a frame from the
 390 Generic Frame Receive process to the Receive Convergence process. The M_DATA.-
 391 request service primitive passes the contents of a frame from the Transmit Convergence
 392 process to the Generic Frame Transmit process. The parameter signatures of the
 393 service primitives are as shown in Algorithm 6.1 and Algorithm 6.2⁴.

394 The parameters are defined as follows:

395 6.2.1.1. DA

396 An array of zero to LEN_ADDR (6.3.3) bits, containing the destination address of a
 397 frame.

398 6.2.1.2. SA

399 An array of zero to LEN_ADDR (6.3.3) bits, containing the source address of a frame.

400 6.2.1.3. SDU

401 An array of zero or more bits, containing a service data unit of a frame. The number
 402 of bits after complete reception of a frame is an integer multiple LEN_OCT (6.3.2).

403 6.2.1.4. FCS

404 An array of zero to LEN_FCS (6.3.4) bits, containing the frame check sequence of a
 405 frame.

406 6.2.2. M_UNITDATA.indication and M_UNITDATA.request

407 As specified in IEEE Std 802.1AC[3, 11.1], with the identical parameter signatures as
 408 shown in Algorithm 6.3 and Algorithm 6.4.

⁴The parameters in this version of this document limit to those introduced in Roger Marks' GSCF slides [7]. Future versions may introduce more flexibility (e.g., for IEEE Std 802.11 [9, 9.2]).

Algorithm 6.3 Signature of the M_UNITDATA.indication service primitive.

```

M_UNITDATA.indication(
    destination_address,
    source_address,
    mac_service_data_unit,
    priority, drop_eligible,
    frame_check_sequence,
    service_access_point_identifier,
    connection_identifier
)

```

Algorithm 6.4 Signature of the M_UNITDATA.request service primitive.

```

M_UNITDATA.request(
    destination_address,
    source_address,
    mac_service_data_unit,
    priority, drop_eligible,
    frame_check_sequence,
    service_access_point_identifier,
    connection_identifier
)

```

409 6.3. Global Constants

410 6.3.1. PREAMBLE

411 A lower layer-dependent array of zero⁵ or more bits, containing the expected preamble
 412 of each frame.

413 6.3.2. LEN_OCT

414 The integer number eight (8), indicating the number of bits per octet.

415 6.3.3. LEN_ADDR

416 An integer denoting the length of the DA and SA parameters of M_DATA.indication
 417 parameters, in bits. For example,

$$\text{LEN_ADDR} = 48 \quad (6.1)$$

418 indicates an EUI-48 addresses.

⁵Including length zero permits to support lower layers that do not expose a preamble to the Generic Data Receive process.

419 **6.3.4. LEN_FCS**

420 An integer denoting the length of frame check sequence and the length FCS parameter
421 of M_DATA.indication parameter, respectively, in bits. For example,

$$\text{LEN_FCS} = 32 \quad (6.2)$$

422 indicates a four octet frame check sequence.

423 **6.3.5. LEN_MIN**

424 A lower layer-dependent integer, denoting the minimum length of a frame, in bits.
425 Invocation of the M_DATA.indication service primitive starts once the Generic Frame
426 Receive process received the first LEN_MIN bits of a frame. Values for LEN_MIN
427 with

$$\text{LEN_MIN} \geq \text{PREAMBLE.length} + \text{LEN_FCS} \quad (6.3)$$

428 are valid.

429 **6.3.6. LEN_MAX**

430 A lower layer-dependent integer, denoting the maximum length of a frame, in bits. In-
431 vocation of the M_DATA.indication service primitive ends at latest once the Generic
432 Frame Receive process received at most LEN_MAX bits of a frame. Values for
433 LEN_MIN with

$$\text{LEN_MAX} \geq \text{PREAMBLE.length} + 2\text{LEN_ADDR} + \text{LEN_FCS} \quad (6.4)$$

434 are valid.

435 **6.3.7. LEN_DATA**

436 A lower layer-dependent integer, denoting the data width of the RxData and TxData
437 variables, in bits.

438 **6.4. Global Variables**

439 **6.4.1. RxBitEnable**

440 A Boolean variable, set by the Generic Data Receive process and reset by the Generic
441 Frame Receive process, which indicates an update of the RxBit variable, RxBitStatus
442 variable, or both.

443 **6.4.2. RxBit**

444 A bit variable used to pass a single bit value to the Generic Frame Receive process.

Algorithm 6.5 Definition of data type `low_data_t`.

```
typedef struct {
    Boolean start;
    Boolean end;
    bit[] value;
} low_data_t;
```

445 **6.4.3. RxBitStatus**

446 An enumeration variable used to pass the receive status from the Generic Data Receive
447 process to the Generic Frame Receive process. The valid enumeration literals are as
448 follows:

449 **RECEIVING** Indicates that the Generic Data Receive process received data from lower
450 layers in a serial stream without knowledge of the remaining length of the overall
451 data stream.

452 **TRAILER** Indicates that the Generic Data Receive process received data from lower
453 layers in a serial stream with the knowledge that LEN_FCS or less bits follow.

454 **6.4.4. RxDataEnable**

455 A Boolean variable, set by a lower layer and reset by the Generic Data Receive process,
456 which indicates an update of the RxData variable, RxDataStatus variable, or both.

457 **6.4.5. RxData**

458 A variable of composite data type `low_data_t`, used for serially passing data words of
459 frames from a lower layer to the Generic Data Receive process. Type `low_data_t` is
460 defined in Listing 6.5. The semantics of the constituent parameters is as follows:

461 **start** Indicates whether the data word is the first word of a frame (TRUE) or not
462 (FALSE).

463 **end** Indicates whether the data word is the last word of a frame (TRUE) or not
464 (FALSE).

465 **value** A lower layer-dependent non-empty array of up to LEN_DATA (6.3.7) bits,
466 containing a data word of a frame. An array length less than LEN_DATA bits
467 is only valid if end is TRUE.

468 **6.4.6. RxDataStatus**

469 An enumeration variable used to pass the receive status from lower layers to the Generic
470 Data Receive process. The valid enumeration literals are as follows:

471 **IDLE** Indicates that data stream reception from lower layers is not active.

472 **RECEIVING** Indicates that data stream reception from lower layers is active.

473 **6.4.7. TxBitEnable**

474 A Boolean variable, set by the Generic Frame Transmit process and reset by the
475 Generic Data Transmit process, which indicates an update of the TxBit variable.

476 **6.4.8. TxBit**

477 A bit variable used to pass a single bit value of a frame's bit stream to the Generic
478 Data Transmit process.

479 **6.4.9. TxBitStatus**

480 An enumeration variable that indicates the transmission state from the Generic Frame
481 Transmit process to the Generic Data Transmit process. The valid enumeration literals
482 are as follows:

483 **IDLE** Indicates that the Generic Frame Transmit process is not generating the bit
484 stream of a frame.

485 **TRANSMITTING** Indicates that the Generic Frame Transmit process is generating
486 the bit stream of a frame.

487 **6.4.10. TxDataEnable**

488 A Boolean variable, set by the Generic Data Transmit process a lower layer and reset
489 by the lower layer, which indicates an update of the TxData variable.

490 **6.4.11. TxData**

491 A variable of composite datatype `low_data_t` (6.5), used for serially passing data
492 words of frames from the Generic Data Transmit process to a lower layer.

493 **6.4.12. TxDataStatus**

494 An enumeration variable that indicates the transmission state from the Generic Data
495 Transmit process to the lower layer. The valid enumeration literals are as follows:

496 **IDLE** Indicates that the Generic Data Transmit process is not generating the data
497 stream of a frame.

498 **TRANSMITTING** Indicates that the Generic Data Transmit process is generating the
499 data stream of a frame.

500 6.5. Generic Data Receive process

501 The Generic Data Receive process translates a lower layer-dependent⁶ serial data
502 stream into a uniform bit stream. In addition, it realizes the following functions:

- 503 – Determine the position in the serial data stream of a frame at which the frame
504 check sequence begins (delay line modeling).
- 505 – Truncate excess bits to satisfy the frame length requirements implied by the
506 parameter definition of the M_DATA.indication primitive (6.2.1).

