

TCLink: A Fully Integrated Open Core for Timing Compensation in FPGA-Based High-Speed Links

Eduardo Mendes¹, Sophie Baron², Jeroen Hegeman², Jan Troska², and Nikitas Loukas

Abstract—The high luminosity expected in the second phase of the upgrades of the Large Hadron Collider (LHC phase-2 upgrades) will pose unprecedented challenges to its four experiments in terms of collisions density—also known as pile-up—per beam crossing. Disentangling the vertices of 200 simultaneous collisions every 25 ns requires high granularity in the detectors, as well as extremely precise and stable timing. While short-term timing stability is usually a concern addressed in timing distribution systems, long-term variations due to changing environmental conditions can accumulate through distribution chains and can dominate the overall timing stability of the systems they serve. Timing distribution systems in LHC experiments typically use high-speed links and clock recovery. This article presents a logic core that can be used to mitigate long-term temperature variations in high-speed links. The timing compensated link (TCLink) is an open-source firmware core fully integrated in Xilinx Ultrascale Field Programmable Gate Arrays (FPGAs). It demonstrates picosecond-level phase precision over timing distribution systems, improving the overall timing stability in physics experiments.

Index Terms—Fast timing, field programmable gate arrays (FPGAs), high energy physics (HEP) instrumentation, optical links, timing circuits.

I. INTRODUCTION

THE EUROPEAN Council for Nuclear Research (CERN) has been operating the world's largest, most powerful particle accelerator, the Large Hadron Collider (LHC) since 2009. The LHC collides bunches of particles at a frequency of 40.0789 MHz (derived from the accelerator radio frequency and known as the LHC Bunch Clock) [1]. This signal, crucial for the detectors, is distributed to the four experiments located around the 27-km ring of the LHC. Within each of these experiments, the clock is re-distributed to thousands of end-nodes located in a harsh environment.

A major requirement of the distribution system is the bunch clock reaching each end-node with a fixed and deterministic phase relationship to both the clock source and the other end-nodes. We refer to a fixed and deterministic relation when the phase of the clock between two nodes changes as little as possible between system startups and remains fixed over time, regardless of its absolute value.

Manuscript received 20 December 2022; accepted 17 January 2023. Date of publication 30 January 2023; date of current version 15 February 2023.

Eduardo Mendes, Sophie Baron, Jeroen Hegeman, and Jan Troska are with the Experimental Physics Department, CERN, 1211 Geneva, Switzerland (e-mail: eduardo.brandao.de.souza.mendes@cern.ch).

Nikitas Loukas is with the Department of Physics, University of Notre Dame, Notre Dame, IN 46556 USA.

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TNS.2023.3240539>.

Digital Object Identifier 10.1109/TNS.2023.3240539

While current LHC detectors can tolerate phase variations of the order of a hundred picoseconds, their upgrades to match the high-luminosity upgrade of the LHC (HL-LHC) bring new challenges in terms of timing distribution. To disentangle the large number of collisions occurring simultaneously in their centers, some experiments like A Toroidal LHC Apparatus (ATLAS) and Compact Muon Solenoid (CMS) have foreseen the installation of high-precision timing detectors [2], [3]. This poses severe constraints on the timing distribution systems stability, limiting the tolerance of phase variations to a few tens of picoseconds. With such a tight specification, all sources of uncertainty must be mitigated along the entire distribution chain.

Although no common solution is employed by the experiments, a tree-like structure constructed of a cascade of point-to-multipoint networks such as passive optical networks (PONs) [4] or point-to-point optical links are typically used, along with modern field programmable gate arrays (FPGAs) as shown in Fig. 1. This is followed by one final stage consisting of an FPGA connected to an ASIC via an optical link. For the phase-2 upgrades, the connection to the front-end consists of the CERN-specific point-to-point Versatile Link+ [5] coupled to the low-power Gigabit Transceiver (lpGBT) ASIC [6] inside the detectors.

The short-term stability is mainly ensured by this last hop of the network (back-end FPGA, jitter cleaning PLL, and lpGBT). Previous studies have shown that the short-term stability specifications can be comfortably met with such a solution [7].

The long-term instability due to restarts or temperature variations, however, accumulates over the full-chain. Part of this instability is already addressed in [8].

To mitigate long-term variations, several solutions exist in literature exploiting the bidirectionality of an optical link. The White-Rabbit project [9], which is a timing link where the clock characteristics (frequency and phase) are encoded in the data and transmitted over fibers to FPGAs, offers phase and absolute timing compensation with sub-nanosecond accuracy and tens of picoseconds of precision. This project employs an external VCXO to implement a phase-shifter and it is based on the 1-Gb/s ethernet protocol with wavelength division multiplexing. Although efficient, this solution requires a specific protocol, network switch, and dedicated hardware which are not compatible with the low space occupancy and the radiation tolerance requirements of parts of the HL-LHC systems.

In [10], a White-Rabbit inspired solution is demonstrated, which targets very long-haul systems using a specific protocol.

