

RECEIVED: November 2, 2024

REVISED: December 23, 2024

ACCEPTED: January 29, 2025

PUBLISHED: March 11, 2025

TOPICAL WORKSHOP ON ELECTRONICS FOR PARTICLE PHYSICS
UNIVERSITY OF GLASGOW, SCOTLAND, U.K.
30 SEPTEMBER–4 OCTOBER 2024

Ensuring clock phase repeatability by preventing loss of the 40.078 MHz clock in time-critical detectors: a non-disruptive clock switching approach

N. Loukas^{a,*} L. Moreno^{b,*} T. Anderson^c G. Cucciati^a Z. Eberle^d S. Goadhouse^{a,c} T. Gorski^e and A. Svetek^e on behalf of the CMS collaboration

^aDepartment of Physics and Astronomy, University of Notre Dame,
225 Nieuwland Science Hall, Notre Dame, IN 46556-5670, U.S.A.

^bDepartment of Physics, Princeton University,
Jadwin Hall, Washington Road, Princeton, NJ 08540, U.S.A.

^cDepartment of Physics, University of Virginia,
382 McCormick Road, Charlottesville, VA 22904, U.S.A.

^dSchool of Physics and Astronomy, University of Minnesota,
John T. Tate Hall, 130, 116 Church St SE, Minneapolis, MN 55455, U.S.A.

^eDepartment of Physics, University of Wisconsin-Madison,
Thomas C Chamberlin Hall, 1150 University Ave 2320, Madison, WI 53706, U.S.A.

E-mail: nikitas.loukas@cern.ch, lperezmo@cern.ch

ABSTRACT: For the High Luminosity phase of the LHC (HL-LHC), the CMS detector electronics require precise and accurate timing to discriminate pileup events. The upgraded Back-End electronics of the Barrel Electromagnetic Calorimeter supply the required high-precision clock (Time Interval Error $\sigma < 5$ ps) to the Front-End. However, after a reset is applied on the timing distribution system, the resulting phase of the clock is not accurately repetitive but it gives an up to 10 ps offset for each FPGA node of the timing distribution chain. That is resulting to the uncertainty on the total clock skew with the maximum being $N \times 10$ ps between neighbor channels of the system (where N is the number of FPGA-nodes of the subdetector branch). We have studied a solution minimizing the need to reset the timing distribution system by switching on-the-fly and seamlessly the source clock at the top of the tree to a local LHC-frequency clock. This approach prevents the loss of the links towards the on-detector electronics ensuring clock phase repeatability and offers stability to the on-detectors electronics.

KEYWORDS: Control and monitor systems online; Digital electronic circuits; Calorimeters

*Corresponding author.

Contents

1	Introduction	1
2	The challenge of achieving picosecond level clock phase accuracy	1
3	A non-disruptive clock switching approach	4
4	Testing seamless switching between global clock and local clock	4
5	Conclusions	5

1 Introduction

In order to achieve the required HL-LHC physics performance [1], the CMS Electromagnetic CALorimeter (ECAL) Barrel [2, 3] needs to give an overall precision of 30 ps on the arrival time of photons produced by Higgs bosons decays. To achieve this overall system requirement, each ADC channel of the subdetector requires even better precision as they are only one part of the overall timing precision budget. The FPGA-based electronics recover the source LHC clock with a Time Interval Error (TIE) typically better than $\sigma = 5$ ps, which meets the low jitter clock distribution requirement [4]. However, after a reset is applied to the receiver of the Multi-Gigabit Transceiver (MGT), the repeatability of the output clock phase may vary up to 10 ps peak-to-peak per FPGA node in the clock distribution system. These resets are required whenever the source clock becomes unstable, which can happen at every new physics run or during intervention, when the LHC machine clock is being reset. Since the timing system consists of several nodes in series, after a reconfiguration, the total skew between neighbor endpoints of the subsystem accumulates to a multiple of 10 ps. In the case of CMS ECAL Barrel, which has 3 FPGA nodes in series, the accuracy error between two endpoint channels can vary from 0 up to $3 \times 10 = 30$ ps. This can, in principle, be reduced through calibration with online physics data analysis at the beginning of each acquisition run, however, that introduces complexity and dead time on the data acquisition. In addition, due to the glitch on the LHC reference clock, the Front-End optical links stop transmitting which require the Front-Ends to be reconfigured. This causes one additional minute of delay from starting each CMS run.

2 The challenge of achieving picosecond level clock phase accuracy

To understand the clock phase repeatability of the FPGA’s MGT deserializer after a reset and its behavior versus the operating temperature, we installed all Phase-2 Back-End and Front-End prototypes of the CMS ECAL Barrel into a climatic chamber (figure 2) and we controlled the ambient temperature to vary no more than 0.1 °C. This ensures no phase drift due to temperature variance in the fibers and optics. We used an 1.3 ps TIE jitter clock generator to supply the Master card (AMD’s VU118 FPGA dev. card). The Master fed a clock to the Back-End (Barrel Calorimeter Processor, BCP) which fed a clock to the Front-End (IpGBT [5]) which, finally, synchronized the Very-Front-End electronics [6] and its LiTeDTU ASICs [7] whose PLLs clocked their internal ADCs (figure 1).




