DOI 10.5162/etc2016/1.2

Nanosecond Synchronous Analog Data Acquisition over Precision Time Protocol

Matthias Braun, Milan Juranek, András Széll, Péter Szántó, Dr. Cruz Marín
Zodiac Data Systems GmbH, Friedrich-Ebert-Str. 75, 51429 Bergisch Gladbach, Germany
matthias.braun@zodiacaerospace.com
milan.juranek@zodiacaerospace.com
szell@fpgart.hu
szanto@fpgart.hu
cruz.marin@zodiacaerospace.com

Abstract

This paper describes an implementation of the Precision Time Protocol – IEEE™1588-2002 and IEEE™1588-2008 – in order to reach maximum accuracy in time synchronization as a requirement for simultaneous analog data sampling.

Exact cross-correlation calculation between various measurements calls for a very small phase error margin between the sampled signals. Besides, there is an increasing need to synchronize the analog sampling carried out at different places without utilizing additional synchronization connections between different devices.

Nowadays, Ethernet is universally used and therefore present almost everywhere. Data acquisition systems typically deliver their data via networks. For time synchronization, the Precision Time Protocol (IEEE1588™) is basically a mandate. This protocol provides the high accuracy required for synchronous and simultaneous analog sampling.

The performance evaluation of the present implementation shows that an absolute time synchronization better than 10 nanoseconds can be achieved between two given data acquisition systems. Furthermore, this paper elaborates on how the analog sampling can be synchronized to the absolute time with this same accuracy.

Key words: IEEE1588, Time Synchronization, Analog Data Acquisition, Synchronous and Simultaneous Analog Sampling

Introduction

In flight test instrumentation, reducing and simplifying wiring is a continuous effort. One way to achieve this goal is to use distributed data acquisition systems, where connecting only power and some serial digital communication lines (most commonly Ethernet technology) obviates the need of long cabling for all sensors over long distances. Beyond data transfer, Ethernet is used for the distribution of absolute time for coherent time stamping, enabling more and more accepted methods to guarantee synchronous analog sampling between multiple data acquisition units. Because no additional hardware or software components are required, the method described in the following, which is based on the standard Precision Time Protocol (PTP), is data acquisition hardware manufacturer independent.

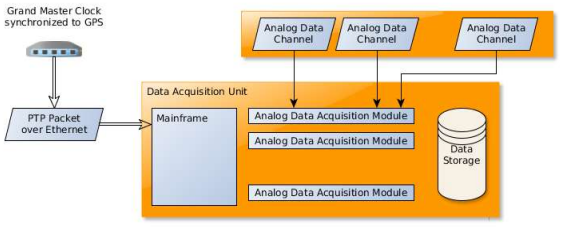

Fig.1 Data Acquisition Unit

Analog Synchronization

In order to be able to correlate analog data from different acquisition systems it is necessary to synchronize the time and the analog sampling points in multiple systems to each other. Using the PTP protocol is one approach conducive to synchronizing the acquisition systems to a GPS based PTP time server. To avoid the need of a digital resampling of the analog data the target is to ensure that the real analog samples are

taken at the same time by means of synchronizing both the frequency and phase of the sampling clocks.

The first step is to establish a precise absolute time base for the Data Acquisition Devices based on the PTP Synchronization Protocol. To synchronize the sampling clock frequency of all analog to digital converters (ADCs) to this absolute time base, the digital frequency generator devices have to be tuned in a second step. This tuning can only be done in very small steps in order to avoid jitter in the analog sampling – so the generated absolute time from the first step must be jitter-free as much as possible. In the last step the same phase of the sampling has to be guaranteed. To be able to sample the analog signals simultaneously in multiple systems it is necessary to take into account the delays effective along the entire signal path, on which a signal propagates from the input connector to the place where digital sampling takes place, including the pass through amplifiers and filters – the analog sampling needs to be delayed accordingly.

On one hand it is possible to design digital filters with arbitrary delays so that the total propagation delay of the ADCs and digital filters is typically an integer number of samples. But on the other hand, when it is not possible, the sampling signals shall be adjusted to compensate also fractions of the delays.

