Upgrading the Timing System of ASDEX Upgrade

M. Gehring, M. Astrain, H. Eixenberger, B. Sieglin, W. Treutterer, T. Zehetbauer,
the ASDEX Upgrade Team

***Abstract*—During more than 30 years of operation at ASDEX Upgrade, the timing system evolved with the requirements of its users as well as technical feasibility. Currently, the next revision of the timing system is being prepared and implemented. It aims at replacing custom build hardware with modern and open standards like the Precision Time Protocol and Synchronous Ethernet as well as off-the-shelf hardware. With tight schedules between experiment campaigns, replacing the hardware timing system all at once is not possible. Instead, network time synchronization will be offered in parallel to allow a step-wise migration of the timing system. Operating this intermediate *hybrid timing system* requires to keep the two time domains synchronous within the time accuracy of the current hardware timing system. Here, the implementation of the hybrid timing system as well as its performance will be discussed. Also, an outlook on the next steps for the full migration of the timing system will be given.**

***Index Terms*—ASDEX Upgrade, CODAC, Precision Time Protocol (PTP), Synchronous Ethernet (SyncE), Timing System**

I. INTRODUCTION

MANY devices at ASDEX Upgrade (AUG) including actuators, the Discharge Control System (DCS) [1] as well as real-time plasma diagnostics [2] need to operate time-synchronous for successful plasma discharges. Also, a common time base is required for the post-discharge evaluation of data acquired by all in all more than a hundred diagnostic systems, where timing information allows to align and correlate measurement data with events that occurred during the discharge. When ASDEX Upgrade was designed and commissioned, no off-the-shelf solution was available and therefore custom hardware was designed for its timing system. Maintaining this hardware gets increasingly difficult, as many of its electronic components are no longer in stock and finding spare parts or suitable replacements is becoming difficult. Today, network time synchronization can achieve sub-nanosecond accuracy using open standards like Precision Time Protocol (PTP) and Synchronous Ethernet (SyncE). So with a scheduled renewal of the network infrastructure, it was decided to migrate to network time synchronization.

However, two main reasons prevent an all at once replacement of the timing system. The timing system is mandatory for the operation of ASDEX Upgrade and schedules are tight between experiment campaigns. With the amount of involved hardware, replacing the full timing system including appropriate testing at one moment seems not feasible. Also, many systems in operation require specialized hardware, that on its own is not prepared for network time synchronization. So instead, network time will be provided time-synchronous to the hardware timing system as additional source of synchronization. The successful implementation and operation of this setup, internally referred to as hybrid timing system, now allows a step-wise migration from the current, custom build timing system to network time synchronization.

Section II provides a brief introduction to the quantities used within this work to evaluate the performance of the hybrid timing system. Section III will first give an overview of the hardware timing system of ASDEX Upgrade, before introducing the hybrid timing system. Section IV discusses the performance of the hybrid timing system and Section V will end with a discussion of the results as well as an outlook on next steps in the full migration of the timing system.

II. TIME MEASUREMENT AND TIMING ACCURACY

Time synchronization as well as the evaluation of the accuracy of timing devices is an extensive field of study. For the purposes of this work, however, time error and time interval error provide sufficient metrics to evaluate the performance of the hybrid timing system. They will be introduced below. For detailed information on time and frequency analysis, refer to standards or handbooks such as [3].

In a nutshell, a clock monitors a periodic process of cycle time T or frequency $f = 1/T$. After N cycles of the periodic process have completed, a time interval of $\Delta t = N \cdot T$ has passed. If the device also keeps a record of time t_0, when counting was started, the time measured by the clock device is $t = t_0 + N \cdot T$. The cycle time T_{clk} or clock frequency f_{clk} of real-world devices will deviate from the expected value due to manufacturing tolerances, temperature dependent material properties, etcetera, such that $T_{\mathrm{clk}} = \big(1 + \varepsilon_T(t)\big) \cdot T$ or $f_{\mathrm{clk}} = \big(1 + \varepsilon_f(t)\big) \cdot f$. With cycle time and frequency being dependent variables, so are the error terms $\varepsilon_f = -\varepsilon_T/(1+\varepsilon_T)$ and cycle time (error) and frequency (error) will be used interchangeably from here on.

