

DVJ Perspective on: Timing and synchronization for time-sensitive applications in bridges local area networks

Draft 0.718

Contributors:
See page xx.

Abstract: This working paper provides background and introduces possible higher level concepts for the development of Audio/Video bridges (AVB).

Keywords: audio, visual, bridge, Ethernet, time-sensitive

1 **IEEE Standards** documents are developed within the IEEE Societies and the Standards Coordinating Committees of
2 the IEEE Standards Association (IEEE-SA) Standards Board. The IEEE develops its standards through a consensus
3 development process, approved by the American National Standards Institute, which brings together volunteers repre-
4 senting varied viewpoints and interests to achieve the final product. Volunteers are not necessarily members of the Insti-
5 tute and serve without compensation. While the IEEE administers the process and establishes rules to promote fairness
6 in the consensus development process, the IEEE does not independently evaluate, test, or verify the accuracy of any of
7 the information contained in its standards.

8 Use of an IEEE Standard is wholly voluntary. The IEEE disclaims liability for any personal injury, property or other
9 damage, of any nature whatsoever, whether special, indirect, consequential, or compensatory, directly or indirectly
10 resulting from the publication, use of, or reliance upon this, or any other IEEE Standard document.

11 The IEEE does not warrant or represent the accuracy or content of the material contained herein, and expressly disclaims
12 any express or implied warranty, including any implied warranty of merchantability or fitness for a specific purpose, or
13 that the use of the material contained herein is free from patent infringement. IEEE Standards documents are supplied
14 “**AS IS.**”

15 The existence of an IEEE Standard does not imply that there are no other ways to produce, test, measure, purchase, mar-
16 ket, or provide other goods and services related to the scope of the IEEE Standard. Furthermore, the viewpoint expressed
17 at the time a standard is approved and issued is subject to change brought about through developments in the state of the
18 art and comments received from users of the standard. Every IEEE Standard is subjected to review at least every five
19 years for revision or reaffirmation. When a document is more than five years old and has not been reaffirmed, it is rea-
20 sonable to conclude that its contents, although still of some value, do not wholly reflect the present state of the art. Users
21 are cautioned to check to determine that they have the latest edition of any IEEE Standard.

22 In publishing and making this document available, the IEEE is not suggesting or rendering professional or other services
23 for, or on behalf of, any person or entity. Nor is the IEEE undertaking to perform any duty owed by any other person or
24 entity to another. Any person utilizing this, and any other IEEE Standards document, should rely upon the advice of a
25 competent professional in determining the exercise of reasonable care in any given circumstances.

26 Interpretations: Occasionally questions may arise regarding the meaning of portions of standards as they relate to spe-
27 cific applications. When the need for interpretations is brought to the attention of IEEE, the Institute will initiate action
28 to prepare appropriate responses. Since IEEE Standards represent a consensus of concerned interests, it is important to
29 ensure that any interpretation has also received the concurrence of a balance of interests. For this reason, IEEE and the
30 members of its societies and Standards Coordinating Committees are not able to provide an instant response to interpre-
31 tation requests except in those cases where the matter has previously received formal consideration.

32 Comments for revision of IEEE Standards are welcome from any interested party, regardless of membership affiliation
33 with IEEE. Suggestions for changes in documents should be in the form of a proposed change of text, together with
34 appropriate supporting comments. Comments on standards and requests for interpretations should be addressed to:

35 Secretary, IEEE-SA Standards Board
36 445 Hoes Lane
37 P.O. Box 1331
38 Piscataway, NJ 08855-1331
39 USA.
40

41 **Note:** Attention is called to the possibility that implementation of this standard may require use of subject matter
42 covered by patent rights. By publication of this standard, no position is taken with respect to the existence or validity
43 of any patent rights in connection therewith. The IEEE shall not be responsible for identifying patents for which a
44 license may be required by an IEEE standard or for conducting inquiries into the legal validity or scope of those
45 patents that are brought to its attention.

46
47 Authorization to photocopy portions of any individual standard for internal or personal use is granted by the Institute of
48 Electrical and Electronics Engineers, Inc., provided that the appropriate fee is paid to Copyright Clearance Center. To
49 arrange for payment of licensing fee, please contact Copyright Clearance Center, Customer Service, 222 Rosewood
50 Drive, Danvers, MA 01923 USA; (978) 750-8400. Permission to photocopy portions of any individual standard for
51 educational classroom use can also be obtained through the Copyright Clearance Center.
52
53
54

Editors' Foreword

Comments on this draft are encouraged. **PLEASE NOTE: All issues related to IEEE standards presentation style, formatting, spelling, etc. should be addressed, as their presence can often obfuscate relevant technical details.**

By fixing these errors in early drafts, readers can devote their valuable time and energy to comments that materially affect either the technical content of the document or the clarity of that technical content. Comments should not simply state what is wrong, but also what might be done to fix the problem.

Information on 802.1 activities, working papers, and email distribution lists etc. can be found on the 802.1 Website:

<http://ieee802.org/1/>

Use of the email distribution list is not presently restricted to 802.1 members, and the working group has had a policy of considering ballot comments from all who are interested and willing to contribute to the development of the draft. Individuals not attending meetings have helped to identify sources of misunderstanding and ambiguity in past projects. Non-members are advised that the email lists exist primarily to allow the members of the working group to develop standards, and are not a general forum.

Comments on this document may be sent to the 802.1 email reflector, to the editors, or to the Chairs of the 802.1 Working Group and Interworking Task Group.

This draft was prepared by:

David V James
JGG
3180 South Court
Palo Alto, CA 94306
+1.650.494.0926 (Tel)
+1.650.954.6906 (Mobile)
Email: dvj@alum.mit.edu

Chairs of the 802.1 Working Group and Audio/Video Bridging Task Group:

Michael Johas Teener
Chair, 802.1 Audio/Video Bridging Task
Broadcom Corporation
3151 Zanker Road
San Jose, CA
95134-1933
USA
+1 408 922 7542 (Tel)
+1 831 247 9666 (Mobile)
Email: mikejt@broadcom.com

Tony Jeffree
Group Chair, 802.1 Working Group
11A Poplar Grove
Sale
Cheshire
M33 3AX
UK
+44 161 973 4278 (Tel)
+44 161 973 6534 (Fax)
Email: tony@jeffree.co.uk

Introduction to IEEE Std 802.1AS™

(This introduction is not part of P802.1AS, IEEE Standard for Local and metropolitan area networks—Timing and synchronization for time-sensitive applications in bridged local area networks.)

This standard specifies the protocol and procedures used to ensure that the synchronization requirements are met for time sensitive applications, such as audio and video, across bridged and virtual bridged local area networks consisting of LAN media where the transmission delays are fixed and symmetrical; for example, IEEE 802.3 full duplex links. This includes the maintenance of synchronized time during normal operation and following addition, removal, or failure of network components and network reconfiguration. The design is based on concepts developed within the IEEE Std 1588, and is applicable in the context of IEEE Std 802.1D and IEEE Std 802.1Q.

Synchronization to an externally provided timing signal (e.g., a recognized timing standard such as UTC or TAI) is not part of this standard but is not precluded.

Version history

Version	Date	Edits by	Comments
0.082	2005-04-28	DVJ	Updates based on 2005Apr27 meeting discussions
0.085	2005-05-11	DVJ	– Updated list-of-contributors, page numbering, editorial fixes.
0.088	2005-06-03	DVJ	– Application latency scenarios clarified.
0.090	2005-06-06	DVJ	– Misc. editorials in bursting and bunching annex.
0.092	2005-06-10	DVJ	– Extensive cleanup of Clause 5 subscription protocols.
0.121	2005-06-24	DVJ	– Extensive cleanup of clock-synchronization protocols.
0.127	2005-07-04	DVJ	– Pacing descriptions greatly enhanced.
0.200	2007-01-23	DVJ	Removal of non time-sync related information, initial layering proposal.
0.207	2007-02-01	DVJ	Updates based on feedback from Monterey 802.1 meeting. – Common entity terminology; Ethernet type code expanded.
0.216	2007-02-17	DVJ	Updates based on feedback from Chuck Harrison: – linkDelay based only on syntonization to one’s neighbor. – Time adjustments based on observed grandMaster rate differences.
0.224	2007-03-03	DVJ	Updates for whiplash free PLL cascading.
0.230	2007-03-05	DVJ	Major changes: – simplified back-interpolation – first iteration on an Ethernet-PON interface – client-level clock-master and clock-slave interfaces defined
0.243	2007-04-20	DVJ	– Revised GrandSync entity illustrations – General cleanup
0.708	2007-05-30	DVJ	– Simulation results provided within an annex – Extensive code revisions for simplicity & clarity. – Interpolation better described.
—	TBD	—	—

Formats

In many cases, readers may elect to provide contributions in the form of exact text replacements and/or additions. To simplify document maintenance, contributors are requested to use the standard formats and provide checklist reviews before submission. Relevant URLs are listed below:

- General: <http://grouper.ieee.org/groups/msc/WordProcessors.html>
- Templates: <http://grouper.ieee.org/groups/msc/TemplateTools/FrameMaker/>
- Checklist: <http://grouper.ieee.org/groups/msc/TemplateTools/Checks2004Oct18.pdf>

TBDs

Further definitions are needed in the following areas:

- a) Should low-rate *leapSeconds* occupy space in timeSync frames, if this information rarely changes?
- b) What other (than *leapSeconds*) low-rate information should be transferred between stations?
- c) When the grandmaster changes, how should the new grandmaster affect change:
 - 1) Transition immediately to the rate of its reference clock.
 - 2) Transition slowly (perhaps 1ppm/s) between previous and reference clock rates.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

Contents

1		
2		
3	List of figures.....	9
4		
5	List of tables.....	11
6		
7	1. Overview.....	13
8		
9	1.1 Scope	13
10	1.2 Purpose	13
11	1.3 Introduction	13
12		
13	2. References.....	15
14		
15	3. Terms, definitions, and notation	16
16		
17	3.1 Conformance levels	16
18	3.2 Terms and definitions	16
19	3.3 State machines	17
20	3.4 Arithmetic and logical operators	19
21	3.5 Numerical representation.....	19
22	3.6 Field notations	20
23	3.7 Bit numbering and ordering.....	21
24	3.8 Byte sequential formats	22
25	3.9 Ordering of multibyte fields	22
26	3.10 MAC address formats.....	23
27	3.11 Informative notes.....	24
28	3.12 Conventions for C code used in state machines	24
29		
30	4. Abbreviations and acronyms	25
31		
32	5. Architecture overview	27
33		
34	5.1 Application scenarios	27
35	5.2 Design methodology.....	28
36	5.3 Network topologies	29
37	5.4 Information parameters	30
38	5.5 Grandmaster selection	31
39	5.6 Synchronized time distribution.....	34
40	5.7 Distinctions from IEEE Std 1588	39
41		
42	6. GrandSync operation	41
43		
44	6.1 Overview	41
45	6.2 Service interface primitives.....	42
46	6.3 GrandSync state machine	46
47		
48	7. ClockMaster/ClockSlave state machines.....	50
49		
50	7.1 Overview	50
51	7.2 ClockMaster service interfaces.....	51
52	7.3 ClockMaster state machine.....	52
53	7.4 ClockSlave service interfaces.....	55
54	7.5 ClockSlave state machine.....	56

8. Ethernet full duplex (EFDX) state machines.....	59	1
8.1 Overview	59	2
8.2 efdxClockSync frame format	62	3
8.3 TimeSyncRxEfdx state machine	64	4
8.4 TimeSyncTxEfdx state machine.....	67	5
		6
		7
9. Wireless state machines.....	71	8
		9
9.1 Overview	71	10
9.2 Service interface definitions	73	11
9.3 TimeSyncRxR11v state machine	76	12
9.4 TimeSyncTxR11v state machine.....	79	13
		14
10. Ethernet passive optical network (EPON) state machines	82	15
		16
10.1 Overview	82	17
10.2 timeSyncEpon frame format.....	84	18
10.3 TimeSyncRxEpon service interface primitives	85	19
10.4 TimeSyncRxEpon state machine.....	86	20
10.5 TimeSyncTxEpon service interface primitives	89	21
10.6 TimeSyncTxEpon state machine.....	90	22
		23
Annex A (informative) Bibliography	93	24
		25
Annex B (informative) Time-scale conversions	94	26
		27
B.1 Overview	94	28
B.2 TAI and UTC.....	94	29
B.3 NTP and GPS	95	30
B.4 Time-scale conversions	96	31
B.5 Time zones and GMT	97	32
		33
Annex C (informative) Reclocked clockSync requirements.....	98	34
		35
C.1 Cascaded clock considerations	98	36
C.2 Sampling offset/rate conversion	99	37
		38
Annex D (informative) Simulation results (preliminary).....	102	39
		40
D.1 Simulation environment	102	41
D.2 Initialization transients	103	42
D.3 Steady-state interpolation errors.....	104	43
D.4 Steady-state extrapolation errors	105	44
		45
Annex E (informative) Bridging to IEEE Std 1394.....	106	46
		47
E.1 Hybrid network topologies	106	48
		49
Annex F (informative) Time-of-day format considerations	108	50
		51
F.1 Possible time-of-day formats.....	108	52
		53
Annex G (informative) C-code illustrations.....	111	54

- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9
- 10
- 11
- 12
- 13
- 14
- 15
- 16
- 17
- 18
- 19
- 20
- 21
- 22
- 23
- 24
- 25
- 26
- 27
- 28
- 29
- 30
- 31
- 32
- 33
- 34
- 35
- 36
- 37
- 38
- 39
- 40
- 41
- 42
- 43
- 44
- 45
- 46
- 47
- 48
- 49
- 50
- 51
- 52
- 53
- 54

List of figures

		1
		2
Figure 1.1—Topology and connectivity	14	3
Figure 3.1—Bit numbering and ordering	21	4
Figure 3.2—Byte sequential field format illustrations	22	5
Figure 3.3—Multibyte field illustrations	22	6
Figure 3.4—Illustration of fairness-frame structure	23	7
Figure 3.5—MAC address format	23	8
Figure 3.6—48-bit MAC address format.....	24	9
Figure 5.1—Garage jam session.....	27	10
Figure 5.2—Possible looping topology	28	11
Figure 5.3—Synchronized-system topologies	29	12
Figure 5.4—GrandSync service-interface components.....	30	13
Figure 5.5—ClockSelect parameters	31	14
Figure 5.6—Clock-time synchronization flows	34	15
Figure 5.7—Intermediate-bridge responsibilities.....	36	16
Figure 5.8—CLOCK_SLAVE link-delay compensation	37	17
Figure 5.9—CLOCK_MASTER time-sample interpolation	38	18
Figure 5.10—1588 grandmaster preference	39	19
Figure 6.1—GrandSync interface model.....	41	20
Figure 6.2—The clockSelect/clockSync service-interface components	42	21
Figure 6.3—Global-time subfield format	44	22
Figure 6.4—precedence subfields.....	44	23
Figure 6.5— <i>clockID</i> format.....	44	24
Figure 6.7— <i>extraTime</i> format	45	25
Figure 6.8— <i>localTime</i> format	45	26
Figure 6.6— <i>grandTime</i> format.....	45	27
Figure 7.1—ClockMaster interface model	51	28
Figure 8.1—EFDX-link interface model.....	59	29
Figure 8.2—Contents of rxSync/txSync indications	59	30
Figure 8.3—Link-delay compensation.....	60	31
Figure 8.4— <i>efdxClockSync</i> frame format	62	32
Figure 9.1—R11v interface model	71	33
Figure 9.2—Formats of wireless-dependent times.....	71	34
Figure 9.3—802.11v time-synchronization interfaces	72	35
Figure 10.1—EPON topology	82	36
Figure 10.2—EPON interface model	83	37
Figure 10.3—Format of EPON-dependent times	83	38
Figure 10.4— <i>timeSyncEpon</i> frame format	84	39
Figure 10.5— <i>tickTime</i> format	84	40
Figure C.1—Cascading causes sync-interval bunching	98	41
Figure C.2—Reclocking eliminates sync-interval bunching.....	98	42
Figure 3.3—Reclocking localizes sync-interval properties.....	99	43
		44
		45
		46
		47
		48
		49
		50
		51
		52
		53
		54

1	Figure 3.4—Extrapolation for <i>grandTime</i>	99
2	Figure C.5—Extrapolation for <i>grandTime</i>	100
3	Figure C.6—Interpolation for <i>grandTimeA</i>	100
4	Figure C.7—Interpolation of <i>extraTimeD</i>	101
5	Figure D.1—Time-synchronization flows	102
6	Figure D.2—Startup transients with 8 stations	103
7	Figure D.3—Startup transients with 64 stations	103
8	Figure D.4—Time interpolation with 8 stations	104
9	Figure D.5—Time interpolation with 64 stations	104
10	Figure D.6—Time extrapolation with 8 stations	105
11	Figure D.7—Time extrapolation with 64 stations	105
12	Figure E.1—IEEE 1394 leaf domains	106
13	Figure E.2—IEEE 802.3 leaf domains	106
14	Figure E.3—Time-of-day format conversions.....	107
15	Figure E.4—Grandmaster precedence mapping	107
16	Figure F.1—Global-time subfield format	108
17	Figure F.2—IEEE 1394 timer format	108
18	Figure F.3—IEEE 1588 timer format	109
19	Figure F.4—EPON timer format	109
20		
21		
22		
23		
24		
25		
26		
27		
28		
29		
30		
31		
32		
33		
34		
35		
36		
37		
38		
39		
40		
41		
42		
43		
44		
45		
46		
47		
48		
49		
50		
51		
52		
53		
54		

List of tables

Table 3.1—State table notation example 18

Table 3.2—Special symbols and operators..... 19

Table 3.3—Names of fields and sub-fields 20

Table 3.4—*wrap* field values 21

Table 6.1—GrandSync state table 49

Table 7.1—ClockMaster state machine table 54

Table 7.2—ClockSlave state table 58

Table 8.1—Clock-synchronization intervals 63

Table 8.2—TimeSyncRxEfdx state machine table 66

Table 8.3—TimeSyncTxEfdx state machine table 70

Table 9.1—TimeSyncRxR11v state machine table 78

Table 9.2—TimeSyncTxR11v state table 81

Table 10.1—TimeSyncRxEpon state machine table 88

Table 10.2—TimeSyncTxEpon state machine table 92

Table B.1—Time-scale parameters 94

Table B.2—Time-scale conversions 96

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

DVJ Perspective on: Timing and synchronization for time-sensitive applications in bridges local area networks

1. Overview

1.1 Scope

This draft specifies the protocol and procedures used to ensure that the synchronization requirements are met for time sensitive applications, such as audio and video, across bridged and virtual bridged local area networks consisting of LAN media where the transmission delays are fixed and symmetrical; for example, IEEE 802.3 full duplex links. This includes the maintenance of synchronized time during normal operation and following addition, removal, or failure of network components and network reconfiguration. It specifies the use of IEEE 1588 specifications where applicable in the context of IEEE Std 802.1D and IEEE Std 802.1Q. Synchronization to an externally provided timing signal (e.g., a recognized timing standard such as UTC or TAI) is not part of this standard but is not precluded.

1.2 Purpose

This draft enables stations attached to bridged LANs to meet the respective jitter, wander, and time synchronization requirements for time-sensitive applications. This includes applications that involve multiple streams delivered to multiple endpoints. To facilitate the widespread use of bridged LANs for these applications, synchronization information is one of the components needed at each network element where time-sensitive application data are mapped or demapped or a time sensitive function is performed. This standard leverages the work of the IEEE 1588 WG by developing the additional specifications needed to address these requirements.

1.3 Introduction

1.3.1 Background


Ethernet has successfully propagated from the data center to the home, becoming the wired home computer interconnect of choice. However, insufficient support of real-time services has limited Ethernet's success as a consumer audio-video interconnects, where IEEE Std 1394 Serial Bus and Universal Serial Bus (USB) have dominated the marketplace. Success in this arena requires solutions to multiple topics:

- a) Discovery. A controller discovers the proper devices and related streamID/bandwidth parameters to allow the listener to subscribe to the desired talker-sourced stream.
- b) Subscription. The controller commands the listener to establish a path from the talker. Subscription may pass or fail, based on availability of routing-table and link-bandwidth resources.
- c) Synchronization. The distributed clocks in talkers and listeners are accurately synchronized. Synchronized clocks avoid cycle slips and playback-phase distortions.
- d) Pacing. The transmitted classA traffic is paced to avoid other classA traffic disruptions.

1 This draft covers the “Synchronization” component, assuming solutions for the other topics will be devel-
2 oped within other drafts or forums.

3 4 **1.3.2 Interoperability**

5
6 AVB time synchronization interoperates with existing Ethernet, but the scope of time-synchronization is
7 limited to the AVB cloud, as illustrated in Figure 1.1; less-precise time-synchronization services are
8 available everywhere else. The scope of the AVB cloud is limited by a non-AVB capable bridge or a
9 half-duplex link, neither of which can support AVB services.



10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26 **Figure 1.1—Topology and connectivity**

27
28 Separation of AVB devices is driven by the requirements of AVB bridges to support subscription (bandwidth
29 allocation) and pacing of time-sensitive transmissions, as well as time-of-day clock-synchronization.

30 31 **1.3.3 Document structure**

32
33 The clauses and annexes of this working paper are listed below.

- 34 — Clause 1: Overview
 - 35 — Clause 2: References
 - 36 — Clause 3: Terms, definitions, and notation
 - 37 — Clause 4: Abbreviations and acronyms
 - 38 — Clause 5: Architecture overview
 - 39 — Clause 6: GrandSync operation
 - 40 — Clause 7: ClockMaster/ClockSlave state machines
 - 41 — Clause 8: Ethernet full duplex (EFDX) state machines
 - 42 — Clause 9: Wireless state machines
 - 43 — Clause 10: Ethernet passive optical network (EPON) state machines
 - 44 — Annex A: Bibliography
 - 45 — Annex B: Time-scale conversions
 - 46 — Annex C: Reclocked clockSync requirements
 - 47 — Annex D: Simulation results (preliminary)
 - 48 — Annex E: Bridging to IEEE Std 1394
 - 49 — Annex F: Time-of-day format considerations
 - 50 — Annex G: C-code illustrations
- 51
52
53
54

2. References

The following documents contain provisions that, through reference in this working paper, constitute provisions of this working paper. All the standards listed are normative references. Informative references are given in Annex A. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this working paper are encouraged to investigate the possibility of applying the most recent editions of the standards indicated below.

ANSI/ISO 9899-1990, Programming Language-C.^{1,2}

IEEE Std 802.1D-2004, IEEE Standard for Local and Metropolitan Area Networks: Media Access Control (MAC) Bridges.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

¹Replaces ANSI X3.159-1989

²ISO documents are available from ISO Central Secretariat, 1 Rue de Varembe, Case Postale 56, CH-1211, Geneve 20, Switzerland/Suisse; and from the Sales Department, American National Standards Institute, 11 West 42 Street, 13th Floor, New York, NY 10036-8002, USA

3. Terms, definitions, and notation

3.1 Conformance levels

Several key words are used to differentiate between different levels of requirements and options, as described in this subclause.

3.1.1 may: Indicates a course of action permissible within the limits of the standard with no implied preference (“may” means “is permitted to”).

3.1.2 shall: Indicates mandatory requirements to be strictly followed in order to conform to the standard and from which no deviation is permitted (“shall” means “is required to”).

3.1.3 should: An indication that among several possibilities, one is recommended as particularly suitable, without mentioning or excluding others; or that a certain course of action is preferred but not necessarily required; or that (in the negative form) a certain course of action is deprecated but not prohibited (“should” means “is recommended to”).

3.2 Terms and definitions

For the purposes of this working paper, the following terms and definitions apply. The Authoritative Dictionary of IEEE Standards Terms [B2] should be referenced for terms not defined in the clause.

3.2.1 bridge: A functional unit interconnecting two or more networks at the data link layer of the OSI reference model.

3.2.2 clock master: A bridge or end station that provides the link clock reference.

3.2.3 clock slave: A bridge or end station that tracks the link clock reference provided by the clock master.

3.2.4 cyclic redundancy check (CRC): A specific type of frame check sequence computed using a generator polynomial.

3.2.5 grandmaster: Refers to the station that is selected to provide the network time reference.

3.2.6 link: A unidirectional channel connecting adjacent stations (half of a span).

3.2.7 listener: A sink of a stream, such as a television or acoustic speaker.

3.2.8 local area network (LAN): A communications network designed for a small geographic area, typically not exceeding a few kilometers in extent, and characterized by moderate to high data transmission rates, low delay, and low bit error rates.

3.2.9 MAC client: The layer entity that invokes the MAC service interface.

3.2.10 medium (plural: media): The material on which information signals are carried; e.g., optical fiber, coaxial cable, and twisted-wire pairs.

3.2.11 medium access control (MAC) sublayer: The portion of the data link layer that controls and mediates the access to the network medium. In this working paper, the MAC sublayer comprises the MAC datapath sublayer and the MAC control sublayer.

3.2.12 network: A set of communicating stations and the media and equipment providing connectivity among the stations.

3.2.13 plug-and-play: The requirement that a station perform classA transfers without operator intervention (except for any intervention needed for connection to the cable).

3.2.14 protocol implementation conformance statement (PICS): A statement of which capabilities and options have been implemented for a given Open Systems Interconnection (OSI) protocol.

3.2.15 span: A bidirectional channel connecting adjacent stations (two links).

3.2.16 station: A device attached to a network for the purpose of transmitting and receiving information on that network.

3.2.17 topology: The arrangement of links and stations forming a network, together with information on station attributes.

3.2.18 transmit (transmission): The action of a station placing a frame on the medium.

3.2.19 unicast: The act of sending a frame addressed to a single station.

3.3 State machines

3.3.1 State machine behavior

The operation of a protocol can be described by subdividing the protocol into a number of interrelated functions. The operation of the functions can be described by state machines. Each state machine represents the domain of a function and consists of a group of connected, mutually exclusive states. Only one state of a function is active at any given time. A transition from one state to another is assumed to take place in zero time (i.e., no time period is associated with the execution of a state), based on some condition of the inputs to the state machine.

The state machines contain the authoritative statement of the functions they depict. When apparent conflicts between descriptive text and state machines arise, the order of precedence shall be formal state tables first, followed by the descriptive text, over any explanatory figures. This does not override, however, any explicit description in the text that has no parallel in the state tables.

The models presented by state machines are intended as the primary specifications of the functions to be provided. It is important to distinguish, however, between a model and a real implementation. The models are optimized for simplicity and clarity of presentation, while any realistic implementation might place heavier emphasis on efficiency and suitability to a particular implementation technology. It is the functional behavior of any unit that has to match the standard, not its internal structure. The internal details of the model are useful only to the extent that they specify the external behavior clearly and precisely.

3.3.2 State table notation

NOTE—The following state machine notation was used within 802.17, due to the exactness of C-code conditions and the simplicity of updating table entries (as opposed to 2-dimensional graphics). Early state table descriptions can be converted (if necessary) into other formats before publication.

Each row of the table is preferably provided with a brief description of the condition and/or action for that row. The descriptions are placed after the table itself, and linked back to the rows of the table using numeric tags.

State machines may be represented in tabular form. The table is organized into two columns: a left hand side representing all of the possible states of the state machine and all of the possible conditions that cause transitions out of each state, and the right hand side giving all of the permissible next states of the state machine as well as all of the actions to be performed in the various states, as illustrated in Table 3.1. The syntax of the expressions follows standard C notation (see 3.12). No time period is associated with the transition from one state to the next.

Table 3.1—State table notation example

Current		Row	Next	
state	condition		action	state
START	sizeofMacControl > spaceInQueue	1	—	START
	passM == 0	2		
	—	3	TransmitFromControlQueue();	FINAL
FINAL	SelectedTransferCompletes()	4	—	START
	—	5	—	FINAL

Row 3.1-1: Do nothing if the size of the queued MAC control frame is larger than the PTQ space.

Row 3.1-2: Do nothing in the absence of MAC control transmission credits.

Row 3.1-3: Otherwise, transmit a MAC control frame.

Row 3.1-4: When the transmission completes, start over from the initial state (i.e., START).

Row 3.1-5: Until the transmission completes, remain in this state.

Each combination of current state, next state, and transition condition linking the two is assigned to a different row of the table. Each row of the table, read left to right, provides: the name of the current state; a condition causing a transition out of the current state; an action to perform (if the condition is satisfied); and, finally, the next state to which the state machine transitions, but only if the condition is satisfied. The symbol “—” signifies the default condition (i.e., operative when no other condition is active) when placed in the condition column, and signifies that no action is to be performed when placed in the action column. Conditions are evaluated in order, top to bottom, and the first condition that evaluates to a result of TRUE is used to determine the transition to the next state. If no condition evaluates to a result of TRUE, then the state machine remains in the current state. The starting or initialization state of a state machine is always labeled “START” in the table (though it need not be the first state in the table). Every state table has such a labeled state.

Each row of the table is preferably provided with a brief description of the condition and/or action for that row. The descriptions are placed after the table itself, and linked back to the rows of the table using numeric tags.

3.4 Arithmetic and logical operators

In addition to commonly accepted notation for mathematical operators, Table 3.2 summarizes the symbols used to represent arithmetic and logical (boolean) operations. Note that the syntax of operators follows standard C notation (see 3.12).

Table 3.2—Special symbols and operators

Printed character	Meaning
&&	Boolean AND
	Boolean OR
!	Boolean NOT (negation)
&	Bitwise AND
	Bitwise OR
^	Bitwise XOR
<=	Less than or equal to
>=	Greater than or equal to
==	Equal to
!=	Not equal to
=	Assignment operator
//	Comment delimiter

3.5 Numerical representation

NOTE—The following notation was taken from 802.17, where it was found to have benefits:

- The subscript notation is consistent with common mathematical/logic equations.
- The subscript notation can be used consistently for all possible radix values.

Decimal, hexadecimal, and binary numbers are used within this working paper. For clarity, decimal numbers are generally used to represent counts, hexadecimal numbers are used to represent addresses, and binary numbers are used to describe bit patterns within binary fields.

Decimal numbers are represented in their usual 0, 1, 2, ... format. Hexadecimal numbers are represented by a string of one or more hexadecimal (0-9,A-F) digits followed by the subscript 16, except in C-code contexts, where they are written as 0x123EF2 etc. Binary numbers are represented by a string of one or more binary (0,1) digits, followed by the subscript 2. Thus the decimal number “26” may also be represented as “1A₁₆” or “11010₂”.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

MAC addresses and OUI/EUI values are represented as strings of 8-bit hexadecimal numbers separated by hyphens and without a subscript, as for example “01-80-C2-00-00-15” or “AA-55-11”.

3.6 Field notations

3.6.1 Use of italics

All field names or variable names (such as *level* or *myMacAddress*), and sub-fields within variables (such as *thisState.level*) are italicized within text, figures and tables, to avoid confusion between such names and similarly spelled words without special meanings. A variable or field name that is used in a subclause heading or a figure or table caption is also italicized. Variable or field names are not italicized within C code, however, since their special meaning is implied by their context. Names used as nouns (e.g., *subclassA0*) are also not italicized.

3.6.2 Field conventions

This working paper describes fields within packets or included in state-machine state. To avoid confusion with English names, such fields have an italics font, as illustrated in Table 3.3.

Table 3.3—Names of fields and sub-fields

Name	Description
<i>newCRC</i>	Field within a register or frame
<i>thisState.level</i>	Sub-field within field <i>thisState</i>
<i>thatState.rateC[n].c</i>	Sub-field within array element <i>rateC[n]</i>

Run-together names (e.g., *thisState*) are used for fields because of their compactness when compared to equivalent underscore-separated names (e.g., *this_state*). The use of multiword names with spaces (e.g., “This State”) is avoided, to avoid confusion between commonly used capitalized key words and the capitalized word used at the start of each sentence.

A sub-field of a field is referenced by suffixing the field name with the sub-field name, separated by a period. For example, *thisState.level* refers to the sub-field *level* of the field *thisState*. This notation can be continued in order to represent sub-fields of sub-fields (e.g., *thisState.level.next* is interpreted to mean the sub-field *next* of the sub-field *level* of the field *thisState*).

Two special field names are defined for use throughout this working paper. The name *frame* is used to denote the data structure comprising the complete MAC sublayer PDU. Any valid element of the MAC sublayer PDU, can be referenced using the notation *frame.xx* (where *xx* denotes the specific element); thus, for instance, *frame.serviceDataUnit* is used to indicate the *serviceDataUnit* element of a frame.

Unless specifically specified otherwise, reserved fields are reserved for the purpose of allowing extended features to be defined in future revisions of this working paper. For devices conforming to this version of this working paper, nonzero reserved fields are not generated; values within reserved fields (whether zero or nonzero) are to be ignored.

3.6.3 Field value conventions

This working paper describes values of fields. For clarity, names can be associated with each of these defined values, as illustrated in Table 3.4. A symbolic name, consisting of upper case letters with underscore separators, allows other portions of this working paper to reference the value by its symbolic name, rather than a numerical value.

Table 3.4—wrap field values

Value	Name	Description
0	STANDARD	Standard processing selected
1	SPECIAL	Special processing selected
2,3	—	Reserved

Unless otherwise specified, reserved values allow extended features to be defined in future revisions of this working paper. Devices conforming to this version of this working paper do not generate nonzero reserved values, and process reserved fields as though their values were zero.

A field value of TRUE shall always be interpreted as being equivalent to a numeric value of 1 (one), unless otherwise indicated. A field value of FALSE shall always be interpreted as being equivalent to a numeric value of 0 (zero), unless otherwise indicated.

3.7 Bit numbering and ordering

Data transfer sequences normally involve one or more cycles, where the number of bytes transmitted in each cycle depends on the number of byte lanes within the interconnecting link. Data byte sequences are shown in figures using the conventions illustrated by Figure 3.1, which represents a link with four byte lanes. For multi-byte objects, the first (left-most) data byte is the most significant, and the last (right-most) data byte is the least significant.




Figure 3.1—Bit numbering and ordering

Figures are drawn such that the counting order of data bytes is from left to right within each cycle, and from top to bottom between cycles. For consistency, bits and bytes are numbered in the same fashion.

NOTE—The transmission ordering of data bits and data bytes is not necessarily the same as their counting order; the translation between the counting order and the transmission order is specified by the appropriate reconciliation sublayer.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

3.8 Byte sequential formats

Figure 3.2 provides an illustrative example of the conventions to be used for drawing frame formats and other byte sequential representations. These representations are drawn as fields (of arbitrary size) ordered along a vertical axis, with numbers along the left sides of the fields indicating the field sizes in bytes. Fields are drawn contiguously such that the transmission order across fields is from top to bottom. The example shows that *field1*, *field2*, and *field3* are 1-, 1- and 6-byte fields, respectively, transmitted in order starting with the *field1* field first. As illustrated on the right hand side of Figure 3.2, a multi-byte field represents a sequence of ordered bytes, where the first through last bytes correspond to the most significant through least significant portions of the multi-byte field, and the MSB of each byte is drawn to be on the left hand side.




Figure 3.2—Byte sequential field format illustrations

NOTE—Only the left-hand diagram in Figure 3.2 is required for representation of byte-sequential formats. The right-hand diagram is provided in this description for explanatory purposes only, for illustrating how a multi-byte field within a byte sequential representation is expected to be ordered. The tag “Transmission order” and the associated arrows are not required to be replicated in the figures.

3.9 Ordering of multibyte fields

In many cases, bit fields within byte or multibyte objects are expanded in a horizontal fashion, as illustrated in the right side of Figure 3.3. The fields within these objects are illustrated as follows: left-to-right is the byte transmission order; the left-through-right bits are the most significant through least significant bits respectively.

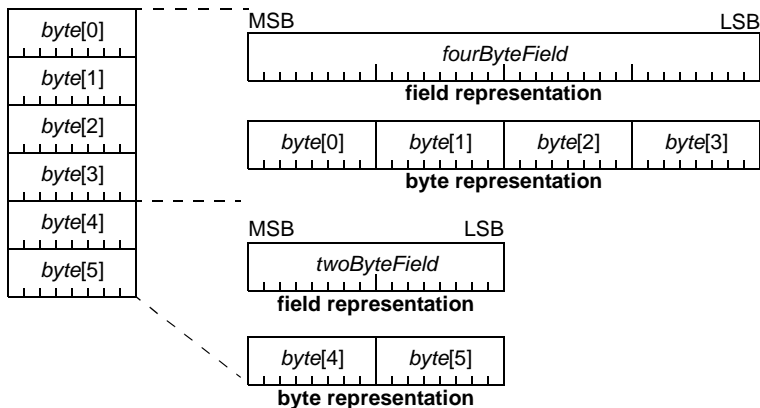


Figure 3.3—Multibyte field illustrations

The first *fourByteField* can be illustrated as a single entity or a 4-byte multibyte entity. Similarly, the second *twoByteField* can be illustrated as a single entity or a 2-byte multibyte entity.

To minimize potential for confusion, four equivalent methods for illustrating frame contents are illustrated in Figure 3.4. Binary, hex, and decimal values are always shown with a left-to-right significance order, regardless of their bit-transmission order.




Figure 3.4—Illustration of fairness-frame structure

3.10 MAC address formats

The format of MAC address fields within frames is illustrated in Figure 3.5.

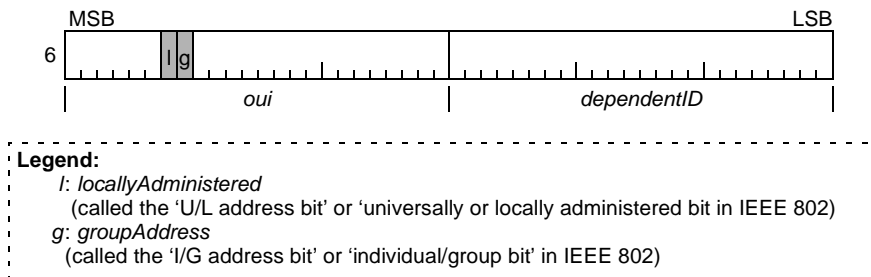


Figure 3.5—MAC address format

3.10.1 oui: A 24-bit organizationally unique identifier (OUI) field supplied by the IEEE/RAC for the purpose of identifying the organization supplying the (unique within the organization, for this specific context) 24-bit *dependentID*. (For clarity, the *locallyAdministered* and *groupAddress* bits are illustrated by the shaded bit locations.)

