Measurement 257 (2026) 118528

Contents lists available at ScienceDirect

Measurement

journal homepage: www.elsevier.com/locate/measurement

# Hardware-agnostic framework for general-purpose data acquisition systems

Grzegorz Korcyl [a,b,*], Maciej Bakalarek [c], Paweł Moskal [b,c]

[a] Institute of Applied Computer Science, Jagiellonian University, Prof. St. Łojasiewicza 11, PL30348, Cracow, Poland
[b] Center for Theranostics, Jagiellonian University, M. Kopernika 40, PL31034, Cracow, Poland
[c] Marian Smoluchowski Institute of Physics, Jagiellonian University, Prof. St. Łojasiewicza 11, PL30348, Cracow, Poland

ARTICLE INFO

*Keywords:*
Data acquisition systems
Time synchronization
Synchronous networks
Field-Programmable Gate Arrays

ABSTRACT

Modern data acquisition systems in physics experiments of various scales rely heavily on precise time synchronization of electronic components used for digitizing analog signals generated by detectors. Scientific collaborations often develop custom, in-house solutions designed for particular use cases that form ecosystems of compatible hardware platforms, protocols, and standards.

In this work, we present and evaluate a set of gateware and software components that enable the implementation of complete data acquisition systems with synchronization, data transport, and control functionalities over multi-gigabit transceivers. The components can be organized to form a logical system structure, including the master, data concentrators, and endpoint modules, on virtually any hardware platform equipped with a Field-Programmable Gate Array with an exposed transceiver. The system allows for a rapid deployment of complete measurement systems with facilitated integration of new elements on both commercially available and custom-designed platforms.

## 1. Introduction

Data Acquisition Systems (DAQ) for modern physics experiments comprise digitizing nodes distributed and installed over the detector system. These nodes must simultaneously register the detector response to a physical event, each taking a snapshot of a particular detector region. These snapshots must be precisely aligned in time to properly reconstruct the complete detector response from events in sub-nanosecond timescale. Hence, the nodes must share a common reference timing signal.

Typically, dedicated circuitry would have been set up to provide a common reference signal to the nodes. This signal would often be the output of an advanced hardware processor analyzing certain detector channels in real time and thus providing preliminary event selection, called a trigger. Nowadays, the signatures of the physics phenomena to be registered have become much more complex, rendering hardware trigger systems obsolete and shifting the data selection to the online software computing facilities [1]. The trigger-less or continuous readout systems reuse the data links to distribute the common clock and synchronization signals by employing synchronous protocols and methods. The bitstreams transmitted over data links between the nodes of the system can be used to extract the sender's clock by the receiver [2]. By designing a hierarchical system architecture with a single master node, one can distribute the master clock to the endpoints using this clock recovery feature of the network transceivers.

The quality of this process will reflect on the performance of the entire measurement system as any degradation of the synchronization on any endpoint will result in an inability to reconstruct physical events properly. The systems must ensure the recovered clocks on the endpoints remain stable, fixed in phase between each other, do not drift over long periods or due to external conditions, and are repetitive after system power cycles.

Scientific collaborations develop dedicated hardware solutions, where the clock handling is carefully designed to achieve the best synchronization performance. On top of this, custom network protocols perform hand-shaking between the nodes to establish and maintain synchronous connections. However optimized, this approach significantly raises the costs and complexity of the systems, simultaneously narrowing their flexibility and adaptability to various detector systems.

### 1.1. Related works

Numerous hardware platforms, protocols, and implementations have been developed by scientific collaborations. However, the system scheme remains common. A single, master electronic module is the source of the common clock and synchronization signal, which is forwarded and distributed by the concentrator cards to the endpoints that digitize the detector data. The synchronization is propagated downstream through synchronous protocols, usually over fiber-optic links. In

* Corresponding author.
*E-mail address:* grzegorz.korcyl@uj.edu.pl (G. Korcyl).

https://doi.org/10.1016/j.measurement.2025.118528
Received 2 April 2025; Received in revised form 25 June 2025; Accepted 23 July 2025
Available online 4 August 2025

the LHCb experiment, a custom FPGA platform called PCIe40 [3] serves multiple functions in the system essentially maintaining a synchronous connection to the endpoints through the GigaBit Transceiver (GBT) links [4] - a standard developed by CERN, with the ability to receive master synchronization from the Large Hadron Collider via Timing distribution and Fast Control (TFC) system [5]. The system is tightly coupled to the LHC bunch clock frequency 40.08 MHz, which is distributed to the endpoints as a dedicated data link clock through a series of components with guaranteed fixed latency to ensure the deterministic phase alignment of the endpoints. Similar concepts are being implemented in the CBM experiment with BNL-712 [6], a different FPGA card with multiple optical link bundles for synchronous connections to the endpoints over GBT protocol and a TFC receiver [7].

Although the hardware platforms are different in those two cases, they must share common solutions required by the chosen protocols, thus limiting their applications to other systems. Both cards are clients of the TFC protocol and require a dedicated receiver and protocol implementation (the BNL-712 has a dedicated Timing Mezzanine Card attached). The synchronous connection to the endpoints from the master clock signal recovered by the TFC is implemented as a common multi-gigabit transceiver reference clock that requires a hardware jitter cleaner, fan-out device, and precise net routing (Fig. 1B) on the Printed Circuit Board (PCB) design.

This clocking scheme is common in recent synchronous systems. However, it has drawbacks, namely, it leads to the increased complexity of the hardware design, narrows the clocking flexibility of the manufactured components, and imposes a significant barrier in the integration of hardware modules not foreseen to operate within the system a priori.

When looking at the protocols for time synchronization IEEE 1588 known as Precision Time Protocol (PTP) comes as a standard in distributed networks [8]. The protocol defines a sequence of messages the nodes exchange, carrying timestamps (hardware timestamping is required for sub-nanosecond precision). Based on those timestamps, the nodes can determine the time offset between them and align their clocks if the hardware is capable (e.g., are equipped with Voltage Controller Oscillators (VCO)). The protocol relies on Ethernet networks and the transport OSI layer utilizing the User Datagram Protocol (UDP) and requires frequent handshakes between the nodes to adjust the clocks. Therefore, both nodes must have active Internet Protocol (IP) stacks implemented to build and exchange packets in an Ethernet environment. For small FPGA devices, often installed in front-end electronics, the complexity and high resource consumption of network stacks make it challenging to implement PTP clients and impose an unnecessary overhead.

While IEEE 1588 provides the means to determine timing offsets between the nodes and adjust their clocking periodically, it does not hold the synchronization. It allows the clocks to drift over time. This is because the protocol operates on a regular Ethernet infrastructure, which does not guarantee the synchronicity of the clock domains between the input and output ports of the commercial equipment.

An extension of the standard called White Rabbit [9] has been developed as a response to this issue, combining time offset determination of the IEEE 1588 and Synchronous Ethernet (SyncE) hardware infrastructure components [10]. This assures clock domain transfer between the input and output ports of the equipment and provides a fixed phase relationship between the clocks on distributed nodes. Despite being widely adopted by the community and becoming a standard in large-scale research facilities, it requires dedicated and expensive hardware for the infrastructure and specific clocking features on the end nodes, narrowing the range of electronics compatible with the standard and making it not feasible to deploy for small or prototype setups.

On the other end of the spectrum are commercial complete DAQ solutions such as NI PXI system [11–13] or digitizers from CAEN [14,15]. They provide excellent signal processing functionalities in convenient form factors with scaling capabilities through precise synchronization features. However, the overall channel density remains low for large-scale detector systems, and the cost per channel is often unreachable. Despite the possibilities of implementing and including custom functionalities in the readout chain, the systems remain closed and unavailable for fine-tuning to particular applications.

### 1.2. Outline

This paper presents a complete DAQ solution based solely on gateware and software components (Fig. 1C). It allows the implementation of synchronous systems on virtually any hardware platform sporting a Field Programmable Gate Array (FPGA) device with a Multi-Gigabit Transceiver (MGT), which is recently a standard in this type of electronics. The key points of this contribution are as follows:

- Decoupling from dedicated hardware platforms and leveraging the functionality to gateware and software makes the systems less complex, less expensive, and easier to set up
- No specific hardware requirements make it suitable to integrate with a variety of both custom and Commercial-Off-The-Shelf (COTS) hardware platforms
- The gateware components are compliant with recent standards, making it easy to integrate, prototype, and develop complete solutions with modern FPGA development techniques and tools
- The synchronization process is adaptable to any master clock frequency and the resulting data link rate (above the minimum rate and supported by all nodes)
- A single physical data link is shared between three logical functions providing synchronization, measurement data transport, and control and monitoring messages exchange
- The system has been implemented and evaluated on multiple platforms, including those equipped with Time-to-Digital and Analog-to-Digital Converters (TDC and ADC)
- The synchronization performance has been measured as the offset jitter of the recovered clocks at the level of 10 ps, making it suitable for measurements with high-end TDC and ADC devices

The system logic architecture is introduced in Section 2, and the essential synchronization mechanism is described in Section 3, including performance evaluation in Section 3.3. Afterward, the mechanisms for transporting the measurement data and the control and monitoring system are discussed in Sections 4 and 5, respectively. By the end, the hardware platforms on which the system has been implemented are demonstrated in Section 6, and the overall summary is included in Section 7.

## 2. System architecture

### 2.1. Logical components

Synchronous systems require a common clock source propagated to the endpoints that digitize signals from the detectors. Therefore, the system's hierarchical, tree-like architecture is inferred, in which a component location defines its logic function. Following the convention from the AXI (Advanced eXtensible Interface) standard (see Section 2.4), we recognize two main logic functions: a master that provides synchronization and a slave that receives it. System functions can be assigned to the physical components that instantiate a combination of these two functions.

