Published by IOP Publishing for Sissa Medialab

Received: *November 2, 2024*
Accepted: *December 11, 2024*
Published: *January 9, 2025*

Topical Workshop on Electronics for Particle Physics
University of Glasgow, Scotland, U.K.
30 September–4 October 2024

An FPGA-agnostic system for achieving picosecond-level phase determinism in timing distribution links for High Energy Physics experiments

E. Orzes,[a,*] **S. Baron**,[a] **E. Mendes**,[a] **R. Rusack**,[b] **Z. Eberle**,[b] **R. Saradhy**,[b] **Y. Tousi**,[b] **E. Frahm**,[b] **F. Martina**[a] **and P. Vichoudis**[a]

[a] *CERN, European Center for Nuclear Research, Geneva, Switzerland*
[b] *The School of Physics and Astronomy, The University of Minnesota, Minneapolis, Minnesota, U.S.A.*

E-mail: edoardo.orzes@cern.ch

Abstract: Picosecond-level phase determinism in timing distribution systems is a requirement for future detectors in High Energy Physics. FPGA transceivers traditionally used to propagate timing signals generally do not meet this stringent requirement, suffering from phase jumps at startup and from drifts due to temperature variations. While *ad hoc* solutions have been developed using FPGA features to measure and correct for phase shifts, they remain FPGA specific and therefore cannot be generalized to all types of timing links. This paper presents a study using discrete components to monitor and correct for phase drift, allowing for an implementation of a generic deterministic link.

Keywords: Digital electronic circuits; Analogue electronic circuits; Analysis and statistical methods

*Corresponding author.

https://doi.org/10.1088/1748-0221/20/01/C01003

2025 JINST 20 C01003

Contents

2025 JINST 20 C01003

1 Introduction

With the high-luminosity upgrade of the LHC the number of particle collisions per bunch crossing (BX) will increase from the current 60 events per bunch crossing to at least 140 (pile up), which all occur within a $\approx$ 325 ps time interval [1]. Currently the detectors distinguish collision events within a BX with reconstructed collision products. To mitigate the increase in background due to pile up, the detectors at the LHC are installing new timing detectors to measure the time of arrival of a particle with a precision of $\approx$ 30 ps, so that particles of interest from a time window in the BX can be separated from background particles outside of the time window [2]. Motivated by these detector developments and in anticipation of future trends, we are addressing the problem of how to achieve picosecond-level phase stability in the back-end links for timing distribution within the detectors. The front-end links use the radiation tolerant lpGBT [3] ASIC which has a phase stability of 3 ps, while the back-end uses commercial FPGAs. Current FPGAs and their transceivers do not meet the phase stability requirement of <5 ps and *ad hoc* transceiver-specific solutions have been implemented to circumvent this limitation. Here we present a fully-agnostic, discrete-logic system that achieves a fixed-latency, is phase-stable, and compensates for run-time phase drifts and shifts in the FPGA. The temperature and phase stability of each component of the system is characterized separately. To demonstrate this we have assembled the system using discrete dedicated PCBs before the design and fabrication of the final PCB, which is on-going.

2 Phase stability of back-end timing links

The timing distribution system used in the back-end of the LHC relies on a cascade of FPGAs, as in figure 1, connected via data-links that carry the timing information. In this architecture the clock recovered by a receiver is used as a reference for the next transmitter of the cascade. Ideally this allows to have a fixed latency from input to output on the FPGA.

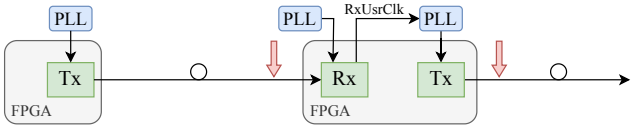

Figure 1. Fixed latency timing link based on clock recovery.

However, different perturbations cause the phase to be unstable over time, regardless of which FPGA type is used.