3.10.2 dependentID: An 24-bit field supplied by the *oui*-specified organization. The concatenation of the *oui* and *dependentID* provide a unique (within this context) identifier.

To reduce the likelihood of error, the mapping of OUI values to the *oui/dependentID* fields are illustrated in Figure 3.6. For the purposes of illustration, specific OUI and *dependentID* example values have been assumed. The two shaded bits correspond to the *locallyAdministered* and *groupAddress* bit positions illustrated in Figure 3.5.




Figure 3.6—48-bit MAC address format

3.11 Informative notes

Informative notes are used in this working paper to provide guidance to implementers and also to supply useful background material. Such notes never contain normative information, and implementers are not required to adhere to any of their provisions. An example of such a note follows.

NOTE—This is an example of an informative note.

3.12 Conventions for C code used in state machines

Many of the state machines contained in this working paper utilize C code functions, operators, expressions and structures for the description of their functionality. Conventions for such C code can be found in Annex G.

4. Abbreviations and acronyms

This working paper contains the following abbreviations and acronyms:

		1
		2
		3
		4
		5
AP	access point	6
AV	audio/video	7
AVB	audio/video bridging	8
AVB network	audio/video bridged network	9
		10
BER	bit error ratio	11
		12
BMC	best master clock	13
		14
BMCA	best master clock algorithm	15
		16
CRC	cyclic redundancy check	17
		18
EFDX	Ethernet full duplex	19
		20
EPON	Ethernet passive optical network	21
		22
FIFO	first in first out	23
		24
IEC	International Electrotechnical Commission	25
		26
IEEE	Institute of Electrical and Electronics Engineers	27
		28
IETF	Internet Engineering Task Force	29
		30
ISO	International Organization for Standardization	31
		32
ITU	International Telecommunication Union	33
		34
LAN	local area network	35
		36
LSB	least significant bit	37
		38
MAC	medium access control	39
		40
MAN	metropolitan area network	41
		42
MSB	most significant bit	43
		44
OSI	open systems interconnect	45
		46
PDU	protocol data unit	47
		48
PHY	physical layer	49
		50
PLL	phase-locked loop	51
		52
PTP	Precision Time Protocol	53
		54
R11V	radio 802.11v	
RFC	request for comment	
RPR	resilient packet ring	
VOIP	voice over internet protocol	

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

5. Architecture overview

5.1 Application scenarios

5.1.1 Garage jam session

As an illustrative example, consider AVB usage for a garage jam session, as illustrated in Figure 5.1. The audio inputs (microphone and guitar) are converted, passed through a guitar effects processor, two bridges, mixed within an audio console, return through two bridges, and return to the ear through headphones.




Figure 5.1—Garage jam session

Using Ethernet within such systems has multiple challenges: low-latency and tight time-synchronization. Tight time synchronization is necessary to avoid cycle slips when passing through multiple processing components and (ultimately) to avoid under-run/over-run at the final D/A converter's FIFO. The challenge of low-latency transfers is being addressed in other forums and is outside the scope of this draft.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

5.1.2 Looping topologies

Bridged Ethernet networks currently have no loops, but bridging extensions are contemplating looping topologies. To ensure longevity of this standard, the time-synchronization protocols are tolerant of looping topologies that could occur (for example) if the dotted-line link were to be connected in Figure 5.2.

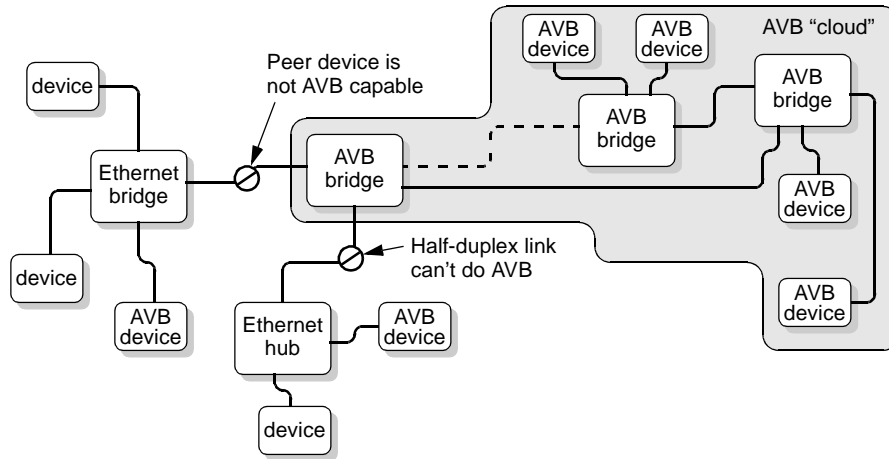


Figure 5.2—Possible looping topology

Separation of AVB devices is driven by the requirements of AVB bridges to support subscription (bandwidth allocation) and pacing of time-sensitive transmissions, as well as time-of-day clock-synchronization.

5.2 Design methodology

5.2.1 Assumptions

This working paper specifies a protocol to synchronize independent timers running on separate stations of a distributed networked system, based on concepts specified within IEEE Std 1588-2002. Although a high degree of accuracy and precision is specified, the technology is applicable to low-cost consumer devices. The protocols are based on the following design assumptions:

- a) Each end station and intermediate bridges provide independent clocks.
- b) All clocks are accurate, typically to within $\pm 100\text{PPM}$.
- c) Details of the best time-synchronization protocols are physical-layer dependent.

5.2.2 Objectives

With these assumptions in mind, the time synchronization objectives include the following:

- a) Precise. Multiple timers can be synchronized to within 10's of nanoseconds.
- b) Inexpensive. For consumer AVB devices, the costs of synchronized timers are minimal. (GPS, atomic clocks, or 1PPM clock accuracies would be inconsistent with this criteria.)
- c) Scalable. The protocol is independent of the networking technology. In particular:
 - 1) Cyclical physical topologies are supported.
 - 2) Long distance links (up to 2 km) are allowed.
- d) Plug-and-play. The system topology is self-configuring; no system administrator is required.

5.2.3 Strategies

Strategies used to meet these objectives include the following:

- a) Precision is achieved by calibrating and adjusting *grandTime* clocks.
 - 1) Offsets. Offset value adjustments eliminate immediate clock-value errors.
 - 2) Rates. Rate value adjustments reduce long-term clock-drift errors.
- b) Simplicity is achieved by the following:
 - 1) Concurrence. Most configuration and adjustment operations are performed concurrently.
 - 2) Firmware friendly. Clock offset and rate adjustments can be performed by low-rate firmware; analog PLLs are unnecessary within bridges (although necessary within some applications).
 - 3) Frequent. Frequent ~100 Hz interchanges eliminates needs for expensive/precise clocks.

5.3 Network topologies

Clock synchronization involves streaming of timing information from a grandmaster clock to one or more slave clocks. Although primarily intended for non-cyclical physical topologies, the synchronization protocols also function correctly on cyclical physical topologies, by activating only a non-cyclical subset of the physical topology.

The grandmaster station (GM) provides a time reference to locally attached clock-slave station (S), as illustrated in Figure 5.3a. Bridges synchronize their clock-master (M) ports to their clock-slave (S) port, via internal communications. These clock-master (M) ports typically connect to another bridge’s clock-slave (S) port, or to an attached end-station clock-slave station (ES).

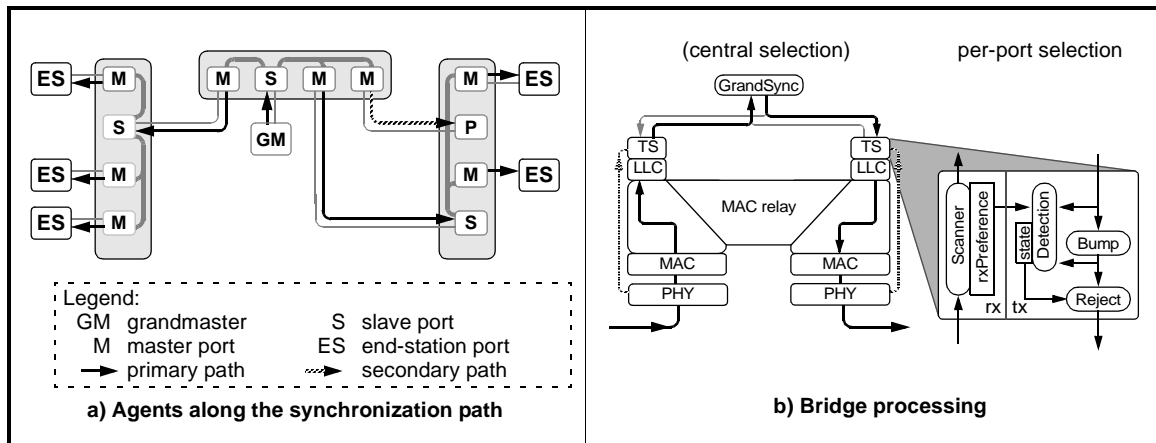


Figure 5.3—Synchronized-system topologies

In concept, network-wide clock-synchronization protocol starts with the selection of the reference-clock station, called a grandmaster station (oftentimes abbreviated as grandmaster). Every AVB-capable station is grandmaster capable, but only one is selected to become the grandmaster station within each network grandmaster selection involves the transmission of ClockSelect messages between grandmaster capable stations.

In the presence of redundant loops, a bridge port can also be placed in a passive (P) state, wherein it monitors ClockSelect messages but does not synchronize its clock to its associated clock-master port.

After grandmaster selection, time-reference information flows from this selected grandmaster station to attached clock-slave stations. Thus, time-synchronization involve two subservices, as follows:

- a) Selection. Topologies are pruned into spanning tree with the grandmaster at the root (see 5.5).
- b) Distribution. Synchronized time is distributed through the spanning tree from the root (see 5.6).

5.4 Information parameters

Clock-synchronization information includes clock-selection parameters (for selecting the grandmaster) and clock-synchronization parameters (for synchronizing between clock-master and clock-slave ports), as illustrated in Figure 5.4. A portion of the clock-synchronization parameters are media-independent and pass across the bridge-internal service interface; the remainder are link-media dependent and used for calibration of clock-master to clock-slave port delays.

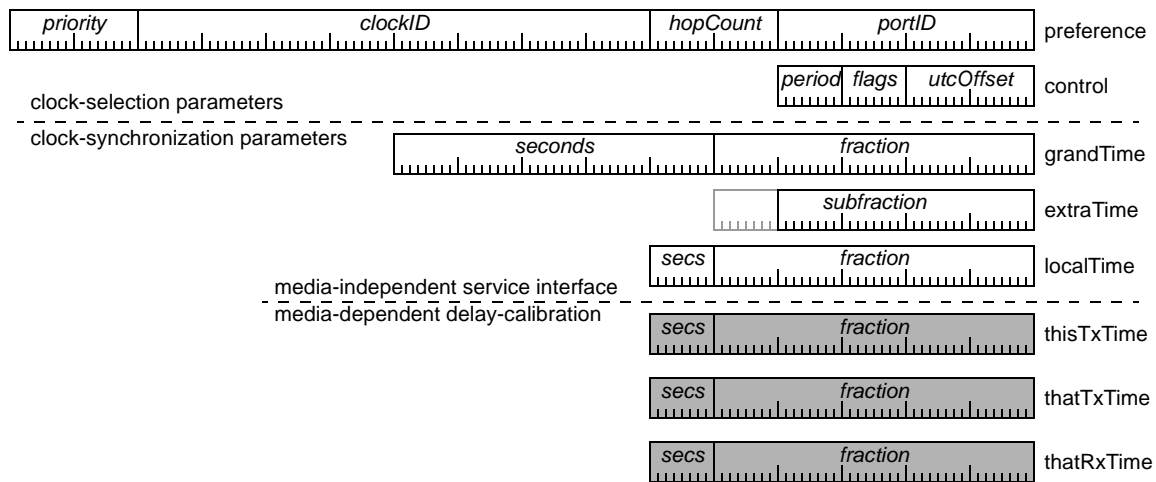


Figure 5.4—GrandSync service-interface components

NOTE—Depending on the media specification, the clock-selection and clock-synchronization parameters could be transmitted in distinct clockSelect and clockSync packets, or could be combined into a single clockSync packet.

The concatenation of *priority*, *clockID*, *hopCount*, and *portID* values forms a preference value for the selection of the preferred grandmaster candidate.

The {*grandTime*,*localTime*,*extraTime*} triad percolates from the grand-master to the clock-slave stations, with adjustments along the way, to provide clock-slave stations with estimates of the grandmaster’s clock time. The *grandTime* represents a station’s estimate of the grandmaster’s clock sampled at that stations’s *localTime* instant. The *extraTime* value represents a cumulative deviation error that (for the purposes of accuracy and responsiveness) is maintained separately from the grandTime value.

The remaining media-dependent parameters allow a bridge’s receive-port to calibrate and compensate for delays introduced by the receive link of the span connecting its neighbor. For full-duplex Ethernet, these include *thatTxTime* and *thatRxTime* (snapshots taken on opposing-link frames), and *thisTxTime* (a snapshot taken on the the neighbor’s transmission). Different values are applicable to alternative media types.

5.5 Grandmaster selection

5.5.1 Selection overview

Intermediate bridge entities are responsible for providing standard/centralized preference-selection and media-specific/per-port rejection entities, as illustrated in Figure 5.3b. The central GrandSync preference-selection entity is responsible for selecting and echoing only the best ClockSelect messages. The per-port rejection entity is responsible for discarding (unnecessary) reverse-flow ClockSelect messages.

5.5.2 ClockSelect-message content

As part of the grandmaster selection process, ClockSelect messages are generated by clock-master capable stations; the ClockSelect messages transport clock-preference values (see Figure 5.5) to other stations. Only ClockSelect messages with the best observed clock-preference value are forwarded to neighbor stations, allowing the overall best-preference value to be ultimately selected and known by all. The station with that best preference value is called the grandmaster.

This ClockSelect information content is conceptually similar to the format of the spanning-tree preference value, but larger identifiers and additional values are provided for the sake of generality, as illustrated in Figure 5.5.

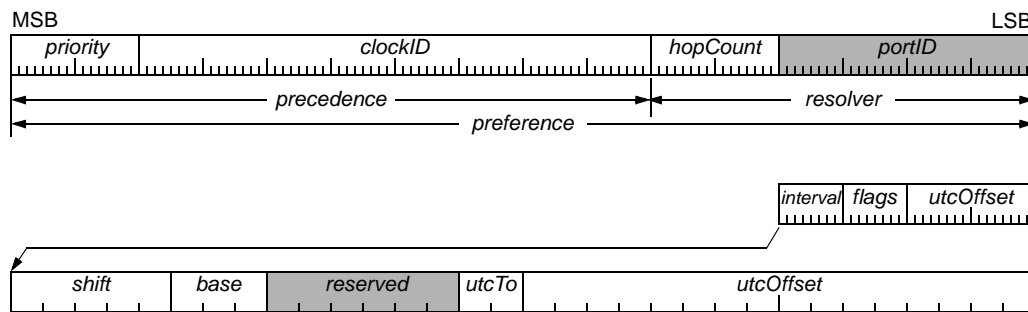


Figure 5.5—ClockSelect parameters

The portID portion of this information (shaded in dark gray) represents an internal field that is only passed between bridge components; the remaining fields are also present within transmitted frames.

The grandmaster *preference* (a concatenation of multiple fields) identifies the best grand-master station. The *interval* value specifies the nominal link-dependent delay between successive ClockSelect messages; *flags* provides control and warning indications; *utcOffset* specifies the difference between UTC and TAI time (commonly called leap-seconds).

5.5.3 ClockSelect preference content

The fields within the preference group of an ClockSelect message (see Figure 5.5) are optimized for spanning-tree topologies, which have no loops or parallel paths. However, the tie-breaker values are sufficient to break loops within shortest-path bridging topologies, so that the clock-distribution paths (based on grandmaster selection) could be optimized independently from the default broadcast distribution paths.

The value of each station's *priority* field can be set by the application, for the purposes of biasing the grand-master preference in an application-specific manner. The *clockID* field is a globally unique number assigned to each station. The *hopCount* represents the distance from the grandmaster. The internal *portID* identifies the port where the ClockSelect message was received and is not transmitted between stations.

5.5.4 Other ClockSelect content

The *interval* field specifies the source-station's nominal interval between ClockSelect/clockSync frames, and is used for timeout purposes. For compactness, its value is encoded as simple floating point value:

```
value = (shift >= 0) ? (base << shift) : (base >> -shift);
```

The control-group fields supports the management of leap-second updates and timeouts. The *utcOffset* field represents the grand-master's vision of the difference between UTC (wall clock) time and TAI (continuous) times. The two-bit difference of *utcTo* and *utcOffset*, $\text{diff} = (\text{utcTo} - \text{utcOffset}) \% 4$, specifies the *utcOffset* properties, as follows:

Value	Property
0	Known and stable
1	Increment at midnight UTC
2	Decrement at midnight UTC
3	Unknown

NOTE—When the value of *utcOffset* is decremented, the event of “midnight” can occur twice, one second apart. Sending an idempotent *utcTo* encoding (that specifies the changed value, not the change amount), eliminates undesirable effects that might otherwise occur due to such as “double-clocking” event.

NOTE—To ensure correctness, an exact equation (not simply an English statement) should be provided to ensure proper/consistent updating of UTC time.

5.5.5 Central ClockSelect-message processing

The central/standard GrandSync entity (see Figure 5.3b) is responsible for processing ClockSelect PDUs provided by attached ports, as summarized in Equation 5.1. If the ClockSelect message preference equals or exceeds the previously saved *sync.preference* value: *sync.preference* is updated with the ClockSelect-message preference, the ClockSelect-message timeout is cleared, and the ClockSelect message is echoed to all bridge ports. A timeout clears the *sync.preference* history in the absence of expected (normally periodic) ClockSelect messages.

```

// Performed by the GrandSync entity                                     (5.1)
// Compare(x,y) operates on multiple-precision arguments x and y
// The returned integer result is:
    1 if x-y > 0
    0 if x-y == 0
   -1 if x-y < 0
while (FOREVER) { // Check constantly
    clockSelect = Dequeue(GS_RX); // Get ClockSelect message
    if (clockSelect != NULL) { // if one is available
        gsPreference = GsPreference(clockSelect); // Form preference value
        if (Compare(gsPreference, sync.preference) <= 0) { // The best preference
            sync.preference = gsPreference; // is observed and saved
            Enqueue(GS_TX, clockSelect); // Echoed the PDU to others
            sync.lastTime = currentTime; // Restart the timeout
            sync.timeout = GsTimeout(clockSelect.interval); // Form preference value
        }
    }
    if (currentTime >= sync.lastTime + sync.timeout) // Inactivity timeout;
        grandPreference = ONES; // clear saved preference
}

```


5.5.6 Per-port received ClockSelect-message processing

The per-port/media-dependent TS.rx entity (see Figure 5.3b) is responsible for processing rxClockSelect PDUs provided by attached ports, as summarized in Equation 5.2. If the rxClockSelect-PDU preference equals or exceeds the previously saved *port.preference* value: *port.preference* is updated with the rxClockSelect-PDU preference and the ClockSelect-message timeout is cleared. Regardless, the (media dependent) rxClockSelect-PDU is converted to a generic clockSelect-PDU and sent to the central GrandSync entity. A timeout clears the *port.preference* history in the absence of expected (normally periodic) ClockSelect messages.

```

// Performed by the TS.receive entity (5.2)
while (FOREVER) {
    rxClockSelect = Dequeue(ES_RX); // Check constantly
    if (rxClockSelect != NULL) { // Get a ClockSelect
        rxPreference = RxPreference(rxClockSelect); // if one is available
        if (Compare(rxPreference, port.preference) < 0) { // Form preference value
            port.preference = rxPreference; // The best preference
            port.lastTime = currentTime; // is observed and saved
        } // Restart the timeout
        Enqueue(GS_RX, EsToGs(rxClockSelect)); // Send PDU to GrandSync
    }
    if (currentTime >= port.lastTime + rxTimeout) // Inactivity timeout;
        port.preference = ONES; // clear saved preference
}

```

5.5.7 Per-port transmit ClockSelect-message processing

The per-port TS.tx entity (see Figure 5.3b) is responsible for setting the port state and selectively forwarding clockSelect PDUs provided by the GrandSync entity, as summarized in Equation 5.3. If *port.preference* compares favorably to the clockSelect-PDU (indicating this is the best-preference port), the *port.state* is set to CLOCK_SLAVE and ClockSelect messages are discard. Otherwise, *port.state* (and its associated behaviors) are dependent on comparisons between port and to-be-transmitted PDU preferences.

```

// Performed by the TS.transmit entity (5.3)
while (FOREVER) {
    clockSelect = Dequeue(GS_TX); // Check constantly
    if (clockSelect != NULL) { // Get received ClockSelect
        if (clockSelect.hopCount == HOP_LIMIT) return; // if one is available
        g0Preference = GsPreference(clockSelect); // Ignore aged frames
        gsClockSelect.hopCount += 1; // Form preference value
        g1Preference = GsPreference(clockSelect); // Increment hop count
        if (Compare(port.preference, g0Preference) <= 0) // Form preference value
            port.state = CLOCK_SLAVE, return; // Rx preference is best;
        switch (Compare(g1Preference, port.preference)) { // Rx ClockSelect and time
            case -1: // Tx preference compare
                port.state = CLOCK_PASSIVE, return; // Rx ClockSelect, ignore time
            case 1:
                port.state = CLOCK_MASTER, break; // Tx ClockSelect and time
            case 0:
                port.state = CLOCK_IDLELED, break; // Tx ClockSelect, block time
        }
        Enqueue(ES_TX, GsToEs(clockSelect)); // Transmit the ClockSelect
    }
}

```

5.6 Synchronized time distribution

5.6.1 Clock-time synchronization flows

As background for understanding, consider the flow of clock-time synchronization information within a simple/stable three-bridge topology, as illustrated in Figure 5.6. The clock-time information flows from the selected application-level ClockSource entity (located on the grandmaster station), through intermediate bridges (when present), and terminates at application-level ClockSink entities.

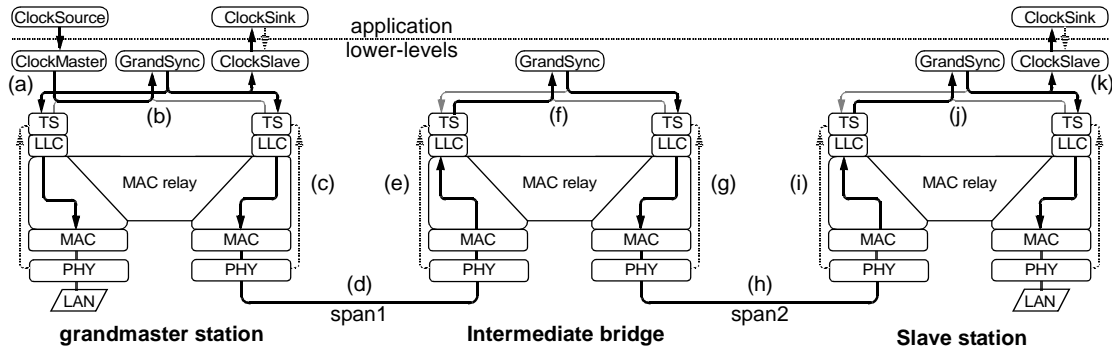


Figure 5.6—Clock-time synchronization flows

The clock-master station (containing the ClockSource entity) and the clock-slave station (containing the ClockSink entity) are illustrated as multipurpose bridges; either could also be end stations (not illustrated).

Time synchronization yields distributed but closely-matched *grandTime* values within stations and bridges. No attempt is made to eliminate intermediate jitter with bridge-resident jitter-reducing phase-locked loops (PLLs) but application-level phase locked loops are expected to filter high-frequency jitter from the supplied *grandTime* values.

Maintaining an accurate time reference relies on the presence of accurate time-stamp hardware capabilities in or near the media-dependent PHY. A bypass path (illustrated as hashed PHY-to-TS lines within Figure 5.6) allows the time-stamp information to be affiliated with the arriving clockSync-frame information, before the PDU is processed by the time-synchronization (TS) entity above the MAC. A similar bypass path is also required at the transmitter, so that an accurate time-stamp of a transmitted frame can be placed into the next-transmitted clockSync frame.

This grandmaster station comprises client-level ClockSink as well as ClockSource entities. The ClockSink entity allows the application-level clock to be synchronized to the network clock, whenever another station becomes the grandmaster station. (The ClockSource entity on the grandmaster station provides the network-synchronized clock-time reference.)

5.6.2 Steps in the *grandTime* flow

Processing steps along the path of the time-synchronization flow of Figure 5.6 include the following:

- a) The ClockSource's time-reference parameters are converted into a standard format, supplemented with additional parameters (such as the station-local arrival time), packaged into a PDU, and sent to the GrandSync entity.
- b) The GrandSync entity echoes the time-synchronization information from (what it determines to be) the preferred port. Information from lower-preference ports is continuously monitored to detect preference changes (typically due to attach or detach of clock-master capable stations). Clock-time information in PDUs from lower-preference ports is ignored.
- c) Clock-time parameters within the echoed PDUs are saved in transmit-port time-value array. When the next clockSync frame is transmitted, the transmission time is computed based on these array values. The computed transmission time is queued for inclusion in the following clockSync transmission.
- d) The clockSync frame travels over span1 to the intermediate bridge, incurring a media-dependent transmission *delay* and arriving at (station-local) time *rxTime*.
- e) The (station-local) frame-arrival snapshot value, *rxTime*, is compensated by subtracting the estimated transmission delay to create an *rcTime* value, where:

$$rcTime = rxTime - delay0$$
 where *delay0* is an estimate of span *delay*.
 The *grandTime* value is extracted from the arriving frame; the {*grandTime*,*rcTime*} time-pair values are encapsulated within a PDU; that PDU is sent to the GrandSync entity.
- f) The GrandSync entity echoes the clockSync PDU provided by the preferred port, as in step (c).
- g) Time parameters within the echoed PDUs are saved and eventually transmitted, as in step (d).
- h) The clockSync frame travels over span2 to the final bridge, incurring a media-dependent transmission *delay* and arriving at (station-local) time *rxTime*.
- i) The (station-local) frame-arrival snapshot value, *rxTime*, is compensated by subtracting the estimated transmission delay to create an *rcTime* value, as in step (d).
 The {*grandTime*,*rcTime*} affiliations are placed in a PDU; that PDU is sent to the GrandSync entity.
- j) The GrandSync entity echoes the clockSync PDU from the preferred port, as in step (c).
- k) The GrandSync's time-reference parameters are converted from the standard PDU format, unnecessary parameters are discarded, and *grandTime* is communicated to the GrandSync entity.

5.6.3 Intermediate bridge entities

Entities within the intermediate bridge (see Figure 5.6) are responsible for performing three distinct (and largely decoupled) functions, as illustrated in Figure 5.7. A port in the CLOCK_SLAVE state is responsible for compensating for link-transmission delays between this station and its neighbor, as described in 5.6.4. The GrandSync entity is responsible for selecting the clockSync PDUs from the grandmaster station; only thus selected PDUs are forwarded to transmitter ports.




Figure 5.7—Intermediate-bridge responsibilities

Ports in the CLOCK_MASTER state are responsible for revising the GrandSync-supplied clockSync PDUs to supply the appropriate media-dependent service-interface parameters and/or frames. Since the transmission times and rates may differ from those on the clock-slave port, the CLOCK_MASTER-state port is responsible for interpolating/extrapolating between previously received time samples to generate parameters corresponding to the recently observed transmit-snapshot, as described in 5.6.4

5.6.4 CLOCK_SLAVE link-delay compensation

When a clockSync frame is received at a CLOCK_SLAVE-state port, a snapshot of the local *stationTime* is saved. This snapshot time is compensated by the measured link delay. Although these time-delay adjustments are media-dependent, the concepts described within this subcaluse are applicable to most transmission media.

Processing of the *rxTime* snapshot to account for link-transmission delays involves the subtraction of a precomputed link *delay* value, as illustrated in Figure 5.8a. In parallel with this subtraction, the received *grandTime* value is extracted from the received clockSync frame. The *grandTime* and *rcTime* (the delay-compensated *txTime*) values provide an accurate time-affiliation pair that is packaged into the PDU and sent to the GrandSync entity.

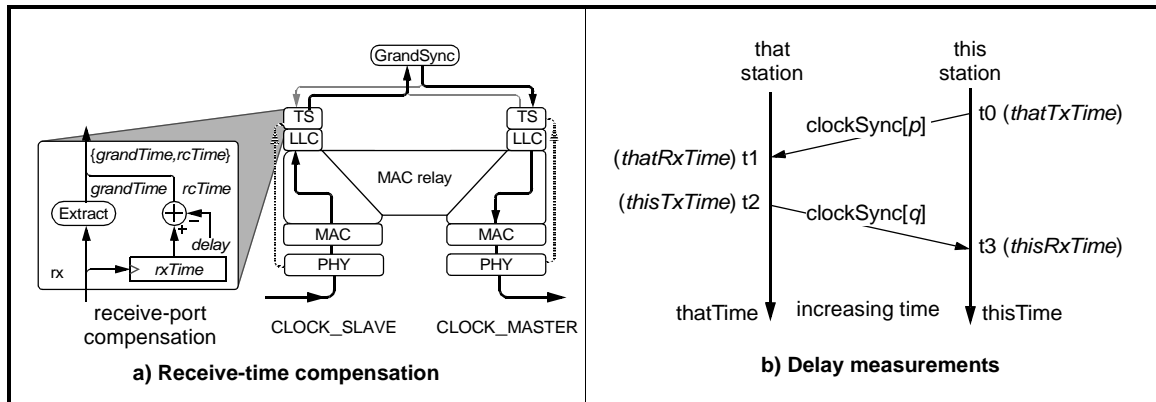


Figure 5.8—CLOCK_SLAVE link-delay compensation

Link delays can be accurately precomputed if the delays are constant, accurate time-snapshot values are present, and the clocks are stable. The computations are based on time-stamp measurements performed on distinct clockSync[*p*] and following clockSync[*q*] frame transmissions, as illustrated in Figure 5.8b.

These computations are based on the free-running station-local timers, illustrated as *thisTime* and *thatTime*. For stability and accuracy, computations use time differences measured with respect to the local timers; constant offset or frequency errors between local and neighbor timers thus have no effect on the accuracy of the computations.

To improve accuracy, a ratio of local-to-neighbor clock-intervals is precomputed over *N* (perhaps 16) frame transmissions, as specified in Equation 5.4. For hardware accuracy and simplicity, the t_2 transmission-time snapshot are held and transmitted in the following frame. Thus, the t_2 snapshot for clockSync[*n*] is transmitted within the following clockSync[*n*+1] frame.

$$\begin{aligned} \#define N 16 & \qquad \qquad \qquad // \text{ A typical value} & \qquad \qquad (5.4) \\ \text{ratio} = (t_3[n+N] - t_3[n]) / (t_2[n+N] - t_2[n]); & \qquad \qquad // \text{ Neighbor's rate ratio} \end{aligned}$$

For symmetry and accounting simplicity, the t_0 transmission-time snapshot is also held and transmitted in the following frame. Affiliated $\{t_0, t_1\}$ pairs are then returned in following clockSync[*q*] frame transmissions. The process is symmetric and thus allows either station to (otherwise independently) compute the average link transmission delay, as specified in Equation 5.5.

$$\text{delay} = ((t_3[q] - t_0[p]) - \text{rateRatio} * (t_2[q] - t_1[p])) / 2; \quad // \text{ Average link delay} \quad (5.5)$$

5.6.5 CLOCK_MASTER time-sample interpolation

A port in the CLOCK_MASTER state receives the $\{grandTime, extraTime, rcTime\}$ triads that are echoed by the GrandSync entity. Since these (oftentimes) cannot be transmitted immediately, the received time-triads are saved in a $times[]$ array, for access after the next clockSync frame transmission, as illustrated in Figure 5.9a.

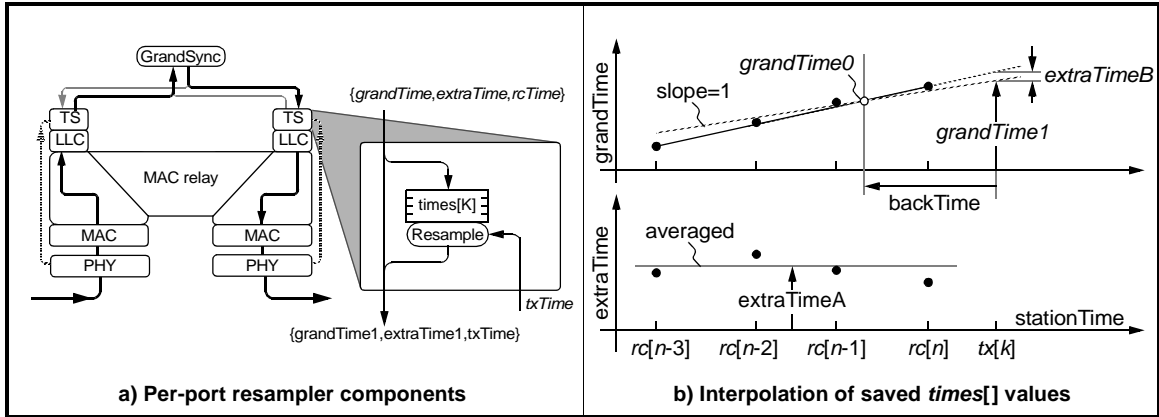


Figure 5.9—CLOCK_MASTER time-sample interpolation

NOTE—The distribution of distinct $times[]$ storage entities is an architectural model; implementations could emulate the specified behavior by providing access to shared storage and intermediate results.

During the next clockSync frame transmission, a $txTime$ snapshot is produced. The Resample logic operates on the $times[]$ array values to compute the $\{grandTime1, extraTime1, txTime\}$ time-triad that is placed within the following clockSync frame transmission, based on interpolation and averaging strategies illustrated in Figure 5.9b. Within Figure 5.9b, rc and tx are abbreviations for $rcTime$ and $txTime$ values, respectively.

To avoid traditional whiplash and gain-peaking problems associated with cascaded PLLs, an interpolation (as opposed to extrapolation) strategy yields distinct $grandTime1$ and $extraTime1$ values. These two times are forwarded independently but are recombined at the ClockSink entity, so that a single time can be passed to the higher-level application.

In concept, the computation of $grandTime1$ involves computing $tbTime = txTime - backTime$, using $tbTime$ to generate $grandTime0$ by interpolating between previously saved $times[]$ values, and then generating $grandTime1$ by extrapolating $grandTime0$ forward to time $txTime$ (assuming a constant $grandTime$ slope of one). The value of $extraTimeB$ is set to the difference between an extrapolated value (based on saved $times[]$ values) and the computed $grandTime1$ value.

Computation of the $extraTime1$ involves computing $extraTimeA$, the average of previously saved $extraTime$ values within the $times[]$ array. The sum of $extraTimeA$ (upstream station's contributions) and $extraTimeB$ (this station's contribution) yields the $extraTime1$ value. These computations are summarized in Equation 5.6.

```

48 #define N 4 // Typical value (5.6)
49 rateRatio= // Relative rates
50 (grandTime[n] - grandTime[n-N]) / (rcTime[n] - rcTime[n-N]); // GM's rate ratio
51 for (extraTimeA= i= 0; i < N; i += 1) // Use samples to
52     extraTimeA += extraTime[n-i]/N; // find an average
53 extraTimeB = backTime * (rateRatio - 1.0); // Contributed value
54 grandTime1 = rc[n] + (txTime - rc[n]) * rateRatio - extraTimeB; // Tx grandTime value
55 extraTime1 = extraTimeA/N + extraTimeB; // Tx extraTime value

```

5.7 Distinctions from IEEE Std 1588

5.7.1 Distinction summary

Advantageous properties of this protocol that distinguish it from other protocols (including portions of IEEE Std 1588) include the following:

- a) Synchronization between grandmaster and local clocks occurs at each station:
 - 1) All bridges have a lightly filtered synchronized image of the grandmaster time.
 - 2) End-point stations have a heavily filtered synchronized image of the grandmaster time.
- b) Time is uniformly represented as scaled integers: seconds and fractions-of-a-second.
- c) Locally media-dependent synchronized networks don't require extra time-snapshot hardware.
- d) Error magnitudes are sublinear with hop distances; PLL-whiplash and $O(n^2)$ errors are avoided.
- e) Multicast (one-to-many) services are not required; only adjacent-neighbor addressing is required.
- f) A relatively frequent 100 Hz (as compared to 1 Hz) update frequency is assumed:
 - 1) This rate can be readily implemented (in today's technology) for minimal cost.
 - 2) The more-frequent rate improves accuracy and reduces transient-recovery delays.
 - 3) The more-frequent rate reduces transient-recovery delays.
- g) Only one frame type simplifies the protocols and reduces transient-recovery times. Specifically:
 - 1) Cable delay computations, based on local clocks, are unaffected by *grandTime* transients.
 - 2) Rogue frames are quickly scrubbed (2.6 seconds maximum, for 256 stations).
 - 3) Drift-induced errors are greatly reduced.

5.7.2 1588 ClockSelect-message preference

Within 1588, the grandmaster preference is based on the concatenation of multiple fields, including clock-master identifiers, as illustrated in Figure 5.10. The gray-shaded values are not supported in this document.




Figure 5.10—1588 grandmaster preference

The *clockClass*, *clockAccuracy*, and *offsetScaledLogVariance* values are sparsely-encoded and application-dependent and thus not supported in this document. Similar functionality can be achieved by higher-level conventions that specify the partitioning/assignment of distinct 16-bit *priority* values.

The *txStation* and *txPort* values are unnecessary and therefore not supported in this document.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

6. GrandSync operation

6.1 Overview

6.1.1 GrandSync behavior

This clause specifies the state machines that specify GrandSync-entity processing. The operations are described in an abstract way and do not imply any particular implementations or any exposed interfaces. There is not necessarily a one-to-one correspondence between the primitives and formal procedures and the interfaces in any particular implementation.

