

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

Draft 0.200

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: residential, Ethernet, isochronous, real time

1 **Editors' Foreword**

2
3
4 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.**

5
6
7
8 By fixing these errors in early drafts, readers will be able to devote their valuable time and energy to
9 comments that materially affect either the technical content of the document or the clarity of that technical
10 content. Comments should not simply state what is wrong, but also what might be done to fix the problem.

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

13
14 <http://ieee802.org/1/>

15
16 Use of the email distribution list is not presently restricted to 802.1 members, and the working group has had
17 a policy of considering ballot comments from all who are interested and willing to contribute to the devel-
18 opment of the draft. Individuals not attending meetings have helped to identify sources of misunderstanding
19 and ambiguity in past projects. Non-members are advised that the email lists exist primarily to allow the
20 members of the working group to develop standards, and are not a general forum.

21
22 Comments on this document may be sent to the 802.1 email reflector, to the editors, or to the Chairs of the
23 802.1 Working Group and Interworking Task Group.

24
25 This draft was prepared by:

26
27 David V James
28 JGG
29 3180 South Court
30 Palo Alto, CA 94306
31 +1.650.494.0926 (Tel)
32 +1.650.954.6906 (Mobile)
33 Email: dvj@alum.mit.edu

34
35 Chairs of the 802.1 Working Group and Audio/Video Bridging Task Group:.

36
37
38
39 Michael Johas Teener
40 Chair, 802.1 Audio/Video Bridging Task
41 Broadcom Corporation
42 3151 Zanker Road
43 San Jose, CA
44 95134-1933
45 USA
46 +1 408 922 7542 (Tel)
47 +1 831 247 9666 (Mobile)
48 Email:mikejt@broadcom.com

39 Tony Jeffree
40 Group Chair, 802.1 Working Group
41 11A Poplar Grove
42 Sale
43 Cheshire
44 M33 3AX
45 UK
46 +44 161 973 4278 (Tel)
47 +44 161 973 6534 (Fax)
48 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 Stds 802.1D and 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
 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

Contributors

This working paper is based on contributions or review comments from the people listed below. Their listing doesn't necessarily imply they agree with the entire content or the author's interpretation of their input.

Jim Battaglia	Pioneer
Alexei Beliaev	Gibson
Dirceu Cavendish	NEC Labs America
George Claseman	Micrel
Feifei (Felix) Feng	Samsung Electronics
John Nels Fuller	Independent
Geoffrey M. Garner	Samsung Electronics
Kevin Gross	Cirrus Logic
Jim Haagen-Smit	HP
David V James	JGG
Dennis Lou	Pioneer
Michael D. Johas Teener	Broadcom
Fred Tuck	EchoStar

Version history

Version	Date	Author	Comments
0.082	2005Apr28	DVJ	Updates based on 2005Apr27 meeting discussions
0.085	2005May11	DVJ	– Updated front-page list of contributors – Updated book for continuous pages (Clause 1 discontinuity fixed) – Miscellaneous editing fixes
0.088	2005Jun03	DVJ	– Application latency scenarios clarified.
0.090	2005Jun06	DVJ	– Misc editorials in bursting and bunching annex.
0.092	2005Jun10	DVJ	– Extensive cleanup of Clause 5 subscription protocols, based on 2005Jun08 teleconference review comments.
0.121	2005Jun24	DVJ	– Extensive cleanup of clock-synchronization protocols, base on 2005Jun22 teleconference review comments.
0.127	2005Jul04	DVJ	– Pacing descriptions greatly enhanced.
0.200	2007Jan23	DVJ	Removal of non time-sync related information. Update based on recent teleconference suggestion (layering), as well as input available from others' drafts.
—	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>

Topics for discussion

Readers are encouraged to provide feedback in all areas, although only the following areas have been identified as specific areas of concern.

- a) Layering. Should be reviewed.

TBDs

Further definitions are needed in the following areas:

- a) Details of the client time-sync services should be defined.
- b) Details of the Ethernet per-port time-sync services should be defined.

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.....	8
4		
5	List of tables.....	9
6		
7	1. Overview.....	11
8		
9	1.1 Scope	11
10	1.2 Purpose	11
11	1.3 Introduction	11
12		
13	2. References.....	13
14		
15	3. Terms, definitions, and notation	14
16		
17	3.1 Conformance levels	14
18	3.2 Terms and definitions	14
19	3.3 Service definition method and notation.....	15
20	3.4 State machines	16
21	3.5 Arithmetic and logical operators	19
22	3.6 Numerical representation.....	19
23	3.7 Field notations	20
24	3.8 Bit numbering and ordering.....	21
25	3.9 Byte sequential formats	22
26	3.10 Ordering of multibyte fields	22
27	3.11 MAC address formats.....	23
28	3.12 Informative notes.....	24
29	3.13 Conventions for C code used in state machines	24
30		
31	4. Abbreviations and acronyms	25
32		
33	5. Architecture overview	27
34		
35	5.1 Application scenarios	27
36	5.2 Design methodology.....	28
37	5.3 Time-synchronization facilities	29
38	5.4 Client processing details.....	32
39	5.5 Key distinctions from IEEE Std 1588	33
40		
41	6. Medium access control (MAC) service and reference model	35
42		
43	6.1 Overview	35
44	6.2 Overview of MAC services.....	35
45	6.3 MAC services to the client layer	35
46		
47	7. Ethernet duplex-cable time synchronization	38
48		
49	7.1 Synchronized time-of-day clocks	38
50	7.2 timeSync frame format.....	41
51	7.3 Per-port processing details	44
52		
53		
54		

Annex A (informative) Bibliography	45	1
		2
Annex B (informative) Bridging to IEEE Std 1394.....	46	3
		4
B.1 Hybrid network topologies	46	5
		6
Annex C (informative) Review of possible alternatives	48	7
		8
C.1 Clock-synchronization alternatives	48	9
		10
Annex D (informative) Time-of-day format considerations	50	11
		12
D.1 Possible time-of-day formats.....	50	13
D.2 Time format comparisons.....	52	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		
3	Figure 1.1—Topology and connectivity	12
4	Figure 3.1—Service definitions	15
5	Figure 3.2—Bit numbering and ordering	21
6	Figure 3.3—Byte sequential field format illustrations	22
7	Figure 3.4—Multibyte field illustrations	22
8	Figure 3.5—Illustration of fairness-frame structure	23
9	Figure 3.6—MAC address format	23
10	Figure 3.7—48-bit MAC address format.....	24
11		
12	Figure 5.1—Garage jam session.....	27
13	Figure 5.2—Possible looping topology	28
14	Figure 5.3—Timing information flows	29
15	Figure 5.4—Grand-master precedence flows	30
16	Figure 5.5—Grand-master precedence	30
17	Figure 5.6—Hierarchical flows	31
18	Figure 5.7—AVB service interface model.....	32
19		
20	Figure 6.1—AVB service and reference model relationship to the ISO/IEC OSI reference model	35
21	Figure 7.1—Rate-adjustment effects	39
22	Figure 7.2—Timer snapshot locations.....	40
23	Figure 7.3—timeSync frame format.....	41
24	Figure 7.4— <i>typeCount</i> format.....	42
25	Figure 7.5— <i>gmSelection</i> format.....	42
26	Figure 7.6— <i>uniqueID</i> format	43
27	Figure 7.7—Complete seconds timer format.....	43
28	Figure B.1—IEEE 1394 leaf domains	46
29	Figure B.2—IEEE 802.3 leaf domains	46
30	Figure B.3—Time-of-day format conversions	47
31	Figure B.4—Grand-master precedence mapping	47
32	Figure 4.1—Complete seconds timer format.....	50
33	Figure D.2—IEEE 1394 timer format	50
34	Figure D.3—IEEE 1588 timer format	51
35	Figure D.4—EPON timer format.....	51
36	Figure D.5—Compact seconds timer format.....	51
37	Figure D.6—Nanosecond timer format	51
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 17

Table 3.2—Called state table notation example 18

Table 3.3—Special symbols and operators..... 19

Table 3.4—Names of fields and sub-fields 20

Table 3.5—*wrap* field values 21

Table 7.1—Clock-synchronization intervals 38

Table C.1 Protocol comparison 48

Table D.1—Time format comparison..... 52

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 Stds 802.1D and 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

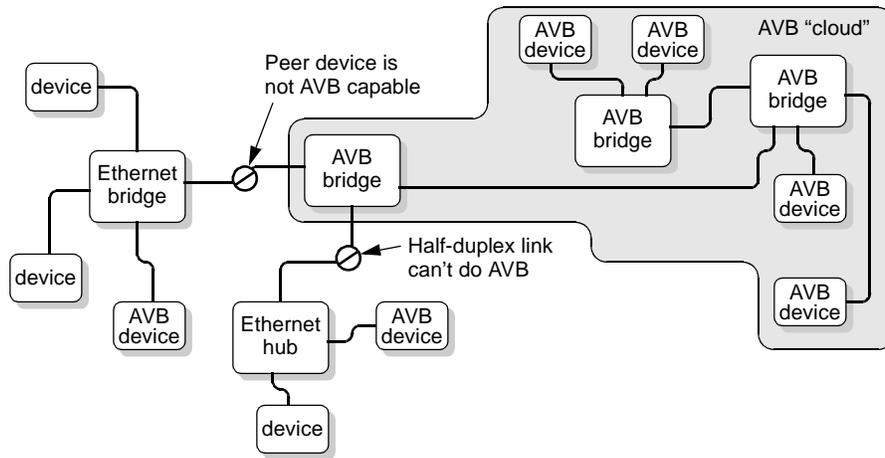
Ethernet has successfully propagated from the data center to the home, becoming the wired home computer interconnect of choice. However, insufficient support of real-time services has limited Ethernet's success as a consumer audio-video interconnects, where IEEE Std 1394 Serial Bus and Universal Serial 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 xx: FrameFormats
 - 40 — Clause xx: xx
 - 41 — Annex A: Bibliography
 - 42 — Annex B: Bridging to IEEE Std 1394
 - 43 — Annex C: Review of possible alternatives
 - 44 — Annex D: Time-of-day format considerations
 - 45 — Annex E: C-code illustrations
- 46
47
48
49
50
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 grand clock master: The clock master 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 Service definition method and notation

The service of a layer or sublayer is the set of capabilities that it offers to a user in the next higher (sub)layer. Abstract services are specified in this working paper by describing the service primitives and parameters that characterize each service. This definition of service is independent of any particular implementation (see Figure 3.1).

