IEEE 1588 Version 2

Geoffrey M. Garner
Consultant

September 24, 2008

Outline - 1

- Acknowledgment
- Purpose of IEEE 1588
- Status and history of IEEE 1588
- New features for Version 2
- Overview of IEEE 1588 V2
- Application service interface
- Some basic PTP concepts and entities
- Synchronization basics
  - Delay request-response mechanism
  - Peer delay mechanism
- End-to-end transparent clocks

Outline - 2

ISPCS Ann Arbor'08

- Peer-to-peer transparent clocks
- Time format
- Architectural choices
- Best master selection
- PTP profiles and conformance
- General optional features
- State configuration options
- Compatibility requirements
- Transport specific field
- Security
- Transport of cumulative frequency offset information

Outline - 3

- This tutorial is an updated version of reference 5

Acknowledgment

- The author would like to acknowledge and thank John Eidson for having provided references 2 – 4, from which many of the slides used here were taken or adapted.

Purpose of IEEE 1588

- IEEE 1588 is a protocol designed to synchronize real-time clocks in the nodes of a distributed system that communicate using a network
  - It does not say how to use these clocks (this is specified by the respective application areas)

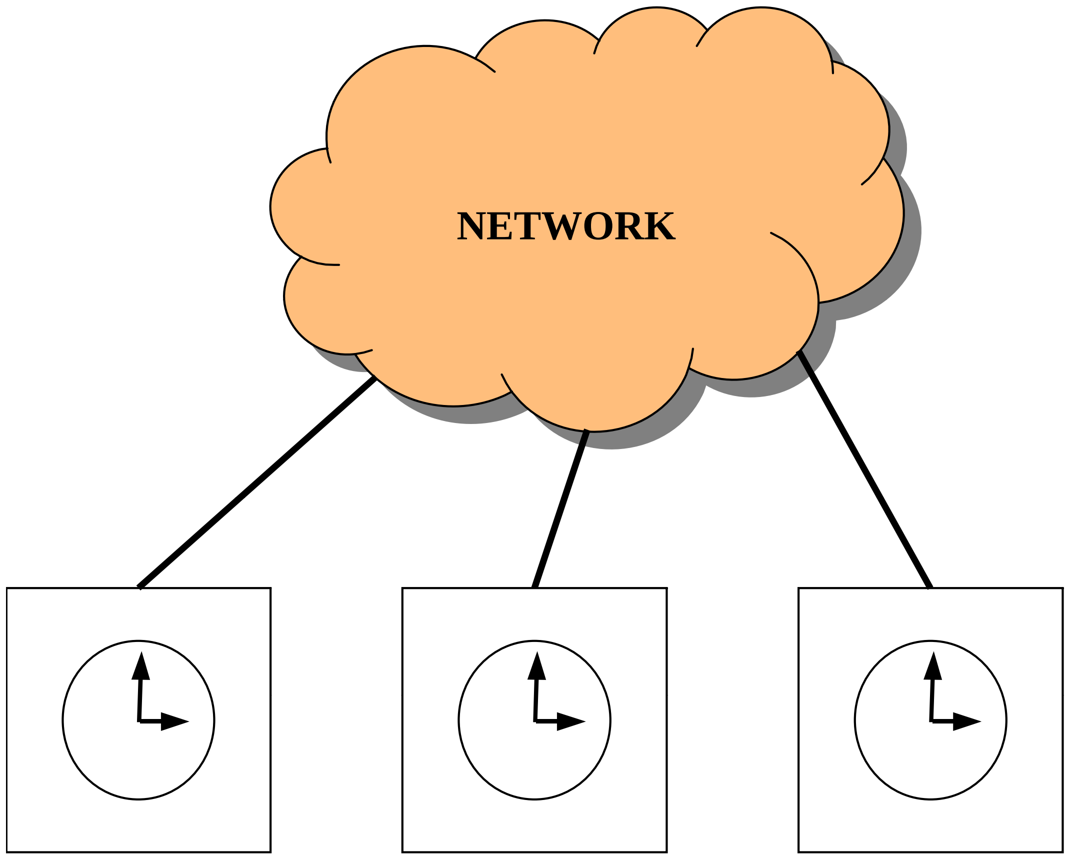

Status and History of IEEE 1588 - 1

ISPCS Ann Arbor'08

- Version 1 published as IEEE Std. 1588™ – 2002 on November 8, 2002
- Version 1 approved as IEC standard IEC 61588 on May 21, 2004
- V1 products and installations began appearing in late 2003
- Conferences on IEEE 1588 held 2003 – 2007
- Version 2 PAR approved March 20, 2005
- Version 2 technical work completed February 9, 2007
- Version 2 sponsor ballot opened July, 2007 and closed August 8, 2007

Status and History of IEEE 1588 - 2

- Version 2 sponsor ballot comment resolution and coordination with IEEE Registration Authority Committee (RAC) occurred during August – December, 2007
- Version 2 recirculation ballot occurred January 14 – 24, 2008
- P1588 committee voted on January 31, 2008 to send version to IEEE RevCom
- Version 2 approved by RevCom on March 26, 2008 and by IEEE Standards Board on March 27, 2008
- Version 2 published as IEEE Std 1588™ – 2008 on July 24, 2008 (reference 1)

Status and History of IEEE 1588 - 3

- Applications
  - Current: industrial automation, T&M, military, power generation and distribution
  - New: consumer electronics, telecommunications
- Products (V1, V2 mainly under development): microprocessors, GPS-linked clocks, boundary clocks, NIC cards, protocol stacks, RF instrumentation, aircraft flight monitoring instruments
- Referenced in: IEEE P802.1AS D3.0, PC37.1, IEEE 1646-2004, LXI Consortium, ODVA
  - IEEE 802.1AS will include a PTP profile (see reference 8 for details)

New Features for Version 2 - 1

- Mappings to UDP/IPv4&6, Ethernet (direct mapping), DeviceNet™, PROFINET, ControlNet™
- Formal mechanisms for message extensions (using TLV)
- Transparent clocks
- Synchronization accuracies better than 1 ns
- Options for redundancy and fault tolerance
- New management capabilities and options
- Higher sampling rates compared to V1; asymmetry corrections

New Features for Version 2 - 2

- Optional unicast messaging (in addition to multicast)
- PTP profiles
- Conformance specifications
- Configuration options
- Security (experimental specification only)
  – Covered very briefly in this tutorial; see reference 6 for more information
- Means to accumulate cumulative frequency scale factor offset relative to grandmaster (experimental specification only)

Overview of IEEE 1588 Version 2 - 1

ISPCS
Ann Arbor'08

Clause	Purpose	Clause	Purpose
1	Scope and purpose	8	Data sets
2	Standards references	9	PTP protocol for ordinary and boundary clocks
3	Definitions	10	PTP protocol for transparent clocks
4	Notation Convention	11	Clock offset, path delay, residence time, and asymmetry corrections
5	Data types	12	Synchronization and syntonization
6	Description of PTP protocol (Informative)	13	Message formats
7	PTP entities	14	TLV entities

Overview of IEEE 1588 Version 2 - 2

Clause	Purpose	Annex	Purpose
15	Management	A	Using the PTP protocol
16	General optional features	B	Timescales and epochs
17	State configuration options	C	Examples of residence and asymmetry corrections
18	Compatibility (V1/V2 and V2/future versions)	D	Transport over UDP/IPv4
19	Conformance and PTP profiles	E	Transport over UDP/IPv6
		F	Transport over IEEE 802.3/Ethernet

Overview of IEEE 1588 Version 2 - 3

Annex	Purpose	Annex	Purpose
G	Transport over DeviceNet	L	Transport of cumulative frequency scale factor offset (experimental)
H	Transport over ControlNet	M	Bibliography
I	Transport over IEC 61158 Type 10		
J	Default PTP profiles		
K	Security protocol (experimental)		

Application Service Interface

- This is outside the scope of IEEE 1588 (V1 and V2), and also outside the scope of this tutorial
  - However, it cannot be ignored
- The timing requirements (jitter, wander, time synchronization) of the respective applications that use the PTP clock must be met
- The quality of the clock delivered to the application depends on both the quality of the PTP clock and the application service interface
- See reference 7 for more information on this topic

Some Basic PTP Concepts and Entities - 1

- PTP Domain: A logical grouping of PTP clocks that synchronize to each other using the PTP protocol, but that are not necessarily synchronized to PTP clocks in another domain
  - Domain concept is carried over from V1
- PTP clock types (will cover in more detail later)
  - See Clause 3 of IEEE Std 1588™ – 2008 (reference 1) for exact wording of definitions

Some Basic PTP Concepts and Entities - 2

- PTP clock types (cont.)
  - Ordinary clock (OC): has a single PTP port in a domain and maintains the timescale used in the domain. It may serve as a source of time, i.e., be a master, or may synchronize to another clock, i.e., be a slave. It may provide time to an application or end device.
  - Boundary clock (BC): has multiple PTP ports in a domain and maintains the timescale used in the domain. It may serve as a source of time, i.e., be a master, or may synchronize to another clock, i.e., be a slave. It may provide time to an application.
    - A boundary clock that is a slave has a single slave port, and transfers timing from that port to the master ports

Some Basic PTP Concepts and Entities - 3

- PTP clock types (cont.)
  - Transparent clock (TC): A device that measures the time taken for a PTP event message to transit the device and provides this information to clocks receiving this PTP event message
    - Peer-to-peer transparent clock (P2P TC): A transparent clock that, in addition to providing PTP event transit time information, also provides corrections for the propagation delay of the link connected to the port receiving the PTP event message. In the presence of peer-to-peer transparent clocks, delay measurements between slave clocks and the master clock are performed using the peer delay measurement mechanism.