A system master is at the root of the system, the common clock's source. It can offer several master interfaces and, hence, become a data concentrator. A regular data concentrator has an instance of the slave receiving synchronization from the master and forwarding it to the master interfaces. Finally, an endpoint has a slave instance, no master interfaces, and is supposed to be the measurement data source. Since no specific hardware components are required, these functions can be implemented as a combination on a single FPGA, multiple FPGAs on a single board, or as an entirely distributed system (multi-gigabit transceiver connectivity is required in the latter two) (Fig. 2).

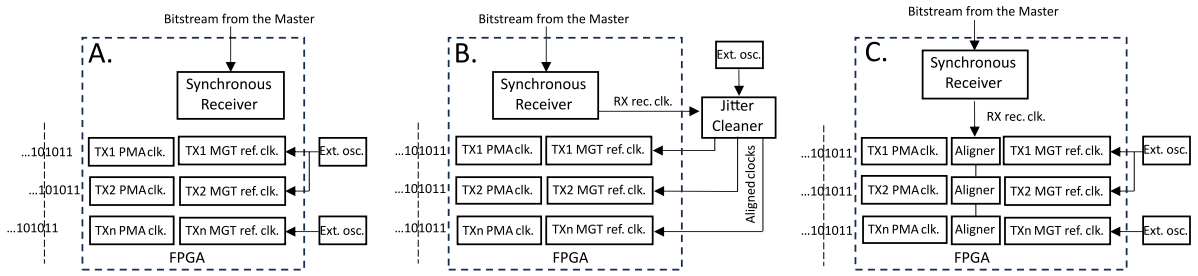


**Fig. 1.** Clocking distribution schemes for transmit channels (TX) in Multi-Gigabit Transceivers (MGT). External oscillators source the input of MGT reference clocks, based on which the high-speed serial clocks (in the Physical Medium Attachment layer (PMA)) are produced in transmission channels. By default, when sourced from the same physical oscillator, the channels are not aligned with each other (A) and, therefore, cannot be used to synchronize the receiving nodes. The master clock domain cannot be forwarded to the transmit channels because of the net routing in the hardware design. A common solution to this problem is implementing a jitter cleaner device (B), which can take the recovered clock from the master (RX) and source the MGT reference clock inputs. However, such a scheme must be implemented during the hardware design phase, and once produced, cannot be altered. In this contribution, we present a method to align the clocks internally, inside the FPGA device (C). Despite the clocking scheme (A or B), until the expected bitrate can be achieved, the hardware module can be integrated with the AuroraSync system.

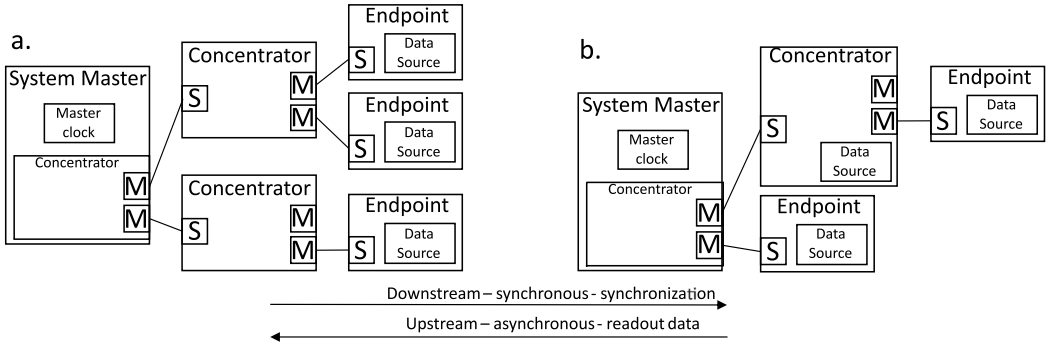


**Fig. 2.** Two of the possible configurations of the system. The master (M) and slave (S) types of interfaces allow the connection of endpoints or data concentrators at any level of the tree-like architecture of the setups. The synchronous channel is provided in the direction from the master towards the endpoints (downstream) to provide synchronization; the return channel (upstream) remains asynchronous and transports measurement data.

### 2.2. Readout process

The continuous readout mode of the system is realized by issuing SYNC commands by the master at a fixed frequency. The endpoints continuously buffer gathered data and transfer the buffer content out upon reception of the SYNC command, making room for new digitized signals. The data can be transferred back to the data concentrators or directly outside the system, depending on the hardware platform capabilities (e.g., a second transceiver). In such a way, the entire measurement period is segmented into equally spaced timeslots.

### 2.3. Data aggregation

The upstream channel, from the endpoints in the direction of the master, is asynchronous to preserve resources and simplify the design. It allows the aggregation of the data from multiple endpoints in a single place (data concentrator). The assembled data fragments can be transformed into network packets and transmitted out. That gives flexibility in the organization of the data links going out of the system to storage or online software processing. Since multiple endpoints can cover a logical segment of the detector system, it also enables the implementation of higher-level pre-processing algorithms on the data-concentrating FPGA.

### 2.4. Standards and protocols

AMD/Xilinx ecosystem has been selected as being the most accessible and the most commonly used FPGA manufacturer in the DAQ systems in physics experiments. This choice entailed the usage of Aurora [16] and AXI as native, point-to-point data transport protocols between the devices and standardized interfaces for component-component interconnection, respectively. A wide selection of ready-to-use components (IP cores) allows for fast system composition and prototyping. Custom logic components can be encapsulated with AXI-compatible wrappers and incorporated into the system. AXI transactions (stream and memory-mapped) can be serialized and transported over Aurora links using the AXI Chip-2-Chip IP component [17], which enables the distribution of the logic functions across remote hardware platforms.

The AXI standard is based on the master–slave connectivity concept. The master board in the system provides the main AXI master (e.g., HPMM interface of the Zynq SoC). Multiple AXI slaves can then be connected to that master via AXI Interconnect/SmartConnect IP. The interconnect multiplexes between single slave and multiple exposed master interfaces. A complex system composed of AXI components can be constructed in such a cascaded way. At any point between a master and a slave interface, the AXI Chip-2-Chip with an Aurora transceiver combo can be inserted allowing protocol transport over a network medium (optical fibers, copper cables, onboard PCB nets, etc.) leaving logically the system intact (Fig. 3) and preserving master clock domain.

All major functional components such as synchronous transceiver with Aurora protocol, data sources, status, and control blocks are enclosed as stand-alone IP core objects with AXI interfaces, which can be imported to Vivado design suite from AMD/Xilinx and extend its built-in IP catalog. Based on these few points, one can relatively easily construct an entire system based on components accessible from that catalog, all with a drag-and-drop method in the graphic mode such as Block Design in Vivado [18]. This is a significant aspect of this project, as it lowers the entry threshold for the developers in collaborations with limited personnel experienced in FPGA technology and low-level VHDL or Verilog programming.

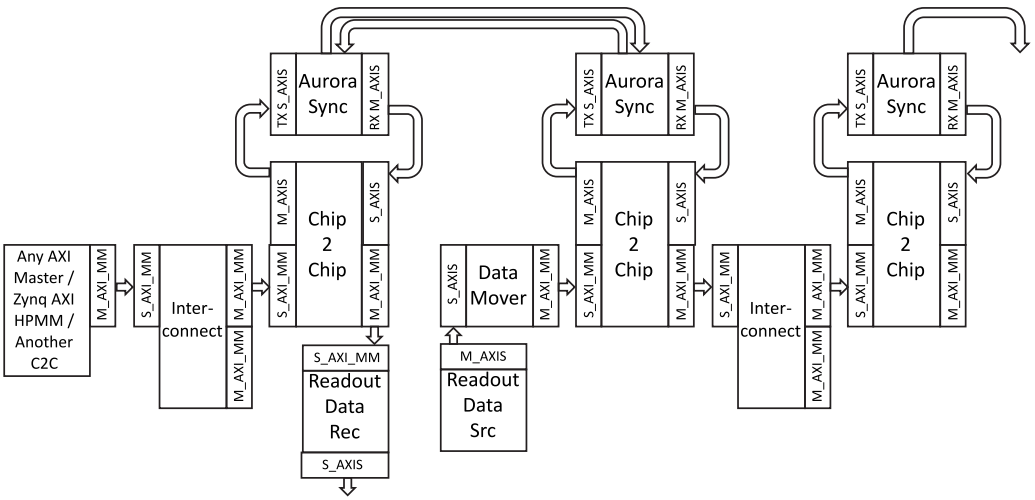


**Fig. 3.** Internal configuration of the logic components required to establish a link. The synchronization source is the single M_AXI_MM interface on the master distributed through the Interconnect and further Chip2Chip components controlling the Aurora Sync transceivers. The system can be scaled up by adding the next levels of Interconnects, opening the possibility for more Aurora Sync channels (if enough transceivers are exposed).

## 3. System synchronization

### 3.1. Implementation

In synchronous systems, distributed endpoints require a common synchronization signal to align the data from the complete detector system in post-processing. To reduce the number of connections and elements, three logical functions have been incorporated into a single Aurora-based data link: measurement data transport (upstream), control and monitoring messages (downstream and upstream), and synchronization pulse delivery (downstream).

The latter requires that all data links are synchronous, meaning the phase relationship between the recovered and the master clock must remain fixed. By default, all MGT channels operate on their own generated local clock domains, based on, supplied by hardware design reference clock (usually from an external oscillator). Essentially, it means that each network partner has its individual clock domain.

The base principle of the clock distribution mechanism via serial data links is that the master clock is encoded within the transmitted bitstream (Fig. 4). The receiver Clock and Data Recovery (CDR) block extracts the original serial clock, which must align to the transitions between consecutive bits to decode their logic state properly. Usually, this extracted clock is used to de-serialize the stream and produce a data word. Transceiver components such as the comma (special, control data words) detector, 8b/10b decoder (encoding used to guarantee constant logic state changes), and elastic buffer automatically manipulate the decoded data word and transition to local clock domain. Hence, these components must be deactivated, and their functionality handled manually (Fig. 5). After analyzing the Aurora protocol, a dedicated RX synchronizer component has been developed that uses the RX slide feature of the transceiver to lock the RX recovered clock to a unique 40b comma sequence. Assuming that all TX links transmit the comma sequence synchronously, all RX partners would always lock in a fixed relationship.

