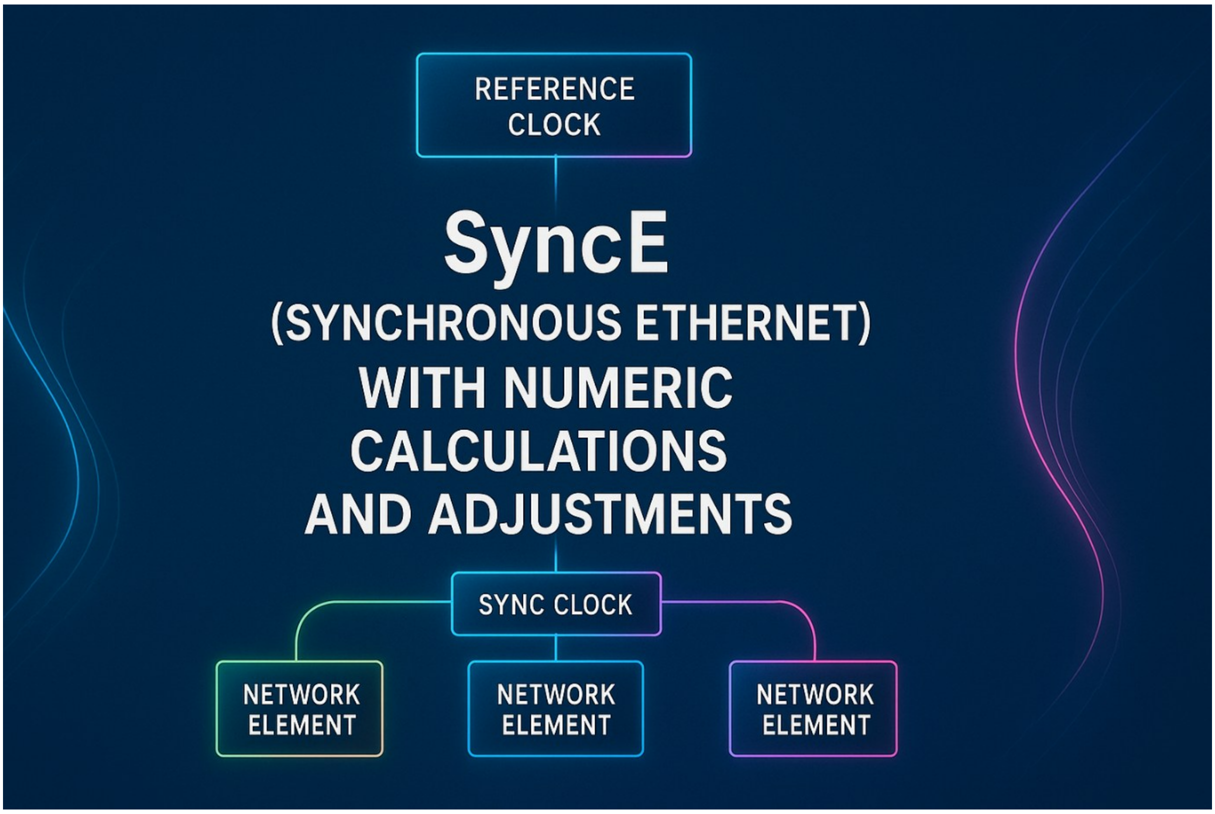

SyncE (Synchronous Ethernet) with Numeric Calculations and Adjustments

NetTimeLogic is a leader in high performance, high accuracy and high resiliency time synchronization solutions.

The Telecom industry has been using SyncE since the late 2000s, early 2010s. So, it is nothing particularly new, however NetTimeLogic takes a slightly different approach when it comes to synchronization of a SyncE Node when in Slave mode.

Most SyncE nodes use a dedicated PLL chip (which NetTimeLogic also offers as a PMOD-compatible module, but which is not used in this scheme). This PLL receives the recovered clock from an Ethernet PHY with RX clock recovery, synchronizes to it, and uses the PLL output clock to run the local time when operating in Slave mode. In Master mode, the local oscillator is used to drive the reference clock of the PHY. The PLL also provides the functionality to generate a clock in the absence of the recovered RX clock. In this setup, switching between SyncE Master and Slave operation, or between different SyncE sources, normally requires switching physical clocks and/or relocking the PLL.

Numeric Calculation and Adjustments

NetTimeLogic's approach is slightly different. We use the recovered RX clock to generate a PPS (Pulse Per Second) which we then just simply timestamp with a high resolution Timestamper (1ns). Reducing the recovered clock from 25/125MHz to 1Hz gives a much better noise ratio. Two consecutive PPS are timestamped and the delta calculated. The delta is supposed to be exactly 10^9 ns. If the delta is, for example, 1 ns more, this means the local oscillator runs 1 ppb (parts per billion) too fast. Conversely, if the delta is 1 ns less than 10^9, the local oscillator runs 1 ppb too slow. This frequency difference is then fed through an outlier filter, a PI servo loop, and an additional rate change limiter before being used to either numerically correct a counter clock or to pull the frequency of the local oscillator. The servo allows not only stepping the frequency in 1 ppb increments, but also applying fractions of 1 ppb (in our implementation 1/2^16 ppb). Since the PI Servo is freely configurable the adjustment rate and stability can be changed as needed. This is the same scheme that NetTimeLogic uses for all its other synchronization sources like PTP, PPS, IRIG etc. This also means that the drift calculation will benefit from the Advanced Holdover and Aging mechanism. Also drift corrections can be combined between e.g. PTP and SyncE (e.g. for ITU-T G.8275.1). Whenever one of them fails, frequency adjustments continue using the remaining mechanism. This switchover is completely seamless and therefore adds extra resilience. NetTimeLogic also offers a Dynamic Source Selector, Voter, and Merger module, which allows voting and merging of numeric adjustments from multiple SyncE sources (e.g. when multiple PRTCs, aka GMs, are connected). Averaging of multiple SyncE sources will not be possible with a classic PLL where one source has to be explicitly chosen and also voting can again increase the resiliency in case a source differs too much.