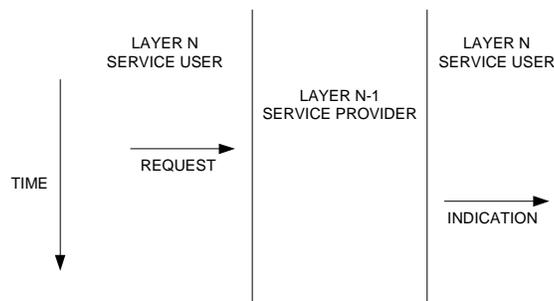


Figure 3.1—Service definitions

Specific implementations can also include provisions for interface interactions that have no direct end-to-end effects. Examples of such local interactions include interface flow control, status requests and indications, error notifications, and layer management. Specific implementation details are omitted from this service specification, because they differ from implementation to implementation and also because they do not impact the peer-to-peer protocols.

3.3.1 Classification of service primitives

Primitives are of two generic types.

- a) REQUEST. The request primitive is passed from layer N to layer N-1 to request that a service be initiated.
- b) INDICATION. The indication primitive is passed from layer N-1 to layer N to indicate an internal layer N-1 event that is significant to layer N. This event can be logically related to a remote service request, or can be caused by an event internal to layer N-1.

The service primitives are an abstraction of the functional specification and the user-layer interaction. The abstract definition does not contain local detail of the user/provider interaction. For instance, it does not indicate the local mechanism that allows a user to indicate that it is awaiting an incoming call. Each primitive has a set of zero or more parameters, representing data elements that are passed to qualify the functions invoked by the primitive. Parameters indicate information available in a user/provider interaction. In any particular interface, some parameters can be explicitly stated (even though not explicitly defined in the primitive) or implicitly associated with the service access point. Similarly, in any particular protocol specification, functions corresponding to a service primitive can be explicitly defined or implicitly available.

3.4 State machines

3.4.1 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.

3.4.2 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.

3.4.2.1 Parallel-execution state tables

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 3.13). 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	—	START
	—	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.

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

3.4.2.2 Called state tables

A RETURN state is the terminal state of a state machine that is intended to be invoked by another state machine, as illustrated in Table 3.2. Once the RETURN state is reached, the state machine terminates execution, effectively ceasing to exist until the next invocation by the caller, at which point it begins execution again from the START state. State machines that contain a RETURN state are considered to be only instantiated when they are invoked. They do not have any persistent (static) variables.

Table 3.2—Called state table notation example

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

Row 3.2-1: The size of the queued MAC control frame is less than the PTQ space.

Row 3.2-2: In the absence of MAC control transmission credits, no action is taken.

Row 3.2-3: MAC control transmissions have precedence over client transmissions.

Row 3.2-4: If the transmission completes with an error, set an error defect indication.

Row 3.2-5: Otherwise, no error defect is indicated.

3.5 Arithmetic and logical operators

In addition to commonly accepted notation for mathematical operators, Table 3.3 summarizes the symbols used to represent arithmetic and logical (boolean) operations. Note that the syntax of operators follows standard C notation (see 3.13).

Table 3.3—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

3.6 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”.

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”.

3.7 Field notations

3.7.1 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.

3.7.2 Field conventions

This working paper describes values that are packetized or MAC-resident, such as those illustrated in Table 3.2.

Table 3.4—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.

3.7.3 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 3.5. 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 3.5—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.

3.8 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 3.2, 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.

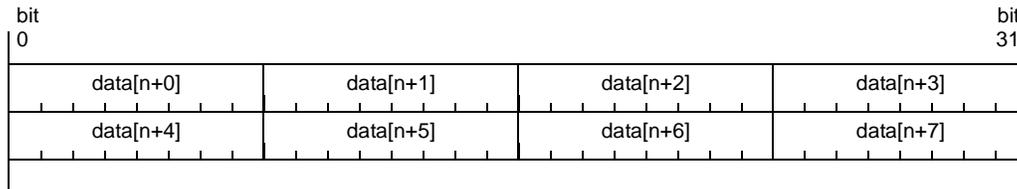


Figure 3.2—Bit numbering and ordering

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.

3.9 Byte sequential formats

Figure 3.3 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 3.3, 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.

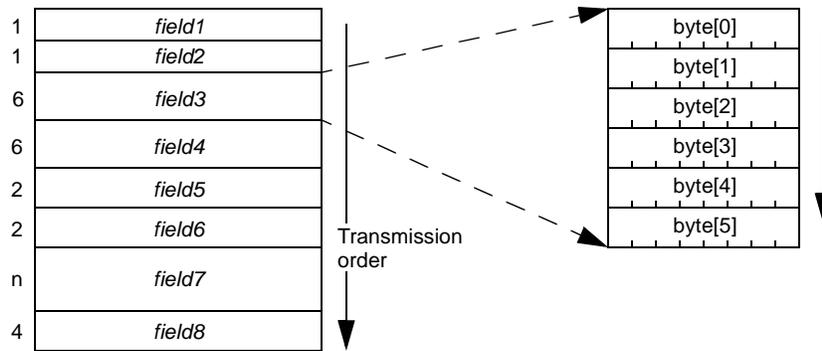


Figure 3.3—Byte sequential field format illustrations

NOTE—Only the left-hand diagram in Figure 3.3 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.

3.10 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 3.4. 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.

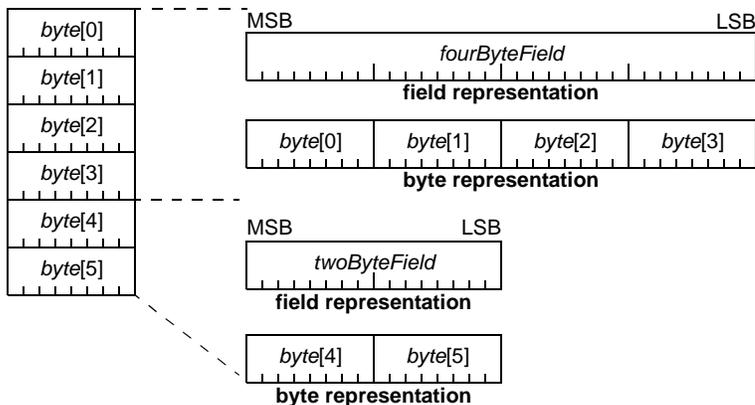


Figure 3.4—Multibyte field illustrations

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.

NOTE—The following text was taken from 802.17, where it was found to have benefits:
The details should, however, be revised to illustrate fields within an AVB frame header *serviceDataUnit*.

To minimize potential for confusion, four equivalent methods for illustrating frame contents are illustrated in Figure 3.5. Binary, hex, and decimal values are always shown with a left-to-right significance order, regardless of their bit-transmission order.

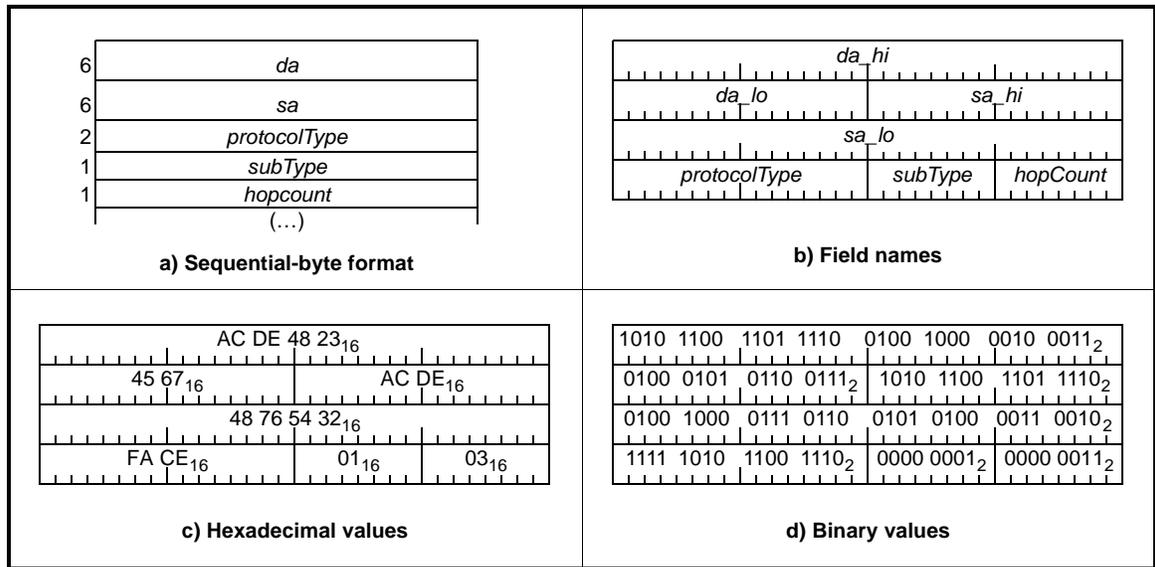


Figure 3.5—Illustration of fairness-frame structure

3.11 MAC address formats

The format of MAC address fields within frames is illustrated in Figure 3.6.