The GrandSync entity is responsible for monitoring time-sync PDUs via the `CLOCK_SYNC.indication` service primitive, selectively echoing a subset of these PDUs via the `CLOCK_SYNC.response` service primitive, as follows:

- a) When a preferred time-sync related `CLOCK_SYNC.indication` arrives:
 - 1) The grandmaster preference and port-timeout parameters are saved.
 - 2) `CLOCK_SYNC.indication` parameters are echoed in `CLOCK_SYNC.response` parameters.
 - 3) The arrival time is recorded, for the purpose of monitoring port timeouts.
- b) Arriving non-preferred `CLOCK_SYNC.indications` are discarded.
The intent is to echo only PDUs from the currently selected grandmaster port.
- c) If the preferred-port timeout is exceeded, the preferred-port parameters are reset.
The intent is to restart grandmaster selection based on the remaining candidate ports.

6.1.2 GrandSync interface model

The time-synchronization service model assumes the presence of one or more time-synchronized AVB ports communicating with a MAC relay, as illustrated in Figure 6.1. All components are assumed to have access to a common free-running (not adjustable) *localTime* value.

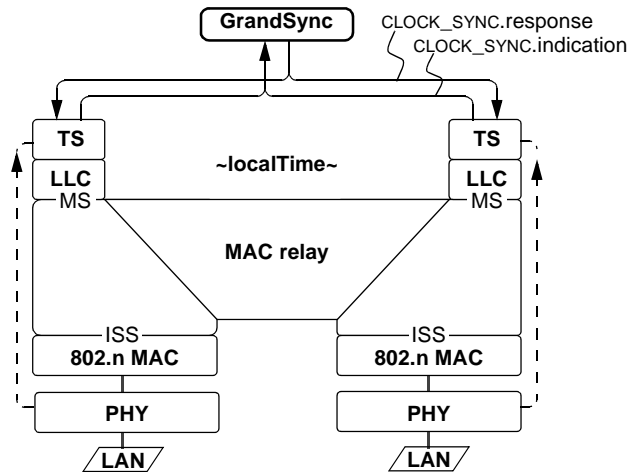


Figure 6.1—GrandSync interface model

A received MAC frame is associated with link-dependent timing information, processed within the TimeSync (TS) state machine, and passed to the GrandSync protocol entity. The GrandSync state machine (illustrated with a darker boundary) is responsible for saving time parameters from observed `CLOCK_SYNC.indication` parameters and generating `CLOCK_SYNC.response` parameters for delivery to other ports.

The clockselect and clockSync information exchanged with the GrandSync entity includes *precedence* information for grandmaster selection, a time-affiliation pair {*grandTime*,*localTime*}, and a cumulative *extraTime*, and control information, as illustrated in Figure 6.2. A clock-slave end-point can filter the sum of *grandTime* and *extraTime* values, thereby yielding its image of the globally synchronized *grandTime* value.




Figure 6.2—The clockSelect/clockSync service-interface components

6.2 Service interface primitives

6.2.1 CLOCK_SYNC.indication

6.2.1.1 Function

Provides the GrandSync protocol entity with clock-synchronization parameters derived from PDUs sent from attached media-dependent ports. The information is sufficient to identify a single clock-slave port (typically the closest-to-grandmaster port) and to disseminate grandmaster supplied clock-synchronization information to other ports.

6.2.1.2 Semantics of the service primitive

The semantics of the primitives are as follows:

```

CLOCK_SYNC.indication {
    destination_address, // Destination address
    source_address,     // Optional
    priority,           // Forwarding priority
    service_data_unit, // Delivered content
    {
        protocolType, // Distinguishes AVB frames from others
        function,     // Distinguishes between clockSync and other AVB frames
        version,      // Distinguishes between clockSync frame versions
        priority,     // Precedence for grandmaster selection
        clockID,      // Precedence for grandmaster selection
        hopCount,     // Distance from the grandmaster station
        portID,       // Identifies the source port
        interval,     // Nominal clockSync transmission interval
        flags,        // Control flags
        utcOffset,    // Difference between UTC and TAI timescales
        grandTime,    // Global-time snapshot (1-cycle delayed)
        extraTime,    // Accumulated grandTime error
        localTime     // Local-time snapshot (1-cycle delayed)
    }
}
    
```

1 NOTE—The *grandTime* field has a range of approximately 36,000 years, far exceeding expected equipment life-spans.
2 The *localTime* and *linkTime* fields have a range of 256 seconds, far exceeding the expected clockSync frame trans-
3 mission interval. These fields have a 1 pico-second resolution, more precise than the expected hardware snapshot
4 capabilities. Future time-field extensions are therefore unlikely to be necessary in the future.

5 The parameters of the CLOCK_SYNC.indication are described as follows:
6

7 **6.2.1.2.1 destination_address:** A 48-bit field that allows the frame to be conveniently stripped by its
8 downstream neighbor. The *destination_address* field contains an otherwise-reserved group 48-bit MAC
9 address (TBD).
10

11 **6.2.1.2.2 source_address:** A 48-bit field that specifies the local station sending the frame. The
12 *source_address* field contains an individual 48-bit MAC address (see 3.10), as specified in 9.2 of IEEE Std
13 802-2001.
14

15 **6.2.1.2.3 priority:** Specifies the (802.3) priority associated with content delivery.
16

17 **6.2.1.2.4 serviceDataUnit:** A multi-byte field that provides information content.
18

19 For GrandSync-entity time-sync interchanges, the *serviceDataUnit* consists of the following subfields:
20

21 **6.2.1.2.5 protocolType:** A 16-bit field contained within the payload that identifies the format and function of
22 the following fields.
23

24 **6.2.1.2.6 function:** An 8-bit field that distinguishes the clockSync frame from other AVB frame type.
25

26 **6.2.1.2.7 version:** An 8-bit field that identifies the version number associated with of the following fields.
27 TBD—A more exact definition of version is needed.
28

29 **6.2.1.2.8 priority:** A 16-bit field that can be configured by the user and overrides the remaining
30 *preference*-group field values (see 6.2.1.4).
31

32 **6.2.1.2.9 clockID:** A 64-bit globally-unique field that ensures a unique precedence value for each potential
33 grandmaster, when the *priority* field happens to have the same value (see 6.2.1.4 and 6.2.1.5).
34

35 **6.2.1.2.10 hopCount:** A 16-bit field that represents the number of hops from the grandmaster. This field
36 applies preference to the shortest path from the grand-master, if multiple paths are provided by the inter-
37 connect (see 6.2.1.4).
38

39 **6.2.1.2.11 portIDt:** A 32-bit field that represents the port number that sourced the PDU (see 6.2.1.4).
40

41 **6.2.1.2.12 interval:** An 8-bit pseudo floating-point field that specifies the nominal interval between
42 clockSync frame transmissions (see xx).
43

44 **6.2.1.2.13 flags:** An 8-bit field that warns/triggers changes to the end-station *utcOffset* value.
45

46 **6.2.1.2.14 utcOffset:** A 16-bit field that represents the current leap-seconds value.
47

48 **6.2.1.2.15 grandTime:** An 80-bit field that specifies a grandmaster synchronized time (see 6.2.1.6).
49

50 **6.2.1.2.16 localTime:** A 48-bit field that specifies the local free-running time within this station, when the
51 previous clockSync frame was received (see 6.2.1.8).
52
53
54

6.2.1.2.17 *extraTime*: A 32-bit field that specifies the cumulative grandmaster synchronized-time error. (Propagating *extraTime* and *grandTime* separately eliminates whiplash associated with cascaded PLLs.)

6.2.1.3 Version format

For compatibility with existing 1588 time-snapshot, a single bit within the version field is constrained to be zero, as illustrated in Figure 6.3. The remaining *versionHi* and *versionLo* fields shall have the values of 0 and 1 respectively.




Figure 6.3—Global-time subfield format

6.2.1.4 preference group

The precedence group includes the concatenation of multiple fields that are used to establish precedence between grandmaster candidates, as illustrated in Figure 6.4.




Figure 6.4—precedence subfields

6.2.1.5 clockID subfields

The 64-bit *clockID* field is a unique identifier. For stations that have a uniquely assigned 48-bit *macAddress*, the 64-bit *clockID* field is derived from the 48-bit MAC address, as illustrated in Figure 6.5.




Figure 6.5—clockID format

6.2.1.5.1 *oui*: A 24-bit field assigned by the IEEE/RAC (see 3.10.1).

6.2.1.5.2 *extension*: A 16-bit field assigned to encapsulated EUI-48 values.

6.2.1.5.3 *ouiDependent*: A 24-bit field assigned by the owner of the *oui* field (see 3.10.2).

6.2.1.6 grandTime format

The *grandTime* (time-of-day) field within a frame are based on seconds and fractions-of-second values, consistent with IETF specified NTP[B7] and SNTP[B8] protocols, as illustrated in Figure 6.6.

6.2.1.6.1 *seconds*: A 40-bit signed field that specifies time in seconds.




Figure 6.6—*grandTime* format

6.2.1.6.2 *fraction*: A 40-bit unsigned field that specifies a time offset within each *second*, in units of 2^{-40} second.

The concatenation of these fields specifies a 96-bit *grandTime* value, as specified by Equation 6.1.

$$grandTime = seconds + (fraction / 2^{40}) \tag{6.1}$$

6.2.1.7 *extraTime*

The error-time values within a frame are based on a selected portion of a fractions-of-second value, as illustrated in Figure 6.7. The 40-bit signed *fraction* field specifies the time offset within a *second*, in units of 2^{-40} second.




Figure 6.7—*extraTime* format

6.2.1.8 *localTime* format

The *localTime* value within a frame is based on seconds and fractions-of-second field values, as illustrated in Figure 6.8. The 48-bit *fraction* field specifies the time offset within the *second*, in units of 2^{-48} second.




Figure 6.8—*localTime* format

6.2.1.9 When generated

The time-sync related `CLOCK_SYNC.indication` service primitive is generated when new time-sync information is available. Such information could change the selection of the grandmaster or could provide a more-recent $\{grandTime, stationTime\}$ time affiliation necessary for maintaining accurate grandmaster synchronized time references.

6.2.1.10 Effect of receipt

Receipt of the service primitive by the GrandSync entity triggers an update of the grandmaster selection information. If the grandmaster selection determines the source-port to be the preferred port, its provided $\{grandTime, localTime\}$ time affiliation is also echoed to the attached entities, via invocation of the `CLOCK_SYNC.response` service primitive.

6.2.2 CLOCK_SYNC.response

6.2.2.1 Function

Communicates GrandSync protocol-entity supplied information to attached media-dependent ports. The information is sufficient for attached ports to update/propagate grandmaster clock-synchronization parameters.

6.2.2.2 Semantics of the service primitive

The semantics of the primitives are as specified for CLOCK_SYNC.indication (see 6.2.1).

6.2.2.3 When generated

Generated by the GrandSync entity upon receipt of a time-sync related CLOCK_SYNC.indication from a preferred (by grandmaster selection protocol) source port.

6.2.2.4 Effect of receipt

Receipt of the service primitive by a ClockSlave or TS entity updates entity storage. This storage update allows the destination-port to provide accurate {*grandTime*, *stationTime*} affiliations during later time-sync information transmissions.

6.3 GrandSync state machine

6.3.1 Function

The GrandSync state machine is responsible for observing CLOCK_SYNC.indication parameters, selecting PDUs with preferred time-sync content, and echoing this content in following CLOCK_SYNC.response parameters.

6.3.2 State machine definitions

AVB identifiers

Assigned constants used to specify AVB frame parameters.

AVB_FUNCTION—The function code that corresponds to a time-sync frame.
value—TBD.

AVB_MCAST—The multicast destination address corresponding to the adjacent neighbor.
value—TBD.

AVB_TYPE—The *protocolType* corresponding that uniquely identifies time-sync SDUs.
value—TBD.

AVB_VERSION—The number that uniquely identifies this version of time-sync SDUs.
value—TBD.

LAST_HOP

A constant that specifies the largest possible *hopCount* value.
value—255

NULL

A constant indicating the absence of a value that (by design) cannot be confused with a valid value.

ONES

A large constant wherein all binary bits of the numerical representation are set to one.

1 queue values

2 Enumerated values used to specify shared FIFO queue structures.

3 Q_MS_IND—Queue identifier for CLOCK_SYNC.indication transfers.

4 Q_MS_REQ—Queue identifier for CLOCK_SYNC.response transfers.

6.3.3 State machine variables

7
8 *ePtr*

9 A pointer to entity-dependent storage, where that storage comprises the following:

10 *lastTime*—Time of the last best-preference update, used for timeout purposes.

11 *rxSaved*—A copy of the best-preference GrandSync PDU parameters.

12 *new, old*

13 Local variables consisting of concatenated *preference*, *hopCount*, and *port* parameters.

14 *rsPtr*

15 A pointer to the service-data-unit portion of *rxInfo* storage.

16 *rxInfo*

17 Parameters associated with an CLOCK_SYNC.indication (see 6.2.1.2), comprising the following:

18 *destination_address*, *source_address*, *service_data_unit*

19 Where *service_data_unit* comprises:

20 *extraTime*, *function*, *grandTime*, *hopCount*, *precedence*,

21 *protocolType*, *localTime*, *syncInterval*, *version*

22 *rxPtr*

23 A pointer to the *rxInfo* storage.

24 *stationTime*

25 A shared value representing current time within each station.

26 Within the state machines of this standard, this is assumed to have two components, as follows:

27 *seconds*—An 8-bit unsigned value representing seconds.

28 *fraction*—An 40-bit unsigned value representing portions of a second, in units of 2^{-40} second.

29 *ssPtr*

30 A pointer to the service-data-unit portion of *ePtr->rxSaved* storage.

31 *sxPtr*

32 A pointer to the *ePtr->rxSaved* storage.

33 *tsPtr*

34 A pointer to the service-data-unit portion of *txInfo* storage.

35 *txInfo*

36 Parameters associated with an CLOCK_SYNC.response (see 6.2.1.2), comprising the following:

37 *destination_address*, *source_address*, *service_data_unit*

38 Where *service_data_unit* comprises:

39 *extraTime*, *function*, *grandTime*, *hopCount*, *precedence*,

40 *protocolType*, *localTime*, *syncInterval*, *version*

41 *txPtr*

42 A pointer to the *txInfo* storage.

6.3.4 State machine routines

45
46 *Dequeue(queue)*

47 Returns the next available frame from the specified queue.

48 *info*—The next available parameters.

49 NULL—No parameters available.

50 *Enqueue(queue, info)*

51 Places the *info* parameters at the tail of the specified queue on all ports.

<i>FormPreference(precedence, hops, port)</i>	1
Forms a 16-byte <i>preference</i> by concatenating the following fields:	2
<i>precedence</i> (14 bytes)	3
<i>hops</i> (1 byte)	4
<i>port</i> (1 byte)	5
<i>StationTime(ePtr)</i>	6
Returns the value of the station's shared local timer, encoded as follows:	7
<i>seconds</i> —A 16-bit unsigned value representing seconds.	8
<i>fraction</i> —A 48-bit unsigned value representing portions of a second, in units of 2^{-40} second.	9
<i>ClockSyncSdu(info)</i>	10
Checks the frame contents to identify MS_DATAUNIT.indication frames.	11
TRUE—The frame is a clockSync frame.	12
FALSE—Otherwise.	13

6.3.5 GrandSync state table

The GrandSync state machine includes a media-dependent timeout, which effectively restarts the grandmaster selection process in the absence of received clockSync frames, as specified by Table 6.1.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

Table 6.1—GrandSync state table

Current		Row	Next	
state	condition		action	state
START	(rxInfo = Dequeue(Q_MS_IND)) != NULL	1	—	TEST
	(stationTime – ePtr->timer) > 4 * ePtr->syncInterval	2	ePtr->lastTime = stationTime; ssPtr->hopCount = ssPtr->sourcePort = ssPtr->precedence = ONES;	START
	—	3	stationTime = StationTime();	
TEST	ClockSyncSdu(rsPtr) && rsPtr->hopCount != LAST_HOP	4	test = FormPreference(rsPtr->precedence, rsPtr->hopCount, rsPtr->port); best = FormPreference(ssPtr->precedence, ssPtr->hopCount, ePtr->sourcePort);	SERVE
	—	5	—	START
SERVE	test <= best	6	ePtr->lastTime = stationTime; *ssPtr = *tsPtr = *rsPtr;	HOPS
	—	7	—	START
LAST	—	8	txPtr->destination_address = AVB_MCAST; txPtr->source_address = MacAddress(ePtr); tsPtr->protocolType = AVB_TYPE; tsPtr->function = AVB_FUNCTION; tsPtr->version = AVB_VERSION; Enqueue(Q_MS_REQ, txPtr);	START

Row 6.1-1: Available indication parameters are processed.

Row 6.1-2: The absence of indications forces a timeout, after a entity-dependent delay

Row 6.1-3: Wait for changes of conditions.

Row 6.1-4: Still-active time-sync PDUs are processed further, based on grandmaster preferences. The *new* and *old* preference values consist of *precedence*, *hopCount*, and *port* components.

Row 6.1-5: Other PDUs and over-aged indications are discarded.

Row 6.1-6: Same-port indications always have preference.

Row 6.1-7: Other indications are discarded.

Row 6.1-8: Reset the timeout timer; broadcast saved parameters to all ports (including the source).

7. ClockMaster/ClockSlave state machines

7.1 Overview

7.1.1 ClockMaster/ClockSlave behaviors

This clause specifies the state machines that specify ClockMaster and ClockSlave entity processing. The operations are described in an abstract way and do not imply any particular implementations or any exposed interfaces. There is not necessarily a one-to-one correspondence between the primitives and formal procedures and the interfaces in any particular implementation.

The ClockMaster entity is responsible for forwarding the grandmaster time supplied by the ClockSource via the CLOCK_SOURCE.request service primitive, as follows:

- a) A count value (that is incremented in sequential CLOCK_SOURCE.request PDUs) is checked.
- b) The grandmaster time parameter within the CLOCK_SYNC.response[$n+1$] PDU is associated with the CLOCK_SYNC.response[n] PDU arrival time.
- c) The CLOCK_SYNC.response parameters are supplemented to form a CLOCK_SYNC.response PDU, which is then passed to the GrandSync entity.

The ClockSlave entity is responsible for extracting the grandmaster time from CLOCK_SYNC.indications and supplying the current value to the ClockTarget entity through the CLOCK_TARGET service interfaces, as follows:

- a) grandmaster time samples are extracted from GrandSync-supplied CLOCK_SYNC.response PDUs, and saved for computing grandmaster times in following CLOCK_TARGET.indication PDUs.
- b) When triggered by a CLOCK_TARGET.request indication, a CLOCK_TARGET.indication PDU is delivered to the ClockTarget state machine. That returned CLOCK_TARGET.indication PDU supplies the grandmaster time associated with the CLOCK_TARGET.request invocation time.

7.1.2 ClockMaster/ClockSlave interface model

The time-synchronization service model assumes the presence of one or more grandmaster capable entities communicating with the GrandSync state machine, as illustrated on the left side of Figure 7.1. A grandmaster capable port is also expected to provide clock-slave functionality, so that any non-selected grandmaster-capable station can synchronize to the selected grandmaster station.

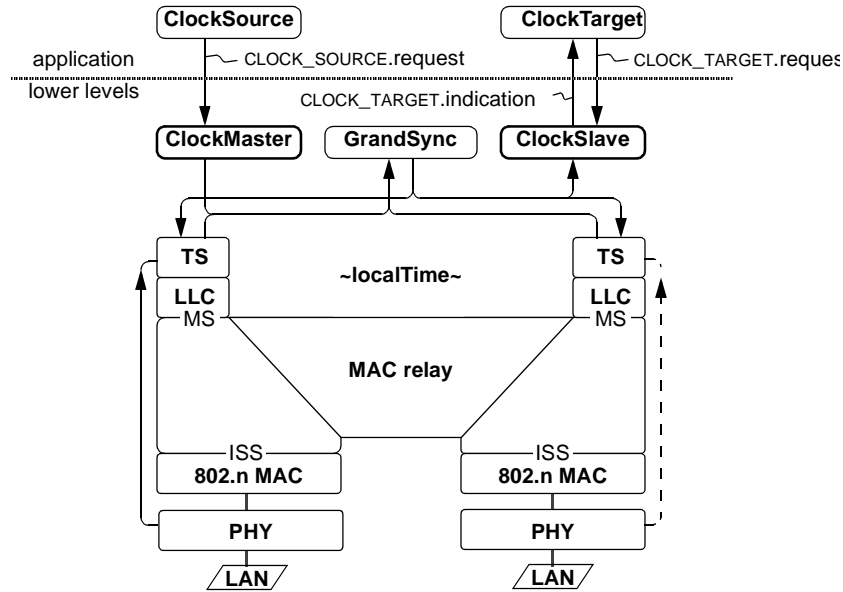


Figure 7.1—ClockMaster interface model

The clock-master ClockMaster state machine (illustrated with an italics name and darker boundary) is responsible for monitoring its port’s CLOCK_SOURCE.request PDUs and sending CLOCK_SYNC.indication PDUs. The sequencing of this state machine is specified by Table 7.1; details of the computations are specified by the C-code of Annex G.

The time-synchronization service model assumes the presence of one or more clock-slave capable time-sync entities communicating with a GrandSync protocol entity, as illustrated on the top-side of Figure 7.1. A non-talker clock-slave capable entity is not required to be grandmaster capable.

The ClockSlave state machine (illustrated with an italics name and darker boundary) is responsible for saving time parameters from echoed CLOCK_SYNC.response frames and servicing CLOCK_SOURCE.request PDUs supplied by the associated clock-slave interface. The sequencing of this state machine is specified by Table 7.2; details of the computations are specified by the C-code of Annex G.

7.2 ClockMaster service interfaces

7.2.1 Shared service interfaces

The ClockMaster entity is coupled to the bridge ports TS entities via the defined time-sync related CLOCK_SYNC.indication service interface (see 6.2.1).

7.2.2 CLOCK_SOURCE.request service interface

7.2.2.1 Function

Provides the ClockMaster entity with clock-synchronization parameters derived from the reference clock. The information is sufficient to provide the ClockMaster with accurate {*grandTime*, *localTime*} associations. The ClockSource entity supplies the reference time for service-interface invocation *n* within the parameters of the next service-interface invocation *n*+1.

7.2.2.2 Semantics of the service primitive

The semantics of the primitives are as follows:

```
CLOCK_SOURCE.request {
    frameCount,      // An integrity-check that is incremented each invocation
    grandTime,      // Global-time snapshot (1-cycle delayed)
}
```

The parameters of the CLOCK_SOURCE.request service-interface primitive are described as follows:

7.2.2.2.1 *frameCount*: An 8-bit field that is incremented on each service-interface invocation.

7.2.2.2.2 *grandTime*: An 80-bit field that specifies the grandmaster synchronized time within the source station, when the previous clockSync frame was transmitted (see 6.2.1.6).

7.2.2.3 When generated

The CLOCK_SOURCE.request service primitive is invoked by a client-resident ClockSource entity. The intent is to provide the ClockMaster with continuous/accurate updates from a ClockSource-resident clock reference.

7.2.2.4 Effect of receipt

Upon receipt by the ClockMaster entity, the encapsulated *grandTime* value is affiliated with the *stationTime* snapshot from the previous invocation; the resulting {*grandTime*, *stationTime*} affiliation is passed to the GrandSync entity for redistribution to other ClockSlave and TS entities.

7.3 ClockMaster state machine

7.3.1 State machine definitions

7.3.1.1 AVB identifiers: Assigned constants used to specify AVB frame parameters (see 6.3.2).

AVB_FUNCTION, AVB_MCAST, AVB_TYPE, AVB_VERSION

7.3.1.2 NULL: A constant value that (by design) cannot be confused with a valid value.

7.3.1.3 COUNT: A numerical constant equal to the range of the *info.frameCount* field value.

7.3.1.4 queue values: Enumerated values used to specify shared FIFO queue structures.

Q_CM_SET—The queue identifier associated with received clock-master sync frames.

Q_MS_IND—A GrandSync queue identifier (see 6.3.2).

1 **7.3.2 State machine variables**

2
3 **7.3.2.1 count:** A transient value representing the expected value of the next *rxInfo.frameCount* field value.

4
5 **7.3.2.2 ePtr:** A pointer to an entity data structure with information comprising the following:
6 *precedence*—A 14-byte field that specifies the grandmaster selection precedence.
7 *rxSaved*—Saved parameters from a received CLOCK_SOURCE.request primitive.
8 *snapShot0*—The *info.snapShot* field value from the last receive-port poke indication.
9 *snapShot1*—The value of the *ePtr->snapShot0* field saved from the last poke indication.
10 *syncInterval*—The expected rate of clockMaster service-interface invocations.

11
12
13 **7.3.2.3 rxInfo:** A contents of a higher-level supplied time-synchronization request, including the following:
14 *frameCount*—A value that increments on each CLOCK_SOURCE.request frame transmission.
15 *grandTime*—The grandmaster time, when the previous CLOCK_SOURCE.request frame was sent.

16
17 **7.3.2.4 rxPtr:** A pointer to *rxInfo* storage.

18
19 **7.3.2.5 stationTime:** See 6.3.3.

20
21 **7.3.2.6 sxPtr:** A pointer to the *ePtr->rxSaved* storage.

22
23 **7.3.2.7 txPtr:** A pointer to the service-data-unit portion of *txInfo* storage.

24
25 **7.3.2.8 txInfo:** Storage for to-be-transmitted CLOCK_SYNC.response parameters (see 6.2.2.2), comprising:
26 *destination_address, source_address, service_data_unit*
27 Where *service_data_unit* comprises:
28 *extraTime, function, grandTime, hopCount, precedence,*
29 *protocolType, localTime, sourcePort, syncInterval, version*

30
31 **7.3.2.9 txPtr:** A pointer to *txInfo* storage.

32
33 **7.3.3 State machine routines**

34
35 *Dequeue(queue)*
36 *Enqueue(queue, info)*
37 *SourcePort(entity)*
38 *StationTime(entity)*
39 *ClockSyncSdu(info)*
40 See 6.3.4.

7.3.4 ClockMaster state table

The ClockMaster state table encapsulates clock-provided sync information into a MAC-relay frame, as illustrated in Table 7.1.

Table 7.1—ClockMaster state machine table

Current		Row	Next	
state	condition		action	state
START	(rxInfo = Dequeue(Q_CM_SET)) != NULL	1	ePtr->snapShot1 = ePtr->snapShot0; ePtr->snapShot0 = stationTime; count = (sxPtr->frameCount + 1) % COUNT; grandTime = rxPtr->grandTime; *sxPtr = rxInfo;	SEND
	—	2	stationTime = StationTime(ePtr);	START
SEND	count == sxPtr->frameCount	3	txPtr->destination_address = AVB_MCAST; txPtr->source_address = MacAddress(ePtr); tsPtr->prototoType = AVB_TYPE; tsPtr->function = AVB_FUNCTION; tsPtr->version = AVB_VERSION; tsPtr->precedence = ePtr->precedence; tsPtr->hopCount = 0; tsPtr->sourcePort = SourcePort(ePtr); tsPtr->grandTime = grandTime; tsPtr->extraTime = 0; tsPtr->localTime = ePtr->snapShot1; tsPtr->syncInterval = ePtr->syncInterval; Enqueue(Q_MS_IND, txInfo);	START
	—	4	—	—

Row 7.1-1: Update snapshot values on CLOCK_SOURCE.request request arrival.

Row 7.1-2: Wait for the next change of state.

Row 7.1-3: Sequential requests are forwarded as a CLOCK_SYNC.response to the GrandSync entity.

Row 7.1-4: Nonsequential requests are discarded.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

7.4 ClockSlave service interfaces

7.4.1 Shared service interfaces

The ClockSlave entity is coupled to the GrandSync entity, via the defined CLOCK_SYNC.response service interface (see 6.2.2).

7.4.2 CLOCK_TARGET.request service interface

7.4.2.1 Function

Triggers the ClockSlave entity to provide a {*grandTime*, *localTime*} association that is synchronized with the grandmaster clock.

7.4.2.2 Semantics of the service primitive

The semantics of the primitives are as follows:

```
CLOCK_TARGET.request {  
    frameCount      // An integrity-check that is incremented each invocation  
}
```

The parameters of the CLOCK_TARGET.request service-interface primitive are described as follows:

7.4.2.2.1 *frameCount*: An 8-bit field that is incremented on each service-interface invocation.

7.4.2.3 When generated

The CLOCK_TARGET.request service primitive is invoked by a client-resident ClockTarget entity. The intent is to trigger the ClockSlave's invocation of a following CLOCK_TARGET.indication primitive, thus providing the ClockTarget entity with a recent {*grandTime*, *stationTime*} affiliation.

7.4.2.4 Effect of receipt

Upon receipt by a ClockSlave entity, a copy of the current *stationTime* value is saved and an invocation of a following CLOCK_TARGET.indication primitive is triggered.

7.4.3 CLOCK_TARGET.indication service interface

7.4.3.1 Function

Provides the ClockSync entity with clock-synchronization parameters derived from the reference clock. The information comprises {*frameCount*, *grandTime*} associations: *frameCount* is supplied by the previous CLOCK_TARGET.request invocation; *grandTime* represents the invocation time of that preceding CLOCK_TARGET.request service primitive.

7.4.3.2 Semantics of the service primitive

The semantics of the primitives are as follows:

```
CLOCK_TARGET.indication {  
    frameCount,      // Identifies the previous CLOCK_TARGET.request invocation  
    grandTime,       // grandmaster synchronized snapshot.  
}
```

The parameters of the CLOCK_TARGET.indication service-interface primitive are described as follows:

7.4.3.2.1 *frameCount*: An 8-bit field that copied from the like-named field of the previous CLOCK_TARGET.request service-interface invocation.

7.4.3.2.2 *grandTime*: An 80-bit field that specifies the grandmaster synchronized time within the ClockSlave entity, when the previous CLOCK_TARGET.request service-interface was invoked.

7.4.3.3 When generated

The invocation of the CLOCK_TARGET.indication service primitive is invoked by the receipt of a ClockTarget supplied CLOCK_TARGET.request PDU. The intent is to provide the ClockTarget entity with a recent {*grandTime,stationTime*} affiliation.

7.4.3.4 Effect of receipt

Upon receipt by a ClockTarget entity, the {*grandTime,stationTime*} affiliation is expected to be saved and (along with previously saved copies) used to adjust the rate of the grandmaster synchronized ClockTarget-resident clock.

7.5 ClockSlave state machine

7.5.1 Function

7.5.2 State machine definitions

7.5.2.1 NULL: A constant that (by design) cannot be confused with a valid value.

7.5.2.2 queue values: Enumerated values used to specify shared FIFO queue structures.

Q_MS_REQ—A GrandSync queue identifier (see 6.3.2).

Q_CS_REQ—The queue identifier associated with CLOCK_TARGET.request requests.

Q_CS_IND—The queue identifier associated with CLOCK_TARGET.indication indications.

7.5.3 State machine variables

7.5.3.1 *ePtr*: A pointer to entity-dependent information, including the following:

rxSaved—A copy of the GrandSync supplied MA_DATAUNIT.request value.

syncInterval—The expected service rate of CLOCK_TARGET.request services.

baseTimer—Recently saved time events, each consisting of the following:

index—Index into the *timed*[] array, where last times were stored.

range—Number of entries within the *timed*[] array

timed[range]—Recently saved time events, each consisting of the following:

grandTime—A previously sampled grandmaster synchronized time.

extraTime—The residual error associated with the sampled *grandTime* value.

stationTime—The station-local time affiliated with the sampled *grandTime* value.

7.5.3.2 *cxInfo*: A contents of a higher-level supplied time-synchronization request, including the following:

frameCount—A value that increments on each CLOCK_SOURCE.request PDU transfer.

7.5.3.3 *nextTime*: Storage representing *grandTime* and *extraTime* values returned from call to *NextTimed*().

7.5.3.4 *rsPtr*: A pointer to the service-data-unit portion of *rxInfo*.

1 **7.5.3.5 rxInfo:** A contents of a GrandSync supplied CLOCK_SYNC.response (see 6.2.2), including:
2 *destination_address, source_address, service_data_unit*

3 Where *service_data_unit* comprises:

4 *extraTime, function, grandTime, protocolType, localTime, version*
5

6 **7.5.3.6 rxPtr:** A pointer to *rxInfo*.
7

8 **7.5.3.7 rxSyncInterval:** The synchronization interval of this station's GrandSync-selected clock-slave port.
9

10 **7.5.3.8 stationTime:** See 6.3.3.
11

12 **7.5.3.9 ssPtr:** A pointer to the service-data-unit portion of the *ePtr->rxSaved* storage
13

14 **7.5.3.10 sxPtr:** A pointer to the *ePtr->rxSaved* storage
15

16 **7.5.3.11 timePtr:** A pointer to the *ePtr->timed[]* array storage
17

18 **7.5.3.12 txInfo:** A contents of a ClockSlave supplied CLOCK_TARGET.indication (see 6.2.2), comprising:
19 *frameCount*—The saved value of the like named field from the previous CLOCK_TARGET.request PDU.
20 *grandTime*—The grandmaster synchronized time sampled during the CLOCK_TARGET.request transfer.
21

22 **7.5.3.13 txPtr:** A pointer to *txInfo* storage.
23

24 **7.5.3.14 txSyncInterval:** The synchronization interval of this ClockSlave entity.
25

26 **7.5.4 State machine routines** 27

28 **7.5.4.1 Dequeue(queue):** See 6.3.4.
29

30 **7.5.4.2 Enqueue(queue, info):** See 6.3.4.
31

32 **7.5.4.3 NextSaved(btPtr, rateInterval, grandTime, extaTime, thisTime):**

33 Saves *grandTime, extraTime* values associated with a snapshot taken at *thisTime*, with the saved values
34 spanning a *rateInterval* specified interval.
35

36 **7.5.4.4 NextTimed(btPtr, stationTime, backInterval):**

37 Returns *grandTime* and *extraTime* values associated with a snapshot taken at *stationTime*, back-interpolated
38 by a *backInterval* time, based on previous received-time information saved in the *btPtr* referenced data
39 structure.
40

41 **7.5.4.5 StationTime(entity):** See 6.3.4.
42

43 **7.5.4.6 ClockSyncSdu(info):** See 7.3.3.
44
45
46
47
48
49
50
51
52
53
54

7.5.5 ClockSlave state table

The ClockSlave state machine includes a media-dependent timeout, which effectively disconnects a clock-slave port in the absence of received clockSync frames, as illustrated in Table 7.2.

Table 7.2—ClockSlave state table

Current		Row	Next	
state	condition		action	state
START	(rxInfo = Dequeue(Q_MS_REQ)) != NULL	1	—	TEST
	((cxInfo = Dequeue(Q_CS_REQ)) != NULL	2	rxSyncInterval = ssPtr->syncInterval; txSyncInterval = ePtr->syncInterval; backInterval = (3 * rxSyncInterval + txSyncInterval) / 2; nextTimes = NextTimed(btPtr, stationTime, backInterval); txPtr->count = cxInfo.count; txPtr->grandTime = nextTimes.grandTime + nextTimes.extraTime; Enqueue(Q_CS_IND, txInfo);	START
	—	3	stationTime = StationTime(ePtr);	
TEST	ClockSyncSdu(rsPtr)	4	*sxPtr = *rxPtr; NextSaved(btPtr, rateInterval, rsPtr->grandTime; rsPtr->extraTime, rsPtr->localTime);	START
	—	5	—	

Row 7.2-1: The received CLOCK_SYNC.response parameters are dequeued for checking.

Row 7.2-2: A clock-slave request generates an affiliated information-providing indication.

The affiliated indication has the sequence-count information provided by the request.

The delivered end-point *grandTime* value is the sum of delivered *grandTime* and *extraTime* values.

The requested content is queued for delivery to the higher-level client.

Row 7.2-3: Wait for the next change-of-conditions.

Row 7.2-4: Validated GrandSync entity requests are accepted; its time parameters are saved.

The back-interpolation time is estimated from the *syncInterval* times of the source and clock slave.

(This back-interpolation time is used by *NextTimed()*, which provides transmission-time estimates.)

Row 7.2-5: Wait for the next change-of-conditions.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

8. Ethernet full duplex (EFDX) state machines

8.1 Overview

This clause specifies the state machines that support 802.3 Ethernet full duplex (EFDX) bridges. The operations are described in an abstract way and do not imply any particular implementations or any exposed interfaces. There is not necessarily a one-to-one correspondence between the formal specification and the interfaces in any particular implementation.

8.1.1 EFDX link indications

The duplex-link TimeSyncRxEfdx state machines are provided with snapshots of clockSync-frame reception and transmission times, as illustrated by the ports within Figure 8.1. These link-dependent indications can be different for bridge ports attached to alternative media.

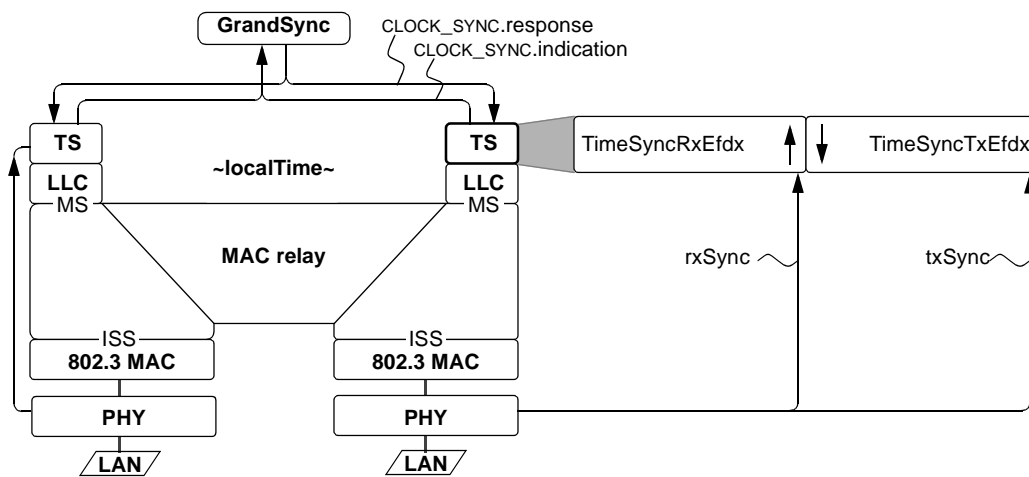


Figure 8.1—EFDX-link interface model

The rxSync and txSync indications provide a tag (to reliably associate them with MAC-supplied clockSync frames) and a *localTime* stamp indicating when the associated clockSync frame was received, as illustrated within Figure 8.2.