With this method it is theoretically feasible to phase correlate the acquired signals even if different types of acquisition hardware components are used. To simplify the analysis process of finding corresponding samples based on the timestamps included in data streams, the start of the data packets can be synchronized to the second's change, which is only doable when using integer sampling rates.

The accuracy of sampling synchronization depends on the following factors:

- The stability of the reference clock.
- The accuracy of the absolute time synchronization in the data acquisition unit.
- The resolution of the sampling frequency control.
- The accuracy of the compensation of the analog and digital delays of the digitalization process.
- The jitter of the sampling clock.

Recording Format

This implementation uses the IRIG106 Chapter 10 data format [1] for its recordings. In this data format (see Fig. 2) all acquired data is timestamped with a free running relative time counter. While every data acquisition unit has its own relative time counter and their values don't necessarily equal each other, it is necessary to convert all timestamps to one time domain to correlate two analog channels from different systems. The following steps are mandatory to convert the time domain of recording 2 to the time domain of recording 1:

1. Calculate the time offset between analog channel 1 and time channel from recording 2. In this case 1480-1478=2. The unit is 100 ns.
2. Calculate the time offset between analog channel 1 and time channel from recording 1. In this case 1236 – 1234 = 2.
3. Calculate the time offset of the relative time counters between the two time channels at time packets containing the same absolute time.1478 – 1234 = 244.
4. Analog channel 1 from recording 2 gets its new relative time by TS - offset1 + offset2 + offset3 = 1480 – 2 – 244 + 2 = 1236.

In this example the analog data was created simultaneously.

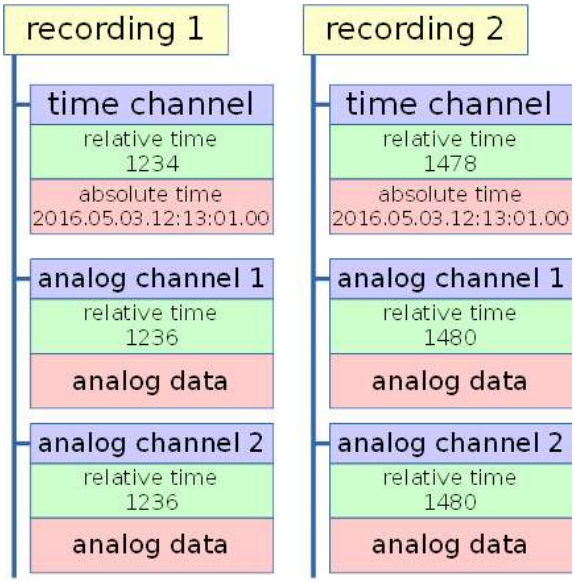

Fig. 2 Packet structure of two IRIG106 CHAPTER 10 recordings.

Precision Time Synchronization Protocol

To synchronize computer clocks over a given network, a protocol called 'Precision Time Synchronization Protocol' exists. Today two versions are available. The first version was published in 2002, and therefore it is named

'IEEE Std 1588™-2002, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems'. Version 2 was introduced in 2008, and it is known as 'IEEE Std 1588™-2008, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems'. It is not backwards compatible and is designed to improve the precision and robustness of the time synchronization while achieving an accuracy better than 1 nanosecond.

In a PTP network domain, different types of clocks can coexist. They are called 'Ordinary Clocks', 'Boundary Clocks' and 'Transparent Clocks'. An 'Ordinary Clock' is a device that uses a single network connection for time synchronization. It can act as a master or as a slave clock. A 'Master Clock' sends out timing information for a 'Slave Clock' to get in sync with it. A 'Boundary Clock' uses multiple network connections and is used to synchronize several network segments. The third type of clock is called 'Transparent Clock'. This one is introduced in IEEE™1588-2008. 'Transparent Clocks' convey PTP messages and modify the timing information with the time slice the PTP messages need to pass through the network node. This allows better compensation of message delays.

