

Limits of Synchronization Accuracy Using Hardware Support in IEEE 1588

Patrick Loschmidt, Reinhard Exel, Anetta Nagy, Georg Gaderer
Research Unit for Integrated Sensor Systems
Austrian Academy of Sciences
Viktor Kaplan Strasse 2
A-2700 Wiener Neustadt
{Patrick.Loschmidt, Reinhard.Exel, Anetta.Nagy, Georg.Gaderer}@OEAW.ac.at

Abstract— Clock synchronization protocols for packet-oriented networks, like IEEE 1588, depend on time stamps drawn from a local clock at distinct points in time. Due to the fact that software-generated time stamps suffer from jitter caused by non-deterministic execution times, many implementations for high precision clock synchronization rely on hardware support. This allows time readings for packets with very low jitter close to the physical layer.

Nevertheless, approaches using hardware support have to carefully consider influences on synchronization accuracy when it comes to the range of nanoseconds. Among others, limits come from the update interval, oscillator stability, or hardware clock frequency. This paper enlightens the limits for such implementations based on an analysis of the influences of the main factors for jitter. The conclusions give hints for efficiently optimizing current implementations.

Keywords— IEEE 1588, accuracy, time stamping, synchronization interval, jitter, oscillator

I. INTRODUCTION

When synchronizing clocks over packet-oriented networks, the main problem hindering high accuracy is the jitter introduced by the uncertain processing time of the various OSI network layers involved in the transmission. Therefore, IEEE 1588 defined the start-of-frame delimiter of a packet as the reference point to draw a time stamp. Up to now, this is usually done in hardware on the MII (Media Independent Interface) in the case of Ethernet.

As shown in figure 1, the MII Scanner (MIIS) is able to supply the Clock Synchronization Cell (CSC) with time stamps which do not contain the jitter of the upper layers. Consequently, the achievable accuracy just depends on the physical layer, the network between the nodes and the properties of the implemented hardware support. The latter are mainly the clock rate (resolution of the time stamp) and the stability of the oscillator used to generate the nodes internal clock.

In this paper a system consisting of two directly connected nodes equipped with hardware support on the MII (see figure 1) is used to analyze the remaining jitter sources. Both nodes have independent clocks and use oversampling (>25 MHz) on the MII. Other interfaces to the physical layer, like RMII (Reduced MII), are not considered since in that case the synchronization between the clock domains of the sender and

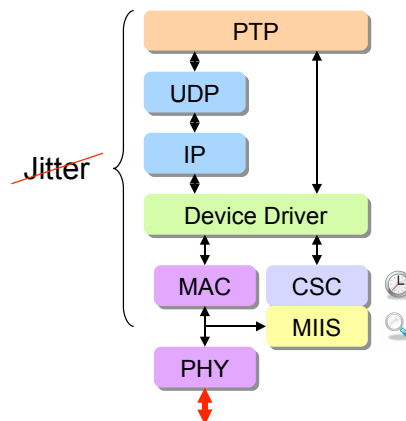


Fig. 1. Jitter sources and cancellation through hardware support

the receiver is done within the physical layer device. Consequently, the advantage of oversampling and using statistical methods to increase accuracy cannot be used.

Besides the ability to accurately time stamp packets relevant for synchronization, also the resolution for generating events heavily influences the usability of hardware solutions for high accuracy. The simplest example of such an event is the generation of the 1 PPS (one pulse per second) signal in order to allow external offset measurement of the nodes.

After an analysis of the remaining jitter sources in the next chapter, the paper further describes the influence of the synchronization interval on the control loop. While the settings of the used hardware approach are easy to control, the ideal parameters of the control loop also depend on the reaction time of the system. Thus, they have to be matched with the hardware parameters as mentioned in the introduction to chapter III. The measurements in this chapter show some achieved results using optimized systems, while the conclusions summarises up the findings on which optimization gives the best results depending on the allowed effort.

II. JITTER SOURCES AND UNCERTAINTIES

As mentioned above, due to the hardware supported detection of start-of-packet on the MII all jitter sources from

higher network layers can be compensated. Nevertheless, there are remaining jitter sources and uncertainties on and below the physical layer, which affect the achievable synchronization accuracy in IEEE 1588 networks. Additionally, the stability of the oscillator in combination with the period of synchronization messages also has major influence on the accuracy of the overall system.