Some Basic PTP Concepts and Entities - 4

- PTP clock types (cont.)
    - End-to-end transparent clock (E2E TC): A device that only measures the time taken for a PTP event message to transit the device and provides this information to clocks receiving this PTP event message; i.e., it does not provide corrections for the propagation delay of the link connected to the port receiving the PTP event message. It does not use the peer delay measurement mechanism but, instead, supports use of the delay request-response mechanism

Some Basic PTP Concepts and Entities - 5

- PTP event messages (time stamped on egress from a node (clock) and ingress to a node)
  - Sync
  - Delay_Req
  - Pdelay_Req
  - Pdelay_Resp
- PTP general messages (not time stamped)
  - Follow_Up
  - Delay_Resp
  - Pdelay_Resp_Follow_Up
  - Announce
  - Management
  - Signaling

Some Basic PTP Concepts and Entities - 6

- PTP communication path: The signaling path portion of a particular network enabling direct communication among ordinary and boundary clocks
  - A communication path may contain transparent clocks
  - Cannot mix end-to-end and peer-to-peer transparent clocks on the same communication path (except in a restricted class of topologies, e.g., linear chain)

Synchronization Basics

ISPCS
Ann Arbor'08

- Step 1: Prior to synchronizing, organize the clocks into a master-slave hierarchy (based on observing the clock property information contained in Announce messages)
  - Announce messages carry information used to establish the master-slave hierarchy; they are not used for synchronization
- Step 2: Each slave synchronizes to its master using the delay request-response or peer delay mechanism, by exchanging messages with its master (and possibly with an upstream peer-to-peer transparent clock, if one is present)

Synchronization Basics – Delay Request-Response Mechanism

- Initially, consider the case where no transparent clocks are present
- Each slave synchronizes to its master using Sync, possibly Follow_Up, Delay_Req, and Delay_Resp messages exchanged between master and its slave

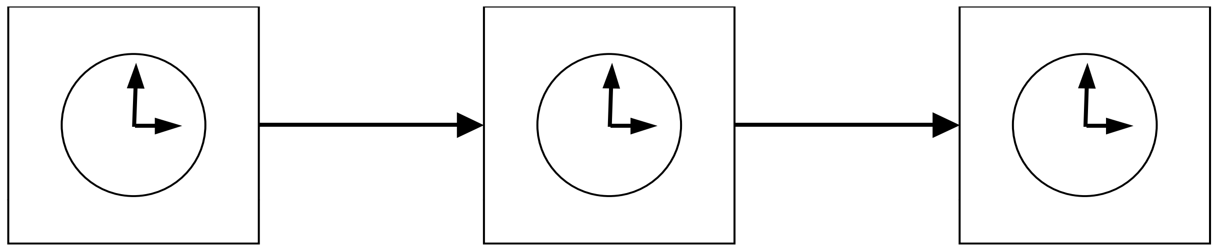

Grandmaster Clock This clock determines the time base for the system

Slave to the Grandmaster Clock and Master to its Slave

Slave to its Master

Synchronization Basics – Delay Request-Response Mechanism

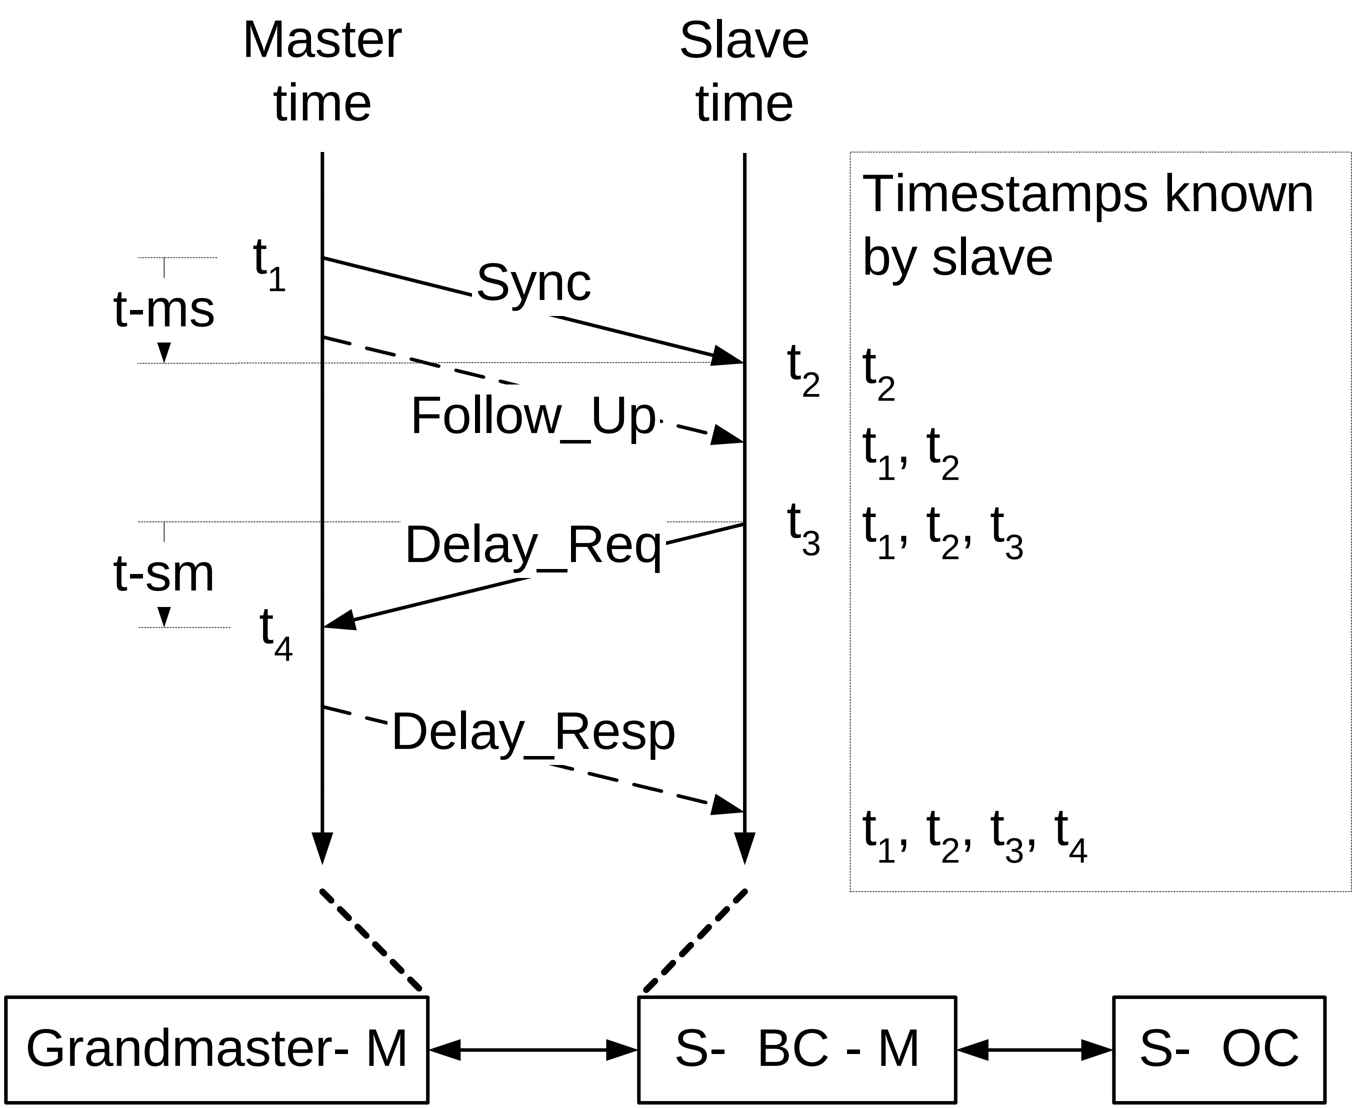

Synchronization Basics – Delay Request-Response Mechanism

ISPCS Ann Arbor'08

- Under the assumption that the link is symmetric (i.e., propagation time from master to slave = propagation time from slave to master)

Offset = (Slave time) – (Master time) = $[(t_2 - t_1) - (t_4 - t_3)]/2$ = [(t-ms) – (t-sm)]/2

(propagation time) = $[(t_2 - t_1) + (t_4 - t_3)]/2$

- Can rewrite the offset as

Offset = $t_2 - t_1$ – (propagation time) = (t-ms) – (propagation time)

- If the link is not symmetric
  - The propagation time computed as above is the mean of the master-to-slave and slave-to- master propagation times
  - The offset is in error by the difference between the actual master-to-slave and mean propagation times

Synchronization Basics – Delay Request-Response Mechanism

- As in 1588 V1, the delay request-response mechanism works on multipoint communication paths
  - Multicast Sync and Follow_Up messages (option for unicast)
  - Delay_Req from each slave sent to master; may be unicast
  - Separate Delay_Resp messages from master to each slave; may be unicast
- One-step clock: the Sync egress time stamp is carried in the Sync message, and Follow_Up is not sent
- Two-step clock: the Sync egress time stamp is carried in the Follow_Up message. This option is useful in cases where it is not possible both to achieve sufficient time stamp accuracy and place the time stamp in the message as it is transmitted.

Synchronization Basics – Peer Delay Mechanism - 1

- Initially, consider the case where no transparent clocks are present
- Each slave synchronizes to its master using Sync, possibly Follow_Up, Pdelay_Req, Pdelay_Resp, and possibly Pdelay_Resp_Follow_Up messages exchanged between master and its slave

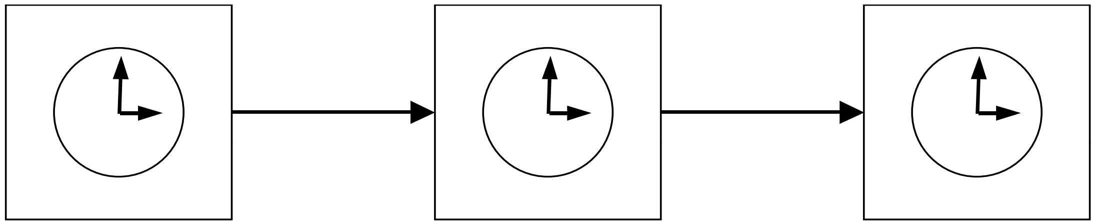

Grandmaster Clock This clock determines the time base for the system	Slave to the Grandmaster Clock and Master to its Slave	Slave to its Master