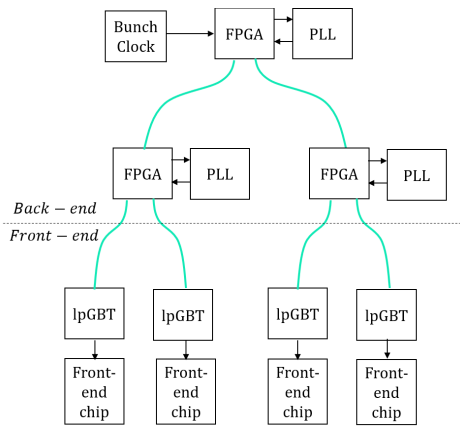


Fig. 1. Generic example of a timing distribution chain in an HL-LHC experiment. The clock is distributed via high-speed optical links based on FPGAs in the back-end and jitter cleaning PLLs. In the front-end, the rad-hard ASIC lpGBT is used to recover the clock.

In this project, identical wavelengths are adopted for the uplink and downlink which require space consuming circulators to ensure a perfect symmetry of the link. Again, such requirements are not compatible with the high density of the front-end electronics in HL-LHC experiments.

In [11], a multi-transceiver synchronization scheme is presented employing TDCs. The goal of this project is the synchronization of multiple transceivers inside an FPGA and not a full high-speed link.

In this article, we propose a fully integrated FPGA core for compensating for long-term timing variations in high-speed optical links. The core is only implemented in one end of the link and the logic resources can be reduced if the control loop is implemented in software. The core is also fully programmable, leaving the user the freedom of bandwidth choice and other parameters. It can be used in multiple links of an FPGA to provide independent compensation in real time.

The core is protocol-agnostic in the sense that it can be configured for different data rates and makes no use of the underlying transmitted data information. Therefore, the experiment designers can employ the most suitable protocol for their application.

This solution is fully compatible with the lpGBT ASIC designed at CERN, but other systems requiring long-term phase-stability can also freely employ the core as the design sources are open-source licensed [12].

The first proof-of-concept of a fully FPGA-integrated timing compensation scheme (on which this work is based) is initially demonstrated by us in [13]. In this article, we consolidate the core implementation and show its performance in different metrics.

II. TIMING COMPENSATION

In the LHC experiments, the Bunch Clock must be distributed with a fixed and deterministic relation to the end nodes. The absolute phase is not relevant since a calibration based on physics is performed in the system. However, it is desirable that minimal phase variations occur over time to avoid the need of re-calibrating the system. Therefore, timing compensated link (TCLink) does not correct for an absolute




Fig. 2. Concept of a timing compensation scheme for high-speed links. The timing compensation here refers to the long-term phase stability rather than an absolute timing distribution.




Fig. 3. TCLink simplified block-diagram—all the blocks are fully integrated in the FPGA requiring no external components.

phase but rather performs adjustments due to phase variations over time, maintaining a long-term stability.

The principle of the TCLink timing compensation method is shown in Fig. 2. A master node transmits information to the follower node in the downlink direction. The follower must, in turn, reuse its recovered clock for the uplink transmission. The roundtrip variation (Δt_{RT}) is given by the sum of the downlink and uplink variations (Δt_D and Δt_U). An assumption must be made linking the downlink variation and the roundtrip variation for the correction ($\Delta t_D = \alpha \times \Delta t_{RT}$). Typically, we set $\alpha = 0.5$, which assumes a perfect symmetry of the link but the user has the flexibility of choosing their own α coefficient. The timing compensation employed here does not target an absolute delay of the timing but rather a relative phase-variation with respect to an initial set-point.

III. IMPLEMENTATION

A high-level overview of the TCLink core is shown in Fig. 3. The core is only present on the master side, while the follower side simply requires careful configuration of its transceiver IP. All the VHDL cores of TCLink are publicly available on Gitlab [14]. A phase detector, controller, and phase-shifter are integrated in the Master FPGA as shown in Fig. 3. The phase detector measures the round-trip variation by comparing the returned clock to the master and a controller adjusts the phase-shifter to compensate for the measured variation.

A. Phase Detector

The phase detector used in the TCLink is based on the digital dual mixer time difference (DDMTD) core employed in the White-Rabbit project [9]. The DDMTD measures the phase difference between two clocks of equal frequency, here noted f_c , using a third clock with a frequency (f_{dmtd}) close to the carrier to sample the signals being measured. This auxiliary clock is generated inside the FPGA using a mixed-mode clock manager (MMCM). The phase-resolution obtained in this technique depends on this offset clock frequency and can reach a ps-resolution.




Fig. 4. Phase-shifter model similar to a DCO where the output phase is the input control signal accumulated.

For an averaging factor given by N_{avg} , the sampling frequency is given by (1). The DDMTD resolution is shown in the following equation:

$$f_s = \frac{f_{dmtd} - f_c}{N_{avg} + 1} \quad (1)$$

$$DMTD_{res} = \frac{f_{dmtd} - f_c}{N_{avg} f_c f_{dmtd}}. \quad (2)$$

In addition to being a highly used core, there are several motivations in using the DDMTD as a phase detector. This circuit provides a good phase measurement linearity and resolution. The DDMTD can also have a sampling frequency in the order of kHz which is enough for our application. In this implementation, a bit-median deglitcher is used with a metastability window of 500 ps [15].

B. Phase-Shifter

The phase-shifter is the hard-IP phase-interpolator, which is only available in Xilinx Ultrascale and Ultrascale+ transceivers [16], [17]. This feature is the key ingredient of the high-precision timing distribution (HPTD) IP [8], which provides a fixed-phase of the transceiver across start-ups while using the transmitter FIFO (a requirement when using the phase-interpolator). The phase-shifter resolution is in the order of ps.

