Advancing Charged Particles Cancer Treatment by Real-Time Measurement on FPGA

Arash Amini Bardpareh[1], Eleonora Vacca[1], Corrado De Sio[1], Sarah Azimi[1], Luca Sterpone[1]
Elisa Fiorina[2], Emanuele Maria Data[2], Felix Mas Millian[2]
[1]*Dipartimento di Automatica e Informatica, Politecnico di Torino,* Turin, Italy
[2]*Istituto Nazionale di Fisica Nucleare,* Turin, Italy

***Abstract*—Charged particle therapy offers significant advantages in cancer treatment due to its ability to deliver highly localized doses to tumors. However, the effectiveness and safety of this treatment rely critically on precise control and real-time monitoring of beam parameters. One of the key parameters is the particle Time-of-Flight (ToF), which directly informs about the energy and range of the particles within the patient. Accurate ToF measurement requires the integration of particle detectors with high-resolution Time-to-Digital Converters (TDCs) with picosecond-level precision. In this work, we present a compact FPGA-based TDC implementation, designed for seamless integration with the ESA-ABACUS oncology particle sensor. By combining calibration with place & route optimization, our TDC achieves 6.2 ps standard deviation, outperforming the current state-of-the-art. The in-laboratory characterization was successfully validated through beam tests conducted at sub-clinical dose rates.**

Keywords—FPGA, Particle Therapy, Time-to-Digital Converter.

I. Introduction

Charged Particle Therapy (CPT) is an advanced form of cancer treatment that uses beams of charged particles, typically protons or carbon ions, to target tumors with high precision [1]. Unlike conventional radiotherapy, CPT allows for more localized energy deposition, reducing the risk of damage to surrounding healthy tissues. This precision is primarily due to the Bragg peak phenomenon, where charged particles deposit most of their energy at a specific depth, which can be finely tuned based on tumor location and patient-specific anatomy.

To ensure treatment efficacy, it is essential to monitor the beam parameters and measure the particle range in real time [2]. This requires accurately tracking both primary particles (those directly accelerated and delivered to the tumor) and secondary particles (produced from interactions within the patient's body) [3]. The Time-of-Flight (ToF) [4] and energy of these particles provide critical information for beam monitoring, range verification, and adaptive treatment planning, especially as the tumor mass may shrink or shift during therapy. Accurate ToF measurements are vital to avoid irradiating healthy tissue and to allow clinicians to dynamically adjust the beam per patient and treatment session [5]. To achieve this, a complete and highly precise Data Acquisition (DAQ) system is required. This system typically relies on dedicated sensors [6][7][8] to detect particle events. A key role in the DAQ chain is played by the Time-to-Digital Converter (TDC), which is responsible for converting the fine time intervals between consecutive particle detections into digital values suitable for subsequent processing and analysis.

Traditionally, TDCs are implemented using Application-Specific Integrated Circuits (ASICs) due to their high accuracy and performance. However, ASIC-based TDCs suffer from high costs, long development cycles, and a lack of flexibility needed for evolving experimental setups. To address these limitations, Field-Programmable Gate Arrays (FPGAs) have emerged as a viable alternative for implementing TDC. Thanks to continuous technology scaling, FPGAs now offer higher performance, lower power consumption, and the ability to integrate close to the computational core responsible for beam control algorithms. Their reconfigurability and fast prototyping make them well-suited for medical applications requiring rapid adaptation and innovation. Nevertheless, implementing high-resolution TDCs on FPGAs remains challenging since they are not inherently optimized for fine-grained timing measurement. Issues such as variability in logic cell delays, routing path non-uniformities, jitter, and metastability affect the achievable resolution and reliability. These factors, combined with manufacturing process variations, necessitate careful architectural design and rigorous calibration to achieve performance comparable to ASICs.