Synchronization Basics – Peer Delay Mechanism

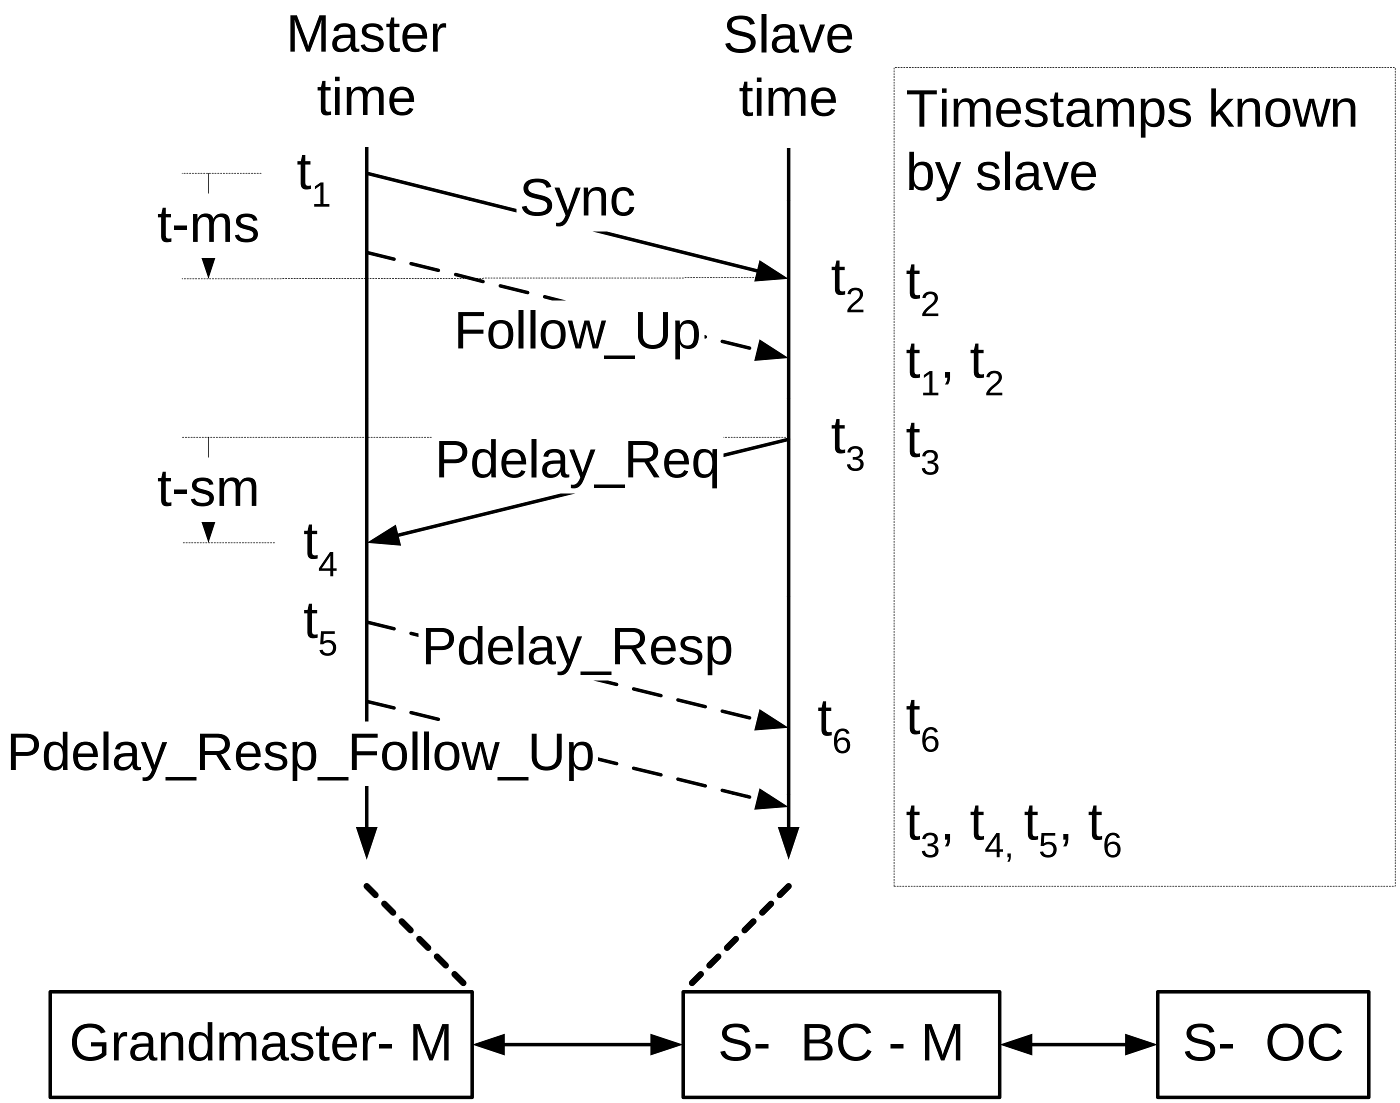

Synchronization Basics – Peer Delay Mechanism - 2

- Under the assumption that the link is symmetric (i.e., propagation time from master to slave = propagation time from slave to master)

  Offset = $t_2 - t_1$ – (propagation time) = (t-ms) – (propagation time)

  (propagation time) = $[(t_4 - t_3) + (t_6 - t_5)]/2$

- If the link is not symmetric
  - The propagation time computed as above is the mean of the master-to-slave and slave-to- master propagation times
  - The offset is in error by the difference between the actual master-to-slave and mean propagation times

Synchronization Basics – Peer Delay Mechanism - 3

- The peer delay mechanism is limited to point-to-point links between two ordinary clocks, boundary clocks, and/or peer-to-peer transparent clocks
  – This is because the protocol does not provide for the clock that receives Pdelay_Req to keep track of which clock it receives the message from, and respond separately to each clock
- The mechanism is symmetric, i.e., it operates separately in both directions on a link

Synchronization Basics – Peer Delay Mechanism - 4

- One-step clock: The Sync timestamp is carried in the Sync message, and Follow_Up is not sent. The Pdelay_Resp message carries the turnaround time, $t_5 – t_4$, and Pdelay_Resp_Follow_Up is not sent.

Synchronization Basics – Peer Delay Mechanism - 5

- Two-step clock: The Sync timestamp is carried in the Follow_Up message.
  - Either
    - (a) the turnaround time, $t_5 – t_4$, is carried in Pdelay_Resp_Follow_Up, or
    - (b) t_5 is carried in Pdelay_Resp_Follow_Up and t_4 is carried in Pdelay_Resp.
  - This option is useful in cases where it is not possible both to achieve sufficient time stamp accuracy and place the time stamp in the message as it is transmitted.

End-to-End Transparent Clocks - 1

- One alternative to placing an OC or BC at every node in a 1588 network is to use E2E TCs
- An E2E TC is not part of the master-slave hierarchy, i.e., it does not synchronize to the grandmaster (GM)
  - Instead, an E2E TC forwards Sync and corresponding Follow_Up messages on all ports not blocked by any underlying network protocols (e.g., rapid spanning tree protocol (RSTP) operating in an IEEE 802 bridged LAN)
- The E2E TC time stamps the Sync message on ingress and egress, and computes the time taken for the message to traverse the node (termed the residence time)

End-to-End Transparent Clocks - 2

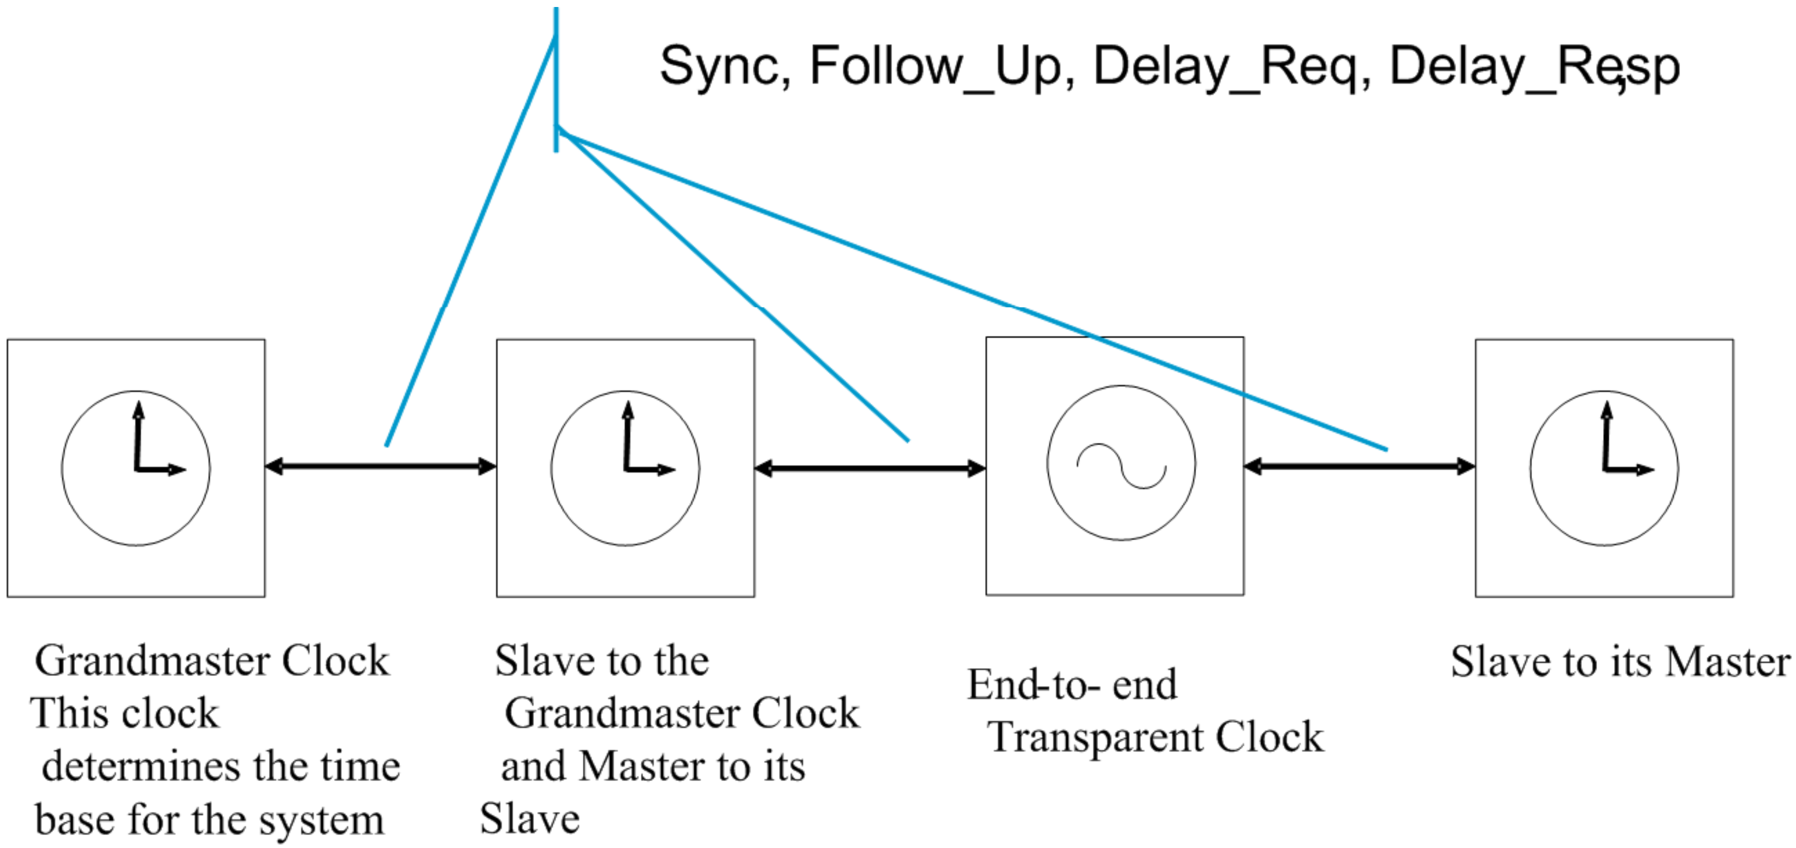

End-to-End Transparent Clocks - 3

- The residence time is accumulated in a field of the Sync (one-step clock) or Follow_Up (two-step clock) messages