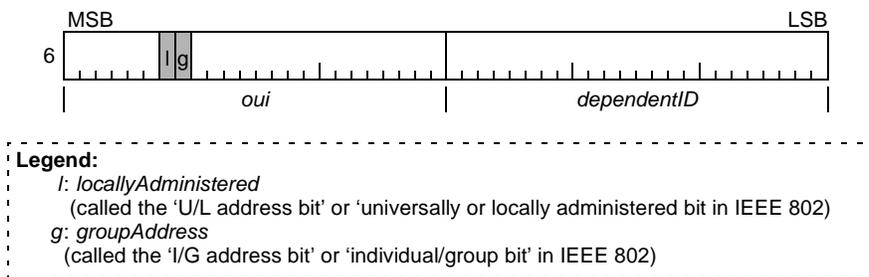
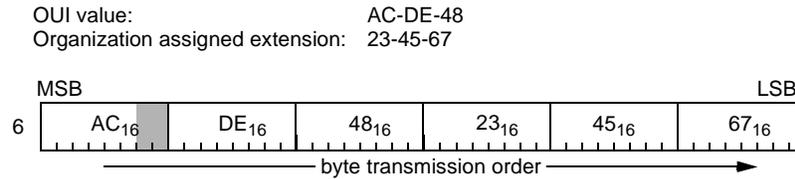


Figure 3.6—MAC address format

3.11.1 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.)

1 **3.11.2 dependentID:** An 24-bit field supplied by the *oui*-specified organization. The concatenation of the *oui*
2 and *dependentID* provide a unique (within this context) identifier.

3
4 To reduce the likelihood of error, the mapping of OUI values to the *oui/dependentID* fields are illustrated in
5 Figure 3.7. For the purposes of illustration, specific OUI and *dependentID* example values have been
6 assumed. The two shaded bits correspond to the *locallyAdministered* and *groupAddress* bit positions illus-
7 trated in Figure 3.6.



10
11
12
13
14
15
16 **Figure 3.7—48-bit MAC address format**

17
18
19 **3.12 Informative notes**

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

23
24
25 NOTE—This is an example of an informative note.

26
27 **3.13 Conventions for C code used in state machines**

28
29 Many of the state machines contained in this working paper utilize C code functions, operators, expressions
30 and structures for the description of their functionality. Conventions for such C code can be found in
31 Annex E.

4. Abbreviations and acronyms

NOTE—This clause should be skipped on the first reading (continue with Clause 5).
This text has been lifted from the P802.17 draft standard, which has a relative comprehensive list.
Abbreviations/acronyms are expected to be added, revised, and/or deleted as this working paper evolves.

This working paper contains the following abbreviations and acronyms:

AP	access point
AV	audio/video
AVB	audio/video bridging
AVB network	audio/video bridged network
BER	bit error ratio
BMC	best master clock
BMCA	best master clock algorithm
CRC	cyclic redundancy check
FIFO	first in first out
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
ISO	International Organization for Standardization
ITU	International Telecommunication Union
LAN	local area network
LSB	least significant bit
MAC	medium access control
MAN	metropolitan area network
MSB	most significant bit
OSI	open systems interconnect
PDU	protocol data unit
PHY	physical layer
PLL	phase-locked loop
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.

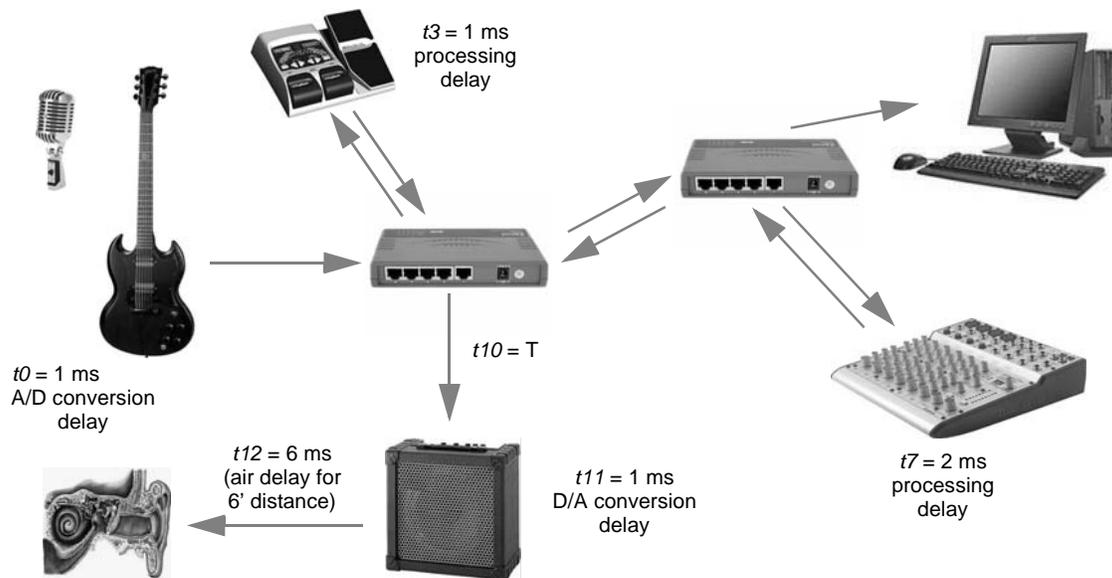


Figure 5.1—Garage jam session

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 overrun/underrun at the 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.

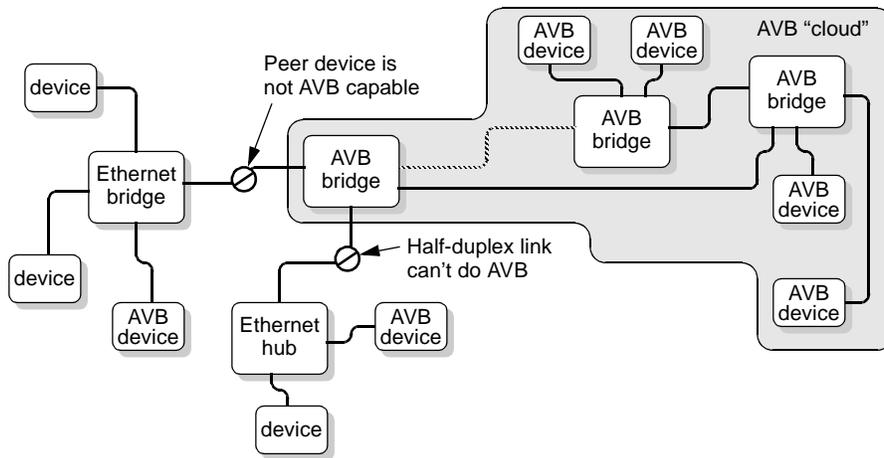


Figure 5.2—Possible looping topology

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 *timeOfDay* 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) Feed-forward. PLLs are unnecessary within bridges, but possible within applications.
 - 3) Frequent. Frequent (nominally 100 Hz) interchanges reduces needs for overly precise clocks.

5.3 Time-synchronization facilities

In concept, the clock-synchronization protocol starts with the selection of the reference-timer station, called a grand-master station (oftentimes abbreviated as grand-master). Every AVB-capable station is grand-master capable, but only one is selected to become the grand-master station within each network. To assist in the grand-master selection, each station is associated with a distinct preference value; the grand-master is the station with the “best” preference values. Thus, time-synchronization services involve two subservices, as listed below and described in the following subclauses.

- a) Selection. Looping topologies are isolated (from a time-synchronizatin perspective) into a spanning tree. The root of the tree, which provides the time reference to others, is the grand master.
- b) Distribution. Synchronized time is distributed through the the grand-master’s spanning tree.

5.3.1 Grand-master selection

5.3.1.1 Grand-master responsibilities

Clock synchronization involves streaming of timing information from a grand-master timer to one or more slave timers. Although primarily intended for non-cyclical physical topologies (see Figure 5.3a), the synchronization protocols also function correctly on cyclical physical topologies (see Figure 5.3b), by activating only a non-cyclical subset of the physical topology.