Figure 8.2—Contents of rxSync/txSync indications

8.1.2 Link-delay compensation

Synchronization accuracies are affected by the transmission delays associated with transmissions over links between bridges. To compensate for these transmission delays, the receive port is responsible for compensating {*grandTime*, *stationTime*} affiliations by the (assumed to be constant) frame-transmission delay.

The clock-slave entity uses the computed cable-delay measurement and is therefore (in concept) responsible for initiating such measurements. Cable-delay measurements begin with the transmission of frame F1 between the clock-slave and clock-master stations and conclude with the a clock-master response, a transmission of frame F2 between the clock-master to clock-slave stations, as illustrated in Figure 8.3.

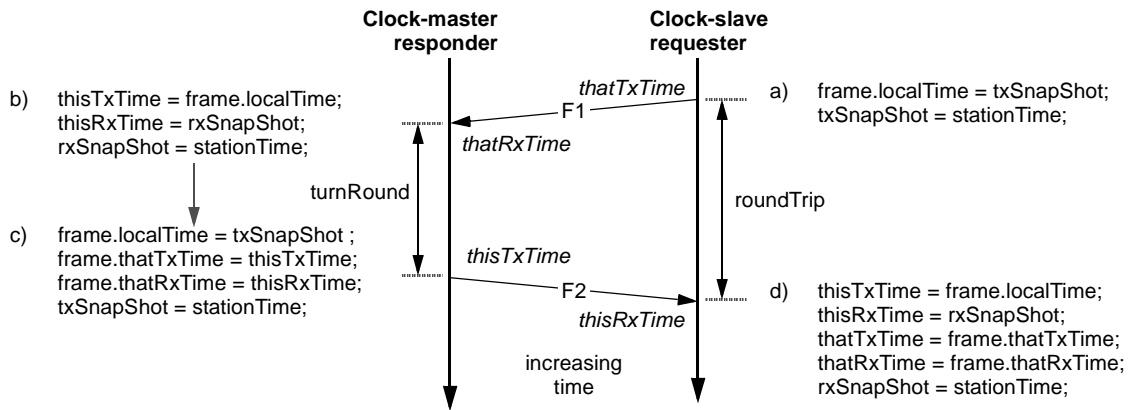


Figure 8.3—Link-delay compensation

The cable-delay computations are performed in multiple steps, as follows:

- a) The F1-frame transmission involves multiple steps:
 - 1) The *txSnapShot* value (time of the last F1 transmission) is copied to *frame.localTime* storage.
 - 2) Remaining fields are copied into *frame* storage; the *frame*-storage content is transmitted.
 - 3) The *txSnapShot* value is set to the frame-F1 transmission time, for next step (a) usage.
- b) The F1-frame reception involves multiple steps:
 - 1) The *frame.localTime* value is copied to a port-local *thisTxTime* field, for next step (c) usage.
 - 2) The *rxSnapShot* value (time of the last F1 reception) is copied to a port-local *thisRxTime* field.
 - 3) The *rxSnapShot* value is set to the F1-frame reception time, for next step (b) usage.
- c) The F2-frame transmission involves multiple steps:
 - 1) The *txSnapShot* value (time of the last F1 transmission) is copied to *frame.localTime* storage.
 - 2) The receive-port *thisTxTime* value is copied to *frame.thatTxTime* storage.
 - 3) The receive-port *thisRxTime* value is copied to *frame.thatRxTime* storage.
 - 4) Remaining fields are copied into *frame* storage; the *frame*-storage content is transmitted.
 - 5) The *txSnapShot* value is set to the frame-F2 transmission time, for next step (c) usage.
- d) The F2-frame reception involves multiple steps:
 - 1) The *frame.localTime* value is copied to a port-local *thisTxTime* field.
 - 2) The *rxSnapShot* value (time of the last F2 reception) is copied to a port-local *thisRxTime* field.
 - 3) The *frame.thatTxTime* value is copied to a port-local *thatTxTime* field.
 - 4) The *frame.thatRxTime* value is copied to a port-local *thatRxTime* field.
 - 5) The *rxSnapShot* value is set to the F2-frame reception time, for next step (d) usage.

1 At the conclusion of these steps, the values returned to the clock-slave requester include the values below.
2 (Within Figure 8.3, these values are also illustrated in the center, at their source, using a distinct italic font.)

- 3 — *thatTxTime*. The clock-slave transmit time.
- 4 — *thatRxTime*. The clock-master receipt time.
- 5 — *thisTxTime*. The clock-master transmit time.
- 6 — *thisRxTime*. The clock-slave receipt time.

7
8 Based on the preceding listed values, Equation 8.1 defines the computations for computing *linkDelay*.
9 Although not explicitly stated, the best accuracy can be achieved by performing these computation every
10 cycle.

$$\begin{aligned} \text{linkDelay} &= (\text{roundTrip} - \text{turnRound}) / 2; & (8.1) \\ \text{roundTrip} &= \text{thisRxTime} - \text{thatTxTime}; \\ \text{turnRound} &= (\text{thisTxTime} - \text{thatRxTime}) * \text{ratesRatio}; \end{aligned}$$

11
12
13
14
15 Where:

$$\text{ratesRatio} \cong (\text{deltaRxTime} / \text{deltaTxTime});$$

16
17 The value of *ratesRatio* is necessary to maintain tight accuracies in the presence of significant (± 200 PPM)
18 differences in clock-master/clock-slave timing references and significant (multiple milliseconds) *turnRound*
19 delays. This value is also readily computed from the preceding listed values, as specified by Equation 8.2.
20

$$\text{ratesRatio}[n] = (\text{thisRxTime}[n] - \text{thisRxTime}[n-N]) / (\text{thisTxTime}[n] - \text{thisTxTime}[n-N]); \quad (8.2)$$

21
22
23 NOTE—For 802.3 and other inexpensive interconnects, the processing of slow-rate PDUs is oftentimes performed by
24 firmware and (due to interrupt and processing delays) the turn-around delays can be much larger than the packet-trans-
25 mission times.

26
27 The cable-delay computations assume the transmission delays associated with frame F1 and frame F2 are
28 equal and constant. If the duplex links within a span have different propagation delays, these *linkDelay*
29 calculations do not correspond to the different propagation delays, but represent the average of the two link
30 delays. Implementers have the option of manually specifying the link-delay differences via MIB-accessible
31 parameters, within tightly-synchronized systems where this inaccuracy might be undesirable.

32
33 This cable-delay calculation does not rely on the particular timings of F1 and F2 frame transmissions. These
34 transmissions can be triggered independently (as opposed to one triggered by the other) and could occur at
35 different rates (although the accuracies are limited by the slower rate). As a direct benefit of these
36 independence properties, distinct interlocks or timeouts for expected-but-corrupted-and-not-delivered trans-
37 missions are unnecessary.

38
39 Furthermore, there is no need to transport F1 and F2 content in distinct frames. The contents of clock-slave
40 affiliated F1 and clock-master affiliated F2 frames can be merged and transported within the same frame.
41 Thus, distinct frame types and/or transmission timings are unnecessary; the link-delay calibration protocols
42 do nothing to prevent the same frame from communicating master-to-slave and slave-to-master link delays,
43 in addition to the baseline grandmaster timing and selection parameters.
44
45
46
47
48
49
50
51
52
53
54

8.2 efdxClockSync frame format

8.2.1 efdxClockSync fields

EFDX clock-synchronization (efdxClockSync) frames facilitate the synchronization of neighboring clock-master and clock-slave stations. The frame, which is normally sent at 10ms intervals, includes time-snapshot information and the identity of the network’s clock master, as illustrated in Figure 8.4. The gray boxes represent physical layer encapsulation fields that are common across Ethernet frames.

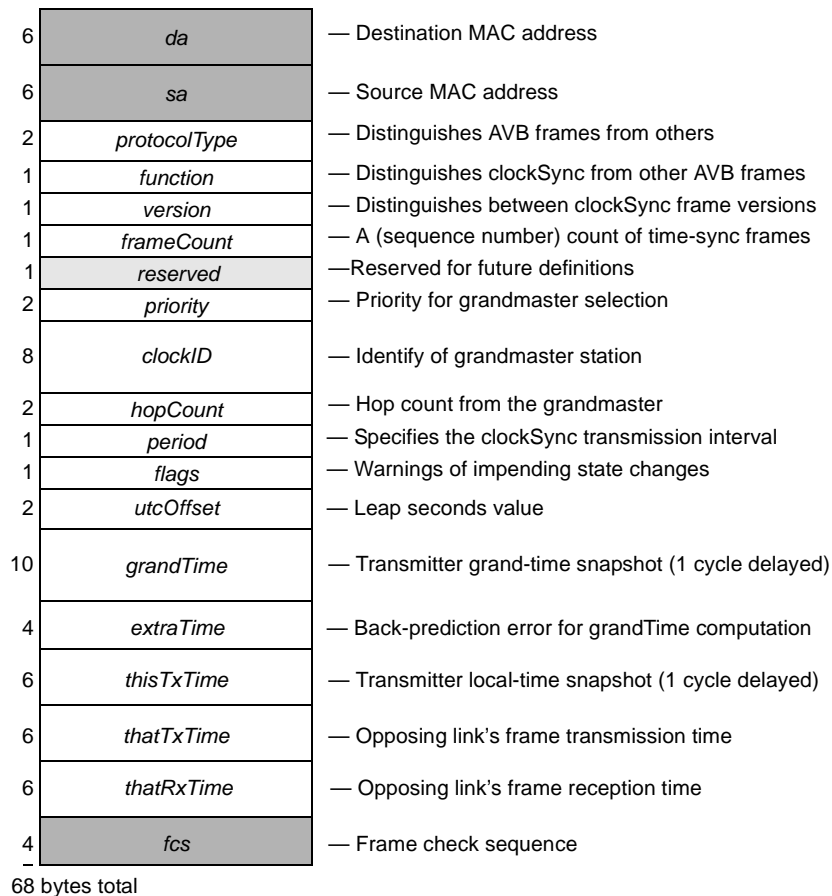


Figure 8.4—efdxClockSync frame format

NOTE— Existing 1588 time-snapshot hardware captures the values between byte-offset 34 and 45 (inclusive). The location of the *frameCount* field (byte-offset 44) has been adjusted to ensure this field can be similarly captured for the purpose of unambiguously associating clockSync-packet snapshots (that bypass the MAC) and clockSync-packet contents (that pass through the MAC).

The 48-bit *da* (destination address), 48-bit *sa* (source address) field, 16-bit *protocolType*, 8-bit *function*, 8-bit *version*, 2-byte *priority*, 8-byte *clockID*, 2-byte *hopCount*, 1-byte *priority*, 1-byte *flags*, 2-byte *utcOffset*, 80-bit *grandTime*, and 32-bit *extraTime* field are specified in 6.2.1.2.

8.2.1.1 *frameCount*: An 8-bit field that is incremented by one between successive clockSync frame transmission.

8.2.1.2 *thatTxTime*: A 48-bit field that specifies the local free-running time within the source station, when the previous clockSync frame was transmitted on the opposing link (see 6.2.1.8).

1 **8.2.1.3 *thatRxTime*:** A 48-bit field that specifies the local free-running time within the target station, when
2 the previous clockSync frame was received on the opposing link (see 6.2.1.8).

3
4 **8.2.1.4 *thisTxTime*:** A 48-bit field that specifies the local free-running time within the neighbor station,
5 when the previous clockSync frame was transmitted on the incoming link (see 6.2.1.8).

6
7 **8.2.1.5 *fccs*:** A 32-bit (frame check sequence) field that is a cyclic redundancy check (CRC) of the frame.

8
9 **8.2.2 Clock-synchronization intervals**

10
11 Clock synchronization involves synchronizing the clock-slave clocks to the reference provided by the grand
12 clock master. Tight accuracy is possible with matched-length duplex links, since bidirectional frame
13 transmissions can cancel the cable-delay effects.

14
15 Clock synchronization involves the processing of periodic events. Multiple time periods are involved, as
16 listed in Table 8.1. The clock-period events trigger the update of free-running timer values; the period affects
17 the timer-synchronization accuracy and is therefore constrained to be small.

18
19
20 **Table 8.1—Clock-synchronization intervals**

21
22

Name	Time	Description
clock-period	< 50 ns	Resolution of timer-register value updates
send-period	10 ms	Time between sending of periodic clockSync frames between adjacent stations
slow-period	100 ms	Time between computation of clock-master/clock-slave rate differences

23
24
25
26
27
28

29
30 The send-period events trigger the interchange of clockSync frames between adjacent stations. While a
31 smaller period (1 ms or 100 μ s) could improve accuracies, the larger value is intended to reduce costs by
32 allowing computations to be executed by inexpensive (but possibly slow) bridge-resident firmware.

33
34 The slow-period events trigger the computation of timer-rate differences. The timer-rate differences are
35 computed over two slow-period intervals, but recomputed every slow-period interval. The larger 100 ms (as
36 opposed to 10 ms) computation interval is intended to reduce errors associated with sampling of
37 clock-period-quantized slow-period-sized time intervals.

8.3 TimeSyncRxEfdx state machine

8.3.1 Function

The TimeSyncRxEfdx state machine is responsible for monitoring its port's rxSync indications, receiving MAC-supplied frames, and sending MAC-relay frames. The sequencing of this state machine is specified by Table 8.2; details of the computations are specified by the C-code of Annex G.

8.3.2 State machine definitions

8.3.2.1 LAST_HOP: A constant representing the largest-possible frame.hopCount value.
value—255.

8.3.2.2 NULL: A constant that (by design) cannot be confused with a valid value.

8.3.2.3 queue values: Enumerated values used to specify shared FIFO queue structures.
 Q_MS_IND—The queue identifier associated with MAC frames sent into GrandSync.
 Q_ES_IND—The queue identifier associated with the received MAC frames.
 Q_RX_SYNC—The queue identifier associated with rxSync, sent from the lower levels.

8.3.3 State machine variables

8.3.3.1 cableDelay: Values (possibly scaled integers) representing cable-delay times.

8.3.3.2 count: A transient value representing the expected value of the next *rxInfo.frameCount* value.

8.3.3.3 cxInfo: A contents of a lower-level supplied time-synchronization poke indication, including:
frameCount—The value of the like-named field within the last clockSync packet arrival.
localTime—The value of *stationTime* associated with the last clockSync packet arrival.

8.3.3.4 cxPtr: A pointer to *cxInfo* storage.

8.3.3.5 ePtr: A pointer to a data structure that contains port-specific information comprising the following:
frameCount—The value of *frameCount* within the last received frame.
rated—The ratio of the local-station and remote-station local-timer rates.
snapCount—The value of *frameCount* saved from the last snapshot indication.
snapShot0—The *info.snapShot* field value from the last receive-port snapshot indication.
snapShot1—The value of the *ePtr->snapShot0* field at the snapshot indication.
times[N]—An array of time groups, where each array elements consists of:
thisTime—The local receive time associated with received time-sync frames.
thatTime—The remote transmit time associated with received time-sync frames.

8.3.3.6 ratesRatio: A variable representing the ratio of this station's timer to this port's neighbor timer.

8.3.3.7 roundTrip: The time between transmit-to-neighbor and receive-from-neighbor events.

8.3.3.8 rsPtr: A pointer to the service-data-unit portion of *rxInfo* storage.

8.3.3.9 rxInfo: Storage for received time-sync PDUs, comprising:
destination_address, *source_address*, *service_data_unit*

Where *service_data_unit* comprises:

extraTime, *frameCount*, *function*, *grandTime*, *hopCount*, *localTime*,
protocolType, *precedence*, *thatTxTime*, *thatRxTime*, *version*

8.3.5 TimeSyncRxEfdx state machine table

The TimeSyncRxEfdx state machine associates PHY-provided sync information with arriving clockSync frames and forwards adjusted frames to the MAC-relay function, as illustrated in Table 8.2.

Table 8.2—TimeSyncRxEfdx state machine table

Current		Row	Next	
state	condition		action	state
START	(cxInfo = Dequeue(Q_RX_SYNC)) != NULL	1	ePtr->snapShot1 = ePtr->snapShot0; ePtr->snapShot0 = cxPtr->localTime; ePtr->snapCount = cxPtr->frameCount; match= (rxPtr->frameCount==ePtr->snapCount);	START
	(rxInfo=Dequeue(Q_ES_IND)) != NULL	2	count = (ePtr->rxFrameCount + 1) % COUNT; ePtr->rxFrameCount = rxPtr->frameCount;	TEST
	match	3	—	PAIR
	—	4	stationTime = StationTime(ePtr);	START
TEST	!ClockSyncSdu(rsPtr)	5	Enqueue(Q_CS_IND, rxPtr);	START
	rxPtr->hopCount == LAST_HOP	6	—	
	count != rxPtr->frameCount	7	—	
	—	8	match= (rxPtr->frameCount==ePtr->snapCount);	
PAIR	—	9	ePtr->times[0].thisTime = ePtr->snapShot1; ePtr->times[1].thatTime = rsPtr->localTime; ratesRatio = RemoteRate(ePtr->times); roundTrip = localTime - ePtr->thatTxTime; turnRound = rsPtr->localTime - rsPtr->thatRxTime; cableDelay = Min(0, roundTrip - (turnRound * ratesRatio)); txPtr->destination_address = rxPtr->destination_address; txPtr->source_address = rxPtr->source_address; tsPtr->protocolType = rsPtr->protocolType; tsPtr->function = rsPtr->function; tsPtr->version = rsPtr->version; tsPtr->grandTime = rsPtr->grandTime; tsPtr->extraTime = rsPtr->extraTime; tsPtr->localTime = ePtr->snapShot1 - cableDelay; tsPtr->sourcePort = SourcePort(ePtr); tsPtr->hopCount = rsPtr->hopCount; tsPtr->syncInterval = ePtr->syncInterval; Enqueue(Q_MR_HOP, relayFrame);	START

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

- 1 **Row 8.2-1:** Update snapshot values on clockSync frame arrival.
2 **Row 8.2-2:** Initiate inspection of frames received from the lower-level MAC.
3 **Row 8.2-3:** Generate a GrandSync PDUs using matching snapshot and frame information.
4 **Row 8.2-4:** Wait for the next change-of-state.
5
6 **Row 8.2-5:** The non-clockSync frames are passed through.
7 **Row 8.2-6:** Over-aged clockSync frames are discarded.
8 **Row 8.2-7:** Non-sequential clockSync frames are ignored.
9 **Row 8.2-8:** Associated snapshot and frame information trigger a GrandSync indication generation.
10
11 **Row 8.2-9:** Generate a time-sync GrandSync indication from saved snapshot and frame information.
12

13 **8.4 TimeSyncTxEfdx state machine**

14 **8.4.1 Function**

15
16
17 The TimeSyncTxEfdx state machine is responsible for saving time parameters from relayed timedSync
18 frames and forming clockSync frames for transmission over the attached link.
19

20 **8.4.2 State machine definitions**

21
22 **8.4.2.1 NULL:** A constant that (by design) cannot be confused with a valid value.
23

24 **8.4.2.2 queue values:** Enumerated values used to specify shared FIFO queue structures.

25 Q_MR_HOP—The queue identifier associated with frames sent from the relay.

26 Q_ES_REQ—The queue identifier associated with frames sent to the MAC.

27 Q_TX_SYNC—The queue identifier associated with txSync, sent from the lower levels.
28

29 **8.4.2.3 T10ms:** A constant that represents a 10 ms value.
30

31 **8.4.3 State machine variables**

32
33 **8.4.3.1 backInterval:** A variable that represents the back-interpolation interval for transmit-time affiliations.
34

35 **8.4.3.2 cxInfo:** A contents of a lower-level supplied time-synchronization poke indication, including:

36 *snapCount*—The value of the like-named field within the last clockSync packet arrival.

37 *localTime*—The value of *stationTime* associated with the last clockSync packet arrival.
38

39 **8.4.3.3 dPtr:** A pointer this port's associated TimeSyncRxEfdx-entity storage.
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

8.4.3.4 ePtr: A pointer to a data structure that contains port-specific information comprising the following:	1
<i>baseTimer</i> —Recently saved time events, each consisting of the following:	2
<i>index</i> —Index into the <i>timed[]</i> array, where last times were stored.	3
<i>range</i> —Number of entries within the <i>timed[]</i> array	4
<i>timed[range]</i> —Recently saved time events, each consisting of the following:	5
<i>grandTime</i> —A previously sampled grandmaster synchronized time.	6
<i>extraTime</i> —The residual error associated with the sampled <i>grandTime</i> value.	7
<i>stationTime</i> —The station-local time affiliated with the sampled <i>grandTime</i> value.	8
<i>frameCount</i> —A consistency-check identifier that is incremented on each transmission.	9
<i>lastTime</i> —The last transmit time, saved for timeout purposes.	10
<i>rxSaved</i> —A copy of the last received GrandSync parameters.	11
<i>syncInterval</i> —The expected interval between successive time-sync transmissions.	12
<i>txSnapCount</i> —The <i>frameCount</i> value associated with the last transmission.	13
<i>txlocalTime</i> —The <i>stationTime</i> value associated with the last transmission.	14
	15
8.4.3.5 rsPtr: A pointer to service-data-unit portion of <i>rxInfo</i> storage.	16
	17
8.4.3.6 rxInfo: Storage for received time-sync PDUs from the GrandSync entity, comprising:	18
<i>destination_address, source_address, service_data_unit</i>	19
Where <i>service_data_unit</i> comprises:	20
<i>extraTime, function, grandTime, hopCount,</i>	21
<i>precedence, protocolType, localTime, syncInterval, version</i>	22
	23
8.4.3.7 rxPtr: A pointer to <i>rxInfo</i> storage.	24
	25
8.4.3.8 rxSyncInterval: Represents the sync-interval associated with this station’s clock-slave port.	26
	27
8.4.3.9 stationTime: See 6.3.3.	28
	29
8.4.3.10 ssPtr: A pointer to the service-data-unit portion of the <i>ePtr->rxSaved</i> storage	30
	31
8.4.3.11 sxPtr: A pointer to the <i>ePtr->rxSaved</i> storage.	32
	33
8.4.3.12 tsPtrI A pointer to service-data-unit portion of <i>txInfo</i> storage.	34
	35
8.4.3.13 txInfo: Storage for to-be-transmitted time-sync PDUs, comprising:	36
<i>destination_address, source_address, service_data_unit</i>	37
Where <i>service_data_unit</i> comprises:	38
<i>extraTime, function, frameCount, grandTime, hopCount, localTime,</i>	39
<i>precedence, protocolType, thatRxTime, thatTxTime, version</i>	40
	41
8.4.3.14 txPtr: A pointer to <i>txInfo</i> storage.	42
	43
8.4.3.15 txSyncInterval: A variable that represents the sync-interval associated with this clock-master port.	44
	45
	46
	47
	48
	49
	50
	51
	52
	53
	54

Table 8.3—TimeSyncTxEfdx state machine table

Current		Row	Next	
state	condition		action	state
START	(rxInfo = Dequeue(Q_MS_REQ)) != NULL	1	—	TEST
	(stationTime – ePtr->lastTime) > T10ms	2	ePtr->lastTime = stationTime;	SEND
	(cxInfo = Dequeue(Q_TX_SYNC)) != NULL	3	ePtr->txlocalTime = cxPtr->localTime; ePtr->txSnapCount = cxPtr->frameCount;	START
	—	4	stationTime = StationTime(ePtr); rxSyncInterval = ssPtr->syncInterval; txSyncInterval = ePtr->syncInterval; backInterval = (3 * rxSyncInterval + txSyncInterval) / 2; rateInterval = backInterval + (3 * txSyncInterval) / 2;	
TEST	ClockSyncSdu(rsPtr)	5	ePtr->rxSaved = rxInfo; NextSaved(btPtr, rateInterval,rsPtr->grandTime, rsPtr->extraTime, rsPtr->localTime);	START
	—	6	Enqueue(Q_ES_REQ, rxPtr);	
SEND	—	7	dPtr = PortPair(ePtr); nextTimes = NextTimed(btPtr, stationTime, backInterval); ePtr->syncInterval, ePtr->timed); ePtr->txFrameCount = (ePtr->txSnapCount + 1) % COUNT; txPtr->destination_address = sxPtr->destination_address; txPtr->source_address = sxPtr->source_address; tsPtr->protocollID = ssPtr->protocollID; tsPtr->function = ssPtr->function; tsPtr->version = ssPtr->version; tsPtr->hopCount = ssPtr->hopCount; tsPtr->frameCount = ssPtr->frameCount; tsPtr->grandTime = nextTimes.grandTime; tsPtr->extraTime = nextTimes.extraTime; tsPtr->localTime = ePtr->txlocalTime; tsPtr->thatTxTime = dPtr->thisTxTime; tsPtr->thatRxTime = dPtr->thisRxTime; Enqueue(Q_ES_REQ, txPtr);	START

Row 8.3-1: Relayed frames are further checked before being processed.

Row 8.3-2: Transmit periodic clockSync frames.

Row 8.3-3: Update snapshot values on clockSync frame departure.

Row 8.3-4: Wait for the next change-of-state.

Row 8.3-5: The clockSync PDUs are checked further.

Row 8.3-6: The non-clockSync PDUs are passed through.

Row 8.3-7: Active clockSync frames are cable-delay compensated and passed through.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

9. Wireless state machines

EDITOR DVJ NOTE—This clause is based on indirect knowledge of the 802.11v specifications, as interpreted by the author, and have not been reviewed by the 802.1 or 802.11v WGs. The intent was to provide a forum for evaluation of the media-independent MAC-relay interface, while also triggering discussion of 802.11v design details. As such, this clause is highly preliminary and subject to change.

Specifically, we have not resolved the grouping of information that is transferred through the service interfaces (currently written as all) and the information that would be transferred through standard MAC frames (currently written as none).

9.1 Overview

This clause specifies the state machines that support wireless 802.11v-based bridges. The operations are described in an abstract way and do not imply any particular implementations or any exposed interfaces. There is not necessarily a one-to-one correspondence between the formal specification and the interfaces in any particular implementation.

9.1.1 Link-dependent indications

The wireless 802.11v TimeSyncR11v state machines are provided with MAC service-interface parameters, as illustrated within Figure 9.1. These link-dependent indications can be different for bridge ports attached to alternative media.

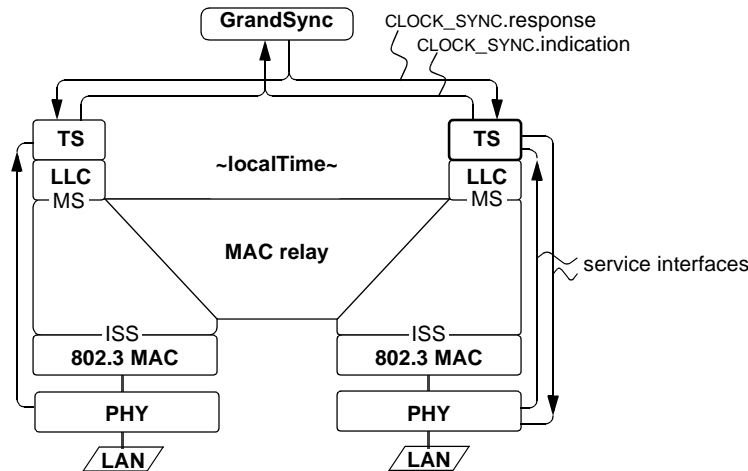


Figure 9.1—R11v interface model

The rxSync and txSync indications are localized communications between the MAC-and-PHY and are not directly visible to a TimeSync state machines. Client-level interface parameters include the timing information, based on the formats illustrated within Figure 9.2.




Figure 9.2—Formats of wireless-dependent times

9.1.2 Service interface overview

A sequence of 802.11v TimeSync service interface actions is illustrated in Figure 9.3 and summarized below:

- a) A periodic clock-slave trigger initiates the initial MLME_PRESENCE_REQUEST.request action.
- b) The clock-master gets an MLME_PRESENCE_REQUEST.indication upon request receipt.
The clock-slave gets an MLME_PRESENCE_REQUEST.confirm when the ack is returned.
- c) The clock-master processes the MLME_PRESENCE_REQUEST.indication parameters, returning them in MLME_PRESENCE_RESPONSE.request parameters for the clock-slave station.
- d) The clock-slave gets an MLME_PRESENCE_RESPONSE.indication upon response receipt.
The clock-master gets an MLME_PRESENCE_RESPONSE.confirm when the ack is returned.

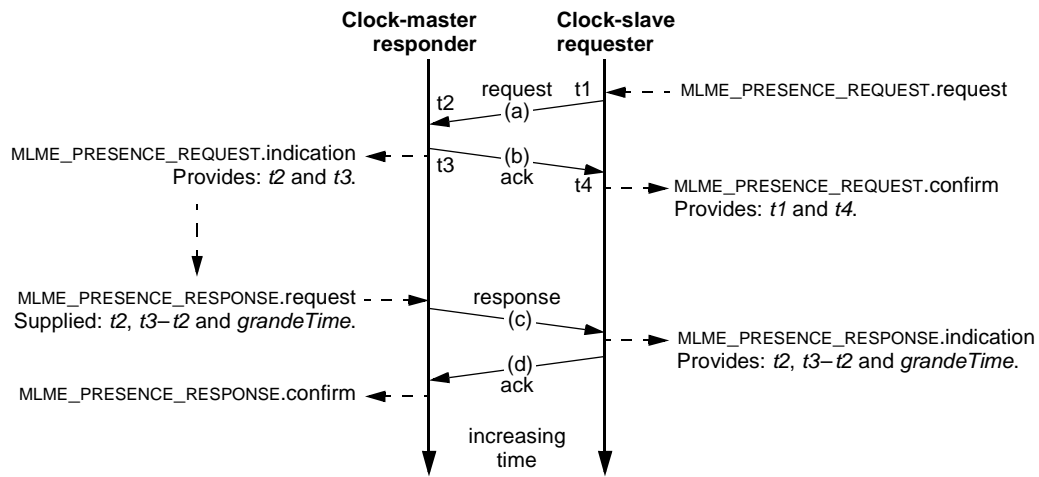


Figure 9.3—802.11v time-synchronization interfaces

The properties of these service interfaces are summarized below:

- MLME_PRESENCE_REQUEST.request**
Generated periodically by the clock-slave entity.
Triggers a (Figure 9.3a) request to fetch clock-master resident timing parameters.
- MLME_PRESENCE_REQUEST.indication**
Generated after receiving a (Figure 9.3a) request.
Provides t_2 and t_3 timing information to the clock-master entity.
- MLME_PRESENCE_REQUEST.confirm**
Generated after the (Figure 9.3b) request-ack is returned.
Provides $time1$ and $time4$ timing information to the clock-slave entity.
Confirms completion of the request transmission.
- MLME_PRESENCE_RESPONSE.request**
Generated shortly after processing a received (Figure 9.3a) request
Triggers a (Figure 9.3-c) response to update clock-slave resident timing parameters.
- MLME_PRESENCE_RESPONSE.indication**
Generated in response to receiving a (Figure 9.3c) response.
Provides $time2$, $time3 - time2$, and $grandeTime$ information to the clock-slave entity.
- MLME_PRESENCE_RESPONSE.confirm**
Generated after the (Figure 9.3d) ack is returned.
Confirms completion of the response transmission.

9.2 Service interface definitions

9.2.1 MLME_PRESENCE_REQUEST.request

9.2.1.1 Function

The service interface triggers the sending of a (Figure 9.3a) request from the clock-slave requester to the clock-master responder. A snapshot of the transmit time is also saved for deferred transmission/processing.

9.2.1.2 Semantics of the service primitive

The semantics of the primitives are as follows:

```
MLME_PRESENCE_REQUEST.request {  
    other_arguments    // Arguments for other purposes  
}
```

9.2.1.3 When generated

Generated periodically by a receive TS port as the first phase of a time-sync information transfer.

9.2.1.4 Effect of receipt

Upon receipt by a transmit TS port, an MLME_PRESENCE_REQUEST.indication is invoked; times of the arriving (Figure 9.3a) request and departing (Figure 9.3b) request-ack are both passed within this indication.

9.2.2 MLME_PRESENCE_REQUEST.indication

9.2.2.1 Function

The receipt of a (Figure 9.3a) request from the clock-slave requester triggers the return of an (Figure 9.3b) request-ack from the clock-master port. The transfer of an MLME_PRESENCE_REQUEST.indication to the clock-master provides snapshots of the (Figure 9.3a) request-receipt time as well as the following (Figure 9.3b) ack-transmit time.

9.2.2.2 Semantics of the service primitive

The semantics of the primitives are as follows:

```
MLME_PRESENCE_REQUEST.indication {  
    other_arguments,    // Arguments for other purposes  
    time_t2,           // Arrival time of request  
    time_t3            // Departure time of request-ack  
}
```

9.2.2.3 When generated

Generated by the receipt of a (Figure 9.3a) request during the first phase of a time-sync transfer.

9.2.2.4 Effect of receipt

Upon receipt, the times of the arriving (Figure 9.3a) request and (Figure 9.3b) request-ack are both saved for deferred processing.

9.2.3 MLME_PRESENCE_REQUEST.confirm

9.2.3.1 Function

The receipt of a (Figure 9.3b) request-ack at the clock-slave requester triggers the invocation of the MLME_PRESENCE_REQUEST.confirm service interface. The transmit time of the original (Figure 9.3a) request and the receive time of the recent (Figure 9.3b) request-ack are both provided.

9.2.3.2 Semantics of the service primitive

The semantics of the primitives are as follows:

```
MLME_PRESENCE_REQUEST.indication {
    other_arguments,    // Arguments for other purposes
    time_t1,           // Departure time of request
    time_t4            // Arrival time of confirm
}
```

9.2.3.3 When generated

Generated by the receipt of a (Figure 9.3b) request-ack during the initial phases of a time-sync transfer.

9.2.3.4 Effect of receipt

Upon receipt, the transmit time of the previous (Figure 9.3a) request and receive time of the recent (Figure 9.3b) request-ack are both saved for deferred processing.

9.2.4 MLME_PRESENCE_RESPONSE.request

9.2.4.1 Function

After the initial phases, a clock-slave requester triggers the transfer of an MLME_PRESENCE_RESPONSE.request. The transmit time of the original (Figure 9.3a) request, the transmit time of the recent (Figure 9.3b) request-ack, and the current time-sync related information are all included in the service primitives.

9.2.4.2 Semantics of the service primitive

The semantics of the primitives are as follows:

```
MLME_PRESENCE_RESPONSE.request {
    other_arguments,    // Arguments for other purposes
    time_t2            // Arrival time of request
    time_t32,         // Turn-round time
    grande_time       // Current media-dependent time
}
```

9.2.4.3 When generated

Triggered at the clock-master by the servicing of an MLME_PRESENCE_REQUEST.indication, indicating the completion of the initial time-sync phase.

9.2.4.4 Effect of receipt

Upon receipt, an MLME_PRESENCE_RESPONSE.indication is invoked, to provide the clock-slave with sufficient information to send a GrandSync PDU.

9.2.5 MLME_PRESENCE_RESPONSE.indication

9.2.5.1 Function

Additional information is provided to a clock-slave port. Along with previous information (saved earlier for deferred processing), the clock-slave has sufficient information to send a GrandSync PDU.

9.2.5.2 Semantics of the service primitive

The semantics of the primitives are as follows:

```
MLME_PRESENCE_RESPONSE.indication {  
    other_arguments,    // Arguments for other purposes  
    time_t2             // Arrival time of request  
    time_t32,          // Turn-round time  
    level_time,        // Current media-dependent time  
}
```

9.2.5.3 When generated

Triggered at the clock-slave by the receipt of a (Figure 9.3c) response, nearing the completion of the final time-sync phases.

9.2.5.4 Effect of receipt

Upon receipt, the clock-slave is provided with sufficient information to send a GrandSync PDU.

9.2.6 MLME_PRESENCE_RESPONSE.confirm

9.2.6.1 Function

Confirmation is provided to the clock-master, confirming clock-slave has sufficient information to send a GrandSync PDU.

9.2.6.2 Semantics of the service primitive

The semantics of the primitives are as follows:

```
MLME_PRESENCE_RESPONSE.confirm {  
    other_arguments,    // Arguments for other purposes  
}
```

9.2.6.3 When generated

Triggered at the clock-master by the receipt of a (Figure 9.3-d) response-ack, at the completion of the final time-sync phases.

9.2.6.4 Effect of receipt

Upon receipt, the clock-master is provided with a time-sync success status.

9.3 TimeSyncRxR11v state machine

9.3.1 Function

The TimeSyncRxR11v state machine consumes primitives provided by the MAC service interface and (in response) generates frames for the GrandSync entity.

9.3.2 State machine definitions

9.3.2.1 NULL: A constant that (by design) cannot be confused with a valid value.

9.3.2.2 queue values: Enumerated values used to specify shared FIFO queue structures.

Q_MS_IND—Queue identifier associated with the GrandSync receive port.

Q_S1_REQ—Queue identifier for MLME_PRESENCE_REQUEST.request parameters.

Q_S1_CON—Queue identifier for MLME_PRESENCE_REQUEST.confirm parameters.

Q_S2_IND—Queue identifier for MLME_PRESENCE_RESPONSE.indication parameters.

9.3.3 State machine variables

9.3.3.1 backTime: A variable representing the lapsed time since the remote request-ack transmission.

9.3.3.2 con1: Values returned within the MLME_PRESENCE_REQUEST.request service primitive:

time_t1—A local-timer snapshot at the (Figure 9.3a) request transmission.

time_t4—A local-timer snapshot at the (Figure 9.3b) request-ack reception.