As illustrated in Fig. 1, the time interval Δt_{clk} measured by the clock will deviate from the ideal time interval Δt:

$$\Delta t_{\mathrm{clk}} = T \cdot N = T \cdot \left\lfloor \int_0^{\Delta t} \frac{\mathrm{d}t}{T_{\mathrm{clk}}(t)} \right\rfloor \approx \int_0^{\Delta t} \frac{\mathrm{d}t}{1 + \varepsilon_T(t)} \quad . \tag{1}$$

With the ideal time interval Δt not being measurable, clock devices are compared against reference devices instead. For the

---

This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No 101052200 — EUROfusion). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

M. Gehring is with the Max-Planck-Institute for Plasma Physics, 85748 Garching, Germany (matthias.gehring@ipp.mpg.de).

M. Astrain, H. Eixenberger, B. Sieglin, W. Treutterer and T. Zehetbauer are with the Max-Planck-Institute for Plasma Physics, 85748 Garching, Germany.

for the ASDEX Upgrade Team see author list of T. Pütterich et al., 2026 Nucl. Fusion 66 116002

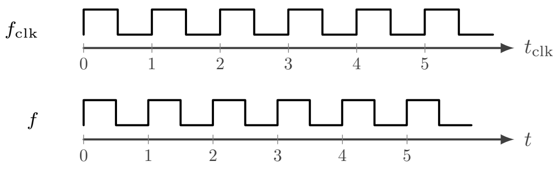

Fig. 1. Time interval error resulting from a clock frequency deviation.

reference device, the frequency is expected to be more accurate by orders of magnitude. The deviation between Δt_{clk} and a reference interval τ is called time interval error

$$\mathrm{TIE}(\tau) = \Delta t_{\mathrm{clk}} - \tau \approx \int_0^{\tau} \varepsilon_f(t)\,\mathrm{d}t = \langle \varepsilon_f \rangle \cdot \tau \quad . \tag{2}$$

Varying the length of τ allows to investigate different properties of the clock. For short intervals τ, the mean error $\langle \varepsilon_f \rangle$ of the clock frequency is dominated by random jitter, characterizing the intrinsic precision of the timing device. For long intervals, systematic deviations will dominate, allowing to characterize the long term accuracy of the device.

The deviation between time t_{clk} measured by the clock and the reference time t_{ref} is called time error

$$\mathrm{TE} = t_{\mathrm{clk}} - t_{\mathrm{ref}} \approx \underbrace{t_{0,\mathrm{clk}} - t_{0,\mathrm{ref}}}_{\text{initial offset}} + \underbrace{\int_{t_0}^{t_{\mathrm{ref}}} \varepsilon_f(t)\,\mathrm{d}t}_{\text{accumulated drift}} \quad . \tag{3}$$

For a clock device to provide an accurate measurement of time, a mechanism to minimize the initial offset on start-up as well as long term stability of the clock frequency are required.

III. THE ASDEX UPGRADE TIMING SYSTEM

When ASDEX Upgrade went into operation in 1991, neither off the shelf software nor hardware was available, that met the timing and synchronization requirements of the experiment. So a custom trigger-timer system, the Local Timer System, was developed [4], [5]. It uses a per discharge relative timescale aligned to plasma breakdown and runs on a cycle time of 100 ns. With a rework of the control and data acquisition (DAQ) infrastructure around 2002 [6], Time to Digital Converter (TDC) hardware was developed to build the TDC Timing System [7]. Its cycle time is 20 ns and the experiment time migrated from per discharge relative timescales to an absolute timestamp. The so called TDC time is the UNIX epoch expressed in nanoseconds. To avoid the additional work of retrofitting specialized hardware to the TDC Timing System, that did not require the new features provided by TDCs, the Local Timer System was integrated into the TDC Timing System instead of being replaced.

With the 2026 campaign at ASDEX Upgrade, the hybrid timing system went into operation. Its successful commissioning starts the step-wise migration from the current hardware timing system to network time synchronization.