A. Physical Layer Device

The physical layer device (PHY) is the link to the analogue data on the transmission line. On the MII, the receive and transmit data paths are operated with independent clocks. For the transmitter the data is transferred in the form of nibbles at 25 MHz via the MII and then transmitted using 125 MSymbols/s on the line. The clock for transmitting is formed out of the externally applied 25 MHz clock.

On the receiver side, the clock of the transmission line is recovered and downsampled in the PHY to derive the interface clock on the receive path. This downsampling causes an unpredictable, but constant delay, since there are five possible clock edges in the 125 MHz clock domain to derive the 25 MHz receive interface clock. The behaviour of this jitter is heavily dependent on the implementation. The selected edge can change on a new link establishment or even with each packet. Therefore, some PHY manufacturers [1] align the MII clock to the data on the transmission line to keep the delay constant.

Concerning jitter, the primary source may be found in the generation and recovery of the 125 MHz transmission clock. The latter is done using internal PLLs, which add approximately 100 ps jitter, even on specially designed PHYs.

B. Hardware Time Stamping

Using an Adder Based Clock (ABC) design [2], in general there is no inaccuracy when adjusting the clock, since just the clock's rate gets updated. Nevertheless, the granularity of the clock readout is limited due to the clock frequency of the ABC. Hence, all time stamps will inherently jitter equally distributed over one clock period T_{clock} , which adds the standard variance of an equally distributed event $\sigma_{\text{TS}}^2(T_{\text{clock}}) = \frac{T_{\text{clock}}^2}{12}$ to the system. The jitter can be reduced by boosting the clock frequency or by using other schemes to increase the effective resolution.

One feasible approach is to run a small counter with an even multiple of the clock frequency of the ABC. Consequently, its count can be used together with the ABC's last time value to achieve fine grained time stamps. This solution has the advantage, that only a small part of the system has to run at high clock frequency.

C. Oscillator

The physical properties of oscillators are an essential parameter for the accuracy of a clock synchronization node. In general, the stability of standard Crystal Oscillators (XOs), as well as the optimized versions like Temperature Compensated Oscillators (TCXOs) and Oven Controlled Oscillators (OCXOs) is dependent on the sampling period.

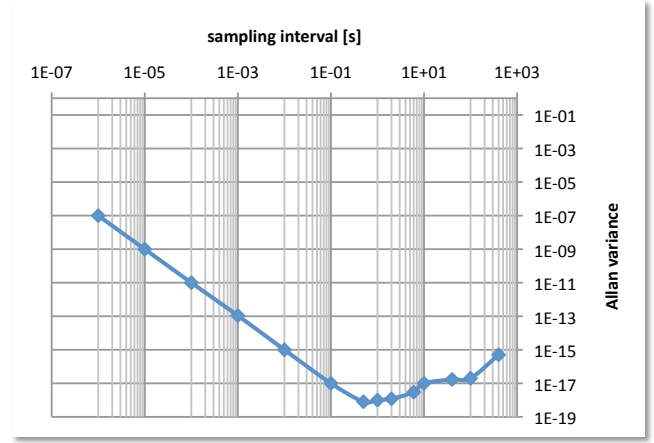


Fig. 2. Typical Allan variance of a standard crystal oscillator (measured)

In order to account for this fact, the Allan variance [3], which is defined by $\sigma_y^2(\tau) = \frac{1}{2} \langle (y_{n+1} - y_n)^2 \rangle$, where y denotes the normalized frequency averaged departure over one sample period and τ the sampling period, is used to characterize oscillators. A low Allan variance means good stability over the measured period. The frequency can be calculated as $y_n = \frac{1}{\tau}(x_{n+1} - x_n)$, with x_n as the time error, at sample number n . Hence, the Allan variance can be calculated by the measured period time of a clock as $\sigma_y^2(\tau) = \frac{1}{2\tau^2} \langle (x_{n+2} - 2x_{n+1} + x_n)^2 \rangle$.