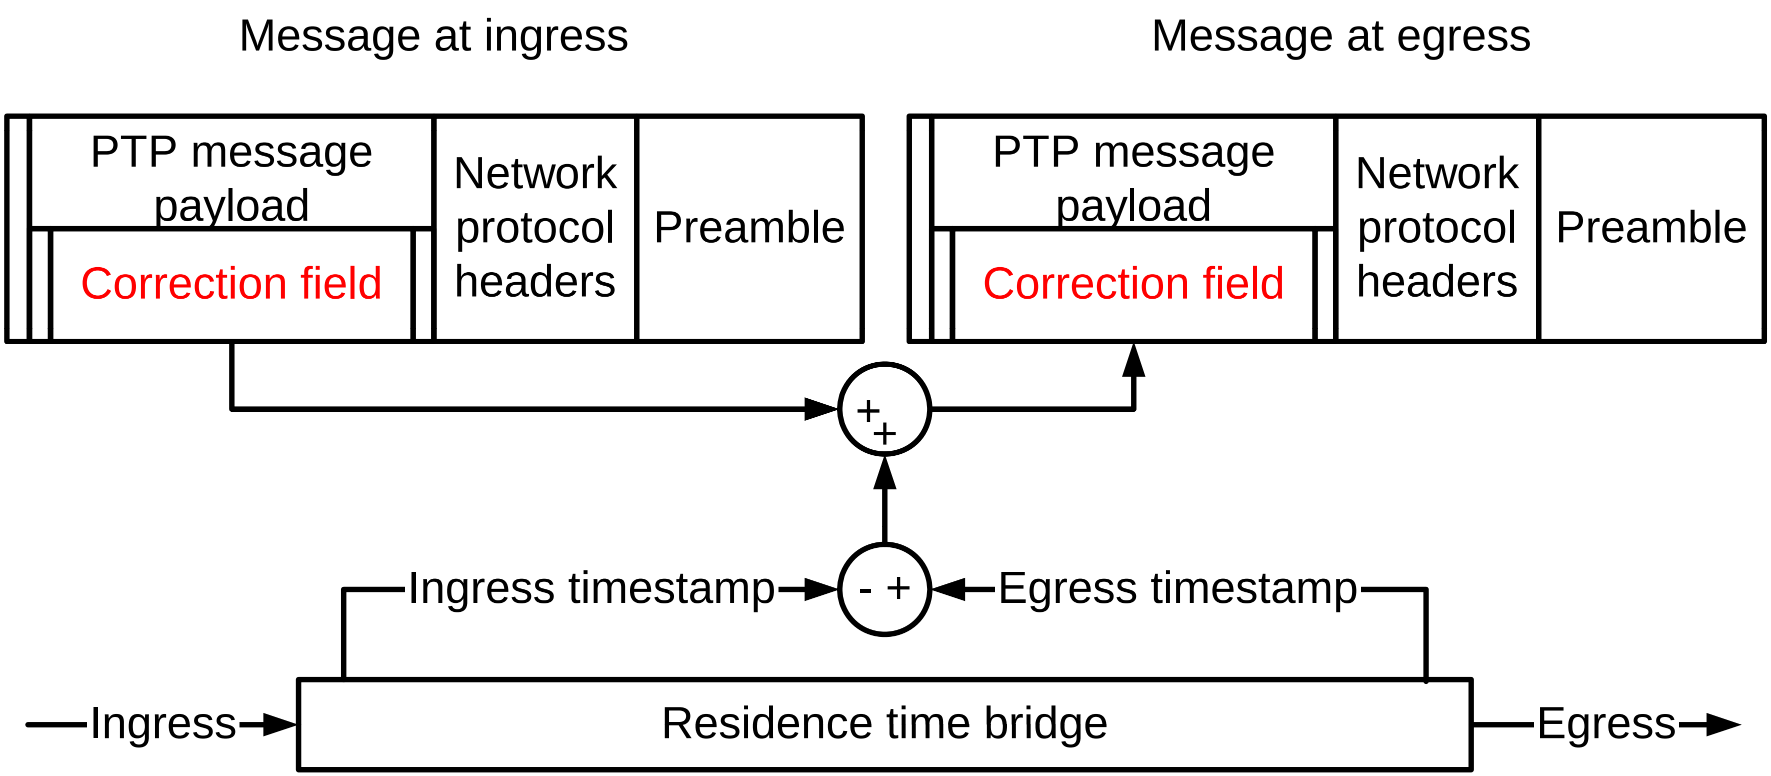

End-to-End Transparent Clocks - 4

- Each slave uses the cumulative residence time in the computation of propagation time from the master
  - In the previous formula, the correction fields of the Sync, Follow_Up, and Delay_Resp messages are subtracted in the numerator
  - The offset is given by t_2 minus t_1 minus propagation time minus Sync correction field minus Follow_Up correction field
- The residence time measurement does not require that the E2E TC be time-synchronized to the master
  - This is because the residence time is a time interval
- However, the residence time measurement will be in error if the rate of the E2E TC differs from the GM rate
  - The resulting synchronization error is equal to the fractional frequency offset multiplied by the residence time

End-to-End Transparent Clocks - 5

- If this error is unacceptably large, the E2E TC clock may optionally be syntonized, i.e., synchronized in frequency, to the GM
  - This may be done by using the egress time stamps (i.e., egress at master) and corrections fields in the successive Sync and Follow_Up messages, and the ingress time stamps of these Sync messages, to measure elapsed time of the master and corresponding elapsed time of the E2E TC
  - These two elapsed times may be used to obtain an estimate of the frequency offset of the TC relative to the master

End-to-End Transparent Clocks - 6

- IEEE 1588 does not specify how the information contained in the egress time stamps and correction fields of the successive Sync and Follow_Up messages, and in the ingress time stamps of those Sync messages, is to be used to accomplish the Syntonization
  - This is consistent with the fact that IEEE 1588 (V1 and V2) does not specify how the computed offset is used to synchronize OCs and BCs

End-to-End Transparent Clocks - 7

- Detailed specifications of whether to syntonize, and how, is the subject of PTP profiles and application specifications
  - Syntonization may be done in hardware, e.g., by physically adjusting the TC oscillator frequency, or in software, e.g., by using the measured frequency offset in the residence time computation

Peer-to-Peer Transparent Clocks - 1

- The E2E TC measures residence time of a Sync and Delay_Req message through the TC
  - However, propagation time between the master and slave is measured using the delay request-response mechanism
  - If the master changes, propagation time from the slave to the new master must be measured
  - This results in a longer transient during the network reconfiguration
- The P2P TC, in addition to measuring the residence time of the Sync message, measures path delay to the adjacent P2P TC, BC, or OC using the peer delay mechanism

Peer-to-Peer Transparent Clocks - 2

- The delay measurement is made on every link in both directions, including links that are blocked by lower layer network protocols (e.g., rapid spanning tree protocol)
  - This allows a delay measurement to be immediately available on links newly used after a change of GM or network reconfiguration
- The cumulative residence plus propagation time is accumulated in a field of the Sync (one-step clock) or Follow_Up (two-step clock) messages

Peer-to-Peer Transparent Clocks - 3

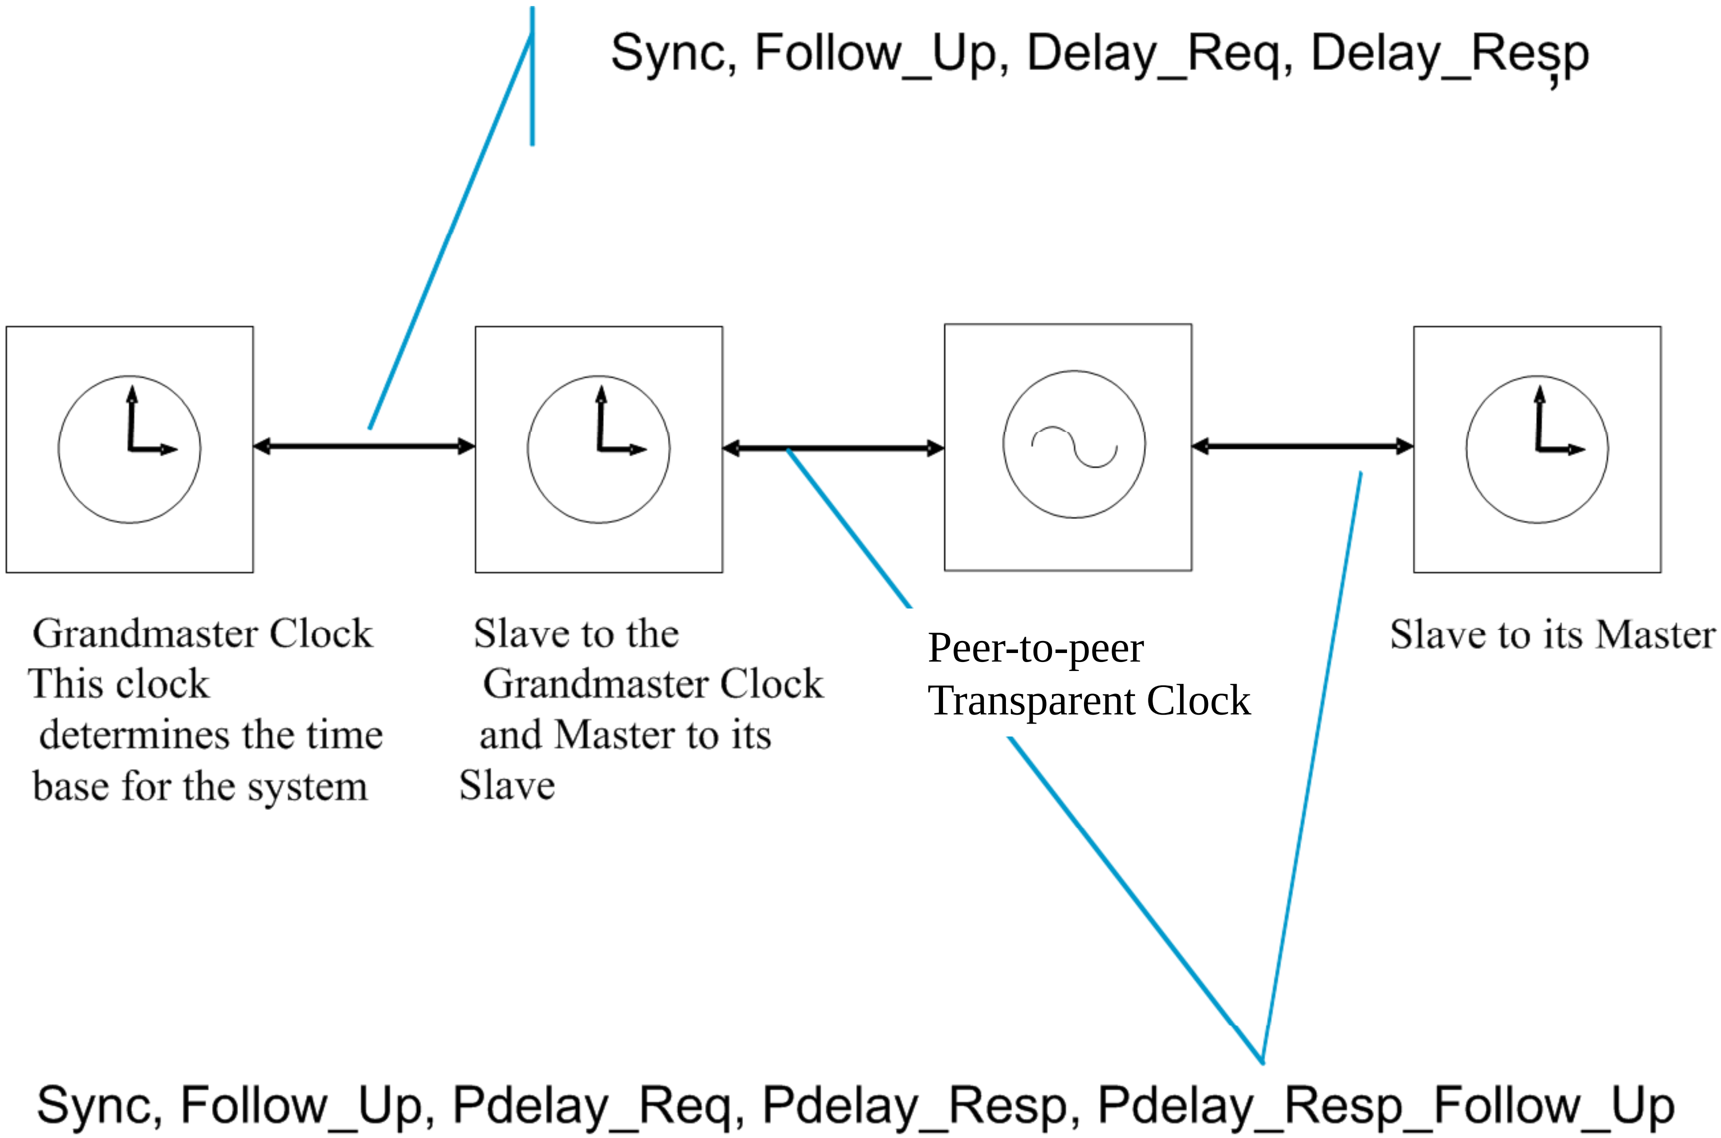

Peer-to-Peer Transparent Clocks - 4

- The offset between the slave and master is given by

  Offset = $t_2 - t_1$ – (Sync correction field) – (Follow_Up correction field) – (propagation time on final link to slave)