A. The Hardware Timing System

The hardware timing system as well as its components are depicted in Fig. 2. So far, the majority of time clients still relies

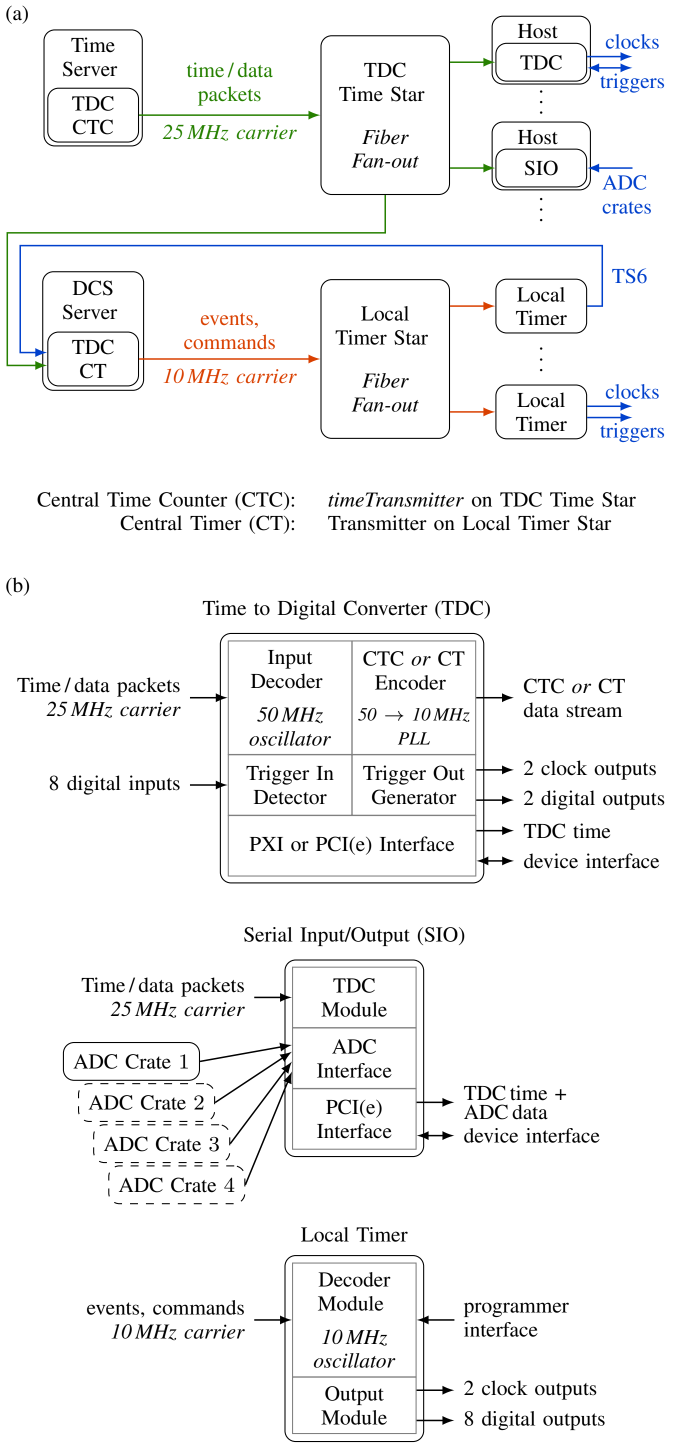

Fig. 2. The ASDEX Upgrade hardware timing system (a) and its timing devices (b). The ratio of carrier to local oscillator frequency (TDC: 25/50 MHz, Local Timer: 10/10 MHz) depends on the electronic component used for clock data recovery.

on Local Timer or Time to Digital and Serial Input/Output (SIO) hardware for:

1) Synchronization
   - Provided clock signals are used to syntonize to the reference clock
   - Distributed event notifications are used to synchronize to the discharge timeline
2) Timestamping
   - Hardware timestamping during data acquisition
   - Hardware timestamping of trigger inputs
   - Software timestamping on host devices

To eliminate drift within the hardware timing system, all clocks are traced to a common reference. This is shown in Fig. 3. The Central Time Counter (CTC) is the *timeTransmitter*, providing time information to the TDC Time Star as Manchester encoded data stream on a 25 MHz carrier. All TDC and SIO devices use the data stream for clock recovery and syntonize their local 50 MHz oscillators to the CTC. The Central Timer (CT) is the transmitter on the Local Timer Star. To embed the Local Timer System into the TDC Timing System, the former hardware CT module was reimplemented into TDC firmware. This allows to derive the 10 MHz carrier frequency for the Manchester encoded data stream on the Local Timer Star from the TDC's clock. Again, the Local Timers use the data stream for clock recovery to syntonize their 10 MHz oscillators. Therefore, all clocks within the hardware timing system get syntonized to the 50 MHz frequency of the voltage-controlled oscillator (VCXO) onboard the CTC.