9.3.3.3 ind2: Values returned within the MLME_PRESENCE_RESPONSE.indication service primitive:

grandTime—A remote snapshot of grandTime at the request-ack transmission.

time_t2—A remote-timer snapshot at the (Figure 9.3a) request reception.

time_t3—A remote-timer snapshot at the (Figure 9.3b) request-ack transmission.

9.3.3.4 ePtr: Points to entity-specific storage, comprising the following:

lastTime—The time of the last request transmission, for pacing periodic transmissions.

roundTrip—Saved (*con1.time_t4*–*con1.time_t4*) value.

rsInfo—Saved grandmaster selection values.

rxFastTimed—Saved *args2.fastTimed* value.

localTime—Saved *con1.time_t4* value.

syncInterval—The sync-interval associated with this clock-slave port.

9.3.3.5 grandTime: An variable representing the normalized/synchronized grandmaster time.

9.3.3.6 lapseTime: An variable representing the lapsed time since the request-ack reception.

9.3.3.7 localTime: A variable representing the calibrated one-way link delay.

9.3.3.8 radioDelay: An variable representing the round-trip transmission delay.

9.3.3.9 radioTime: An variable representing the current time, in media-specific units.

1 **9.3.3.10 req1:** A set of values returned within the MLME_PRESENCE_REQUEST.request service primitive,
2 consisting of other (unrelated) parameters.

3
4 **9.3.3.11 stationTime:** See 6.3.3.

5
6 **9.3.3.12 rsPtr:** A pointer to the *rsInfo* portion of *ePtr* referenced storage.

7
8 **9.3.3.13 turnRound:** Represents the difference between local time-sync transmit and receive times.

9
10 **9.3.3.14 turnStart:** An variable representing the remote time-sync transmit time.

11
12 **9.3.3.15 txInfo:** Storage for information sent to the GrandSync entity, comprising:

13 *destination_address, source_address, service_data_unit*

14 Where *service_data_unit* comprises:

15 *extraTime, sourcePort, function, grandTime, hopCount,*
16 *localTime, protocolType, precedence, syncInterval, version*

17
18 **9.3.3.16 txPtr:** A pointer to *txInfo* storage.

19
20 **9.3.4 State machine routines**

21
22 **9.3.4.1 Dequeue(queue):** See 6.3.4.

23
24 **9.3.4.2 Enqueue(queue, info):** See 6.3.4.

25
26 **9.3.4.3 R11vTime(entity):** Returns the local media-dependent free-running timer.

27
28 **9.3.4.4 SourcePort(entity):** See 7.3.3.

29
30 **9.3.4.5 StationTime(entity):** See 6.3.4.

31
32 **9.3.4.6 ClockSyncSdu(info):** See 6.3.4.

9.3.5 TimeSyncRxR11v state table

The TimeSyncRxR11v state machine consumes MAC-provided service-primitive information and forwards adjusted frames to the MAC-relay function, as illustrated in Table 9.1.

Table 9.1—TimeSyncRxR11v state machine table

Current		Row	Next	
state	condition		action	state
START	(stationTime – ePtr->lastTime) > ePtr->syncInterval	1	ePtr->lastTime = stationTime; req1 = SetupReq1(); Enqueue(Q_S1_REQ, req1);	START
	(con1 = Dequeue(Q_S1_CON)) != NULL	2	ePtr->localTime = con1.time_t4; ePtr->roundTrip = con1.time_t4 – con1.time_t1;	
	(ind2 = Dequeue(Q_S2_IND)) != NULL	3	turnStart = rxPtr->time_t2; turnRound = rxPtr->time_t32;	SINK
	—	4	stationTime = StationTime(ePtr); radioTime = R11vTime(ePtr);	START
SINK	—	5	linkDelay = (ePtr->roundTrip – turnRound) / 2; lapseDelay = (radioTime – ePtr->localTime); backTime = R11vToStation(lapseDelay + linkDelay); grandTime = GrandeToGrand(rxptr->grandeTime); txPtr->destination_address = AVB_MCAST; txPtr->source_address = TBD; tsPtr->protocolType = AVB_TYPE; tsPtr->function = AVB_FUNCTION; tsPtr->version = AVB_VERSION; tsPtr->precedence = rsPtr->precedence; tsPtr->sourcePort = SourcePort(ePtr); tsPtr->hopCount = rsPtr->hopCount; tsPtr->grandTime = grandTime; tsPtr->extraTime = rsPtr->extraTime; tsPtr->localTime = stationTime – backTime; Enqueue(Q_MS_IND, txPtr);	START

Row 9.1-1: Requests are sent at a periodic rate.

Row 9.1-2: Save the times that are available when the request-ack returns.

Row 9.1-3: Capture the parameters when the MLME_PRESENCE_RESPONSE.indication returns.

Row 9.1-4: Update times while waiting for state changes.

Row 9.1-5: Send accumulated/supplemented information to the GrandSync entity.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

9.4 TimeSyncTxR11v state machine

9.4.1 Function

The TimeSyncTxR11v state machine consumes GrandSync-generated frames, to maintain estimates of the current (*grandTime*, *stationTime*) and (*errorTime*, *stationTime*) affiliations. The TimeSyncTxR11v state machine also provides time-synchronization information through the MAC service interface, in response to clock-slave initiated requests.

9.4.2 State machine definitions

9.4.2.1 NULL: A value that (by design) cannot be confused with a valid value.

9.4.2.2 queue values: Enumerated values used to specify shared FIFO queue structures.

Q_MS_REQ—The queue identifier associated with frames sent from the GrandSync entity.

Q_S1_IND—Queue identifier for MLME_PRESENCE_REQUEST.indication parameters.

Q_S2_REQ—Queue identifier for MLME_PRESENCE_RESPONSE.request parameters.

Q_S2_ACK—Queue identifier for MLME_PRESENCE_RESPONSE.confirm parameters.

9.4.3 State machine variables

9.4.3.1 backInterval: A variable that represents the back-interpolation interval for transmit-time affiliations.

9.4.3.2 btPtr: A pointer to the ePtr->baseTimer storage.

9.4.3.3 con2: Values returned within the MLME_PRESENCE_RESPONSE.confirm service primitive:

9.4.3.4 ePtr: A pointer to the entity-specific storage containing the following:

baseTimer—Recently saved time events, each consisting of the following:

index—Index into the *timed[]* array, where last times were stored.

range—Number of entries within the *timed[]* array

timed[range]—Recently saved time events, each consisting of the following:

grandTime—A previously sampled grandmaster synchronized time.

extraTime—The residual error associated with the sampled *grandTime* value.

stationTime—The station-local time affiliated with the sampled *grandTime* value.

lastTime—The last transmit time, saved for pacing transmissions.

rxSaved—A copy of the last received GrandSync parameters.

syncInterval—The expected interval between successive time-sync transmissions.

9.4.3.5 ind1: Values returned within the MLME_PRESENCE_REQUEST.indication service primitive:

time_t2—A local snapshot at the time of (Figure 9.3a) request reception.

time_t3—A local snapshot at the time of (Figure 9.3b) request-ack transmission.

9.4.3.6 radioTime: A variable representing the media-dependent station-local time.

9.4.3.7 rateInterval: A variable representing the time interval over which the grandTime rate is measured.

9.4.3.8 req2: Values provided to the MLME_PRESENCE_REQUEST.request service primitive:

grandTime—A local snapshot of the grandTime as the request-ack transmission.

time_t2—Previously saved *ind1.time_t2* value.

time_t32—Previously saved (*ind1.time_t3*–*ind1.time_t2*) value.

9.4.3.9 rsPtr: A pointer to service-data-unit portion of *rxInfo* storage.

9.4.3.10 <i>rxInfo</i> : Storage for received time-sync PDU from the GrandSync entity, comprising:	1
<i>destination_address, source_address, service_data_unit</i>	2
Where <i>service_data_unit</i> comprises:	3
<i>extraTime, function, grandTime, hopCount,</i>	4
<i>precedence, protocolType, localTime, syncInterval, version</i>	5
	6
9.4.3.11 <i>rxPtr</i> : A pointer to <i>rxInfo</i> storage.	7
	8
9.4.3.12 <i>rxSyncInterval</i> : Represents the sync-interval associated with this station's clock-slave port.	9
	10
9.4.3.13 <i>ssPtr</i> : A pointer to the service-data-unit portion of the <i>ePtr->rxSaved</i> storage	11
	12
9.4.3.14 <i>stationTime</i> : A shared value representing current time. There is one instance of this variable for each station. Within the state machines of this standard, this is assumed to have two components, as follows:	13
<i>seconds</i> —An 8-bit unsigned value representing seconds.	14
<i>fraction</i> —An 40-bit unsigned value representing portions of a second, in units of 2^{-40} second.	15
	16
	17
9.4.3.15 <i>sendTime</i> : Represents the local time estimate of the remote request-ack transmission time.	18
	19
9.4.3.16 <i>sxPtr</i> : A pointer to the <i>ePtr->rxSaved</i> storage.	20
	21
9.4.3.17 <i>timeT2</i> : A variable that represents the request receipt time.	22
	23
9.4.3.18 <i>timeT3</i> : A variable that represents the request-ack transmit time.	24
	25
9.4.3.19 <i>txSyncInterval</i> : A variable that represents the sync-interval associated with this port.	26
	27
9.4.4 State machine routines	28
	29
9.4.4.1 <i>Dequeue(queue)</i> : See 6.3.4.	30
	31
9.4.4.2 <i>Enqueue(queue, info)</i> : See 6.3.4.	32
	33
9.4.4.3 <i>NextSaved(btPtr, rateInterval, grandTime, extaTime, thisTime)</i> : See 7.5.4.	34
	35
9.4.4.4 <i>NextTimed(btPtr, stationTime, backInterval)</i> : See 7.5.4.	36
	37
9.4.4.5 <i>StationTime(entity)</i> : See 7.3.3.	38
	39
9.4.4.6 <i>ClockSyncSdu(info)</i> : See 6.3.4.	40
	41
	42
	43
	44
	45
	46
	47
	48
	49
	50
	51
	52
	53
	54

9.4.5 TimeSyncTxR11v state table

NOTE—This state machine is preliminary; sequence timeouts have not been considered.

The TimeSyncTxR11v state machine includes a media-dependent timeout, which effectively disconnects a clock-slave port in the absence of received timedSync frames, as illustrated in Table 9.2.

Table 9.2—TimeSyncTxR11v state table

Current		Row	Next	
state	condition		action	state
START	(rxInfo = Dequeue(Q_MS_REQ)) != NULL	1	—	SINK
	(ind1 = Dequeue(Q_S1_IND)) != NULL	2	timeT2 = ind1.time_t2; timeT3 = ind1.time_t3;	SEND
	(con2 = Dequeue(Q_S2_CON)) != NULL	3	—	START
	—	4	stationTime = StationTime(ePtr); radioTime = R11vTime(ePtr); rxSyncInterval = ssPtr->syncInterval; txSyncInterval = ePtr->syncInterval; backInterval = (3 * rxSyncInterval + txSyncInterval) / 2; rateInterval = backInterval + (3 * txSyncInterval) / 2;	START
SINK	ClockSyncSdu(rsPtr)	5	ePtr->rxSaved = rxInfo; NextSaved(btPtr, rateInterval, rsPtr->grandTime, rsPtr->extraTime, stationTime);	SERVE
	—	6	Enqueue(Q_ES_REQ, rxPtr);	START
SEND	—	7	sendTime = stationTime – ((radioTime – timeT2) * RADIO_TIME); nextTimes = NextTimed(btPtr, stationTime, backInterval); req2.time_t2 = timeT2; req2.time_t32 = timeT3 – timeT2; req2.grandTime = GrandToR11v(nextTimes.grandTime); // If possible for extraTime. req2.extraTime = nextTimes.extraTime; Enqueue(Q_S2_REQ, req2);	WAIT2

Row 9.2-1: GrandSync generated frames are further checked before being processed.

Row 9.2-2: Save parameters from a service-interface primitive call.

Row 9.2-3: The final acknowledge provides a completion indication.

Row 9.2-4: Wait for the next change-of-state.

Row 9.2-5: Parameters from clockSync PDUs are saved.

Row 9.2-6: The contents of non-clockSync PDUs are passed through.

Row 9.2-7: Provide parameters for the MLME_PRESENCE_RESPONSE.response interface.

10. Ethernet passive optical network (EPON) state machines

NOTE—This clause is based on indirect knowledge of the Ethernet-PON (EPON) specifications, as interpreted by the author, and have not been reviewed by the 802.1 or 802.3 WGs. The intent was to provide a forum for evaluation of the GrandSync interfaces, while also triggering discussion of EPON design details. As such, the contents are highly preliminary and subject to change.

10.1 Overview

This clause specifies the state machines that support Ethernet passive optical network (EPON) based bridges. The operations are described in an abstract way and do not imply any particular implementations or any exposed interfaces. There is not necessarily a one-to-one correspondence between the formal specification and the interfaces in any particular implementation.

The EPON topology, as illustrated in Figure 10.1.




Figure 10.1—EPON topology

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

10.1.1 Link-dependent indications

The TimeSyncEpon state machines have knowledge of network-local synchronized *ticksTime* timers. With this knowledge, the TimeSyncEpon state machines can operated on frames received from the LLC, as illustrated in Figure 10.2. Link-dependent indications could be required for bridge ports attached to alternative media.

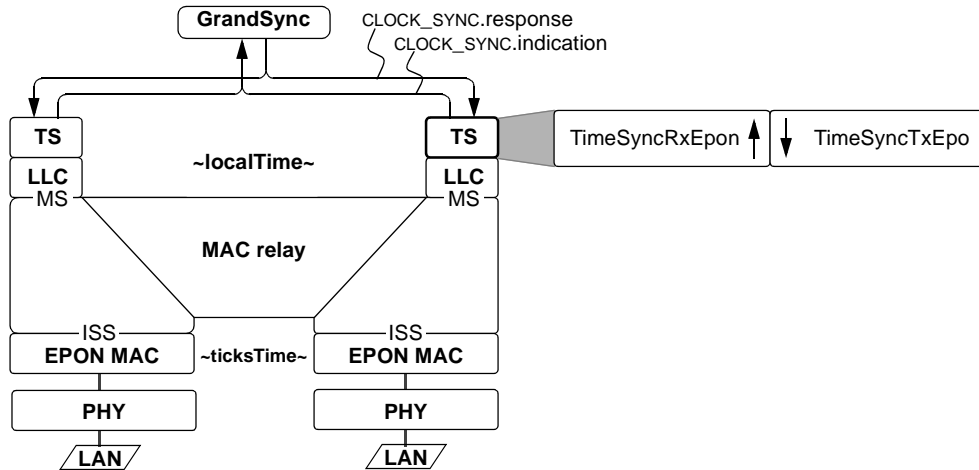


Figure 10.2—EPON interface model

The *localTime* values are represented as timers that are incremented once every 16 ns interval, as illustrated on the left side of Figure 10.3. Each synchronized local timer is roughly equivalent to a 6-bit *sec* (seconds) field and a 26-bit *fraction* (fractions of second) field timer, as illustrated on the right side of Figure 10.3.



Figure 10.3—Format of EPON-dependent times

The EPON MAC is supplied with frame transmit/receive snapshots, but these are transparent-to and not-used-by the TimeSync state machine. Instead, these are used to synchronize the *ticksTime* values in associated MACs and the TimeSyncEpon state machines have access to these synchronized *ticksTime* values.

10.1.2 Link-delay compensation

The synchronized-clock accuracies are influenced by the transmission delays between ports. To compensate for these transmission delays, the receive port is normally responsible for compensating {*grandTime*, *ticksTime*} affiliations by the (assumed to be constant) frame transmission delay.

The EPON MAC provides access to a subnet-synchronized media-dependent *ticksTime* timer. Thus, the {*grandTime*, *ticksTime*} affiliation specified the transmitter remains valid within the receiver and transmission-delay compensation (in this sense) is unnecessary.

However, each time-sync related GrandSync PDU includes an {*grandTime*, *stationTime*} affiliation, wherein *stationTime* represents a recent snapshot of a shared station-local clock. To provide such an affiliation, the transmission delay (measured as a *ticksTime* difference) is scaled and subtracted from the *stationTime* that is sampled when the conversion is performed. Thus, no additional receiver snapshot hardware is required.

10.2 timeSyncEpon frame format

The timeSyncEpon frames facilitate the synchronization of neighboring clock-master and clock-slave stations. The frame, which is normally sent at 10 ms intervals, includes time-snapshot information and the identity of the network’s clock master, as illustrated in Figure 10.4. The gray boxes represent physical layer encapsulation fields that are common across Ethernet frames.

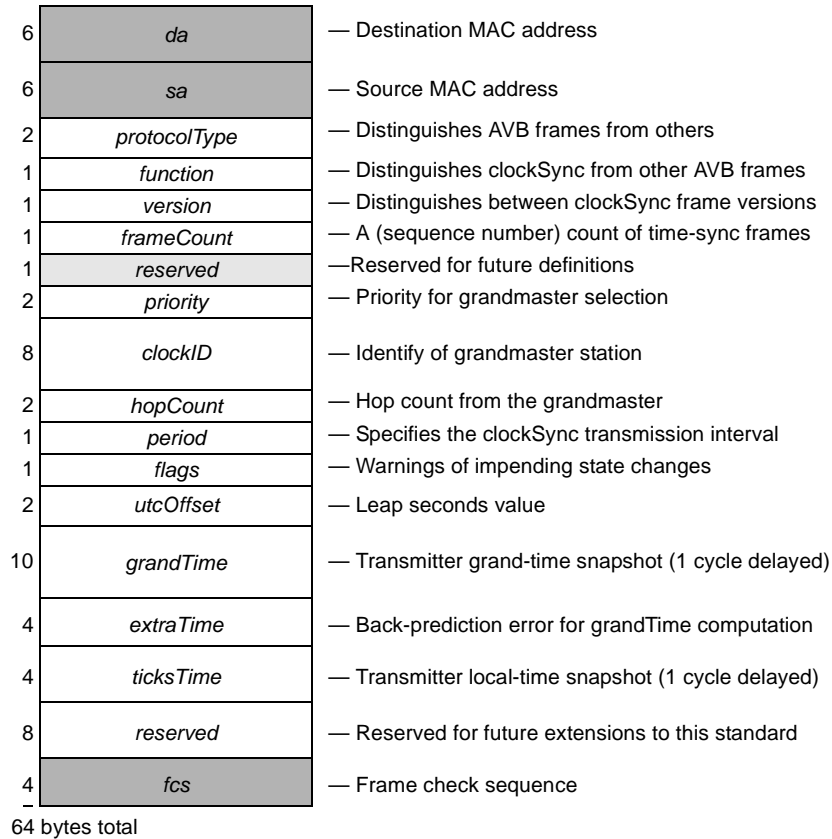


Figure 10.4—timeSyncEpon frame format

The 48-bit *da* (destination address), 48-bit *sa* (source address) field, 16-bit *protocolType*, 8-bit *function*, 8-bit *version*, 2-byte *priority*, 8-byte *clockID*, 2-byte *hopCount*, 1-byte *period*, 1-byte *flags*, 2-byte *utcOffset*, 80-bit *grandTime*, and 32-bit *extraTime* field are specified in 6.2.1.2..

10.2.1 ticksTime: A value representing local time in units of a 16 ns timer ticks, as illustrated in Figure 10.5.




Figure 10.5—tickTime format

10.3 TimeSyncRxEpon service interface primitives

10.3.1 ES_UNITDATA.indication

10.3.1.1 Function

Provides the TimeSyncRxEpon entity with clock-synchronization parameters derived from arriving time-sync frames.

10.3.1.2 Semantics of the service primitive

The semantics of the primitives are as follows:

```
ES_UNITDATA.indication {
    destination_address,    // Destination address
    source_address,        // Optional
    priority,               // Forwarding priority
    service_data_unit,     // Delivered content
    {
        protocolType,      // Distinguishes AVB frames from others
        function,          // Distinguishes between clockSync and other AVB frames
        version,           // Distinguishes between clockSync frame versions
        precedence,        // Precedence for grandmaster selection
        clockID,           // Precedence for grandmaster selection
        hopCount,          // Distance from the grandmaster station
        portID,            // Identifies the source port
        interval,          // Nominal clockSync transmission interval
        flags,              // Control flags
        utcOffset,         // Difference between UTC and TAI timescales
        grandTime,         // Global-time snapshot (1-cycle delayed)
        extraTime,         // Accumulated grandTime error
        ticksTime          // Local-time snapshot (1-cycle delayed)
    }
}
```

The parameters of the ES_UNITDATA.indication are described as follows:

The 48-bit **destination_address**, 48-bit **source_address**, and 8-bit **priority** field are specified in 6.2.1.2.

The service_data_unit consists of subfields; for content exchanged with the GrandTime protocol entity, these fields include the following.

The 16-bit **protocolType**, 8-bit **function**, 8-bit **version**, 14-byte **precedence**, 80-bit **grandTime**, 32-bit **extraTime**, and 8-bit **hopCount** fields are specified in 6.2.1.2.

10.3.1.2.1 frameCount: An 8-bit consistency-check field that increments on successive frames.

10.3.1.2.2 ticksTime: A 32-bit field that specifies the local free-running time within this subnet, when the previous clockSync frame was received (see 10.2.1).

10.3.1.3 When generated

The service primitive is generated upon the receipt of a time-sync related frame delivered from the MAC. The intent is to facilitate reformatting and snapshot-time adjustment before the content of that frame is delivered to the ClockMaster and TS entities.

10.3.1.4 Effect of receipt

The service primitive invokes processing of time-sync related content and forwarding of unrelated content. For time-sync related content, the processing included reformatting and compensation for receive-link transmission delays.

10.4 TimeSyncRxEpon state machine**10.4.1 Function**

The TimeSyncRxEpon state machine is responsible for receiving MAC-supplied frames, converting their media-dependent parameters, and sending normalized MAC-relay frames. The sequencing of this state machine is specified by Table 10.1; details of the computations are specified by the C-code of Annex G.

10.4.2 State machine definitions

10.4.2.1 NULL: A value that (by design) cannot be confused with a valid value.

10.4.2.2 queue values: Enumerated values used to specify shared FIFO queue structures.

Q_MS_IND—Associated with the GrandSync entity (see 6.3.2).

Q_ES_IND—The queue identifier associated with the received MAC frames.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

10.4.5 TimeSyncRxEpon state machine table

The TimeSyncRxEpon state machine associates PHY-provided sync information with arriving clockSync frames and forwards adjusted frames to the MAC-relay function, as illustrated in Table 8.2.

Table 10.1—TimeSyncRxEpon state machine table

Current		Row	Next	
state	condition		action	state
START	(rxInfo = Dequeue(Q_RX_MAC)) != NULL	1	—	TEST
	—	2	stationTime = StationTime(ePtr); ticksTime = EponTime(ePtr);	START
TEST	ClockSyncSdu(rsPtr)	3	*rxPtr = rxInfo;	SYNC
	—	4	Enqueue(Q_MS_IND, txInfo);	START
SYNC	rsPtr->hopCount != LAST_HOP	5	backTime = ticksTime - rsPtr->ticksTime; compTime = stationTime - TicksToTime(backTime); txPtr->destination_address = rxPtr->destination_address; txPtr->source_address = rxPtr->source_address; tsPtr->protocolType = rsPtr->protocolType; tsPtr->function = rsPtr->function; tsPtr->version = rsPtr->version; tsPtr->precedence = rsPtr->precedence; tsPtr->grandTime = rsPtr->grandTime; tsPtr->extraTime = rsPtr->extraTime; tsPtr->localTime = compTime; tsPtr->sourcePort = SourcePort(ePtr); tsPtr->hopCount = frame.hopCount; tsPtr->syncInterval = ePtr->syncInterval; Enqueue(Q_MS_IND, txInfo);	START
	—	6	—	—

Row 10.1-1: Initiate inspection of frames received from the lower-level MAC.

Row 10.1-2: Wait for the next frame to arrive.

Row 10.1-3: The clockSync frames are checked further.

Row 10.1-4: The non-clockSync frames are passed through.

Row 10.1-5: Active clockSync frames are adjusted for transfer delays and passed through.

Row 10.1-6: Overly-aged clockSync frames are discarded.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

10.5 TimeSyncTxEpon service interface primitives

10.5.1 ES_UNITDATA.request

10.5.1.1 Function

Provides the EPON entity with clock-synchronization parameters for constructing departing time-sync frames.

10.5.1.2 Semantics of the service primitive

The semantics of the primitives are as follows:

```
ES_UNITDATA.request
{
    destination_address, // Destination address
    source_address,     // Optional
    priority,           // Forwarding priority
    service_data_unit,  // Delivered content
    {                  // Contents of the service_data_unit
        protocolType,  // Distinguishes AVB frames from others
        function,      // Distinguishes between clockSync and other frames
        version,       // Distinguishes between clockSync frame versions
        precedence,    // Precedence for grandmaster selection
        grandTime,     // Global-time snapshot (1-cycle delayed)
        extraTime,     // Accumulated grandTime error
        hopCount,      // Distance from the grandmaster station
        ticksTime      // Local-time snapshot (1-cycle delayed)
    }
}
```

The parameters of the MA_UNITDATA.request are described in 10.3.1.2.

10.5.1.3 When generated

The service primitive is generated at a periodic rate, for the purposes of synchronizing the *grandTime* values resident in other stations.

10.5.1.4 Effect of receipt

The service primitive triggers the transmission of a clockSync frame on the affiliated port.

10.6 TimeSyncTxEpon state machine

10.6.1 Function

The TimeSyncTxEpon state machine is responsible for modifying time-sync CLOCK_SYNC.response parameters to form clockSync frames for transmission over the attached link.

10.6.2 State machine definitions

10.6.2.1 NULL: A value that (by design) cannot be confused with a valid value.

10.6.2.2 queue values: Enumerated values used to specify shared FIFO queue structures.

Q_MS_REQ—Associated with the GrandSync entity (see 6.3.2).

Q_ES_REQ—The queue identifier associated with frames sent to the MAC.

10.6.2.3 T10ms: A constant that represents a 10 ms value.

10.6.3 State machine variables

10.6.3.1 backInterval: Represents the back-interpolation interval for transmit-time affiliations.

10.6.3.2 ePtr: A pointer to an entity-specific data structure comprising the following:

baseTimer—Recently saved time events, each consisting of the following:

index—Index into the *timed[]* array, where last times were stored.

range—Number of entries within the *timed[]* array

timed[range]—Recently saved time events, each consisting of the following:

grandTime—A previously sampled grandmaster synchronized time.

extraTime—The residual error associated with the sampled *grandTime* value.

stationTime—The station-local time affiliated with the sampled *grandTime* value.

lastTime—The last PDU-transmit time; used to space periodic transmissions.

rxSaved—A copy of the last received GrandSync parameters.

syncInterval—The expected interval between time-sync frame transmissions.

10.6.3.3 rsPtr: A pointer to the service-data-unit portion of *rxInfo* storage.

10.6.3.4 rxInfo: Storage for the contents of GrandSync PDUs, comprising:

destination_address, *source_address*, *service_data_unit*

Where *service_data_unit* comprises:

extraTime, *function*, *grandTime*, *hopCount*, *precedence*,

protocolType, *localTime*, *syncInterval*, *version*

10.6.3.5 rxPtr: A pointer to the *rxInfo* storage.

10.6.3.6 rxSyncInterval: Represents the sync-interval associated with this station's clock-slave port.

10.6.3.7 stationTime: See 6.3.3.

10.6.3.8 ssPtr: A pointer to the service-data-unit portion of the *ePtr->rxSaved* storage

10.6.3.9 sxPtr: A pointer to the *ePtr->rxSaved* storage.

10.6.3.10 tsPtr: A pointer to the service-data-unit portion of *txInfo* storage.

10.6.3.11 <i>txInfo</i> : Storage for a to-be-transmitted MAC frame, comprising:	1
<i>destination_address, source_address, service_data_unit</i>	2
Where <i>service_data_unit</i> comprises:	3
<i>extraTime, function, grandTime, hopCount,</i>	4
<i>protocolType, precedence, ticksTime, version</i>	5
	6
10.6.3.12 <i>txPtr</i> : A pointer to the <i>txInfo</i> storage.	7
	8
10.6.3.13 <i>ticksTime</i> : A 32-bit shared value representing EPON media-dependent time; incremented every 16 ns.	9
	10
	11
10.6.4 State machine routines	12
	13
10.6.4.1 <i>Dequeue(queue)</i> : See 6.3.4.	14
	15
10.6.4.2 <i>Enqueue(queue, info)</i> : See 6.3.4.	16
	17
10.6.4.3 <i>NextTimed(btPtr, stationTime, backInterval)</i> : See 7.5.4.	18
	19
10.6.4.4 <i>SourcePort(entity)</i> : See 7.3.3.	20
	21
10.6.4.5 <i>StationTime(entity)</i> : See 6.3.4.	22
	23
10.6.4.6 <i>TicksTime(entity)</i> : See 10.4.4.	24
	25
10.6.4.7 <i>ClockSyncSdu(info)</i> : See 6.3.4.	26
	27
	28
	29
	30
	31
	32
	33
	34
	35
	36
	37
	38
	39
	40
	41
	42
	43
	44
	45
	46
	47
	48
	49
	50
	51
	52
	53
	54

10.6.5 TimeSyncTxEpon state machine table

The TimeSyncTxEpon state machine includes a media-dependent timeout, which effectively disconnects a clock-slave port in the absence of received timeSyncEpon frames, as illustrated in Table 10.2.

Table 10.2—TimeSyncTxEpon state machine table

Current		Row	Next	
state	condition		action	state
START	(rxInfo = Dequeue(Q_MS_REQ)) != NULL	1	—	SINK
	(stationTime – ePtr->lastTime) > T10ms	2	ePtr->lastTime = stationTime;	SEND
	—	3	stationTime = StationTime(ePtr); ticksTime = TicksTime(ePtr);	START
SINK	ClockSyncSdu(rsPtr)	4	ePtr->rxSaved = rxInfo;	START
	—	5	Enqueue(Q_ES_REQ, rxPtr);	
SEND	—	6	rxSyncInterval = ssPtr->syncInterval; txSyncInterval = ePtr->syncInterval; backInterval = (3 * rxSyncInterval + txSyncInterval) / 2; nextTimes = NextTimed(btPtr, stationTime, backInterval); txPtr->destination_address = sxPtr->destination_address; txPtr->source_address = sxPtr->source_address; tsPtr->protocolType = ssPtr->protocolType; tsPtr->function = ssPtr->function; tsPtr->version = ssPtr->version; tsPtr->precedence = ssPtr->precedence; tsPtr->hopCount = ssPtr->hopCount; tsPtr->grandTime = nextTimes.grandTime; tsPtr->extraTime = nextTimes.extraTime; tsPtr->ticksTime = ticksTime; Enqueue(Q_ES_REQ, txPtr);	START

Row 10.2-1: Relayed frames are further checked before being processed.

Row 10.2-2: Transmit periodic clockSync frames.

Row 10.2-3: Wait for the next change-of-state.

Row 10.2-4: The clockSync PDU is saved and processed further.

Row 10.2-5: Non-clockSync PDUs are retransmitted in the standard fashion.

Row 10.2-6: Format and transmit the media-specific clockSync frame.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

Annexes

Annex A

(informative)

Bibliography

- [B1] IEEE 100, The Authoritative Dictionary of IEEE Standards Terms, Seventh Edition.¹
- [B2] IEEE Std 802-2002, IEEE Standards for Local and Metropolitan Area Networks: Overview and Architecture.
- [B3] IEEE Std 801-2001, IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture.
- [B4] IEEE Std 802.1D-2004, IEEE Standard for Local and Metropolitan Area Networks: Media Access Control (MAC) Bridges.
- [B5] IEEE Std 1394-1995, High performance serial bus.
- [B6] IEEE Std 1588-2002, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems.
- [B7] IETF RFC 1305: Network Time Protocol (Version 3) Specification, Implementation and Analysis, David L. Mills, March 1992²
- [B8] IETF RFC 2030: Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI, D. Mills, October 1996.

¹IEEE publications are available from the Institute of Electrical and Electronics Engineers, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331, USA (<http://standards.ieee.org/>).

²IETF publications are available via the World Wide Web at <http://www.ietf.org>.

Annex B

(informative)

Time-scale conversions

B.1 Overview

For historical reasons, time is specified in a variety of ways as listed in Table B.1. GPS, PTP, and TAI times are based on values yielded by atomic clocks and advance on each second. NTP and UTC times are similar, but are occasionally adjusted by one leap-second, to account for differences between the atomic clocks and the rotation time of the earth.

Table B.1—Time-scale parameters

Parameter	Time scale				
	GPS	PTP	TAI	NTP	UTC
approximate epoch	1980-01-06 1999-08-22	1970-01-01	1972-01-01*	1900-01-01	1972-01-01*
representation	weeks.seconds	seconds	YYYY-MM-DD hh:mm:ss	seconds	YYYY-MM-DD hh:mm:ss
rollover (years)	19.7	8,925,513	10,000	136.19	10,000
leapSeconds	no			yes	

Notes:

* The TAI time when TAI and UTC were first specified to deviate by only integer seconds.
(There is no true epoch for the TAI and UTC time scales.)

GPS global positioning satellite

NTP Network Time Protocol

PTP Precision Time Protocol (commonly used in POSIX)

TAI International Atomic Time (from the French term *Temps Atomique International*)

UTC Coordinated Universal Time (a compromise between the English and French):

English speakers wanted the initials of their language: CUT for "coordinated universal time"

French speakers wanted the initials of their language: TUC for "temps universel coordonné".

B.2 TAI and UTC

TAI and UTC are international standards for time based on the SI second; both are expressed in days, hours, minutes and seconds. TAI is implemented by a suite of atomic clocks and forms the timekeeping basis for other time scales in common use. The rate at which UTC time advances is normally identical to the rate of TAI. An exception is an occasion when UTC is modified by adding or subtracting leap seconds.

Prior to 1972-01-01, corrections to the offset between UTC and TAI were made in fractions of a second. After 1972-01-01, leap-second corrections are applied to UTC preferably following second 23:59:59 of the last day of June or December. As of 2006-01-01, TAI and UTC times differed by +33 seconds.

In POSIX based computer systems, the common time conversion algorithms can produce the correct ISO 8601-2004 printed representation format "YYYY-MM-DD hh:mm:ss" for both TAI and UTC.

The PTP epoch is set such that a direct application of the POSIX algorithm to a PTP time-scale timestamp yields the ISO 8601-2004 printed representation of TAI. Subtracting the current *leapSeconds* value from a PTP timestamp prior to applying the POSIX algorithm yields the ISO 8601-2004 printed representation of UTC. Conversely, applying the inverse POSIX algorithm and adding *leapSeconds* converts from the ISO 8601-2004 printed form of UTC to the form convenient for generating a PTP timestamp.

Example: The POSIX algorithm applied to a PTP timestamp value of 8 seconds yields 1970-01-01 00:00:08 (eight seconds after midnight on 1970-01-01 TAI). At this time the value of *leapSeconds* was approximately 8 seconds. Subtracting this 8 seconds from this time yields 1970-01-01 00:00:00 UTC.

Example: The POSIX algorithm applied to a PTP timestamp value of 0 seconds yields 1970-01-01 00:00:00 TAI. At this time the value of *leapSeconds* was approximately 8 seconds. Subtracting this 8 seconds from this time yields 1969-12-31 23:59:52 UTC.

B.3 NTP and GPS

Two standard time sources of particular interest in implementing PTP systems: NTP and GPS. Both NTP and GPS systems are expected to provide time references for calibration of the grandmaster supplied PTP time.

NTP represents seconds as a 32 bit unsigned integer that rolls-over every 2^{32} seconds \approx 136 years, with the first such rollover occurring in the year 2036. The precision of NTP systems is usually in the millisecond range.

NTP is a widely used protocol for synchronizing computer systems. NTP is based on sets of servers, to which NTP clients synchronize. These servers themselves are synchronized to time servers that are traceable to international standards.

NTP provides the current time. In NTP version 4, the current *leapSeconds* value and warning flags marking indicating when a *leapSecond* will be inserted at the end of the current UTC day. The NTP clock effectively stops for one second when the leap second is inserted.

GPS time comes from a global positioning satellite system, GPS, maintained by the U.S. Department of Defense. The precision of GPS system is usually in the 10-100 ns range. GPS system transmissions represent the time as $\{weeks, secondsInWeek\}$, the number of weeks since the GPS epoch and the number of seconds since the beginning of the current week.

GPS also provides the current *leapSeconds* value, and warning flags marking the introduction of a leap second correction. UTC and TAI times can be computed solely based the information contained in the GPS transmissions.

GPS timing receivers generally manage the epoch transitions (1024-week rollovers), providing the correct time (YYYY-MM-DD hh:mm:ss) in TAI and/or UTC time scales, and often also local time; in addition to providing the raw GPS week, second of week, and leap second information.

B.4 Time-scale conversions

Previously discussed representations of time can be readily converted to/from PTP *time* based on a constant offsets and the distributed *leapSeconds* value, as specified in Table B.2. Within Table B.2, all variables represent integers; '/' and '%' represent a integer divide and remainder operation, respectively.

Table B.2—Time-scale conversions

<i>ta</i>		PTP value <i>tb</i> :
name	format	
GPS	weeks:seconds	$tb = ta.seconds + 315\,964\,819 + (gpsRollovers * 1024 + ta.weeks) * (7 * DAYSECS);$
		$ta.weeks = (tb - 315\,964\,819) / (7 * DAYSECS);$ $ta.days = (tb - 315\,964\,819) \% (7 * DAYSECS);$
TAI	date{YYYY,MM,DD}:time{hh,mm,ss}	$tb = DateToDays("1970-01-01", ta.date) * DAYSECS + ((ta.time.hh * 24) + ta.time.mm) * 60 + ta.time.ss;$
		$secs = tb \% DAYSECS;$ $ta.date = DaysToDate("1970-01-01", tb / DAYSECS);$ $ta.time.hh = secs / 3600;$ $ta.time.mm = (secs \% 3600) / 60;$ $ta.time.ss = (secs \% 60);$
NTP	seconds	$tb = (ta + leapSeconds) - 2\,208\,988\,800;$
		$ta = (tb - leapSeconds) + 2\,208\,988\,800;$
UTC	date{YYYY,MM,DD}:time{hh,mm,ss}	$tb = DateToDays("1970-01-01", ta.date) * DAYSECS + ((ta.time.hh * 24) + ta.time.mm) * 60 + ta.time.ss + leapSeconds;$
		$tc = tb - leapSeconds;$ $secs = tc \% DAYSECS;$ $ta.date = DaysToDate("1970-01-01", tc / DAYSECS);$ $ta.time.hh = secs / 3600;$ $ta.time.mm = (secs \% 3600) / 60;$ $ta.time.ss = (secs \% 60);$

Note:

gpsRollovers Currently equals 1; changed from 0 to 1 between 1999-08-15 and 1999-08-22.

