





Optimization of the Upgraded Timing Distribution System of the LHCb experiment at CERN

Mauricio Feo , Federico Alessio , Frédéric Hachon, Jean-Pierre Cachemiche , Paolo Durante  and Guillaume Vouters

Abstract—The Timing Distribution System of the LHCb Experiment at CERN employs FPGA firmware cores to ensure synchronicity across all Front-End and Back-End modules of the Detector Readout System. The control signals are centrally generated in an FPGA-based card, the Readout Supervisor (SODIN), which receives the LHC synchronous clock, and packs timing and control signals into data words (TFC data), distributing it downstream with fixed and deterministic latency. The TFC system is divided into two levels of hierarchy. SODIN transmits the TFC data to FPGA-based Control Cards, which recovers the clock and control signals and forwards them to the Readout Cards and the Front-End modules. While the clock is recovered at the Control and Readout Cards with a phase uncertainty of 76ps, at the Front-End modules the phase is recovered within a window of 4ns, causing a time-alignment shift of up to 4ns after each clock loss. The reason is that the FPGA internal PLL used for generating the reference clock for the transceivers and the external PLL used for jitter-cleaning don't keep the same input-to-output phase after a loss of lock. This paper focuses on the components responsible for timing distribution between the Control Cards and the Front-End modules, achieved through individual bidirectional control links implemented using the GBT-FPGA firmware core. Different firmware and clock routing alternatives were studied for such a link, reducing the window from 4ns to the order of tenths of picoseconds. This paper describes such implementations and the results obtained with each of them.

Index Terms—Data acquisition, Real-time systems, Field programmable gate arrays, FPGA, Timing distribution

I. INTRODUCTION

THE Large Hadron Collider Beauty experiment (LHCb) [1] underwent a major upgrade from 2019 to 2022, requiring a complete redesign of the LHCb Online system [2], which integrates all subsystems responsible for control and data acquisition.

Within the Online system, the Timing and Fast Control system (TFC) [3] guarantees the synchronicity of all different elements of the detector. It is the subsystem responsible for distributing the 40 MHz LHC bunch clock with a fixed phase, centrally generating and distributing real-time control signals with fixed and deterministic latency with respect to the LHC bunch clock, and monitoring the system in real time. In

other words, the TFC system is responsible for the real-time operations of the LHCb Online system. These functionalities are implemented in FPGA firmware on the back-end cards.

As the back-end electronics, the Online system utilizes an FPGA-based card named PCIe40 [4], which interfaces with the detector front-end cards, implements data acquisition and control functionalities, and can assume different specific roles depending on the FPGA firmware.

The initial implementation of the TFC system for the LHC Run 3 upgrade provided a clock recovery phase uncertainty of 76 ps among the back-end cards; however, the clock phase recovered on the front-end could shift within a window of 4.4 ns after a loss of clock. The time alignment of the subdetectors is directly related to the clock phase, and a shift in the order of nanoseconds does not meet the LHCb requirements. Therefore, studies were conducted to optimize the clock distribution from back-end to front-end electronics (FEE), aiming to reduce the clock recovery phase uncertainty by an order of magnitude on the front-end boards.

II. THE LHCb TFC SYSTEM

As described in [3], the LHCb Timing and Fast Control system is responsible for generating and distributing real-time commands with fixed latency to control all the stages of the data readout between the front-end electronics and the DAQ cards. Additionally, it controls, monitors, and distributes the LHC bunch clock with a fixed and deterministic phase.

A. The LHCb PCIe40 Card

This subsection (II-A) is adapted from [3].

The LHCb PCIe40 is an FPGA-based PCIe Gen 3 interface card that characterizes the back-end of the Online system. It is equipped with an Intel Arria10 FPGA, 4x12 channels organized in 4 pairs of MiniPOD optical transmitters and receivers that are connected to 4 respective pairs of MPO connectors, two SFP+ modules, and two Si5345 Rev B. jitter attenuator PLLs, among other components dedicated to power, services and clock multiplexing. A functional diagram of the card is shown in Fig. 1.