We model the phase-shifter similar to a digitally controlled oscillator (DCO) shown in Fig. 4. Here the variable z denotes the z -transform domain where z^{-1} indicates a delay. To simplify the loop modeling, we also include the DDMTD resolution in the DCO.

The DCO transfer function is therefore given by the following equation where K_{dco} is the phase-shifter resolution and K_{dmtd} is the inverse of the DDMTD resolution:

$$DCO(z) = K_{dmtd} K_{dco} \frac{z^{-1}}{1 - z^{-1}}. \quad (3)$$

C. Controller

The controller is a custom VHDL block shown in Fig. 5. Its input is the data from the phase detector. It maintains the phase around an offset point set by the user during a calibration procedure in which this initial offset is measured. An integral-proportional controller is employed, followed by an oversampling block based on sigma-delta modulation which provides a set of pulse streams to either advance or retard the phase. A feedback path, named mirror path, is implemented digitally and accumulates the current phase. The mirror path is employed to partially compensate for the total roundtrip variation with a given α coefficient set by the user.

In addition, the controller contains a tester for testing the loop dynamics. The tester contains a numerically controlled oscillator (NCO) which is used to inject a digital sinusoidal




Fig. 5. TCLink controller is a full digital implementation inspired in the world of All Digital PLLs. A mirror path is included to allow for different α coefficients.




Fig. 6. Loop control model containing an integral-proportional controller.

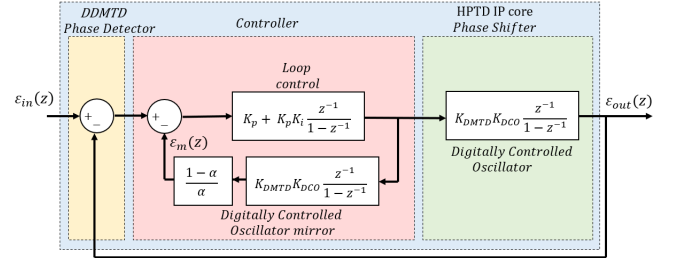


Fig. 7. TCLink loop model including the phase detector, loop control, and phase-shifter.

input to the phase detector. By measuring the accumulated output rms phase for different frequencies of the NCO, the transfer function of the loop can be plotted. This transfer function can then be used to verify that the loop dynamics are behaving as expected.

The model for the loop control is demonstrated in Fig. 6 and its transfer function is shown in the following equation:

$$L(z) = K_p + K_p K_i \frac{z^{-1}}{1 - z^{-1}}. \quad (4)$$

The full mathematical model of the TCLink loop (including the phase detector, controller, and phase-shifter) is shown in Fig. 7. This model is a simplification of the real system. The internal PLLs and the link latency are not included: this is a reasonable hypothesis since the TCLink loop bandwidth (in the order of 100 Hz) is much slower than the jitter cleaning PLL bandwidth (in the order of kHz) or the clock and data recovery bandwidth (in the order of MHz). The sigma delta is also not included here since it operates at a frequency considerably higher (at least 16 times) than the loop sampling frequency.

The TCLink closed-loop transfer function is given by the following equation, where $K = K_{dmtd} K_{dco} K_p$:

$$\frac{\epsilon_{out}(z)}{\epsilon_{in}(z)} = \alpha \frac{K z^{-1} (1 - z^{-1} + K_i z^{-1})}{(1 - z^{-1})^2 + K z^{-1} (1 - z^{-1} + K_i z^{-1})}. \quad (5)$$

To calculate the loop parameters, we apply a time-continuous approximation for the discrete model [18]. This is a valid approach since the loop bandwidth is considerably lower than the loop sampling frequency f_s . Under this consideration, we substitute (6) and (7) in (5) to obtain (8). The variable s here denotes the variable used in

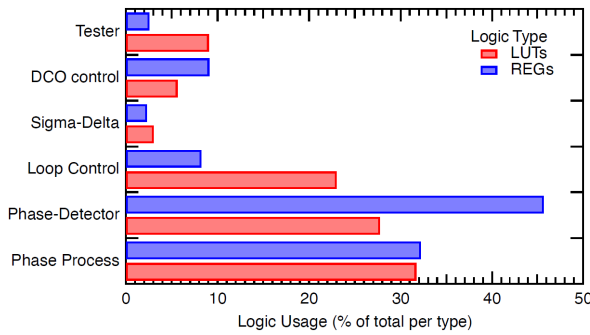


Fig. 8. TCLink Master: Logic usage. The usage is expressed as a percentage of the total number of LUTs and REGs.

the Laplace transform domain

$$z^{-1} \approx 1 \quad (6)$$

$$1 - z^{-1} \approx \frac{s}{f_s} \quad (7)$$

$$\frac{\epsilon_{\text{out}}(s)}{\epsilon_{\text{in}}(s)} = \alpha \frac{s K f_s + K K_i f_s^2}{s^2 + s K f_s + K K_i f_s^2}. \quad (8)$$

The poles of (8) can be calculated in terms of the damping coefficient ζ and the loop bandwidth f_n using the following equation:

$$s^2 + s K f_s + K K_i f_s^2 \leftrightarrow s^2 + 2\zeta(2\pi f_n) + (2\pi f_n)^2. \quad (9)$$

