Summary of IEEE Std 1588™ – 2008 Optional Features Related to Redundancy and Potentially Improved Performance

Geoffrey M. Garner
Consultant

IEEE 802.1 TSN TG
2013.05.14

gmgarner@alum.mit.edu

Acknowledgement

- The author would like to thank John Eidson and Doug Arnold for having provided answers to several questions and useful discussions, on the master cluster and alternate master features of IEEE Std™ 1588 – 2008.

Introduction – 1

- Two of the features for P802.1ASbt pertain to redundancy and improved performance
- The features are (from the PAR Scope (section 5.2 of PAR form))
  - Support for redundant paths
  - Improved performance (e.g., improved grandmaster changeover time, longer chains of time-aware systems)
- IEEE Std™ 1588 – 2008 has two optional features (in clause 17) that can provide redundancy, improved grandmaster changeover time, and possibly better performance for longer chains of time aware systems. The features are:
  - Grandmaster clusters (subclause 17.3)
    - Makes use of the optional unicast message negotiation feature (subclause 16.1)
  - Alternate master (subclause 17.4)

Introduction – 2

- The purpose of this presentation is to summarize these features, as they may not be familiar to the 802.1 TSN TG
  - Note that IEEE 1588 has other optional features not described here, some of which are already used in 802.1AS (e.g., path trace TLV, acceptable master table)
- Note that there has been discussion in the IEEE P1588 committee (currently a study group; PAR for the next version has been approved by the study group and is expected to be submitted to NesCom) on deprecating the grandmaster cluster feature because it appears to not have been used (or hardly used)
  - If there were interest in using this feature in 802.1ASbt, this should be communicated to the P1588 committee
- **This presentation is for information; it is not proposing these features, but only briefly describing them**

Grandmaster Clusters – 1

- In this feature, a group of clocks is designated as a grandmaster cluster
- Presumably, the clocks are known in advance to be better than the other clocks in the network
- The members (i.e., clocks) of the cluster communicate with each other via unicast Announce messages
- Using these unicast Announce messages, each member of the cluster can know which member is best, which is second best, which is third best, etc.
- At any given time, the best clock of the cluster sends ordinary, multicast Announce messages in the normal way (i.e., according to clause 9 of IEEE 1588)
- The other clocks of the cluster, and clocks not in the cluster, process the multicast Announce messages in the normal way (i.e., according to the default or an alternate BMCA, depending on what is specified in the PTP profile)

Grandmaster Clusters – 2

- If the current GM goes away (due to either failure of the GM or failure of the link(s) connecting it to the network) it stops sending both multicast and unicast Announce messages
- The other members of the cluster immediately stop receiving unicast Announce messages, and know that the current GM is no longer present
- The next best member of the cluster immediately begins sending multicast Announce messages on all its ports
  - The reconfiguration time is reduced because the next best grandmaster does not need to wait for the information that the best grandmaster has gone away to propagate to it via multicast Announce messages
  - The reduction depends on how many hops the next best grandmaster is from the best grandmaster

Grandmaster Clusters – 3

- Each member of the cluster maintains a configured table of potential grandmasters, known as the <grandmasterClusterTable>
- The table contains
  - maxTableSize
  - logQueryInterval
    - This is the log to base 2 of the mean interval (in s) between unicast Announce messages from other cluster members
  - actualTableSize
  - Array of size actualTableSize containing the port address of each member of the cluster
    - Note: The port address, for the case of transport over layer 2 (Ethernet), is the MAC address. The use of port address might need to be changed to portIdentity in the next version of IEEE 1588, to address any issues arising from TCs sending PTP messages using the source MAC address of the sending port (but this does not affect any of the points of the current contribution)

Grandmaster Clusters – 4

- The priority1 attribute of each clock in the cluster must be less than (i.e., better than) that of all the other clocks in the network
- Each cluster member uses the unicast negotiation option of IEEE Std 1588 – 2008 to periodically request unicast Announce messages from all the other ports in the <grandmasterClusterTable>
  - Announce messages are requested from a port (of a clock) by sending a REQUEST_UNICAST_TRANSMISSION TLV to that port, in a Signaling message
  - The TLV indicates the messageType (Announce in this case), logInterMessagePeriod (the logQueryInterval in this case, which is the log to base 2 of the desired Announce interval), and the duration that it is desired to receive Announce messages
  - A port receiving a REQUEST_UNICAST_TRANSMISSION responds with a GRANT_UNICAST_TRANSMISSION, to either grant or deny the request
  - A port can cancel a request for unicast messages using the CANCEL_UNICAST_TRANSMISSION TLV
  - See subclause 16.1 of IEEE 1588 for details