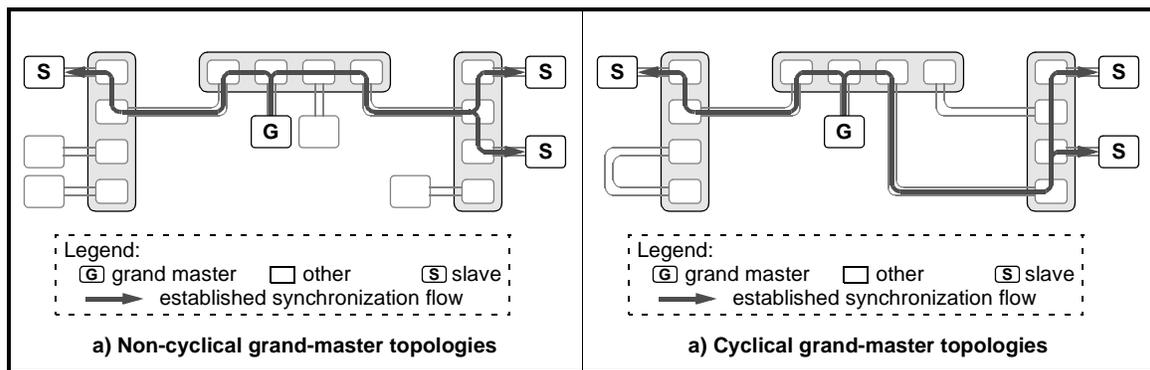


Figure 5.3—Timing information flows

5.3.1.2 Grand-master behavior

As part of the grand-master selection process, stations forward the best of their observed preference values to neighbor stations, allowing the overall best-preference value to be ultimately selected and known by all. The station whose preference value matches the overall best-preference value ultimately becomes the grand-master.

The grand-master station observes that its precedence is better than values received from its neighbors, as illustrated in Figure 5.4a. A slave stations observes its precedence to be worse than one of its neighbors and forwards the best-neighbor precedence value to adjacent stations, as illustrated in Figure 5.4b. To avoid cyclical behaviors, a *hopsCount* value is associated with preference values and is incremented before the best-preference value is communicated to others.

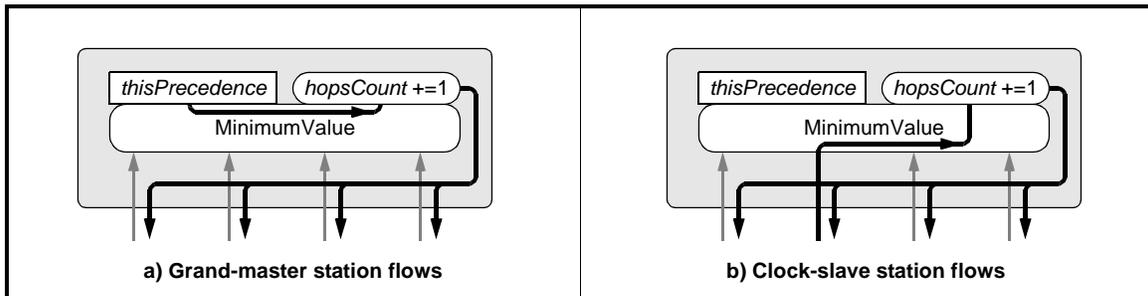


Figure 5.4—Grand-master precedence flows

5.3.1.3 Grand-master precedence

Grand-master precedence is based on the concatenation of multiple fields, as illustrated in Figure 5.5. The *port* value is used within bridges, but is not transmitted between stations.

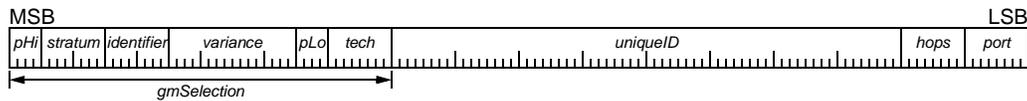


Figure 5.5—Grand-master precedence

This format is similar to the format of the spanning-tree precedence value, but a wider *uniqueID* is provided for compatibility with interconnects based on 64-bit station identifiers.

5.3.2 Synchronized-time distribution

Clock-synchronization information conceptually flows from a grand-master station to clock-slave stations, as illustrated in Figure 5.6a. A more detailed illustration shows pairs of synchronized clock-master and clock-slave components, as illustrated in Figure 5.6b. The active clock agents are illustrated as black-and-white components; the passive clock agents are illustrated as grey-and-white components.

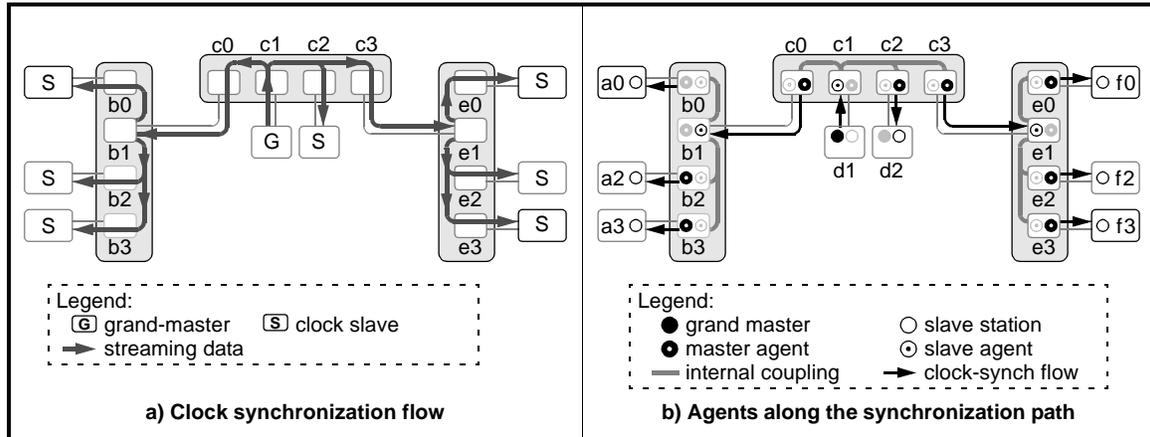


Figure 5.6—Hierarchical flows

Internal communications distribute synchronized time from clock-slave agents b1, c1, and e1 to the other clock-master agents on bridgeB, bridgeC, and bridgeE respectively. Within a clock-slave, precise time synchronization involves adjustments of timer value and rate-of-change values.

Time synchronization yields distributed but closely-matched *timeOfDay* values within stations and bridges. No attempt is made to eliminate intermediate jitter with bridge-resident jitter-reducing phase-lock loops (PLLs,) but application-level phase locked loops (not illustrated) are expected to filter high-frequency jitter from the supplied *timeOfDay* values

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

5.3.3 MAC service model

The MAC service model assumes the presence of one or more time-synchronized AVB ports communicating with a higher-level MAC client, as illustrated in Figure 5.7. The receive portion of each port provides grand-master precedence information, which assists in the grand-master selection process. The transmit portion of each port provides the higher-level client with a local time and (in response) receives status containing the grand-master precedence and the client's synchronized version of the grand-master time. All components are assumed to have access to a common free-running (not adjustable) local timer. There is not necessarily a one-to-one correspondence between the primitives and formal procedures and the interfaces in any particular implementation.

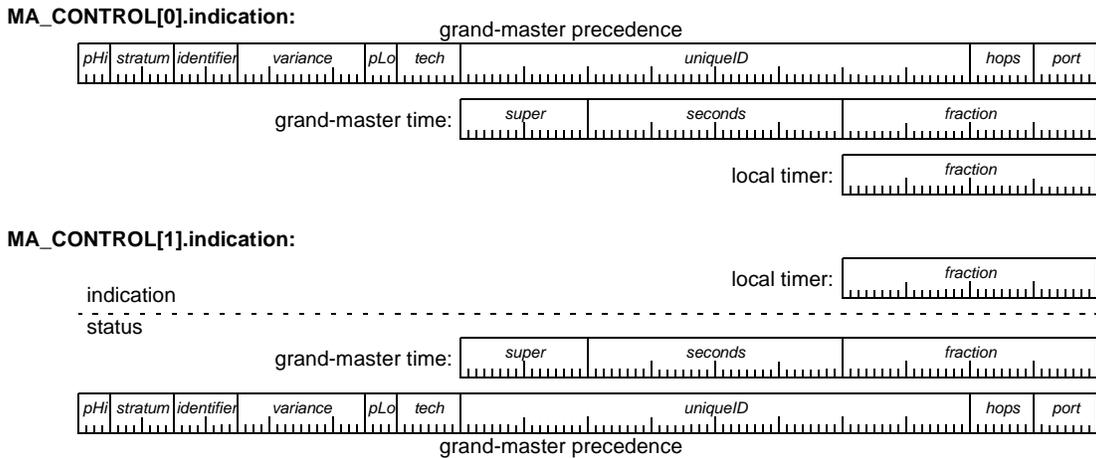
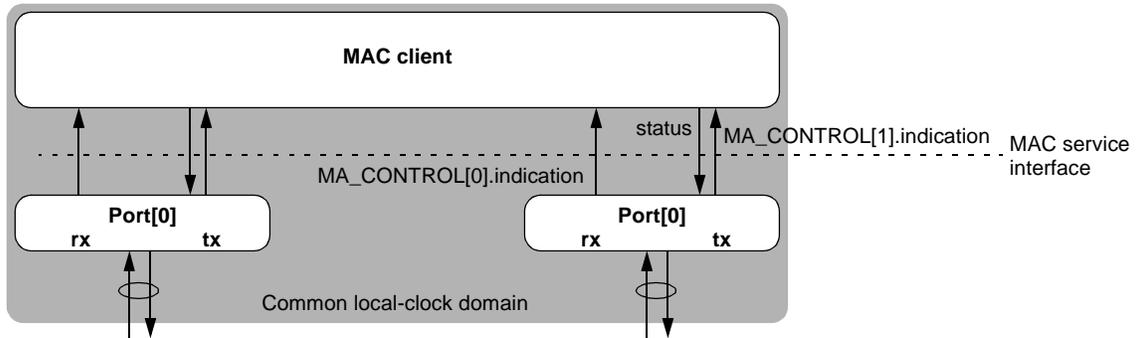


Figure 5.7—AVB service interface model

5.4 Client processing details

TBD.