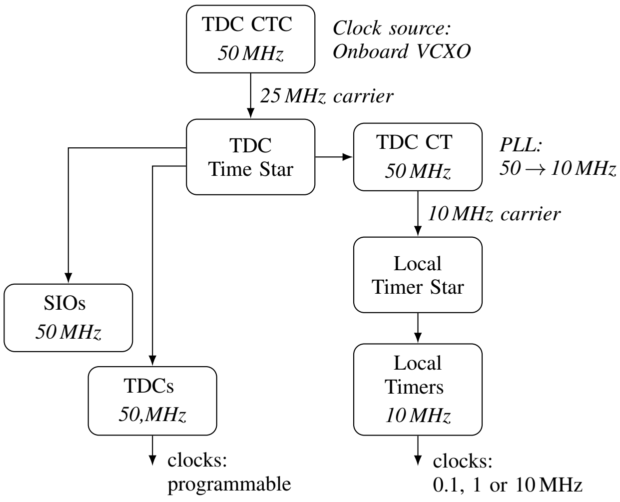

Fig. 3. Tracing all clocks to a common reference.

A (plasma) shot combines preparation tasks, the plasma discharge as well as final tasks. The timeline of shots follows recurring patterns of events. These events can be used to synchronize systems to the progress of the shot and are distributed by the Discharge Control System. Examples are:

DV2: diagnostics prepare for data acquisition
TS6: plasma breakdown
TS8: no more coil currents in the poloidal field coils allowed (guaranteed end of plasma discharge)

When there was only the Local Timer System, its hardware triggers were the only means to receive such event notifications, hence they are called Timer System Events (TSE). Today, Timer System Events are also distributed via the network (DCS signals, MQTT and multicast). Figure 2(a) shows the event synchronization with the Local Timer System as well as the alignment with TDC time. Each Local Timer is individually programmed to react to the distributed Timer System Events by setting or resetting a subset of its eight trigger outputs on some of the TSEs, while ignoring others. To align Timer System Events with TDC time, DCS keeps a record of TDC time when it sends a TSE. By timestamping TS6 on one of its TDC trigger inputs and comparing the arrival time to the time when TS6 was send, the processing delay on the Local Timer Star is accounted for. Where required for accuracy, also the transmission delay on the fiber runs can be accounted for in the alignment of TSE triggers with TDC time.

Last but not least, timestamping of measurement data and providing TDC time to host devices is what the TDC and SIO hardware was designed for. To elevate a TDC to the Central Time Counter, first its time register has to be initialized (initial offset) and then the CTC mode is started. The time register is then incremented by 20 ns for each tick of the onboard VCXO of the CTC and the current value of the time register is distributed to the TDC Time Star in configurable intervals. Each TDC or SIO device on the TDC Time Star receives the time packets and synchronization to the distributed time has to be triggered manually. After synchronization each TDC / SIO continues to increment its time register based on its own (syntonized) local oscillator. In case the value of its time register deviates from the distributed TDC time the TDC or SIO will flag de-synchronization, but will continue to use its local time to for example tag measurement data. This design was chosen to guarantee continuous timestamping of measurement data during plasma discharges, even in case of failure of the CTC or the TDC Time Star.

B. The Hybrid Timing System

For the hybrid timing system shown in Fig. 4, two redundant reference clocks were installed. They are equipped with high-

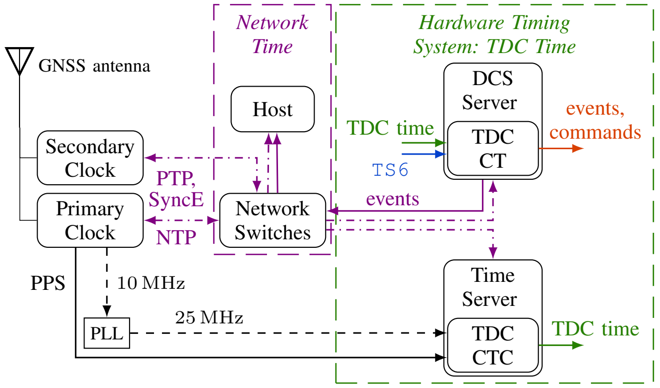

Fig. 4. The hybrid timing system merging the network time domain and the TDC time domain (hardware timing system).

quality oven-controlled oscillators (OCXO) and use time information distributed by the Global Navigation Satellite System (GNSS) to align their time registers with International Atomic Time (TAI) as well as to further condition the frequency of their oscillators. The reference clocks provide Network Time Protocol (NTP), Precision Time Protocol and Synchronous Ethernet to account for different accuracy requirements of time clients on the network. This is the network time domain of the hybrid timing system.