To ensure that all outgoing TX links transfer bits synchronously, their local serial clocks must be synchronized to a common reference. Usually, it is achieved by doing hardware implementation: an external jitter cleaner chip (e.g. LMK04610 [19]) is installed, which takes one input (the master clock or the recovered clock from the upstream link) and drives the MGT reference clock inputs of all transceivers (Fig. 1B). In such a way, all transceivers operate on the same physical clock signal and do not require additional synchronization; however, as mentioned before, it requires specific hardware design and components.

All families of AMD/Xilinx FPGAs feature transceivers with TX Phase Interpolator (PI). It allows for the manual adjustment of the phase of the TX clock and, hence, the alignment of multiple TX channels independently from their reference clocks, assuming the links operate at the same bit rate. With this feature, the transceivers can operate on local, not synchronized MGT reference clocks, and produce synchronized outgoing streams of bits, therefore not requiring off-the-chip clock synchronization infrastructure (Fig. 1C).

To adjust the phase of the TX clock using the Phase Interpolator, a technique described in [20] has been employed. A PICXO component consisting of a high-performance, gateware-based Digital Phase-Locked-Loop (DPLL) is instantiated for each TX channel and can be directly connected to the TX Phase Interpolator control ports (Fig. 5). By providing a common, external reference clock to all instances of PICXO, they will individually adjust and synchronize the TX channels. However, the DPLL consumes considerable resources and significantly extends the compilation time due to the complex clock signal routing and strict timing constraints. An alternative technique, proposed by E. Mendes in [21], could be employed and will be evaluated in the future, where the elastic buffer fill level on the TX path is used to determine the phase relationship between two clock domains and provides feedback to the adjustment of the TX Phase Interpolator.

Additionally, to improve the quality of the recovered clock signal, an instance of PICXO, which has clock cleaning capabilities, has also been instantiated on the receiver channel. Hence, the clock provided to the TX channels as a reference is more stable, neglecting the negative effects of clock forwarding over consecutive layers of electronics. This effect is visible on the phase noise curves in Section 3.3.4.

### 3.2. System extensibility

All the above-mentioned mechanisms have been implemented for all major transceiver families (GTP, GTX, GTH, GTY) as enhancements to the regular transceiver IP. Therefore, it can be easily integrated with the Aurora IPs and incorporated in designs just as regular components. The Aurora protocol is a soft gateware that does not rely on any hardware features of a particular FPGA family and, hence, can be migrated to other devices.

Specific to AMD/Xilinx devices are the TX Phase Interpolator and the PICXO component, which align the outgoing bitstreams that are essential for the concentrator modules in the system. Hardware boards with multiple transceivers but FPGAs from other manufacturers can be incorporated into the system as long as they employ the classic clock forwarding technique with the external jitter cleaner (Fig. 1).

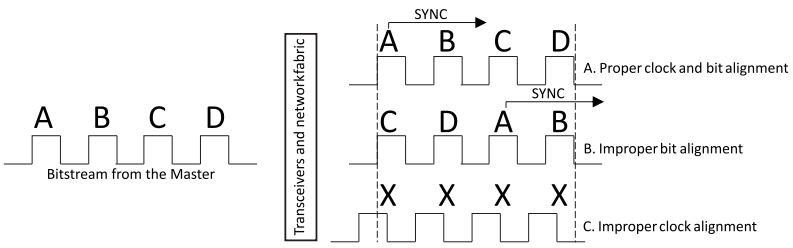


**Fig. 4.** Pictorial representation of the transceiver clocking alignment. Properly operating CDR allows it to align into the incoming bitstream and extract the original data word and, therefore, the master clock and, consequently, the SYNC bit (A). Any deviations of that case result in either spurious decoding of the SYNC bit (B) or the incapability to decode any valid data from the channel (C).

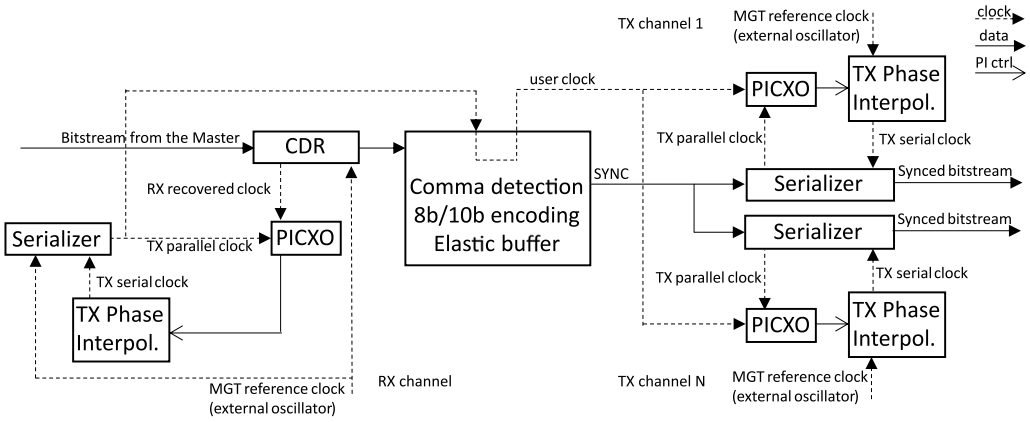


**Fig. 5.** Clocking network on the example concentrator node. The original clock phase is extracted from the bitstream from the Master node by the CDR component of the RX channel. A loopback through a PICXO component is used to reduce the noise on the raw, recovered clock. The subsequent components would typically affect the timing, hence losing the capability to extract the SYNC pulse properly. The developed block allows the preservation of the RX recovered clock domain throughout the processing chain and an automatic reaction to the unique comma code, allowing repetitive bit locking (as in Fig. 4A and B). The extracted clock can be further used to synchronize multiple TX channels. Each TX channel operates autonomously on its reference clock unless they are aligned to a common clock by the Phase Interpolators. The PICXO components monitor the relation between the TX parallel clocks and the common clock, providing feedback to the PIs.

The clock recovery feature of the RX channel from the input bitstream is a common feature in all major FPGA families from many manufacturers. Therefore, most hardware modules can play the role of endpoints in the system as they only need to extract the SYNC pulse without further forwarding. The recovered clock can be further used on the endpoint to produce other clock frequencies that remain fixed in phase by the internal clocking features of the FPGAs. This is important in the case the digitization devices (e.g., Analog-to-Digitial Converters) have operational frequencies other than the one derived from the established communication link (e.g., the SADC platform 6.3.2).

Moreover, the Phase Interpolator operates on the TX PMA clock, the frequency of which is derived from the MGT reference clock and adjusted to the selected link bitrate. This means that the frequency of local crystal oscillators installed on nodes is irrelevant as long as the desired bitrate can be achieved. That way, virtually every modern FPGA-based board with a transceiver exposed can be integrated into the system.

The project's versatility in clocking scenarios, employment of well-established protocols and standards, and a long FPGA product lifecycle ensure that the solution will remain functional for a typical system lifetime of 10–20 years [22]. The fact that the solution is hardware agnostic — not bound to any particular hardware platform but consisting of gateware components, possibly to be migrated to other platforms — makes it even more future-proof and independent from the market availability.

### 3.3. Measurements

When evaluating the stability of the synchronization, the key figures of merit are the period and the long-term jitter of individual clocks, then the relation between the clocks in the system expressed as the phase offset jitter between pairs, the reproducibility of these characteristics over a series of complete system power-cycles and influence of environmental factors such as temperature. These values allow the assessment of whether the system will provide a platform that fulfills particular measurement accuracy goals, not considering the actual signal digitization stage, which is subject to specific endpoint hardware.

#### 3.3.1. Test system

A minimalistic setup composed of three ZCU102 boards has been arranged in a cascade as a single master board, one concentrator, and one endpoint (similar to Fig. 11). A 5 Gbps link has been established, and the 125 MHz (8 ns period) clock signals have been exposed on general purpose pins to be probed with a Keysight DSO-404 A oscilloscope (sampling rate 20 GSps, analog probe bandwidth 500 MHz). All jitter measurements have been performed according to the JEDEC Standard 65B [23] and following the methodology described in [24], taking an average of 25 sets of 10 000 measurements of a particular value. It is essential to note that the measured clock signals are routed outside the FPGA device to be accessible on a physical measurement pin, which already affects the signal quality. Therefore, the system's actual performance is expected to be better than the one presented below.

#### 3.3.2. Baseline

At first, we attempt to determine how the measurement system's intrinsic jitter and random noise affect the results. Two probes have been put on the same test point driven by a clock signal, and the phase offset histogram is recorded to extract the standard deviation value. As the clock signal is naturally periodic, the averaging mode

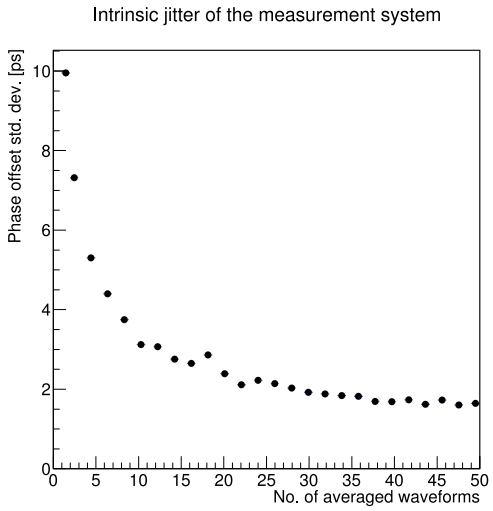


**Fig. 6.** Standard deviation of the phase offset histogram produced as a time difference recorded by two probes measuring the same signal, on the same test point as a function of the number of waveforms taken into averaging acquisition.

of the acquisition can be applied to reduce the impact of the random noise [25]. In Fig. 6, the standard deviation of the phase offset depending on the number of waveforms to be averaged is presented. It can be concluded that the normal acquisition mode is significantly affected by the noise and the jitter of the measurement system, starting from 10 ps RMS for a single waveform and leveling at 2 ps RMS for 30 averaged waveforms onwards. Based on this reference measurement, we can provide a baseline value for the jitter measurements and configure the acquisition mode to average 32 waveforms.