The two links connected to the SFP+ modules are intended for timing distribution among back-end cards, whereas the 48 links of the MiniPODs are the interface to the front-end cards, which can be used for control or DAQ. At LHCb, 520 PCIe40s are used in total, and the cards perform three different roles depending on the firmware with which they are programmed:

- **SODIN**: it is the Readout Supervisor, where the real-time control commands are centrally originated from, and it is

Manuscript received May 29, 2024; revised Month? xx, 2024; accepted Month? xx, 2024. Date of publication Month? xx, 2024; date of current version Month? xx, 2024. This work was supported in part by CERN and in part by CNRS (LAPP/CPPM) France.

M. Feo, F. Alessio, P. Durante and F. Hachon are with CERN, Geneva, 1211 Switzerland. (e-mail: m.feo@cern.ch).

J-P. Cachemiche is with Centre de Physique des Particules de Marseille, Marseille, 13288 France.

G. Vouters is with Laboratoire d'Annecy de Physique des Particules, Annecy, 74941 France.

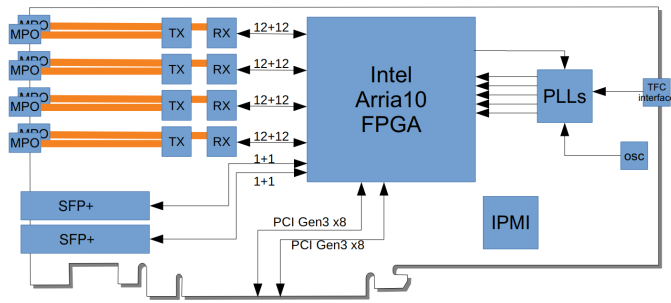


Fig. 1. Diagram of the LHCb PCIe40 card showing its main components.

the single point of entry for the LHC bunch clock for the LHCb experiment.

- **SOL40:** it is responsible for providing the required interfaces for control of the FEE as well as of the DAQ cards.
- **TELL40:** it is responsible for acquiring data from the FEE and delivering it to the Event Builder through the PCIe interface.

B. Timing Distribution

The timing distribution network of LHCb is depicted in Figure 2. At the core of the TFC system is the Readout Supervisor, housed on a single PCIe40 card programmed with a SODIN firmware. This firmware is divided into ten independent cores, each capable of controlling different detector partitions simultaneously. Each core generates commands, which are packed into a word that is transmitted to the Control cards (SOL40) via the SFP+ modules, through 1:32 passive optical network (PON) splitters. The SOL40 cards relay these commands to the appropriate DAQ cards (TELL40) in the same way. Such links that interconnect the back-end cards are called TFC links and employ the TTC-PON protocol [9] for data transmission and clock recovery.

The SOL40 cards are also responsible for the control of the front-end cards, using the GBT protocol from the CERN-GBT project [11], to embed timing and control information that is sent through the 48 MiniPOD optical links, which are referred to as GBT control links.

a) Clock recovery at the back-end: The SOL40 (control) and TELL40 (DAQ) cards recover the clock from the TFC data stream by shifting the 240 MHz recovered parallel clock in the FPGA transceiver until it aligns with a header, which is a fixed pattern inserted in the transmitted word. The TTC-PON firmware core outputs a strobe pulse aligned with the header, which is then used by the SOL40 firmware to generate the 40 MHz clock in phase with the LHC bunch clock. As described in [3], this technique allows the recovery of the clock among the back-end cards with a phase uncertainty of 76 ps.

b) Clock recovery at the front-end: On the front-end cards, the clock is recovered by the GBTx ASIC using a similar header alignment technique in accordance with the GBT protocol. The GBTx chip is capable of recovering the clock with a phase uncertainty of around 15 ps with respect to the transmitter reference clock, therefore, for the front-end to




Fig. 2. Architecture of the Timing Distribution System.

be in phase with the LHC clock, the SOL40 must guarantee a proper alignment between the transceiver reference clock and the 40 MHz recovered TFC clock. When measuring the uncertainty of the recovered clock phase at the FEE using the initial architecture of the SOL40, it was discovered that the phase could be randomly recovered within a window of 4.4 ns. This measurement was performed as described in II-D. The result indicates that a component along the chain would not keep a fixed phase when propagating the 240 MHz reference clock used by the FPGA transceivers. Given the LHCb requirement of 500 ps [3] for the maximum phase difference between the front-end and the LHC clock, a closer look at the initial SOL40 architecture was needed to understand the problem. Such architecture is described in Section II-C.