In a PTP domain, more than one master clock can be present. By evaluating the 'Announce Messages' every master clock generates, a 'Grand Master Clock' also known as GMC, is elected by the 'Best Master Clock Algorithm' (BMC). The GMC is the root timing reference which transmits synchronization information for all other clocks in the PTP network. The slave clocks synchronize to this 'Grand Master Clock' by exchanging different timing messages.

The GMC periodically broadcasts 'Sync Messages' over the network to let the slave clocks know about its timing information. Therefore, the instant of time when the 'Sync Message' leaves the GMC hardware has to be embedded into this message. This is called a one-step synchronization and requires hardware processing for highest accuracy. Master clocks without such hardware follow a two-step synchronization protocol. Here the 'Sync Message' contains an estimated timestamp, but in addition, a separate 'Follow_Up Message' is broadcasted with the accurate timestamp of the instant the 'Sync Message' leaves the GMC. For the implementation of the delay compensation, the calculation of the predicted propagation time of the PTP message transfer over the wire ('Mean Path Delay') is required. Two more PTP messages have to be sent over the network: The slave clock sends out a 'Delay_Req Message' for which the GMC notes the timestamp when it arrives and after that sends back a 'Delay_Resp Message' containing this timestamp. For an exact description of how this message exchange happens please have a look at: 'IEEE1588™-2008, 11.3 Delay request-response mechanism' [3].

The PTP messages are divided into 'General Messages' and 'Event Messages'. 'General Messages' provide more general information whereas 'Event Messages' are time critical and provide timing information for calculations in the PTP stack. Therefore they have to be timestamped. This can be done in software or in hardware. While software timestamping is not deterministic, hardware timestamping allows a much more precise accuracy in timing information for the offset calculations in the PTP stack.

In order to timestamp arriving and departing 'Event Messages', the hardware layer needs to detect these as close as possible to the physical wire using the physical interface chip also known as PHY. Once a PTP Ethernet packet is identified as an 'Event Message', the PHY core stores the value of the PHY internal global counter and the ID of the PTP packet into specific registers. Also, it informs the CPU respectively the kernel PHY driver that such an event happened. As shown in figure 3, a '100MHz Relative Time Counter' (rtc100) is implemented in every Data Acquisition Unit which is driven by a ‘Controlled Oscillator’ (CXO). To allow correct calculations of the value of the timestamp in nanoseconds when an 'Event Message' arrives or departs, the PHY of the given hardware is programmed to continuously trigger an impulse to the FPGA to capture the actual '100MHz Relative Time Counter' value and match them against the PHY internal global time counter. The PTP stack can now grab the timestamp information from the PHY driver for further calculations. In order to sync a Slave Clock to a given Grand Master Clock, the Time Processing Unit has to be informed by the PTP stack of the elapsed nanoseconds since the start of the PTP epoch (January, 1st 1970 00:00:00 UTC) at a given rtc100 counter value. This timestamp is the instant of time when the 'Sync Message' leaves the GMC (also known as the 'Precise Origin Timestamp'), compensated with the 'Mean Path Delay'. For further details refer to 'IEEE1588™-2008, 11.3 Delay request-response mechanism' [3].