For the syntonization and synchronization of the hardware timing system with the reference clocks, their pulse per second (PPS) signal and 10 MHz frequency output are used. The reference frequency is converted to the 25 MHz carrier of the TDC data stream and fed to the CTC's input. Its clock data recovery stage will discard the decoded all zero data stream, but use the carrier to condition the VCXO of the CTC. With syntonization and therefore the long-term accuracy of time synchronization taken care of, aligning the time register of the CTC during startup will provide accurate time synchronization between the hardware timing system and the network time domain. Due to delays and jitter in setting the time register and starting the CTC mode, an iterative process is implemented to remove the offset:

1) Synchronize the time server using PTP.
2) Initialize the CTC's time register with the current system time and start the CTC mode.
3) Timestamp the PPS signal of the primary reference clock to determine the offset between TDC time and GNSS time.
4) Exit, if the offset is within a target range or too many iterations were reached.
   Else, calculate a correction term to add to the system time for the next iteration and go back to 2.

To guarantee the operation of the hybrid timing system even in case of issues with the PPS signal, a software-only process for offset removal is implemented as well. Instead of reading the timestamp of the PPS trigger input, it will compare the current, PTP-synchronized system time and TDC after starting the CTC to calculate the correction term. This, however, requires careful analysis of latencies in reading the TDC time as well as tuning of the host hardware and operating system to minimize latency and jitter introduced from that end.

IV. PERFORMANCE OF THE HYBRID TIMING SYSTEM

Analysis of the performance will follow the same steps that were outlined when introducing the timing system. First, the syntonization of the time domains will be investigated. Then, the offset removal will be analyzed, before additional topics will be looked at.

As shown in (2) and (3), usually a timing device will be compared to a higher accuracy reference clock. Here, however, the synchronization between the network time domain and the TDC time domain is the main focus. Hence, time error and time interval error will be expressed in terms of the host's system clocks (`CLOCK_REALTIME` for absolute time t_{sys}, `CLOCK_MONOTONIC` for measuring intervals Δt_{mono}) as well as TDC time t_{TDC}:

$$\mathrm{TIE} = (t_{\mathrm{TDC},i+1} - t_{\mathrm{TDC},i}) - (t_{\mathrm{sys},i+1} - t_{\mathrm{sys},i}) \quad \text{and} \tag{4}$$

$$\mathrm{TE} = t_{\mathrm{TDC},i} - t_{\mathrm{sys},i} \quad . \tag{5}$$

For all results shown below, time samples were taken in 1 s intervals, where Δt_{mono} is recorded to account for jitter in the acquisition loop. If estimates for a quantity are given below, either the percentile (e.g. Q_1 for 1 % percentile) is specified or a $3\,\sigma$ interval is given.

With only few spare TDC devices available, the analysis of syntonization and offset removal was carried out on one TDC card and a test server. Further testing requires the use of production hardware and therefore was and will be limited to the time between experiment campaigns.

A. Eliminating Drift

The most important goal for the hybrid timing system is the syntonization of both time domains. If successful, drift between clocks in either domain will be eliminated. The results of syntonizing the test system is shown in Fig. 5 (a). After setting the TDC into CTC mode, the VCXO was first left free-running counting time based on its intrinsic cycle time. In free-run, that TDC exhibits an average drift rate of 107 μs/s or 107 ppm frequency error compared to the PTP-synchronized system clock. After roughly 1.5 hours, the frequency reference from one of the GNSS clocks was connected to the TDC. Immediately, the drift between the two time domains disappears.

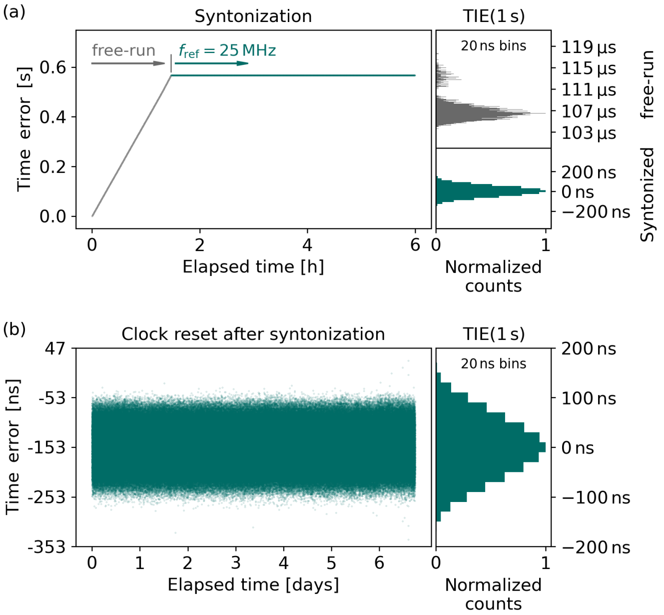

Fig. 5. Removing drift from the TDC in the test setup (a) and long-term stability of the clock syntonization (b).