### 3.3.3. Peak-to-peak jitter

To maintain synchronization of the endpoints, the clocks of the TX channels on the data concentrator are continuously adjusted to a reference clock signal by the PICXO component. Therefore, its stability can be affected. Similarly, the recovered clock is the local oscillator clock signal on the endpoint, adjusted to the received stream of bits by the CDR block of the RX channel. In Fig. 7, the stability of both the master (top row) and the recovered clocks (middle and bottom rows) is presented as the peak-to-peak period jitter (first column) and the long-term jitter (second column).

The peak-to-peak values are below 10 ps, satisfying the required 21.75 ps for applications such as 10 Gigabit Ethernet, where Bit-Error Rate (BER) is expected to be below $10^{-12}$ [26], to ensure proper operation for bitrates below 10 Gbps. For reference, the SI570 programmable oscillator commonly used on the AMD/Xilinx evaluation boards to drive the multi-gigabit transceivers has 14 ps peak-to-peak period jitter, as stated in the datasheet [27]. Interestingly, the recovered clock is significantly more stable (Fig. 7C, E) due to the CDR and PICXO circuitry that works as an additional filter.

Long-term jitter measurements were collected over 10 h and show approximately 10 times wider distribution of the period jitter measured as the time interval between 10 000 clock cycles [24]. In this case, the results are worse as the measurement is more prone to the noise (averaging mode is not applicable in this case) and could also be affected by environmental factors such as daily temperature changes.

### 3.3.4. Phase noise

Phase noise measurements have been performed by inspecting frequencies that affect clock signals at various stages of the system. The curves presented in Fig. 8 have been collected with the DSO-404 A oscilloscope and the EZJIT analysis package [28] over the frequency offset range between 1 kHz and 100 MHz and the default spur sensitivity (value 3.0). The RMS phase jitter was calculated by integrating the phase noise curves within the integration bandwidth between 1.875 MHz and 20 MHz (vertical, dashed lines), which is typical for similar protocols and bit rates, such as 10 GB Ethernet or XAUI [29].

First, the raw master 125 MHz clock (red curve) is measured to provide a reference for others and to visualize the intrinsic noise of the oscillator, as well as other components that can cause crosstalk, the power supply, and the measurement system itself (A). Secondly, the green curve represents the clock on the master board, processed by the PICXO component, which is effectively used for transmission.

The blue curve shows the recovered clock by the RX channel on the slave board. Significant distortions are visible on almost the entire frequency offset range. Since this clock is further used to synchronize and align the transmission channels, it must maintain high quality. Hence, the additional PICXO component is instantiated on the RX channel as described in Section 3.1 and presented on the left side of Fig. 5. The cleaned clock, presented as the magenta curve, shows a significant improvement, almost recovering the noise levels of the master clock.

Theoretical phase noise levels of the Si570 programmable oscillator are displayed as the black line for reference, and the RMS phase jitter value calculated over the selected range is equal to 262 fs [27]. The value measured for the raw master clock (green curve), originating from such an oscillator, is 2.2 ps, which indicates the overhead of the measurement system and the baseline. Similarly, however, including the spurs, the RMS jitter was calculated for both the master and slave clocks, and ranges from 2 ps to 4.26 ps for the worst-case scenario of the unprocessed, recovered clock. This value is above the theoretical limit for the 5 Gbps link, and BER below $10^{-12}$, which is 580 fs [30]. However, these clocks are used to align the transmission channels and are not the hardware MGT reference clocks used to generate serial clocks, for which the limit is defined. An important observation is that the processed clocks remain at a noise level similar to the raw master clock, which, considering the measurement system noise overhead, is satisfactory. Not to be neglected is the deterministic jitter that appears on the curves as the spurs on clocks processed by the PICXO components. Those appearing close to the half-carrier frequency (B) can have a significant impact on performance and should be investigated to optimize the system further [31].

### 3.3.5. Phase offsets

The relation of the clocks is presented in Fig. 9 with a single run (A, C, E) and including complete system reboots in each of the 25 measurement cycles (B, D, F), respectively. The latter is essential, as it presents the deterministic time offsets between the components of the system that are stable over multiple runs. For a given set of FPGA configurations and interconnections, such as medium and link lengths, the offsets remain fixed and can be compensated with a single set of calibration parameters. Slightly worse results for the reboot cycles are due to the prolonged period of the measurement and the environmental effect described above.

Similarly to the period jitter (Section 3.3.3), one can notice the filtering effect on the recovered clocks and a visibly more stable relation of the slave nodes versus each other (E). This allows us to draw conclusions about the scalability of the system, where the introduction of additional concentrator board layers (more hops) does not degrade the synchronization performance.

Considering the system's hardware components (described in Section 6), the measured stability of the synchronization is satisfactory for typical cases. The SADC platform for the PANDA detector system is designed to sample analog signals with an 80 MHz clock derived from the recovered clock. The 5 ps jitter between the endpoints is at the level of 0.4 permils, covering the sampling period of 12.5 ns. It is also aligned with the binning of the TDC on the MTAB platform for the J-PET tomography scanner, where the single element of the delay chain introduces approximately 25 ps latency.

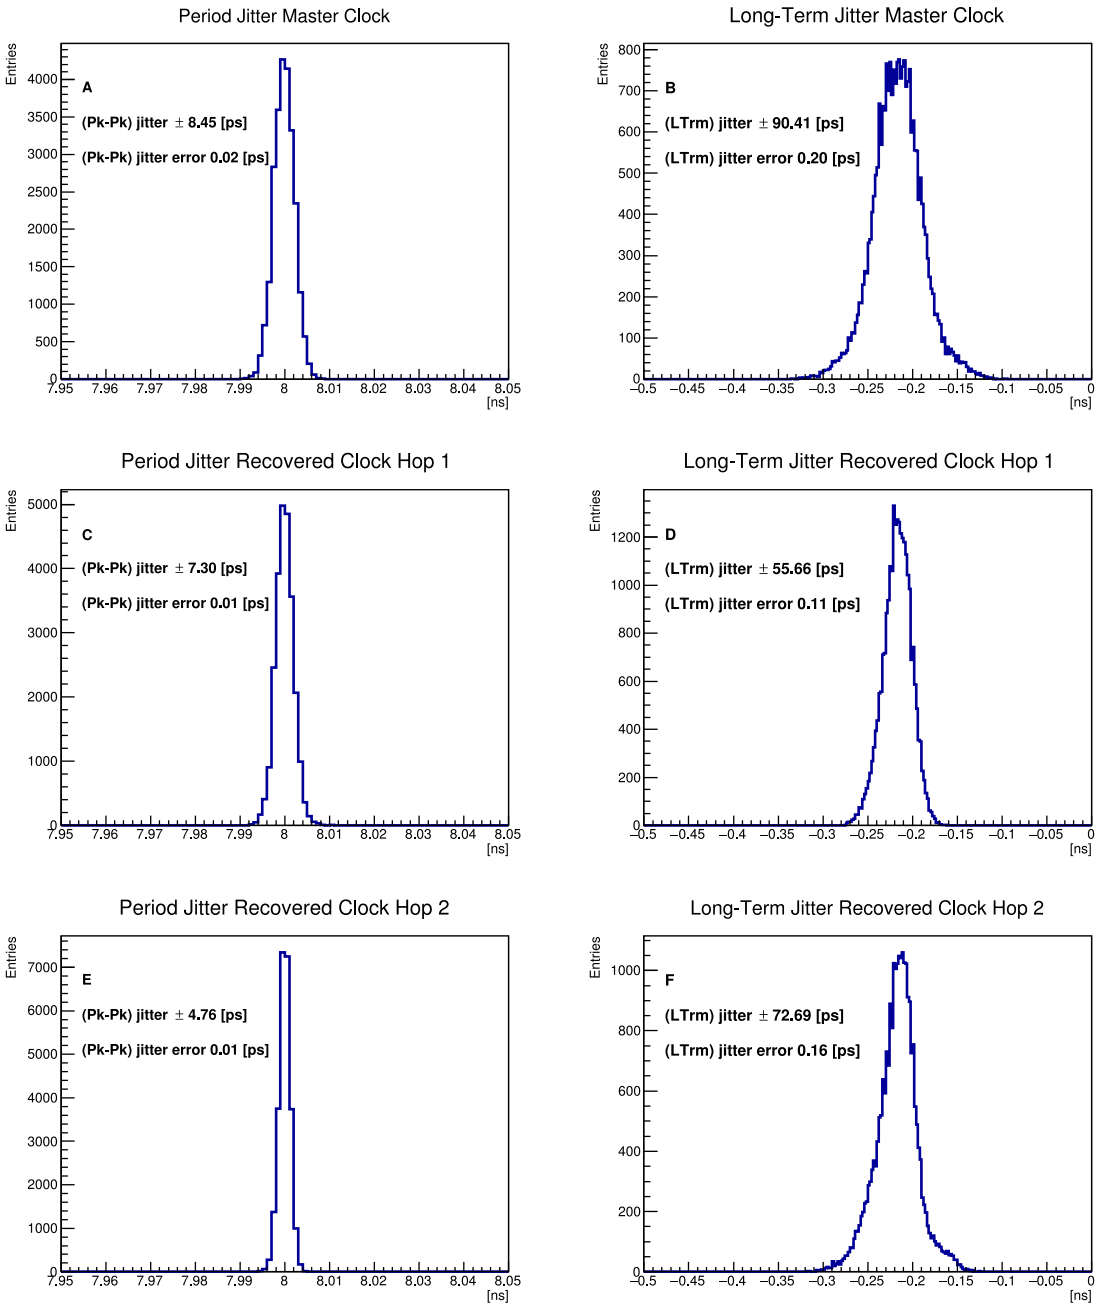


**Fig. 7.** Characterization of the master and the recovered clocks. In the left column (panels A, C, E), the peak-to-peak jitter of the master and the recovered clocks are presented. The second column (panels B, D, and F) analogously shows the long-term jitter presented as a difference from the ideal 80 μs.