To overcome these challenges, this work introduces a custom FPGA-based TDC architecture specifically designed to meet the picosecond-level resolution demands of CPT. Implemented on an AMD Kintex UltraScale FPGA, manufactured using a 20 nm CMOS technology node, the proposed TDC achieves a time resolution of 6 ps and a precision of 30 ps. Unlike previous approaches that rely only on extensive and time-consuming calibration procedures, our design reaches this performance by performing a meticulous optimization of the place-and-route strategy dedicated to reducing non-uniformity in the on-chip implementation of the circuit netlist. Developed within the framework of the Hybrid ONline tEchnologY for particle therapy (HONEY) project, the TDC is tailored for seamless integration with the ESA-ABACUS particle detection system [7][8]. A preliminary experimental beam test campaign using carbon ions successfully validated the effective integration of the FPGA-based TDC with the detector, as well as its capability to accurately measure particle ToF.

II. State of the Art

TDCs are essential components in systems that require precise time measurements, such as CPT, LiDAR, medical imaging, and quantum communication. Over the years, significant efforts have been devoted to developing high-resolution and efficient TDCs, especially in ASIC form, to reach fine-grain resolutions. Notable examples include CERN's picoTDC on 65 nm CMOS technology, achieving 3.7 ps RMS running at 320 MHz, and the solution proposed by [9] on 0.35 μm CMOS, based on a Self-Timed Ring Oscillator, which achieves 72.5 ps resolution. While ASIC-based TDCs offer excellent precision, TDCs on FPGAs have recently grown in popularity. Indeed, the flexibility and reconfigurability of FPGAs make them ideal for building application-specific and cost-effective TDC solutions, especially in research environments or systems that require real-time data processing close to the sensing front-end. The most widely adopted TDC architecture in FPGA implementations is the Tapped Delay Line (TDL), shown in

Fig. 1 [10][11][12]. The delay line consists of a series of logic elements that introduce incremental propagation delays. When a hit signal propagates through this chain and is sampled by a clock edge, it generates a thermometer code representing how far the signal traveled, i.e., the fine time interval between the clock and event. However, implementing a TDC on an FPGA is not trivial. Unlike ASICs, FPGAs are inherently optimized for synchronous sequential logic, whereas TDL-based TDCs rely on purely combinational and asynchronous signal propagation. Furthermore, the logic elements used to construct the delay line do not provide uniform delays, resulting in non-linearity in the TDC output. This variability arises from process variation, intra-chip inconsistencies, and routing delays, which can all significantly degrade TDC resolution and accuracy.

To tackle the known limitations of TDC resolution and linearity, several methods have been explored in the literature. Dual-sampling techniques [13] improve resolution by capturing two signals propagating across the delay line, while wave union methods [14] enhance linearity and precision by reusing signal transitions within the delay line. What remains largely unexplored is the root cause of these nonlinearities at the place-and-route level, and this is where our work makes its key contribution. By carefully controlling the placement and routing of a TDL-TDC implemented on a Kintex UltraScale FPGA, we demonstrate that it is possible to achieve a measurement standard deviation (σ) of 30 ps without any calibration. This can be further improved to 6.2 ps σ and an RMS of 9.8 ps with calibration.

Notably, our results outperform prior state-of-the-art implementations such as [13], which report a 15 ps RMS and 14 ps standard deviation despite requiring a dedicated calibration module composed of ring oscillators, FSM, and dual tapped delay lines (one for measurement, one for standby). We also surpass the 17 ps post-calibration performance reported in [16]. Importantly, our results are achieved without area overhead or relying on a complex calibration procedure or infrastructure, but through fine-grained place-and-route design decisions. This makes our method both effective and resource-efficient, highlighting the latent potential of layout-aware optimization in TDC design.

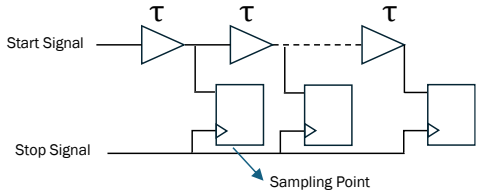

Fig 1. The tapped delay line architecture.