The effect of these perturbations can be divided into two categories, phase jumps and phase drifts. Phase jumps on the data stream of one bit period, the unit interval (UI), can be caused by resets of the Rx or Tx transceivers, and phase drifts can be caused by temperature variations in either the FPGA or the interconnecting fibers. For the AMD GTYe4 and GTYe5 transceivers, Rx resets have been observed to lead to sub-UI phase jumps of up to 10 ps and 30 ps respectively. Current solutions to mitigate the UI phase jumps are transceiver specific, while solutions compensating sub-UI phase jumps are currently under investigation and not optimal. TCLink [4], developed by the CERN HPTD group, corrects for fiber drift achieving a stability of ≈5 ps, assuming that the link is symmetric. All these solutions are transceiver-specific, and currently there is no solution to the FPGA temperature-induced phase drifts. One of the difficulties is the complex temperature dependence of the dielectric, which, depending on the material, can exhibit hysteresis and non-linearities [5], precluding phase compensation based on temperature monitoring.

For example (figure 2), a new technique [6] to cancel out sub-UI phase jumps on the Rx of AMD transceivers is based on a phase measurement between the data stream and the recovered clock (RxUsrClk). The measurement is performed within the FPGA and is used to shift (with AMD Phase Interpolator – PI) the phase of the transmitted data stream in the next transceiver of the cascade to compensate for the phase jumps. This solution can only be applied at startup, since it requires a clock pattern to be sent on the data stream to perform the phase measurement. It is also limited by temperature variations within the FPGA that distort the phase measurement.

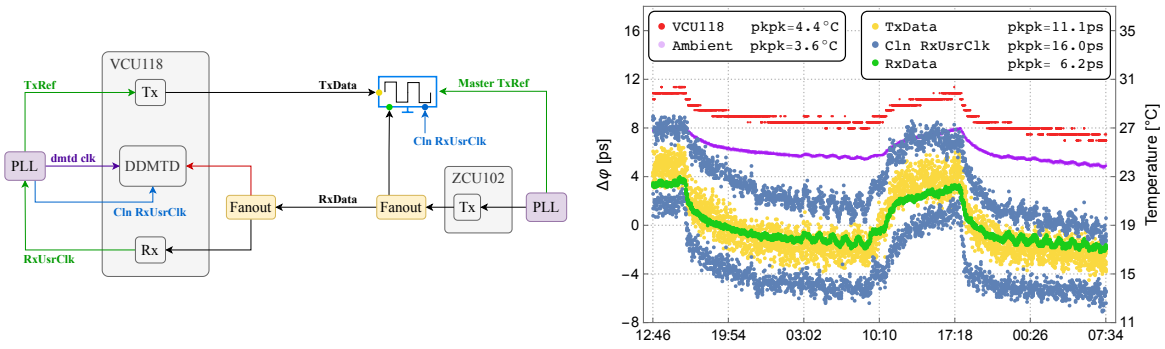

Figure 2. Phase stability measured in Virtex US+ FPGA. RxUsrClk (blue) has a double-peak distribution due to Rx-resets, the DDMTD in the FPGA measures the phase between RxUsrClk and the Rx clock pattern from the fanout (green) and compensates on the Tx data stream (yellow) on the same FPGA.

3 A design for tracking and compensating phase variations

To remove the instabilities of FPGAs, we are proposing a system to implement phase measurements directly on the received and transmitted data links using discrete components, which include a phase shifter used to compensate for phase variations. In this scheme, the input and output data links of an FPGA are doubled using a fast fanout and the embedded clock is extracted from one branch of each fanout using CDRs, allowing for the phase between the two clocks to be measured inside the FPGA. Finally, a coarse and a fine phase shifter are implemented on the outgoing serial link to compensate

2025 JINST 20 C01003

for the phase variations measured in the FPGA, at startup and during operation respectively. The tool used to perform phase measurements within the FPGA is a Digital Dual Mixer Time Difference (DDMTD) [7]. It samples the synchronous input clocks with an offset clock which has a frequency slightly lower than the one of the inputs. This causes a time amplification of the input clocks, allowing for counting between the edges of two channels. Ideally, the relative phase can be obtained with sub-picosecond accuracy. The offset frequency ν_{dmtd} is a function of the input frequency ν_i and N (eq. (3.1)). For the following characterizations we use $\nu_i = 160\,\text{MHz}$ and $N = 100\,\text{k}$, giving a theoretical resolution $\delta_{dmtd} = 62.5\,\text{fs}$ (eq. (3.2)) [7].

$$\nu_{dmtd} = \frac{N}{N+1}\,\nu_i \tag{3.1}$$

$$\delta_{dmtd} = \frac{1}{N+1}\,T_{dmtd} \tag{3.2}$$

However, temperature variations on the FPGA and on the evaluation board can alter the phase measurement. We there have used a DDMTD with the timing critical front-end (i.e. the D-flip-flop samplers) in the PCB. This has less noise than the same circuit implemented in firmware on the FPGA, and also allows for symmetrical routing, resulting insensitive to temperature.