### 3.3.6. Temperature influence

The performance of the electronics can be affected by temperature changes, especially when considering sub-nanosecond scale. A closer look at the system behavior under a significant temperature change (about 15 °C) is presented in Fig. 10. The phase offsets between the test system nodes' clocks were recorded over 20 h, with two periods when the cooling was manually switched off (time samples 5–30 and 100–125). One can see that the system remains stable with constant temperature, however the spikes are visible at the moments of temperature change. This indicates that the phase offset values change with the temperature, but the PICXO and CDR mechanisms maintain the synchronization of the nodes. The latencies of internal circuits affected by the temperature are not compensated automatically, and in the test case, the offsets can shift by 80 ps due to 15 °C change.

To avoid this negative effect, a proper temperature conditioning system should be considered, despite the measurements usually being performed in well-controlled environment areas. Otherwise, the FPGA temperatures should be recorded along with the measurement data for an offline correction.

## 3.4. System validation

The most spectacular application of the system is the Jagiellonian-PET (J-PET) tomography scanner. Consisting of 24 detector modules, equipped with 48 digitizing endpoint boards, aggregated by 4 concentrator boards (VCU108), and controlled with a single master module (ZCU102), it is being used to carry out breakthrough research in medical imaging and fundamental physics [32–36]. Particular hardware platforms employed in this system are described in more detail in Section 6.

Continuous readout and high-resolution signal digitization with high-precision synchronization allowed the recording of multiphoton

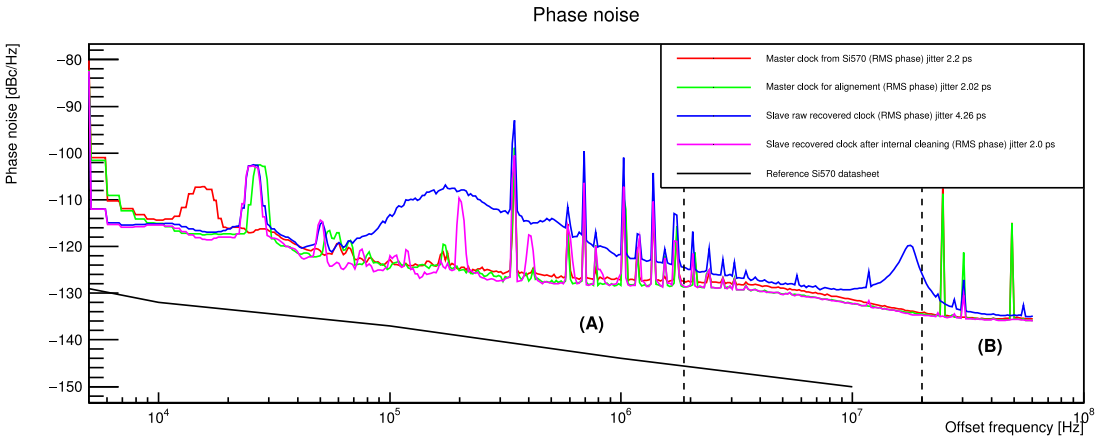


**Fig. 8.** Phase noise curves collected for raw master clock (red), processed master clock (green), recovered clock (blue), and recovered processed clock (magenta). The black line shows reference values from the Si570 datasheet. Integrated RMS jitter is calculated for spectrum range between 1.85 MHz and 20 MHz (vertical, dashed lines). Two kinds of noise are identified: in the region (A) are spurs common to all measured clocks, therefore, these effects originate from the original clock signal, the measurement system, and the environment; in the region (B) are spurs that appear for the signals processed by CDR and PICXO components.

coincidences used to determine positronium lifetimes in the human brain for the first time in vivo. That, in turn, can improve diagnostics of the tissue pathology.

Multiple measurement campaigns have been carried out, including those with hospital patients and phantoms in laboratories. Single measurements ranged from a dozen minutes to several months of continuous data taking. During this time, the system experienced typical issues, such as endpoints losing connectivity due to damaged optical fibers or overheating of modules. None of these have significantly affected the measurement process, and a power cycle would usually bring the system back to operation.

## 4. Measurement data transport

The endpoints continuously register data and fill up the buffers. Upon reception of the synchronization pulse SYNC, the buffer content is stamped with a unique identifier of the timeslot generated by the system master and transmitted from the endpoint. A default channel for the measurement data is the link upstream in the hierarchy, back to the master interface of the data concentrator. If the endpoint is capable (e.g., features a second network interface), the data can also be directly packaged and transmitted out of the system at that level.

The data concentrators monitor activated links to the endpoints and collect the returned packages with measurement data. The packages are stored in local buffers and await a round-robin state machine to fetch the content and compose an aggregated data package that can be transferred out of the system or, again, forwarded one level upstream in the system hierarchy. The case where one endpoint would not deliver its data package in response to the SYNC command is critical. It causes the readout procedure to stop as it typically means the endpoint lost connection to the concentrator and requires an intervention. Such events are immediately visible on the operator control panel and require a system reboot.

Such multi-level architecture allows the control of the data flow. It opens possibilities for implementing online pre-processing mechanisms that require data from a particular detector system segment instead of a single endpoint. Since the data is stored in buffers and transferred between the logic components as AXI Stream, it makes it natural for algorithmics to be implemented with High-Level Synthesis tools, significantly facilitating contributions from non-FPGA experts.

Moreover, the flexibility in arranging the layers of data concentrators and gateways to the network directing to storage allows for managing load balancing in the system. The data funnel on the data concentrators, collecting data from the endpoints one by one, is a critical local bottleneck. Considering the continuous readout of the system at a fixed frequency, this process must be completed before the next SYNC pulse is delivered. Therefore, the data flow in the system must be carefully planned, analyzed, and segmented with additional data concentrators when required. Additionally, the mentioned buffers are equipped with configurable data limit levels, allowing dynamic trimming of the size of data packages produced on various system levels.

The data gateways are implemented as Gigabit or 10 Gigabit Ethernet interfaces with UDP as the data transport protocol by default. They are required to transmit data in the TX direction only and usually operate in controlled, local network environments where packet loss is not an issue as long as enough bandwidth is secured. Network stack with basic protocols such as ARP, DHCP, ICMP for visibility in the network, and UDP are implemented as FPGA logic IP cores and therefore can be easily instantiated on capable endpoints and data concentrators [37].

## 5. Control and monitoring

The starting point for the control and monitoring system is a single AXI-Lite master interface on the system master, which is the root for all connected components. If a Zynq MPSoC is used, it is usually the HPMM interface that stands between the Processing System (PS) and the Programmable Logic (PL) parts of the device. Therefore, this interface is a bridge between the software and gateware domains.

### 5.1. Gateware

As mentioned in Section 3.1, the AXI Interconnect component allows the connection of multiple slaves to a single master interface, and the AXI Chip2Chip component allows AXI transactions to be transported over a multi-gigabit transceiver. Due to this cascaded architecture, each remote AXI-compatible component is reachable from the system master interface as if it were connected locally.

Each logic component can be encapsulated in an AXI-Lite container to expose its 32-bit registers to read from and write values to. The AXI-Lite components are memory-mapped, meaning that all components are assigned unique system-wide 32-bit base addresses and ranges. Consequently, the AXI Chip2Chip is also a memory-mapped component, meaning that a connected, physical device with all logic components inside is given a base address and a range. To target a particular register in a specific place in the system, one has to navigate through the connectivity scheme and assigned addresses. An exemplary, simplified scheme of addressing is presented in Fig. 11. Starting from the root level with a complete address range, transiting to a connected device

