IEC/IEEE 60802 Contribution – Time Sync Informative Annex

Version 03 – August 2023

David McCall (Intel Corporation)

1. Overview

IEC/IEEE 60802 (this specification) enables a network of compliant devices to achieve a time synchronisation accuracy, at the application level, of 1 µs over 100 network hops. To achieve this, it allocates the overall error budget of 1 000 ns as described in Table 1.

Table 1: Time Synchronisation Error Budget

Network Aspect	Error Type	Network-Level Error Budget (ns)	Normative or Informative?
All PTP Instances	Constant Time Error	200	Normative
	Dynamic Time Error	600	
All PTP Links	Constant Time Error	200	Informative
	Dynamic Time Error	Negligible	

A chain of 1 Grandmaster PTP Instance, 99 PTP Relay Instances and 1 PTP End Instance (100 network hops) that all comply with the normative requirements of sections 6.2.2 and 6.2.3 will generate a network-level Time Error at or below the Error Budget for All PTP Instances.

Section 2 describes the principles of operation this specification assumes.

Section 3 provides additional information on specific normative requirements.

The principles of operation include the use of crystal oscillators (XOs) as opposed to more accurate, stable and costly options such as temperature-compensated crystal oscillators (TCXOs). The use of XOs means that some of the normative requirements are difficult or impossible to meet without employing algorithms that track Neighbor Rate Ratio drift and Rate Ratio drift and compensate for consequent errors in calculating Rate Ratio and Correction Field. Section 4 of this annex provides examples of algorithms that can be used for this purpose, and which have been shown to enable compliance with the normative requirements.

Implementations that employ TCXOs or other more accurate, stable oscillators find some of the normative requirements difficult or impossible to meet without employing algorithms to track and compensate for errors due to clock drift in neighboring and/or Grandmaster PTP Instances that use XOs.

There is no normative requirement to use the algorithms described in section 4; an implementation can employ alternative algorithms provided the normative requirements are met. Section 4 describes the potential risks of deploying a network whose instances employ a mix of different algorithms. It is the responsibility of implementers to mitigate the risks and ensure alternative algorithms deliver the network-level performance.

This specification does not include normative requirements for PTP Links. Section 2.3 describes PTP Link characteristics that influence achieving 1 µs time synchronisation accuracy. It includes some examples using common PTP Link characteristics.

This specification's normative requirements regarding instance-level error generation are necessitated by the need to ensure not just an overall level of dTE generation at each node, but also the performance of drift tracking and error compensation algorithms and the amount of dTE generation due to timestamp error verses clock drift. The algorithms are employed to mitigate errors due to clock drift, but cannot mitigate timestamp errors. Section 5 describes an example approach to testing the normative requirements. It is not a test specification nor the only viable approach.

2. Principles of Operation

Achieving 1 µs time synchronisation accuracy across 100 network hops involves managing the accumulation of errors in the Precise Origin Timestamp + Correction Field and the Rate Ratio as they are passed, via Sync or Follow_Up messages, down the chain of PTP instances and are then used by the PTP End Instance to keep its ClockTarget in line with the ClockSource at the Grandmaster PTP Instance. All significant errors can ultimately be traced back to one of two sources: timestamp error or clock drift. The selection of PTP protocol parameters often involves trading off one source of error against the other. This specification requires specific PTP protocol configurations, and assumes the use of mechanisms (algorithms), that reduce dTE due to timestamp error but would also – without additional measures – increase dTE due to clock drift to the point where the latter exceeds the allocated error budget. However, this specification also assumes additional measures to minimise some sources of dTE due to clock drift and mechanisms and to track and compensate for errors from other sources to a sufficient degree that the error budget is not exceeded.

The specific protocol configurations and other measures, along with their intended effects, are described in Table 2.

Table 2: Protocol configurations & other measures to achieve dTE budget

Configuration or Measure	Description and Intended Effect(s)
Sync Interval 125 ms	Effects: 1. Calibrate the balance between dTE from timestamp error vs error due to clock drift. Larger intervals lead to less timestamp error and more error due to clock drift. 2. Keep below acceptable limits the impact of errors in Rate Ratio and Rate Ratio Drift estimation when keeping ClockTarget in line with ClockSource between arrival of Sync messages. Larger intervals increase the impact of any errors.