Therefore, we obtain the loop coefficients shown in the following equations:

$$K_p K_i = \left(\frac{2\pi f_n}{f_s} \right)^2 \frac{1}{K_{\text{dmtd}} K_{\text{dco}}} \quad (10)$$

$$K_p = 2\zeta \sqrt{\frac{K_p K_i}{K_{\text{dmtd}} K_{\text{dco}}}}. \quad (11)$$

If an integral part is not adopted in the controller (default behavior), the K_p coefficient can be calculated using the following equation:

$$K_p = \left(\frac{2\pi f_n}{f_s} \right) \frac{1}{K_{\text{dmtd}} K_{\text{dco}}}. \quad (12)$$

To save area, no multipliers are used in this implementation. The integral-proportional coefficients are set as a power of two using (13) so the multiplication corresponds to a bit shift. In addition, the DCO mirror accumulator unit is given by its scaled coefficient

$$K_{p2} = 2^{\text{round}(\log_2 K_p)}. \quad (13)$$

D. Resource Usage

The typical implementation resource usage for TCLink (concentrated in the Master FPGA) is around 1200 configurable logic block (CLB) look-up tables (LUTs) and 800 CLB registers for a single core. Logic usage for the different elements of the core are shown in Fig. 8.

In the case of a software-based TCLink control loop implementation, only the phase detector would need to remain in the FPGA together with the transceiver IP, resulting in a saving of around 72% of the LUT and 54% of the registers.

Although not currently implemented in our design, for designs implementing multiple TCLinks in the same FPGA, several blocks can be potentially shared to reduce the resource usage. The phase detector and phase process can be shared using a multiplexed phase detection. In addition, the loop control and tester can also be shared.

IV. EXAMPLE DESIGN

An example design is provided together with the core itself. This example is based on a typical front-end link used in the HL-LHC experiments, called Versatile Link Plus. The links transmits the so-called lpGBT protocol which uses a fixed header combined with Reed-Solomon Forward Error Correction and scrambling at 10.24 Gb/s. A symmetric variant of this protocol is proposed for back-end to back-end communication (FPGA-to-FPGA) while the standard asymmetric data-rate lpGBT protocol is used for the example design proposing the back-end to front-end link (FPGA-to-lpGBT). Despite the lpGBT being an asymmetric protocol (2.56-Gb/s downlink, 10.24-Gb/s uplink), the transceivers implemented are configured at 10.24 Gb/s using a four times bit-folding in the transmitter path to maintain symmetry.

In this configuration, the transceiver user clocks are configured at 320.632 MHz and the DDMTD offset clock at 320 MHz. Such a configuration yields a resolution of 6.160 ps for each DDMTD measurement, which is then further averaged over 64 acquisitions. For this data rate, the transmitter phase-interpolator bin size is 1.523 ps by design. The sigma-delta oversampling factor is set to 32.

In case another data and physical layer is desired, a.csv file is available to the user containing the main parameters for the application related to the link (data rate, clock frequencies) and to the system dynamics (bandwidth, DDMTD resolution, sigma-delta oversampling ratio, α coefficient). This file is read by a high-level model in Python which calculates the transfer-function for a given TCLink implementation. From the user high-level parameters, the model calculates the internal loop controller parameters. The model then performs a numerical simulation which is used to trace the TCLink transfer function. The model also calculates the HDL port values required by the core for proper operation.

V. CHARACTERIZATION RESULTS

A. Loop Dynamics

The full loop (including the fiber and follower) dynamics results are shown in Fig. 9 where the bandwidth values for different TCLink core settings are demonstrated and compared to their Python high-level model. The gain was calculated as the ratio (in decibels) of the output phase (in rms) to the tester input (in rms). At low frequencies, the transfer function is around -6 dB, corresponding to the α coefficient of 0.5 used in the example design. The default value of the bandwidth is now set at 100 Hz.

B. Temperature

The setup used for characterizing the TCLink under temperature variation is shown in Fig. 10. An external High

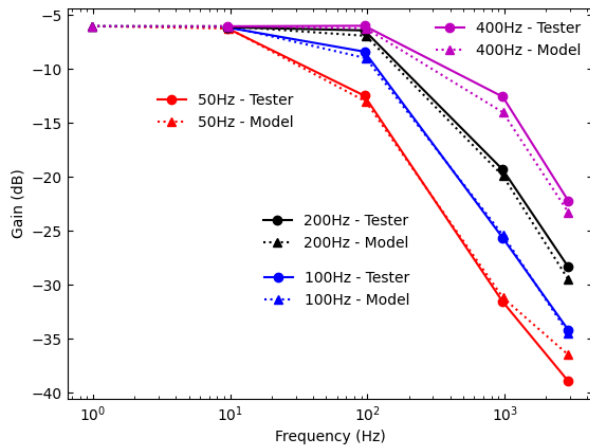


Fig. 9. TCLink bandwidth adaptability. A good agreement in the transfer function was obtained between the model and the FPGA tester for different bandwidth tested.




Fig. 10. Setup for temperature characterization tests. Each part of the setup (KCU105, fibers, and VLDB+) was placed inside a climate chamber for performing temperature variations.

Precision Timing Clock generator [19] provided the reference clock for the measurements. Eight lpGBT-FPGA cores (each containing a TCLink master) were implemented on a single Xilinx KCU105 evaluation board. These were connected to a Samtec Firefly FMC.