- As with E2E TCs, P2P TCs may optionally be syntonized to the master (in hardware or in software)
- If the P2P TCs are not syntonized to the GM, there will be an error in the measured propagation time equal to the sum of two terms:
  - turnaround time (Pdelay_Resp egress time minus Pdelay_Req ingress time) multiplied by the fractional frequency offset between the adjacent clocks that exchange Pdelay messages
  - Fractional frequency offset of the sender of Pdelay_Req relative to the master, multiplied by the measured propagation time

Peer-to-Peer Transparent Clocks - 4

- If the P2P TCs are not syntonized, the first term may be computed by measuring the frequency offset between the adjacent nodes using the egress and ingress time stamps of the successive Pdelay_Resp messages received by the sender of Pdelay_Req
  - The resulting propagation time will still be in error by an amount equal to the second term in this case
    - For many applications, this error (i.e., relative propagation delay error) may be sufficiently small (e.g., 10^{-4} or smaller, since the free-run frequency accuracy of PTP clocks is specified as $\pm$100 ppm (10^{-4}))

Time Format

```
struct Timestamp
{
    UInteger48 seconds;
    UInteger32 nanoseconds;
}

Integer64 correctionField;  /* nanoseconds × 2^16  */
```

Architectural choices - 1

- IEEE 1588 V2 allows for a variety of architectures, each based on choices of clock(s) and delay mechanism
- As stated earlier
  - P2P TCs use the peer delay mechanism
  - E2E TCs use the delay request-response mechanism
  - OCs can use either mechanism within a domain (but a vendor is not required to implement both mechanisms)
  - BCs can use either mechanism on any port within a domain (but a vendor is not required to implement both mechanisms on all ports)
  - The two delay mechanisms do not mix
  - P2P and E2E TCs cannot be mixed on the same communication path, except in very special cases, e.g., linear chains where each E2E TC has only 2 ports and no collocated OC

Architectural choices - 2

- Hierarchical topology example (figure taken from reference 1)
  - Uses BCs and OCs
  - The cyclic path (indicated by the dashed line) is blocked by a lower layer network protocol and appears to not be present to PTP

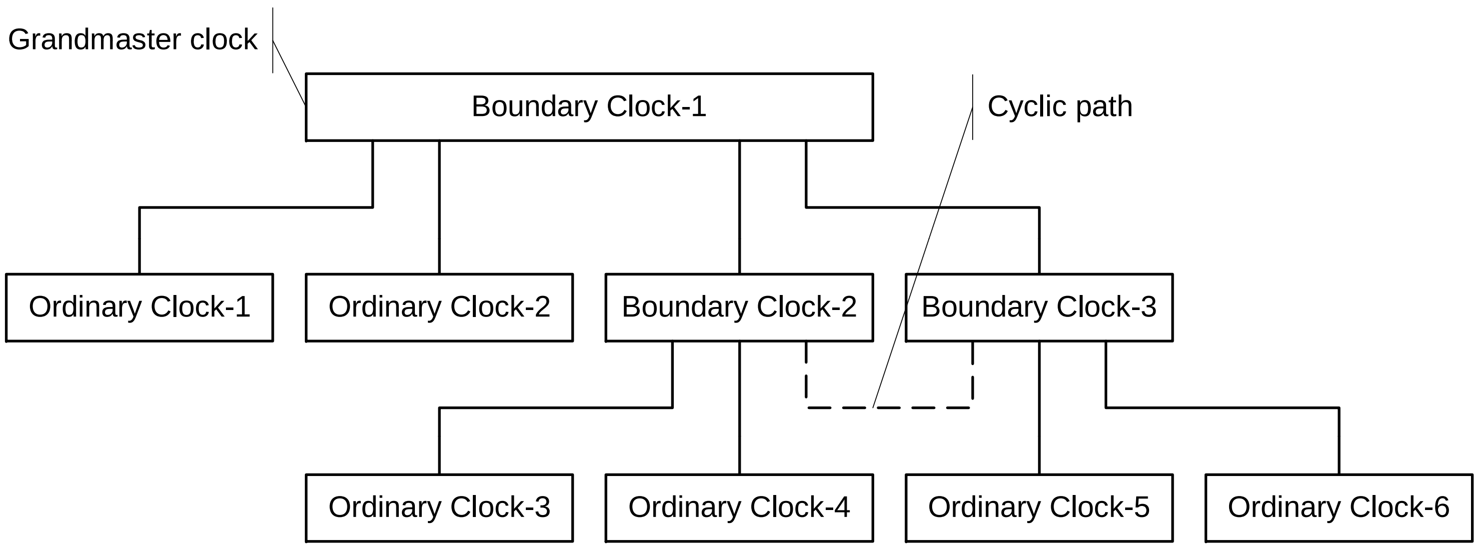

Architectural choices - 3

- Linear topology example (figure taken from reference 1)
  - Uses BC, OCs, and E2E TCs, arranged in linear chains
  - Propagation time measured using delay request-response mechanism separately between each OC and BC-1

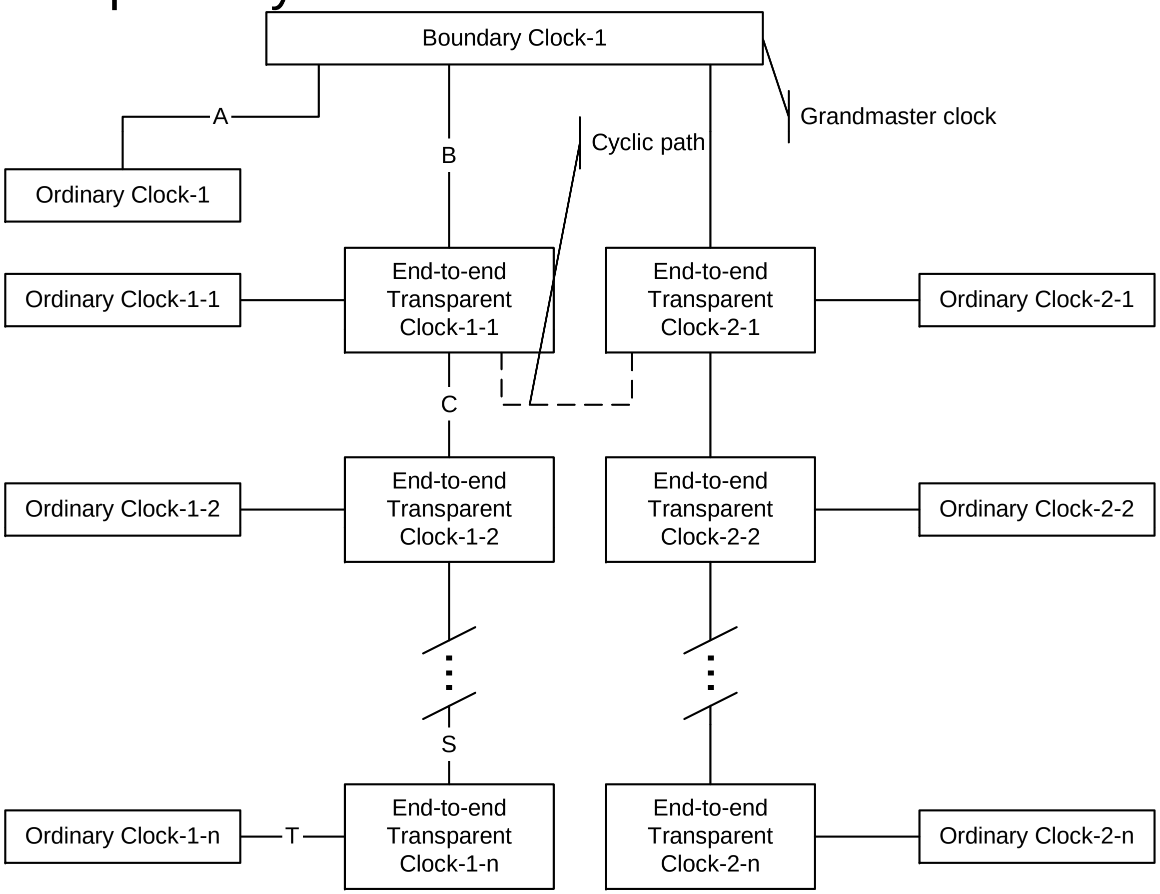

Architectural choices - 4

- Multiply connected topology example (figure taken from reference 1)
  - Uses OCs and P2P TCs, arranged in a mesh
  - Propagation time measured on each link using peer delay mechanism