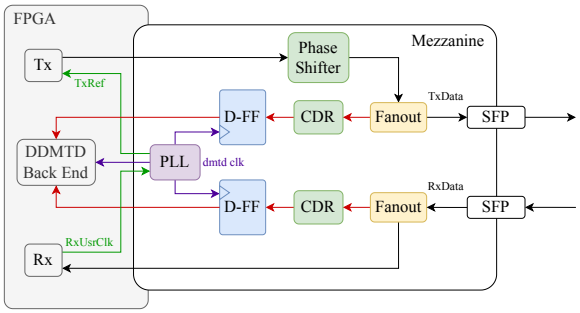

Figure 3. Proposed mezzanine concept for data stream phase monitoring and compensation at run-time.

The circuit will be implemented on an FMC mezzanine (figure 3) equipped with one SFP per channel. In both channels a fanout buffer splits the data stream and feeds it to a CDR recovering the embedded clock, which is then fed to the flip-flop in the front-end stage of the DDMTD. The flip-flop period-amplified clocks are fed to the FPGA for the DDMTD processing. The measured phase is then used for controlling the phase shifter, which is placed on the Tx data channel.

A Skyworks PLL is used as a jitter attenuator, it locks to the transceiver recovered clock, RxUsrClk, and its output is used as a reference for the transmitter. The PLL is in Zero Delay Mode [8], allowing for fixed latency among resets. This PLL is also used to generate the DDMTD offset clock.

Use of the CDR chips allows the phase to be measured and corrected during operation and not just at startup as in the original concept. It also enables compensation of both temperature induced phase drifts and for the UI and sub-UI phase jumps at startup.

4 Components characterization

The success of this concept relies on the phase stability of each of the discrete components, and we therefore evaluated them separately. The University of Minnesota designed a PCB [9] which

2025 JINST 20 C01003

implements the timing critical front-end of a DDMTD; the board provides a PLL for the offset clock and two flip-flops for the clocks sampling. The DDMTD back-end is implemented in an FPGA.

The first two tests compare the performances of a DDMTD fully implemented in the FPGA (including the input flip-flops) with the discrete (PCB) one. The two DDMTDs measure the phase between two synchronous clocks at 160 MHz.

In the temperature stability test (figure 4) the performance of the discrete DDMTD (in green) has a stability of 0.26 ps RMS and an uncertainty of ≈ 0.7 ps, without any evident temperature variations. In the FPGA version (implemented on a Zynq ZCU102) the phase measurement had a temperature dependency that was as much as 0.6 ps/degC. This was attributed to differences in the input clock routings. In the linearity test (figure 5) a phase shift is applied on one of the two input clocks by transmitting a clock pattern from an AMD GTHe4 transceiver, using its phase interpolator. In all 2730 steps of 6.1 ps were used to cover the full period of the clock.

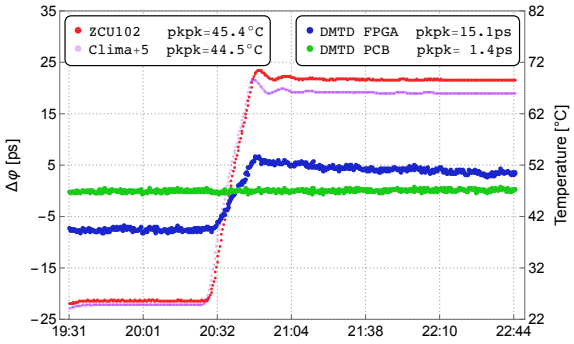

Figure 4. Temperature stability test. DDMTD comparison on PCB and FPGA.