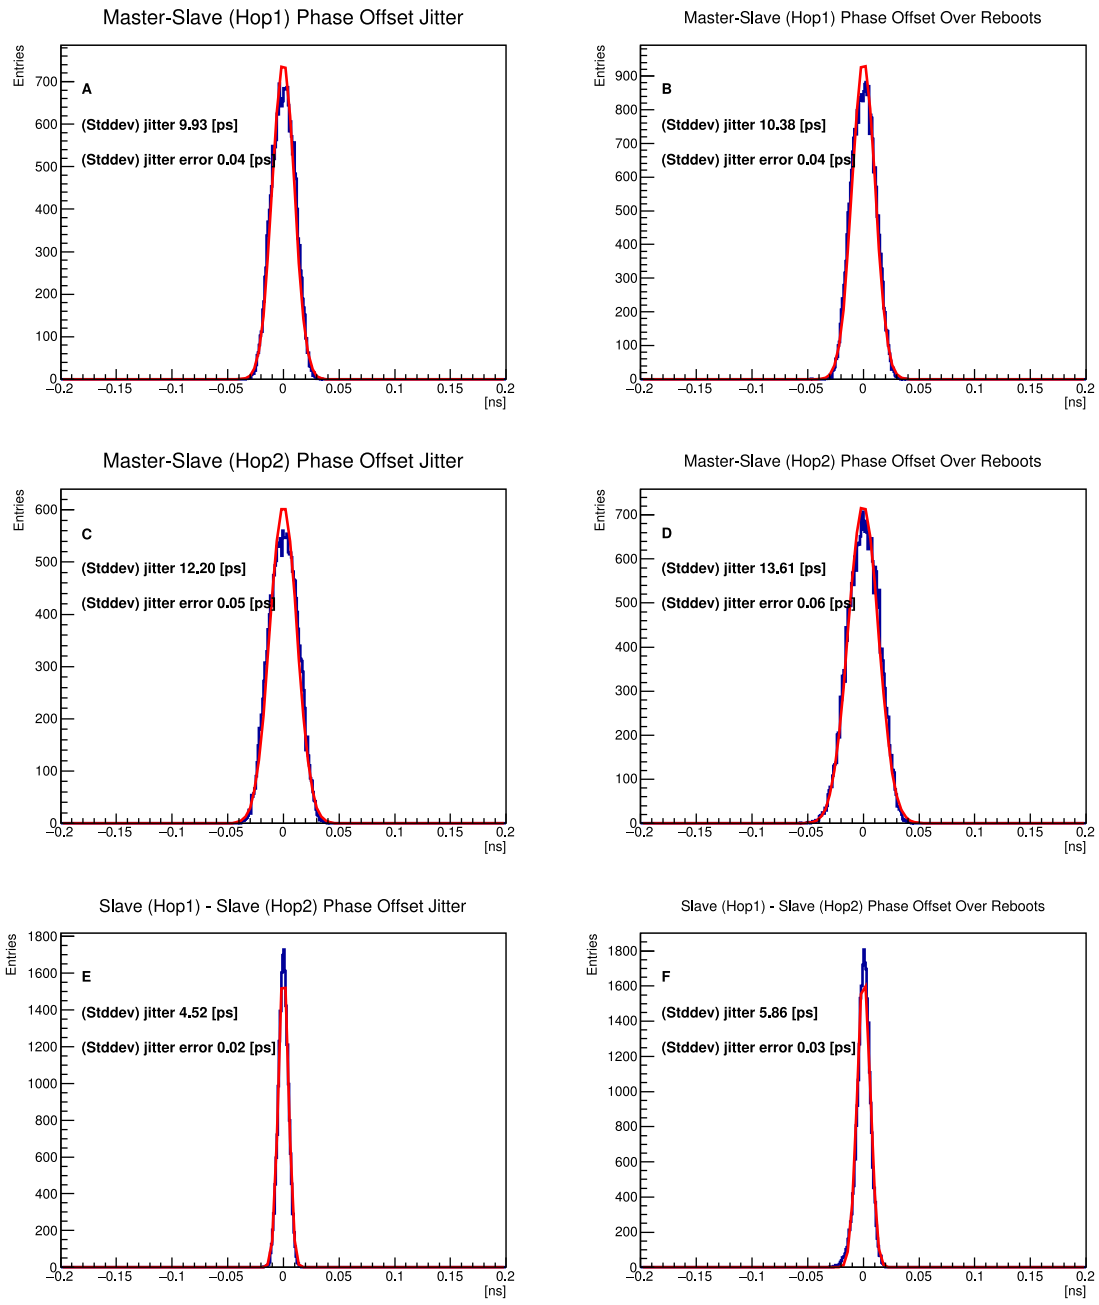


**Fig. 9.** Phase offsets jitter between the clocks of the system nodes measured with the JEDEC Standard 65B procedure (A, C, E) and including a power-cycle of the system in between the runs (B, D, F).

narrows down the available address ranges with each step. The default 32-bit address space can be extended to 64-bit on Zynq MPSoC for large systems [38].

Typically, the complete 32-bit addresses of logic components are assigned during the gateware design phase and thus become hardcoded into a generated FPGA configuration file (bitfile). This means that each node in the system requires an individually generated bitfile with carefully assigned addresses and must always connect to a specific location in the network. Instead, a configurable top-bits address mask is applied to each slave interface, allowing for the assignment of local component address ranges only, the generation of a single bitfile for a specific set of nodes, and the ability to swap the nodes without needing to rebuild FPGA designs.

### 5.2. Software

The above-presented addressing scheme is mapped into the virtual memory of the Petalinux running on ARM cores, which are part of the Zynq Processing System (bare-metal applications can also be developed). It essentially means that all logic components exposing registers through AXI-Lite interfaces are accessible to be read or written by simple read/write operations to a given memory address.

Having a Linux OS interfacing with the logic components is a convenient solution for building high-level tools for control and monitoring purposes. It can natively be connected to an Ethernet network and accessed remotely, and modern programming frameworks can be used to set up the system in operation rapidly. A Python script with ZeroRPC has been developed to expose generic register read and write functions to connected clients. The clients can be Python scripts run locally or remotely that require interaction with logic components. For each

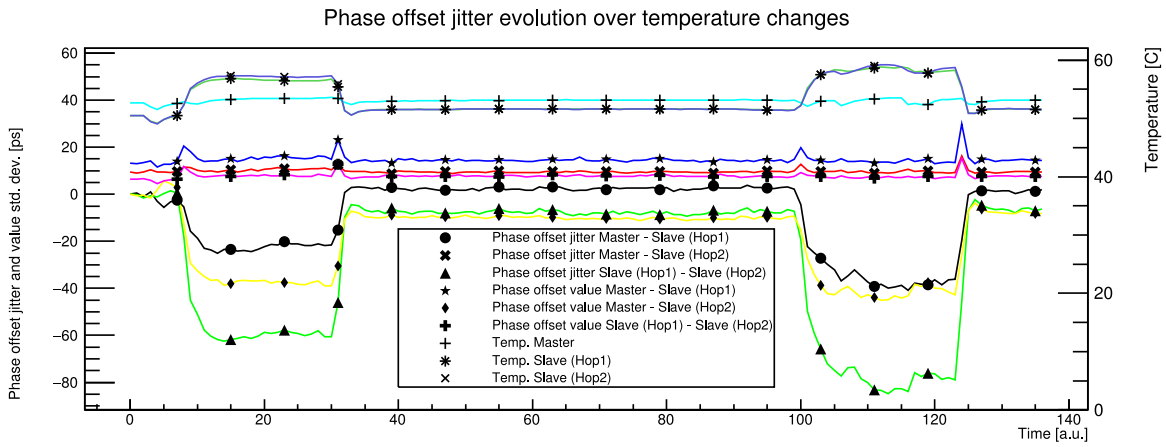


**Fig. 10.** Phase offset jitter and value between the clocks of the nodes in the system as a function of the FPGA temperatures. The graphs on top present the temperatures; only the concentrator and the endpoint cooling were altered during the measurement, while the master remained constant. The three graphs in the middle present the jitter with visible spikes at the temperature changing moments, and the three bottom graphs present the offset values. Each data point was collected for 8 min. Hence, the spikes in jitter graphs show accumulated phase drifts as a spread of the jitter distribution.

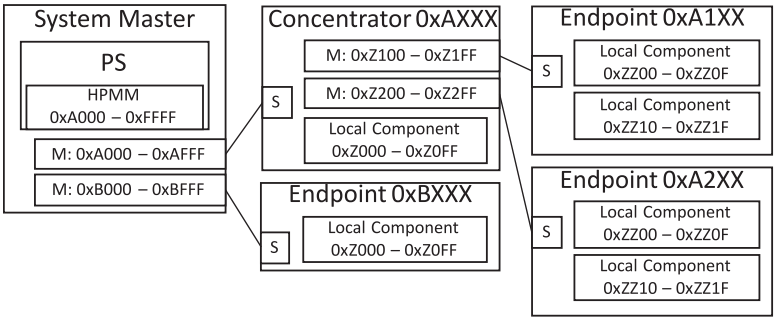


**Fig. 11.** The addressing scheme begins in the master AXI interface (HPMM) with a full address space, which is segmented by each forthcoming master interface (M), creating a sub-address space for all connected components (S), which can be further partitioned. Note the 'Z' in the sub-region addresses; these are masked bits to maintain unified FPGA designs between system nodes.

component, a class is implemented that provides high-level functions encapsulating the atomic register transactions. On top of that, one can implement a Command Line Interface (CLI) for fast, administrator-level system access or a modern Graphical User Interface (GUI) for operators using popular frameworks like Plotly or Grafana.

## 6. Hardware system implementation

The system is presented as a hardware platform agnostic, meaning it does not require dedicated and custom hardware components. The developed set of IP cores can be instantiated on any AMD/Xilinx FPGA with an exposed MGT transceiver. It opens a path to integrate a wide range of boards available from both the research and commercial sectors, significantly expanding accessibility and allowing groups to deploy complex, synchronous DAQ systems rapidly. Below are examples of the platforms used to build test setups, as well as those described in Section 3.4 for the prototype tomography scanner data acquisition system.

### 6.1. System master

For convenience, the system master platform in all tested setups was the ZCU102 Zynq Ultrascale+ MPSoC evaluation board (Fig. 12A). The System-on-Chip device with integrated ARM Cortex-A53 allows for accelerated system deployment and prototyping. As described in the previous section, the board can be easily operated via Ethernet, which provides a gateway for system control. The board alone is suitable for small setups, with 4 SFP+ ports (GTH), and it is possible to expand with up to 16 additional transceivers (GTH) routed through 2 FPGA Mezzannine Card (FMC) connectors in Vita 57.1 standard. The transceivers operate on the programmable low-jitter SI570 oscillator as an MGT reference clock, offering flexibility in the selected operating frequency for the synchronization system.

### 6.2. Data concentrator

The ZCU102 board, used as a system master, enables the connection of up to 20 endpoints. An intermediate layer of data concentrators must be added to expand the system. The system flexibility would allow the reusing of another ZCU102 board as a data concentrator, with minor modifications of the gateware design, such as replacing the synchronization generator with a synchronization slave receiver. However, the VCU108 Virtex Ultrascale evaluation board was used (Fig. 12B), offering up to 20 transceivers (GTH) on two FMC connectors. The Hitech Global HTG-FMC-X10SFP+ FPGA mezzanine cards have been used to expose the transceivers in the form of SFP+ ports.