C. Initial SOL40 Architecture

A diagram of the initial SOL40 architecture is presented in Figure 3. The PON Rx block implements the TTC-PON protocol, recovering the 40 MHz TFC clock and injecting it into the FPGA global clock network to be used as the system clock. The internal fractional PLL then uses the 40

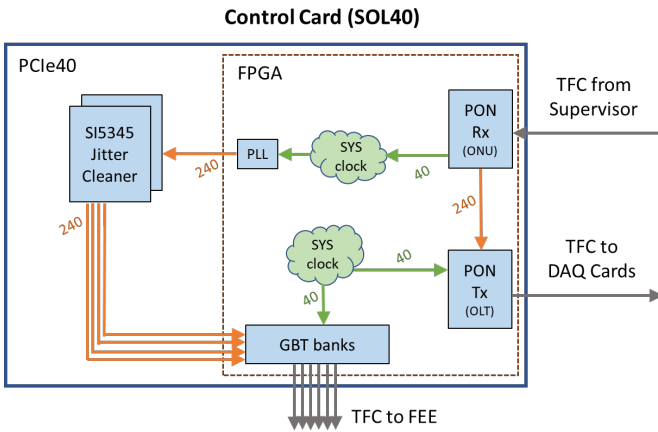


Fig. 3. Initial architecture of the SOL40.

MHz system clock to generate two 240 MHz signals that serve as inputs for both Si5345 PLLs, respectively, which attenuate jitter and provide 2x4 240 MHz reference clock signals for the 8 transceiver banks that interface with the front-end cards, indicated in Figure 3 as GBT banks.

The internal PLL used to supply the clock to the Si5345 PLLs does not maintain a fixed input-to-output phase, resulting in a clock with a random phase distributed over the 4.16 ns period of a 240 MHz clock. As presented in [3], the maximum measured skew between the recovered TFC clock and the original LHC bunch clock was 254 ps. Adding the 4.16 ns window of the internal PLL to the 254 ps skew of the TFC clock results in a 4.4 ns window, justifying the clock phase uncertainty measured on the front-end cards. Figure 4 shows the distribution of the recovered clock phase after 800 resets of the clock to the SOL40, with the filled area ranging from -2200 ps to 2200 ps.

In order to reduce the front-end phase uncertainty to meet the LHCb requirements, alternative solutions were studied based on two principles: ensuring fixed latency along the clock paths, as detailed in Section III, or implementing a mechanism to measure and adjust the phases accordingly, as detailed in Section IV.

D. Measuring the FEE Clock Recovery Phase Uncertainty

The uncertainty of the recovered clock phase at the front-end cards discussed in this article was measured in a laboratory setup reproducing the same architecture as in the experiment but with fewer cards and shorter optical fibers. A SODIN card clocked by its internal 40 MHz oscillator transmits the TFC stream through a 1:32 PON splitter to a SOL40 card, which is connected to a front-end card through a GBT control link. The internal SODIN clock and the recovered clock at the front-end card are connected to an oscilloscope configured to measure the delay between the rising edges of both signals. The measurement is then performed through the following steps:

- Perform a one-second reset of the SODIN TFC transceiver block (PON Tx) to induce the loss of link