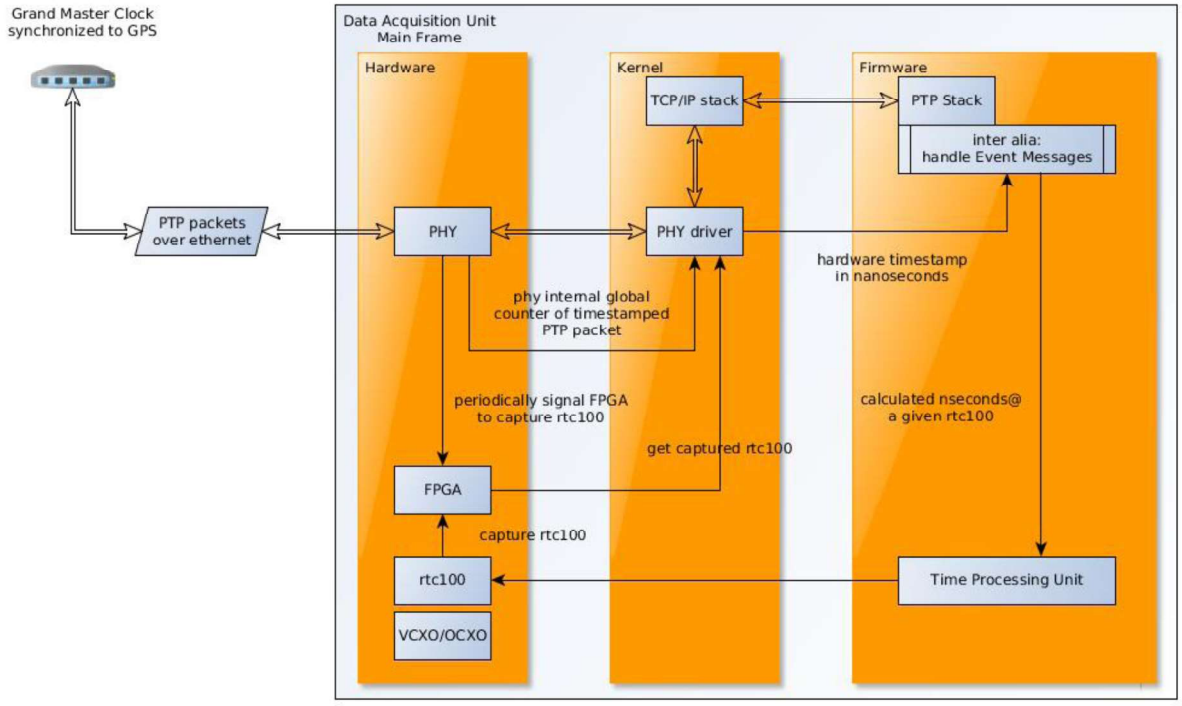

Fig. 3 PTP Message Flow and Timestamping in a Data Acquisition Unit

Time Processing Unit

The time synchronization subsystem or Time Processing Unit (TPU) supports synchronization to several source types including PTP. For any of these time sources, pulse-per-second (PPS) may also be connected to increase time precision. High precision (better than 100 ns) time synchronization to most sources is possible either with the PPS pulses or using long-term filtering in order to reduce measurement noise. Longer filtering (in the range of minutes, up to half an hour) is sensible only if proper time keeping is possible by means of using a stable oscillator. Otherwise, even though measurement noise is reduced by long-term averages or filtered values, temperature changes and other factors affecting the internal clock will introduce mid-term and long-term frequency deviations, degrading the performance of any filters applied to the input source.

The TPU is a time source itself. There is a trade-off between the filtering/tuning parameters: By applying strong filters on the input time source, the clock output will appear more stable and less jittery (assuming an oscillator with good short-term stability), but it will lock in and follow frequency drifts more slowly, resulting in large temporary offset errors against the source.

Due to the different characteristics of the external clock sources, the TPU uses clock source specific filtering. For some sources it is possible to record the raw, unfiltered input timestamps (for error detection) or filtered ones. Filtering, outlier detection and averaging of the input signals is necessary to reduce noise and sampling errors.

Even the best algorithms fail to perform well if the underlying hardware architecture doesn`t allow to properly sample the incoming timestamps and to keep the internal clock in sync with the external source. The base clock used in the TPU is either a voltage controlled oscillator (VCXO) or an oven controlled oscillator (OCXO) with better than 10 ppb stability. The VCXO has frequency fluctuations in the range of 50 ns/sec (50 ppb) even under stable temperature conditions. This is partly due to the imprecision of the crystal and partly due to the higher tuning range of the VCXO, thus it is more susceptible to minor tuning voltage noises (of mV level). In order to synchronize within 30 ns to the master clock, the presented synchronization method uses an OCXO. Under stable ideal conditions this setup can be fine-tuned to be within 1-2 nanoseconds from the reference clock.