III. The Proposed TDC Architecture

As previously mentioned, most TDC architectures are based on the TDL structure. However, TDL alone can only measure a limited range of time intervals, which may not be sufficient for the target application. Increasing the measurable range by increasing the delay line length, while feasible in ASIC implementation, where layout and routing can be tightly controlled, it cannot be adopted as a solution on FPGA-based TDC implementation. FPGA-based TDLs are typically constructed using carry logic as delay elements, which offer the finest and most consistent delay granularity available within the device. Considering that the FPGA area is divided into clock regions, the amount of available carry logic blocks belonging to the same clock distribution area is limited (100 CARRY4 logics for Artix-7 family, 60 CARRY8 logics for Ultrascale and Ultrascale+ family). Extending the delay line across multiple clock regions typically increases non-linearity issues. This is due to both clock jitter and the use of specialized routing paths required to cross clock region boundaries, which are generally longer and less uniform than intra-region connections. These variations further amplify delay inconsistencies already present due to process variation. As a result, to achieve a wider measurement range, the TDL is commonly combined with a synchronous counter.

Since the target application requires that the TDC's ultimate goal is to measure the ToF between primary and secondary particles in CPT, our proposed TDC implementation consists of two independent delay lines coupled with a synchronous counter. The proposed TDC behavior can be summarized by Eq. (1) and Fig. 2, where the ToF is obtained as the sum of three measurement contributions.

$$T = T_1 + T_2 - T_3 \tag{1}$$

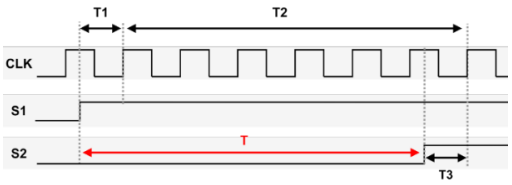

Fig 2. The measurement principle of the proposed TDC architecture.

When the primary particle detector registers an event (S1), this signal enters the first delay line (DL1) and simultaneously enables a synchronous counter. DL1 remains active until the next rising edge of the system clock, at which point it captures the fine time measurement T1. Meanwhile, the counter continues counting clock cycles until a second event (S2) is generated by the secondary particle detector. Upon the arrival of S2, the counter is stopped, and its value represents the coarse time contribution T2. At the same moment, S2 enables the second delay line (DL2), which captures the fine time T3 by stopping at the first rising edge of the system clock following S2. Since both S1 and S2 are asynchronous with respect to the system clock, the delay lines are crucial for achieving high-resolution timing in the picosecond range by measuring fractions of a clock cycle. In contrast, the synchronous counter extends the overall measurement range by counting full clock periods between S1 and S2, typically spanning several nanoseconds.

The output of a delay line is a thermometer code, i.e., a binary sequence of 1s followed by 0s, indicating how far a signal has propagated through the delay elements. In an ideal scenario, each delay element contributes an equal delay, and the signal progresses sequentially through the chain, producing an output like "111100...0" if it reaches the fourth element before being sampled. However, in FPGA-based implementations, this ideal behavior is disrupted by process variations and non-uniform routing delays. Each delay element may introduce a slightly different delay, and the routing paths between elements are not consistent. As a result, the output often deviates from the ideal thermometer code, producing a pseudo-thermometer code affected by bubble errors, unexpected 0s within the sequence of 1s. For instance, instead of "111100...0", the output may be "11011...0", indicating that the signal did not propagate uniformly. These errors can be caused by several factors,

including temperature variations, metastability, and inherent fabrication inconsistencies. While previous approaches have attempted to mitigate this issue through various sampling techniques and time-consuming online or offline calibration, we observed that improving the physical layout, specifically by maximizing symmetry in the design, can significantly reduce the disturbance of bubble errors. Therefore, it is essential to carefully manage the placement and routing of the two delay lines in order to achieve the most equivalent and symmetrical behavior possible. This aspect will be further detailed in the next section. Then, each delay line output is followed by an encoder whose purpose is to convert the DL outputs into the proper binary value, taking into account the offline calibration performed to reduce the impact of the remaining errors. The encoded delay lines outputs are then merged with the coarse-grain measurement of the counter through a simple processing unit that multiplies the counter value by the system clock period, and computes Eq. (1), finally producing the time measurement.