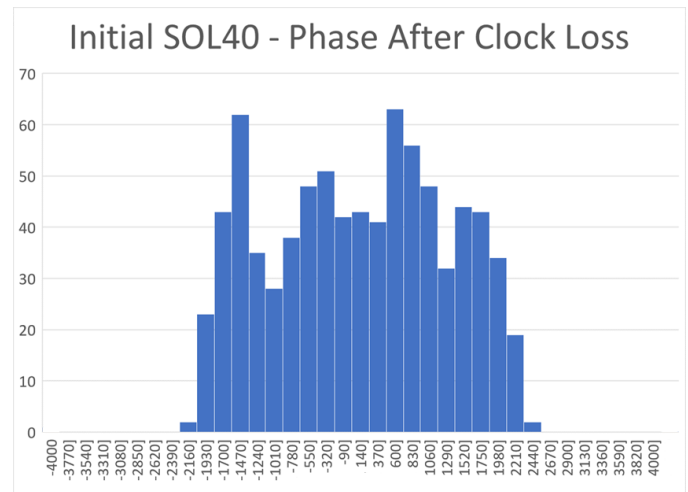


Fig. 4. Histogram of the recovered clock phase at the front-end, when using the initial SOL40 architecture, after 800 clock resets.

lock on the SOL40 and consequently disrupt the system clock.

- Reconfigure the SOL40 according to the firmware needs, by recalibrating the PON Rx transceiver, the PLLs in use, and reconfiguring any clock-related block on the firmware under test.
- Wait for the front-end clock edge to be present and stable.
- Measure the delay between the SODIN and front-end clock edges a thousand times and take the average.
- This represents one measurement. Repeat the steps until enough statistics are recorded.

III. SOL40 FIRMWARE WITH FIXED LATENCY

This solution consists of avoiding any component with no fixed and deterministic input-to-output phase relation. In the initial SOL40 architecture, these components included the internal PLL and the external Si5345 PLLs, depending on the mode of operation. Two alternatives following this principle were implemented.

A. Alternative 1: Bypassing the PLLs

This alternative consists of keeping the reference clocks inside the FPGA: instead of using the external Si5345 PLLs to attenuate the jitter and provide the reference clock to the GBT banks, the 240 MHz parallel clock recovered by the PON Rx transceiver is injected into the FPGA global clock network and used as the reference clock for the GBT transceiver banks. The diagram describing Alternative 1 is presented in Figure 5. This solution is not a recommended practice as, by skipping the Si5345 PLLs, the jitter of the reference clocks for the GBT banks would be considerably higher, affecting the quality of the control data transmitted to the front-end cards.

Performing 100 iterations of the procedure described in Section II-D results in the clock phase distribution seen in Figure 6. The phase measured at the front-end was consistently within a window of 15 ps. This is close to the limit of the GBTx chip capabilities in terms of phase recovery uncertainty.

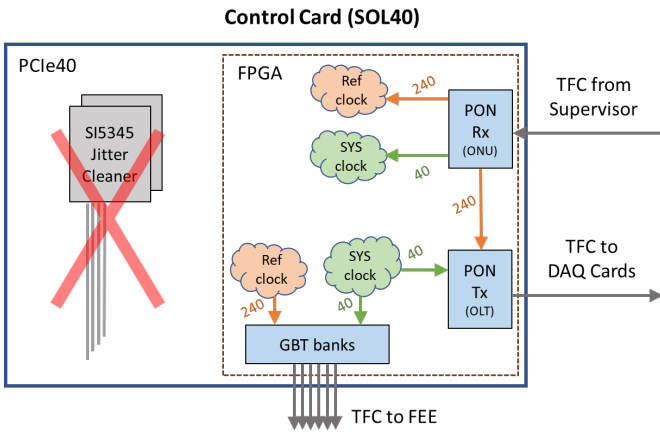


Fig. 5. Architecture of the SOL40 firmware with fixed latency by keeping the clocks inside the FPGA.

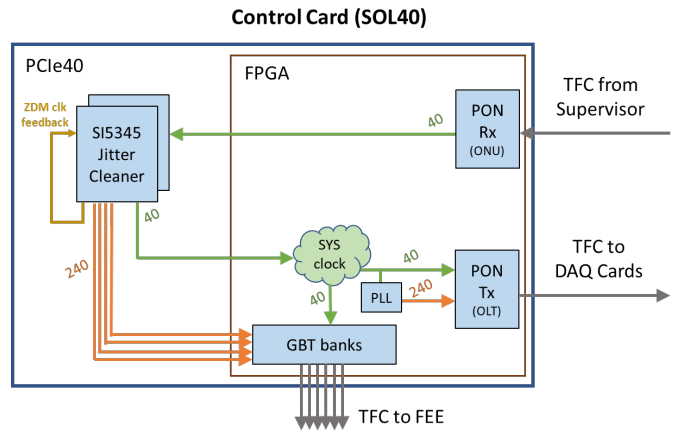


Fig. 7. Architecture of the SOL40 firmware with fixed latency by configuring the Si5345 PLLs in zero-delay mode.