5.5 Key distinctions from IEEE Std 1588

Although based on the concepts of IEEE Std 1588, this draft is different in multiple ways:

- a) All bridges are boundary clocks, since they compensate their received time by pass-through delays, by setting the client time on received frames, then transmitting the current client time when frames are transmitted. There are no transparent (in 1588 terminology) bridges.
- b) The processing by the client is distinguished by the processing performed in the MAC, thus isolating the client from that multiple (and sometimes strange) wireless and PON MACs protocols.
- c) To simplify computations, time is uniformly represented as a simple 80-bit scaled signed integer.
- d) For Ethernet, a higher update frequency of 100 Hz is assumed. This reduces timeouts for failed grand masters, and worst-case times for clear the network of rogue packets, while also reducing timer-value drifts between updates.
- e) For Ethernet, only one frame type simplifies the protocols and reduces transient-recovery times.

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

6. Medium access control (MAC) service and reference model

6.1 Overview

This clause provides an overview of the MAC sublayer and the reconciliation sublayer (see shaded portions of Figure 6.1), specifies the services provided by the MAC sublayer, including the MAC control sublayer and the MAC datapath sublayer, and provides a reference model for the MAC sublayer. Higher-layer clients can include the logical link control (LLC) sublayer, bridge relay entity, or other users of ISO/IEC LAN international standard MAC services. The services 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.

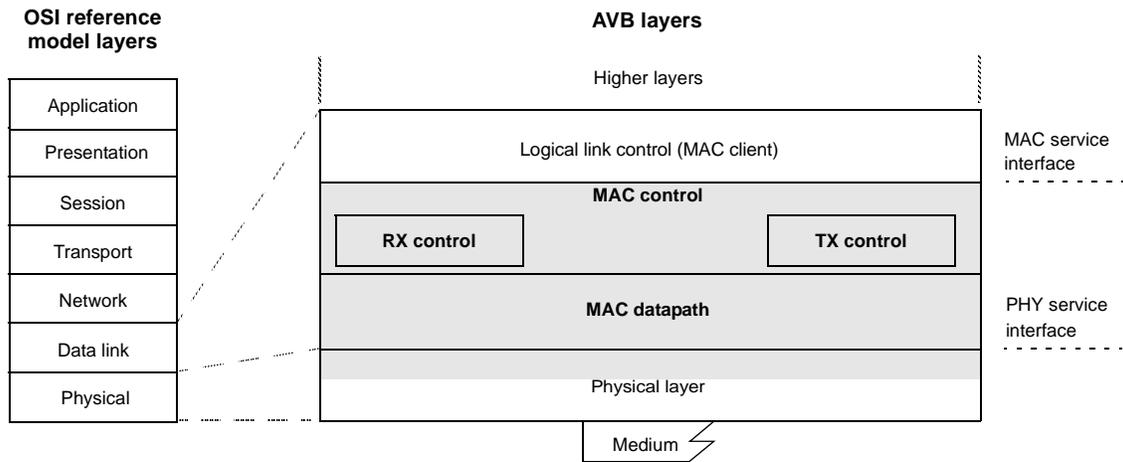


Figure 6.1—AVB service and reference model relationship to the ISO/IEC OSI reference model

6.2 Overview of MAC services

The services provided by the MAC sublayer are listed below. For each port, these services are invoked periodically at modest processing rates (100Hz, for Ethernet MACs).

- a) The receive port provides grand-master precedence and time, as well as local time, to the client.
- b) The transmit port provides the client with a recent-past local time and (in response) the client provides the grand-master precedence and time as status.

6.3 MAC services to the client layer

The services of a layer or sublayer are the set of capabilities that it offers to the next higher (sub)layer. The services specified in this standard are described by abstract service primitives and parameters that characterize each service. This definition of a service is independent of any particular implementation.

The following two service primitives are defined for the client interfaces, and shall be implemented.

- MA_CONTROL[0].indication
- MA_CONTROL[1].indication

6.3.1 MA_CONTROL[0].indication

6.3.1.1 Function

The MA_CONTROL[0].indication primitive provides the client with an association of local time and the grand-master time (as received from the closer-to grand-master station), as well as the grand-master precedence.

6.3.1.2 Semantics of the service primitive

The semantics of the primitives are as follows:

```
MA_CONTROL[0].indication
(
    local_time,
    gm_precedence,
    gm_time
)
```

The parameters of the MA_CONTROL[0].indication are described below.

local_time

Provides the value of a local time when the value of gm_time was sampled.

gm_precedence

Provides the recently observed grand-master precedence, or NULL if such as precedence has not been received within a physical-layer dependent timeout interval. Fields within this component include the following (more-significant fields are listed first):

pHi	A 4-bit user-selectable overriding precedence.
stratum	An 8-bit identifier that identifies the type of grand-master.
identifier	An 8-bit identifier that identifies ...
variance	A 16-bit value that identifies the grand-master clock quality.
pLo	A 4-bit user-selectable tie-breaking precedence.
tech	An 8-bit value that, when appended with the uniqueID, is globally unique.
uniqueID	An 8-bit value that, when prepended with the tech, is globally unique.
hops	An 8-bit value that counts visible repeaters from the grand master.
port	An 8-bit value that identifies the port that received the gm_precedence value.

gm_time

Provides the value of the distributed grand-master time, compensated for intermediate delays.

Fields within this component include the following (more-significant fields are listed first):

super	A signed 16-bit value representing 2^{32} seconds.
seconds	An unsigned 32-bit value representing seconds.
fraction	An unsigned 32-bit value representing fractions-of-seconds.

6.3.1.3 When generated

The MA_CONTROL[0].indication primitive is invoked by the MAC entity whenever new knowledge of time or the grand-master precedence is generated.

6.3.1.4 Effect of receipt

The receipt of the MA_CONTROL[0].indication primitive causes the client entity to adjust its image of grand-master time and rate. Details are TBD.

6.3.2 MA_CONTROL[1].indication**6.3.2.1 Function**

The MA_CONTROL[1].indication primitive provides the client with a recent value of local time; the client returns the associated grand-master time (as synchronized to the closer-to grand-master station), as well as the grand-master precedence.

6.3.2.2 Semantics of the service primitive

The semantics of the primitives are as follows:

```
MA_CONTROL[1].indication
```

```
(
    local_time,
)
```

```
MA_CONTROL[1].status
```

```
(
    gm_precedence,
    gm_time
)
```

The parameters of the MA_CONTROL[1].indication are described below.

```
local_time
```

Provides the value of a local time when the value of gm_time was sampled.

The returned values for the MA_CONTROL[1].indication are described below.

```
gm_precedence
```

Provides the client's recently observed best grand-master precedence. See xx for details.

```
gm_time
```

Provides the value of the distributed grand-master time, compensated for intermediate delays. See xx for details.

6.3.2.3 When generated

The MA_CONTROL[1].indication primitive is periodically invoked by the MAC entity whenever new knowledge of time and/or the grand-master precedence is needed to update attached clock-slave neighbors.

6.3.2.4 Effect of receipt

The receipt of the MA_CONTROL[0].indication primitive causes the client entity to return the desired status values.

7. Ethernet duplex-cable time synchronization

7.1 Synchronized time-of-day clocks

7.1.1 Assumptions

Support of duplex-link Ethernet is 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) Point-to-point transmit/receive duplex connections are provided.
- b) Transmit/receive propagation delays within duplex cables are well matched.

7.1.2 Strategies

Strategies used to meet these objectives include the following:

- a) Precision is achieved by calibrating and adjusting *timeOfDay* 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) Symmetric. Clock-master/clock-slave computations are similar (only slave results are saved).
 - 2) Periodic. Messages are sent periodically, rather than in timely response to other requests.
 - 3) Frequent. Frequent (typically 1 kHz) interchanges reduces needs for precise clocks.
- c) Balanced functionality.
 - 1) Low-rate. Complex computations are infrequent and can be readily implemented in firmware.
 - 2) High-rate. Frequent computations are simple and can be readily implemented in hardware.

7.1.3 Clock-synchronization updates

7.1.3.1 Clock-synchronization intervals

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

Clock synchronization involves the processing of periodic events. Multiple time periods are involved, as listed in Table 7.1. The clock-period events trigger the update of free-running timer values; the period affects the timer-synchronization accuracy and is therefore constrained to be small.