IV. Place-and-Route Design Optimization

In our implementation, we employ two separate delay lines, DL1 and DL2, to independently process signals from the primary and secondary particle sensors. For the system to function correctly, it is critical that these two delay lines exhibit highly similar behavior, a requirement that is far from trivial. Although DL1 and DL2 are identically described at the RTL level, differences arise during physical implementation. If no specific placement and routing constraints are applied, the FPGA toolchain may adopt to significantly different placement strategies and use a different routing resource type. This can introduce substantial mismatches in terms of linearity (i.e., weight of delay elements), delay offset, and measurement range. As a result, the system becomes extremely difficult to calibrate effectively. Figure 3a shows the post-implementation layout of the two delay lines, generated without applying any design constraints. As illustrated, the routing paths from the input pins to DL's entry point (S1 to DL1 in light orange and S2 to DL2 in light blue) differ significantly. These differences introduce a measurement offset between the two delay lines, which is further influenced by their physical placement on the FPGA. Additionally, the counter begins counting upon the arrival of S1 and stops when S2 is sampled. However, any mismatch in routing delays from the S1 and S2 input pins to the counter's input ports introduces an additional, hard-to-predict offset. Moreover, inconsistencies in routing length between the S1 path to DL1 and the S1 path to the counter (and likewise for S2) may cause the counter to register an incorrect number of clock cycles, either counting many cycles more or too few. As a result, leaving the TDC layout entirely to the toolchain's discretion leads to unpredictable behavior that can deviate significantly from RTL simulations. Therefore, applying physical design constraints becomes essential to achieving reliable and consistent performance. The optimal layout is typically found through a trial-and-error process, requiring full implementation up to bitstream generation and real-world testing, as only in-field measurements reveal the true behavior of the delay line propagation and associated errors. The layout solution must account not only for non-uniformities within each individual delay line but also for discrepancies between the two delay lines. To address this, a symmetric layout strategy was adopted, where the delay lines are placed in separate clock regions and designed symmetrically with respect to S1 and S2 I/O pins, while in the unconstrained design, the delay lines were placed in the same clock region with highly different routing paths for the two signals. While the placement customization is easy, the routing paths optimization is far from trivial. Indeed, while commercial toolchains provide a wizard to customize the routing path connecting two logic cells, the delay associated with a routing resource with respect to another is vendor-specific information not available. In Fig. 3b, it is possible to appreciate the highly symmetrical design achieved from the device's post-implementation view, while in Fig. 4, the Input/Output characteristics of the unconstrained design are compared to the proposed symmetric solution.

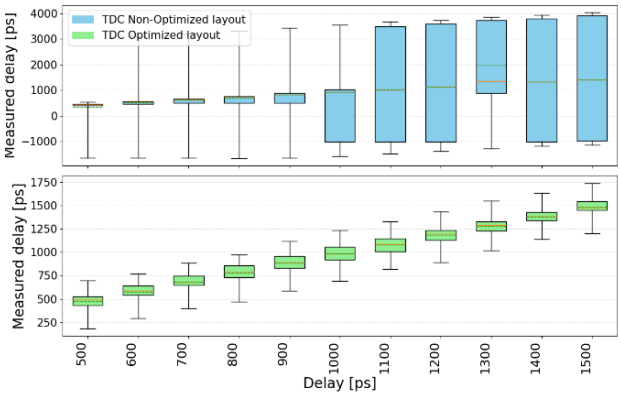

Fig 4. Delay line Input-Output characteristics with and without layout optimization.