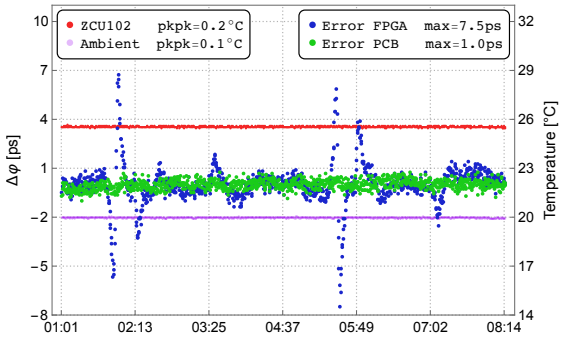

Figure 5. Linearity test. DDMTD comparison on PCB and FPGA, DDMTD Vs oscilloscope residual.

The linearity test was repeated with different DDMTD implementations and routings: eleven times on a Virtex VCU118 and four on a Zynq ZCU102. Using the Vivado device viewer it was determined that the routing inside the FPGA fabric could lead to a deterministic phase distortion. This occurred when one of the two DDMTD input clocks is routed close to another clock of the same, or a multiple of the frequency. If the two clocks are in phase or in phase opposition, a noise injection from one edge into the other can alter the clocks phase. In figure 5 the phase of the distortion is shown, leading to an uncertainty of up to 7.5 ps in the two peaks, corresponding to the in-phase and the phase-opposition edges. We found that using PBlocks for manual placement this crosstalk could be eliminated in half of the cases. However, it is not always possible to obtain a proper routing, for example, locking the DDMTD sampler in the logic block of the input port is not possible for the transceiver reference input buffers.

The chosen CDR is the ADN2817, it is able to lock to any data stream up to 2.7 Gbps. We have designed a PCB evaluation board and tested its stability. Two CDR chips were tested and revealed to be fixed in latency between relocking cycles, with a stability better than 0.2 ps RMS and 0.9 ps peak-peak. The temperature sensitivity is about 3 ps/degC for both chips. Given the symmetry of the final concept (one CDR on each path), the high temperature sensitivity is not an issue.

The phase shifter (DCPS3 [10]) also demonstrated excellent performance; it has negligible temperature dependency and its step error is below the noise-floor of 0.1 ps RMS. This ASIC, used for fine phase shifting, allows for 300 fs steps with a range of 20 ps, while the AMD PI has a step of 1.6 ps, an error of up to ± 2 ps and no range limitation.

2025 JINST 20 C01003

The following test (figure 6) exploits all previously tested devices in proof of concept made of evaluation boards. A Zynq ZCU102 is used as a Tx master and a Virtex VCU118 is receiving the data stream (RxData), recovering the clock and transmitting a second data stream (TxData) (figure 1). The DDMTD phase measurement is performed on the two CDR recovered clocks and the phase-shift is applied by the PI of the VCU118 transmitter, since the DCPS3 board would require a clock as input. This setup is exposed to temperature variations using a climate chamber. The test assess the quality of the system, designed to compensate for phase drifts between the input and output data streams of the Virtex FPGA. The system measures and compensates phase between data links carrying a PRBS pattern. Its stability is evaluated by an oscilloscope which requires clock signals. The PRBS pattern is then briefly switched to a clock in order to perform this measurement.

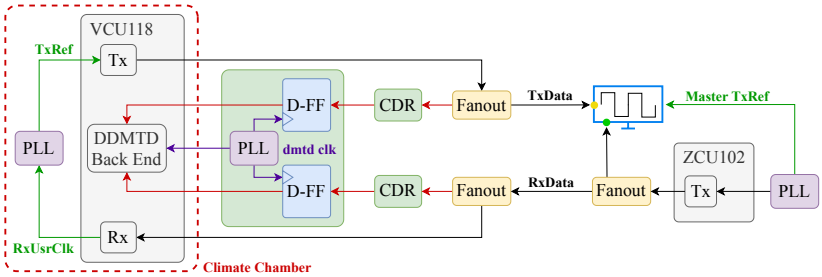

Figure 6. Test-bench for the temperature stability test.