A typical graph in double logarithmic scale for the Allan variance is shown in figure 2. For short periods it will decrease by 100, if the sample time is increased by a factor 10. For large periods long term physical effects (e. g., temperature changes, aging) increase the value of the Allan variance.

Figure 3 shows the distribution of 2000 samples of the frequency of two different oscillators (XO: 50 ppm, OCXO: 0.3 ppm). The frequency was obtained by calculating the mean value over 1 s. Both were measured in a temperature stabilized environment (± 0.05 K). Nevertheless, the measurement points out that the behaviour of the oven controlled oscillator is still better than the XO, although temperature effects were canceled out. Unfortunately, the resolution of the measurement device does not allow to accurately determine the stability of the OCXO.

D. Synchronization Interval

Another interesting parameter in a synchronization system is the equidistant interval T_{sync} for exchanging synchronization messages. Both, the master and the slave are equipped with an oscillator that has a drift δ . The presence of the varying drift of the oscillator will increase the offset between the nodes with the progression of time. If we assume that two oscillators were perfectly synchronous at a given time, the offset α after T_{sync} is $\alpha = \int_{T_{\text{sync}}} (\delta_M + \delta_S) dt$.

If the Allan variance of the oscillator is known, the standard variance of α can be estimated. The two-sample Allan variance can be interpreted as a variance of the frequency

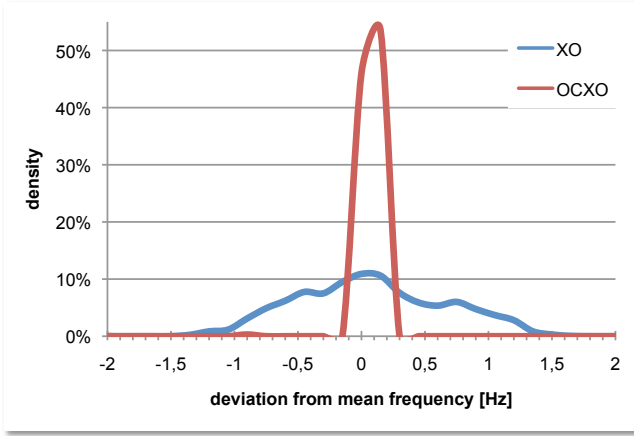


Fig. 3. One second mean frequency distribution of two oscillators

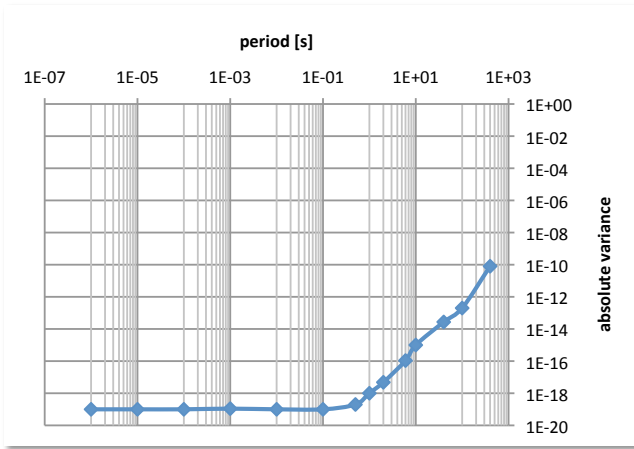


Fig. 4. Absolute variance σ_{abs}^2 of the oscillator over the sampling time τ

change rate. Therefore, the frequency variance is the integral of the Allan variance over T_{sync} plus a frequency offset, the standard variance σ_{δ}^2 for the drift. Another integration results in the time variance plus a time offset. If we assume that the frequency and time offsets are compensated in the synchronization system's servo, the variance of the time error is just the denormalized Allan variance (the normalized Allan variance multiplied by T_{sync}^2). Since the master and the slave oscillator are statistically independent, the resulting variance of α is the sum of the standard variances of both nodes.

For a typical oscillator with an Allan variance as shown in figure 2, the absolute variance for short intervals remains constant (as depicted in figure 4). Thus, just from the oscillator's point of view one gains no additional accuracy if the system synchronizes more often. For long synchronization intervals the standard deviation rises significantly, which is mainly due to long term effects like temperature changes.