Table 7.1—Clock-synchronization intervals

Name	Time	Description
clock-period	< 20 ns	Resolution of timer-register value updates
send-period	10 ms	Time between sending of periodic timeSync frames between adjacent stations
slow-period	100 ms	Time between computation of clock-master/clock-slave rate differences

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

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

The rate-synchronization protocols interchange parameters periodically, but less frequently than the offset-synchronization protocols, possibly every 100 ms.

7.1.4 Rate-difference effects

If the absence of rate adjustments, significant *timeOfDay* errors can accumulate between send-period updates, as illustrated on the left side of Figure 7.1. The 2 μ s deviation is due to the cumulative effect of clock drift, over the 10 ms send-period interval, assuming clock-master and clock-slave crystal deviations of -100 PPM and +100 PPM respectively.

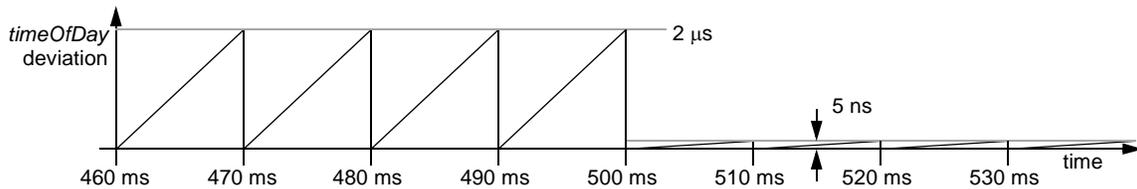


Figure 7.1—Rate-adjustment effects

While this regular sawtooth is illustrated as a highly regular (and thus perhaps easily filtered) function, irregularities could be introduced by changes in the relative ordering of clock-master and clock-slave transmissions, or transmission delays invoked by asynchronous frame transmissions. Tracking peaks/valleys or filtering such irregular functions are thought unlikely to yield similar *timeOfDay* deviation reductions.

The differences in rates could easily be reduced to less than 1 PPM, assuming a 200 ms measurement interval (based on a 100 ms slow-period interval) and a 100 ns arrival/departure sampling error. A clock-rate adjustment at time 100 ms could thus reduce the clock-drift related errors to less than 5 ns. At this point, the timer-offset measurement errors (not clock-drift induced errors) dominate the clock-synchronization error contributions.

7.1.5 Timer snapshot locations

Mandatory jitter-error accuracies are sufficiently loose to allow transmit/receive snapshot circuits to be located with the MAC, as illustrated in Figure 7.2a. Vendors may elect to further reduce timing jitter by latching the receive/transmit times within the PHY, where the uncertain FIFO latencies can be best avoided.

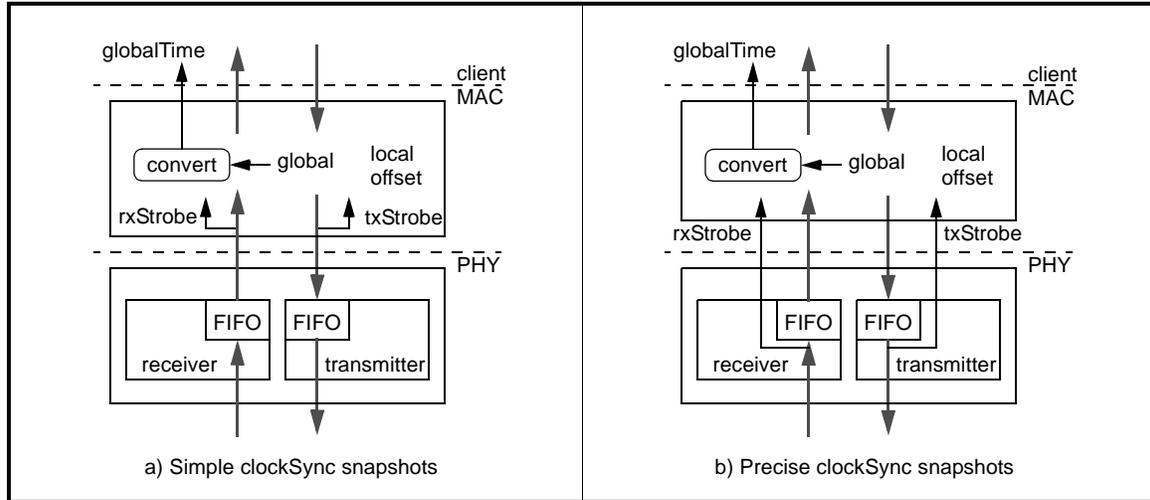


Figure 7.2—Timer snapshot locations

7.2 timeSync frame format

7.2.1 timeSync fields

Clock synchronization (timeSync) frames facilitate the synchronization of neighboring clock span-master and clock span-slave stations. The frame, which is normally sent once each isochronous cycle, includes time-snapshot information and the identity of the network’s clock master, as illustrated in Figure 7.3. The gray boxes represent physical layer encapsulation fields that are common across Ethernet frames.

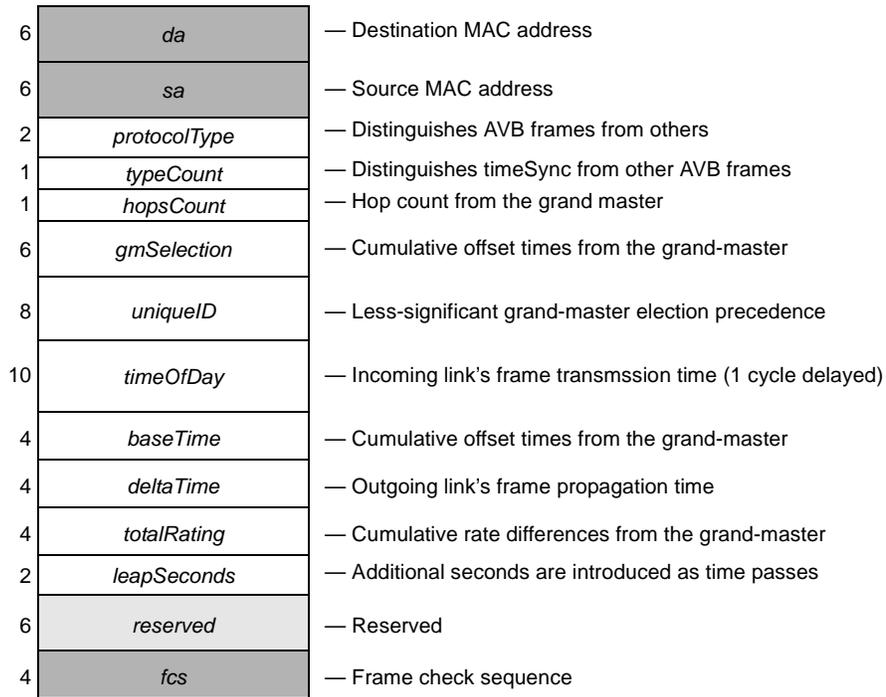


Figure 7.3—timeSync frame format

7.2.1.1 *da*: A 48-bit (destination address) field that specifies the station(s) for which the frame is intended. The *da* field contains either an individual or a group 48-bit MAC address (see ?3.11?), as specified in 9.2 of IEEE Std 802-2001.

7.2.1.2 *sa*: A 48-bit (source address) field that specifies the local station sending the frame. The *sa* field contains an individual 48-bit MAC address (see ?3.11?), as specified in 9.2 of IEEE Std 802-2001.

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

7.2.1.4 *typeCount*: An 8-bit field that identifies the format and function of the following fields (see 7.2.2).

7.2.1.5 *hopsCount*: An 8-bit field that identifies the maximum number of hops between the talker and associated listeners.

7.2.1.6 *gmSelection*: A 6-byte field that has specifies precedence in the grand-master selection protocols (see 7.2.3).

7.2.1.7 *uniqueID*: A 64-bit field that uniquely identifies the contending grand clock master (see ?7.1.3?).

1 **7.2.1.8 *timeOfDay*:** An 96-bit field that specifies the time within the source station when the previous
2 *timeSync* frame was transmitted (see ?7.1.4?).

3
4 **7.2.1.9 *baseTime*:** A 32-bit field that specifies the fractions-off-second offset time within the source station
5 (see ?7.1.4?).

6
7 **7.2.1.10 *deltaTime*:** A 48-bit field that specifies the differences between *timeSync* receive and transmit
8 times, as measured in fractions-of-second on the opposing link (see ?7.1.4?).

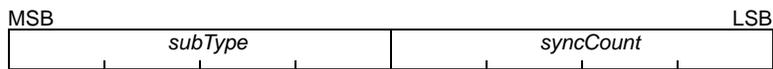
9
10 **7.2.1.11 *totalRating*:** A 32-bit field that specifies a scaled version of the *diffRate* value. The *diffRate* value
11 transmitted by station *n* represents the rate difference between station *n* and the grand master.

12
13 **7.2.1.12 *leapSeconds*:** A 16-bit field that specifies the number of seconds that should be added to the
14 *timeOfDay* value, when converting between *xx* and *yy* values. (On IEEE-1588, this is called the *UTCOffset*
15 field.)