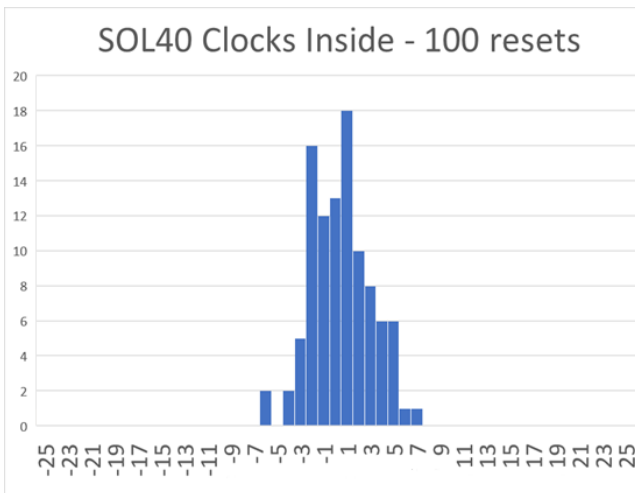


Fig. 6. Histogram of the recovered clock phase at the front-end, when bypassing the PLLs, after 100 clock resets.




Fig. 8. Relative timing jitter across the TFC system.

However, when performing a functional test with production SOL40 cards in the experiment connected to many front-ends, the control links to the front-end presented instabilities, counting FECs in some links after a few minutes with the system running, and eventually losing the lock of a control link. This is expected and most likely due to the increased jitter, which may cause bit errors in the control data fed into the GBT banks during the clock-domain-crossing from the 40 MHz system clock to the 240 MHz transceiver reference clock. Unfortunately, the loss of a control link interrupts the data-taking and requires a reconfiguration of the front-end card, disrupting LHCb operations and causing considerable dead time. Given the frequency in which such link failures would occur, this alternative could not be considered.

B. Alternative 2: PLLs in Zero-Delay mode

The idea behind Alternative 2 is to bypass the internal PLL but still use the Si5345 PLLs to attenuate the jitter and avoid the problem experienced with Alternative 1. This can be obtained by configuring the Si5345 PLLs in zero-delay

mode with a 40 MHz input and providing it directly from the 40 MHz TFC recovered clock. Figure 7 presents such architecture.

A great advantage of this solution is that a spare output of the Si5345s which is routed to the FPGA can be used as a jitter-attenuated system clock since this output can now be configured to 40 MHz. Because of that, the relative timing jitter between back-end cards and back-end to front-end is reduced by at least a factor of two. Figure 8 shows the reduction across different parts of the system.

The phase uncertainty was measured from 300 iterations of the procedure described in Section II-D. The front-end consistently recovered the clock within a window of 15 ps following the same distribution as Alternative 1, but the control links were stable and rarely presented FECs or losses of lock. During a functional test at the experiment using all SOL40s and available front-ends, however, the SOL40s of a specific server exhibited abnormal behavior: after configuring the PLLs in zero-delay mode, the GBT control links were unstable. Some links would lose lock constantly while others would not lock at all. Once the SOL40 enters this state, it can only

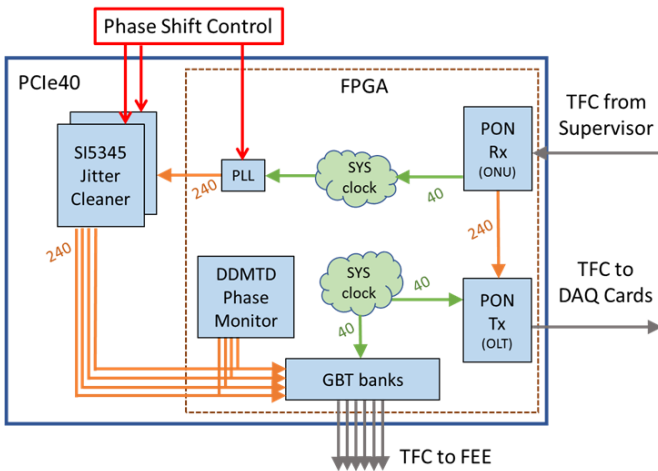


Fig. 9. SOL40 Firmware with Measurable and Adjustable Phases.