Consequently, the synchronization period should be maximized in a way to still be before the flicker floor of the Allan deviation is reached. This is equal to the end of the zero slope part of the absolute variance. From the oscillator's point of

view there is no difference on which point of the zero slope part the synchronization period is chosen, which means that there is no further accuracy gain by exchanging more timing messages than necessary. Furthermore, long synchronization intervals are not recommended. PTP version 1 (IEEE 1588 2002) has a default synchronization interval of 2 seconds and allows synchronization intervals between 1 and 16 seconds, incrementing by a factor of 2. As pointed out by figure 4, even the shortest possible period, according to the IEEE 1588 standard, of 1 s is quite high for the tested oscillator resulting in a degradation of accuracy.

III. MEASUREMENTS

For experimental verification a test system was set up. It is capable of drawing time stamps with 2 ns and generating the 1 PPS signal with 1 ns granularity. To compare the stability, two oscillators were used as a clock source: an OCXO (0.3 ppm) and a XO (50 ppm), respectively. All results concerning the overall system accuracy were obtained by monitoring the 1 PPS pulse on an oscilloscope with 20 GS/s. For the synchronization interval T_{sync} various values ranging from 7.8 ms up to 8 s were used.

If we assume that all jitter sources described in last section are statically independent, the resulting variance of the jitter is just the sum of all the variances. Consequently, the effective variance that can be measured on the oscilloscope (not taking the measurement system with 50 ps granularity ($\sigma_m = 14$ ps) into account, can be calculated as

$$\sigma^2 = 2 [\sigma_{1\text{PPS}}^2 + \sigma_{\text{osc}}^2(T_{\text{sync}}) + \kappa(T_{\text{sync}})(\sigma_{\text{TS}}^2 + \sigma_{\text{PHY}}^2)].$$

Due to the 2 ns resolution of the time stamps the corresponding factor has a value of $\sigma_{\text{TS}} = 577$ ps, whereas the 1 PPS signal adds $\sigma_{1\text{PPS}} = 289$ ps. Both values are calculated under the assumption, that the actual occurrence of the event is equally distributed over the interval between the two instances in time to generate / detect the event.

For the SMSC LAN 8700 physical layer device [4] used in the experimental setup a standard deviation of $\sigma_{\text{PHY}} = 124$ ps using the standard oscillator and $\sigma_{\text{PHY}} = 100$ ps for the OCXO respectively was taken. Figure 5 shows the distribution of the delay between two PHYs measured between the transmit enable (TX.EN) and the receive data valid (RX.DV) signal of the MII. In order to achieve appropriate results, a frequency counter with 25 ps resolution supplied by a GPS disciplined Rubidium frequency standard was used.

As modern synchronization protocols, like the mentioned PTP, feature delay compensation, the mean value of the distribution is compensated by evaluating the round trip time. This is of high importance, since every time a link is (newly) established, the PHY's internal receive and transmit filter parameters may be recalculated, which leads to changes in multiples of 8 ns in the absolute line delay, due to the lock-in of the receiver PLL to the 125 MHz clock on the line.

It has to be mentioned, that the value of σ_{PHY} has been chosen as a combination of jitter sources introduced by the physical layer. These include the oscillator supplying the

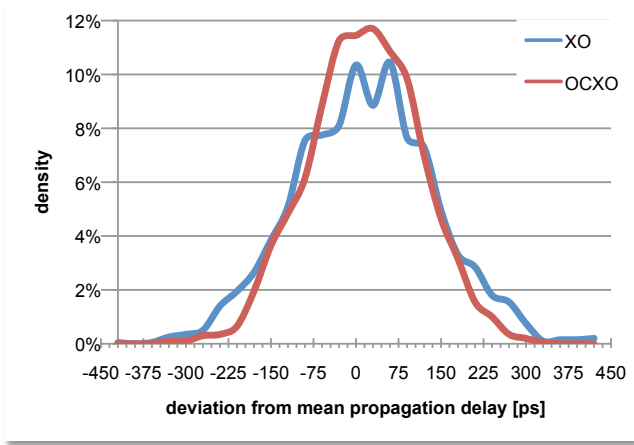


Fig. 5. PHY to PHY delay and it's dependency on the oscillator source

transmitting PHY on the transmitting side as well as the chip internal PLL and the loop filter on the receiver side.