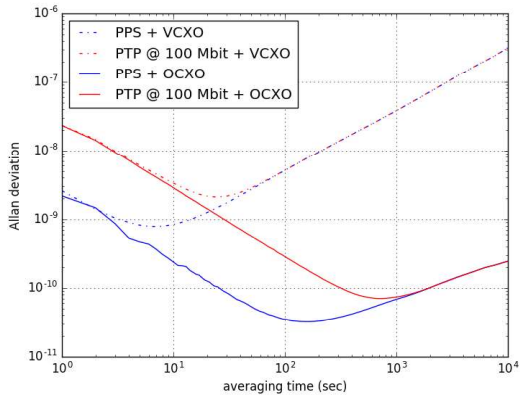

Fig. 4 Oscillators compared: Allan deviation of a TPU synchronized over PTP or PPS featuring an OCXO (solid lines) or VCXO (dashed)

Time stamps in the TPU are generated at 200 MHz resolution for PPS source and at 100 MHz for PTP sources (the rtc100 counter), meaning an inherent time stamping jitter of 5-10 ns. In case of PTP, the actual timing inaccuracy resulting from the finite resolution of the Ethernet PHY's time stamps is larger if the clock source is connected via 100 Mbit Ethernet, as the 25 MHz frequency used for data transmission means 40 ns resolution at best (in contrast to the 1 Gbit Ethernet's 125 MHz clock and 8 ns resolution). To reach the best analog synchronization performance, a Gbit Ethernet was used for the present work

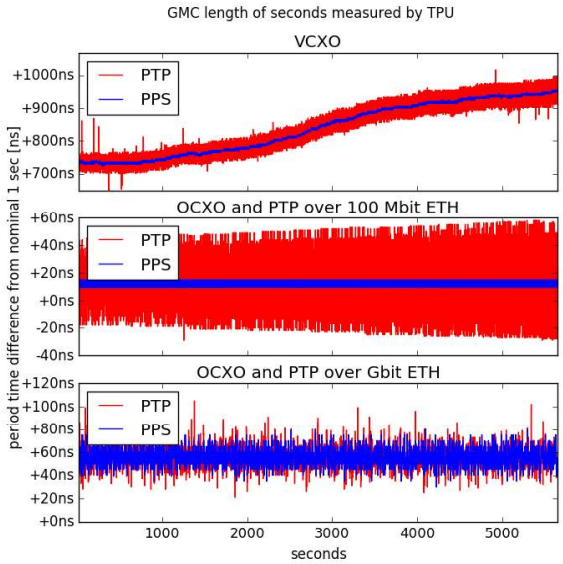

Fig. 5 Effect of Ethernet speed and oscillator type on PPS/PTP timestamping jitter and frequency drift

In Figure 4 the performance of the voltage controlled oscillator is compared to the oven controlled, temperature-stabilized oscillator. In this Allan deviation plot it is apparent that the short-term PTP precision is below the accuracy of PPS input due to the 40 ns resolution, but for longer averages (in the range of 10^3 seconds) their variances match, and the OCXO based solution is one order of magnitude better than the other system utilizing a VCXO. The VCXO frequency stability already deteriorates in the range of 100 seconds mostly due to temperature changes, as seen in Figure 5.

During the initial phase of the synchronization the oscillator is tuned to match the short-term averaged frequency of the GMC (3-5 seconds) and the offset is compensated with period time modification (as if the seconds are a few extra microseconds longer or shorter than 10^6). Once the oscillator frequency and offset errors are within bounds, only oscillator frequency tuning is applied afterwards. Synchronous analog sampling begins only in this second phase. The filter coefficients gradually change to reduce tuning the jitter as the observed offset and frequency error gets smaller. Offset errors are reduced by dithering (temporarily mistuning the frequency similarly to how one car keeps distance from the other just by minor accelerations and decelerations).

After successful synchronization the synchronized internal RTC clock is recorded into IRIG 106-09 Chapter 10 time packets and sent out to all other modules taking part in the data acquisition. The time packets contain the UTC time recorded at the given relative recording time, which is by default time stamped with 100 ns resolution. Later Chapter 10 extensions will increase the precision to nanoseconds.