be recovered by a power cycle. Hard resetting or reconfiguring the PLL might return the expected status flags but the links continue unstable. Configuring the PLL back to normal mode and reprogramming the FPGA with stable firmware also does not help. This state seems to be triggered once you configure the PLLs in zero-delay mode, but this only happens in the 6 SOL40s of a specific server. Unfortunately, it was not possible to continue debugging the problem as it only occurred on a production server during detector operations and could not be reproduced in laboratory setups. It was decided to prioritize the architecture from Section IV until the problem can be properly investigated and understood.

IV. SOL40 FIRMWARE WITH ADJUSTABLE PHASES

In this approach, instead of avoiding the internal PLL where the phase shift occurs, the SOL40 firmware is capable of measuring the reference clock phases and using the phase-shifting capabilities of the PLLs to realign them. The random phase shifts within a 4.4 ns window will still occur after every clock loss, but it will be monitored and corrected by control software. The architecture of SOL40 Firmware with Adjustable Phases can be seen in Figure 9.

A. Phase-Adjusting Procedure

This solution uses the same clock routing as in the initial SOL40 architecture presented in II-C with the difference of the addition of a phase monitoring block that measures the GBT clock phases w.r.t. the TFC recovered clock using the Digital Dual Mixer Time Difference (DDMTD) [12] firmware block.

The phase-adjusting procedure is controlled by software. The control software can reset the DDMTD, trigger a new measurement, and read the result via register access. To achieve a precision of picoseconds, the control software performs 100 phase measurements and takes the average of the results. The phase of each of the 8 GBT transceiver reference clocks is read and the average phase of the 4 clocks from each of the two Si5345 PLLs is taken. The software then shifts the

two output clocks from the internal PLL individually, causing the 4 outputs of each Si5345 to shift together according to their input clocks, and aligning the first output of each Si5345 with a pre-determined setpoint. With the first output clock of each Si5345 aligned to such setpoint, the software proceeds to shift individually the 3 remaining outputs of each Si5345 to as close as it can get to the first clock. Once the procedure is complete, the control software verifies if all 8 GBT reference clocks are within the acceptable range and repeat the adjusting procedure until the clocks are aligned with the setpoint within a reasonable margin.

The precision with which the clocks can be aligned is limited by the shift resolution of the PLLs. The internal one being 104 ps and the Si5345 being 72 ps.

B. Compensating for Internal FPGA Delay

The result of the DDMTD phase measurement will correspond to the phase the probed clock has at the input of the DDMTD block, and not at the transceiver input. Because of signal propagation delays, the phase will have an offset at the DDMTD input. This offset can be provided by the timing analyzer tool of the FPGA vendor software used to compile the firmware and then compensated by software. Every new compilation will have different phase offsets for each of the clocks. Given the many different flavors used in LHCb, the process was automated. Scripts automatically generate text files with the phase offsets after a compilation, that are stored in the network in a way that the control software can identify which firmware is on an SOL40 and use the corresponding phase offset files.

C. Choosing the Phase Setpoint

The phase setpoint to which the SOL40 clocks are aligned is chosen to be the center of the stability window of the clock domain crossing between the system clock and the GBT reference clocks. This is achieved by manually shifting the clocks through the 4.16 ns period to identify the position where the GBT links start experiencing FECs or loss of lock, which defines the edges of the stability window. Once this is found, the setpoint for such firmware is set to the middle of the stability window. Such a window is usually around 2 ns wide and a new setpoint needs to be found for every new compilation.