The PANDA Data Concentrator platform (Fig. 12C), a uTCA-compatible board with Kintex Ultrascale+ and 5 Samtec FireFly optical bundles, has been integrated and evaluated. Each bundle contains 12 individual lanes, summing up to 60 possible slave connections (GTH and GTY). The platform is currently under development, and only a prototype was available, of which 24 transceivers have been put into operation.

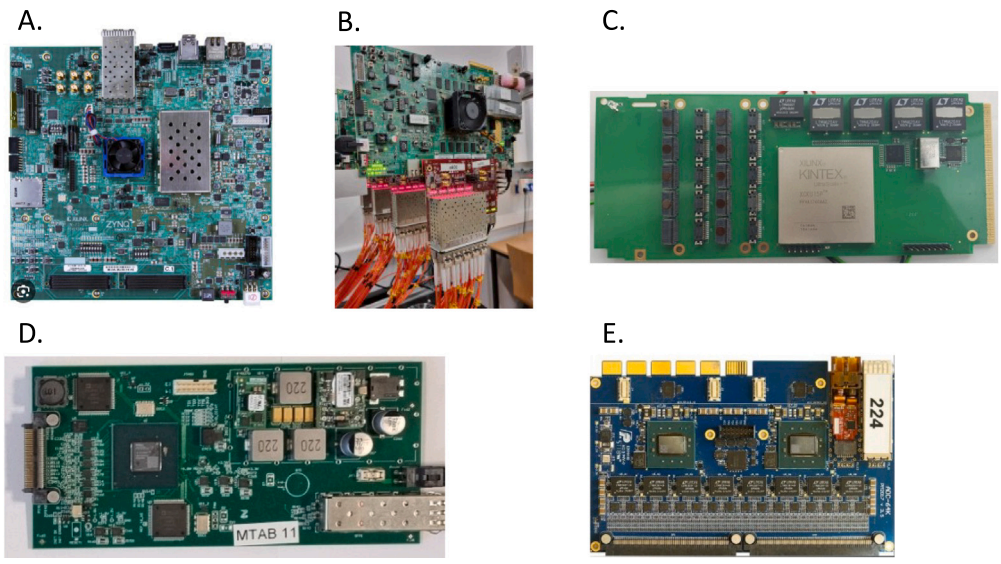


**Fig. 12.** Selection of the few hardware platforms evaluated with the Aurora-Sync framework. The ZCU102 system master is presented on A, the VCU108 as the data concentrator is presented on B, the PANDA Data Concentrator platform on C, and endpoint examples: the J-PET MTAB TDC platform and the SADC on D and E, respectively.

### 6.3. Endpoints

AC701 and KC705 evaluation boards have been used during the system development phase with Artix7 and Kintex7 FPGAs, respectively. They both offer a single SFP+ port with a GTP transceiver in the case of the first one and the GTX in the second. The boards were treated as generic endpoints with dummy data generators, allowing system synchronization, data transport, and control and monitoring mechanisms to be evaluated.

#### 6.3.1. Time-to-digital conversion

The first actual endpoint was the J-PET MTAB board - a custom platform for the digitization of Silicon PhotoMultiplier (SiPM) signals in Positron Emission Tomography (PET) applications (Fig. 12D). The board has the same Artix7 FPGA as the AC701 evaluation board, a single SFP+ port, and AD8000 amplifiers for 13 analog input channels. Each channel is then divided into 4 for multi-thresholding and routed to the FPGA's LVDS buffers (negative inputs) together with threshold levels (positive inputs). The buffers are used as discriminators, and the generated digital pulses are fed into tapped delay chains composed of the carry elements of the FPGA fabric to register the precise signal arrival and time-over-threshold values.

#### 6.3.2. Analog-to-digital conversion

The second production data source is the SADC platform designed initially for the PANDA experiment but due to its universality became popular in many other applications (Fig. 12E) [39]. The board offers 64 channels of 14-bit 80 Msps LTM9009 ADC operated by two Kintex7 FPGAs interconnected with each other with a GTX transceiver routed on the PCB and each FPGA having a dedicated SFP+ port.

Such hardware configuration showcases the flexibility of the system. In one scenario, each FPGA on the board can act as an independent endpoint and have its individual link connected to the data concentrator. Thanks to the inter-FPGA connection, it is also possible to designate one device as a synchronization receiver and a local data concentrator while the second device remains a regular endpoint. In this way, one physical data link can be spared.

## 7. Conclusions

A versatile system for multi-purpose, synchronous data acquisition systems has been developed and evaluated. The system is composed of gateware and software components that can be implemented on any hardware platform with an FPGA and a multi-gigabit transceiver, without requiring any additional clocking circuitry, such as an external jitter cleaner. No dedicated hardware platforms are taking the role of the master, data concentrators, or endpoints. Instead, the logic function in the system is defined by the configuration of logic component instances and hardware capabilities, namely, the number of exposed transceivers. The key element of the system is the data link, which provides master clock and synchronization pulse distribution in the downstream direction, measurement data transport in the upstream direction, and bidirectional exchange of control and monitoring messages.

The system has been successfully deployed for a tomography scanner prototype, allowing unique studies to be performed. The performance of the system has also been thoroughly evaluated in a laboratory setup to provide key figures.

The phase offset jitter of 10 ps enables us to perform measurements with distributed, synchronized, and high-precision TDCs and ADCs. The offset stability over multiple cycles of the system reboots allows reliable reproduction of the system state throughout measurement campaigns. The temperature's influence on the phase offset values has been characterized, showing no significant increase in jitter. The drift of the offsets can be stabilized by either offline corrections or control of the environment.

Inspection of the phase noise curves suggests limited clock signal disturbance over the processing stages between the master and the slave nodes. The RMS phase jitter of the recovered clock at the level of 2 ps, remains high as the value is affected by the noise of the measurement system, however, it is similar to the raw, unprocessed clock from the master node.

The above values can be put in context by comparing them to state-of-the-art systems such as TFC and White Rabbit (Table 1). Below is a table that summarizes the three systems:

High-precision clock distribution systems are commonly used in various measurement systems, including those that register the most fleeting physics phenomena. Usually implemented on highly specialized custom hardware with dedicated clocking circuitry, they provide

**Table 1**
Features and performance comparison.

| | AuroraSync | CERN TFC lpGBT [40] | White Rabbit [41][42] |
|---|---|---|---|
| Features | | | |
| Link rate | Configurable, default bidirectional 5 Gbps | Uplink: 5.12/10.24 Gbps<br>Downlink: 2.56 Gbps | Default: 1 Gbps, 10 Gbps under development |
| Hardware dependency<br>Clock distribution | Hardware agnostic<br>Single Aurora data link | Custom hardware<br>Dedicated optical link | WR compliant switches and nodes<br>Synchronous Ethernet link |
| Synchronization method | TX Phase Interpolator, PICXO DPLL | External jitter cleaner | Digital dual mixer time difference, SoftPLL, external VCXO |
| Metrics | | | |
| Peak-to-peak jitter<br>RMS phase jitter<br>Long-term drift | 10 ps<br>2 ps<br>Below 100 ps | 5 ps<br>Below 5 ps<br>Below 5 ps | HW dependent, typical 100 ps<br>Below 5.9 ps<br>Below 100 ps |

low jitter and long-term stability. In this contribution, we demonstrate that a system with similar capabilities can be implemented on Commercial-Off-The-Shelf platforms by utilizing the FPGA's built-in features.

## CRediT authorship contribution statement

**Grzegorz Korcyl:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Maciej Bakalarek:** Validation, Software, Methodology, Investigation, Formal analysis. **Paweł Moskal:** Funding acquisition.

## Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Grzegorz Korcyl reports equipment, drugs, or supplies was provided by Advanced Micro Devices Inc Austin. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgments

This work was supported by the National Science Centre of Poland through grants no. 2021/42/A/ST2/00423 and no. 2021/43/B/ST2/02150, the Ministry of Science and Higher Education through grant no. IAL/SP/596235/2023, the SciMat and qLife Priority Research Areas budget under the program Excellence Initiative–Research University at the Jagiellonian University.

I want to acknowledge the valuable support from P. Strzempek and K. Farbaniec throughout the project's development.

The J-PET MTAB hardware platform was developed by M. Kajetanowicz, and the TDC gateware was developed by P. Kapusta.

The PANDA SADC and PANDA Data Concentrator platforms were developed by P. Marciniewski.

The project could be realized thanks to the AMD University Program's support and donations.

## Data availability

No data was used for the research described in the article.

## References

[1] J. Albrecht, L. Bozianu, L. Calefice, S. Cella, C.E. Cocha Toapaxi, C. Doglioni, V.V. Gligorov, J.A. Gooding, K.E. Iversen, P. Inkaew, D. Magdalinski, A. Sopasakis, D.J. Wilson-Edwards, Summary of the trigger systems of the Large Hadron Collider experiments ALICE, ATLAS, CMS and LHCb, J. Phys. G: Nucl. Part. Phys. 52 (3) (2025) 030501, http://dx.doi.org/10.1088/1361-6471/adaadc.

[2] P. Muller, Y. Leblebici, Clock and data recovery circuit, in: CMOS Multichannel Single-Chip Receivers for Multi-Gigabit Optical Data Communications, Springer Netherlands, Dordrecht, 2007, pp. 127–180, http://dx.doi.org/10.1007/978-1-4020-5912-4.

[3] J. Cachemiche, P. Duval, F. Hachon, R.L. Gac, F. Réthoré, The PCIe-based readout system for the LHCb experiment, J. Instrum. 11 (02) (2016) P02013, http://dx.doi.org/10.1088/1748-0221/11/02/P02013.