507 6.6. Generic Frame Receive process

508 6.6.1. Description

509 The Generic Frame Receive process transforms a serial bit streams of frames from the
510 Generic Data Receive process into invocations of the M_DATA.indication primitive.

511 6.6.2. State Machine Diagram

512 The operation of the Generic Frame Receive process is specified by the state machine
513 diagram in Figure 6.2 , using the variables and functions defined in subsequent sub-
514 clauses.

515 6.6.3. Variables

516 6.6.3.1. cnt

517 An integer counter variable, used to count the number of bits in a parameter of a
518 frame under reception.

519 6.6.3.2. len

520 An integer variable holding the actual length of a frame under reception, in bits.

521 6.6.3.3. status

522 An enumeration variable holding the current status of the Generic Frame Receive
523 process. The valid enumeration literals are as follows:

524 **Ok** Indicates that no error has been discovered prior or during frame reception.

525 **FrameTooLong** Indicates that a frame under reception exceeded LEN_MAX bits.

526 **FCSInvalid** Indicates inconsistency between the FCS parameter and the remaining pa-
527 rameters of a frame under reception.

⁶Such a lower layer may be an entity on the physical layer (PHY), but the generalized receive operations are not limited to this.

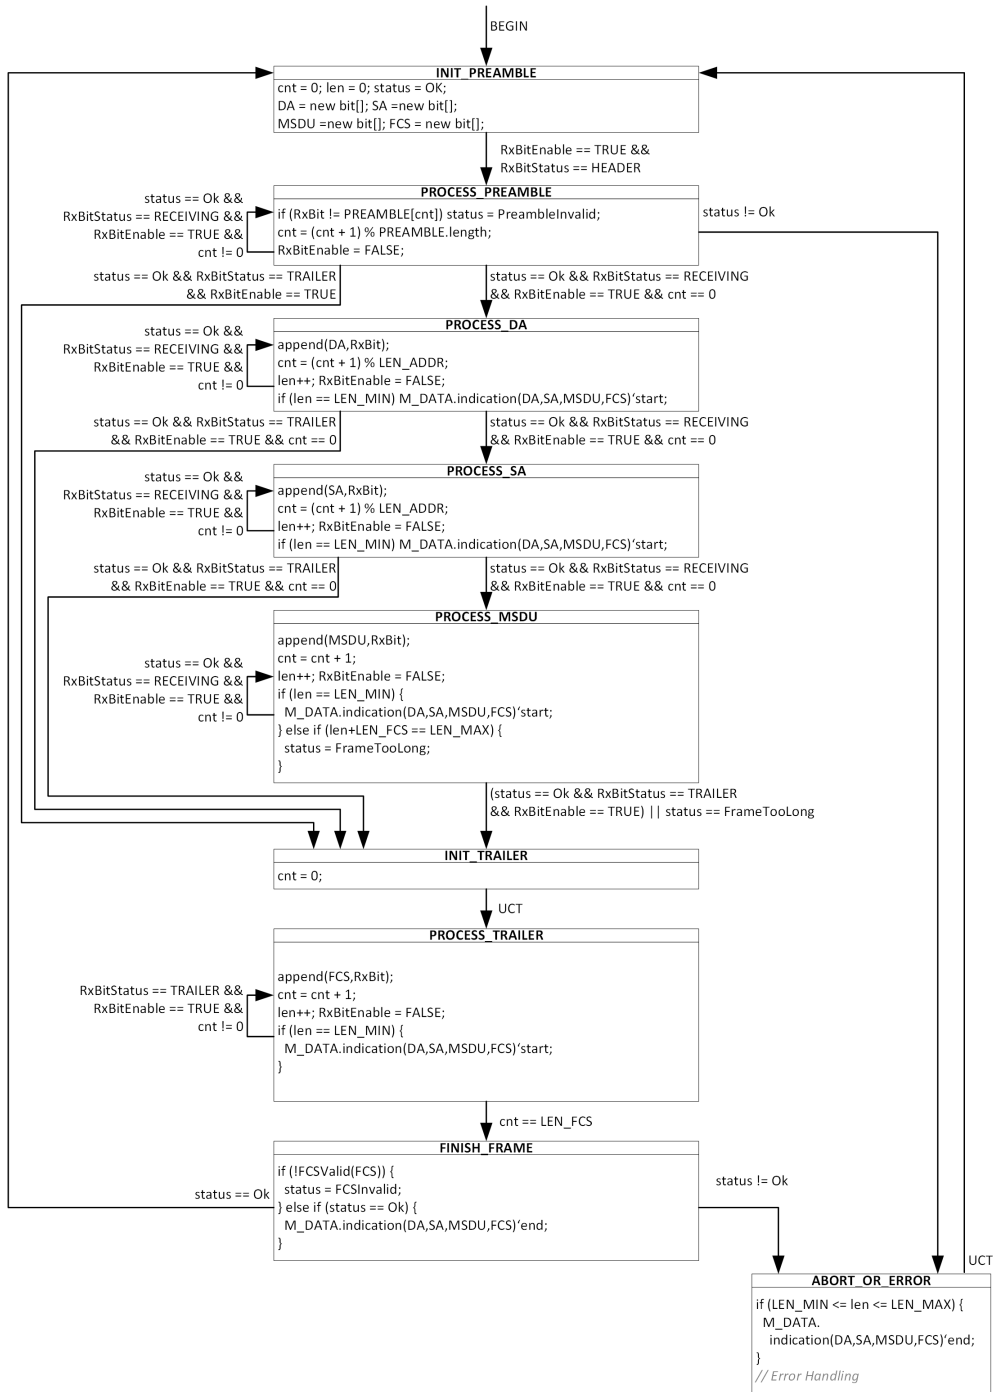


Figure 6.2.: State Machine Diagram of the Generic Frame Receive process.

528 6.6.4. Functions

529 6.6.4.1. `append(parameter,bit)`

530 The `append` function appends a given bit at the end of the passed parameter and
531 increases the length of the variable by one.

532 6.6.4.2. `FCSValid(FCS)`

533 The `FCSValid` function determines if the FCS parameter consistent with the remaining
534 parameters of the `M_DATA.indication` service primitive (TRUE) or not (FALSE). A
535 late error associated with the frame under reception is raised (5.3.2) if the function
536 returns FALSE.

537 6.7. Receive Convergence process

538 The Receive Convergence process implements the translation of `M_DATA.indication`
539 invocations to `M_UNITDATA.indication` invocations. The supported translations are
540 lower layer-dependent and include, but are not limited to, those specified in clause 13
541 of IEEE Std 802.1AC[3].

542 Each `M_DATA.indication` invocation results in an associated `M_UNITDATA.-`
543 `indication` invocation. During the translation, the `M_UNITDATA.indication` param-
544 eters are extracted from the `M_DATA.indication` parameters according to the rules
545 defined for the underlying lower layer.

546 6.8. Generic Data Transmit process

547 PLACEHOLDER, for descriptions symmetrical to 6.5.

548 6.9. Generic Frame Transmit process

549 6.9.1. Description

550 The Generic Frame Transmit process transforms invocations of the `M_DATA.request`
551 primitive from the Transmit Convergence Process into bit streams of frames.

552 6.9.2. State Machine Diagram

553 The operation of the Generic Frame Transmit process is specified by the state machine
554 diagram in Figure 6.3 , using the variables subsequently defined.