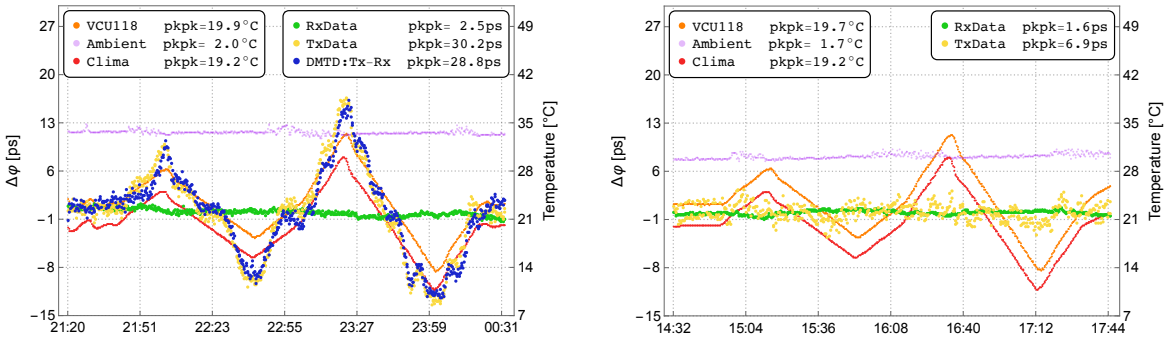

Figure 7. Temperature stability test measuring input-output phase outside the FPGA and compensating.

In figure 7 (left) the DDMTD measures the variations of TxData versus RxData with an error up to 4.9 ps and 1.5 ps RMS. In figure 7 (right) the phase shift is performed and the TxData stability reaches 1.2 ps RMS. When compensation is not activated the temperature sensitivity of the FPGA is 1.5 ps/degC. This is completely canceled out when activating the compensation.

5 Conclusions

HL-LHC upgrade requires a ps-level phase stability in time distribution system which is not naturally provided by current FPGA implementations. Temperature variations within the boards and sub-UI phase jumps remain the greatest obstacles to solve. This proof-of-concept presented in this paper proposes a solution for a transceiver-agnostic and fixed-latency timing link. The current limitations are the phase shifter and CDR bandwidths of 2.7 Gbps and as well as the bulkiness represented by the

2025 JINST 20 C01003

number of discrete components required per link. Each component has been characterized for phase stability and has demonstrated excellent performances, revealing to be compliant for the concept. The design will start this year and the results are expected for the first half of 2025.

Acknowledgments

Authors from the University of Minnesota are grateful for the support of the US DOE Office of High Energy Physics for their support under the awards DE-SC0020185 and DE-SC0011845.

References

[1] I. Zurbano Fernandez et al., *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*, CERN-2020-010 (2020) [DOI:10.23731/CYRM-2020-0010].

[2] A. Colaleo et al., *The 2021 ECFA Detector Research and Development Roadmap*, CERN-ESU-017 (2021) [DOI:10.17181/CERN.XDPL.W2EX].

[3] lpGBT Design Team, *lpGBT Documentation*, CERN (2024) [https://cds.cern.ch/record/2809058].

[4] E. Mendes et al., *TCLink: A Fully Integrated Open Core for Timing Compensation in FPGA-Based High-Speed Links*, *IEEE Trans. Nucl. Sci.* **70** (2023) 156.

[5] K. Czuba and D. Sikora, *Temperature Stability of Coaxial Cables*, *Acta Phys. Polon. A* **119** (2011) 553.

[6] E. Orzes, *Picosecond-Level Phase Stability with Xilinx Transceivers for Timing Distribution Systems in High Energy Physics Experiments*, M.Sc. Thesis, Dept. Information Eng., University of Pisa, Pisa, Italy (2023) [https://etd.adm.unipi.it/t/etd-10102023-180452].

[7] P. Moreira et al., *Digital dual mixer time difference for sub-nanosecond time synchronization in Ethernet*, in the proceedings of the *2010 IEEE International Frequency Control Symposium*, Newport Beach, CA, U.S.A. (2010), p. 449–453 [DOI:10.1109/FREQ.2010.5556289].

[8] Skyworks, *Implementing Zero Delay Mode Using the Si5340/41/42/44/45/80 – AN947*, Skyworks (2022).

[9] R. Saradhy, E. Frahm, E.B.S. Mendes and R. Rusack, *A sub-picosecond digital clock monitoring system*, 2023 *JINST* **18** T01003 [arXiv:2210.05764].

[10] D. Dehmeshki et al., *A sub-picosecond digitally-controlled phase delay*, 2022 *JINST* **17** C03014 [arXiv:2111.13548].

2025 JINST 20 C01003