Drift_Tracking TLV - syncEgressTimestamp	Description: <Is this needed?> Effect: Enables calculation of NRR using Sync message timestamps, which eliminates error due to NRR clock drift that would otherwise occur between calculation of NRR using Pdelay_Resp messages and use during Sync message processing (i.e. calculation of Rate Ratio and output Correction Field values)
NRR Smoothing	Description: Algorithm to use timestamps from multiple past Sync messages when estimating NRR. Effect: Reduce the amount of error in the estimate of NRR due to timestamp error while increasing the amount of error due to clock drift.
NRR Drift Tracking & Compensation	Description: Algorithm to use timestamps from multiple past Sync messages to estimate NRR drift then apply compensation to correct for consequent errors in NRR Smoothing calculation. Effect: Mitigate the effect of errors due to clock drift when calculating and using the estimated NRR.
Drift_Tracking TLV – rateRatioDrift	Description: Carries estimate of Rate Ratio drift rate from one node to the next. Effect: Allows each node to estimate its own Rate Ratio drift rate by combining the incoming Rate Ratio drift rate with the local estimate of NRR drift rate.
RR Drift Compensation	Description: Algorithm that uses the estimate of RR drift rate to compensate for that drift, adjusting the estimated RR over time according to the drift rate. Effect: For PTP Relay Instance, minimises errors in the Correction Field caused by Rate Ratio drift. For PTP End Instances, minimises errors in keeping ClockTarget in line with ClockSource between arrival of Sync messages.
Sync Interval Consistency	Description: This specification requires tighter control of the interval between Sync messages generated at the Grandmaster PTP Instance than the defaults in IEEE 802.1AS-2020. Effect: Errors due to clock drift at Relay Instances have a tendency to cancel out. A clock drift which generates a positive error in NRR measurement on receipt of a Sync message generates a negative error in NRR measurement at the next node. The degree of cancellation depends on the consistency of the intervals over which NRR is measured at neighboring nodes.

	<Is this still necessary given that Sync Interval messages are inherently more consistent node-to-node, varying only according to differences in Residence Time?>
Pdelay Interval Consistency	<Should we add this back in given the conversation on 13th July 2023, i.e. where NRR based on Pdelay_Resp messages may be used when NRR based on Sync messages is not available, e.g. at startup or after a reconfiguration?>
Mean Residence Time	Description: This specification defines an average Residence Time requirement, where the average is significantly lower than the default maximum Residence Time in IEEE 802.1AS-2020. Effect: The amount of error in the Correction Field at the PTP End Instance due to clock drift is proportional to the cumulative meanLinkDelay and residenceTime experienced by a Sync message during transit from the Grandmaster PTP Instance to the PTP End Instance. Specifying a lower average residenceTime reduces this source of error.

2.1. Grandmaster PTP Instance Implementation

Depending on implementation, a Grandmaster PTP Instance may…

1. Contain a single oscillator used for both Local Clock and Clock Source,
2. Contain separate oscillators for Local Clock and Clock Source, or
3. Contain only an oscillator for Local Clock and accept an external input for Clock Source.

In some cases a Grandmaster PTP instance may support more than one mode of operation and transition between then depending on changes in system configuration (see Splitting, Joining and Aligning Time Domains).

In the first case the rateRatio and rateRatioDrift fields transmitted by the Grandmaster PTP Instance will be zero, reflecting the fact there is no difference between the Local Clock and Clock Source frequencies.

In the second and third cases there may be differences between the Local Clock and Clock Source frequencies. Any differences will be reflected in the rateRatio and rateRatioDrift fields transmitted by the Grandmaster PTP Instance. This means that Grandmaster PTP instances will track rateRatio over time in order to calculate rateRatioDrift, similarly to PTP Relay Instances and PTP End Instances. The exact implementation may vary.

2.2. Splitting, Joining and Aligning Time Domains

<Material from 6.2.5 and 6.2.13>

2.3. PTP Link Characteristics

A vast majority of time synchronisation error due PTP link characteristics is cTE due to asymmetrical path delay in one direction verses the other. The mechanism to measure path delay assumes the link is symmetrical and cannot detect asymmetry, thus asymmetry causes an error. The potential maximum asymmetry and thus error typically scales linearly with physical path length.

The error budget for cTE due to PTP link characteristics for an entire system is 200 ns. In any specific system this budget can be allocated as required with some links allocated a higher budget (typically longer length) than others.

A typical specified maximum length asymmetry for Ethernet cables is <TBC>, which translates to a maximum path delay asymmetry of <TBD>. If such cables are used a maximum total cable length between Clock Source and Clock Target with 99 PTP Relay Instances between them (i.e. 100 network hops) is <TBC>.