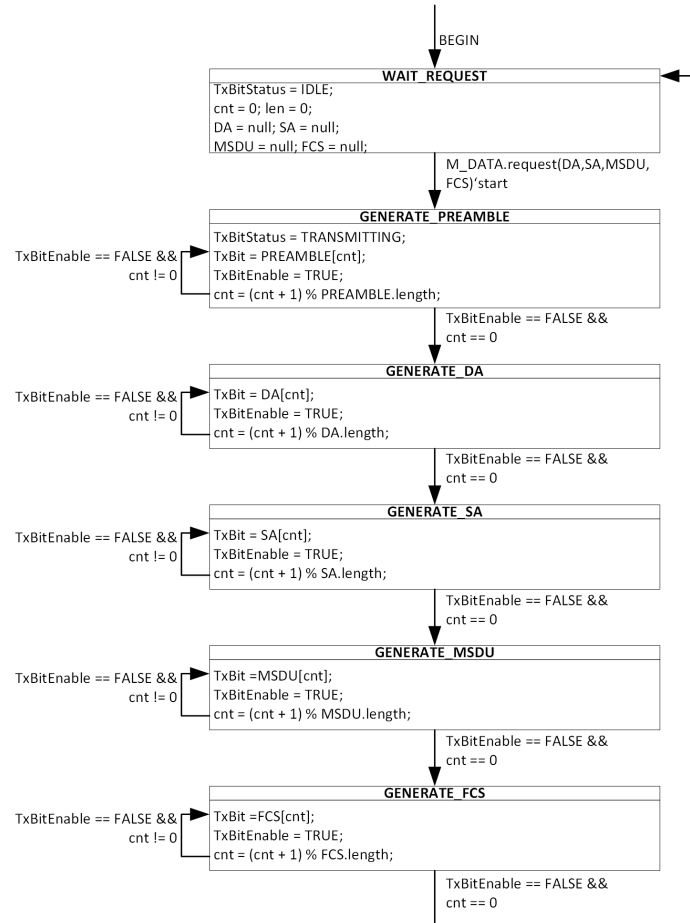


Figure 6.3.: State Machine Diagram of the Generic Frame Transmit process.

555 **6.9.3. Variables**

556 **6.9.3.1. cnt**

557 An integer counter variable, used to count the number of bits in a parameter of a
558 frame under transmission.

559 **6.10. Transmit Convergence process**

560 PLACEHOLDER, for descriptions symmetrical to 6.7.

7. Bridge Port Transmit and Receive Operations

7.1. Overview

The architecture of the bridge port transmit and receive operations in CTF bridges is identical to the architecture of S&F bridges. The architecture is shown in Figure 7.1 and Figure 7.2 for VLAN-unaware and VLAN-aware CTF bridges, respectively.

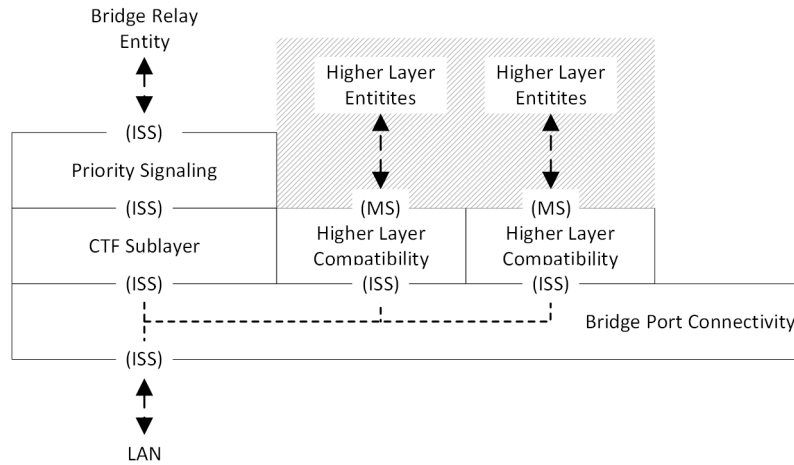


Figure 7.1.: Bridge Port Transmit and Receive (VLAN-unaware).

The elements shown are as follows:

1. Bridges Port Connectivity (7.2) between the access points of the ISS.
2. Priority Signaling in VLAN-unaware CTF bridges (7.4).
3. Translations between ISS and EISS in VLAN-aware CTF bridges (7.4).
4. Higher Layer Compatibility (7.5).
5. CTF Sublayer (7.6).




Figure 7.2.: Bridge Port Transmit and Receive (VLAN-aware).

7.2. Bridge Port Connectivity

Bridge Port connectivity in a CTF bridge is identical to S&F bridges specified in IEEE Std 802.1Q [2, 8.5.1] with the additions described in this section.

For frames under reception originating from the LAN, a copy of such frames for each upper access point is created prior to passing each copy towards the respective upper access point. Frames from the upper access points towards the LAN are passed instantaneously. The multiplexing rules towards the LAN are identical to those of S&F bridges with the addition that frames under reception originating from the bridge relay entity are treated as received frames.

7.3. Priority Signaling

7.3.1. Receive path operations

For VLAN-unaware CTF bridges, the shim for support of the ISS with signaled priority [2, 6.20] is used to determine the drop_eligible and priority parameter (6.2.2) values of tagged frames destined towards the bridge relay entity, with the following additional definitions for frames under transmission and frames under reception.

Frames under reception are stalled pending the initial two octets of the mac_service_data_unit. Dependent on the value of these octets, the processing is as follows:

1. If the octets indicate a Customer VLAN Tag [2, Table 9-1], the frame is stalled pending the PCP and DEI fields of the VLAN Tag Control Information [2, 9.6], the priority and drop_eligible parameters are instantaneously assigned to the

593 frame according to IEEE Std 802.1Q [2, 6.9.3] and the frame is passed towards
594 the bridge relay entity.

595 2. If the octets indicate any other VLAN Tag [2, Table 9-1], processing falls back
596 to S&F prior to passing the frame towards the bridge relay entity¹.

597 3. In all other cases, the frame is passed towards the bridge relay entity instantane-
598 ously.

599 For frames under reception, the invocation of M_UNITDATA.indication (M_UNIT-
600 DATA.indication'start) towards the bridge relay entity starts when the frame is passed
601 to the bridge relay entity according to the aforesaid definitions, and ends when the origi-
602 nating invocation of M_UNITDATA.indication ends (M_UNITDATA.indication'end)².

603 7.3.2. Transmit path operations

604 All frames originating from the bridge relay entity are passed towards bridge Port
605 connectivity (7.2) instantaneously.

606 7.4. Translations between Internal Sublayer Service 607 (ISS) and Enhanced Internal Sublayer Service 608 (EISS)

609 7.4.1. Receive path operations

610 The translations from ISS to EISS on the receive path can discard untagged frames,
611 and decode and remove VLAN tags from the mac_service_data_unit parameter. The
612 receive path operations are as specified in IEEE Std 802.1Q[2, 9.6.1], with the following
613 additional definitions for frames under reception.

614 Each frame under reception is stalled pending the first two octets of the mac_service_data_unit
615 parameter containing that may indicate a VLAN tag, before processing as follows:

616 1. If no VLAN tag is indicated but only tagged frames are accepted [2, item a) in
617 6.9.1], the frame is discarded.

618 2. If no VLAN tag is indicated and untagged frames are accepted [2, items c)2), c)3)
619 and d) in 6.9], the frame is passed towards the bridge relay entity instantaneously.

620 3. If a VLAN tag other than a Customer VLAN Tag [2, Table 9-1] is indicated,
621 processing falls back to S&F prior to processing as specified in IEEE Std 802.1Q
622 and passing the frame towards the bridge relay entity.

¹This fall back condition is introduced to limit the scope of this document. The same rationale applies in 7.4

²This definition is rather for illustration in the model. It provides a means in subsequent processing stages for distinguishing between frames under reception and received frames.

623 4. If a Customer VLAN Tag is indicated, processing is stalled pending the 3rd and
 624 4th octet of the `mac_service_data_unit`, the initial four octets are removed,
 625 and the `vlan_identifier`, `priority` and `drop_eligible` parameters are determined
 626 from the removed octets as specified in IEEE Std 802.1Q. Whether the frame
 627 under reception is then passed towards the bridge relay entity or discarded is
 628 determined according to IEEE Std 802.1Q [2, item b) in 6.9.1].

629 For frames under reception, the invocation of `EM_UNITDATA.indication` (`EM_UNIT-`
 630 `DATA.indication'start`) towards the bridge relay entity starts when the frame is passed
 631 to the bridge relay entity according to the aforesaid definitions, and ends when the orig-
 632 inating invocation of `M_UNITDATA.indication` ends (`EM_UNITDATA.indication'end`).

633 **7.4.2. Transmit path operations**

634 The translations from EISS to ISS on the transmit path of S&F bridges can discard
 635 tagged frames, encode and insert VLAN tags into the `mac_service_data_unit` param-
 636 eter, and adjust the `mac_service_data_unit` parameter in accordance with ISO/IEC
 637 11802-5, IETF RFC 1042 (1988), and IETF RFC 1390 [2, 9.6.2].