Sampling Clock Generation

In a distributed system it is often not possible for the analog data acquisition units to communicate with each other, only their clocks can be synchronized. A 1 Hz 'analog frame signal' populated from the time synchronization subsystem via a dedicated clock pulse distribution network to the analog sub-systems is used to ensure simultaneous sampling.

The analog sampling clock is derived from the time subsystem's 100 MHz clock with a digital synthesizer. This clock value is an integer multiple of the analog sampling clock rate only in a limited set of frequencies, so in most cases it is necessary to fine-tune the synthesizer by means of a PI controller (proportional-integral controller) to ensure that the 1 Hz analog frame signals match the PPS signal of the TPU. The regulation of the PI controller can be seen in Figure 6. In this case an Analog Master Clock of 19.1979 MHz will result in a standard deviation of 3.519 ns. The time offset is measured between the falling edge of the analog frame signal C3 (which is an active low signal) and the rising edge of the PPS out signal C1.

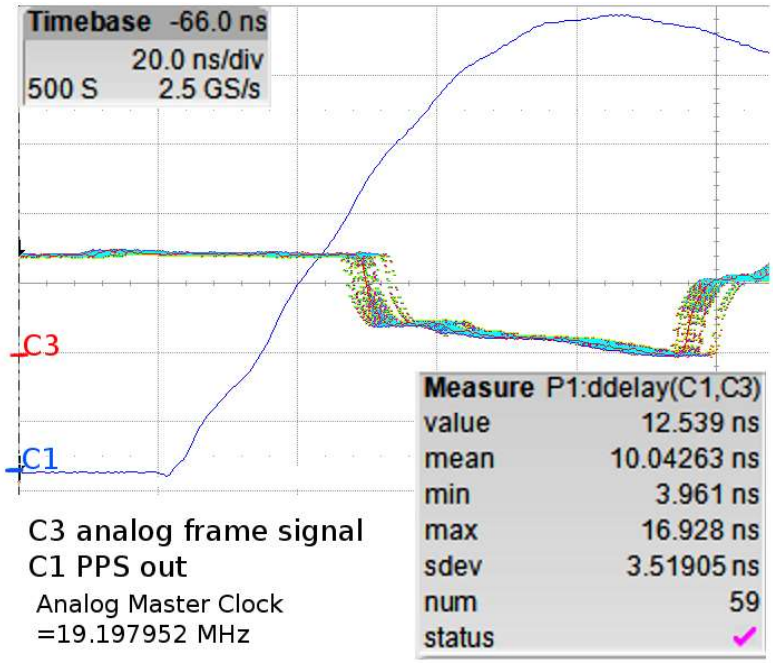

Fig. 6 Deviation measurement between the PPS output and the low active Analog Frame

Properly compensating signal delays in each input channel will allow a correlation between the latter even if their sampling rates are different.

Evaluation of Results

Figure 7 shows the Data Acquisition Unit's synchronization to a Grand Master Clock. The synchronization was carried out via a dedicated 100 Mbit network connection, the recorded data was transferred over a separate Ethernet interface.

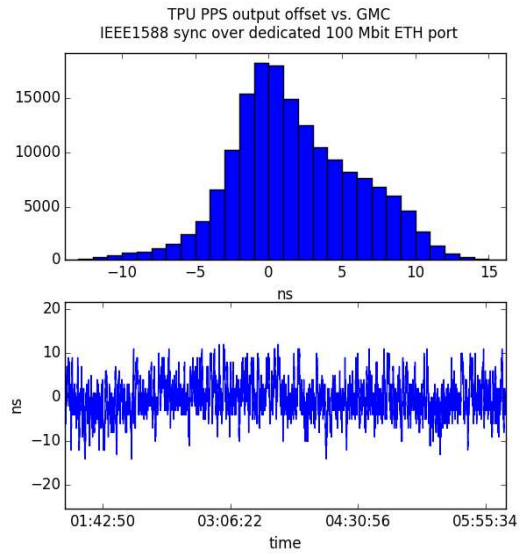

Fig. 7 IEEE1588-2008 direct synchronization against Grand Master Clock: offset error distribution