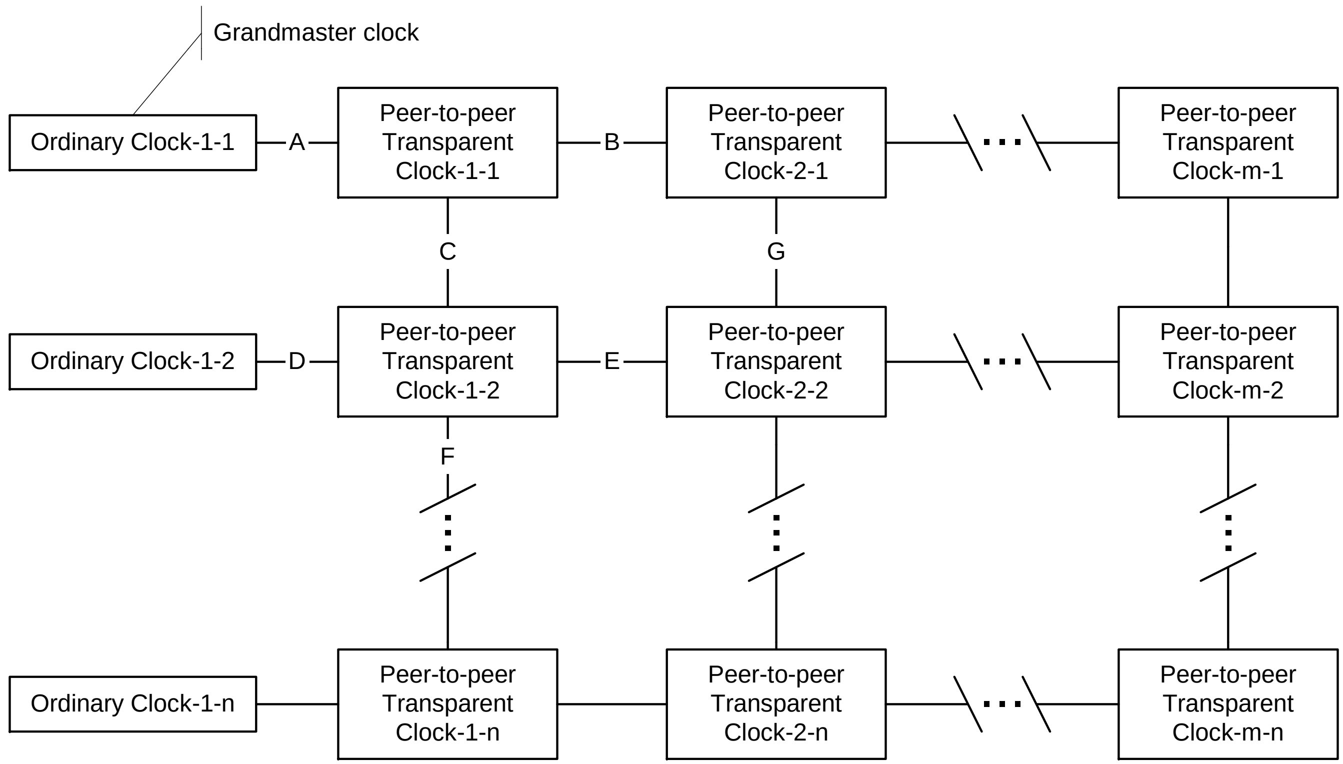

Architectural choices - 5

- Bridging disparate technologies using BC (figure taken from reference 1)
  - Links A use UDP/IPv4, links B use DeviceNet
  - BC-2 bridges the two technologies

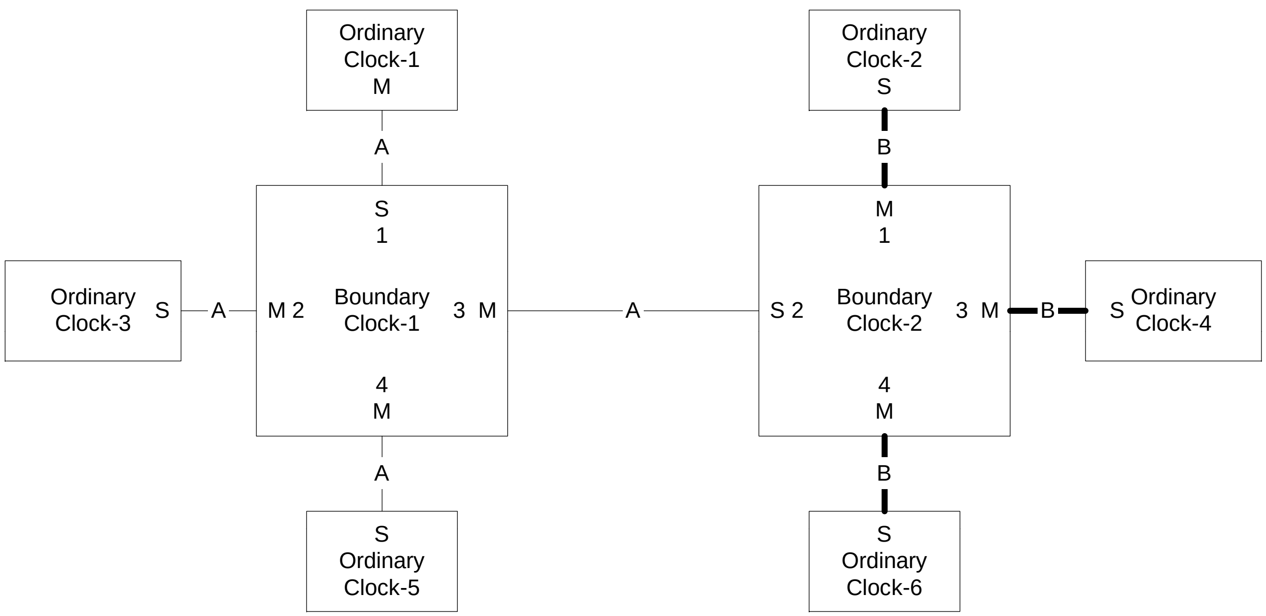

Best Master Selection - 1

- IEEE 1588 V2 uses a default Best Master Clock algorithm (BMCA) that is similar to the V1 BMCA
- Main differences
  - In V2, best master selection information is sent separately, in an Announce message, from synchronization-related information
    - This enables a much smaller Sync message, with higher Sync message rates using less bandwidth (important for some applications)
  - V1 stratum is renamed clock class (to distinguish from use of 'stratum' in telecom); many more of the 256 classes are defined and/or are usable (i.e., not reserved)

Best Master Selection - 2

- Main differences (cont.)
  - V1 clock identifier is replaced by clock accuracy, which indicates time accuracy without regard to clock technology
    - An information-only attribute, time source, that indicates the nature of the source of time, is added
  - Variance is replaced by (offset) scaled log variance, but is essentially the same
  - V1 notion of a clock being preferred for GM is replaced by 2 priority fields, each of which has 256 levels
    - Priority1 takes precedence over all other attributes
    - Priority2 is considered after class, clock accuracy, and scaled log variance

Best Master Selection - 3

- Main differences (cont.)
  - Whereas V1 would partition the network into multiple subnets if more than one potential GM of low enough stratum was present, V2 will not partition the network
    - instead, all clocks of higher class will synchronize to a single GM, and other potential GMs of low enough class (class 128 or lower) will free-run but not synchronize any other clocks

Best Master Selection - 4

Clock attributes (considered in the order given)

1. priority1(256 levels)
2. class (clockClass)
3. accuracy (clockAccuracy)
4. PTP variance (based on Allan variance) (offsetScaledLogVariance)
5. priority2 (256 levels)
6. identifier (IEEE EUI-64) (clockIdentity)

Notes: (i) PTP variance is equal to Allan variance multiplied by $\tau^2/3$, where τ is the sampling interval. (ii) clockIdentity may be formed from EUI-48 using IEEE mapping rules

Best Master Selection - 5

clockClass (decimal)	Specification
0	Reserved to enable compatibility with future versions.
1-5	Reserved
6	Shall designate a clock that is synchronized to a primary reference time source. The timescale distributed shall be PTP. A clockClass 6 clock shall not be a slave to another clock in the domain.
7	Shall designate a clock that has previously been designated as clockClass 6 but that has lost the ability to synchronize to a primary reference time source and is in holdover mode and within holdover specifications. The timescale distributed shall be PTP. A clockClass 7 clock shall not be a slave to another clock in the domain.
8	Reserved
9-10	Reserved to enable compatibility with future versions.
11-12	Reserved
13	Shall designate a clock that is synchronized to an application specific source of time. The timescale distributed shall be ARB. A clockClass 13 clock shall not be a slave to another clock in the domain.
14	Shall designate a clock that has previously been designated as clockClass 13 but that has lost the ability to synchronize to an application specific source of time and is in holdover mode and within holdover specifications. The timescale distributed shall be ARB. A clockClass 14 clock shall not be a slave to another clock in the domain.
15-51	Reserved
52	Degradation alternative A for a clock of clockClass 7 that is not within holdover specification. A clock of clockClass 52 shall not be a slave to another clock in the domain.
53-57	Reserved

Best Master Selection - 6

clockClass (decimal)	Specification
58	Degradation alternative A for a clock of clockClass 14 that is not within holdover specification. A clock of clockClass 58 shall not be a slave to another clock in the domain.
59-67	Reserved
68-122	For use by alternate PTP profiles
123-127	Reserved
128-132	Reserved
133-170	For use by alternate PTP profiles
171-186	Reserved
187	Degradation alternative B for a clock of clockClass 7 that is not within holdover specification. A clock of clockClass 187 may be a slave to another clock in the domain.
188-192	Reserved
193	Degradation alternative B for a clock of clockClass 14 that is not within holdover specification. A clock of clockClass 193 may be a slave to another clock in the domain.
194-215	reserved
216-232	For use by alternate PTP profiles

Best Master Selection - 7

clockClass (decimal)	Specification
233-247	Reserved
248	Default. This clockClass shall be used if none of the other clockClass definitions apply.
249-250	Reserved
251	Reserved for version 1 compatibility, see Clause 18.
252-254	Reserved
255	Shall be the clockClass of a slave-only clock, see 9.2.2.

Best Master Selection - 8

ISPCS Ann Arbor'08

Clock accuracy

Value (hex)	Specification
00-1F	reserved
20	The time is accurate to within 25 ns
21	The time is accurate to within 100 ns
22	The time is accurate to within 250 ns
23	The time is accurate to within 1 us
24	The time is accurate to within 2.5 us
25	The time is accurate to within 10 us
26	The time is accurate to within 25 us
27	The time is accurate to within 100 us
28	The time is accurate to within 250 us
29	The time is accurate to within 1 ms
2A	The time is accurate to within 2.5 ms
2B	The time is accurate to within 10 ms
2C	The time is accurate to within 25 ms
2D	The time is accurate to within 100 ms
2E	The time is accurate to within 250 ms
2F	The time is accurate to within 1 s
30	The time is accurate to within 10 s
31	The time is accurate to >10 s
32-7F	reserved
80-FD	For use by alternate PTP profiles
FE	Unknown
FF	reserved

Best Master Selection - 9

ISPCS Ann Arbor'08

Time source (information-only attribution; not used in BMCA)

Value (hex)	Time source	Description
10	ATOMIC_CLOCK	Any device, or device directly connected to such a device, that is based on atomic resonance for frequency and that has been calibrated against international standards for frequency and, if the PTP timescale is used, time
20	GPS	Any device synchronized to a satellite systems that distributes time and frequency tied to international standards
30	TERRESTRIAL_RADIO	Any device synchronized via any of the radio distribution systems that distribute time and frequency tied to international standards
40	PTP	Any device synchronized to a PTP-based source of time external to the domain
50	NTP	Any device synchronized via NTP or Simple Network Time Protocol (SNTP) to servers that distribute time and frequency tied to international standards
60	HAND_SET	Used for any device whose time has been set by means of a human interface based on observation of an international standards source of time to within the claimed clock accuracy
90	OTHER	Other source of time and/or frequency not covered by other values
A0	INTERNAL_OSCILLATOR	Any device whose frequency is not based on atomic resonance nor calibrated against international standards for frequency, and whose time is based on a free-running oscillator with epoch determined in an arbitrary or unknown manner
F0-FE	For use by alternate PTP profiles	
FF	Reserved	