638 The transmit path operations in this section limit on encoding and insertion of
 639 VLAN tags due to the definitions for queuing (8.1) for frames under reception. The
 640 definitions for queuing prevent against buffer under runs, insertion and encoding of
 641 VLAN-Tag in this section is as specified in IEEE Std 802.1Q.

642 **7.5. Higher Layer Compatibility**

643 Higher layer compatibility ensures that only frames with consistent FCS are passed
 644 via the MAC Service Interface to higher layer entities. Therefore, a CTF bridge falls
 645 back to S&F prior to passing copies of frames under reception towards higher layer
 646 entities.

647 **7.6. CTF Sublayer**

648 **7.6.1. Receive Path Operations**

649 On the receive path, the CTF sublayer can emit late errors for frames under reception
 650 evaluates the `CTFReceptionEnable` parameter (9.2.4).

651 If frame under reception is destined towards the bridge relay entity and the `CTFRe-`
 652 `ceptionEnable` is `FALSE`, processing fall-back to S&F for this frames prior to passing
 653 it to the ISS towards the relay.

654 If frame under reception is destined towards the bridge relay entity and the `CTFRe-`
 655 `ceptionEnable` is `TRUE`, this frame is passed instantaneously to the translation from
 656 ISS towards the relay (7.4 and 7.3). The CTF sublayer maintains reference to frames
 657 under reception after passing these frames towards the bridge relay. If a frame with
 658 inconsistent FCS appears, the following steps are performed:

- 659 1. A late error associated with this frame is raised.
- 660 2. An frame error counter is increased (7.6.3).

661 7.6.2. Transmit Path Operations

662 The transmit path of the CTF sublayer passes frames from the bridge relay entity
 663 towards the LAN instantaneously. For any frame under transmission that is also
 664 frame under reception (i.e., Cut-Through), the transmit path operations of the CTF
 665 sublayer marks (7.6.3) each of these for which a late error has been raised on the CTF
 666 sublayer receive path (7.6.1) or by the bridge relay (8).

667 7.6.3. Inconsistent frame handling

668 Handling of inconsistent frames increases on of two diagnostic error counters on the
 669 receive path, CTFReceptionDiscoveredErrors (9.4.1) and CTFReceptionUndiscovered-
 670 Errors (9.4.2), as follows:

- 671 – If the frame has been marked by an upstream bridge and this mark was identified
 672 as such on the receive path (7.6.1), CTFReceptionDiscoveredErrors is increased.
- 673 – In all other cases, CTFReceptionUndiscoveredErrors is increased.

674 Marking inconsistent frames introduces assigns a externally visible indicator to such
 675 frames, usually at the end of serial transmission. In existing implementations of CTF,
 676 the marking mechanism varies. For example, an implementation may apply a modified
 677 FCS determined as follows:

- 678 1. Calculate a consistent FCS for the frame.
- 679 2. Modify the calculated consistent FCS in a deterministic manner. Examples:
 680 a) Exchange bits of the FCS at known positions.
 681 b) Invert bits of the FCS known positions.
 682 c) Perform an XOR operation between the FCS and a known constant value.
- 683 3. Replace the frame_check_sequence parameter of the associated M_UNITDATA.-
 684 request invocation with the modified FCS.

685 8. Bridge Relay Operations

686 8.1. Overview

687 The structure of the bridge relay entity of CTF bridges is aligned with that of an S&F
688 bridge. Additional definitions for supporting frames under reception for Cut-Through
689 exist primarily in the forwarding process. The structure of the forwarding process in
690 CTF bridges, in terms of processing stages passed by frames, is likewise aligned with
691 that of S&F bridges. It comprises processing stages symmetrical to those found in S&F
692 bridges [2, 8.6 and Figure 8-12] with incorporated processing stages for Frame Repli-
693 cation and Elimination for Reliability [4, 8.1 and Figure 8-2]. The forwarding process
694 of a CTF bridge, additional elements in the bridge relay and indicated interactions
695 between them are shown in Figure 8.1.

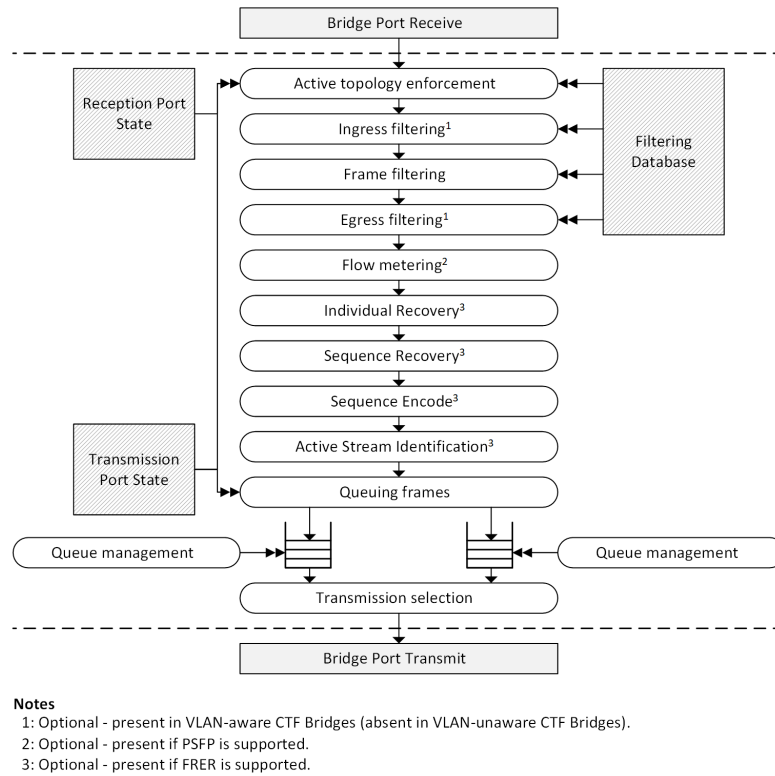


Figure 8.1.: Forwarding process of a CTF bridge.

696 The processing stages and the respective sections in this document are as follows:

- 697 1. Active topology enforcement (8.2)
- 698 2. Ingress filtering (8.3)
- 699 3. Frame filtering (8.4)
- 700 4. Egress filtering (8.5)
- 701 5. Flow classification and metering (8.6)
- 702 6. Individual recovery (8.7)
- 703 7. Sequence recovery (8.8)
- 704 8. Sequence encode (8.9)
- 705 9. Active stream identification (8.10)

706 10. Queuing frames (8.11), and associated additional definitions for queue manage-
707 ment (8.12)

708 11. Transmission selection (8.13)

709 The sections of the processing stages are written in a manner that avoids replicating
710 contents of the corresponding sections in the published IEEE 802.1 Standards. Instead,
711 section provide reference to the corresponding section(s) in the published standards,
712 followed by additional definitions for processing frames under reception. While the
713 emphasis is on processing frames under reception, the stages are equally capable for
714 processing received frames. In the latter case, the behavior of the processing stages is
715 identical to that of an S&F bridge.

716 8.2. Active Topology Enforcement

717 8.2.1. Overview

718 The active topology enforcement stage determines if frames from reception Ports are
719 used for learning, and determines the initial set of potential transmission Ports for each
720 frame. Both operations are as specified in IEEE Std 802.1Q [2, 8.6.1] in CTF bridges,
721 with the additions described in the following for learning (8.2.2) and the initial set of
722 potential transmission Ports (8.2.3) separately.

723 8.2.2. Learning

724 Learning is based on the the source address and VID parameters of frames for adding
725 entries in the forwarding database (FDB) as specified in IEEE Std 802.1Q [2, 8.7].
726 In CTF bridges, the source address and VID parameters are used for learning the
727 following conditions are satisfied:

- 728 1. A frame under reception associated with the parameters reached the end of
729 reception.
- 730 2. This frame's FCS is consistent.
- 731 3. All conditions of an S&F bridge for using the parameters for learning are satisfied
732 [2, 8.4 and 8.6.1].