Not only the oscillator's variance is dependent on the observed synchronization interval T_{sync} , but also the last term representing the quality of the time stamps with a factor κ . The latter describes the implication of the control loop on the standard deviation. Since the control loop parameters must match the behaviour of the oscillator, which depends on the synchronization interval, κ itself is a function of T_{sync} .

A. Oscillator Quality

Figure 6 depicts the observed standard deviation for a 50 ppm XO and an 0.3 ppm OCXO for different synchronization intervals. The appropriate proportional factors for the PI-controller adjusting the ABC's rate are shown as well.

For short synchronization intervals the control loop parameters can be set rather aggressive to compensate quickly for any change of the oscillator's frequency, but for long intervals (more than 1 s) such servo parameters lead to oscillations and instability of the control loop. These oscillation are partially caused by the large absolute offset which can occur due to the drift during the long synchronization interval and the inability of a PI-controller to accurately predict the future behaviour of the oscillator. The combination of the relaxed control loop parameters and the high time error variance of the oscillator result in an increased standard deviation.

B. Time Stamp Accuracy

The granularity of the time stamps has a major influence on the accuracy, since it defines the resolution, with which a node is able to detect the point in time of the transmission or reception of a network packet. If the generation of the time stamps is neither correlated between sender and receiver nor to its local adder based clock, then the effective time stamp variance is $\sigma_{\text{TS}}^2(T_{\text{clock}}) = \frac{T_{\text{clock}}^2}{12} T_{\text{sync}}$.

The statistical independence may not be satisfied in every case. For instance, if the master has a clock increment, which is a multiple of the time stamp granularity, the observed and

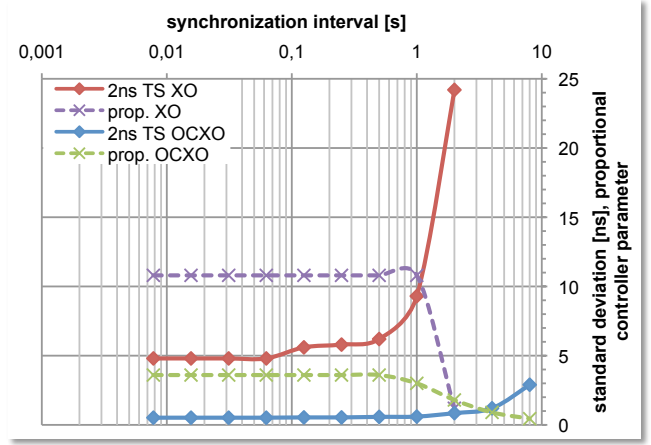


Fig. 6. Time error standard deviation for a XO and an OCXO including the proportional factors for the PI-controller

transmitted time stamp is variance free (assuming that the PHY's send clock is supplied by the same clock source as the ABC). On the other hand, master and slave may be correlated, if both round up or down the time stamp at the same time due to the almost same clock rate.

For short intervals the granularity is of minor importance, since the effective time stamp variance is low since the control loop can average over several time stamps and therefore reduce the jitter by statistical means. If highly stable oscillators or long synchronization intervals are used, low variance time stamps are getting increasingly important due to the relative higher effect on the system variance.

The influence of the time stamp variance becomes dominant, if $\kappa(T_{\text{sync}})(\sigma_{\text{TS}}^2 + \sigma_{\text{PHY}}^2)$ gets in the range of the variance of the oscillator $\sigma_{\text{osc}}^2(T_{\text{sync}})$. For long synchronization intervals the impact on the accuracy with different time stamp granularities is shown in figure 7.

If a certain time error variance is required, there are several possible ways to achieve this: a better oscillator, shorter synchronization interval or a higher granularity of the observed time stamps. The measurements point out that a good oscillator and a short synchronization interval will always help to improve the observed clock variance, whereas fine granular time stamps improve the variance not in every case by a significant amount.