There are a multitude of possible error/jitter sources, among others:

- Frequency drifts due to temperature changes
- Noise/jitter in the signal (on PTP the network traffic, the carrier/PHY buffering jitter and performance of other network components; delay measurement issues)
- Time stamping inaccuracies caused by finite time stamping resolution ('time domain quantization errors')
- Analog sampling clock synchronization against internal (synchronized) clock

In suboptimal cases a lot depends on the input filtering. It is necessary to filter PTP Sync message timestamps and delay measurements for outliers; the tuning coefficients are modified if incoming timestamp jitter/offset is above threshold. Under ideal real world conditions (proper PTP-aware network devices, analog recorders are on the same subnode/switch) some of the described effects are marginal, and others are mitigated by the filtering capabilities of the tuning algorithm. This way, a 30 ns synchronization precision can be achieved against the GMC. In case of a direct connection an even better 5-10 ns precision was measured between two modules synchronizing against the same GMC.

Figure 8 shows the measurement picture of an oscilloscope sampling the 1 PPS signals of the

two Data Acquisition Units (C2 and C3) and the Grand Master Clock (C1). The trend-line F1 shows the delta delay of C2 against C3 with 100 measurements per division in the direction of the x axis. The time base in y axis is 5 nanoseconds per division. A jitter of the delta delay of ±5 ns is measured. For this measurement 3711 samples were taken at a measurement rate of 1 Hz. Within this one hour of measurement a mean delta delay of the 1 PPS signals of the two acquisition units of 7.65871 ns is achieved with a standard deviation of 2.17535 ns.