733 8.2.3. Initial set of potential transmission Ports

734 The initial set of potential transmission Ports is determined by CTF bridges as specified
735 in IEEE Std 802.1Q [2, 8.6.1]. If this determination depends on the VID parameter of
736 a frame under reception, processing stalls pending this parameter prior to passing the
737 frame under reception to the next processing stage:

- 738 – Ingress filtering (8.3) for VLAN-aware CTF bridges

- 739 – Frame filtering (8.4) for VLAN-unaware CTF bridges

740 In absence of this dependency, the frame under reception is passed to the next pro-
 741 cessing stage instantaneously.

742 8.3. Ingress Filtering

743 The ingress filtering stage discards frames originating from reception Ports based on
 744 the VID parameters associated with these frames. The conditions under which a frame
 745 is discarded by a CTF bridge are identical to those specified in IEEE Std 802.1Q [2,
 746 8.6.2]. Frames under reception are stalled by VLAN-aware CTF bridges pending the
 747 VID parameter and passed to the next processing stage (8.4) unless they are discarded
 748 and therefore not passed, either due to the ingress filtering operation or due to the
 749 implicit discarding rule while stalled (5.3).

750 The ingress filtering stage is only present in VLAN-aware CTF bridges.

751 8.4. Frame Filtering

752 The frame filtering stage reduces the set of potential transmission Ports associated
 753 with a frame based on parameters associated with this frame (destination address,
 754 VID, etc.) and querying the FDB of a bridge. The exact set of parameters of a frame
 755 is determined as specified in IEEE Std 802.1Q [2, 8.6.3]. If necessary, a CTF bridge
 756 stalls processing pending all necessary parameters of a frame under reception before
 757 performing an FDB query for this frame [2, 8.8.9].

758 Dependent on the query’s evaluation by the FDB, processing of a frame under
 759 reception falls back to S&F or passes the frame to the next stage instantaneously as
 760 follows:

- 761 – Whenever the query evaluation by the FDB results in flooding (i.e., query eval-
 762 uation hits an “ELSE Forward” branch in 8.8.9 of IEEE Std 802.1Q), processing
 763 of the frame falls back to S&F¹.
- 764 – In all other cases, a frame under reception is passed to the next processing stage
 765 instantaneously.

766 8.5. Egress Filtering

767 The egress filtering stage reduces the set of potential transmission Ports associated with
 768 a frame based on this frame’s VID parameter. The rules under which transmission
 769 Ports are removed from this set are identical to those specified in IEEE Std 802.1Q
 770 [2, 8.6.4]. Frames under reception are passed to the next processing stage once this

¹This fall back is intended to reduce the cases for circulation of inconsistent frames in topological loops, assuming that the performance benefits of CTF traffic that is subject to flooding are of little real-world use.

771 reduction finished². The egress filtering stage is only present in VLAN-aware CTF
772 bridges.

773 8.6. Flow Classification and Metering

774 8.6.1. General

775 The flow classification and metering stage can can apply flow classification and meter-
776 ing to frames that are received on a Bridge Port and have one or more potential trans-
777 mission ports. This processing stage is structured into multiple internal (sub)stages in
778 CTF bridges, identical to the structure specified in IEEE Std 802.1Q [2, 8.6.5]. The
internal stages and their relationships are shown in Figure 8.2 .

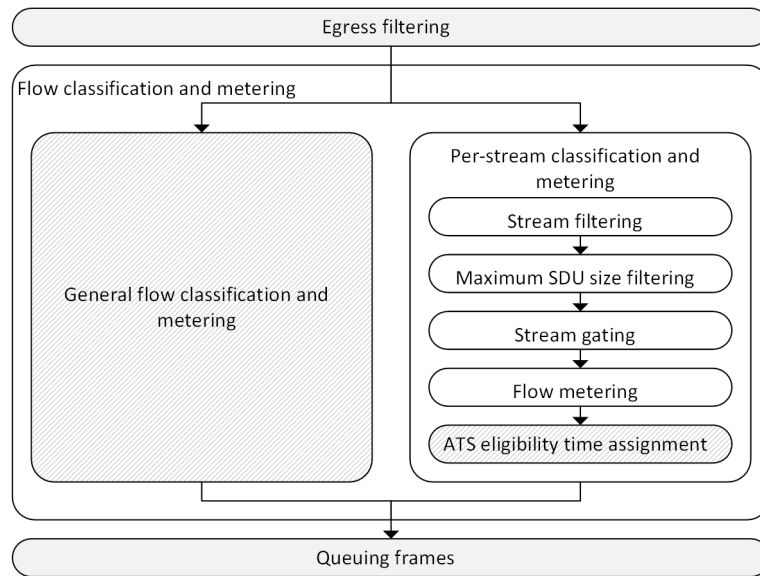


Figure 8.2.: Flow classification and metering.

779 Support for frames under reception is provided by CTF bridges for the following
780 internal stages:
781

- 782 1. Stream filtering
- 783 2. Maximum SDU size filtering
- 784 3. Stream gating
- 785 4. Flow metering

²It is not required to stall processing pending a frame's VID, because this already happened during ingress filtering (8.3).

786 Processing in CTF bridges falls back to S&F immediately if a frame under reception
 787 reaches any other internal stage prior to being processed by this stage. The operation
 788 of stages with support for frames under reception is described in 8.6.2, 8.6.3, 8.6.4 and
 789 8.6.5. With the exception of stream filtering, all subsequently described stages process
 790 frames under reception instantaneously (i.e., stall-free operation). When one of these
 791 stages passed a frame under reception to a subsequent processing stage, the associated
 792 frame counters of the stream filtering [2, items h) through m) in 8.6.5.3] are increased
 793 according to the rules specified in IEEE 802.1Q at the instant of time the frame is
 794 passed.

795 8.6.2. Stream Filtering

796 Frames under reception are associated with stream filters according to the rules spec-
 797 ified in IEEE Std 802.1Q [2, 8.6.5.3]. If this association depends on a *stream_handle*
 798 parameter specified in IEEE Std 802.1CB [4], processing is stalled pending on this
 799 parameter prior to associating a stream filter. An associated stream filter then per-
 800 forms all necessary associations with subsequent internal stages passes these to the
 801 first associated internal stage instantaneously.

802 8.6.3. Maximum SDU size filtering

803 The operation of maximum SDU size filtering for frames under reception is as specified
 804 in IEEE Std 802.1Q [2, 8.6.5.3.1] with the additions in this section. When a frame
 805 under reception reaches maximum SDU size filtering, an initial number of octets of this
 806 frame is already received. This number of octets is used by maximum SDU size filtering
 807 for the decision on whether or not this frame is passed to a subsequent processing stage
 808 or discarded. If a frame under reception already passed frame maximum SDU size
 809 filtering and the associated maximum SDU size limit is exceeded prior to the frame's
 810 end of reception, a late error for that frame is indicated for handling by subsequent
 811 processing stages in a CTF bridge.

812 8.6.4. Stream Gating

813 The operation of stream gates for frames under reception is as specified in IEEE Std
 814 802.1Q [2, 8.6.5.4] with the additions in this section. Once a frame under reception
 815 reaches a stream gate, this frame is only passed to the next processing stage if the
 816 gate is in an open state. The frame is discard otherwise prior to being passed to the
 817 next processing stage. If a stream If a stream gate closes prior to the end of the frame
 818 under reception, a late error for this frame is indicated immediately for handling by
 819 subsequent processing stages in a CTF bridge.

820 8.6.5. Flow Metering

821 The operation of stream gates for frames under reception is as specified in IEEE Std
 822 802.1Q [2, 8.6.5.5] with the additions in this section. When a frame under reception

823 reaches flow metering, an initial number of octets of this frame is already received.
 824 This number of octets is used by the associated flow meter for the decision on whether
 825 or not this frame is passed to a subsequent processing stage or immediately discarded.
 826 If a frame under reception already passed flow metering and the limit of the flow
 827 meter is subsequently exceeded prior to the frame's end of reception, a late error for
 828 this frame is indicated for handling by subsequent processing stages in a CTF bridge.

829 **8.7. Individual Recovery**

830 The individual recovery stage can associate frames belonging to individual Member
 831 streams [4, 7.4.2] with therefore configured instances of the Base recovery function [4,
 832 7.4.3], which then discard frames with repeating sequence_number parameters[4, item
 833 b) in 6.1] on a per Member stream resolution. The operation of the individual recovery
 834 stage is as specified in IEEE Std 802.1CB [4, 7.5], with the following additions for CTF
 835 bridges.