Best Master Selection - 10

Execution of the BMCA is triggered by a ‘state decision event’ (referred to as a STATE_CHANGE_EVENT in V1)

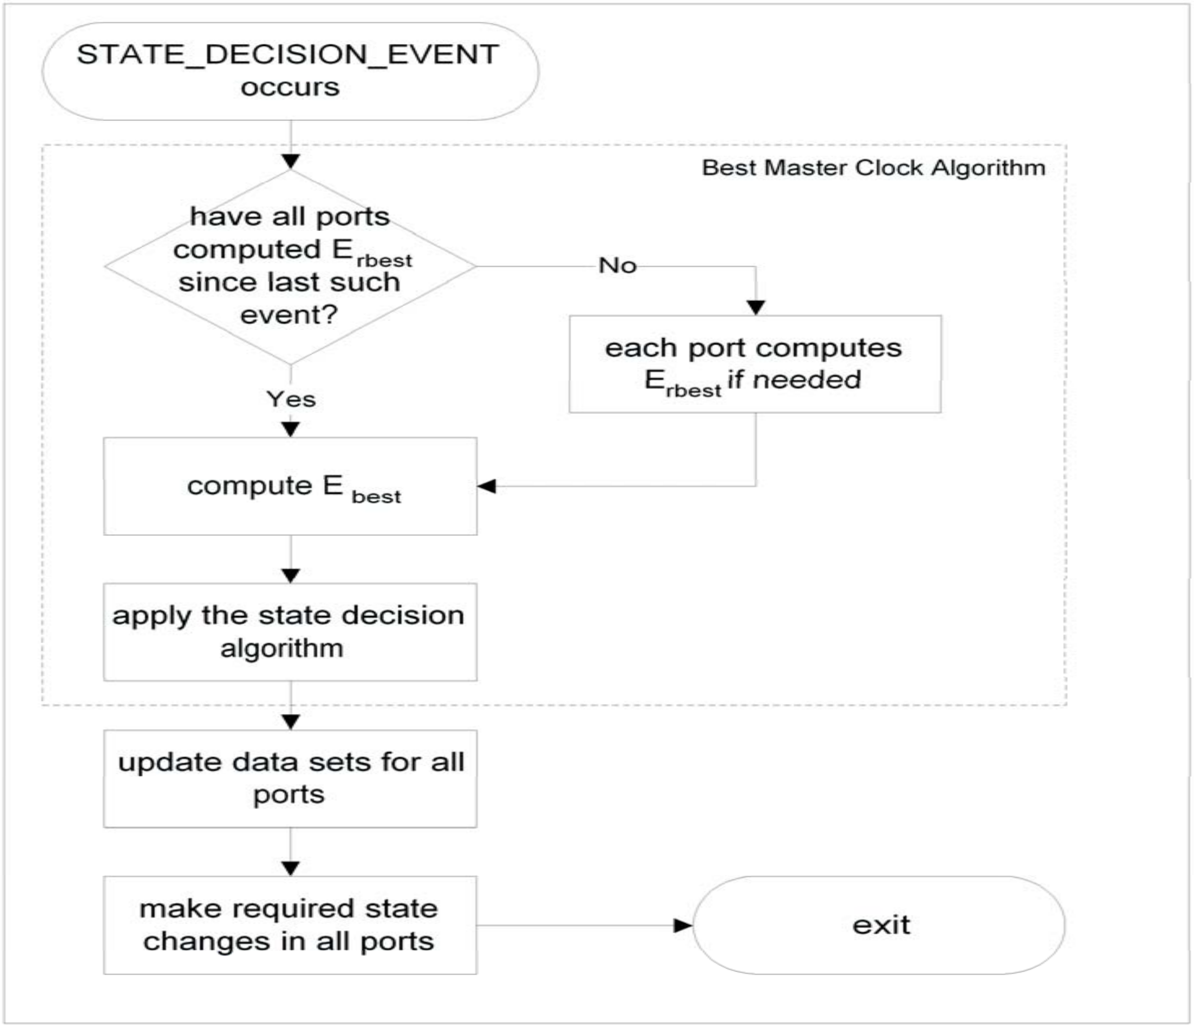

Best Master Selection - 11

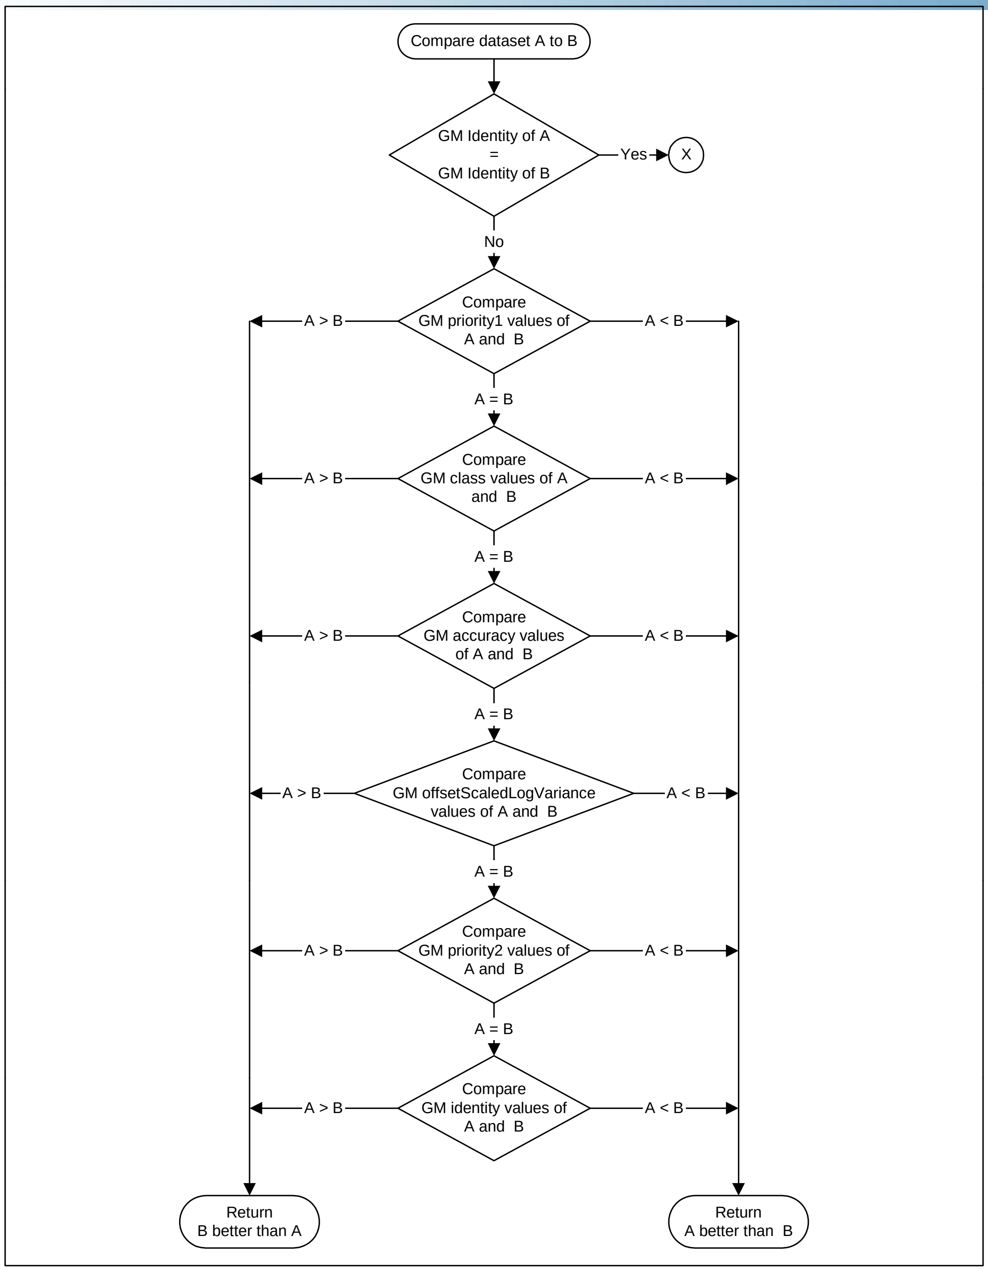

Dataset comparison algorithm - 1

Best Master Selection - 12

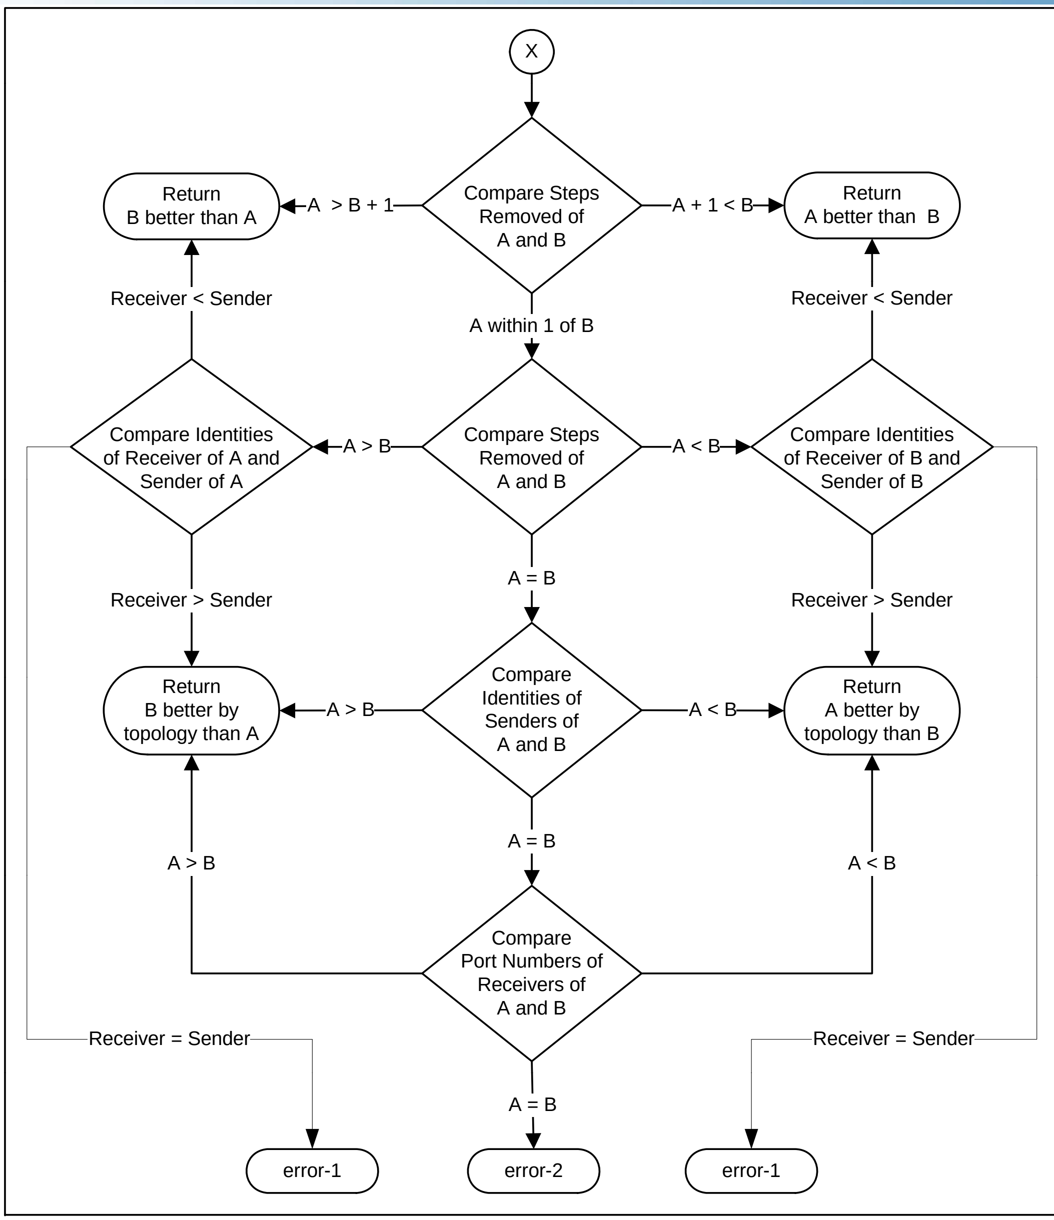

Dataset comparison algorithm - 2

Best Master Selection - 13

- In addition, IEEE 1588 V2 allows a PTP profile to specify an alternate BMCA
  - There are requirements on any alternate BMCA regarding updating of states and data sets (see reference 1 for details)
- IEEE 1588 V2 uses the same BC and OC state machines as V1, with some corrections that were identified since V1 was published
- IEEE 1588 V2 does not specify detailed state machines for TCs
  - Many aspects of TC behavior can be specified in a PTP profile