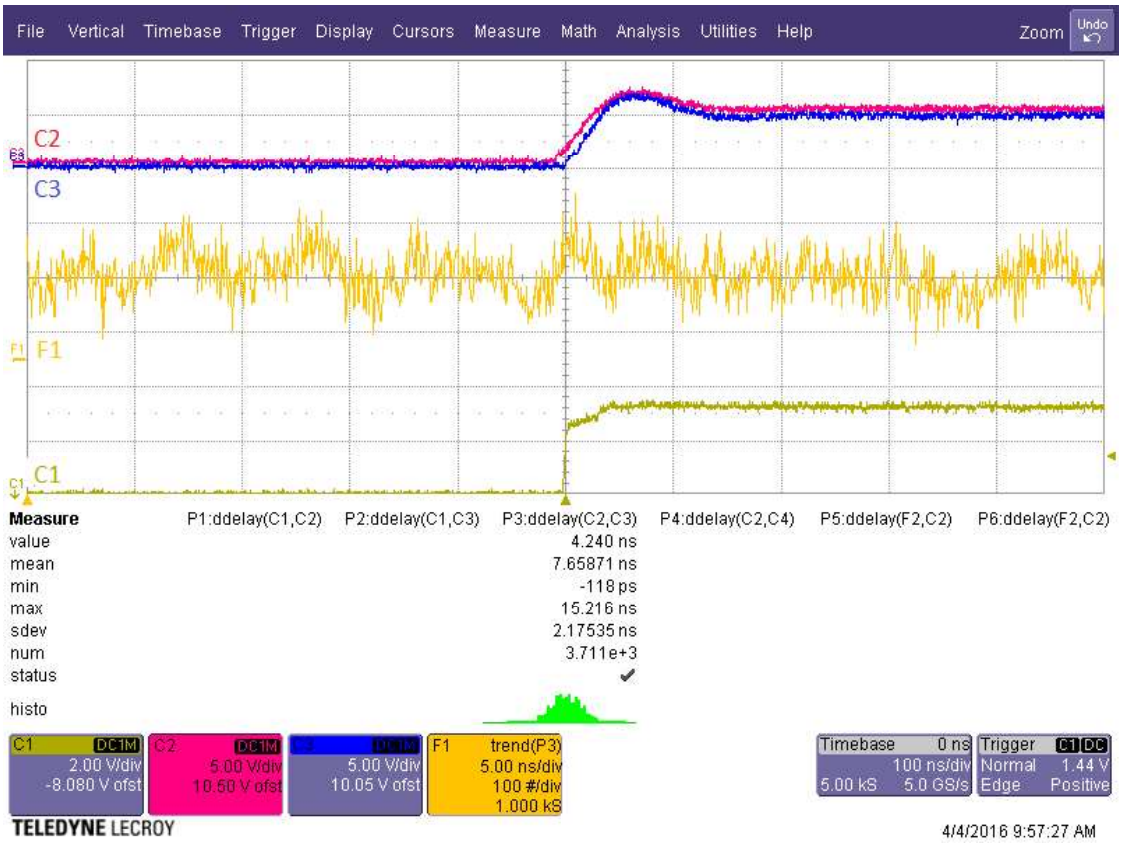

Fig. 8 1 PPS Offset between GMC and two synchronized Data Acquisition Units

Conclusion

It was demonstrated that with precise onboard oscillators and a good quality time source a distributed data acquisition system can achieve synchronous analog sampling accuracy of the 10 ns range by the means of the IEEE™1588-2008 Precision Clock Synchronization Protocol.

The precise synchronization is not instantaneous. In the current implementation it takes up to 5 minutes to achieve stable sync but there are several possibilities to reduce the time the different filters need for settling.

It should be noted that even though module-to-module the analog sampling can be extremely synchronous, this remains so only as long as the same algorithm with same parameters is used for synchronization. E.g. if one of the modules uses PID controller for PTP synchronization and the other uses a Kalman filter based approach [4] or a very simple FIR filter (e.g. present in the earlier IEEE™1588-2002 implementation), these will result in very different responses to the GMC clock changes and a 10 ns analog sampling precision as shown in Fig. 9, may become impossible in a heterogeneous recording environment.

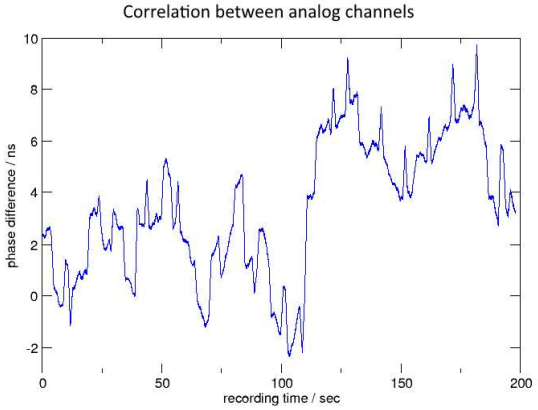

Fig. 9 Phase correlated PTP synchronized analog channels of two Data Acquisition Units

There are possible improvements over the basic implementations of the PTP utilizing Synchronous Ethernet (e.g. White Rabbit [5]) that target the sub-nanosecond synchronization range. These systems use expensive hardware and guarantee that the Ethernet clock frequency is the same throughout the network, removing some noise sources.

References

[1] Telemetry Standards, IRIG Standard 106-13 (Part1), Chapter10, http://irig106.org/docs/106-13/chapter10.pdf

[2] IEEE Std 1588™-2002, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, Print: ISBN 0-7381-3369-8, SH95023; PDF: ISBN 0-7381-3370-1, SS95023

[3] IEEE Std 1588™-2008, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, Print: ISBN 978-0-7381-5401-5, STDPD95773; PDF: ISBN 978-0-7381-5400-8, STD95773

[4] G. Giorgi, C. Narduzzi, Performance analysis of Kalman filter-based clock synchronization in IEEE 1588 networks, *International IEEE Symposium on Precision Clock Synchronization for Measurement, Control and Communication* 2009, doi: 10.1109/ISPCS.2009.5340221

[5] P. Moreira, J. Serrano, T. Wlostowski, P. Loschmidt, G. Gaderer, White Rabbit: Sub-Nanosecond Timing Distribution over Ethernet, *International IEEE Symposium on Precision Clock Synchronization for Measurement, Control and Communication* 2009, doi: 10.1109/ISPCS.2009.5340196