DAYSECS The number of seconds within a day: (60*60*24).

leapSeconds Extra seconds to account for variations in the earth-rotation times: 33 on 2006-01-01.

DateToDays For arguments *DateToDays(past, present)*, returns days between *past* and *present* dates.

DaysToDate For arguments *DaysToDate(past, days)*, returns the current date, *days* after the *past* date.

B.5 Time zones and GMT

The term Greenwich Mean Time (GMT) once referred to mean solar time at the Royal Observatory in Greenwich, England. GMT now commonly refers to the time scale UTC; or the UK winter time zone (Western European Time, WET). Such GMT references are strictly speaking incorrect; but nevertheless quite common. The following representations correspond to the same instant of time:

18:07:00 (GMT), commonplace usage	13:07:00 (Eastern Standard Time, EST)
18:07:00 (UTC)	1:07 PM (Eastern Standard Time, EST)
18:07:00 (Western European Time, WET)	10:07:00 (Pacific Standard Time, PST)
6:07 PM (Western European Time, WET)	10:07 AM (Pacific Standard Time, PST)

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

Annex C

(informative)

Reclocked clockSync requirements

C.1 Cascaded clock considerations

C.1.1 Cascading causes sync-interval bunching

The naive approach towards forwarding time-synchronization information is to quickly propagate time-reference snapshots through successive stations. Unfortunately, relatively small ($\frac{1}{4}$ interval) residence-time delays per station can cause significant bunching, as illustrated in Figure C.1.




Figure C.1—Cascading causes sync-interval bunching

Techniques for avoiding such bunching are well known and practiced in the form of relocked synchronous circuits. For example, Ethernet stations accept (baud-rate) information at a closely matched input clock rate, relock the data with a local reference, and regenerate information without degraded jitter performance.

C.1.2 Reclocking eliminates sync-interval bunching

Applying these techniques to clock-sync transmission is straightforward. Rather than quickly forwarding these frames, their information is saved. That saved information is then forwarded in the same periodic fashion, based on local-station timing, as illustrated in Figure C.2. While such relocked systems more susceptible to gain-peaking/whiplash effects, inherent design and verification simplicities favor their use.




Figure C.2—Reclocking eliminates sync-interval bunching

C.1.3 Reclocking localizes sync-interval properties

The relocked sync-interval strategy is compatible with bridged mixed-media systems. The persistent or transient sync-interval rate of an intermediate (perhaps longer or more power sensitive) link could be less than the rate assumed for the clock-master, as illustrated in the center of Figure 3.3. Similarly, wireless links could base their timing events on triggers initiated by the clock-slave station, as illustrated in the right side of Figure 3.3.




Figure 3.3—Reclocking localizes sync-interval properties

Other flow-through clocking designs would require special “boundary clock” architectures to support such mixed systems. With the interval retiming strategy, the additional (specification and implementation) complexities of such boundary-clock architectures are easily avoided.

C.2 Sampling offset/rate conversion

Each clock-master port is responsible for using its received $\{grandTime, extraTime, rcTime\}$ affiliations to derive distinct $\{grandTimeI, extraTimeI, txTime\}$ affiliations that are transmitted to its neighbor. Since the values of $rcTime$ and $txTime$ are (by convention) coupled to the receive and transmit times, this update involves generation of $\{grandTime, extraTime, rcTime\}$ triads by resampling within the array of previously saved $\{grandTime, extraTime, rcTime\}$ triads.

C.2.1 Forward extrapolation inaccuracies

A typical design approach (and that used by IEEE Std 1588) views the received $\{grandTime, rcTime\}$ affiliations as points on a curve, sampled at received-snapshot times $rc[n]$. The objective is to generate the distinct set of $\{grandTimeI, txTime\}$ affiliations by extrapolating from a distinct set of receive-snapshot times $rc[n]$, as illustrated in Figure 3.4.





Figure 3.4—Extrapolation for grandTime

1 Extrapolation techniques exhibit gain peaking at frequencies whose wavelength is twice the $\{rc[n-N],rc[n]\}$
 2 slope-averaging interval, because the extrapolated value can exceed what would have been the sampled time
 3 value. A cascade of multiple stations emphasizes the gain-peaking inaccuracies, allowing errors to
 4 accumulate in an $O(N^2)$ fashion.

6 C.2.2 Forward interpolation inaccuracies


8 To reduce gain-peaking effects, the resampling computation can be migrated to a safe-interpolation domain.
 9 This involves subtracting a *backTime* constant from *txTime*, yielding a new time *tbTime*, for which a less
 10 gain-peaking sensitive interpolation is viable, as illustrated in Figure C.5. In concept, the stale (but not
 11 incorrect) $\{grandTime,rcTime\}$ affiliations could be passed to the terminal clock-slave stations, wherein a
 12 single extrapolation-to-the-future accumulation could be performed. A preferred technique is to compensate
 13 the interpolation result on an per-station basis as the time-reference flows towards the clock-slave station, as
 14 discussed in the following subclasses.



27 **Figure C.5—Extrapolation for *grandTime***

30 C.2.3 Backward *grandTime* interpolation

31 A more-scalable backward-interpolation approach also views the received $\{grandTime,rcTime\}$ affiliations
 32 as points on a curve. The objective is to generate the distinct set of $\{grandTimeI,txTime\}$ affiliations by
 33 interpolating within a distinct set of $\{grandTime, rcTime\}$ points on the curve, as illustrated in Figure C.6.



46 **Figure C.6—Interpolation for *grandTimeA***

48
$$rateRatio = (grandTime[n] - grandTime[n-N]) / (rc[n] - rc[n-N]) \quad (3.1)$$

49
$$grandTimeI[m] = grandTime[n] + rateRatio * ((txTime - backTime) - rc[n]) + backTime; \quad (3.2)$$

 50 *backTime* is a constant (sync-interval dependent) value.

51
$$extraTimeI[m] = (rateRatio - ONE) * backTime; \quad (3.3)$$

The advantage of this technique is the separation of *grandTime[m]* and *extra[m]* components. The interpolation process eliminates gain-peaking for the *grandTime[m]* value, thus reducing error effects when passing through multiple bridges. The sideband *extraTime* signal remains significant, and is therefore carried through bridges, so that the cumulative *grandTime[m]+extraTime[m]* value can be passed to the end-point application.

From an intuitive perspective, the whiplash-free nature of the back-in-time interpolation is attributed to the use of interpolation (as opposed to extrapolation) protocols. Interpolation between input values never produces a larger output value, as would be implied by a gain-peaking (larger-than-unity gain) algorithm. A disadvantage of back-in-time interpolation is the requirement for a side-band *extraTime* communication channel, over which the difference between nominal and rate-normalized *backTime* values can be transmitted.

C.2.4 Backward *extraTime* Averaging

An averaging (rather than backward-interpolation) approach is applied to the received {*extraTime*, *rcTime*} affiliations as points on a curve, sampled at received-snapshot times *rc[n]*. The {*extraTime*, *tx[m]*} affiliations are produced by averaging recently observed *extraTime* values, as illustrated in Figure C.7.




Figure C.7—Interpolation of *extraTimeD*

$$extraTimeA[m] = (extraTime[n-(N-1)] + \dots + extraTime[n]) / N \tag{3.4}$$

$$extraTimeI[m] = extraTimeA[m] + extraTimeB[m]; \tag{3.5}$$

The to-be-transmitted value of *extraTimeI[m]* consists of a contribution *extraTimeA* (accumulated from previous stations’s *grandTime* interpolations) and a contribution *errorTimeB* (coming from this station’s *grandTime* interpolation). Note that the averaging of *extraA* values is effectively a low-pass filtering process that removes noise without causing a gain-peaking frequency response.

NOTE—For simplicity and scalability, the computed *extraTimeI* time is based on *N*, a fixed number of samples, where *N* is a convenient power-of-two in size.

Annex D

(informative)

Simulation results (preliminary)

D.1 Simulation environment

This annex describes several simulations performed with the intent of comparing time-extrapolation and time-interpolation algorithms. To reduce possibilities of code-conversion errors, the simulation model executes the C code of Annex G. Simulation time is based on a 128-bit *systemTime*, represented by 64-bit seconds and fractions-of-second components, to ensure that precision and range are not constraining factors.

The simulation consists of *bridgeCount* identical super-bridge components, as illustrated in Figure D.1. For generality and uniformity, each bridge includes ClockMaster and ClockSlave entities. The smallest MAC address is assigned to the left-most station; for other stations, the address is incremented for each sequential right-side bridge. The simulations assumed *bridgeCount* values of 8 (the assumed AVB diameter) and 64 (a reasonable IEEE 802.17 ring diameter).




Figure D.1—Time-synchronization flows

The transmit portion of the TS component (emulated by the DuplexTxExec routine) introduces a random delay of no more than 2.5 ms, thus emulating delays consistent with the 10 ms sync-frame transmission rate. A 20 ns sampling clock ambiguity (corresponding to 25 MHz) is incorporated into the MAC component (emulated by the DupMacTxExec routine).

The cable is modeled as a symmetric 500ns delay, corresponding to a cable length of approximately 100 meters.

Station clock accuracies are assigned randomly/uniformly within the range of the allowed ± 100 PPM deviation from the simulation's emulated/exact *systemTime* reference.

NOTE—Please be tolerant of the editor of this document, who just downloaded the gnuplot application and fft4 library today. These initial cut-and-paste of plots are primitive (to be improved, when EPS or other formats are understood) and no noise-spectrum plots (to better illustrate gain peaking) are currently available. Improvements expected soon...

D.2 Initialization transients

D.2.1 Cascaded 8 stations

A significant expected initialization transient is observed when all stations simultaneously start operations, as illustrated in Figure D.2. This can be contributed to inaccurate initial estimates of receiver's link-delay and transmitter's rate estimations. The transient delays (although significant) are much less than expected from designs based on many-sample grandmaster rate-syntonization delays within bridges.




Figure D.2—Startup transients with 8 stations

D.2.2 Cascaded 64 stations

The length of the initialization transient increases when the number of bridges is increased to 64, as illustrated in Figure D.3. The much-longer duration of such transients is perhaps tolerable, but illustrates the desire to avoid extrapolation-based on many-sample grandmaster rate-syntonization delays within bridges.




Figure D.3—Startup transients with 64 stations

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

D.3 Steady-state interpolation errors

D.3.1 Time interpolation with 8 stations

Simulations indicate modest peak-to-peak errors for 8-bridge topologies when interpolation-based protocols are used, as illustrated in Figure D.4.




Figure D.4—Time interpolation with 8 stations

D.3.2 Time interpolation with 64 stations

Simulations indicate modest peak-to-peak error increases for 64-bridge topologies (as expected to 8-bridge topologies) when interpolation-based protocols are used, as illustrated in Figure D.5. The data is consistent with less-than-linear expectations, due to statistical averaging and intermediate interpolation filtering.




Figure D.5—Time interpolation with 64 stations

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

D.4 Steady-state extrapolation errors

D.4.1 Time extrapolation with 8 stations

Simulations indicate approximately twice the errors for 8-bridge topologies when extrapolation-based protocols (as opposed to interpolation-based protocols) are used, as illustrated in Figure D.6.




Figure D.6—Time extrapolation with 8 stations

D.4.2 Time extrapolation with 64 stations

Simulations indicate significantly larger peak-to-peak errors for 64-bridge topologies when extrapolation-based protocols (as opposed to interpolation-based protocols) are used, as illustrated in Figure D.7.




Figure D.7—Time extrapolation with 64 stations

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

Annex E

(informative)

Bridging to IEEE Std 1394

To illustrate the sufficiency and viability of the AVB time-synchronization services, the transformation of IEEE 1394 packets is illustrated.

E.1 Hybrid network topologies

E.1.1 Supported IEEE 1394 network topologies

This annex focuses on the use of AVB to bridge between IEEE 1394 domains, as illustrated in Figure E.1. The boundary between domains is illustrated by a dotted line, which passes through a SerialBus adapter station.




Figure E.1—IEEE 1394 leaf domains

E.1.2 Unsupported IEEE 1394 network topologies

Another approach would be to use IEEE 1394 to bridge between IEEE 802.3 domains, as illustrated in Figure E.2. While not explicitly prohibited, architectural features of such topologies are beyond the scope of this working paper.




Figure E.2—IEEE 802.3 leaf domains

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

E.1.3 Time-of-day format conversions

The difference between AVB and IEEE 1394 time-of-day formats is expected to require conversions within the AVB-to-1394 adapter. Although multiplies are involved in such conversions, multiplications by constants are simpler than multiplications by variables. For example, a conversion between AVB and IEEE 1394 involves no more than two 32-bit additions and one 16-bit addition, as illustrated in Figure E.3.

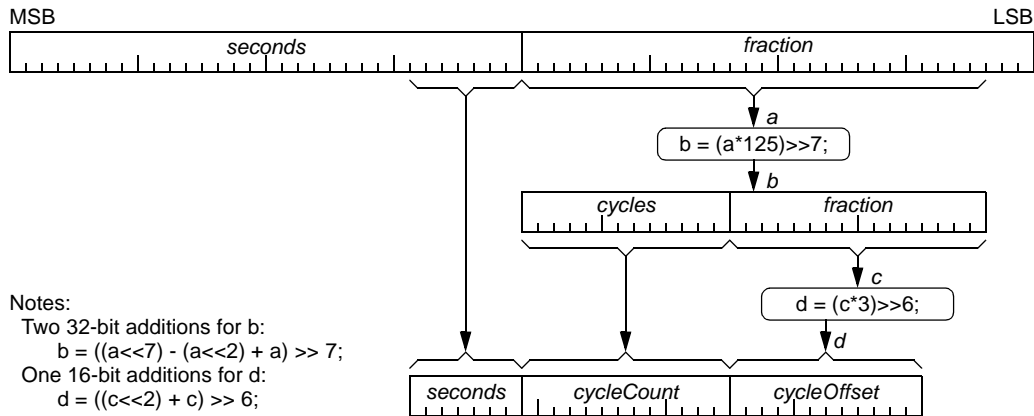


Figure E.3—Time-of-day format conversions

E.1.4 Grandmaster precedence mappings

Compatible formats allow either an IEEE 1394 or IEEE 802.3 stations to become the network’s grandmaster station. While difference in format are present, each format can be readily mapped to the other, as illustrated in Figure E.4:




Figure E.4—Grandmaster precedence mapping

Annex F

(informative)

Time-of-day format considerations

To better understand the rationale behind the ‘extended binary’ timer format, various possible formats are described within this annex.

F.1 Possible time-of-day formats

F.1.1 Extended binary timer formats

The extended-binary timer format is used within this working paper and summarized herein. The 64-bit timer value consist of two components: a 40-bit *seconds* and 40-bit *fraction* fields, as illustrated in Figure F.1.




Figure F.1—Global-time subfield format

The concatenation of 40-bit *seconds* and 40-bit *fraction* field specifies an 80-bit *time* value, as specified by Equation F.1.

$$time = seconds + (fraction / 2^{40}) \tag{F.1}$$

Where:

- seconds* is the most significant component of the time value.
- fraction* is the less significant component of the time value.

F.1.2 IEEE 1394 timer format

An alternate “1394 timer” format consists of *secondCount*, *cycleCount*, and *cycleOffset* fields, as illustrated in Figure F.2. For such fields, the 12-bit *cycleOffset* field is updated at a 24.576MHz rate. The *cycleOffset* field goes to zero after 3071 is reached, thus cycling at an 8kHz rate. The 13-bit *cycleCount* field is incremented whenever *cycleOffset* goes to zero. The *cycleCount* field goes to zero after 7999 is reached, thus restarting at a 1Hz rate. The remaining 7-bit *secondCount* field is incremented whenever *cycleCount* goes to zero.




Figure F.2—IEEE 1394 timer format

F.1.3 IEEE 1588 timer format

IEEE Std 1588-2002 timer format consists of seconds and nanoseconds fields components, as illustrated in Figure F.3. The nanoseconds field must be less than 10^9 ; a distinct *sign* bit indicates whether the time represents before or after the epoch duration.




Figure F.3—IEEE 1588 timer format

F.1.4 EPON timer format

The IEEE 802.3 EPON timer format consists of a 32-bit scaled nanosecond value, as illustrated in Figure F.4. This clock is logically incremented once each 16 ns interval.




Figure F.4—EPON timer format

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

Annex G

(informative)

C-code illustrations

NOTE—This annex is provided as a placeholder for illustrative C-code. Locating the C code in one location (as opposed to distributed throughout the working paper) is intended to simplify its review, extraction, compilation, and execution by critical reviewers. Also, placing this code in a distinct Annex allows the code to be conveniently formatted in 132-character landscape mode. This eliminates the need to truncate variable names and comments, so that the resulting code can be better understood by the reader.

This Annex provides code examples that illustrate the behavior of AVB entities. The code in this Annex is purely for informational purposes, and should not be construed as mandating any particular implementation. In the event of a conflict between the contents of this Annex and another normative portion of this standard, the other normative portion shall take precedence.

The syntax used for the following code examples conforms to ANSI X3T9-1995.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37

```

// *****
//
//      1      2      3      4      5      6      7      8      9      0      1      2      3
// 345678901234567890123456789012345678901234567890123456789012345678901234567890123456789012
// *****
// ***** by Dr. David V James, 2007-05-18 *****
// *****

// NOTE--The following code is portable with respect to endian ordering,
// but (for clarity and simplicity) assumes availability of 64-bit integers.

// TBDs:
// Active:
//   Suggested changes for clarity (anonymous reviewer 2007-05-11):
//   Separate partition for the bulk of header
//   Make TBDs explicit
//   Consider name change: state machines => queue service routines
//   Better use of spaces and comments
//   FFT and spectrum analysis via utilities or piped process
// Completed:
//   Initial checks to be more descriptive, as in "GroupAsserts"
//   Consistent terminology: backInterval
//   Sequence of tests within looks, possibly with "serviced", as in:
//   for (checkForMore = TRUE; checkForMore == TRUE; ) {
//     if (something) {
//       checkForMore = TRUE;
//     }
//     ..
//   } // for(;;) ends here
//   Ports=>queues, from a naming perspective
//   Separate initialization from routines
//   rating => rateRatio0, etc. for similar name usage
//   matched to something more descriptive, as in "countsAreEqual"
//   Ethernet-duplex, Ethernet-pon, more descriptive names to be used

#include <assert.h>
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include "avbHead.h"

// *****
// Time formats used within this simulation are listed below.
// The layout is half scale; each '+' mark represents a byte boundary,
// not a bit-boundary (as is true in other narrow-format conventions).
// The high-level timings are based on largeTime and smallTime values.
//
//
//           largeTime
// +-----+-----+-----+-----+-----+-----+-----+-----+
// |           seconds           |           fraction           |
// +-----+-----+-----+-----+-----+-----+-----+
// Used for: Simulation time base :
// Features: Near-infinite resolution and range
//           :

```

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37


```

//
//           :
//           smallTime
//           +-----+-----+-----+-----+
//           |seconds|      fraction      |
//           +-----+-----+-----+-----+
// Used for: Station-local time base
// Features: Nearly-a-day range; near femto-second resolution
//
//           :
// *****:*****
// Time formats used within application-specific frames include a
// * grandTime - never-overflows globally-synchronized time
// * localTime - per-station locally-managed time
// * tinyTime - extra part of grandTime (a small value)
// * ticksTime - an application-specific variant of localTime
//
//           :
//           grandTime
//           +-----+-----+-----+-----+
//           |      seconds      |      fraction      |
//           +-----+-----+-----+-----+
// Used for: Frames, grandmaster time synchronization
// Features: Thousands-of-years range; pico-second resolution
//
//           :
//           localTime
//           +-----+-----+-----+-----+
//           |sec|      fraction      |
//           +-----+-----+-----+-----+
// Used for: Frames, cable-delay measurements
// Features: Minutes of range; pico-second resolution
//
//           :
//           tinyTime
//           + - +-----+-----+
//           : | subfraction |
//           + - +-----+-----+
//
// Used for: Sideband grandTime error accumulation
// Features: Minutes of range; pico-second resolution
//
//           :
//           ticksTime
//           +-----+-----+-----+-----+
//           |      tickCounts      |
//           +-----+-----+-----+-----+
// Used for: Exists on 802.3-EPON, 802.11v-Radio
// Features: Sufficient range; nano-second-like resolution
//
// *****
//
// *****
// ***** Assumed integer values *****
// *****
//
// typedef unsigned char      uint8_t;          // 1-byte unsigned integer
// typedef unsigned short     uint16_t;         // 2-byte unsigned integer
// typedef unsigned int       uint32_t;         // 4-byte unsigned integer
// typedef unsigned long long uint64_t;         // 8-byte unsigned integer

```

```

// typedef signed char      int8_t;           // 1-byte signed integer
// typedef signed short     int16_t;          // 2-byte signed integer
// typedef signed int       int32_t;          // 4-byte signed integer
// typedef signed long long int64_t;          // 8-byte signed integer

// *****
// ***** Revise timeSync frame parameters as the actual values become known *****
// *****

// Unique identifier values
#define AVB_MCAST 0 // Neighbor multicast address
#define AVB_PROTOCOL 0 // The protocolType for AVB
#define AVB_FUNCTION 0 // The timeSync function
#define AVB_VERSION 1 // The timeSync version

#define DEFAULT_DUPLEX_RX_RANGE 16

#define DEBUG 0

// Generic macro definitions
#define ARRAY_SIZE(x) (sizeof(x)/sizeof(x[0]))
#define BITS(type) (8 * sizeof(type))
#define CLIP_RATE(x, y) (((x) > SMALL_ONE + (y)) ? \
    (SMALL_ONE + (y)) : (((x) < (SMALL_ONE - (y)) ? SMALL_ONE - (y) : (x)))) // Clip within specified rate
#define CLIP_SIZE(x, y) ((x) > (y) ? (y) : ((x) < (-y)) ? (-y) : (x)) // Clip within specified value
#define CLOCK_MASTER_PORT_ID 255 // Clock-master port identifier
#define COUNT 256 // Number of frameCount values
#define EXTRA 16 // Queue-full error status
#define FULL 1 // Largest hop-count value
#define LAST_HOP 255
#define LARGE_10ms SmallToLarge(SMALL_10ms)
#define LARGE_HALF (ONE << 63)
#define LARGE_TOCK (ONE << 62)
#define MASK(bits) ((ONE << bits) - 1)
#define MASK32 (ONES >> 32)
#define MAX(a, b) ((a) > (b) ? (a) : (b)) // Maximum value definition
#define MIN(a, b) ((a) > (b) ? (b) : (a)) // Minimum value definition
#define MTU_SIZED 2048 // Maximum-sized transfer
#define OK 0 // Non-error status
#define ONE ((uint64_t)1) // Wide "1" constant
#define ONES (~uint64_t)0 // Wide "FF..FF" constant
#define PLUS(a, b, c) (((a) + (b) + (c)) % (c))
#define PON_TICK_TIME (DivideSmall(16 * (ONE << 32), 1000000000))
#define PPM100 ((SMALL_ONE * 100) / 1000000) // Scaled 100PPM fraction.
#define PPM250 ((SMALL_ONE * 250) / 1000000) // Scaled 250PPM fraction.
#define RADIO_TICK_TIME DivideSmall(1 << (32 - 9), 1000000000 >> 9) // Ratio radio-ns to localTime
#define RESIDENCE_DELAY ((SMALL_ONE + RandomMagOne()) / 800) // A 2.5ms max residence time
#define SMALL_10ms (SMALL_ONE / 100) // A 10ms smallTime interval
#define SMALL_ONE ((int64_t)(ONE << 48)) // Scaled fraction for 1.0
#define TESTING_OUI ((uint64_t)0Xcabled << 24)
#define TLIMIT 255
#define WIDE_MIN(a, b) (WideCompare((a), (b)) <= 0 ? (a) : (b))

```

```

#define CommonCheck(comPtr) (assert(comPtr != NULL), \
    assert(comPtr->rootLink != NULL), assert(comPtr->pairLink != NULL))

#define SetRxQueue1Ptrs(comPtr, ptr0) (CommonCheck(comPtr), \
    assert(comPtr->rxPortCount == 1), ptr0 = &(comPtr->rxPortPtr[0]) )
#define SetRxQueue2Ptrs(comPtr, ptr0, ptr1) (CommonCheck(comPtr), \
    assert(comPtr->rxPortCount == 2), \
    ptr0 = &(comPtr->rxPortPtr[0]), ptr1 = &(comPtr->rxPortPtr[1]) )
#define SetRxQueue3Ptrs(comPtr, ptr0, ptr1, ptr2) (CommonCheck(comPtr), \
    assert(comPtr->rxPortCount == 3), ptr0 = &(comPtr->rxPortPtr[0]), \
    ptr1 = &(comPtr->rxPortPtr[1]), ptr2 = &(comPtr->rxPortPtr[2]) )
#define SetTxQueue1Ptrs(comPtr, ptr0) (CommonCheck(comPtr), \
    assert(comPtr->txPortCount == 1), ptr0 = &(comPtr->txPortPtr[0]) )
#define SetTxQueue2Ptrs(comPtr, ptr0, ptr1) (CommonCheck(comPtr), \
    assert(comPtr->txPortCount == 2), \
    ptr0 = &(comPtr->txPortPtr[0]), ptr1 = &(comPtr->txPortPtr[1]) )

#define RxPortPtr(comPtr, count) (&(comPtr->rxPortPtr[count]))
#define StationTime(comPtr) (comPtr->smallTime)
#define SystemTime(comPtr) (comPtr->largeTime)
#define TxPortPtr(comPtr, count) (&(comPtr->txPortPtr[count]))

#define PrecedenceToEui64(a) (a.lower)

#define SizePlus(set) (sizeof(set) + EXTRA)

#define LargeToSmall(a) WideExtract(a, 16)
#define SmallToLarge(a) WideShift(SignedToWide(a), -16)
#define SmallToGrand(a) WideShift(SignedToWide(a), 8)
#define SmallToLocal(a) ((a) >> 8)
#define SmallAsLocal(a) ((a) & (ONES >> 8))
#define TinyToGrand(x) (SignedToWide(((int64_t)(x))))
#define TinyToSmall(a) (((int64_t)(a)) << 8)
#define TinyToLarge(a) (SignedToWide(((int64_t)(a)) << 24))

typedef enum {
    INTERPOLATE,
    EXTRAPOLATE
} GuessMode;

enum {
    FALSE,
    TRUE,
    WAIT
};

enum {
    BODY,
    LIST,
    BOTH
};

enum {
    TALK_QUIET,

```

```

TALK_GSYNC,
TALK_FRAME
};

enum {
VOCAL_QUIET,
VOCAL_DEBUG,
VOCAL_PAIRS
};

// Field extract/deposit definitions
#define FieldToSigned(fPtr, field) \
  FrameToValue((uint8_t *)&(fPtr->field), sizeof(fPtr->field), TRUE)           // Convert field to signed
#define FieldToUnsign(fPtr, field) \
  FrameToValue((uint8_t *)&(fPtr->field), sizeof(fPtr->field), FALSE)         // Convert field to unsigned
#define WideToFrame(value, fPtr, field) \
  ValueToFrame(value, (uint8_t *)&(fPtr->field), sizeof(fPtr->field))         // Convert field to unsigned
#define LongToFrame(value, fPtr, field) \
  ValueToFrame(SignedToWide(value), (uint8_t *)&(fPtr->field), sizeof(fPtr->field))

#define DeQueue(a, b) Dequeue(a, (uint8_t *)b, sizeof(*b))
#define EnQueue(a, b) Enqueue(a, (uint8_t *)b, sizeof(*b))

typedef struct {
    uint64_t upper;           // Double-precise integers
    uint64_t lower;         // More significant portion
} WideUnsigned;           // Less significant portion

#ifndef AVB_TIMES

#define NLIMIT 63

typedef struct {
    int64_t upper;           // Double-precise integers
    uint64_t lower;         // More significant portion
} WideSigned;           // Less significant portion

typedef int32_t    TicksTime;           // Link-dependent time
typedef int64_t    SmallTime;          // Local time reference
typedef WideSigned LargeTime;          // General 128-bit timers
#endif

typedef uint8_t    Boolean;            // True or false
typedef uint8_t    Port;              // Received port number
typedef uint8_t    Class;             // 1588: clock class
typedef uint8_t    HopCount;          // 1588: distance from GM
typedef uint8_t    Variance;          // 1588: clock error variance
typedef int32_t    TinyTime;           // Extra part of GM time
typedef int64_t    LocalTime;         // Compacted SmallTime
typedef WideSigned GrandTime;         // 1588: grandmaster time
typedef WideSigned Precedence;        // Fields {priorities, clockID}
typedef WideUnsigned Preference;      // Fields {precedence, hops, port}

```

```

// *****
// ***** Communication components *****
// *****

typedef struct {
    LargeTime    largeTime;           // Grandmaster synchronized
    SmallTime    extraTime;           // Extra part for largeTime
    SmallTime    smallTime;           // Station's free-running
    uint16_t     extraCount;           // Count of extra-values sampling
} BaseTimes;

typedef struct {
    GrandTime    grandTime;           // Time-result collection
    TinyTime     extraTime;           // Grandmaster synchronized
    GrandTime    totalTime;           // Side-band extra values
    GrandTime    totalTime;           // Precise grandTime+extraTime
} NextTimes;

// *****
// ***** Formal interface exchanges *****
// *****

typedef struct {
    uint8_t      frameCount[1];       // Sequential consistency check
    uint8_t      grandTime[10];       // Received grandmaster time
} ClockMasterSet;

typedef struct {
    uint8_t      frameCount[1];       // Sequential consistency check
} ClockSlaveReq;

typedef struct {
    uint8_t      frameCount[1];       // Sequential consistency check
    uint8_t      grandTime[10];       // Provided grandmaster time
} ClockSlaveRes;

typedef struct {
    uint8_t      protocolType[2];     // Time-sync frame parameters
    uint8_t      function[1];         // Protocol identifier
    uint8_t      version[1];          // Identifies timeSync frame
    uint8_t      precedence[14];      // Specific format identifier
    uint8_t      grandTime[10];       // Grandmaster precedence
    uint8_t      extraTime[4];        // Grandmaster time
    uint8_t      sourcePort[1];       // Extra part of grandTime
    uint8_t      hopCount[1];         // Transmit sequence number
    uint8_t      smallTime[8];        // GM hop-count distance
    uint8_t      syncInterval[6];     // Local-time reference
    uint8_t      syncInterval[6];     // Opposite-link transmit time
} SyncSduData;

typedef struct {
    uint8_t      destination_address[6]; // MS_UNITDATA.request
    uint8_t      source_address[6];     // Destination address
    uint8_t      priority[1];           // Source address
    SyncSduData  service_data_unit;     // Delivery priority
} Data content

```

```

} GrandSyncReq;
1

typedef struct {
2
    uint8_t destination_address[6]; // MS_UNITDATA.indication
3
    uint8_t source_address[6]; // Destination address
4
    uint8_t priority[1]; // Source address
5
    SyncSduData service_data_unit; // Delivery priority
6
} GrandSyncInd; // Data content

7
typedef struct { // Time-sync frame parameters
8
    uint8_t protocolType[2]; // Protocol identifier
9
    uint8_t function[1]; // Identifies timeSync frame
10
    uint8_t version[1]; // Specific format identifier
11
    uint8_t precedence[14]; // Grandmaster precedence
12
    uint8_t grandTime[10]; // Grandmaster time
13
    uint8_t extraTime[4]; // Extra part of grandTime
14
    uint8_t frameCount[1]; // Transmit sequence number
15
    uint8_t hopCount[1]; // GM hop-count distance
16
    uint8_t localTime[6]; // Transmitted timeSync time
17
    uint8_t thatTxTime[6]; // Opposite-link transmit time
18
    uint8_t thatRxTime[6]; // Opposite-link received time
19
} EfdxSduData;

20
typedef struct { // Time-sync frame parameters
21
    uint8_t da[6]; // Destination address
22
    uint8_t sa[6]; // Source address
23
    uint8_t protocolType[2]; // Protocol identifier
24
    uint8_t function[1]; // Identifies timeSync frame
25
    uint8_t version[1]; // Specific format identifier
26
    uint8_t precedence[14]; // Grandmaster precedence
27
    uint8_t grandTime[10]; // Grandmaster time
28
    uint8_t extraTime[4]; // Extra part of grandTime
29
    uint8_t frameCount[1]; // Transmit sequence number
30
    uint8_t hopCount[1]; // GM hop-count distance
31
    uint8_t localTime[6]; // Transmitted timeSync time
32
    uint8_t thatTxTime[6]; // Opposite-link transmit time
33
    uint8_t thatRxTime[6]; // Opposite-link received time
34
    uint8_t fcs[4]; // Opposite-link received time
35
} EfdxMacFrame;

36
typedef struct { // Time-sync frame parameters
37
    uint8_t destination_address[6]; // Destination address
38
    uint8_t source_address[6]; // Source address
39
    uint8_t priority[1]; // Delivery priority
40
    EfdxSduData service_data_unit; // Efdx service-data-unit
41
} EfdxMacInd;
42
typedef EfdxMacInd EfdxMacReq;

43
typedef struct {
44
    uint8_t frameCount[1]; // Sequential consistency check
45
    uint8_t smallTime[8]; // Common station-local time
46
} EfdxRxInfo;

47
typedef struct {
48
    uint8_t frameCount[1]; // Sequential consistency check
49
}

```

```

uint8_t  smallTime[8];           // Common station-local time
} EfdxTxInfo;

typedef struct {
    uint8_t  protocolType[2];    // Time-sync frame parameters
    uint8_t  function[1];        // Protocol identifier
    uint8_t  version[1];         // Identifies timeSync frame
    uint8_t  precedence[14];     // Specific format identifier
    uint8_t  grandTime[10];      // Grandmaster precedence
    uint8_t  extraTime[4];       // Grandmaster time
    uint8_t  frameCount[1];      // Extra part of grandTime
    uint8_t  hopCount[1];        // Transmit sequence number
    uint8_t  ticksTime[4];       // GM hop-count distance
    uint8_t  ticksTime[4];       // Local timing reference
} SyncSduEpon;

typedef struct {
    uint8_t  destination_address[6]; // Time-sync frame parameters
    uint8_t  source_address[6];      // Destination address
    uint8_t  priority[1];            // Source address
    SyncSduEpon  service_data_unit;  // Delivery priority
} EponMacInd;
typedef EponMacInd EponMacReq;

typedef struct {
    uint32_t  reserved;           // Reserved
} R11vInfo1Req;

typedef struct {
    uint32_t  ticksTime2;        // Received snapshot
    uint32_t  ticksTime3;        // Transmit snapshot
} R11vInfo1Ind;

typedef struct {
    uint32_t  ticksTime1;        // Transmit snapshot
    uint32_t  ticksTime4;        // Received snapshot
} R11vInfo1Con;

typedef struct {
    uint32_t  ticksTime4;        // Received snapshot
    uint32_t  roundTrip;         // Duration snapshot
    GrandTime levelTime;        // Grandmaster like
    TinyTime  extraTime;        // Extra part of levelTime
    Precedence precedence;      // Grandmaster error
    HopCount  hopCount;         // Grandmaster error
} R11vInfo2Req;

typedef R11vInfo2Req R11vInfo2Ind;

typedef struct {
    uint32_t  reserved;         // Reserved
} R11vInfo2Con;

// *****
// ***** Defined entities *****

```

```

// *****
enum {
    Q_RX00_LAST,
    Q_TX00_LAST = 0
};
enum {
    Q_RX01_BASE,
    Q_RX11_LAST,
    Q_TX01_BASE = 0,
    Q_TX11_LAST
};
enum {
    Q_RX02_BASE,
    Q_RX12_NEXT,
    Q_RX22_LAST,
    Q_TX02_BASE = 0,
    Q_TX12_NEXT,
    Q_TX22_LAST
};
enum {
    Q_RX03_BASE,
    Q_RX13_NEXT,
    Q_RX23_PLUS,
    Q_RX33_LAST,
    Q_TX03_BASE = 0,
    Q_TX13_NEXT,
    Q_TX23_NEXT,
    Q_TX33_LAST
};
// ***** GrandSync entity *****
typedef struct {
    Precedence    precedence;           // GrandSync entity state
    Port          sourcePort;          // Grandmaster precedence
    HopCount      hopCount;            // Source-port identifier
    SmallTime     syncInterval;        // Synchronization interval
} GrandSyncSaved;
typedef struct {
    Common        common;               // GrandSync entity state
    LocalTime     lastTime;            // Common simulation state
    GrandSyncSaved rxSaved;            // Timeout, set on activity
} GrandSyncEntity;
// ***** ClockMaster entity *****
typedef struct {
    Common        common;               // Client-clock master
    uint8_t       rxFrameCount;        // Common simulation state
}