A Versatile Link+ Demonstrator Board (VLDB+) [20] was used to emulate the front-end side. This board contains a Versatile Transceiver+ (VTRx+) and an lpGBT. The VTRx+ was connected to a Firefly using two fibers.

Tests were performed by successively putting each of the parts of the setup (KCU105, fibers, VLDB+) in turn inside a climate chamber. The phase measurement was made by a Keysight DSO9254A oscilloscope. Edge-Edge measurements were performed during single-shot acquisitions using the maximum resolution available in the instrument, yielding an average of 82 000 values. The amplitude of the temperature variation was much larger than typical HL-LHC operating conditions, to properly assess its impact on the link.

The KCU105 (Master) was initially installed inside the chamber. As shown in Fig. 11, the impact of the compensation provided by the TCLink can be observed but is quite limited. Such performance was expected as part of the FPGA logic was not included in the feedback loop of the TCLink.

In the second testing stage, the VLDB+ (playing the role of the Follower) was placed inside the chamber. As shown in Fig. 12, the same effect can be observed, as the additional logic (multiplexers, clock-trees) of the lpGBT was not included in the feedback loop and therefore cannot be compensated by TCLink.

Finally the results of the fibers subject to temperature variation are shown in Fig. 13. This time an almost perfect




Fig. 11. Results for KCU105 inside climate chamber. The TCLink control loop was not able to correct for FPGA temperature variations.

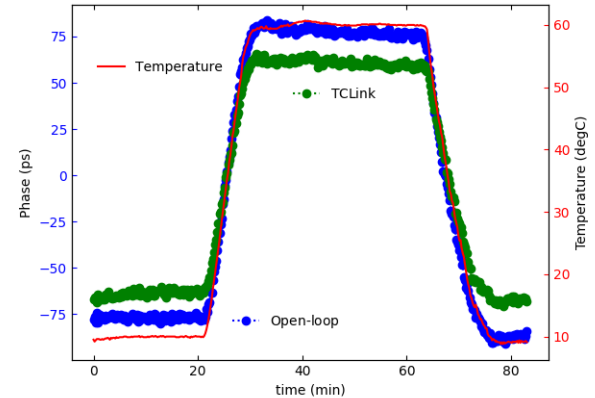


Fig. 12. Results for VLDB+ inside climate chamber. The TCLink control loop was not able to correct for lpGBT temperature variations.




Fig. 13. Results for fiber inside climate chamber. The TCLink control loop was able to correct for fiber temperature variations.

compensation was observed, as anticipated: the fibers were fully included in the loop and the upstream and downstream path were considered as perfectly symmetric as a first estimation.

In high-energy physics (HEP) experiments, Master and Follower electronics will be located in cooled places where the environment temperature is well controlled and is not expected to vary by more than a few degrees. The fibers, however, could be quite long and subject to large environmental changes (day-night or seasonal). Therefore, the performance presented above is fully compatible with the intended application.




Fig. 14. Setup for phase-determinism tests. A master and follower were implemented in the same FPGA evaluation board. The phase was measured using an oscilloscope for different system restarts.

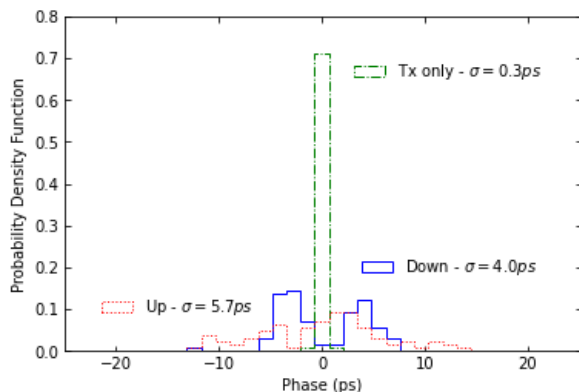


Fig. 15. Phase-determinism results for ZCU102 in open-loop mode. The full downlink and uplink phase variation had the same order of magnitude. The transmitter phase variation was much smaller indicating the phase uncertainty does not come from the transceiver transmitter.

C. Phase Determinism

Phase determinism is an important figure of merit for the HL-LHC experiments. It shows how much the recovered clock phase of a link changes with system start-ups (such as an FPGA reload or a fiber being disconnected). The setup used for characterizing the phase-determinism of the TCLink is shown in Fig. 14. Here also, an external High-Precision Timing Clock generator [19] provided the reference clock for the measurements. For this test, a TCLink was established between a master and follower implemented in the same FPGA but with a fully independent clocking. The test was performed for the example design migrated to four different Xilinx evaluation boards: ZCU102, VCU118, and KCU116 (two different boards noted here as B1 and B2). In addition, the test was also performed on one TCLink user implementation: the CMS barrel calorimeter processor (BCP) board—Version 1 (based on a Kintex Ultrascale FPGA) [21]. In the BCP, the design is implemented together with a realistic user logic occupancy and a similar performance is observed.

The test was performed for 100 system start-ups per board. For each startup, a phase measurement was performed similar to what is explained in Section V-B. The downlink phase was measured between the 40.0789-MHz reference and the downlink recovered clock; the uplink phase was measured between the downlink recovered clock and the uplink recovered clock. For this test, the link was in open-loop mode. We also measured the TCLink internal phase-measurement circuitry. The temperature was measured at each acquisition and for all tests the variation was smaller than 1° for the total duration of the test.