Figure 5 (b) shows the long-term stability over a 7 day period. Before acquisition was started, the TDC was reset into CTC mode using a roughly estimated and preprogrammed, constant offset instead of offset minimization. During the test period, the time error stayed within $-153 \pm 123\,\mathrm{ns}$. With no long-term drift, this jitter is of random nature and the syntonization of the TDC to the reference clock works as expected.

B. Minimizing the Initial Offset

Next, the implemented hardware assisted (PPS) as well as software-only offset removal processes were investigated. With the $20\,\mathrm{ns}$ cycle time of the TDCs, offsets between GNSS time and TDC within this interval cannot be resolved. Therefore, the target offset for the iterative process is set to $\pm 10\,\mathrm{ns}$.

As shown in Fig. 6, resetting the TDC into CTC mode with the PPS signal available reliably reaches to target offset for 500 consecutive runs. If the software-only process is used, the final offset is not guaranteed to stay within the target interval, but does not exceed $\pm 40\,\mathrm{ns}$. As mentioned before, the results of the software-only process do highly depend on careful tuning of the host system and characterization of the involved TDC device. Therefore, it is only implemented as fall-back solution, in case the CTC has to be reset, while no PPS signal is available.

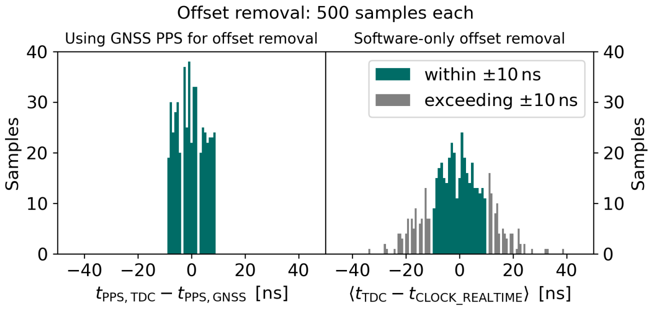

Fig. 6. Removing the initial offset with hardware support and purely in software.

C. Jitter Analysis

With no long-term drift remaining and the offset removal working, the focus moved to the jitter in the time (interval) error between TDC time and the PTP-synchronized system clock. The analysis aims to answer whether the jitter is caused primarily by the TDC, e.g. the conditioning of the VCXO, or the host system.

First, the PPS signal was continuously timestamped and its arrival recorded for the TDC and host used in the test setup. The offset between full second of TDC time and GNSS time is shown in Fig. 7. This offset only varies $\pm 20\,\mathrm{ns}$ or ± 1 clock tick around the average offset and for the majority of samples it does not deviate from the mean value. Therefore, the conditioning of the TDC's oscillator does not seem to be a major cause for the jitter seen in Fig. 5.

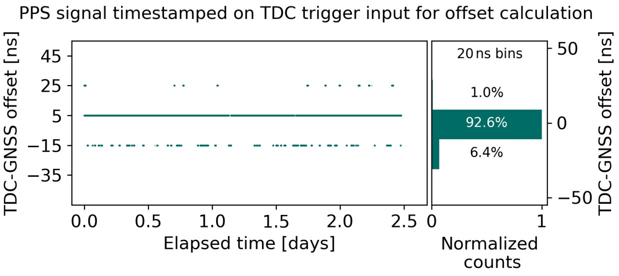

Fig. 7. Offset between TDC time and GNSS time (PPS timestamping).