Grandmaster Clusters – 5

- If a port in the cluster that transmits a unicast Announce message is in the master state (i.e., has port role of MasterPort), the alternateMasterFlag is set to FALSE; otherwise, it is set to TRUE (in the unicast Announce message)
- Each port in the cluster uses the received unicast Announce messages from the other cluster members to determine which port is best, second best, third best, etc.
  - This determination could be made by, e.g., running the BMCA; however, it actually is necessary to only compare the clocks, using part 1 of the dataset comparison algorithm
- The best clock in the cluster can immediately put its ports in the master state and send multicast Announce messages
  - Other clocks in the cluster receive, process, and send multicast Announce messages based on the ordinary operation of the BMCA (either default or alternate)
  - All members of the cluster continue to send unicast Announce messages to each other

Grandmaster Clusters – 6

- In IEEE 1588, the <grandmasterClusterTable> is populated via management messages
- In 802.1AS, the table would become part of the MIB (similar to the acceptable master table in 802.1AS)

Alternate Master – 1

- In this feature, alternate masters (i.e., potential masters) that are not currently the best master can exchange information with slave ports, so that a slave can acquire information on the transmission path between itself and a master
- IEEE 1588 does not specify what a slave does with the information, but the intent is that the slave could use the information to maintain the time of the alternate master, to lessen the changeover time to that master if the current best master goes away
- An alternate master transmits multicast Announce messages, and multicast Sync messages, with the alternate master flag set to TRUE
  - The fact that the flag is TRUE enables each slave port that receives the message to know that the message is from an alternate master

Alternate Master – 2

- Each alternate master maintains the following attributes, which are configurable via management messages (in 802.1AS, these would be part of the MIB)
  - numberOfAlternateMasters (UInteger8) – the maximum number of alternate masters
    - A port that is not in the master state shall transmit multicast Announce messages, with the alternate master flag set to TRUE, when the number of other ports that are
      - –transmitting qualified Announce messages with the alternate master flag set to true, AND
      - – that would be chosen as the best master in preference to the port

      is less than the number of alternate masters
  - transmitAlternateMulticastSync (Boolean) – If this is set to TRUE and the port is currently transmitting multicast Announce messages with the alternate master flag set to TRUE, the port shall also transmit multicast Sync and, if two-step, multi-cast Follow_Up
  - logAlternateMasterSyncInterval – the log to base 2 of the mean interval, in s, between Sync messages transmitted with the alternate master flag set to TRUE

Alternate Master – 3

- Note that the use of the information transmitted by the alternate masters is not specified in IEEE Std 1588 – 2008; only the conditions and format of the transmission of the information is specified
- For example, if it were desired in 802.1AS to use the alternate master information to construct multiple synchronization spanning trees and/or maintain the times of multiple, alternate masters, this would have to be specified in 802.1AS
  - In particular, IEEE 1588 does not specify how boundary clocks process received Announce messages with the alternate master flag set to TRUE
  - It seems that if there were to be multiple alternate masters, and a separate spanning tree for each master, the Sync (and Follow_Up) message would need to carry indication on which of the alternate masters originated the time synchronization information
- Therefore, it seems that if this feature is used, some amount of development will be needed in 802.1AS

Summary – 1

- This presentation has summarized the grandmaster clusters and alternate master features of IEEE Std 1588 – 2008
- The grandmaster cluster feature allows a pre-specified cluster of potential grandmasters to exchange information with each other via unicast messages so that, if the current grandmaster goes away, the next best member of the cluster can take over immediately
  - The result is shorter reconfiguration time
- The alternate master feature allows potential alternate masters to send information to slave ports, which can use the information to maintain the time relative to the alternate master if there is a reconfiguration, or improve performance by using both the main master and alternate master information
  - The result is shorter reconfiguration time and/or improved performance

Summary – 2

- Reference [1] indicates a main difference between the grandmaster cluster and alternate master features:
  - In the master cluster feature, the masters take care of redundancy so that the slaves can be simpler
  - In the alternate master feature, the slaves are more complex so they can be robust

Possible Action Items

- If the 802.1 TSN TG desires to use the grandmaster cluster feature, it should let the P1588 committee know this as soon as possible
  - One item discussed for the PAR for the next version of IEEE 1588 is to deprecate (or even eliminate) the grandmaster cluster feature because up to now it apparently has not been used
- If the 802.1 TSN TG uses either (or both) of the features, it should keep the P1588 committee informed of the work, as possibly other applications/profiles would be interested and it would be desirable to not have multiple developments of essentially the same feature

References

[1] Doug Arnold, email of April 29, 2013