TABLE I
PHASE-DETERMINISM RESULTS FOR ALL BOARDS

Evaluation Board	Downlink (ps)	Uplink (ps)
ZCU102	4.0	5.7
VCU118	1.3	4.8
KCU116-B1	2.8	2.2
KCU116-B2	3.2	7.6
CMS BCP-B1	4.6	Not measured
CMS BCP-B2	4.0	Not measured




Fig. 16. Correlation between phase-measurement circuitry and downlink phase. The correlation between the DDMDT measurement and downlink scope measurement was poor for all boards measured here.

Fig. 15 presents the results for the ZCU102 board in open-loop mode. We measured a standard deviation of 4.0 ps for the downlink and 5.7 ps for the uplink. We also measured the master transmitter phase determinism and we obtained a 0.3-ps standard deviation. Table I shows the results for all the boards. There was a variation of 1.3–4.6 ps for the downlink and 2.2–7.6 ps for the uplink.

To evaluate whether the activation of the TCLink compensation loop could correct for the phase determinism observed here, we compared the downlink phase variation measured by the scope with half of the phase variation measured by the phase-measurement circuitry (DDMDT). Fig. 16 highlights that there is a poor correlation between these two variables. The Pearson coefficient (ρ) was calculated for each board and the values were found to be far from the ideal value of 1.

This result implies that the TCLink loop is unsuitable for correcting for phase-determinism. This is somehow expected because the transmitter and receiver of master and follower are not symmetrical. This can be further observed in Fig. 17. In open-loop, the downlink standard deviation was 4.0 ps and when TCLink was active the standard deviation was 4.2 ps.

D. Jitter Impact

A small penalty in terms of jitter on the recovered clock was observed when closing the TCLink loop. Results measured with a Rohde and Schwarz FSWP-8 are shown in Fig. 18. The integrated phase-noise from 100 Hz to 10 MHz for the open-loop configuration is about 1.5 ps while for the TCLink it is 1.6 ps. In addition, new spurs at around 150 kHz are present. These spurs are related to the delta-sigma modulator and they appear at half the sigma-delta modulation frequency and its odd harmonics.

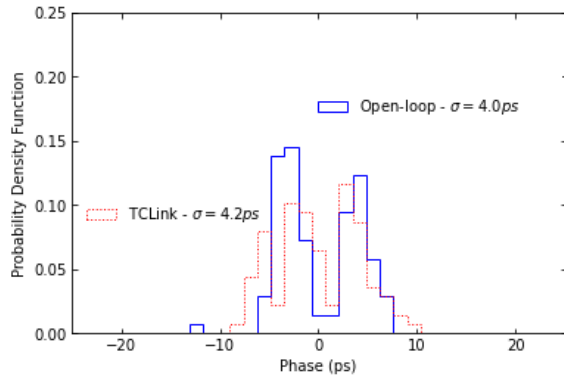


Fig. 17. Phase-determinism downlink results for ZCU102 in open-loop and closed-loop. The TCLink loop was not able to improve the phase determinism of a link.

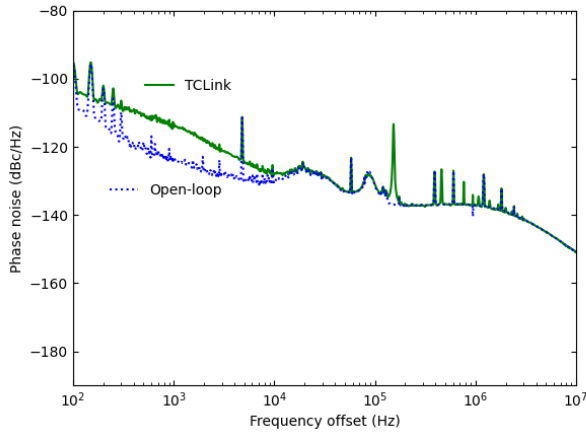


Fig. 18. Phase-noise measurements for open-loop and closed-loop. When the TCLink loop was closed, a small penalty in terms of phase-noise was observed.




Fig. 19. Setup for cascaded tests. Three cascaded TCLink were implemented (two optical links and one backplane link) using the CMS DTH board. The phase was measured between the reference clock and the clock recovered in the last branch.

This results shows that TCLink is not effective for correcting jitter in the recovered clock. The reason for this is because the jitter added in the downlink and uplink paths are not correlated.

If the jitter penalty is prohibitive for certain applications, the TCLink control could be operated in a semi-automatic mode. In such scenario, the loop is by default open, and during beam gaps (or when the measured phase variation is high) the loop is closed allowing the compensation. This way, for most of the time, the jitter penalty will not be observed.

E. Cascaded System Test

For the cascaded system test, the realistic setup of three cascaded TCLinks shown in Fig. 19 was used. The data and trigger hub (DTH) board [22] based on a Virtex Ultrascale+ from CMS was used to build the setup. The setup consisted of two optical links and a back-plane link. Fig. 20 shows the




Fig. 20. Results in open-loop for realistic experiment setup with a single branch. For the small temperature variation measured (around 1°), a phase variation of 30 ps was measured.