D. Phase Uncertainty Results

The phase uncertainty when using the phase shifting mechanism was measured from 1100 iterations of the procedure described in Section II-D. The front-end for the vast majority of times recovered the clock within a window of 220 ps. Occasionally, the clock-shifting mechanism seems to shift beyond the setpoint, however, in these cases, the control software identifies the offset and adjusts the phase to the correct point. Figure 10 shows the distribution of the recovered phases at the front-end when using the SOL40 firmware with adjustable phases.

During the functional tests, the front-end links were stable, not presenting any loss of lock or FECs over many hours of




Fig. 10. Histogram of the recovered clock phase at the front-end after 1100 clock resets, when using the SOL40 Firmware with adjustable phases.

test. The SOL40 firmware with adjustable phases is currently the official firmware version for the SOL40 cards at LHCb.

V. CONCLUSION

In this paper, we presented an analysis and optimization of the Timing and Fast Control (TFC) system of the LHCb experiment's Online system. The initial implementation of the TFC system, while functional, exhibited a clock recovery phase uncertainty of 76 ps among back-end cards and a phase uncertainty of up to 4.4 ns at the front-end electronics, which did not meet the stringent requirements of the LHCb experiment.

To address these issues, we explored and implemented two primary solutions: ensuring fixed latency along the clock paths and implementing a mechanism to measure and adjust clock phases dynamically. The fixed latency approach, while effective in reducing phase uncertainty to 15 ps, encountered practical challenges related to jitter and link stability, particularly in a production environment. The adjustable phase approach, however, demonstrated a robust solution by allowing precise real-time alignment of clock phases using phase-shifting capabilities and a control software-driven adjustment process.

Our results show that the adjustable phase solution consistently aligns the front-end clock within a 220 ps window, maintaining stable links without loss or FECs over extended testing periods. This firmware has been adopted as the official solution for the SOL40 cards at LHCb, ensuring compliance with the phase alignment requirements essential for accurate data acquisition and control.

The findings and methodologies presented here not only enhance the performance of the LHCb TFC system but also provide a framework for future upgrades and similar implementations in other high-energy physics experiments requiring precise timing and control synchronization. Further investigations and optimizations can continue to refine these solutions, contributing to the overall efficiency and reliability of particle detection systems.

VI. CONCLUSION

This paper presents an analysis and optimization of the timing distribution between the SOL40 cards and the front-end electronics of the LHCb experiment. The initial implementation of the SOL40 firmware exhibited a clock recovery phase uncertainty of up to 4.4 ns at the front-end electronics, which did not meet the requirements of the LHCb experiment.

To address these issues, two primary solutions were explored: ensuring fixed latency along the clock paths and implementing a mechanism to measure and adjust clock phases dynamically. The fixed latency approach with the Si5345 PLLs in zero-delay mode, although effective in reducing phase uncertainty to 15 ps and decreasing relative timing jitter by at least a factor of two, encountered practical challenges related to link stability, particularly in a production environment when further investigation is very difficult. The adjustable phase approach demonstrated a robust solution by allowing real-time alignment of clock phases using the PLL phase-shifting capabilities, that can be monitored and controlled by software.

The results show that the adjustable phase solution consistently aligns the front-end clock within a 220 ps window, maintaining stable links without loss of lock or FECs over extended testing periods. This firmware has been adopted as the official solution for the SOL40 cards at LHCb, ensuring compliance with the phase alignment requirements essential for accurate time alignment of the various subdetectors.

The findings and methodologies presented here enhance the performance of the LHCb TFC system and provide a framework for future upgrades. The current version of LHCb's detector can comfortably cope with a timing uncertainty of hundreds of picoseconds, however, the next upgrade will require results in the order of the 15 ps achieved with the fixed latency approach and further investigations and optimizations can continue to refine these solutions or provide answers on how to extract the best performance from such versatile electronics cards like the PCIe40.