Depending on how cables are manufactured and deployed it is feasible for the asymmetries of each cable segment to be additive.

3. Notes on Normative Requirements and PTP Link Recommendations

3.1. Oscillator Requirements

Clock drift at the Grandmaster PTP Instance causes greater dTE than the same amount of clock drift at a PTP Relay Instance or the PTP End Instance. This specification therefore requires tighter limits on maximum fractional frequency offset for an oscillator at the Grandmaster PTP Instance than at other instances.

This specification does not place requirements on operational temperature range or other environmental factors. The required oscillator behaviour is delivered for the operational conditions across which a device claims it is compliant. These conditions typically include temperature range but may also include rate of change of ambient temperature, supply voltage stability, amount of vibration and others.

3.2. Timestamp Granularity Error

Timestamp Granularity Error (TSGE) is the error in timestamping each incoming and outgoing message due to the maximum timestamp resolution of which an implementation is capable. It is typically directly related to an implementation’s clock rate. For example, a clock rate of 125 MHz typically results in a maximum resolution of 8 ns while a clock rate of 500 MHz typically results in a maximum resolution of 2 ns.

It is assumed that TSGE varies between -4 ns and +4 ns with an average of 0 ns. Lower levels of TSGE are better. Implementations where TSGE is higher will find some of the normative requirements difficult or impossible to meet.

3.3. Dynamic Timestamp Error

Dynamic Timestamp Error (DTSE) is the, effectively random, error in timestamping each incoming and outgoing messages due to an implementation’s inherent inaccuracies, excluding TSGE. It is assumed to vary between a minimum of -6 ns and a maximum of +6 ns with an average of 0 ns.

If an implementation timestamps an incoming or outgoing message at a point other than the PHY, any variability in delay between that point and the PHY (PHY delay) will translate to DTSE. Some common implementations were not designed to limit this variability. If care is not taken to avoid implementations with high variability, the assumed DTSE range is easily exceeded. Such implementations will find some of the normative requirements difficult or impossible to meet.

3.4. Grandmaster PTP Instance Error Generation Requirements

Table 12 sets normative requirements for error generation at a Grandmaster PTP Instance that ensure the relevant fields in the Sync and Sync_Followup messages it transmits are sufficiently accurate to deliver the network-level performance.

- preciseOriginTimestamp + correctionField
    - Timestamp Error
- rateRatioDrift
    - Performance of rateRatioDrift measurement
- rateRatio
    - Performance of rateRatio measurement
- syncEgressTimestamp
    - Timestamp Error

3.5. PTP Relay Instance

Table 13 sets normative requirements for error generation at a PTP Relay Instance that ensure the relevant fields in the Sync and Sync_Followup messages it transmits as part of Sync processing are sufficiently accurate to deliver the network-level time sync performance.

- preciseOriginTimestamp + correctionField
    - Measurement of Residence Time
- rateRatioDrift
    - Performance of NRR drift measurement when there is no NRR drift
    - Performance of NRR drift measurement when there is NRR drift
- rateRatio
    - Performance of rateRatio calculation when there is no RR drift
    - Performance rateRatio calculation and error compensation when there is RR drift due to GM clock drift (incoming rateRatioDrift field)
    - Performance of rateRatio calculation and error compensation when there is RR drift due to Local Clock drift (NRR drift measurement)
- syncEgressTimestamp
    - TimestampError

3.6. PTP End Instance

Table 14 sets normative requirements for error generation at a PTP Relay Instance that ensure the relevant fields in the Sync and Sync_Followup messages it transmits as part of Sync processing are sufficiently accurate to deliver the network-level time sync performance.

- Performance of ClockTarget generation when there is no clock drift.
- Performance of ClockTarget generation when there is RateRatio drift due to GM clock drift. (incoming rateRatioDrift field)
- Performance of ClockTarget generation when there is RateRatio drift due to Local Clock drift. (NRRdrift measurement)

4. Example Algorithms

<Imported from "60802 Time Sync - Monte Carlo and Time Series Simulation Configuration Including NRR and RR Drift Tracking & Error Compensation v3" but adjusted to use mNRRcompNAP 8 and mNRRsmoothing 4>

5. Approach to Testing Normative Requirements

<Import from "60802 Time Sync – Rate Ratio Drift Tracking & Error Compensation" with adjustment for use of RRdriftTracking field and additional text to explain test approach>