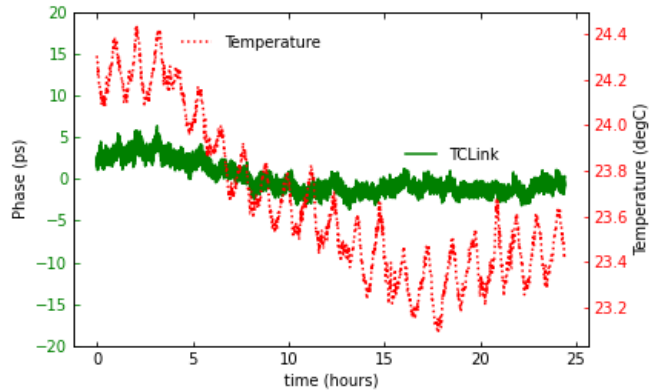


Fig. 21. Results in closed-loop for realistic experiment setup with a single branch. For the small temperature variation measured (around 1°), a phase variation of 6 ps is measured.

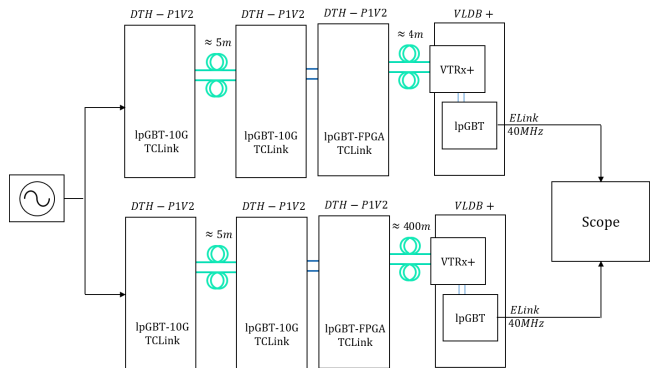


Fig. 22. Setup for cascaded tests including two branches. For each branch, three cascaded TCLink were implemented (two optical links and one backplane link) using the CMS DTH board. The phase was measured between the clock recovered in the two branches.

results from one day of phase measurement when the TCLink loop was open. We observe variations of around 30-ps peak to peak for one day in a relatively stable laboratory environment (temperature variations were around 1°). Fig. 21 shows the results when the TCLink loop was closed—we still observe some variations due to the environment, but the peak-to-peak value is much smaller.

The experiment was performed simultaneously with a duplicated setup as shown in Fig. 22, where the measurement was performed between the two end-nodes' lpGBT Elink clocks. The results for the open-loop and closed-loop TCLink are shown in Figs. 23 and 24, respectively. There was a much

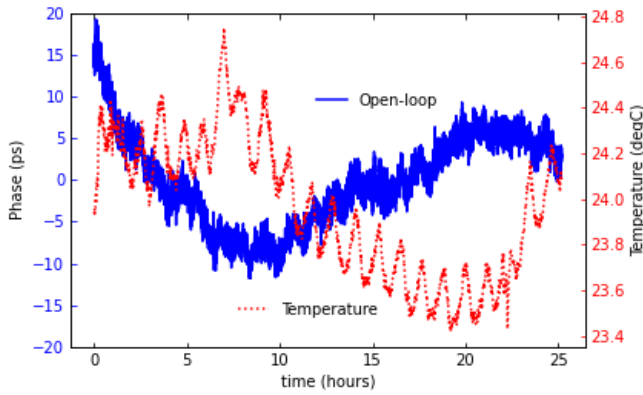


Fig. 23. Results in open-loop for realistic experiment setup with two branches. For the small temperature variation measured (around 1°), a phase variation of 30 ps was measured.




Fig. 24. Results in closed-loop for realistic experiment setup with two branches. For the small temperature variation measured (around 1°), a phase variation of only 5 ps was measured.

smaller phase variation over the day when the TCLink was employed.

VI. CONCLUSION

In this article, we have demonstrated a TCLink IP core which is, to the best of our knowledge, the first such implementation that is fully integrated in an FPGA. The core has low-resource usage and for users wishing to optimize occupancy, the control loop can be implemented in software. The proposed core can be used to mitigate long-term phase variations within HEP experiments.

An additional innovation is the protocol-agnostic nature of the core, giving the user the freedom of employing their own protocol. The user also has the freedom of choosing their own bandwidth for the timing compensation. The phase correction over time can be read-out and saved for offline analysis.

It was shown that the TCLink core improves the overall timing stability when temperature variations are present in an experiment.

Also, the results demonstrated that the techniques employed here cannot be used to correct for jitter or phase determinism. These require additional design optimization.

Finally, the design is publicly available on GIT under the Open Hardware License and despite being developed for the LHC experiments, other applications can also benefit from this development.

ACKNOWLEDGMENT

The authors would like to thank their colleague Philippa Hazell from CERN for reviewing the manuscript.