PTP Profiles and Conformance - 1

- Purpose of PTP profiles
  - Allow specific selections of attribute values and optional features to be specified such that, when using the same transport protocol, interworking and desired performance levels will be ensured
    - The interworking and performance levels will meet the requirements of a particular application
- Recommended that a PTP profile define
  - BMCA option
  - Configuration and node management options
  - Path delay measurement option (peer delay or delay request-response)

PTP Profiles and Conformance - 2

- Recommended that a PTP profile define (cont.)
  - Range and default values of all configurable attributes and parameters
  - Transport mechanisms required, permitted, or prohibited
  - Node types required, permitted, or prohibited
  - Options required, permitted, or prohibited
  - Procedure used to verify conformance

PTP Profiles and Conformance - 3

- In addition, IEEE 1588 specifies how a PTP profile can extend the standard (i.e., extensions, if done, are limited to these specifications):
  - TLV mechanism
  - Optional BMCA
  - Optional management mechanism
  - Use of a unicast messaging model provided the behavior of the PTP protocol is preserved
- Profiles are created by standards organizations, industry trade associations, companies, or other “appropriate” organizations

PTP Profiles and Conformance - 4

- Two default profiles are provided in Annex J of IEEE 1588
  - Delay Request-Response Default PTP Profile
  - Peer-to-Peer Default PTP Profile

PTP Profiles and Conformance - 5

- Conformance is specified in terms of nodes (i.e., the standard does not talk about "network-level conformance")
- Conformance is to
  - IEEE 1588 standard
  - PTP profile (i.e., a node claiming conformance shall conform to at least one PTP profile)
- Nodes conform to all normative sections of IEEE 1588 except those that are optional
  - If an option is implemented, the node shall conform to the clause that specifies the option (i.e., the option must be implemented in its entirety as specified)

PTP Profiles and Conformance - 6

- A node shall conform to either the Annex that defines the transport protocol that it uses or, if it uses another transport protocol, to the specification published by the respective standards organization that has jurisdiction over that transport protocol

General Optional Features - 1

- Unicast negotiation (subclause 16.1)
  - Provides for negotiated unicast sessions of finite duration between two clocks
    - Sync, Announce, and Delay_Resp messages
  - Subclause 16.1 defines TLVs that enable a clock to request unicast transmission, grant a requested unicast transmission made by another clock, and cancel a unicast transmission session

General Optional Features - 2

- Path trace (subclause 16.2)
  - Provides a mechanism for recording the identities of all boundary clocks traversed by a PTP message
    - Typically used with Announce messages to detect rogue frames
  - Path trace is enabled or disabled via management message
  - A TLV appended to Announce message contains the path trace list
  - Current path trace list may be retrieved from a clock via management message

General Optional Features - 3

- Alternate timescales (subclause 16.3)
  - One or more alternate timescales may be optionally maintained
  - ALTERNATE_TIME_OFFSET_INDICATOR TLV is attached to each Announce message, and contains current information for the alternate timescale
    - currentOffset
    - jumpSeconds
    - timeOfNextJump
    - displayName
  - Alternate timescale is enabled, disabled, and configured in GM via management messages

State Configuration Options - 1

- Clause 17 contains options that may be used to more explicitly control (i.e., more explicitly compared to the default BMCA)
  - how a PTP system recovers from failure of a clock or a break in the network
  - Which clocks in the network are acceptable masters
- How these options are used with the default BMCA or an alternate BMCA is up to the system integrator

State Configuration Options - 2

- Grandmaster cluster table (subclause 17.3)
  - Allows two or more OCs or BCs to be designated as a grandmaster cluster
  - Provides for more rapid evaluation of the BMCA to speed up change to new GM
  - Each clock in the cluster maintains a master cluster table
    - Priority1 values of clocks in the cluster should be less than those of all other clocks in the domain
  - Each clock in the cluster periodically requests unicast Announce messages from the ports listed in the grandmaster cluster table

State Configuration Options - 3

- Alternate master (subclause 17.4)
  - Allows alternate masters that are not currently the best master to be visible to slave ports
  - An alternate master is configured to send Announce and optionally Sync messages
  - Slave may use the information contained in these messages to acquire knowledge of the characteristics of the transmission path between itself and each alternate master
    - Allows for faster switchover to an alternate master when the current best master fails
  - Alternate masters are configured via management message

State Configuration Options - 4

- Unicast discovery (subclause 17.5)
  - Allows PTP to be used over a network that does not provide for multicast
  - Slave port is configured with addresses of potential masters, and may request unicast Announce, Sync, and Delay_Resp messages from these potential masters
  - Configure via management TLVs

State Configuration Options - 5

- Acceptable master table (subclause 17.6)
  - Allows configuration of a slave port to accept as a master only ports listed in the acceptable master table
  - This option may be used to constrain which clocks can be in the master-slave hierarchy and, in a limiting case, can force a specific master-slave hierarchy
  - Configured via management TLVs

State Configuration Options - 6

- There is a fundamental conflict between a self-configuring system (default BMCA defined in IEEE 1588) and a configured system
- The system integrator must exercise extreme care in using these options
  - Unintended behavior may result when using these options with the default BMCA

Compatibility Requirements - 1

- Clause 18 defines requirements for compatibility between
  - IEEE 1588 V2 and future versions
  - IEEE 1588 V1 and V2
- IEEE 1588 V2 messages contain a versionPTP field
  - V1 messages also contain a versionPTP field
  - V2 versionPTP is port-specific (i.e., applies to the port that sends the message), because in V2 the ports of a node may support V1 and/or V2

Compatibility Requirements - 2

- Compatibility between V2 and future versions
  - If a V2 node receives a message with a versionPTP field greater than 2, it shall
    - Disregard the message if it is not defined in V2
    - Parse and execute the message if it is defined in V2, but disregard any TLV extensions that are not defined in V2
    - Disregard the message if any inconsistencies are detected

Compatibility Requirements - 3

- Compatibility between V1 and V2
  - Recommended that V1 and V2 devices communicate via a BC
  - Communication between V1 and V2 devices via a TC is outside the scope of IEEE 1588
- Clause 18 contains detailed specifications that describe how V1 entities and attributes are translated to V2, and vice-versa, to ensure V1/V2 interworking
  - Domains
  - clock attributes
  - messageType
  - Message transmission intervals
  - Timestamp representation

Transport Specific Field

- IEEE 1588 V2 messages contain a transportSpecific field, that functions as a subtype of the respective Ethertype
- At present, 2 values are defined for the case of transport of PTP directly over IEEE 802.3/Ethernet
  - DEFAULT (value = 0 (hex)): All PTP layer 2 Ethernet transmissions not covered by another enumeration value
  - ETHERNET_AVB (value = 1 (hex)): This value is reserved for use in connection with the standard being developed by the IEEE 802.1 AVB Task Group as P802.1AS
    - See reference 8 for more detail on P802.1AS

Security - 1

- Annex K describes an experimental security protocol
  - The protocol has been reviewed by security experts
  - The protocol is experimental because it was felt experience must be gained in its use
- Provides group source authentication, message integrity, and replay protection for PTP messages
- Does not provide nonrepudiation
- Does not support data confidentiality (encryption)

Security - 2

- Composed of two basic mechanisms
  - Integrity protection mechanism, which uses a message authentication code to verify that a received message was transmitted by an authenticated source, was not modified in transit, and is fresh (i.e., is not a replay). Replay protection is implemented using counters.
  - A challenge-response mechanism, which is used to affirm the authenticity of new sources and to maintain the freshness of trust relations

Security - 3

- Security protocol uses symmetric message authentication code functions
  - Clocks share secret keys
- Flag in PTP message common header indicates whether the security authentication TLV is present
- Any intermediate transparent clocks that are present must also be security-enabled
- See reference 6 for more detail

Transport of Cumulative Frequency Offset Information

- Annex L describes an experimental TLV that may be used to transport cumulative frequency offset information
  - The quantity transported is a scale factor that depends on both the cumulative frequency offset and the frequency change needed to drive the phase error to zero
- This information can be useful in some phase and frequency compensation algorithms (e.g., see [B25] of the IEEE 1588 bibliography (Annex M))
- The TLV may be appended to either Sync or Follow_Up messages
  - Recommended it be appended to Follow_Up messages
  - If appended to Sync, Delay_Req must be padded to the same length to avoid asymmetry effects (and all nodes must support the extension)

Questions

References - 1

1. IEEE Std1588™ – 2008, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*, IEEE Instrumentation and Measurement Society, July 24, 2008.
2. John C. Eidson, *IEEE 1588 Version 2*, Agilent Laboratories, Measurement Research Lab, March 5, 2007.
3. John C. Eidson, *IEEE-1588 Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems – A Tutorial*, Agilent Technologies, October 10, 2005.

References - 2

4. John C. Eidson, *Implementing IEEE 1588 in LXI*, LXI Beijing meeting, June, 2007.
5. Geoffrey M. Garner, *IEEE 1588 Version 2*, ISPCS Vienna ‘07, Tutorial 1, October 1, 2007.
6. Ron Cohen, *Security and IEEE 1588*, Tutorial 2, ISPCS Vienna 07.
7. John Eidson, Geoffrey M. Garner, John Mackay, and Veselin Skendzic, *Provision of Precise Timing via IEEE 1588 Application Interfaces*, ISPCS Vienna 07.
8. Michael D. Johas Teener, Geoffrey M. Garner, Benjamin Hur, and Dawey (Weida) Huang, *Overview and Timing Performance of IEEE 802.1AS*, ISPCS Ann Arbor ‘08.