836 If frames under reception are associated with a Base recovery function for individual
 837 recovery, processing falls back to S&F prior to performing individual recovery³.

838 **8.8. Sequence Recovery**

839 The sequence recovery stage can associate frames belonging to sets of Member streams
 840 with therefore configured instances of the Base recovery function [4, 7.4.3], which then
 841 remove frames with repeating sequence_number parameters[4, item b) in 6.1] on a
 842 per Member stream set resolution. The operation of the sequence recovery stage is as
 843 specified in IEEE Std 802.1CB [4, 7.4.2], with the following additions for CTF bridges.

844 If frames under reception are associated with a Base recovery function for sequence
 845 recovery, processing falls back to S&F prior to performing sequence recovery.

846 **8.9. Sequence Encode**

847 The sequence encode stage can insert externally visible tags with sequence numbers
 848 into frames that represent the sequence_number parameter associated with these
 849 frames. The operations of the sequence encode stage and the tag formats for frames
 850 under reception are as specified in IEEE Std 802.1CB [4, 7.6 and 7.8].

851 **8.10. Active Stream Identification**

852 PLACEHOLDER, for describing differences and additions to 6.2 of IEEE Std 802.1CB.
 853 This processing stage may be placed differently (in conjunction with incorporating

³Falling back to S&F ensures that individual recovery does not falsely discard a frame with correct sequence_number parameter (and consistent FCS) after accepting a frame with incorrect but identical sequence_number (and inconsistent FCS) earlier. The same rationale applies in 8.8.

Algorithm 8.1 Queuing rules for frames under reception.

IF

(the associated CTFTransmissionEnable parameter [9.2.2] is FALSE) **OR**
 (the associated transmission selection algorithm is not strict priority [2, 8.6.8.1])

THEN

Processing of the frame falls back to S&F before queuing it instantaneously.

ELSE IF

(the associated CTFTransmissionEnable parameter [9.2.2] is TRUE) **AND**
 (the nominal transmit duration of the at the associated transmission Port
 would be less than the nominal duration of it's reception)

THEN

The frame is discarded before queuing.

ELSE

The frame is queued instantaneously.

854 stages for passive stream identification, sequence decoding and sequence generation [4,
 855 Figure 8-2]), subject to ongoing work in IEEE WG 802.1 at time of writing.

8.11. Queuing Frames

857 The queuing frames stage queues each received frame to a per-traffic class queue of
 858 each remaining potential transmission Port associated with the frame (8.2, 8.4 and
 859 8.5). The rules to determine the correct per-traffic queues for frames under reception
 860 are identical to the rules specified in IEEE Std 802.1Q [2, 8.6.6] with the following
 861 additions.

862 Before a frame under reception is queued, a per-queue copy of a frame before queu-
 863 ing. Each frame under reception resulting from this copy operation is then considered
 864 separately to allow for consistent transmission (8.13) as shown in Algorithm 8.1.

8.12. Queue Management

866 The rules for removing frames from IEEE Std 802.1Q [2, 8.6.7] remain unaltered in
 867 CTF bridges.

868 In addition to this, CTF bridges may remove a frame from a queue if all of the
 869 following conditions are satisfied⁴:

- 870 1. The frame was queued while it was under reception.
- 871 2. A processing stage before queuing(8.11) raised a late error for that frame.
- 872 3. the end of reception of the frame was reached before the frame was selected for
 873 transmission (8.13).

⁴Erroneous frames removed according to this additional rule will not become visible on the LAN of an associated transmission Port.

8.13. Transmission Selection

Transmission selection determines whether frames in per traffic class queues are available for transmission, determines transmission ordering and transmission times, initiates transmission of the frames, and removes transmitted frames from the per traffic class queues. Transmission selection in CTF bridges is as specified in IEEE Std 802.1Q [2, 8.6.8].

880 9. Management Parameters

881 9.1. Overview

882 The management parameters for CTF fall into three categories:

- 883 1. Control Parameters (9.2)
- 884 2. Timing Parameters (9.3)
- 885 3. Error Counters (9.4)

886 The control parameters allow to (i) determine whether CTF is supported on a per Port
887 and per Port per Traffic Class resolution, and if CTF is supported, to (ii) enable and
888 disable CTF on these resolutions. These parameters are available in reception Ports
889 and transmission Ports. For a pair of bridge ports, frames can only be subject to the
890 CTF operation if CTF is supported and enabled on both Ports.

891 The timing parameters expose the delays experienced by frames passing from a
892 particular reception Port to another transmission Port. These parameters are primarily
893 intended for automated network and traffic configuration, for example, by a Centralized
894 Network Controller (CNC) using the associated mechanisms from IEEE Std 802.1Q
895 [2, clause 46].

896 The error counters expose information on frames that were subject to the CTF oper-
897 ation in a bridge, even though such frames have consistency errors (i.e., a frame check
898 sequence inconsistent with the remaining contents of that frame) during reception by
899 this bridge. These counters are primarily intended for manual diagnostic purposes
900 to support identifying erroneous links or stations, for example, by a human network
901 administrator.

902 9.2. Control Parameters

903 9.2.1. CTFTransmissionSupported

904 A Boolean read-only parameter that indicates whether CTF on transmission is sup-
905 ported (TRUE) or not (FALSE). There is one CTFTransmissionSupported parameter
906 for each traffic class of each transmission Port.

907 9.2.2. CTFTransmissionEnable

908 A Boolean parameter to enable (TRUE) and disable (FALSE) CTF on transmission.
909 There is one CTFTransmissionEnable parameter for each traffic class of each transmis-
910 sion Port. The default value of the CTFTransmissionEnable parameter is FALSE for

all traffic classes of all transmission Ports. It is an error if a CTFTransmissionEnable is set to TRUE if the associated CTF Transmission Supported parameter is FALSE.

9.2.3. CTFReceptionSupported

A Boolean read-only parameter that indicates whether CTF on reception is supported (TRUE) or not (FALSE). There is one CTFReceptionSupported parameter for each reception Port.

9.2.4. CTFReceptionEnable

A Boolean parameter to enable (TRUE) and disable (FALSE) CTF on reception. There is one CTFReceptionEnable parameter for each reception Port. The default value of the CTFReceptionEnable parameter is FALSE for all reception Ports. It is an error if a CTFReceptionEnable is set to TRUE if the associated CTFReceptionSupported parameter is FALSE.

9.3. Timing Parameters

9.3.1. CTFDelayMin and CTFDelayMax

A pair of unsigned integer read-only parameters, in units of nanoseconds, describing the delay range for frames that are subject to the CTF operation and encounter zero delay for transmission selection [2, 8.6.8]. This occurs when the queue for the frame's traffic class is empty, the frame's traffic class has permission to transmit, and the egress Port is idle (not transmitting). There is one pair of CTFDelayMin and CTFDelayMax parameters per reception Port per transmission Port traffic class pair.

9.4. Error Counters

9.4.1. CTFReceptionDiscoveredErrors

An integer counter, counting the number of received frames with discovered consistency errors. There is one CTFReceptionDiscoveredErrors parameter for each reception Port. A frame with discovered consistency errors has been identified as such by a bridge on the upstream path from which the frame originates and marked by that an implementation-dependent marking mechanism. The value of the counter always increases by one

1. if
 - a) the upstream bridge that applied the marking,
 - b) all bridges on the path of that bridge to the reception Port associated with the CTFReceptionDiscoveredErrors counter and

- 943 c) the receiving bridge of which the reception Port is a part of are different
 944 instances of the same bridge implementation, and
- 945 2. the underlying marking mechanism is identical for all these instances if multiple
 946 marking mechanisms are supported by these instances.
- 947 If either of the conditions in items 1 through 2 is unsatisfied, `CTFReceptionUndiscoveredErrors` may be increased instead of `CTFReceptionDiscoveredErrors`¹.
 948

949 **9.4.2. CTFReceptionUndiscoveredErrors**

950 An integer counter, counting the number of received frames with undiscovered consistency errors. There is one `CTFReceptionUndiscoveredErrors` parameter for each
 951 reception Port. This counter is increased by one if a frame with consistency errors is
 952 received at the associated reception Port and `CTFReceptionDiscoveredErrors` is not
 953 increased.
 954

¹It is assumed that there is a variety of options for implementing a frame marking mechanism. For example, by using physical layer symbols [10, 1.121 - 1.126] or special frame check sequences [11, p.54, 2.2.][12, p.17]. The current description in this document permits any marking mechanism, but the associated error counters are only consistent in networks with homogeneous implementation instances, and may be inconsistent in heterogeneous networks. However, term (`CTFReceptionDiscoveredErrors` + `CTFReceptionUndiscoveredErrors`) on a reception Port should be identical in several heterogeneous networks. A human network administrator may be able to localize erroneous links or stations solely by considering this term along multiple reception Ports across a network instead of its constituents.