REFERENCES

- [1] S. Baron, T. Mastoridis, J. Troska, and P. Baudrenghien, "Jitter impact on clock distribution in LHC experiments," *J. Instrum.*, vol. 7, no. 12, Dec. 2012, Art. no. C12023.
- [2] C. Allaire, "A high-granularity timing detector in ATLAS: Performance at the HL-LHC," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 924, pp. 355–359, Apr. 2019. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900218306284>
- [3] L. Gray and T. T. de Fatis, "Technical proposal for a MIP timing detector in the CMS experiment phase 2 upgrade," CERN-LHCC, Geneva, Switzerland, Tech. Rep. LHCC-P-009, Nov. 2017. [Online]. Available: <https://cds.cern.ch/record/2296612/files/LHCC-P-009.pdf>
- [4] E. Mendes, S. Baron, C. Soos, and F. Vasey, "A timing, trigger, and control system with picosecond precision based on 10 Gbit/s passive optical networks for high-energy physics," *IEEE Trans. Nucl. Sci.*, vol. 68, no. 4, pp. 447–457, Apr. 2021.
- [5] C. Soós et al., "Versatile link PLUS transceiver development," *J. Instrum.*, vol. 12, no. 3, Mar. 2017, Art. no. C03068, doi: [10.1088/1748-0221/12/03/C03068](https://doi.org/10.1088/1748-0221/12/03/C03068).
- [6] P. Moreira. *LpGBT Specifications*. CERN. Update 18. Accessed: Oct. 2022. [Online]. Available: <https://espace.cern.ch/GBT-Project/LpGBT/Specifications/LpGbtSpecifications.pdf>
- [7] J. Hegeman et al., "First measurements with the CMS DAQ and timing hub prototype-1," *PoS*, vol. 370, p. 111, Apr. 2020. [Online]. Available: <https://pos.sissa.it/370/111>
- [8] E. Mendes, S. Baron, C. Soos, J. Troska, and P. Novellini, "Achieving picosecond-level phase stability in timing distribution systems with Xilinx ultrascale transceivers," *IEEE Trans. Nucl. Sci.*, vol. 67, no. 3, pp. 473–481, Mar. 2020.
- [9] M. Lipiński, T. Wlostowski, J. Serrano, and P. Alvarez, "White rabbit: A PTP application for robust sub-nanosecond synchronization," in *Proc. IEEE Int. Symp. Precis. Clock Synchronization Meas., Control Commun.*, Sep. 2011, pp. 25–30.
- [10] Y. Hu, Y. Wang, J. Kuang, and B. Wu, "A clock distribution and synchronization scheme over optical links for large-scale physics experiments," *IEEE Trans. Nucl. Sci.*, vol. 68, no. 6, pp. 1351–1358, Jun. 2021.
- [11] H.-B. Xie, Y. Li, Q. Shen, S.-K. Liao, and C.-Z. Peng, "A high-precision 2.5-ps RMS time synchronization for multiple high-speed transceivers in FPGA," *IEEE Trans. Nucl. Sci.*, vol. 66, no. 7, pp. 1070–1075, Jul. 2019.
- [12] CERN. *CERN Open Hardware License CERN-OHL-W*. Accessed: Oct. 2022. [Online]. Available: <https://cern-ohl.web.cern.ch>
- [13] E. Brandao De Souza Mendes, S. Baron, and M. Taylor, "TCLink: A timing compensated high-speed optical link for the HL-LHC experiments," in *Proc. Top. Workshop Electron. Part. Phys.*, Mar. 2020, p. 57.
- [14] E. Mendes and S. Baron. (Jun. 2021). GitLab for TCLink project. CERN. [Online]. Available: <http://gitlab.cern.ch/HPTD/tclink>
- [15] P. Moreira, I. Darwazeh, "Digital femtosecond time difference circuit for CERN's timing system," Univ. College London, London, U.K., Aug. 2011. [Online]. Available: www.ee.ucl.ac.uk/lcs/previous/LCS2011/LCS1136.pdf
- [16] Xilinx. (Sep. 2017). *UltraScale Architecture GTY Transceivers User Guide. UG578 (V1.3)*. [Online]. Available: https://www.xilinx.com/support/documentation/user_guides/ug578-ultrascale-gty-transceivers.pdf
- [17] Xilinx. (Aug. 2018). *UltraScale Architecture GTH Transceivers User Guide. UG576 (V1.5.1)*. [Online]. Available: https://www.xilinx.com/support/documentation/user_guides/ug576-ultrascale-gth-transceivers.pdf
- [18] F. M. Gardner, *Phaselock Techniques*, 3rd ed. New York, NY, USA: Wiley, 2005.
- [19] E. Mendes and S. Baron. (Mar. 2020). GitLab for HPTC project. CERN. [Online]. Available: <http://gitlab.cern.ch/HPTD/hptc>
- [20] D. H. Montesinos, S. Baron, N. Guettouche, and J. Mendez, "The versatile link+ demo board (VLDB+)," *J. Instrum.*, vol. 17, no. 3, Mar. 2022, Art. no. C03032, doi: [10.1088/1748-0221/17/03/C03032](https://doi.org/10.1088/1748-0221/17/03/C03032).
- [21] N. Loukas et al., "The CMS barrel calorimeter processor demonstrator (BCPv1) board evaluation," *J. Instrum.*, vol. 17, no. 8, Aug. 2022, Art. no. C08005, doi: [10.1088/1748-0221/17/08/C08005](https://doi.org/10.1088/1748-0221/17/08/C08005).
- [22] V. Amoiridis et al., "CMS phase-2 DAQ and timing hub prototyping results and perspectives," *J. Instrum.*, vol. 17, no. 5, May 2022, Art. no. C05003, doi: [10.1088/1748-0221/17/05/C05003](https://doi.org/10.1088/1748-0221/17/05/C05003).