As a next step, the time (interval) error was measured on additional hosts and TDC cards. As one host / TDC pair showed significantly lower jitter, the TDC in that host was exchanged and the measurement repeated. The results are shown in Fig. 8. So far, the results are not fully conclusive. However, with the jitter of the host clock ($\Delta t_{\mathrm{sys}} - \Delta t_{\mathrm{mono}} \leq 0 \pm 4\,\mathrm{ns}$) being significantly lower than the jitter between t_{sys} and t_{TDC} as well as with the strong variation of jitter between two different TDC cards installed in the same host, also the host device does not seem to be the major cause for the jitter. Currently, the jitter is attributed to the process of reading TDC time on the host system via the PCI(e) interface. Analysis so far showed, that calling the device driver to read the TDC time register takes between 2 and 3 μs to complete. This means, that the returned TDC time is outdated by some amount of time within the runtime of the system call. This is the cause for some of the required tuning necessary for the software-only offset removal process that was mentioned before. An additional source of jitter might be related to the clock domain crossing within the TDC's FPGA. To separate the $50\,\mathrm{MHz}$ TDC clock domain from the PCI(e) clock domain, registers are continuously mirrored between both domains.

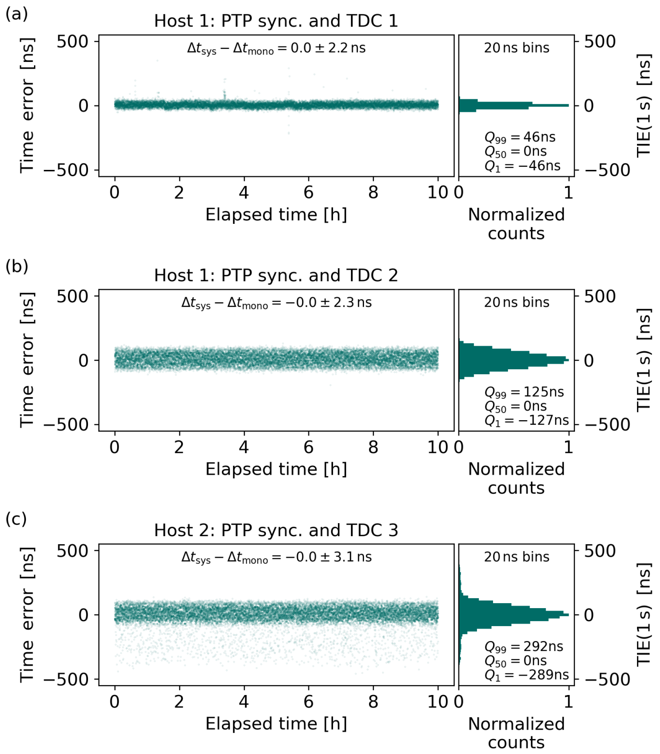

Fig. 8. Comparison of three TDC timing cards (a-c), all hosts using PTP for synchronization of the system clock.

The next opening between experiment campaigns will be used for additional testing. Recently, a network interface card

(NIC) was purchased, that provides hardware support for SyncE and is equipped with hardware outputs for a PPS signal as well as a 10 MHz clock. This will allow the use of dedicated measurement equipment to further investigate the jitter between TDC hardware and network time domain.

D. Comparison of Time Protocols

It is planned to also run NTP-synchronized devices within the network time domain. Although such devices have lower accuracy requirements (e.g. only logging done on host device), also the time (interval) error between a NTP-synchronized host and the TDC time was analyzed. For the results shown in Fig. 9, data was recorded simultaneously for one host NTP-synchronized to the campus NTP pool (a), one host NTP-synchronized to the reference clocks (b) and one host PTP-synchronized to the reference clocks. To no surprise, synchronization to the campus NTP pool results in significant drift between system clock and TDC time. However, using the GNSS clock as Stratum 1 NTP time source within a flat network hierarchy and slightly adjusting NTP parameters on the time client, the time error can be kept within a few microseconds and the time interval error only slightly degrades compared to the PTP-synchronized systems. This level of accuracy while using NTP came to some surprise, but is gladly accepted.