16
17 **7.2.1.13 *fcs*:** A 32-bit (frame check sequence) field that is a cyclic redundancy check (CRC) of the frame.

18
19 **7.2.2 *typeCount* subfields**

20
21 The 8-bit *typeCount* field provides a subtype identifier and a sequence number for detecting lost-frame
22 transmissions, as illustrated in Figure 7.4.



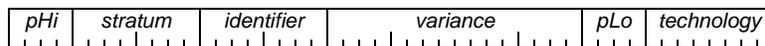
23
24
25
26
27
28 **Figure 7.4—*typeCount* format**

29
30 **7.2.2.1 *subCount*:** A 4-bit field that distinguishes the *timeSync* frame from other frames with the same
31 *protocolType* field.

32
33 **7.2.2.2 *syncCount*:** A 4-bit field that is incremented on each *timeSync*-frame transmission.

34
35 **7.2.3 *gmSelection* subfields**

36
37 The 48-bit *gmSelection* field provides precedence information of the current grand master, as illustrated in
38 Figure 7.5.



39
40
41
42
43 **Figure 7.5—*gmSelection* format**

44
45 **7.2.3.1 *pHi*:** A 4-bit field that can be configured by the user and overrides the remaining
46 *gmSelection*-resident precedence fields.

47
48 **7.2.3.2 *stratum*:** An 8-bit precedence-selection field defined by the like-named IEEE-1588 field.

49
50 **7.2.3.3 *identifier*:** An 8-bit precedence-selection field defined by the like-named IEEE-1588 field.

51
52 **7.2.3.4 *variance*:** A 16-bit precedence-selection field defined by the like-named IEEE-1588 field.

7.2.3.5 pLo: A 4-bit field that can be configured by the user and overrides the remaining *gmSelection*-resident precedence fields.

7.2.3.6 technology: An 8-bit field that identifies the format of the 8-byte *uniqueID* field.

7.2.4 uniqueID subfields

The 64-bit *uniqueID* field is a unique identifier. For stations that have a uniquely assigned 48-bit *macAddress*, the 64-bit *uniqueID* field is derived from the 48-bit MAC address, as illustrated in Figure 7.6.

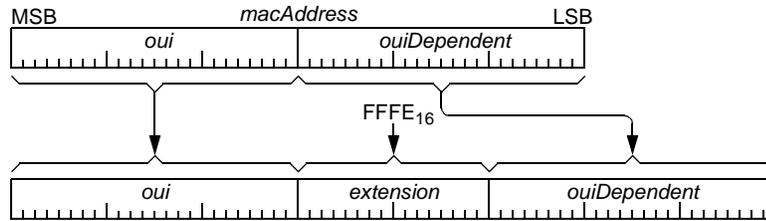


Figure 7.6—uniqueID format

7.2.4.1 oui: A 24-bit field assigned by the IEEE/RAC (see xx).

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

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

7.2.5 timeOfDay subfield formats

Time-of-day values 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 7.7.

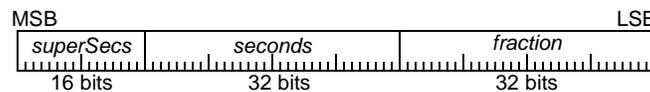


Figure 7.7—Complete seconds timer format

7.2.5.1 superSecs: A 16-bit field that extends the range of the *seconds* field.

7.2.5.2 seconds: A 32-bit field that specifies time in seconds.

7.2.5.3 fraction: A 32-bit field that specified time offset within the *second*, in units of 2^{-32} second.

The concatenation of these fields specifies a 96-bit *timeOfDay* value, as specified by Equation 7.1.

$$time = superSecs * 2^{32} + seconds + (fraction / 2^{32}) \tag{7.1}$$

1 **7.3 Per-port processing details**

2
3 TBD.

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)

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.

B.1 Hybrid network topologies

B.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 B.1. The boundary between domains is illustrated by a dotted line, which passes through a SerialBus adapter station.

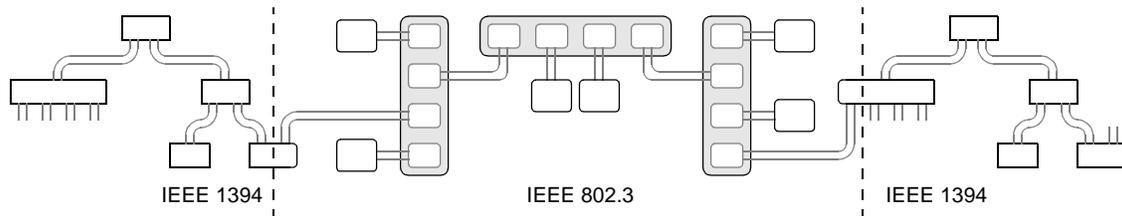


Figure B.1—IEEE 1394 leaf domains

B.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 B.2. While not explicitly prohibited, architectural features of such topologies are beyond the scope of this working paper.

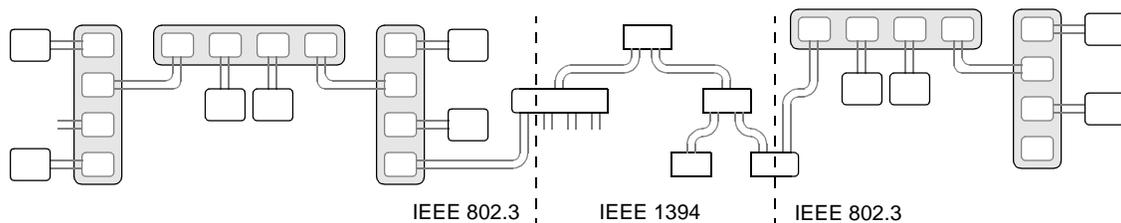


Figure B.2—IEEE 802.3 leaf domains

B.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 B.3.

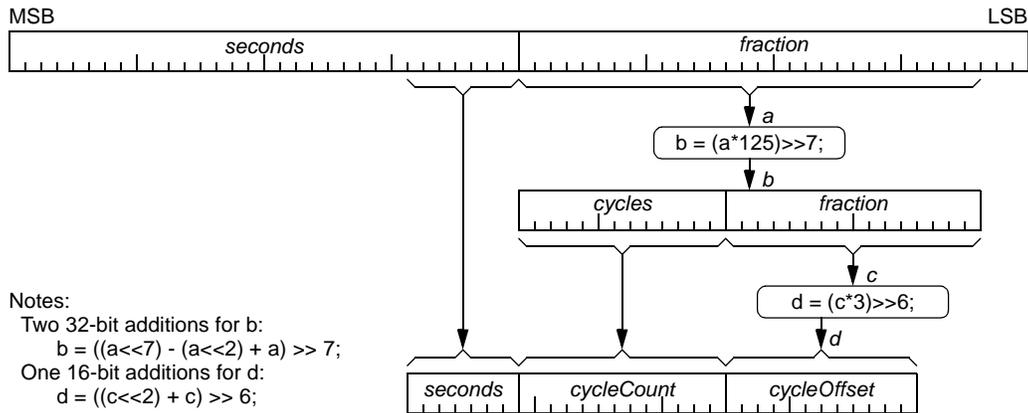


Figure B.3—Time-of-day format conversions

B.1.4 Grand-master precedence mappings

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

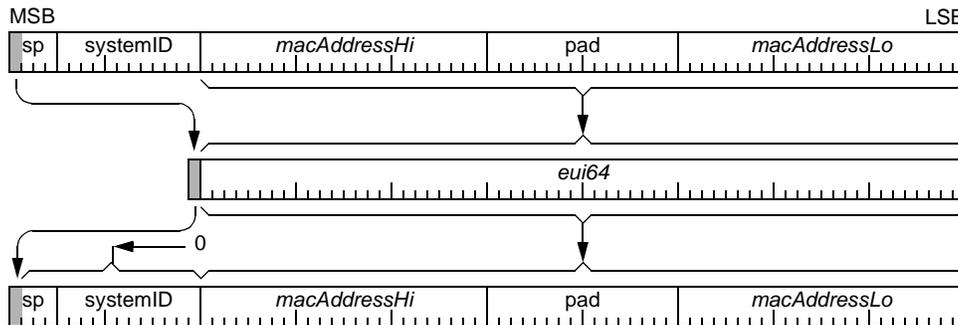


Figure B.4—Grand-master precedence mapping

Annex C

(informative)

Review of possible alternatives

C.1 Clock-synchronization alternatives

NOTE—This tables has not been reviewed for considerable time and is thus believed to be inaccurate. However, the list is being maintained (until it can be updated) for its usefulness as talking points.

A comparison of the AVB and IEEE 1588 time-synchronization proposals is summarized in Table C.1.

Table C.1 Protocol comparison

Properties		Row	Descriptinos	
state			AVB-SG	1588
timeSync MTU <= Ethernet MTU		1	yes	
No cascaded PLL whiplash		2	yes	
Number of frame types		3	1	> 1
Phaseless initialization sequencing		4	yes	no
Topology		5	duplex links	general
Grand-master precedence parameters		6	spanning-tree like	special
Rogue-frame settling time, per hop		7	10 ms	1 s
Arithmetic complexity	numbers	8	64-bit binary	2 x 32-bit binary
	negatives	9	2's complement	signed
Master transfer discontinuities	rate	10	gradual change	
	offset limitations	11	duplex-cable match sampling error	
Firmware friendly	no delay constraints	12	yes	
	n-1 cycle sampling	13	yes	
Time-of-day value precision	offset resolution	14	233 ps	
	overflow interval	15	136 years	

Row 1: The size of a timeSync frame should be no larger than an Ethernet MTU, to minimize overhead.