955

Part III.

956

Cut-Through Forwarding in Bridged Networks

957

958 PLACEHOLDER, for contents on using CTF in networks [11, p.46 – p.49].

959

Part IV.

960

Appendices

A. Interaction of the Lower Layer Interface (LLI) with existing Lower Layers

A.1. PLS Service Interface

A.1.1. Overview

This section summarizes how interfacing between the PLS service primitives on top of the Reconciliation Sublayer [13, clause 22, clause 35, etc.] and LLI (6.1) is possible, similar to the interfacing of the original GSCF [7]. Interfacing between PLS service primitives and LLI can be established by three processes that translate between the LLI global variables (6.4) and the PLS service primitives. The processes and interactions are shown in Figure A.1.




Figure A.1.: Processes and interactions for interfacing between LLI and PLS service primitives.

A.1.2. Service Primitives

The `PLS_DATA.indication`, `PLS_DATA_VALID.indication`, `PLS_CARRIER.indication` and `PLS_DATA.request` service primitives are as specified in IEEE Std 802.3 [13,

975 clause 6] limiting on full duplex mode¹.

976 **A.1.3. Global Variables and Constants**

977 **A.1.3.1. BitTick**

978 A global Boolean variable, used to generate a bit clock for the PLS Transmit process.

979 **A.1.3.2. LEN_FRAMEGAP**

980 An integer constant defining the duration of the Inter-Frame Gap (IFG), in bits.

981 **A.1.4. Global Constraints**

982 The following constraints are introduced for the Global Constants in sections 6.3 and
983 A.1.3:

- 984 1. PREAMBLE = "10101010 10101010 10101010 10101010 10101010 10101010 10101010
985 10101011"²
- 986 2. LEN_MIN = 8*64 + PREAMBLE.length
- 987 3. LEN_MAX = 8*1500 + PREAMBLE.length
- 988 4. LEN_FCS = 32
- 989 5. LEN_DATA = 1
- 990 6. LEN_FRAMEGAP = 8*12

991 **A.1.5. Transmit Bit Clock process**

992 The Transmit Bit Clock process periodically sets the BitTick variable to TRUE, where
993 the period equals the duration of a Bit on the physical layer.

994 **A.1.6. PLS Transmit process**

995 **A.1.6.1. Description**

996 The PLS Transmit process translates between global variables from the Generic Data
997 Transmit process (6.8) and the PLS_CARRIER.indication and PLS_DATA.request
998 service primitives (A.1.2).

¹The PLS_SIGNAL.indication service primitive is effectively not required in this mode [13, 6.3.2.2.2 and 7.2.1.2]

²First bit in quotes is PREAMBLE[0], second bit in quotes is PREAMBLE[1], etc. whitespaces are ignored.

999 A.1.6.2. State Machine Diagram

1000 The operation of the PLS Transmit process is defined by the state machine diagram in Figure A.2.

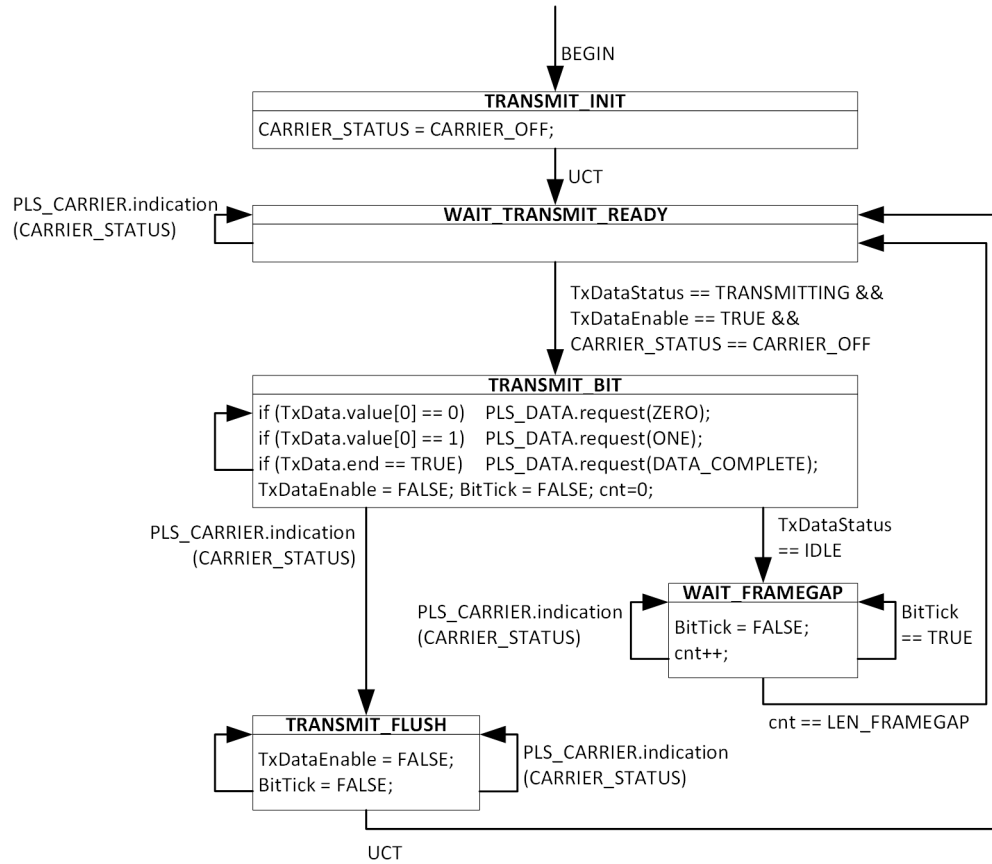


Figure A.2.: State machine diagram of the PLS Transmit process.

1001

1002 A.1.6.3. Variables

1003 A.1.6.3.1. cnt An integer variable for counting bits.

1004 A.1.7. PLS Receive process

1005 A.1.7.1. Description

1006 The PLS Receive process translates between global variables from the Generic Data
 1007 Receive process (6.5) and the PLS_CARRIER.indication and PLS_DATA.request

1008 service primitives (A.1.2).