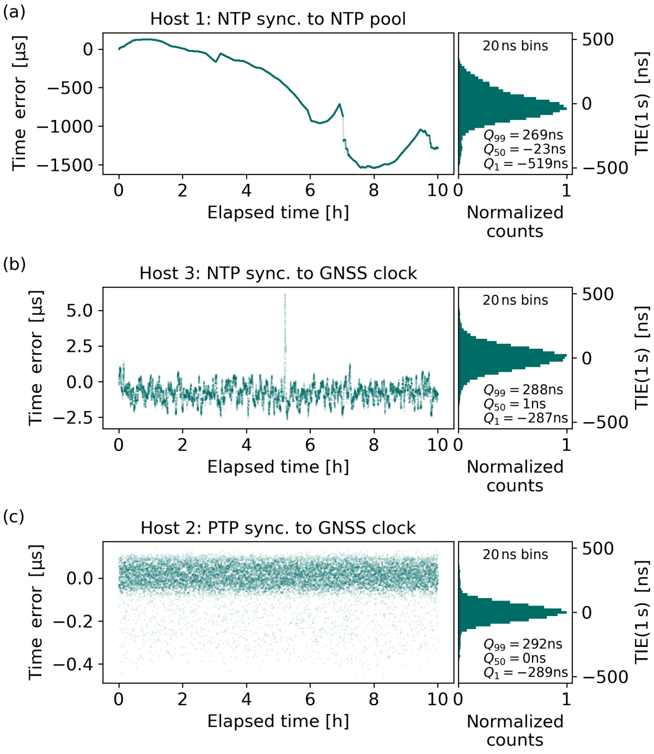

Fig. 9. Comparison of hosts using different methods of network time synchronization: (a) NTP synchronization to the institute NTP pool, (b) NTP synchronization to the GNSS clocks and (c) PTP synchronization to the GNSS clocs.

V. CONCLUSIONS

The hybrid timing system was successfully implemented and commissioned at ASDEX Upgrade. The performance allowed to migrate first demonstration systems to network time synchronization for the currently ongoing 2026 experiment campaign. Depending on the accuracy requirements of the time client, different protocols from NTP to PTP or PTP with SyncE can be used. Further testing is planned to pinpoint the root cause(s) for the remaining jitter as currently measured between host devices and TDC time domain. For the current demonstration systems, the jitter is below the accuracy requirements of those devices. For diagnostics using high sampling rates during data acquisition, however, this behaviour needs to be fully characterized.

Next steps also include contributions to the ongoing network infrastructure upgrade, to have the requirements of the network timing system accounted for during the upgrade. Once the new hardware and network concept is implemented, high accuracy network time synchronization can be provided to all of ASDEX Upgrade, instead of currently a limited set of demonstration systems. Also, a successor to the Local Timer and TDC hardware is planned. These devices will translate between network time domain and time clients using hardware clock and trigger signals for synchronization.

REFERENCES

[1] W. Treutterer, R. Cole, K. Lüddecke, G. Neu, C. Rapson, G. Raupp, et al., "ASDEX Upgrade Discharge Control System – A real-time plasma control framework," *Fusion Engineering and Design*, 89(3), 146-154, 2014, 10.1016/j.fusengdes.2014.01.001.

[2] K. Behler, H. Blank, H. Eixenberger, A. Lohs, K. Lüddecke, R. Merkel, et al., "Real-time diagnostics at ASDEX Upgrade - Architecture and operation," *Fusion Engineering and Design*, 83(2-3), 304-311, 2008, 10.1016/j.fusengdes.2008.01.015.

[3] W. Riley, D. Howe, "Handbook of Frequency Stability Analysis", *Special Publication (NIST SP) National Institute of Standards and Technology, Gaithersburg, MD, [online]*, 2008, https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=50505 (Accessed June 12, 2026)

[4] G. Raupp, H. Richter, C. Aubanel, H. Bruhns, R. Huber, G. Schramm, et al., "The timing system for the ASDEX Upgrade experiment control," *IEEE Transactions on Nuclear Science*, 39(2), 198-204, 1992, 10.1109/23.277482.

[5] H. Richter, G. Neu, G. Raupp, D. Zasche, "System integration of the ASDEX-Upgrade timing system," *IEEE Transactions on Nuclear Science*, 41(1), 178-180, 1994, 10.1109/23.281484.

[6] G. Raupp, K. Behler, G. Neu, R. Merkel, W. Treutterer, D. Zasche, et al., "Integrating discharge control and data acquisition at ASDEX Upgrade," *Fusion Engineering and Design*, 60(3), 353-359, 2002, 10.1016/S0920-3796(02)00032-7

[7] G. Raupp, R. Cole, K. Behler, M. Fitzek, P. Heimann, A. Lohs, K. Lüddecke, G. Neu, J. Schacht, W. Treutterer, D. Zasche, Th. Zehetbauer, M. Zilker, et al., "A "Universal Time" system for ASDEX upgrade," *Fusion Engineering and Design*, Volumes 66–68. 947-951, 2003, 10.1016/S0920-3796(03)00381-8.