Even after determining the optimal layout configuration, calibration is still necessary because non-linearity caused by process variations cannot be eliminated through place-and-route adjustments alone. The calibration procedure involved applying two input signals to the delay lines with precisely controlled time differences, increasing in 10-ps steps. This allowed for accurate characterization of each bin's width. In Fig. 5, the Input/Output characteristics are compared before and after calibration.

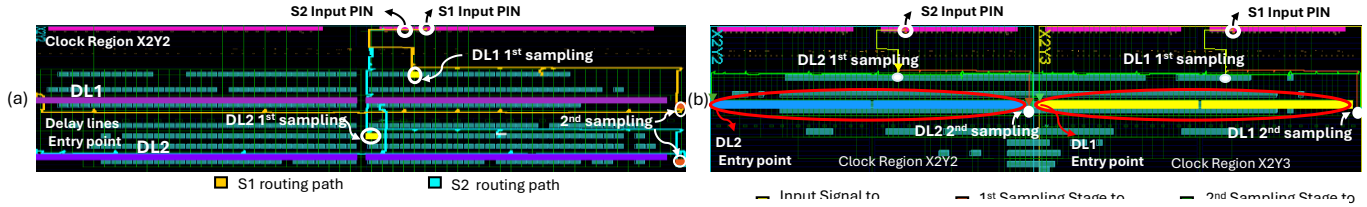

Fig 3. Proposed TDC post implementation. In (a), the plain version, while in (b), the constrained layout.

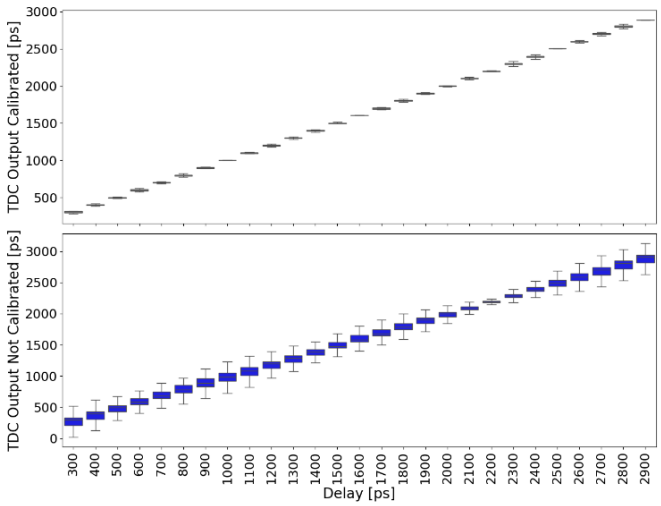

Fig. 5. Delay line calibration process output.

V. Experimental Validation

The proposed TDC is implemented on the Xilinx KCU105 development board, embedding a 20nm CMOS Ultrascale FPGA. The proposed TDC operates at 400 MHz, post-implementation details are provided in Table 1. A key design requirement is the interfacing of the TDC with the ESA-ABACUS [7][8] board, which employs Low Gain Avalanche Detectors (LGADs) for high-energy particle detection and outputs digital signals in the LVDS25 standard. However, most modern FPGAs do not support LVDS25, featuring only a limited number of I/O pins using LVDS18. Consequently, an external conversion board is required to translate LVDS25 signals to LVDS18 for compatibility with the FPGA. Furthermore, since only specific I/O pins can be configured for LVDS18, their selection significantly impacts the place-and-route phase, leading to the final layout proposed in Section IV.

Table 1. Post-implementation details of the proposed TDC

Resource	Available [#]	Utilization [#]	Utilization [%]
LUT	242,400	1459	0.6
REG	484,800	3292	0.68
CARRY8	30,300	285	0.94
BRAM	600	1	0.17
DSP	1920	1	0.05

Prior to calibration, the TDC exhibits an output σ of 30 ps, which is already relatively low given the inherent nonlinearity introduced by the FPGA. Following calibration, the standard deviation is reduced by a factor of 5. Table 2 presents a comparative analysis against state-of-the-art solutions in terms of RMS and σ, where data availability permits. Notably, through the combined effect of layout optimization and calibration, our approach outperforms existing methods, achieving 9.7 ps RMS and 6.2 ps σ.