```



```

Precedence      precedence;          // Grandmaster precedence          1
SmallTime       syncInterval;       // Synchronization interval        2
SmallTime       snapShot0;          // Recent snapshot                  3
SmallTime       snapShot1;          // Remote snapshot                  4
} ClockMasterEntity;                5

// ***** ClockSlave entity ***** 6
typedef struct {                      // Client-clock slave              7
    SmallTime    syncInterval;       8
} ClockSlaveSaved;                    9

typedef struct {                      // Client-clock slave             10
    Boolean      validated;           // Validated; operational         11
    uint16_t     headIndex;           // Recent interval index          12
    uint16_t     tailIndex;           // Oldest interval index          13
    SmallTime    interval;           // Rate-averaging interval        14
    BaseTimes    times[64];          // txTimes value array            15
} BaseTimer;                          16

typedef struct {                      // Client-clock slave             17
    Common       common;              // Common simulation info          18
    uint8_t      frameCount;          // Consistency-check count        19
    SmallTime    syncInterval;       // Synchronization interval       20
    SmallTime    snapShot0;          // Recent snapshot                 21
    SmallTime    snapShot1;          // Remote snapshot                 22
    ClockSlaveSaved rxSaved;         // Saved rx information            23
    BaseTimer    baseTimer;          // Receive-time history            24
} ClockSlaveEntity;                  25

// ***** duplex-Ethernet ***** 26
typedef struct {                      // EFDX receive                    27
    SmallTime    thisTime;           // Saved previous snapshot         28
    SmallTime    thatTime;           // Saved previous thisTxTime       29
} PastTimes;                          30

typedef struct {                      // EFDX receive                    31
    Boolean      validated;           // Validity indication             32
    uint16_t     headIndex;           // Recent interval index          33
    uint16_t     tailIndex;           // Oldest interval index          34
    SmallTime    interval;           // Rate-averaging interval        35
    PastTimes    times[64];          // Larger than ever needed        36
} RxTimer;                             37

typedef struct {                      // EFDX receive                    38
    Common       common;              // Common simulation info          39
    Boolean      txReady;             // Cable-delay valid               40
    LocalTime    syncInterval;        // Clock-master's tockTime        41
    uint16_t     snapCount;           // The indication's frameCount    42
    uint16_t     frameCount;          // The timeSync's frameCount      43
    SmallTime    snapShot0;           // This frame's arrival time      44
    SmallTime    snapShot1;           // Past frame's arrival time      45
}

```

```

LocalTime      thisTxTime;          // Frame transmission time          1
LocalTime      thisRxTime;          // Frame reception time             2
LocalTime      thatRxTime;          // Frame reception time             3
EfdxMacInd     savedInd;            // Received timeSync indication     4
RxTimer        rxTimer;             // For computing rateRatio (new)    5
} EfdxRxEntity;
                                                       6
typedef struct {
    uint64_t     da;                  // EFDX transmit                     7
    uint64_t     sa;                  // destination_address               8
    uint16_t     type;                // source_address                     9
    uint8_t      function;            // Received protocolType             10
    uint8_t      version;             // Specified function (AVB)         11
    WideSigned   precedence;          // Version number in AVB            12
    HopCount     hopCount;            // Grandmaster preference           13
    SmallTime    syncInterval;        // Grandmaster distance             14
} EfdxTxSaved;
                                                       15
typedef struct {
    Common        common;              // EFDX transmit                     16
    GuessMode     guessMode;          // Common simulation info           17
    Boolean        rxReady;            // Estimating next value            18
    Boolean        txReady;            // Sinking rx initialized           19
    SmallTime     execTime;            // Sending tx initialized            20
    SmallTime     lastTime;            // Next transmission time           21
    uint8_t        frameCount;         // Periodic transmission time       22
    SmallTime     syncInterval;        // The timeSync frame count        23
    EfdxTxSaved   rxSaved;             // Sync interval duration           24
    uint8_t        sxSnapCount;        // Received GrandSync request       25
    SmallTime     sxSnapTimed;        // Received MAC snapshot            26
    BaseTimer     baseTimer;           // Received MAC snapshot            27
} EfdxTxEntity;
                                                       28
// ***** 802.11v radio *****
typedef struct {
    Common        common;              // 802.11v wireless receive        29
    LocalTime     syncInterval;        // Common simulation info           30
    TicksTime     turnRound;           // Clock-master's tockTime          31
} R11vRxEntity;
                                                       32
typedef struct {
    WideSigned   precedence;          // 802.11v wireless transmit       33
    HopCount     hopCount;            // Grandmaster preference           34
    SmallTime     syncInterval;        // hopCount                          35
} R11vTxSaved;
                                                       36
typedef struct {
    Common        common;              // 802.11v wireless transmit       37
    LocalTime     syncInterval;        // Common simulation info           38
    LocalTime     pastTime;            // Clock-master's tockTime          39
    LocalTime     lastTime;            // Back-interpolation time          40
    R11vTxSaved   rxSaved;            // Last transmission                41
    BaseTimer     baseTimer;           // Saved parameters                 42
} R11vTxEntity;
                                                       43

```

```

    TicksTime    snapShot1;                // Saved ticksTime1
    TicksTime    roundTrip;                // Saved ticksTime4-ticksTime1
    TicksTime    rxTurnRound;              // Turn-round delay times
    TicksTime    snapShot4;                // Saved ticksTime4
    TicksTime    rxRoundTrip;              // Saved ticksTime4-ticksTime1
    Boolean      respondNow;
} R11vTxEntity;

// ***** Ethernet-PON entity *****

typedef struct {                          // Ethernet-PON receive
    Common      common;                    // Common simulation info
    LocalTime   syncInterval;              // Clock-master's tockTime
} EponRxEntity;

typedef struct {                          // Ethernet-PON transmit
    uint64_t    da;                        // destination_address
    uint64_t    sa;                        // source_address
    Precedence  precedence;                // Grandmaster precedence
    HopCount    hopCount;                  // hopCount
    SmallTime   syncInterval;
} EponTxSaved;

typedef struct {                          // Ethernet-PON transmit
    Common      common;                    // Common simulation info
    LocalTime   syncInterval;              // Clock-master's tockTime
    SmallTime   lastTime;                  // Last time checkpoint
    EponTxSaved rxSaved;
    BaseTimer   baseTimer;                 // Receive-time history
} EponTxEntity;

// External control parameters
GuessMode argGuessMode = INTERPOLATE;
SmallTime argMegaHertz = 0;
uint16_t  argBridgeCount = 10;
uint32_t  argFirstSecs = 0;
uint32_t  argFinalSecs = 150;
uint32_t  argVocalType = VOCAL_DEBUG;

// Standard state-machine routines
Common    *GrandSyncExec(Common *, char *);
Common    *ClockMasterExec(Common *, char *);
Common    *ClockSlaveExec(Common *, char *);
Common    *EfdxRxExec(Common *, char *);
Common    *EfdxTxExec(Common *, char *);
Common    *EfdxSpanExec(Common *, char *);
Common    *R11vRxExec(Common *, char *);
Common    *R11vTxExec(Common *, char *);
Common    *EponRxExec(Common *, char *);
Common    *EponTxExec(Common *, char *);

```

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37

```

void          CommonChecks(Common *, uint16_t, uint16_t);
Boolean      PortsCheck(Common *, char *);
Common       *CommonCreate(Common *, uint16_t,
Common        *(Common *, char *), uint16_t, uint16_t, uint16_t);
Port         PortID(Common *);
TicksTime   EponTime(EponRxEntity *);
TicksTime   R1lvTime(R1lvRxEntity *);

SmallTime   NextRate(RxTimer *, SmallTime, SmallTime, SmallTime, SmallTime);
void        NextSaved(BaseTimer *, SmallTime, SmallTime, LargeTime, SmallTime, SmallTime);
NextTimes   NextTimed(BaseTimer *, SmallTime, SmallTime);

// A minimalist double-width integer library
int         WideCompare(WideSigned, WideSigned);
int         WideCompareUnsigned(WideUnsigned, WideUnsigned);
int64_t     WideExtract(WideSigned, uint8_t);
WideSigned  WideAddition(WideSigned, WideSigned);
WideSigned  WideSubtract(WideSigned, WideSigned);
WideSigned  WideShift(WideSigned, int16_t);
WideSigned  WideNegate(WideSigned);
WideSigned  WideMultiply(int64_t, int64_t);

int64_t     DivideSmall(int64_t, int64_t);
int64_t     MultiplySmall(int64_t, int64_t);
SmallTime   ClockTicks(int64_t, uint64_t);

double      LargeToFloat(LargeTime);
GrandTime   LargeToGrand(LargeTime);
double      LocalToFloat(LocalTime);
SmallTime   LocalToSmall(LocalTime);
double      GrandToFloat(GrandTime);
LargeTime   GrandToLarge(GrandTime);
TinyTime    SmallToTiny(SmallTime);
double      SmallToFloat(SmallTime);
double      TinyToFloat(TinyTime);

// Other routines
uint64_t     Eui48ToEui64(uint64_t);
Precedence  FieldsToPrecedence(uint8_t, uint8_t, uint16_t, uint8_t, uint64_t);
WideSigned  FrameToValue(uint8_t *, uint16_t, Boolean);
Preference  FormPreference(WideSigned, uint8_t, uint8_t);
GrandTime   GrandToLevel(GrandTime);
GrandTime   LevelToGrand(GrandTime);
uint64_t     MacAddress(Common *);
SmallTime   RandomMagOne();
WideSigned  SignedToWide(int64_t);
Boolean     TimeSyncSdu(SyncSduData *);
void        ValueToFrame(WideSigned, uint8_t *, uint16_t);

Entry       *Bequeue(RxPort *);
uint32_t    Dequeue(RxPort *, void *, unsigned);
Boolean     Enqueue(TxPort *, void *, unsigned);
void        SleepOnRoot(Common *, LargeTime);
void        SleepOnBase(Common *, LocalTime);

```

```

char          *StrPair(char *, char *, char *, uint16_t);

// *****
// ***** GrandSync state-machine routine *****
// *****

Common *GrandSyncInit(Common *, char *);

// Sets common state to allow grandTime values to be back-interpolated
// arguments:
//   comPtr - associated state-maintaining data structure
//   name - initialization trigger and assigned entity name
Common *
GrandSyncExec(Common *comPtr, char *name) {
    uint8_t rxInfo[SizePlus(GrandSyncInd)], txInfo[SizePlus(GrandSyncReq)];
    GrandSyncEntity *ePtr = (GrandSyncEntity *)comPtr;
    GrandSyncInd *rxPtr = (GrandSyncInd *)rxInfo;
    GrandSyncReq *txPtr = (GrandSyncReq *)txInfo;
    GrandSyncSaved *sxPtr = &(ePtr->rxSaved);
    SyncSduData *rsPtr = &(rxPtr->service_data_unit);
    SyncSduData *tsPtr = &(txPtr->service_data_unit);
    RxPort *rxQueuePtr;
    TxPort *txQueuePtr;
    GrandTime rxGrandTime;
    Preference test, past;
    Precedence rxPrecedence;
    SmallTime nextTime, rxSmallTime, rxSyncInterval, stationTime;
    TinyTime rxExtraTime;
    HopCount rxHopCount;
    Port rxSourcePort;
    uint64_t macAddress;
    uint32_t sized;
    uint16_t accelerated, count, hopCount;
    Boolean serviced, better;

    if (name != NULL) // Initialization
        return(GrandSyncInit(comPtr, name));
    SetRxQueuePtrs(comPtr, rxQueuePtr);
    SetTxQueuePtrs(comPtr, txQueuePtr);

    do {
        serviced = FALSE;
        // *****
        // ***** Processing arriving GrandSync MA_UNITDATA.indication parameters *****
        // *****
        stationTime = StationTime(comPtr); // Station's localTime
        sized = DeQueue(rxQueuePtr, &rxInfo); // Check rx queue
        if (sized == sizeof(GrandSyncInd) && TimeSyncSdu((SyncSduData *)rsPtr)) { // Verify parameters
            serviced = TRUE;
            rxPrecedence = FieldToUnsign(rsPtr, precedence); // Precedence value
            rxHopCount = FieldToUnsign(rsPtr, hopCount).lower; // Hop-count distance
            rxSourcePort = FieldToUnsign(rsPtr, sourcePort).lower; // Received port identifier
            rxGrandTime = FieldToSigned(rsPtr, grandTime); // Grandmaster time
            rxExtraTime = FieldToSigned(rsPtr, extraTime).lower; // Extra part of grandTime
        }
    } while (!serviced);
}

```

```

rxSmallTime = FieldToSigned(rsPtr, smallTime).lower; // Internal station time 1
rxSyncInterval = FieldToUnsign(rsPtr, syncInterval).lower; // Sync-interval time 2

test = FormPreference(rxPrecedence, rxHopCount, rxSourcePort); // Test preference 3
past = FormPreference(sxPtr->precedence, sxPtr->hopCount, sxPtr->sourcePort); // Past preference 4
better = rxSourcePort == sxPtr->sourcePort || WideCompareUnsigned(test, past) <= 0; // This one is better 5
if (rxHopCount != LAST_HOP && better) { // and is also valid 6
    ePtr->lastTime = stationTime; // Update timeout timer 7
    count = sxPtr->hopCount; // Saved hopCount value 8
    sxPtr->precedence = rxPrecedence; // Saved GM values 9
    sxPtr->sourcePort = rxSourcePort; // ... 10
    sxPtr->hopCount = rxHopCount; // ... 11
    sxPtr->syncInterval = rxSyncInterval; // ... 12

    accelerated = 1 + (LAST_HOP + rxHopCount) / 2; // Accelerated aging 13
    hopCount = MIN(LAST_HOP, rxHopCount > count ? accelerated : rxHopCount + 1); // if receiver is aged 14

    // *****
    // ***** Create and transmit MA_UNITDATA.indication parameters *****
    // *****
    macAddress = MacAddress(comPtr);
    LongToFrame(AVB_MCAST, txPtr, destination_address); // Neighbor multicast address 15
    LongToFrame(macAddress, txPtr, source_address); // This port's MAC address 16
    LongToFrame(AVB_PROTOCOL, tsPtr, protocolType); // The AVB protocol 17
    LongToFrame(AVB_FUNCTION, tsPtr, function); // The timeSync frame in AVB 18
    LongToFrame(AVB_VERSION, tsPtr, version); // This version number 19
    WideToFrame(rxPrecedence, tsPtr, precedence); // Create new precedende 20
    LongToFrame(hopCount, tsPtr, hopCount); // Create GM distance 21
    LongToFrame(rxSourcePort, tsPtr, sourcePort); // Create port indentifier 22
    WideToFrame(rxGrandTime, tsPtr, grandTime); // Echo grandTime 23
    LongToFrame(rxExtraTime, tsPtr, extraTime); // Echo extraTime 24
    LongToFrame(rxSmallTime, tsPtr, smallTime); // Echo localTime 25
    LongToFrame(rxSyncInterval, tsPtr, syncInterval); // Echo synch interval 26
    EnQueue(txQueuePtr, txPtr); // Enqueue the result; 27
}
} 28
} 29

assert(sxPtr->syncInterval != 0); // Consistency check 30
nextTime = ePtr->lastTime + 4 * sxPtr->syncInterval; // Timeout threshold 31
if ((stationTime - nextTime) >= 0) { // Timeout actions 32
    serviced = TRUE;
    sxPtr->precedence.upper = sxPtr->precedence.lower = ONES; // Worst precedence & 33
    sxPtr->hopCount = sxPtr->sourcePort = 0xFF; // worst tie-breakers 34
    ePtr->lastTime = stationTime; // Resets the timeout 35
    // printf("GrandSync (timeout):\n"); 36
} 37
} while (serviced == TRUE); 38
assert((nextTime - stationTime) > 0); 39
SleepOnBase(comPtr, nextTime); 40
return(NULL); 41
} 42

Common *
GrandSyncInit(Common *oldPtr, char *string) { 43
    Common *comPtr; 44
} 45

```

```

GrandSyncEntity *ePtr;
GrandSyncSaved *sxPtr;
char temp[TLIMIT+1], name[NLIMIT+1], data[NLIMIT+1], *nextPtr;

comPtr = CommonCreate(oldPtr, sizeof(GrandSyncEntity), &GrandSyncExec, TYPE_ENTITY, Q_RX11_LAST, Q_TX11_LAST);
if (comPtr != NULL) {
    ePtr = (GrandSyncEntity *)comPtr;
    sxPtr = &(ePtr->rxSaved);
    sxPtr->syncInterval = SMALL_10ms;
    sxPtr->precedence.upper = sxPtr->precedence.lower = ONES;

    nextPtr = StrPair(strncpy(temp, string, TLIMIT), name, data, NLIMIT);
    for ( ; nextPtr != NULL && name[0] != '\0'; nextPtr = StrPair(nextPtr, name, data, NLIMIT)) {
        if (strcmp(name, "name") == 0)
            strcpy(comPtr->name, data);
        else if (strcmp(name, "rx0") == 0)
            assert(comPtr->rxPortCount > 0), strcpy(comPtr->rxPortPtr[0].name, data);
        else if (strcmp(name, "tx0") == 0)
            assert(comPtr->txPortCount > 0), strcpy(comPtr->txPortPtr[0].name, data);
    }
    assert(PortsCheck(comPtr, nextPtr) == TRUE);
}
return(comPtr);
}

// *****
// ***** ClockMaster state-machine routines *****
// *****

Common *ClockMasterInit(Common *, char *);
// Provide the clock source information, retransmits to GrandSync
// arguments:
//   comPtr - associated state-maintaining data structure
//   name - initialization trigger and assigned entity name
Common *
ClockMasterExec(Common *comPtr, char *name) {
    uint8_t cmInfo[SizePlus(ClockMasterSet)], txInfo[SizePlus(GrandSyncInd)];
    ClockMasterEntity *ePtr = (ClockMasterEntity *)comPtr;
    ClockMasterSet *rxPtr = (ClockMasterSet *)cmInfo;
    GrandSyncInd *txPtr = (GrandSyncInd *)txInfo;
    SyncSduData *tsPtr = &(txPtr->service_data_unit);
    RxPort *rxQueuePtr;
    TxPort *txQueuePtr;
    GrandTime grandTime;
    SmallTime stationTime;
    uint64_t macAddress;
    uint32_t sized;
    uint16_t count, frameCount;
    uint8_t portID;
    Boolean serviced;

    if (name != NULL)
        return(ClockMasterInit(comPtr, name));
    SetRxQueuePtrs(comPtr, rxQueuePtr);
    // The entity name
    // for initialization

```

```

SetTxQueuePtrs(comPtr, txQueuePtr);

macAddress = MacAddress(comPtr);
portID = PortID(comPtr);

do {
    serviced = FALSE;
    // *****
    // **** Processing arriving GrandSync clock-master indication parameters ****
    // *****
    stationTime = StationTime(comPtr);
    sized = DeQueue(rxQueuePtr, rxPtr);
    if (sized > 0) {
        serviced = TRUE;
        assert(sized == sizeof(ClockMasterSet));
        ePtr->snapShot1 = ePtr->snapShot0;
        ePtr->snapShot0 = stationTime;
        count = (ePtr->rxFrameCount + 1) % COUNT;
        frameCount = FieldToUnsign(rxPtr, frameCount).lower;
        grandTime = FieldToSigned(rxPtr, grandTime);
        ePtr->rxFrameCount = frameCount;
        if (count == frameCount) {

            // *****
            // ***** Creation and transmit of MA_UNITDATA.indication parameters *****
            // *****

            LongToFrame(AVB_MCAST,          txPtr, destination_address); // Neighbor multicast address
            LongToFrame(macAddress,        txPtr, source_address); // This port's MAC address
            LongToFrame(AVB_PROTOCOL,     tsPtr, protocolType); // The AVB protocol
            LongToFrame(AVB_FUNCTION,     tsPtr, function); // The timeSync frame in AVB
            LongToFrame(AVB_VERSION,     tsPtr, version); // This version number
            WideToFrame(ePtr->precedence,  tsPtr, precedence); // Create new precedence
            LongToFrame(0,                 tsPtr, hopCount); // Initial GM distance
            LongToFrame(portID,           tsPtr, sourcePort); // Create port identifier
            WideToFrame(grandTime,        tsPtr, grandTime); // Report grandTime
            LongToFrame(ePtr->snapShot1,   tsPtr, smallTime); // Report smallTime
            LongToFrame(0,                 tsPtr, extraTime); // Initial extraTime
            LongToFrame(ePtr->syncInterval, tsPtr, syncInterval); // Sync-frame interval
            // TBD: fcs, priority, ...
            EnQueue(txQueuePtr, txPtr); // Enqueue the result;
        }
    }
} while (serviced == TRUE);
return(NULL);
}

Common *
ClockMasterInit(Common *oldPtr, char *string) {
    Common *comPtr;
    ClockMasterEntity *ePtr;
    uint64_t macAddress;
    char temp[TLIMIT+1], name[NLIMIT+1], data[NLIMIT+1], *nextPtr;

    comPtr = CommonCreate(oldPtr, sizeof(ClockMasterEntity), &ClockMasterExec, TYPE_RX_PORT, Q_RX11_LAST, Q_TX11_LAST);

```



```

    if (comPtr != NULL) {
        macAddress = MacAddress(comPtr);
        comPtr->portID = CLOCK_MASTER_PORT_ID;
        comPtr->portLink = comPtr;
        ePtr = (ClockMasterEntity *)comPtr;
        ePtr->syncInterval = SMALL_10ms;
        ePtr->precedence.upper = 0;
        ePtr->precedence.lower = Eui48ToEui64(macAddress);

        nextPtr = StrPair(strncpy(temp, string, TLIMIT), name, data, NLIMIT);
        for ( ; name[0] != '\0'; nextPtr = StrPair(nextPtr, name, data, NLIMIT)) {
            if (strcmp(name, "name") == 0)
                strncpy(comPtr->name, data, NLIMIT);
            else if (strcmp(name, "rx0") == 0)
                assert(comPtr->rxPortCount >= 1), strcpy(comPtr->rxPortPtr[0].name, data);
            else if (strcmp(name, "tx0") == 0)
                assert(comPtr->txPortCount >= 1), strcpy(comPtr->txPortPtr[0].name, data);
        }
        assert(PortsCheck(comPtr, nextPtr) == TRUE);
    }
    return(comPtr);
}

// *****
// ***** ClockSlave state-machine routines *****
// *****

Common *ClockSlaveInit(Common *, char *);
// Generates a GrandSync indication, after being triggered
// arguments:
//   comPtr - associated state-maintaining data structure
//   name - initialization trigger and assigned entity name
Common *
ClockSlaveExec(Common *comPtr, char *name) {
    uint8_t sxInfo[SizePlus(ClockSlaveReq)], rxInfo[SizePlus(GrandSyncReq)], txInfo[SizePlus(ClockSlaveRes)];
    ClockSlaveEntity *ePtr = (ClockSlaveEntity *)comPtr;
    ClockSlaveSaved *rcPtr = &(ePtr->rxSaved);
    BaseTimer *btPtr = &(ePtr->baseTimer);
    ClockSlaveReq *sxPtr = (ClockSlaveReq *)sxInfo;
    GrandSyncReq *rxPtr = (GrandSyncReq *)rxInfo;
    ClockSlaveRes *txPtr = (ClockSlaveRes *)txInfo;
    SyncSduData *rsPtr = &(rxPtr->service_data_unit);
    RxPort *cxPortPtr, *rxQueuePtr;
    TxPort *txQueuePtr;
    NextTimes nextTimes;
    LargeTime systemTime;
    GrandTime grandTime;
    SmallTime backInterval, rateInterval, smallTime, stationTime, syncInterval;
    TinyTime extraTime;
    uint32_t sized;
    uint8_t frameCount;
    Boolean serviced;

    if (name != NULL)

```

// The entity's name

```

return(ClockSlaveInit(comPtr, name)); // for initialization
1
SetRxQueue2Ptrs(comPtr, rxQueuePtr, cxPortPtr);
2
SetTxQueue1Ptrs(comPtr, txQueuePtr);
3
stationTime = StationTime(comPtr);
4
do {
5
    serviced = FALSE;
6
    systemTime = SystemTime(comPtr);
7
    backInterval = (3 * rcPtr->syncInterval + ePtr->syncInterval) / 2;
8
    rateInterval = backInterval + (3 * ePtr->syncInterval) / 2;
9
    // *****
10
    // ***** Processing arriving GrandSync ClockSlave.request parameters *****
11
    // *****
12
    sized = DeQueue(cxPortPtr, sxPtr);
13
    if (sized != 0) {
14
        serviced = TRUE;
15
        assert(sized == sizeof(ClockSlaveReq));
16
        frameCount = FieldToUnsign(sxPtr, frameCount).lower;
17
        nextTimes = NextTimed(btPtr, stationTime, backInterval);
18
        grandTime = nextTimes.totalTime;
19
        LongToFrame(frameCount, txPtr, frameCount); // Tag from the request
20
        WideToFrame(grandTime, txPtr, grandTime); // Associated grandTime
21
        EnQueue(txQueuePtr, txPtr); // Enqueue the result;
22
    }
23
    // *****
24
    // ***** Processing arriving GrandSync MA_UNITDATA.request parameters *****
25
    // *****
26
    sized = DeQueue(rxQueuePtr, rxPtr);
27
    if (sized == sizeof(GrandSyncReq) && TimeSyncSdu((SyncSduData *)rsPtr)) {
28
        serviced = TRUE;
29
        syncInterval = FieldToSigned(rsPtr, syncInterval).lower;
30
        grandTime = FieldToSigned(rsPtr, grandTime);
31
        extraTime = FieldToSigned(rsPtr, extraTime).lower;
32
        smallTime = FieldToSigned(rsPtr, smallTime).lower;
33
        rcPtr->syncInterval = syncInterval;
34
        NextSaved(btPtr, ePtr->syncInterval, rateInterval, GrandToLarge(grandTime), TinyToSmall(extraTime), smallTime);
35
    }
36
} while (serviced == TRUE); // Return tagged indication
return(NULL);
}
}
Common *
ClockSlaveInit(Common *oldPtr, char *string) {
Common *comPtr;
ClockSlaveEntity *ePtr;
BaseTimer *btPtr;
char temp[TLIMIT+1], name[NLIMIT+1], data[NLIMIT+1], *nextPtr;

comPtr = CommonCreate(oldPtr, sizeof(ClockSlaveEntity), &ClockSlaveExec, TYPE_TX_PORT, Q_RX22_LAST, Q_TX11_LAST);
if (comPtr != NULL) {
ePtr = (ClockSlaveEntity *)comPtr;
ePtr->syncInterval = SMALL_10ms;
ePtr->rxSaved.syncInterval = SMALL_10ms;
}
}

```

```

    btPtr = &(ePtr->baseTimer);
    btPtr->validated = 0;
    btPtr->headIndex = btPtr->tailIndex = 0;

    nextPtr = StrPair(strncpy(temp, string, TLIMIT), name, data, NLIMIT);
    for ( ; name[0] != '\0'; nextPtr = StrPair(nextPtr, name, data, NLIMIT)) {
        if (strcmp(name, "name") == 0)
            strcpy(comPtr->name, data, NLIMIT);
        else if (strcmp(name, "rx0") == 0)
            assert(comPtr->rxPortCount > 0), strcpy(comPtr->rxPortPtr[0].name, data);
        else if (strcmp(name, "rx1") == 0)
            assert(comPtr->rxPortCount > 1), strcpy(comPtr->rxPortPtr[1].name, data);
        else if (strcmp(name, "tx0") == 0)
            assert(comPtr->txPortCount >= 1), strcpy(comPtr->txPortPtr[0].name, data);
    }
    assert(PortsCheck(comPtr, nextPtr) == TRUE);
}
return(comPtr);
}

// *****
// ***** Ethernet-duplex state-machine routines *****
// *****

Common *EfdxRxInit(Common *, char *);
// Receives duplex-Ethernet SDUs, retransmits them to GrandSync
// arguments:
//   comPtr - associated state-maintaining data structure
//   name - initialization trigger and assigned entity name
Common *
EfdxRxExec(Common *comPtr, char *name) {
    uint8_t    rxInfo[MTU_SIZED], txInfo[MTU_SIZED], sxInfo[SizePlus(EfdxRxInfo)];
    EfdxRxEntity *ePtr = (EfdxRxEntity *)comPtr;
    RxTimer    *btPtr = &(ePtr->rxTimer);
    EfdxRxInfo *sxPtr = (EfdxRxInfo *)sxInfo;
    EfdxMacInd *rxPtr = (EfdxMacInd *)rxInfo;
    GrandSyncInd *txPtr = (GrandSyncInd *)txInfo;
    EfdxSduData *rsPtr = &(rxPtr->service_data_unit);
    SyncSduData *tsPtr = &(txPtr->service_data_unit);
    EfdxMacInd *dxPtr = &(ePtr->savedInd);
    EfdxSduData *dsPtr = &(dxPtr->service_data_unit);
    RxPort    *rxQueuePtr, *sxQueuePtr;
    TxPort    *txQueuePtr;
    Precedence precedence;
    GrandTime grandTime;
    SmallTime cableDelay, compRxTime, rateRatio, smallTime, stationTime, thisRxTime, thisTxTime;
    LocalTime roundTrip, thatTxTime, thatRxTime, thisTxTime, turnRound, turnRound0;
    TinyTime  extraTime;
    uint64_t  da, sa;
    uint32_t  sized;
    uint16_t  frameCount, protocol, rxHeadSize, snapCount, txHeadSize;
    uint8_t   function, guess, hopCount, portID, version;
    Boolean    countsAreEqual, serviced;

```

```

if (name != NULL)                                     // The entity's name
    return(EfdxRxInit(comPtr, name));                 // for initialization.
SetRxQueue2Ptrs(comPtr, rxQueuePtr, sxQueuePtr);
SetTxQueue1Ptrs(comPtr, txQueuePtr);

portID = PortID(comPtr);
countsAreEqual = 0;
do {
    stationTime = StationTime(comPtr);                // Station's localTime
    serviced = FALSE;

    // ***** Processing arriving timeSync snapshots *****
    // ***** Processing arriving timeSync frames *****
    sized = DeQueue(sxQueuePtr, &sxInfo);
    if (sized != 0) {
        serviced = TRUE;
        assert(sized == sizeof(EfdxRxInfo));
        snapCount = FieldToUnsign(sxPtr, frameCount).lower;
        smallTime = FieldToUnsign(sxPtr, smallTime).lower;
        ePtr->snapShot1 = ePtr->snapShot0;
        ePtr->snapShot0 = smallTime;
        ePtr->snapCount = snapCount;
        countsAreEqual = (snapCount == ePtr->frameCount);
    }

    // ***** Processing arriving timeSync frames *****
    sized = Dequeue(rxQueuePtr, rxPtr, sizeof(rxInfo));
    if (sized != 0) {
        serviced = TRUE;
        if (sized == sizeof(EfdxMacInd) && TimeSyncSdu((SyncSduData *)rsPtr)) {
            frameCount = FieldToUnsign(rsPtr, frameCount).lower;
            guess = (ePtr->frameCount + 1) % COUNT;
            ePtr->frameCount = frameCount;
            if (frameCount == guess) {
                bcopy(rxPtr, dxPtr, sizeof(EfdxMacInd));
                countsAreEqual = (frameCount == ePtr->snapCount);
            }
        } else {
            da = FieldToUnsign(rxPtr, destination_address).lower;
            sa = FieldToUnsign(rxPtr, source_address).lower;
            LongToFrame(da, txPtr, destination_address);
            LongToFrame(sa, txPtr, source_address);
            rxHeadSize = (void *)rsPtr - (void *)rxPtr;
            txHeadSize = (void *)tsPtr - (void *)txPtr;
            bcopy(rsPtr, tsPtr, sized - rxHeadSize);
            Enqueue(txQueuePtr, txPtr, sized + txHeadSize - rxHeadSize);
        }
    }
}

if (countsAreEqual == TRUE) {
    serviced = TRUE;
}

```

```

countsAreEqual = FALSE;
da = FieldToUnsign(dxPtr, destination_address).lower;
sa = FieldToUnsign(dxPtr, source_address).lower;
protocol = FieldToUnsign(dsPtr, protocolType).lower;
function = FieldToUnsign(dsPtr, function).lower;
version = FieldToUnsign(dsPtr, version).lower;
precedence = FieldToUnsign(dsPtr, precedence);
hopCount = FieldToUnsign(dsPtr, hopCount).lower;
frameCount = FieldToUnsign(dsPtr, frameCount).lower;
thatTxTime = FieldToSigned(dsPtr, thatTxTime).lower;
thatRxTime = FieldToSigned(dsPtr, thatRxTime).lower;
grandTime = FieldToSigned(dsPtr, grandTime);
extraTime = FieldToSigned(dsPtr, extraTime).lower;
thisTxTime = FieldToSigned(dsPtr, localTime).lower;
thisRxTime = ePtr->snapShot1;
ePtr->thisTxTime = thisTxTime; // Saved for transmit
ePtr->thisRxTime = SmallToLocal(thisRxTime); // over returning link

if (ePtr->thatRxTime != thatRxTime) {
    ePtr->thatRxTime = thatRxTime;
    ePtr->txReady = TRUE;
}
thisTxTimed = thisRxTime - LocalToSmall(SmallToLocal(thisRxTime) - thisTxTime);
if (ePtr->txReady == TRUE)
    rateRatio = NextRate(btPtr, thisRxTime, thisTxTimed, ePtr->syncInterval, btPtr->interval);
else
    rateRatio = SMALL_ONE;

if (btPtr->validated != TRUE || ePtr->txReady != TRUE)
    hopCount = LAST_HOP;
roundTrip = LocalToSmall(ePtr->thisRxTime - thatTxTime); // Round-trip delay
turnRound0 = LocalToSmall(thisTxTime - thatRxTime); // Turn-around delay
turnRound = MultiplySmall(turnRound0, rateRatio); // Normalized turnRound
cableDelay = MAX(0, roundTrip - turnRound) / 2; // Cable-delay computed
compRxTime = thisRxTime - cableDelay; // Cable-delay adjustment

// *****
// ***** Update revised service-data-unit parameters *****
// *****

LongToFrame(da, txPtr, destination_address); // The destination address
LongToFrame(sa, txPtr, source_address); // The source address
LongToFrame(protocol, tsPtr, protocolType); // The protocol identifier
LongToFrame(function, tsPtr, function); // The function identifier
LongToFrame(version, tsPtr, version); // The function identifier
WideToFrame(precedence, tsPtr, precedence); // GM selection precedence
LongToFrame(hopCount, tsPtr, hopCount); // GM hop-count distance
LongToFrame(portID, tsPtr, sourcePort); // Source-port identifier
WideToFrame(grandTime, tsPtr, grandTime); // grandTime at snapShot
LongToFrame(extraTime, tsPtr, extraTime); // Next extraTime value
LongToFrame(compRxTime, tsPtr, smallTime); // Transmitted frame time
LongToFrame(ePtr->syncInterval, tsPtr, syncInterval); // Sync transmit interval
EnQueue(txQueuePtr, txPtr); // Enqueue the result
}
} while (serviced == TRUE);
return(NULL);

```

```

}
1
Common *
2
EfdxRxInit(Common *oldPtr, char *string) {
3
    Common *comPtr;
4
    EfdxRxEntity *ePtr;
5
    RxTimer *btPtr;
6
    char temp [TLIMIT+1], name [NLIMIT+1], data [NLIMIT+1], *nextPtr;
7
    comPtr = CommonCreate(oldPtr, sizeof(EfdxRxEntity), &EfdxRxExec, TYPE_RX_PORT, Q_RX22_LAST, Q_TX11_LAST);
8
    if (comPtr != NULL) {
9
        ePtr = (EfdxRxEntity *)comPtr;
10
        ePtr->syncInterval = SMALL_10ms;
11
        ePtr->snapCount = ePtr->frameCount - 1;
12
        ePtr->txReady = FALSE;
13
14
        btPtr = &(ePtr->rxTimer);
15
        btPtr->validated = FALSE;
16
        btPtr->interval = SMALL_10ms * 20;
17
        btPtr->headIndex = btPtr->tailIndex = 0;
18
19
        nextPtr = StrPair(strncpy(temp, string, TLIMIT), name, data, NLIMIT);
20
        for ( ; name[0] != '\0'; nextPtr = StrPair(nextPtr, name, data, NLIMIT)) {
21
            if (strcmp(name, "name") == 0)
22
                strcpy(comPtr->name, data);
23
            else if (strcmp(name, "rx0") == 0)
24
                assert(comPtr->rxPortCount > 0), strcpy(comPtr->rxPortPtr[0].name, data);
25
            else if (strcmp(name, "rx1") == 0)
26
                assert(comPtr->rxPortCount > 1), strcpy(comPtr->rxPortPtr[1].name, data);
27
            else if (strcmp(name, "tx0") == 0)
28
                assert(comPtr->txPortCount > 0), strcpy(comPtr->txPortPtr[0].name, data);
29
        }
30
        assert(PortsCheck(comPtr, nextPtr) == TRUE);
31
    }
32
    return(comPtr);
33
}
34
SmallTime
35
NextRate(RxTimer *btPtr, SmallTime thisRxTime, SmallTime thisTxTime, SmallTime interval0, SmallTime interval1)
36
{
37
    PastTimes *timePtr;
38
    SmallTime rateRatio0, rateRatio1, thatDelta, thisDelta;
39
    uint16_t headIndex, tailIndex, lastIndex, limit;
40
    uint8_t i;
41
42
    assert(btPtr != NULL); // Verify the pointer
43
    timePtr = btPtr->times; // Array value pointer
44
    limit = ARRAY_SIZE(btPtr->times); // Array-size limits
45
46
    if (btPtr->headIndex == btPtr->tailIndex) { // Unitialized array
47
        assert(btPtr->validated == FALSE); // has no validated
48
        btPtr->headIndex = 1, btPtr->tailIndex = 0; // Initialize index
49
        timePtr[0].thisTime = thisRxTime; // and tail-indexed
50
        timePtr[0].thatTime = thisTxTime; // data values
51
    }
52
}