1009 A.1.7.2. State Machine Diagram

1010 The operation of the PLS Transmit process is defined by the state machine diagram in Figure A.3.

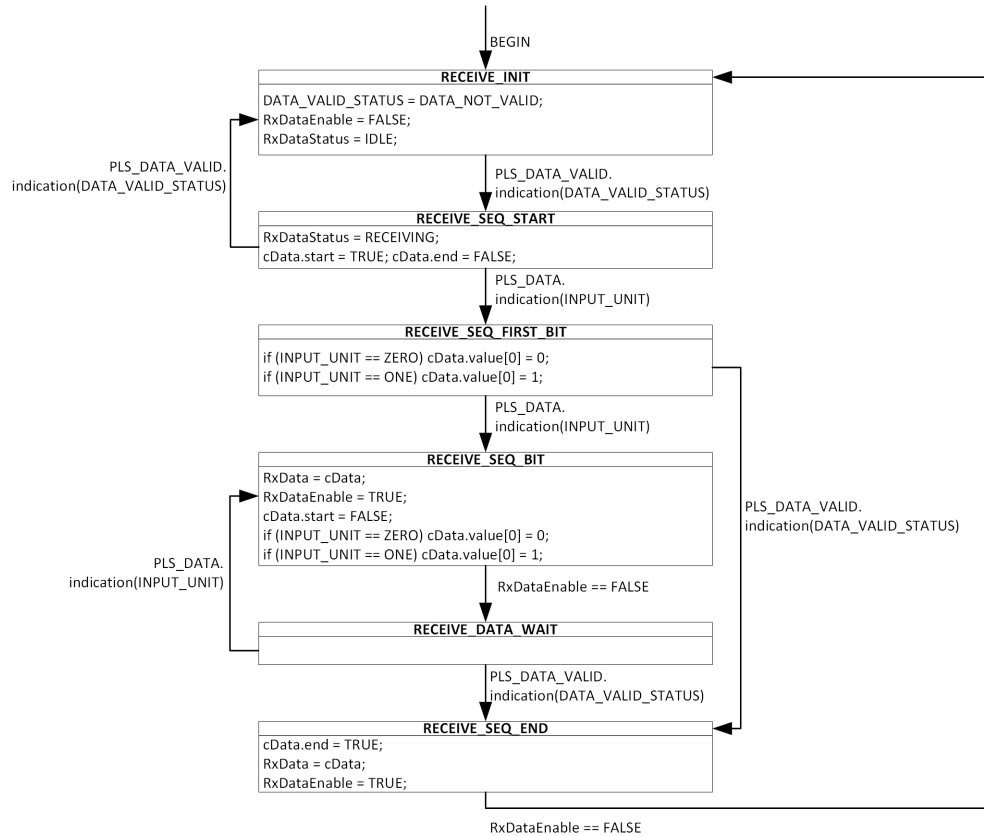


Figure A.3.: State machine diagram of the PLS Receive process.

1011

1012 A.1.7.3. Variables

1013 **A.1.7.3.1. cData** A variable of type `low_data_t` (6.5), used for implementing a
 1014 delay line of a single bit.

1015 Bibliography

- 1016 [1] IEEE Standards Association, *2021 IEEE SA Standards Style Manual*. [Online].
 1017 Available: [https://mentor.ieee.org/myproject/Public/mytools/draft/styleman.](https://mentor.ieee.org/myproject/Public/mytools/draft/styleman.pdf)
 1018 pdf
- 1019 [2] “IEEE Standard for Local and Metropolitan Area Network–Bridges and Bridged
 1020 Networks,” *IEEE Std 802.1Q-2018 (Revision of IEEE Std 802.1Q-2014) and pub-*
 1021 *lished amendments*, pp. 1–1993, 2018.
- 1022 [3] “IEEE Standard for Local and metropolitan area networks – Media Access Con-
 1023 trol (MAC) Service Definition,” *IEEE Std 802.1AC-2016 (Revision of IEEE Std*
 1024 *802.1AC-2012)*, pp. 1–52, 2017.
- 1025 [4] “IEEE Standard for Local and metropolitan area networks–Frame Replication and
 1026 Elimination for Reliability,” *IEEE Std 802.1CB-2017 and published amendments*,
 1027 pp. 1–102, 2017.
- 1028 [5] “IEEE Standard for Local and metropolitan area networks – Media Access Con-
 1029 trol (MAC) Service Definition,” *IEEE Std 802.1AC-2016 (Revision of IEEE Std*
 1030 *802.1AC-2012)*, pp. 1–52, 2017.
- 1031 [6] Johannes Specht (Self; Analog Devices, Inc.; Mitsubishi Electric Corpo-
 1032 ration; Phoenix Contact GmbH & Co. KG; PROFIBUS Nutzerorganisa-
 1033 tion e.V.; Siemens AG; Texas Instruments, Inc.), *An Idealistic Model*
 1034 *for P802.1DU*. [Online]. Available: [https://mentor.ieee.org/802.1/dcn/22/](https://mentor.ieee.org/802.1/dcn/22/1-22-0015-01-ICne-idealistic-model-for-p802-1du.pdf)
 1035 [1-22-0015-01-ICne-idealistic-model-for-p802-1du.pdf](https://mentor.ieee.org/802.1/dcn/22/1-22-0015-01-ICne-idealistic-model-for-p802-1du.pdf)
- 1036 [7] Roger Marks (EthAirNet Associates), *Generic Serial Convergence Function*
 1037 *(GSCF)*, 2022. [Online]. Available: [https://mentor.ieee.org/802.1/dcn/22/](https://mentor.ieee.org/802.1/dcn/22/1-22-0040-02-ICne-generic-serial-convergence-function-gscf.pdf)
 1038 [1-22-0040-02-ICne-generic-serial-convergence-function-gscf.pdf](https://mentor.ieee.org/802.1/dcn/22/1-22-0040-02-ICne-generic-serial-convergence-function-gscf.pdf)
- 1039 [8] Johannes Specht (Self; Analog Devices, Inc.; Mitsubishi Electric Corporation;
 1040 Phoenix Contact GmbH & Co. KG; PROFIBUS Nutzerorganisation e.V.; Siemens
 1041 AG; Texas Instruments, Inc.), *CTF - Considerations on Modelling, Compatibility*
 1042 *and Locations*. [Online]. Available: [https://mentor.ieee.org/802.1/dcn/22/](https://mentor.ieee.org/802.1/dcn/22/1-22-0021-04-ICne-ctf-considerations-on-modelling-compatibility-and-locations.pdf)
 1043 [1-22-0021-04-ICne-ctf-considerations-on-modelling-compatibility-and-locations.](https://mentor.ieee.org/802.1/dcn/22/1-22-0021-04-ICne-ctf-considerations-on-modelling-compatibility-and-locations.pdf)
 1044 pdf
- 1045 [9] “IEEE Standard for Information Technology–Telecommunications and Informa-
 1046 tion Exchange between Systems - Local and Metropolitan Area Networks–Specific

- 1047 Requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Phys-
 1048 ical Layer (PHY) Specifications,” *IEEE Std 802.11-2020 (Revision of IEEE Std*
 1049 *802.11-2016)*, pp. 1–4379, 2021.
- 1050 [10] Astrit Ademaj (TTTech) and Guenter Steindl (Siemens), *Cut-Through –*
 1051 *IEC/IEEE 60802 – V1.1*, 2019. [Online]. Available: [https://www.ieee802.org/1/](https://www.ieee802.org/1/files/public/docs2019/60802-Ademaj-et-al-CutThrough-0919-v11.pdf)
 1052 [files/public/docs2019/60802-Ademaj-et-al-CutThrough-0919-v11.pdf](https://www.ieee802.org/1/files/public/docs2019/60802-Ademaj-et-al-CutThrough-0919-v11.pdf)
- 1053 [11] Johannes Specht, Jordon Woods, Paul Congdon, Lily Lv, Henning
 1054 Kaltheuner, Genio Kronauer and Alon Regev, *IEEE 802 Tutorial:*
 1055 *Cut-Through Forwarding (CTF) among Ethernet networks – DCN 1-21-0037-*
 1056 *00-ICne*, 2021. [Online]. Available: [https://mentor.ieee.org/802.1/dcn/21/](https://mentor.ieee.org/802.1/dcn/21/1-21-0037-00-ICne-ieee-802-tutorial-cut-through-forwarding-ctf-among-ethernet-networks.pdf)
 1057 [1-21-0037-00-ICne-ieee-802-tutorial-cut-through-forwarding-ctf-among-ethernet-networks.](https://mentor.ieee.org/802.1/dcn/21/1-21-0037-00-ICne-ieee-802-tutorial-cut-through-forwarding-ctf-among-ethernet-networks.pdf)
 1058 pdf
- 1059 [12] Peter Jones (Cisco), *802.3 NEA CTF: CTF concerns*, 2022. [Online].
 1060 Available: [https://www.ieee802.org/3/ad_hoc/ngrates/public/calls/22_0427/](https://www.ieee802.org/3/ad_hoc/ngrates/public/calls/22_0427/jones_nea_01_220427.pdf)
 1061 [jones_nea_01_220427.pdf](https://www.ieee802.org/3/ad_hoc/ngrates/public/calls/22_0427/jones_nea_01_220427.pdf)
- 1062 [13] “IEEE Standard for Ethernet,” *IEEE Std 802.3-2018 (Revision of IEEE Std*
 1063 *802.3-2015)*, pp. 1–5600, 2018.