Table 2. Comparison with SOTA

	RMS [ps]	Std. [ps]	Tech. Node [nm]
[11]	12.0	-	28
[15]	68.4	-	28
[13]	15.7	14.0	28
[16]	-	17.0	20
This work	9.7	6.2	20

A. Experimental Results

The functionality of the proposed TDC was evaluated within the framework of the Honey project. An initial experimental campaign was conducted at CNAO (National Center for Oncological Hadrontherapy) in Pavia, Italy, where the TDC was used in a carbon ion beam experiment.

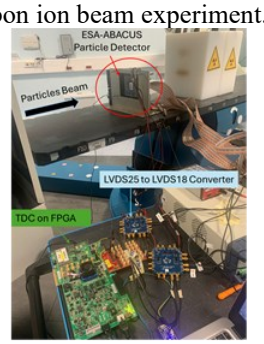

Fig 6. CNAO Radiation Test Setup.

Fig. 6 shows the test setup. The ESA-ABACUS particle detector is positioned in front of the beam. In this validation experiment, a single strip (i.e., one channel) of the detector was connected to the TDC. In this configuration, the TDC measurements are referenced to an external trigger signal provided by a pulse generator. Consequently, the TDC captures the time interval between the rising edge of the trigger and the arrival of particles at the detector, spilled as bursts. The experimental campaign included tests at varying fractions of the clinical beam rate (20%, 50%, and 100%) and across a range of particle energies from 115 MeV to 398 MeV. Experimental data collected were compared against the measurements obtained, under identical experimental conditions, using an ASIC TDC, the PicoTDC [9]. The PicoTDC is a well-established and validated system previously adopted in the framework of the HONEY project. The experiment showed comparable statistics in terms of the number of detected events and consistent inter-arrival time distribution profiles between the two systems, with stable and reliable operation at 20% of the clinical rate and energies up to 240 MeV. However, as the particle flux increased, either due to higher beam rates or reduced interarrival times, the proposed TDC signal quality deteriorated. This degradation was primarily attributed to pulse pile-up and increased noise, which affects the accuracy of time discrimination.

VI. Conclusions

We proposed an FPGA-based TDC for real-time particle ToF measurements, providing implementation details. A preliminary beam test experiment validated its performance at a subclinical rate. Compared to prior studies that focused solely on design or simulation, this work presented both the development and a preliminary beamline validation of the TDC. Even if working at a subclinical rate, such field testing is critical, as laboratory-generated signals represent idealized scenarios that may fail to account for the full range of distortions and noise associated with high-energy particle interactions in realistic conditions. These findings provide valuable insights for future optimization of the system.

Acknowledgment

This research is part of HONEY (Hybrid ONline technology for particle therapy), a PRIN (Progetti di Ricerca di Interesse Nazionale) project funded by the Italian Ministry of University and Research (MIUR), CUP number I53D23001280006 . The project is conducted in collaboration with Politecnico di Torino, INFN, and the University of Torino, with valuable support from CNAO.

REFERENCES

[1] A. R. Smith, "Particle Accelerators for Radiotherapy - A Review," in IEEE Transactions on Nuclear Science, vol. 28, no. 2, pp. 1875-1879, April 1981, doi: 10.1109/TNS.1981.4331545.

[2] L. Gasparini et al., "MONDO: A neutron tracker for particle therapy secondary emission measurements," 2016 IEEE Nuclear Science Symposium, Medical Imaging Conference and Room-Temperature Semiconductor Detector Workshop (NSS/MIC/RTSD), Strasbourg, France, 2016, pp. 1-3, doi: 10.1109/NSSMIC.2016.8069401.