AVB-SG: The size of a timeSync frame is an Ethernet MTU.

1588: The size of a timeSync frame is (to be provided).

Row 2: Cascaded phase-lock loops (PLLs) can yield undesirable whiplash responses to transients.

AVB-SG: There are no cascaded phase-lock loops.

1588: There are multiple initialization phases (to be provided).

Row 3: There number of frame types should be small, to reduce decoding and processing complexities.	1
AVB-SG: Only one form of timeSync frame is used.	2
1588: Multiple forms of timeSync frames are used (to be provided).	3
	4
Row 4: Multiple initialization phases adds complexity, since miss-synchronized phases must be managed.	5
AVB-SG: There are no distinct initialization phases.	6
1588: There are multiple initialization phases (to be provided).	7
	8
Row 5: Arbitrary interconnect topologies should be supported.	9
AVB-SG: Topologies are constrained to point-to-point full-duplex cabling.	10
1588: Supported topologies include broadcast interconnects.	11
	12
Row 6: Grand-master selection precedence should be software configurable, like spanning-tree parameters.	13
AVB-SG: Grand-master selection parameters are based on spanning-tree parameter formats.	14
1588: Grand-master selection parameters are (to be provided).	15
	16
Row 7: The lifetime of rogue frames should be minimized, to avoid long initialization sequences.	17
AVB-SG: Rogue frame lifetimes are limited by the 10 ms per-hop update latencies.	18
1588: Rogue frame lifetimes are limited by (to be provided).	19
	20
Row 8: The time-of-day formats should be convenient for hardware/firmware processing.	21
AVB-SG: The time-of-day format is a 64-bit binary number.	22
1588: The time-of-day format is a (to be provided).	23
	24
Row 9: The time-of-day negative-number formats should be convenient for hardware/firmware processing.	25
AVB-SG: The time-of-day format is a 2's complement binary number.	26
1588: The time-of-day format is a (to be provided).	27
	28
Row 10: The rate discontinuities caused by grand-master selection changes should be minimal.	29
AVB-SG: Smooth rate-change transitions with a 2.5 second time constant is provided.	30
1588: (To be provided).	31
	32
Row 11: The time-of-day discontinuities caused by grand-master selection changes should be minimal.	33
AVB-SG: Maximum time-of-day errors are limited by cable-length asymmetry and time-snapshot errors.	34
1588: (To be provided).	35
	36
	37
Row 12: Firmware friendly designs should not rely on fast response-time processing.	38
AVB-SG: Response processing time have no significant effect on time-synchronization accuracies.	39
1588: (To be provided).	40
	41
Row 13: Firmware friendly designs should not rely on immediate or precomputed snapshot times.	42
AVB-SG: Snapshot times are never used within the current cycle, but saved for next-cycle transmission.	43
1588: (To be provided).	44
	45
Row 14: The fine-grained time-of-day resolution should be small, to faciliate accurate synchronization.	46
AVB-SG: The 64-bit time-of-day timer resolution is 233 ps, less than expected snapshot accuracies.	47
1588: (To be provided).	48
	49
Row 15: The time-of-day extent should be suffiently large to avoid overflows within one's lifetime.	50
AVB-SG: The 64-bit time-of-day timer overflows once every 136 years.	51
1588: (To be provided).	52
	53
	54

Annex D

(informative)

Time-of-day format considerations

To better understand the rationale behind the ‘extended binary’ timer format, other formats are evaluated and compared within this annex.

D.1 Possible time-of-day formats

D.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 32-bit *seconds* and 32-bit *fraction* fields, as illustrated in Figure 4.1.



Figure 4.1—Complete seconds timer format

The concatenation of 32-bit *seconds* and 32-bit *fraction* field specifies a 64-bit *time* value, as specified by Equation D.1.

$$time = seconds + (fraction / 2^{32}) \tag{D.1}$$

Where:

- seconds* is the most significant component of the time value (see Figure 4.1).
- fraction* is the less significant component of the time value (see Figure 4.1).

D.1.2 IEEE 1394 timer format

An alternate “1394 timer” format consists of *secondCount*, *cycleCount*, and *cycleOffset* fields, as illustrated in Figure D.2. For such fields, the 12-bit *cycleOffset* field is updated at a 24.576MHz rate. The *cycleOffset* field goes to zero after 3171 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.

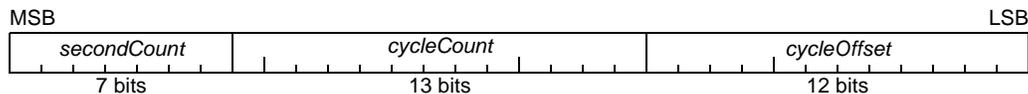


Figure D.2—IEEE 1394 timer format

D.1.3 IEEE 1588 timer format

IEEE 1588 timer format consists of seconds and nanoseconds fields components, as illustrated in Figure D.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.

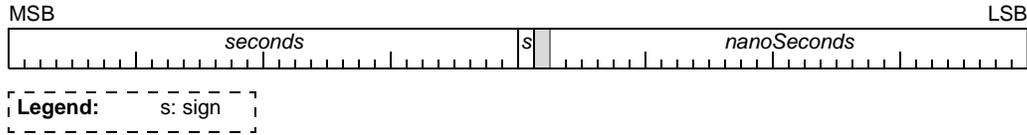


Figure D.3—IEEE 1588 timer format

D.1.4 EPON timer format

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

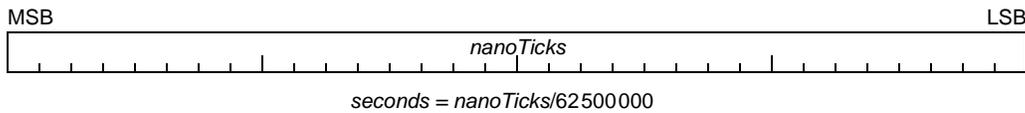


Figure D.4—EPON timer format

D.1.5 Compact seconds timer format

An alternate “compact seconds” format could consist of 8-bit *seconds* and 24-bit *fraction* fields, as illustrated in Figure D.5. This would provided similar resolutions to the IEEE 1394 timer format, without the complexities associated with its binary coded decimal (BCD) like encoding.



Figure D.5—Compact seconds timer format

D.1.6 Nanosecond timer format

An alternate “nanosecond” format could consists of 2-bit *seconds* and 30-bit *nanoSeconds* fields, as illustrated in Figure D.6.

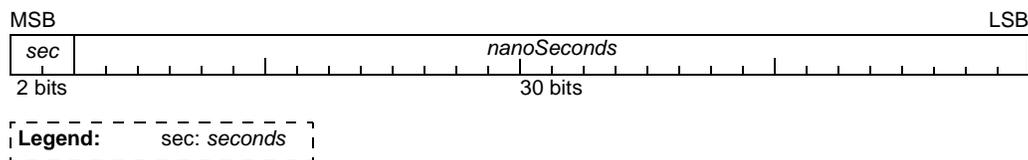


Figure D.6—Nanosecond timer format

D.2 Time format comparisons

To better understand the relative benefits of different time formats, the relevant properties are summarized in Table D.1. Counter complexity is not included in the comparison, since the digital logic complexity is comparable for all formats.

Table D.1—Time format comparison

Name	Subclause	Range	Precision	Arithmetic	Seconds	Defined standards
Column	—	1	2	3	4	5
extended binary	TBD	136 years	232 ps	Good	Good	RFC 1305 NTP, RFC 2030 SNTPv4
IEEE 1394	D.1.2	128 s	30 ns	Poor	Good	IEEE 1394
IEEE 1588	D.1.3	272 years	1 ns	Fair	Good	IEEE 1588
IEEE 802 (EPON)	D.1.4	69 s	16 ns	Good	Poor	IEEE 802.3
compact seconds	D.1.5	256 s	60 ns	Best	Good	—
nanoseconds	D.1.6	4 s	1 ns	Best	Poor	—

Column 1: A desirable property is the support of a wide range of second values, to eliminate the need for defining/coordinating/implementing auxiliary seconds-synchronization protocols. The 136-year range of the extended binary format is sufficient for this purpose.

Column 2: A desirable property is a fine-grained resolution, sufficient to measure each bit-transmission times. The ‘extended binary’ provides the most precision; exceeds the resolution of expected cost-effective time-capture circuits.

Column 3: Computation of time differences involves the subtraction of two timer-snapshot values. Subtraction of ‘extended binary’ numbers involving standard 64-bit binary arithmetic; no special field-overflow compensations are required. Only the less precise ‘compact seconds’ and nanoseconds formats are simpler, due to the reduced 32-bit size of the timer values.

Column 4: Time values must oftentimes be compared to externally provided values (e.g., timers extracted from GPS or stratum-clock sources). For these purposes, the availability of a seconds component is desired. The ‘extended binary’ format provides a seconds component that can be easily extracted or such purposes.