[4] P. Moreira, et al., The GBT project, 2009, http://dx.doi.org/10.5170/CERN-2009-006.342.

[5] D. Campora Perez, A. Falabella, D. Galli, F. Giacomini, V. Gligorov, M. Manzali, U. Marconi, N. Neufeld, A. Otto, F. Pisani, V. Vagnoni, The 40 MHz trigger-less DAQ for the LHCb upgrade, Nucl. Instrum. Methods Phys. Res. Sect. A: Accel. Spectrom. Detect. Assoc. Equip. 824 (2016) 280–283, http://dx.doi.org/10.1016/j.nima.2015.10.047, Frontier Detectors for Frontier Physics: Proceedings of the 13th Pisa Meeting on Advanced Detectors.

[6] K. Chen, H. Chen, J. Huang, F. Lanni, S. Tang, W. Wu, A generic high bandwidth data acquisition card for physics experiments, IEEE Trans. Instrum. Meas. 69 (7) (2020) 4569–4577, http://dx.doi.org/10.1109/TIM.2019.2947972.

[7] J. de Cuveland, D. Emschermann, V. Friese, I. Froehlich, P. Gasik, D. Hutter, W.F.J. Mueller, C. Sturm, Technical Design Report for the CBM Online Systems - Part I DAQ and FLES Entry Stage, Tech. Rep., The CBM Collaboration, 2023.

[8] IEEE standard for a precision clock synchronization protocol for networked measurement and control systems, in: IEEE Std 1588-2008 (Revision of IEEE Std 1588-2002), 2008, pp. 1–269, http://dx.doi.org/10.1109/IEEESTD.2008.4579760.

[9] M. Lipiński, T. Włostowski, J. Serrano, P. Alvarez, White rabbit: a PTP application for robust sub-nanosecond synchronization, in: 2011 IEEE International Symposium on Precision Clock Synchronization for Measurement, Control and Communication, 2011, pp. 25–30, http://dx.doi.org/10.1109/ISPCS.2011.6070148.

[10] J.-L. Ferrant, M. Gilson, S. Jobert, M. Mayer, M. Ouellette, L. Montini, S. Rodrigues, S. Ruffini, Synchronous ethernet: a method to transport synchronization, IEEE Commun. Mag. 46 (9) (2008) 126–134, http://dx.doi.org/10.1109/MCOM.2008.4623717.

[11] E. Barrera, M. Ruiz, S. Lopez, D. Machon, J. Vega, PXI-based architecture for real time data acquisition and distributed dynamical data processing, in: 14th IEEE-NPSS Real Time Conference, 2005, 2005, p. 4, http://dx.doi.org/10.1109/RTC.2005.1547509.

[12] Y. Huang, S. Liu, J. Wang, X. Hu, C. Feng, Q. An, Development of a high resolution PXI based data acquisition system for electron momentum spectrometer, in: 2012 18th IEEE-NPSS Real Time Conference, 2012, pp. 1–4, http://dx.doi.org/10.1109/RTC.2012.6418154.

[13] D.G. Mathews, H. Acharya, C.B. Crawford, M.H. Gervais, A.P. Jezghani, M. McCrea, A. Nelsen, A. Atencio, N. Birge, L.J. Broussard, J.H. Choi, F.M. Gonzalez, H. Li, N. Macsai, A. Mendelsohn, R.R. Mammei, G.V. Riley, R.A. Whitehead, A flexible data acquisition system architecture for the nab experiment, 2024, arXiv:2407.17606, URL https://arxiv.org/abs/2407.17606.

[14] M. Venaruzzo, A. Abba, C. Tintori, Y. Venturini, FERS-5200: a distributed front-end readout system for multidetector arrays, 2020, arXiv:2010.15688, URL https://arxiv.org/abs/2010.15688.

[15] M. Kadziela, B. Jablonski, P. Perek, D. Makowski, Evaluation of the ITER real-time framework for data acquisition and processing from pulsed gigasample digitizers, J. Fusion Energy (2020) http://dx.doi.org/10.1007/s10894-020-00264-3.

[16] Xilinx, Aurora 8b/10b protocol specification, 2010, URL https://docs.amd.com/v/u/en-US/aurora_8b10b_protocol_spec_sp002.
[17] Xilinx, AXI Chip2Chip v5.0 LogiCORE IP product guide, 2022, URL https://docs.amd.com/r/en-US/pg067-axi-chip2chip.
[18] Xilinx, Designing IP subsystems using IP integrator, 2022, URL https://www.xilinx.com/support/documents/sw_manuals/xilinx2022_1/ug994-vivado-ip-subsystems.pdf.
[19] Texas Instruments, LMK04610 ultra-low noise and low power JESD204b compliant clock jitter cleaner with dual-loop PLLs, 2019, URL https://www.ti.com/lit/ds/symlink/lmk04610.pdf.
[20] D. Taylor, M. Klein, V. Vendramin, All digital VCXO replacement for gigabit transceiver applications (7 series/zynq-7000), 2021, XAPP589.
[21] E. Mendes, S. Baron, Member IEEE, C. Soos, J. Troska, P. Novellini, Achieving picosecond-level phase stability in timing distribution systems with xilinx ultrascale transceivers, IEEE Trans. Nucl. Sci. (2020) http://dx.doi.org/10.1109/TNS.2020.2968112.
[22] AMD, AMD supports new, long lifecycle FPGA designs through 2040, 2045, and beyond, 2024, URL https://community.amd.com/t5/adaptive-computing/amd-supports-new-long-lifecycle-fpga-designs-through-2040-2045/ba-p/702533.
[23] Definition of Skew Specifications for Standard Logic Devices, Tech. Rep., JEDEC Solid State Technology Association, 2003.
[24] Clock jitter definitions and measurement methods, 2024, URL https://www.sitime.com/clock-jitter-definitions-and-measurement-methods. (Accessed 08 2024).
[25] Daniel Bogdanoff, Why you should care about oscilloscope acquisition modes, 2024, URL https://www.electronicdesign.com/technologies/test-measurement/article/21801317/why-you-should-care-about-oscilloscope-acquisition-modes. (Accessed 08 2024).
[26] Jitter budget for 10 gigabit ethernet applications with sitime SiT9120/1 oscillators, 2024, URL https://file.elecfans.com/web1/M00/20/E6/ooYBAFmk3NGARmVAASLC37Kfl4161.pdf. (Accessed 08 2024).
[27] SI570 datasheet 10 MHZ to 1.4 GHZ I2C programmable XO/VCXO, 2022.
[28] D9010JITA/D9020JITA EZJIT complete jitter and vertical noise analysis and phase noise analysis for infiniium oscilloscopes, 2023.
[29] Phase noise to jitter calculator, 2025, URL https://www.sitime.com/phase-noise-jitter-calculator. (Accessed 06 2025).
[30] Aimee Kalnoskas, Application relevance of clock jitter, 2025, URL https://www.eeworldonline.com/application-relevance-of-clock-jitter. (Accessed 06 2025).
[31] Estimating period jitter from phase noise, 2021, URL https://www.skyworksinc.com/-/media/Skyworks/SL/documents/public/application-notes/AN279.pdf.
[32] P. Moskal, et al., Positronium imaging with the novel multiphoton PET scanner, Sci. Adv. 7 (42) (2021) http://dx.doi.org/10.1126/sciadv.abh4394.
[33] G. Korcyl, et al., Trigger-less and reconfigurable data acquisition system for positron emission tomography, Bio-Algorithms Med-Syst. 10 (1) (2014) http://dx.doi.org/10.1515/bams-2013-0115.
[34] G. Korcyl, et al., Evaluation of single-chip, real-time tomographic data processing on FPGA SoC devices, IEEE Trans. Med. Imaging 37 (11) (2018) 2526–2535, http://dx.doi.org/10.1109/TMI.2018.2837741.
[35] P. Moskal, E. Czerwinski, J. Raj, et al., Discrete symmetries tested at 10-4 precision using linear polarization of photons from positronium annihilations", Nat. Commun. 15 (2024) 78, http://dx.doi.org/10.1038/s41467-023-44340-6.
[36] P. Moskal, A. Gajos, M. Mohammed, et al., Testing CPT symmetry in ortho-positronium decays with positronium annihilation tomography, 2021, http://dx.doi.org/10.1038/s41467-021-25905-9.
[37] G. Korcyl, A Novel Data Acquisition System Based on Fast Optical Links and Universal Read-Out Boards (Ph.D. thesis), AGH-UST, Cracow, 2015.
[38] Xilinx, Zynq UltraScale+ MPSoC example design: Using 64-bit addressing with AXI dma, 2023, URL https://adaptivesupport.amd.com/s/article/70413?language=en_US.
[39] J.S. Müllers, An FPGA-Based Sampling ADC for the Crystal Barrel Calorimeter (Ph.D. thesis), U. Bonn (main), 2019.
[40] D.H. Montesinos, S. Baron, Phase-Noise and Phase Variation Tests of the Lpgbt, Tech. Rep., CERN, EP-ESE Department, 1211 Geneva 23, Switzerland, 2023, URL https://vldbplus.web.cern.ch/template/static/docs/Phase_stability_tests_of_the_lpGBT.pdf.
[41] J. Savory, J. Sherman, S. Romisch, White Rabbit-Based Time Distribution at NIST, Tech. Rep., National Institute of Standards and Technology (NIST), Boulder, CO, USA, 2023, URL https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=925954.
[42] M. Rizzi, M. Lipiński, T. Włostowski, J. Serrano, G. Daniluk, P. Ferrari, S. Rinaldi, White Rabbit Clock Characteristics, Tech. Rep., CERN and University of Brescia, Geneva, Switzerland, 2023, URL https://white-rabbit.web.cern.ch/documents/White_Rabbit_Clock_Characteristics.pdf.