Finally, it has to be stated, that an accurate mathematical model of the investigated problem is difficult to derive. This is mainly due to the behaviour of the oscillator and analogue parts within the transmission chain (e.g., PHYs, PLLs) because of their hidden internal parameters, which are only known to the chip manufacturer.

For the chosen system the theoretical limit is $\sigma = 410$ ps, if $\kappa = 0$ and a perfect oscillator is assumed. Nevertheless, the lowest standard deviation that was possible to achieve was 520 ps using the OCXO with a synchronization interval of $1/128$ s.

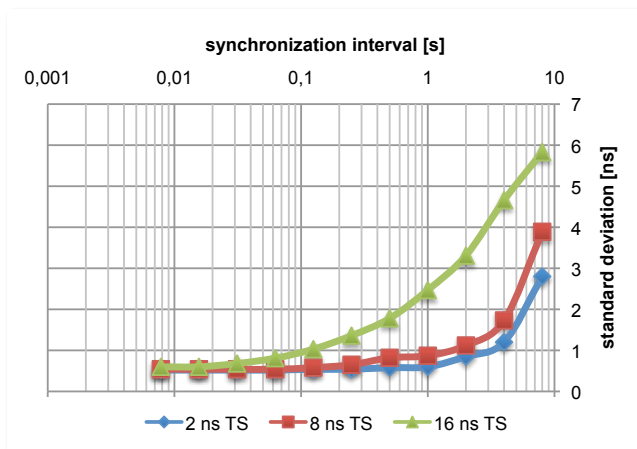


Fig. 7. Overall standard deviation σ of two IEEE 1588 synchronizing nodes, measured via the 1 PPS signals

IV. CONCLUSION

The accuracy of an IEEE 1588 synchronization system is affected by several jitter sources: obviously the most important of all, the local oscillator. For the chosen synchronization interval its frequency jitter must be as low as possible. Otherwise the system drifts away during the synchronization interval in an unpredictable way.

In order to optimize an existing synchronization system, the best way is to characterize the oscillator driving the node's internal clock. If the Allan variance for the source oscillator is known, the absolute variance can be calculated. The latter typically shows a zero slope part for short measurement periods τ .

Consequently, the synchronization interval should be chosen long enough to minimize the network traffic for the system, but short enough to operate the oscillator in its optimum accuracy range. As shown in figure 6, the synchronization interval has to be lowered till a certain optimal period, which is dependent on the quality of the oscillator. Further decreasing of the message exchange interval doesn't give any reasonable additional improvement. Based on the measurement, this optimal sampling time was chosen to be at an interval of 0.5 s which gives appropriate results for both of the two oscillators very much varying in their internal jitter behaviour.

Additionally, it has been shown, that for a typical interval of 2 s the quality of the time stamps is an essential parameter. For such a system time stamping with high granularity and a control loop optimised for the oscillator are possible improvements. In any case, the characterisation of the used oscillators as well as the dimension of the other jitter sources is essential for optimising networked clock synchronization systems using the formula given in section III.

V. ACKNOWLEDGMENTS

The work is partially co-funded via the Austrian FIT-IT Project ϵ -WiFi (Nr. 913310) and from the European Commission under the European Regional Development Fund according to article 4. Further, the authors would like to thank Oregano Systems - Design & Consulting Ltd. for their project support.

REFERENCES

- [1] N. Semiconductor, *DP83848 and DP83849 Data Latency Application Note*, Aug. 2006.
- [2] U. Schmid, M. Horauer, and N. Kerö, "How to distribute gps-time over cots-based lans," in *Proceedings of the 34th IEEE Precise Time and Time Interval Systems and Application Meeting (PTTI'02)*, 2002.
- [3] D. B. Sullivan, D. W. Allan, D. A. Howe, and F. Walls, "Characterization of clocks and Oscillators," NIST, Technical Note 1337, 1190.
- [4] SMSC, *SMSC LAN8700/LAN8700I datasheet*, Aug. 2006. [Online]. Available: <http://www.smsc.com/main/catalog/lan8700.html>