All these mechanisms, which have already been proven to allow superior holdover, can be used without extra effort or additional resource usage. If no synchronization source is available anymore (including SyncE), the local time continues to run with excellent holdover performance and can still provide SyncE synchronization in the downstream direction for quite a long time.

Test

NetTimeLogic did a simple test by synchronizing a Slave device by a combination of PTP and SyncE from a GNSS synchronized GM. After a couple of minutes (when the local time was frequency and phase locked for a while) the PTP synchronization was disabled and the node continued syntonization purely on SyncE. The PPS of the GM and Slave node was constantly monitored for 48h. The two PPS stayed within 50 ns over the whole duration which not only means it stayed accurately frequency aligned but also pretty much phase aligned. How far the PPS differs mainly depends on how quickly the SyncE Slave node reacts to frequency changes of the SyncE Master. Since the frequency of the SyncE Master is also fluctuating due to synchronization to the GNSS average frequency, the SyncE Slave is constantly adjusting its frequency to follow the Master.

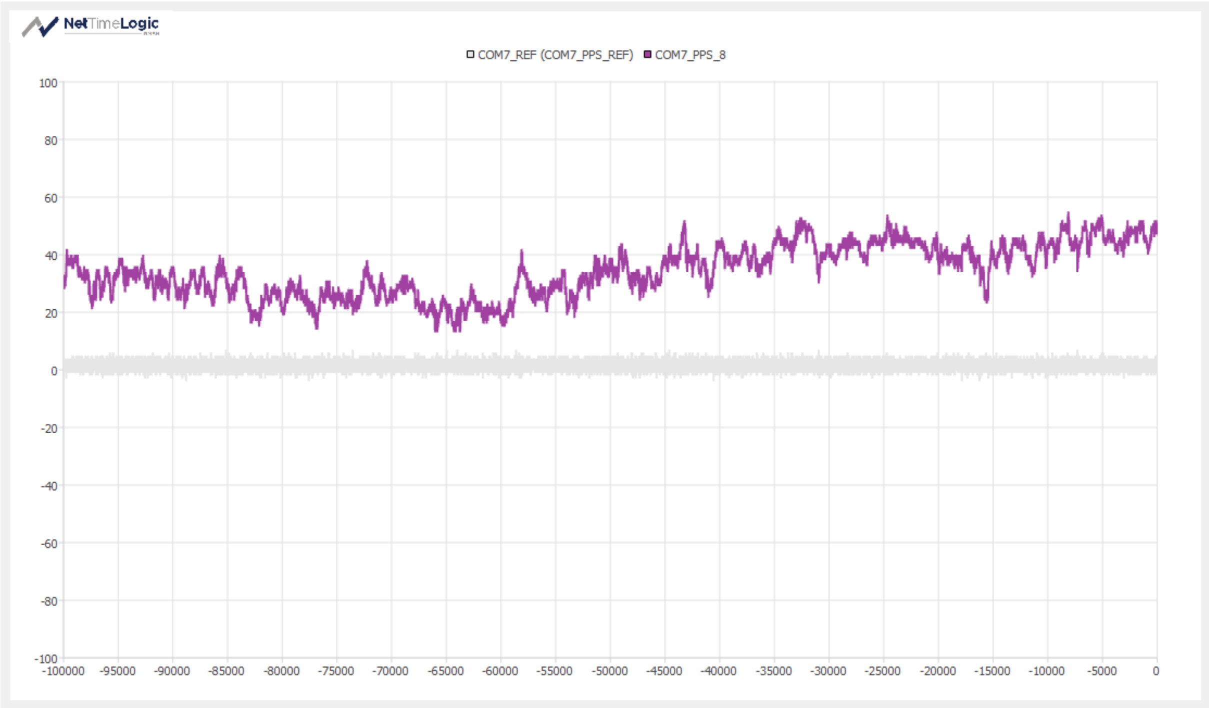

SyncE-only syntonization was run for 48h, this shows the PPS of the latest 100000 seconds (~27h).

Advantage of using a numeric calculation over a dedicated PLL

- You can save on the cost of a SyncE-compliant PLL, which is usually quite expensive
- Better holdover due to full control of the PI servo, filter, holdover, and aging mechanisms
  - Adaptive PI Servos possible to synchronize fast at the beginning and slower once frequency locked
- Same synchronization scheme as for any other mechanism
  - No special PLL chip just for SyncE
- Can be easily combined with other synchronization mechanisms.
- Can average and vote on multiple possible SyncE sources
  - A PLL can only follow one SyncE source

If you would like more information about our SyncE scheme, send us a message or leave a comment.