[3] Newhauser, W., Durante, M. Assessing the risk of second malignancies after modern radiotherapy. Nat Rev Cancer 11, 438–448 (2011). https://doi.org/10.1038/nrc3069.

[4] A. El-Hadbi et al., "Precise Time-of-Flight Measurements Using a Novel High-Resolution Time-to-Digital Converter Architecture," 2024 European Frequency and Time Forum (EFTF), Neuchâtel, Switzerland, 2024, pp. 24-27, doi: 10.1109/EFTF61992.2024.10722727.

[5] K. Lee *et al*., "34.1 A 64×64 Implantable Real-Time Single-Charged-Particle Radiation Detector for Cancer Therapy," *2020 IEEE International Solid-State Circuits Conference - (ISSCC)*, San Francisco, CA, USA, 2020, pp. 506-508, doi: 10.1109/ISSCC19947.2020.9063125.

[6] J. Jakubek *et al*., "High Resolution Particle Tracker for In-vivo Carbon-ion Radiotherapy Monitoring," *2024 IEEE Nuclear Science Symposium (NSS), Medical Imaging Conference (MIC) and Room Temperature Semiconductor Detector Conference (RTSD)*, Tampa, FL, USA, 2024, pp. 1-1, doi: 10.1109/NSS/MIC/RTSD57108.2024.10657006.

[7] E.M. Data, et al," A novel detector for 4D tracking in particle therapy, Nuclear Instruments and Methods" in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment.

[8] A. Vignati, et al., Innovative thin silicon detectors for monitoring of therapeutic proton beams: preliminary beam tests, J. Instrum. 12 (12) (2017) C12056.

[9] S. Altruda, et al., PicoTDC: a flexible 64 channel TDC with picosecond resolution, J. Instrum. 18 (07) (2023) P07012.

[10] C. Liu, Y. et al, "A 3.9 ps RMS resolution time-to-digital converter using dual-sampling method on Kintex UltraScale FPGA," 2016 IEEE-NPSS Real Time Conference (RT), Padua, Italy, 2016, pp. 1-3.

[11] F. Garzetti et al., "Time-to-Digital Converter IP-Core for FPGA at State of the Art," in IEEE Access, vol. 9, pp. 85515-85528, 2021.

[12] M. Fishburn et al.,"A 19.6 ps, FPGA-Based TDC With Multiple Channels for Open Source Applications," in IEEE Transactions on Nuclear Science, vol. 60, no. 3, pp. 2203-2208, June 2013.

[13] Z. Jia, Y. Wang, J. Ding, Q. Xu and Y. Zhu, "An FPGA-Based Time-to-Digital Converter With Online Dual-Chain Calibration," in IEEE Transactions on Instrumentation and Measurement, vol. 74, pp. 1-8, 2025, Art no. 2008508, doi: 10.1109/TIM.2025.356139

[14] J. Qi, Z. Deng, H. Gong and Y. Liu, "A 20ps resolution wave union FPGA TDC with on-chip real time correction," IEEE Nuclear Science Symposium & Medical Imaging Conference, Knoxville, TN, USA, 2010, pp. 396-399, doi: 10.1109/NSSMIC.2010.5873788.

[15] N. Lusardi, F. Garzetti, A. Costa, E. Ronconi and A. Geraci, "From Multiphase to Novel Single-Phase Multichannel Shift-Clock Fast Counter Time-to-Digital Converter," in *IEEE Transactions on Industrial Electronics*, vol. 71, no. 8, pp. 9886-9894, Aug. 2024, doi: 10.1109/TIE.2023.3322007.

[16] M. Makhlouf et al, "Photon-Counting Optical Time Domain Reflectometer Using Superconducting Nanowire Single Photon Detector Based on Time to Digital Converter," *2025 7th International Youth Conference on Radio Electronics, Electrical and Power Engineering (REEPE)*, Moscow, Russian Federation, 2025, pp. 1-5, doi: 10.1109/REEPE63962.2025.10971062.