```

```

headIndex = btPtr->headIndex; // The head and last 1
lastIndex = PLUS(headIndex, -1, limit); // index values within 2
assert(headIndex < limit && lastIndex < limit); // the circular buffer 3

if ((timePtr[headIndex].thatTime - timePtr[lastIndex].thatTime) > (interval0 / 2)) { // Time to advance 4
    btPtr->headIndex = headIndex = PLUS(headIndex, 1, limit); // increment headIndex 5
    btPtr->validated = TRUE; // Set when ready 6
}
timePtr[headIndex].thisTime = thisRxTime; // Save received time 7
timePtr[headIndex].thatTime = thisTxTime; // Save transmit time 8
if (btPtr->validated == FALSE) // Until times change, 9
    return(SMALL_ONE); // assume slope==1

for (i = 0; i < 2; i += 1, btPtr->tailIndex = tailIndex) { // Check tailIndex twice 10
    tailIndex = PLUS(btPtr->tailIndex, 1, limit); // Next tailIndex value 11
    if (tailIndex == headIndex) // The tailIndex can 12
        break; // never equal headIndex 13
    if (thisTxTime - timePtr[tailIndex].thatTime <= interval1) // Update tailIndex if 14
        break; // range is maintained 15
}
tailIndex = btPtr->tailIndex; 16

thisDelta = thisRxTime - timePtr[tailIndex].thisTime; // Received interval 17
thatDelta = thisTxTime - timePtr[tailIndex].thatTime; // Transmit interval 18
assert(thatDelta != 0); // Must have changed 19
rateRatio0 = DivideSmall(thisDelta, thatDelta); // Compute the rate 20
rateRatio1 = CLIP_RATE(rateRatio0, PPM250); // Clip within 250PPM 21
return(rateRatio1); 22
} 23

Common *EfdxTxInit(Common *, char *); 24
// Receives GrandSync SDUs, retransmits them as duplex-Ethernet SDUs 25
// arguments: 26
// comPtr - associated state-maintaining data structure 27
// name - initialization trigger and assigned entity name 28
Common * 29
EfdxTxExec(Common *comPtr, char *name) { 30
    uint8_t txInfo[MTU_SIZED], rxInfo[MTU_SIZED], sxInfo[SizePlus(EfdxTxInfo)]; 31
    EfdxTxEntity *ePtr = (EfdxTxEntity *)comPtr; 32
    BaseTimer *btPtr = (ePtr->baseTimer); 33
    GrandSyncReq *rxPtr = (GrandSyncReq *)rxInfo; 34
    EfdxTxSaved *rcPtr = (ePtr->rxSaved); 35
    EfdxMacReq *txPtr = (EfdxMacReq *)txInfo; 36
    SyncSduData *rsPtr = (rxPtr->service_data_unit); 37
    EfdxSduData *tsPtr = (txPtr->service_data_unit);
    EfdxTxInfo *sxPtr = (EfdxTxInfo *)sxInfo;
    EfdxRxEntity *dPtr;
    RxPort *rxQueuePtr, *sxQueuePtr;
    TxPort *txQueuePtr;
    NextTimes nextTimes;
    GrandTime grandTime;
    SmallTime backInterval, execTime, nextTime, rateInterval, smallTime, stationTime, wakeTime;
    LocalTime localTime;
    TinyTime extraTime;
    uint16_t rxHeadSize, txHeadSize, size, sized;

```

```

Boolean    serviced;

if (name != NULL)
    return(EfdxTxInit(comPtr, name));
SetRxQueue2Ptrs(comPtr, rxQueuePtr, sxQueuePtr);
SetTxQueue1Ptrs(comPtr, txQueuePtr);
assert((dPtr = (EfdxRxEntity *)comPtr->pairLink) != NULL);

do {
    serviced = FALSE;
    stationTime = StationTime(comPtr);
    backInterval = (3 * rcPtr->syncInterval + ePtr->syncInterval) / 2;
    rateInterval = backInterval + (3 * ePtr->syncInterval) / 2;

    // ***** Processing arriving timeSync snapshots *****
    sized = DeQueue(sxQueuePtr, &sxInfo);
    if (sized != 0) {
        serviced = TRUE;
        assert(sized == sizeof(EfdxTxInfo));
        ePtr->sxSnapCount = FieldToUnsign(sxPtr, frameCount).lower;
        ePtr->sxSnapTimed = FieldToUnsign(sxPtr, smallTime).lower;
        ePtr->txReady = TRUE;
    }

    // ***** Processing arrived MS_DATAUNIT.request frames *****
    sized = DeQueue(rxQueuePtr, &rxInfo);
    if (sized != 0) {
        serviced = TRUE;
        if (sized != sizeof(GrandSyncReq) || !TimeSyncSdu(rsPtr)) {
            LongToFrame(rcPtr->da, txPtr, destination_address);
            LongToFrame(rcPtr->sa, txPtr, source_address);
            rxHeadSize = (void *)rsPtr - (void *)rxPtr;
            txHeadSize = (void *)tsPtr - (void *)txPtr;
            bcopy(rsPtr, tsPtr, sized - rxHeadSize);
            size = sized + txHeadSize - rxHeadSize;
            Enqueue(txQueuePtr, txPtr, size);
        } else {
            rcPtr->da = FieldToUnsign(rxPtr, destination_address).lower;
            rcPtr->sa = FieldToUnsign(rxPtr, source_address).lower;
            rcPtr->type = FieldToUnsign(rsPtr, protocolType).lower;
            rcPtr->function = FieldToUnsign(rsPtr, function).lower;
            rcPtr->version = FieldToUnsign(rsPtr, version).lower;
            rcPtr->hopCount = FieldToUnsign(rsPtr, hopCount).lower;
            rcPtr->precedence = FieldToUnsign(rsPtr, precedence);
            rcPtr->syncInterval = FieldToUnsign(rsPtr, syncInterval).lower;
            grandTime = FieldToSigned(rsPtr, grandTime);
            extraTime = FieldToSigned(rsPtr, extraTime).lower;
            smallTime = FieldToSigned(rsPtr, smallTime).lower;
            NextSaved(btPtr, ePtr->syncInterval, rateInterval,
                GrandToLarge(grandTime), TinyToSmall(extraTime), smallTime);
        }
    }
}

```



```

//          if (ePtr->guessMode == EXTRAPOLATE)
//          ePtr->execTime = stationTime + RESIDENCE_DELAY;
    }
}

// *****
// ***** Preparing transmitted timeSync frames *****
// *****
nextTime = ePtr->lastTime + ePtr->syncInterval;
if ((stationTime - nextTime) >= 0) { // Next sync transmission
    serviced = TRUE;
    ePtr->lastTime = nextTime; // Restart 10ms timer
    ePtr->execTime = stationTime + RESIDENCE_DELAY;
}
execTime = ePtr->execTime;
if ((stationTime - execTime) >= 0) {
    serviced = TRUE;
    ePtr->execTime = stationTime + LARGE_TOCK; // Indefinite future
    ePtr->frameCount = (ePtr->frameCount + 1) % COUNT; // Incremented counter
    if (!ePtr->txReady)
        ePtr->sxSnapTimed = (stationTime - ePtr->syncInterval);
    localTime = SmallToLocal(ePtr->sxSnapTimed);
    switch(ePtr->guessMode) {
    case INTERPOLATE:
        nextTimes = NextTimed(btPtr, ePtr->sxSnapTimed, backInterval);
        break;
    case EXTRAPOLATE:
        nextTimes = NextTimed(btPtr, ePtr->sxSnapTimed, (SmallTime)0);
        break;
    }

    LongToFrame(rcPtr->da, txPtr, destination_address); // The destination_address
    LongToFrame(rcPtr->sa, txPtr, source_address); // The source_address
    LongToFrame(AVB_PROTOCOL, tsPtr, protocolType); // The protocol identifier
    LongToFrame(AVB_FUNCTION, tsPtr, function); // The basic function
    LongToFrame(AVB_VERSION, tsPtr, version); // and version identifier
    WideToFrame(rcPtr->precedence, tsPtr, precedence); // Grandmaster precedence
    LongToFrame(rcPtr->hopCount, tsPtr, hopCount); // The ~GM distance.
    LongToFrame(ePtr->frameCount, tsPtr, frameCount); // Source-port identifier
    WideToFrame(nextTimes.grandTime, tsPtr, grandTime); // grandTime at snapShot
    LongToFrame(nextTimes.extraTime, tsPtr, extraTime); // Next extraTime value
    LongToFrame(localTime, tsPtr, localTime); // Transmitted frame time
    LongToFrame(dPtr->thisTxTime, tsPtr, thatTxTime); // Opposing transmit time
    LongToFrame(dPtr->thisRxTime, tsPtr, thatRxTime); // Opposing received time
    EnQueue(txQueuePtr, txPtr); // Enqueue the result
}
wakeTime = (execTime - nextTime) > 0 ? nextTime : execTime;
} while (serviced == TRUE);
assert((wakeTime - stationTime) > 0);
SleepOnBase(compPtr, wakeTime);
return(NULL);
}

Common *
EfdxTxInit(Common *oldPtr, char *string) {

```

```

Common      *comPtr;
EfdxTxEntity *ePtr;
BaseTimer   *btPtr;
EfdxTxSaved *rcPtr;
char temp[TLIMIT+1], name[NLIMIT+1], data[NLIMIT+1], *nextPtr;

comPtr = CommonCreate(oldPtr, sizeof(EfdxTxEntity), &EfdxTxExec, TYPE_TX_PORT, Q_RX22_LAST, Q_TX11_LAST);
if (comPtr != NULL) {
    ePtr = (EfdxTxEntity *)comPtr;
    rcPtr = &(ePtr->rxSaved);
    ePtr->syncInterval = rcPtr->syncInterval = SMALL_10ms;
    ePtr->guessMode = INTERPOLATE;

    rcPtr->precedence.upper = rcPtr->precedence.lower = ONES;
    rcPtr->hopCount = 0xFF;

    btPtr = &(ePtr->baseTimer);
    btPtr->validated = 0;
    btPtr->headIndex = btPtr->tailIndex = 0;

    nextPtr = StrPair(strncpy(temp, string, TLIMIT), name, data, NLIMIT);
    for ( ; name[0] != '\0'; nextPtr = StrPair(nextPtr, name, data, NLIMIT)) {
        if (strcmp(name, "name") == 0)
            strncpy(comPtr->name, data, NLIMIT);
        else if (strcmp(name, "rx0") == 0)
            assert(comPtr->rxPortCount > 0), strcpy(comPtr->rxPortPtr[0].name, data);
        else if (strcmp(name, "rx1") == 0)
            assert(comPtr->rxPortCount > 1), strcpy(comPtr->rxPortPtr[1].name, data);
        else if (strcmp(name, "tx0") == 0)
            assert(comPtr->txPortCount > 0), strcpy(comPtr->txPortPtr[0].name, data);
        else if (strcmp(name, "guessMode") == 0)
            ePtr->guessMode = atoi(data);
    }
    assert(PortsCheck(comPtr, nextPtr) == TRUE);
}
return(comPtr);
}

// The NextTimed() routine computes grandTime based on current txTime and
// previously sampled rxTimes information. The computation effect is:
// 1) Step back in time by a duration backInterval, to tbTime
// 2) Interpolate between rxTimes[n-N] and rxTimes[n+0], yielding tiTime
// 3) Extrapolate the tiTime forward, assuming slope==1, yielding grandTime
// 4) Extrapolate the tiTime forward, assuming rateRatio, yielding totalTime
// The rateRatio is the ratio of grandTime to stationTime changes.
// 5) Forward {extraTime = totalTime - grandTime} along with grandTime.
// The incoming extraTime is also filtered, but not extrapolated forward:
// 1) Step back in time by a duration backInterval, to tbTime
// 2) Interpolate between rxTimes[n-N] and rxTimes[n+0], yielding extraTime
// The value of backInterval is based on worst-case latencies:
// backInterval = (3 * thatInterval + thisInterval) / 2
// thatInterval - is the syncInterval for the selected clock-slave
// thisInterval - is the syncInterval for this clock-master port

NextTimes

```

```

NextTimed(BaseTimer *btPtr, SmallTime txTime, SmallTime backInterval) {
    BaseTimes *timePtr = btPtr->times;
    BaseTimes  thisTimes, pastTimes;
    NextTimes  nextTimes;
    LargeTime  largeTime0, largeTime1;
    SmallTime  deltaTime, extraDelta, extraTime0, extraTime1, extraTime2, extraTime3,
              grandDelta, rateRatio0, rateRatio1, smallDelta, weight;
    uint16_t   headIndex, tailIndex, extraCount, i;

    assert(btPtr != NULL);
    headIndex = btPtr->headIndex;
    tailIndex = btPtr->tailIndex;
    thisTimes = timePtr[headIndex];
    pastTimes = timePtr[tailIndex];

    grandDelta = LargeToSmall(WideSubtract(thisTimes.largeTime, pastTimes.largeTime));
    extraDelta = thisTimes.extraTime - pastTimes.extraTime;
    smallDelta = thisTimes.smallTime - pastTimes.smallTime;
    if (smallDelta == 0) {
        grandDelta = smallDelta = 2 * SMALL_10ms;
        extraDelta = 0;
    }

    weight = DivideSmall((txTime - backInterval) - thisTimes.smallTime, smallDelta);
    rateRatio0 = DivideSmall(grandDelta, smallDelta);
    rateRatio1 = CLIP_RATE(rateRatio0, PPM250);
    deltaTime = MultiplySmall((rateRatio1 - SMALL_ONE), backInterval);
    if (rateRatio1 != rateRatio0)
        grandDelta = MultiplySmall(rateRatio1, smallDelta);

    largeTime0 = WideAddition(thisTimes.largeTime, SmallToLarge(MultiplySmall(grandDelta, weight)));
    largeTime1 = WideAddition(largeTime0, SmallToLarge(backInterval));

    // Average the accumulated extraTime values...
    extraTime0 = extraCount = 0;
    for (i = tailIndex; ; i = PLUS(i, 1, ARRAY_SIZE(btPtr->times))) {
        extraTime0 += timePtr[i].extraTime;
        extraCount += timePtr[i].extraCount;
        if (i == headIndex)
            break;
    }
    assert(headIndex == tailIndex || extraCount != 0);
    extraTime1 = (extraCount != 0) ? (extraTime0 / extraCount) : 0;
    extraTime2 = extraTime1 + deltaTime;
    extraTime3 = CLIP_SIZE(extraTime2, (SMALL_ONE / 256) - 1);

    nextTimes.grandTime = LargeToGrand(largeTime1);
    nextTimes.extraTime = SmallToTiny(extraTime3);
    nextTimes.totalTime = LargeToGrand(WideAddition(largeTime1, SmallToLarge(extraTime3)));
    return(nextTimes);
}

void
NextSaved(BaseTimer *btPtr, SmallTime interval0, SmallTime interval1,
          LargeTime largeTime, SmallTime extraTime, SmallTime stationTime) {

```

```

BaseTimes *timePtr = btPtr->times;
uint16_t headIndex, tailIndex, lastIndex, limit;
uint8_t i;

assert(btPtr != NULL);
if (btPtr->headIndex == btPtr->tailIndex) {
    assert(btPtr->validated == FALSE);
    btPtr->headIndex = 1, btPtr->tailIndex = 0;
    timePtr[0].largeTime = largeTime;
    timePtr[0].smallTime = stationTime;
    timePtr[0].extraTime = extraTime;
    timePtr[0].extraCount = 1;
}
limit = ARRAY_SIZE(btPtr->times);
headIndex = btPtr->headIndex;
lastIndex = PLUS(headIndex, -1, limit);
assert(headIndex < limit && lastIndex < limit);
if (timePtr[headIndex].smallTime == stationTime)
    return;

if ((timePtr[headIndex].smallTime - timePtr[lastIndex].smallTime) > (interval0 / 2)) {
    btPtr->headIndex = headIndex = PLUS(headIndex, 1, limit);
    timePtr[headIndex].extraCount = timePtr[headIndex].extraTime = 0;
    btPtr->validated = TRUE;
}
timePtr[headIndex].largeTime = largeTime;
timePtr[headIndex].smallTime = stationTime;
timePtr[headIndex].extraTime += extraTime;
timePtr[headIndex].extraCount += 1;

for (i = 0; i < 2; i += 1, btPtr->tailIndex = tailIndex) {
    tailIndex = PLUS(btPtr->tailIndex, 1, limit);
    if (tailIndex == headIndex)
        break;
    if (stationTime - timePtr[tailIndex].smallTime <= interval1)
        break;
}
}

// *****
// ***** Wireless 802.11v wireless state-machine routines *****
// *****

Common *R11vRxInit(Common *, char *);
// Receives radio service-interface parameters, sends GrandSync an SDU
// arguments:
//   comPtr - associated state-maintaining data structure
//   name - initialization trigger and assigned entity name
Common *
R11vRxExec(Common *comPtr, char *name) {
    uint8_t r1Info[SizePlus(R11vInfo1Ind)], r2Info[SizePlus(R11vInfo2Ind)], txInfo[MTU_SIZED];
    R11vRxEntity *ePtr = (R11vRxEntity *)comPtr;
    R11vInfo1Ind *r1Ptr = (R11vInfo1Ind *)r1Info;
    R11vInfo2Ind *r2Ptr = (R11vInfo2Ind *)r2Info;
}

```

```

GrandSyncInd *txPtr = (GrandSyncInd *)txInfo;
SyncSduData *tsPtr = &(txPtr->service_data_unit);
RxPort *r1QueuePtr, *r2QueuePtr;
TxPort *txQueuePtr;
GrandTime grandTime;
SmallTime stationTime, cableDelay, totalDelay, smallTime;
TinyTime extraTime;
TicksTime ticksTime;
uint64_t da, sa;
uint32_t sized;
uint8_t hopCount, portID;
Boolean serviced;

if (name != NULL)
    return(R1lvRxInit(comPtr, name));
SetRxQueue2Ptrs(comPtr, r1QueuePtr, r2QueuePtr);
SetTxQueue1Ptrs(comPtr, txQueuePtr);

portID = PortID(comPtr);
do {
    stationTime = StationTime(comPtr);
    serviced = FALSE;

    // ***** Processing arriving MLME_PRESENCE_REQUEST.indication snapshots *****
    sized = DeQueue(r1QueuePtr, &r1Info);
    if (sized != 0) {
        assert(sized == sizeof(R1lvInfo1Ind));
        ePtr->turnRound = r1Ptr->ticksTime3 - r1Ptr->ticksTime2;
    }

    // ***** Processing arriving MLME_PRESENCE_RESPONSE.indication snapshots *****
    sized = DeQueue(r2QueuePtr, &r2Info);
    if (sized != 0) {
        assert(sized == sizeof(R1lvInfo2Ind));
        serviced = TRUE;
        ticksTime = R1lvTime(ePtr);
        cableDelay = MIN(0, r2Ptr->roundTrip - ePtr->turnRound) / 2;
        totalDelay = cableDelay + (ticksTime - r2Ptr->ticksTime4);
        grandTime = LevelToGrand(r2Ptr->levelTime);
        hopCount = r2Ptr->hopCount;
        extraTime = r2Ptr->extraTime;
        smallTime = stationTime - MultiplySmall(totalDelay, RADIO_TICK_TIME);

        // ***** Creation of service-data-unit parameters *****
        LongToFrame(da, txPtr, destination_address);
        LongToFrame(sa, txPtr, source_address);
        LongToFrame(hopCount, tsPtr, hopCount);
        LongToFrame(portID, tsPtr, sourcePort);
        WideToFrame(grandTime, tsPtr, grandTime);
    }
} while (!serviced);

// The entity name
// for initialization.

// Station's localTime

// Station local times
// Cable delay ticks
// Elapsed time
// Grandmaster time
// Grandmaster time
// Grandmaster time
// Local-time stamp

// Destination address
// Source-port identifier
// The ~GM distance.
// Source-port identifier
// grandTime at snapShot

```

```

        LongToFrame(extraTime,          tsPtr, extraTime);           // Next extraTime value
        LongToFrame(smallTime,         tsPtr, smallTime);           // Transmitted frame time
        LongToFrame(ePtr->syncInterval, tsPtr, syncInterval);       // Sync transmit interval
        EnQueue(txQueuePtr, txPtr);                                   // Enqueue the result
    }
} while (serviced == TRUE);
return(NULL);
}

Common *
R11vRxInit(Common *oldPtr, char *string) {
    Common *comPtr;
    R11vRxEntity *ePtr;
    char temp[TLIMIT+1], name[NLIMIT+1], data[NLIMIT+1], *nextPtr;

    comPtr = CommonCreate(oldPtr, sizeof(R11vRxEntity), &R11vRxExec, TYPE_RX_PORT, Q_RX22_LAST, Q_TX11_LAST);
    if (comPtr != NULL) {
        ePtr = (R11vRxEntity *)comPtr;                                // Setup entity pointer
        ePtr->syncInterval = SMALL_10ms;                             // Set default interval

        nextPtr = StrPair(strncpy(temp, string, TLIMIT), name, data, NLIMIT);
        for ( ; name[0] != '\0'; nextPtr = StrPair(nextPtr, name, data, NLIMIT)) {
            if (strcmp(name, "name") == 0)
                strcpy(comPtr->name, data, NLIMIT);
            else if (strcmp(name, "rx0") == 0)
                assert(comPtr->rxPortCount > 0), strcpy(comPtr->rxPortPtr[0].name, data);
            else if (strcmp(name, "rx1") == 0)
                assert(comPtr->rxPortCount > 1), strcpy(comPtr->rxPortPtr[1].name, data);
            else if (strcmp(name, "tx0") == 0)
                assert(comPtr->txPortCount > 0), strcpy(comPtr->txPortPtr[0].name, data);
        }
        assert(PortsCheck(comPtr, nextPtr) == TRUE);
    }
    return(comPtr);
}

Common *R11vTxInit(Common *, char *);

// Receives radio GrandSync SDUs, retransmits as service-interface parameters
// arguments:
//   comPtr - associated state-maintaining data structure
//   name - initialization trigger and assigned entity name
Common *
R11vTxExec(Common *comPtr, char *name) {
    uint8_t rxInfo[MTU_SIZED], c2Info[SizePlus(R11vInfo2Con)], c1Info[SizePlus(R11vInfo1Con)];
    R11vTxEntity *ePtr = (R11vTxEntity *)comPtr;
    BaseTimer *btPtr = &(ePtr->baseTimer);
    GrandSyncReq *rxPtr = (GrandSyncReq *)rxInfo;
    R11vTxSaved *sxPtr = &(ePtr->rxSaved);
    SyncSduData *rsPtr = &(rxPtr->service_data_unit);
    R11vInfo2Con *c2Ptr = (R11vInfo2Con *)c2Info;
    R11vInfo1Con *c1Ptr = (R11vInfo1Con *)c1Info;
    R11vRxEntity *dPtr;
    R11vInfo1Req r1Info, *r1Ptr = &r1Info;

```

```

R1lvInfo2Req *r2Ptr, r2Info;
TxPort      *r1QueuePtr, *r2QueuePtr;
RxPort      *c1QueuePtr, *c2QueuePtr, *rxQueuePtr;
NextTimes   nextTimes;
GrandTime   grandTime;
SmallTime   backInterval, lapseTime, nextTime, rateInterval, stationTime, smallTime;
TicksTime   ticksTime;
TinyTime    extraTime;
uint32_t    sized;
Boolean     serviced;

if (name != NULL) // The entity name
    return(R1lvTxInit(comPtr, name)); // for initialization
SetRxQueue3Ptrs(comPtr, c1QueuePtr, c2QueuePtr, rxQueuePtr);
SetTxQueue2Ptrs(comPtr, r1QueuePtr, r2QueuePtr);

assert((dPtr = (R1lvRxEntity *) (comPtr->pairLink)) != NULL); // Receiver pair
do { // Station's localTime
    stationTime = StationTime(comPtr);
    serviced = TRUE;
    backInterval = (3 * sxPtr->syncInterval + ePtr->syncInterval) / 2;
    rateInterval = backInterval + (3 * ePtr->syncInterval) / 2;

    // *****
    // ***** Preparing transmitted MLME_PRESENCE_RESPONSE.request information *****
    // *****
    nextTime = ePtr->lastTime + SMALL_10ms;
    if ((stationTime - nextTime) >= 0) { // Next sync transmission
        serviced = TRUE;
        ePtr->lastTime = nextTime; // Restart 10ms timer
        EnQueue(r2QueuePtr, r1Ptr); // Enqueue the trigger
    }

    // *****
    // ***** Processing arrived MS_DATAUNIT.request frames *****
    // *****
    sized = DeQueue(rxQueuePtr, rxPtr);
    if (sized != 0) {
        serviced = TRUE;
        assert(sized == sizeof(GrandSyncReq));
        sxPtr->hopCount = FieldToUnsign(rsPtr, hopCount).lower;
        sxPtr->precedence = FieldToUnsign(rsPtr, precedence);
        sxPtr->syncInterval = FieldToUnsign(rsPtr, syncInterval).lower;
        grandTime = FieldToSigned(rsPtr, grandTime);
        extraTime = FieldToSigned(rsPtr, extraTime).lower;
        smallTime = FieldToSigned(rsPtr, smallTime).lower;
        NextSaved(btPtr, ePtr->syncInterval, rateInterval, GrandToLarge(grandTime), TinyToSmall(extraTime), smallTime);
    }

    // *****
    // ***** Processing arriving MLME_PRESENCE_REQUEST.confirm information *****
    // *****
    sized = DeQueue(c1QueuePtr, &c1Info);
    if (sized != 0) {
        serviced = TRUE;

```

```

    assert(sized == sizeof(R1lvInfo1Con));
    ePtr->snapShot1 = c1Ptr->ticksTime1;
    ePtr->snapShot4 = c1Ptr->ticksTime4;
    ePtr->respondNow = TRUE;
}
// *****
// ***** Preparing transmitted MLME_PRESENCE_RESPONSE.request information *****
// *****
if (ePtr->respondNow == TRUE) {
    serviced = TRUE;
    ePtr->respondNow = TRUE;
    ticksTime = R1lvTime(dPtr);
    lapseTime = ticksTime - ePtr->snapShot4;
    smallTime = stationTime - MultiplySmall(lapseTime, RADIO_TICK_TIME);
    nextTimes = NextTimed(btPtr, smallTime, backInterval);

    r2Ptr = &r2Info;
    r2Ptr->ticksTime4 = ePtr->snapShot4;
    r2Ptr->roundTrip = ePtr->roundTrip;
    r2Ptr->levelTime = GrandToLevel(nextTimes.grandTime);
    r2Ptr->extraTime = nextTimes.extraTime;
    r2Ptr->precedence = sxPtr->precedence;
    r2Ptr->hopCount = sxPtr->hopCount;
    EnQueue(r2QueuePtr, r2Ptr);
}
// *****
// ***** Processing arriving MLME_PRESENCE_RESPONSE.confirm information *****
// *****
sized = DeQueue(c2QueuePtr, &c2Info);
if (sized != 0) {
    serviced = TRUE;
    assert(sized == sizeof(R1lvInfo2Con));
    assert(c2Ptr != NULL);
}
} while (serviced == TRUE);
assert((nextTime - stationTime) > 0);
SleepOnBase(comPtr, nextTime);
return(NULL);
}

Common *
R1lvTxInit(Common *oldPtr, char *string) {
    Common *comPtr;
    R1lvTxEntity *ePtr;
    BaseTimer *btPtr;
    char temp[TLIMIT+1], name[NLIMIT+1], data[NLIMIT+1], *nextPtr;

    comPtr = CommonCreate(oldPtr, sizeof(R1lvTxEntity), &R1lvTxExec, TYPE_TX_PORT, Q_RX33_LAST, Q_TX22_LAST);
    if (comPtr != NULL) {
        ePtr = (R1lvTxEntity *)comPtr;
        btPtr = &(ePtr->baseTimer);
        btPtr->validated = 0;
        btPtr->headIndex = btPtr->tailIndex = 0;
    }
}

```



```

nextPtr = StrPair(strncpy(temp, string, TLIMIT), name, data, NLIMIT);
for ( ; name[0] != '\0'; nextPtr = StrPair(nextPtr, name, data, NLIMIT)) {
    if (strcmp(name, "name") == 0)
        strncpy(comPtr->name, data, NLIMIT);
    else if (strcmp(name, "rx0") == 0)
        assert(comPtr->rxPortCount > 0), strcpy(comPtr->rxPortPtr[0].name, data);
    else if (strcmp(name, "rx1") == 0)
        assert(comPtr->rxPortCount > 1), strcpy(comPtr->rxPortPtr[1].name, data);
    else if (strcmp(name, "rx2") == 0)
        assert(comPtr->rxPortCount > 2), strcpy(comPtr->rxPortPtr[2].name, data);
    else if (strcmp(name, "tx0") == 0)
        assert(comPtr->txPortCount > 0), strcpy(comPtr->txPortPtr[0].name, data);
    else if (strcmp(name, "tx1") == 0)
        assert(comPtr->txPortCount > 1), strcpy(comPtr->txPortPtr[1].name, data);
}
assert(PortsCheck(comPtr, nextPtr) == TRUE);
return(comPtr);
}

Common *EponRxInit(Common *, char *);

// *****
// ***** Ethernet-PON state-machine routines *****
// *****

// Receives Ethernet-PON SDUs, reformats and sends to GrandSync
// arguments:
//   comPtr - associated state-maintaining data structure
//   name - initialization trigger and assigned entity name
Common *
EponRxExec(Common *comPtr, char *name) {
    uint8_t rxInfo[MTU_SIZED], txInfo[MTU_SIZED];
    EponRxEntity *ePtr = (EponRxEntity *)comPtr;
    EponMacInd *rxPtr = (EponMacInd *)rxInfo;
    GrandSyncInd *txPtr = (GrandSyncInd *)txInfo;
    SyncSduEpon *rsPtr = &(rxPtr->service_data_unit);
    SyncSduData *tsPtr = &(txPtr->service_data_unit);
    RxPort *rxQueuePtr;
    TxPort *txQueuePtr;
    GrandTime grandTime;
    SmallTime smallTime, stationTime;
    TinyTime extraTime;
    TicksTime lapseTime, ticksTime, ponTime;
    uint64_t da, sa;
    uint32_t sized;
    uint8_t hopCount, portID;
    Boolean serviced;

    if (name != NULL)
        return(EponRxInit(comPtr, name));
    SetRxQueuePtrs(comPtr, rxQueuePtr);
    SetTxQueuePtrs(comPtr, txQueuePtr);
}

```

```

// The entity name
// for initialization.

```

```

portID = PortID(comPtr);
do {
    serviced = FALSE;
    stationTime = StationTime(comPtr); // Station's localTime

    // *****
    // ***** Processing arriving timeSync frames *****
    // *****
    sized = DeQueue(rxQueuePtr, rxPtr);
    if (sized != 0) {
        serviced = TRUE;
        assert(sized == sizeof(EponMacInd));
        ponTime = EponTime(ePtr);

        // *****
        // ***** Extract frame parameters and perform basic consistency checks *****
        // *****
        da = FieldToUnsign(rxPtr, destination_address).lower;
        sa = FieldToUnsign(rxPtr, source_address).lower;
        hopCount = FieldToUnsign(rsPtr, hopCount).lower;
        grandTime = FieldToSigned(rsPtr, grandTime);
        extraTime = FieldToSigned(rsPtr, extraTime).lower;
        ticksTime = FieldToSigned(rsPtr, ticksTime).lower;
        lapseTime = ponTime - ticksTime;
        smallTime = stationTime - MultiplySmall(lapseTime, PON_TICK_TIME);

        // *****
        // ***** Update revised service-data-unit parameters *****
        // *****
        LongToFrame(da, txPtr, destination_address); // Destination address
        LongToFrame(sa, txPtr, source_address); // Source-port identifier
        LongToFrame(hopCount, tsPtr, hopCount); // The ~GM distance.
        LongToFrame(portID, tsPtr, sourcePort); // Source-port identifier
        WideToFrame(grandTime, tsPtr, grandTime); // grandTime at snapShot
        LongToFrame(extraTime, tsPtr, extraTime); // Next extraTime value
        LongToFrame(smallTime, tsPtr, smallTime); // Transmitted frame time
        LongToFrame(ePtr->syncInterval, tsPtr, syncInterval); // Sync transmit interval
        EnQueue(txQueuePtr, txPtr); // Enqueue the result
    }
} while (serviced == TRUE);
return(NULL);
}

Common *EponRxInit(Common *oldPtr, char *string) {
    Common *comPtr;
    EponRxEntity *ePtr;
    char temp[TLIMIT+1], name[NLIMIT+1], data[NLIMIT+1], *nextPtr;

    comPtr = CommonCreate(oldPtr, sizeof(EponRxEntity), &EponRxExec, TYPE_RX_PORT, Q_RX11_LAST, Q_TX11_LAST);
    if (comPtr != NULL) {
        ePtr = (EponRxEntity *)comPtr;
        ePtr->syncInterval = SMALL_10ms;
        nextPtr = StrPair(strncpy(temp, string, TLIMIT), name, data, NLIMIT);
        for ( ; name[0] != '\0'; nextPtr = StrPair(nextPtr, name, data, NLIMIT)) {
            if (strcmp(name, "name") == 0)

```

```

        strncpy(comPtr->name, data, NLIMIT);
    else if (strcmp(name, "rx0") == 0)
        assert(comPtr->rxPortCount > 0), strcpy(comPtr->rxPortPtr[0].name, data);
    else if (strcmp(name, "tx0") == 0)
        assert(comPtr->txPortCount > 0), strcpy(comPtr->txPortPtr[0].name, data);
    }
    assert(PortsCheck(comPtr, nextPtr) == TRUE);
}
return(comPtr);
}

Common *EponTxInit(Common *, char *);

// Receives GrandSync SDU, reformats/resends as Ethernet-PON SDU
// arguments:
//   comPtr - associated state-maintaining data structure
//   name - initialization trigger and assigned entity name
Common *
EponTxExec(Common *comPtr, char *name) {
    uint8_t rxInfo[MTU_SIZED], txInfo[MTU_SIZED];
    EponTxEntity *ePtr = (EponTxEntity *)comPtr;
    BaseTimer *btPtr = &(ePtr->baseTimer);
    GrandSyncReq *rxPtr = (GrandSyncReq *)rxInfo;
    EponTxSaved *sxPtr = &(ePtr->rxSaved);
    EponMacReq *txPtr = (EponMacReq *)&txInfo;
    SyncSduData *rsPtr = &(rxPtr->service_data_unit);
    SyncSduEpon *tsPtr = &(txPtr->service_data_unit);
    RxPort *rxQueuePtr;
    TxPort *txQueuePtr;
    NextTimes nextTimes;
    GrandTime grandTime;
    SmallTime backInterval, smallTime, rateInterval;
    TicksTime ticksTime;
    TinyTime extraTime;
    LocalTime nextTime, stationTime;
    uint32_t sized;
    Boolean serviced;

    if (name != NULL)
        return(EponTxInit(comPtr, name));
    SetRxQueuePtrs(comPtr, rxQueuePtr);
    SetTxQueuePtrs(comPtr, txQueuePtr);

    do {
        serviced = FALSE;
        stationTime = StationTime(comPtr);
        backInterval = (3 * sxPtr->syncInterval + ePtr->syncInterval) / 2;
        rateInterval = backInterval + (3 * ePtr->syncInterval) / 2;

        // *****
        // ***** Processing arrived MS_DATAUNIT.request frames *****
        // *****
        sized = DeQueue(rxQueuePtr, rxPtr);
        if (sized != 0) {

```

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37

```

serviced = TRUE;
assert(sized == sizeof(GrandSyncReq));
sxPtr->da = FieldToUnsign(rxPtr, destination_address).lower;
sxPtr->sa = FieldToUnsign(rxPtr, source_address).lower;
sxPtr->hopCount = FieldToUnsign(rsPtr, hopCount).lower;
sxPtr->precedence = FieldToUnsign(rsPtr, precedence);
sxPtr->syncInterval = FieldToUnsign(rsPtr, syncInterval).lower;
grandTime = FieldToSigned(rsPtr, grandTime);
extraTime = FieldToSigned(rsPtr, extraTime).lower;
smallTime = FieldToSigned(rsPtr, smallTime).lower;
NextSaved(btPtr, ePtr->syncInterval, rateInterval, GrandToLarge(grandTime), TinyToSmall(extraTime), smallTime);
}

// *****
// ***** Preparing transmitted timeSync frames *****
// *****
nextTime = ePtr->lastTime + SMALL_10ms;
if ((stationTime - nextTime) >= 0) { // Next sync transmission
    serviced = TRUE;
    ePtr->lastTime = nextTime; // Restart 10ms timer
    ticksTime = EponTime((EponRxEntity *)comPtr); // Get localTime values
    nextTimes = NextTimed(btPtr, stationTime, backInterval);
    LongToFrame(sxPtr->da, txPtr, destination_address); // The destination_address
    LongToFrame(sxPtr->sa, txPtr, source_address); // The source_address
    WideToFrame(sxPtr->precedence, tsPtr, precedence); // GM precedence
    LongToFrame(sxPtr->hopCount, tsPtr, hopCount); // GM distance
    WideToFrame(nextTimes.grandTime, tsPtr, grandTime); // grandTime at snapShot
    LongToFrame(nextTimes.extraTime, tsPtr, extraTime); // Next grandTime value
    LongToFrame(ticksTime, tsPtr, ticksTime); // Next extraTime value
    EnQueue(txQueuePtr, txPtr); // Enqueue the result
}
} while (serviced == TRUE);
assert((nextTime - stationTime) > 0);
SleepOnBase(comPtr, nextTime);
return(NULL);
}

Common *
EponTxInit(Common *oldPtr, char *string) {
    Common *comPtr;
    EponTxEntity *ePtr;
    BaseTimer *btPtr;
    char temp[TLIMIT+1], name[NLIMIT+1], data[NLIMIT+1], *nextPtr;

    comPtr = CommonCreate(oldPtr, sizeof(EponTxEntity), &EponTxExec, TYPE_TX_PORT, Q_RX11_LAST, Q_TX11_LAST);
    if (comPtr != NULL) {
        ePtr = (EponTxEntity *)comPtr;
        btPtr = &(ePtr->baseTimer);
        btPtr->validated = 0;
        btPtr->headIndex = btPtr->tailIndex = 0;

        nextPtr = StrPair(strncpy(temp, string, TLIMIT), name, data, NLIMIT);
        for (; name[0] != '\0'; nextPtr = StrPair(nextPtr, name, data, NLIMIT)) {
            if (strcmp(name, "name") == 0)
                strncpy(comPtr->name, data, NLIMIT);
        }
    }
}

```

```
    else if (strcmp(name, "rx0") == 0)
        assert(comPtr->rxPortCount > 0), strcpy(comPtr->rxPortPtr[0].name, data);
    else if (strcmp(name, "tx0") == 0)
        assert(comPtr->txPortCount > 0), strcpy(comPtr->txPortPtr[0].name, data);
    }
    assert(PortsCheck(comPtr, nextPtr) == TRUE);
}
return(comPtr);
}
```

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54