REFERENCES

- [1] The LHCb Collaboration, "The LHCb Detector at the LHC," *JINST* 3, (2008) S08005 [Online]. Available: <https://cds.cern.ch/record/1129809>. Accessed on: July 30, 2022.
- [2] The LHCb Collaboration, "LHCb trigger and online upgrade technical design report," CERN, Meyrin, Switzerland, Tech. Rep. CERN-LHCC-2014-016; LHCb-TDR-016, 2014. [Online]. Available: <https://cds.cern.ch/record/170136>. Accessed on: July 30, 2022.
- [3] Feo, M., Alessio, F., Durante, P., Cardoso, L. and Vouters, G. "The Real-Time System for Distribution of Clock, Control and Monitoring Commands With Fixed Latency of the LHCb Experiment at CERN," *IEEE Trans. Nuclear Science*, vol. 70, no. 6, pp. 985–992, June 2023, doi: 10.1109/TNS.2023.3273086 [Online]. Available: <https://ieeexplore.ieee.org/document/10115510>. Accessed on: May 20, 2024.
- [4] Cachemiche, J-P., Duval, P. Y., Hachon, F., Le Gac, R. and Réthoré, F., "The PCIe-based readout system for the LHCb experiment," *JINST* 11, (2016) P02013 [Online]. Available: <https://cds.cern.ch/record/2262859>. Accessed on: July 30, 2022.
- [5] Mendes, E., Baron, S., Soos, C. and Vasey, F., "A Timing, Trigger, and Control System With Picosecond Precision Based on 10 Gbit/s Passive Optical Networks for High-Energy Physics," *IEEE Trans. Nuclear Science*, vol. 68, no. 4, pp. 447–457, April 2021, doi: 10.1109/TNS.2021.3065128. [Online]. Available: <https://ieeexplore.ieee.org/document/9374437>. Accessed on: July 30, 2022.

- [6] Brarda, L., Jost, B., Lacarrere, D., Lindner, R., Neufeld, N., Roy, L. and Thomas, E., "A new data-centre for the LHCb experiment," *J. Phys.: Conf. Ser.*, volume 396, pages 012009. 6 p, 2012. [Online]. Available: <https://cds.cern.ch/record/1565928>. Accessed on: July 30, 2022.
- [7] Alessio, F., Durante, P. and Vouters, G., "The readout supervisor firmware for controlling the upgraded LHCb detector and readout system," arXiv, 1806.08626, 2018. [Online]. Available: <http://cds.cern.ch/record/2629423>. Accessed on: July 30, 2022.
- [8] Alessio, F. and Jacobsson, R., "System-level Specifications of the Timing and Fast Control System for the LHCb Upgrade," CERN, Geneva, Switzerland. Tech. Rep. CERN-LHCb-PUB-2012-001, Feb/2014. [Online]. Available: <https://cds.cern.ch/record/1424363>. Accessed on: July 30, 2022.
- [9] Mendes, E., Baron, S., Soos, C. and Vasey, F., "A Timing, Trigger, and Control System With Picosecond Precision Based on 10 Gbit/s Passive Optical Networks for High-Energy Physics," *IEEE Trans. Nuclear Science*, vol. 68, no. 4, pp. 447–457, April 2021, doi: 10.1109/TNS.2021.3065128. [Online]. Available: <https://ieeexplore.ieee.org/document/9374437>. Accessed on: July 30, 2022.
- [10] Amaral, L., Dris, S., Gerardin, A., *et al.* "The versatile link, a common project for super-LHC," *JINST* 4, (2009) P12003 [Online]. Available: <http://cds.cern.ch/record/1265172>. Accessed on: July 30, 2022.
- [11] Moreira, P., Marchioro, A. and Kloukinas, "The GBT: A proposed architecture for multi-Gb/s data transmission in high energy physics," in *TWEPP*, Prague, Czech Republic, 2007, pp.332-336. [Online]. Available: <http://cds.cern.ch/record/1091474>. Accessed on: July 30, 2022.
- [12] Moreira, P., Alvarez, P., Serrano, J., Darwezeh, I. and Wlostowski, T., "Digital dual mixer time difference for sub-nanosecond time synchronization in Ethernet," *IEEE International Frequency Control Symposium*, pp. 449-453, 2010, doi:10.1109/FREQ.2010.5556289 [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/5556289>. Accessed on: May 20, 2024.