Figure 1. Emulating ECAL Barrel clock distribution.

With the use of a 10 GHz bandwidth oscilloscope, we measured the clock phase difference (skew) of the Clock Generator with respect to the recovered clock on each stage (BCP, lpGBT, LiTeDTU). As we aimed to measure the accuracy and not the precision of the clock, we ignored the jitter and focused on clock accuracy. Therefore, we acquired the skew a thousand times and we used their mean as measurement points. We performed the following tests and the results are shown in the plots of figure 3.




Figure 2. CMS-ECAL Barrel prototypes in the climatic chamber.

- a) The first test (figure 3(a)) shows the distribution of the mean of the skew between Master and BCP (in blue), Master and lpGBT (in orange) and Master and LiTeDTU (in green). This test represents a stable system as there are no resets. The standard deviation of those Gaussian distributions is: Master versus BCP 1.4 ps, versus lpGBT 1.5 ps and versus LiTeDTU 1.6 ps. In this test, we had only one FPGA to FPGA node (Master → BCP) but in ECAL there will be two more FPGA to FPGA nodes. Nevertheless, the observed degradation is of 0.1 ps rms per node, and that therefore we are expecting only 0.5 ps in total.
- b) In the second test (figure 3(b)), all hardware was reset before every measurement. As in test a), the FPGA temperature (in Super Logic Region 0, SLR0 [8, p. 13]) was stable at 38 °C and the ambient temperature at 25 °C. This test gave a dual-Gaussian distribution on the Master versus BCP (blue plot). We can clearly see that the two-peaks propagated downstream to the on-detector electronics (orange and green plots) and the spread of the skew widened. Focusing




Figure 3. Measuring clock phase jumps.

on the FPGA-node (blue plot), we measured $\sigma = 1.4$ ps on the right Gaussian (same as in test a)) and $\sigma = 1.7$ ps on the left Gaussian. The right holds 50.5 % of the entries and the left holds 49.5 % of the total entries. The difference between these two peaks is 10.2 ps. Hence the phase offset which occurs due to the reset is 0 (no offset) or 10 ps (offset) with approximately 50 % probability of having the 10 ps phase jump.

c) In the third test, we did not plot the distribution of the mean skew but instead the mean skew versus the number of measurements. In this case, we gradually increased the internal temperature of the FPGA using internal heaters. This raised the internal FPGA temperature from 35 °C to 76 °C while the climatic chamber kept the ambient temperature stable to 25 °C. We increased the temperature every 100 measurements resetting every 5 measurements. As can be seen in plot 3(d), the clock skew reduces and then widens and reduces again as we heated the FPGA. Around $SLR0 = 38$ °C (as in test (b)) and around 55 °C, the clock phase jumps between two zones with a probability of about 50%. However around 45 °C and 60 °C the clock phase stays within a signal narrow zone (no phase jumps). At higher temperatures, the behavior changes and the clock phase jumps again but this time the probability being even at 90% for one of

the two zones. These results illustrate the complexity of these phase jumps especially when we consider that outside of the climatic chamber in an ATCA crate, the temperature on the Back-End system cannot be controlled.

- d) In this test, we attempted to find the cause of these phase jumps. We probed the recovered clock at three points of the FPGA [9, p. 217]: the RXUSRCLK (blue), the RXOUTCLKPCS (purple) and the RXRECLKOUT (red). As we can see in figure 3(c), the blue RXUSRCLK clock suffers from the phase jumps but not the other two. RXUSRCLK makes use of the internal Delay Aligner circuit which compensates for temperature variations, avoids metastability and ensures the integrity of the recovered data and therefore its use is mandatory.

3 A non-disruptive clock switching approach

From section 2, it becomes clear that a reset on the FPGA deserializer, which recovers the LHC clock, causes a phase jump whose width depends on the operating conditions. One approach to mitigate this problem is to detect the jump when it happens and compensate for it. Given the difficulty of detecting picosecond differences in the clock and then accurately adjusting while the temperature fluctuates, this is a challenge. Another approach, which we are proposing in this paper, is a non-disruptive clock switch technique to a local clock when the global source clock is unstable. This prevents the loss of the downstream 40.078 MHz clock. The key point of the proposed method is that it avoids resetting the downstream system while the upstream clock is unstable. Instead, the system switches, on-the-fly, to a local clock until the source LHC clock becomes stable. If such a technique can be applied at the top of the clock distribution chain, none of the downstream parts of the subdetector will need to be reset when the source clock is unstable and therefore those downstream endpoints will not see a phase jump.

In general, a random clock switch can cause a glitch which upsets the logic of the FPGA and its high speed links. In our approach, we use the jitter cleaner [10] in Zero Delay Mode configuration to seamlessly switch the input clock to the FPGA without glitches. In section 4, we evaluate whether or not this clock switch is feasible and if it causes a negative impact to the system.

4 Testing seamless switching between global clock and local clock

In order to evaluate the non-disruptive clock switching method, we performed the two following tests. The tests were all run with all PLLs locked to ensure signal integrity and synchronization. Skyworks Si53xx family jitter cleaners [10, 11] were used for their hitless switching feature. We used the BCP, or equivalent ATCA FED (Front-End Driver), along with a DAQ and TCDS Hub (DTH) hardware [12] installed in an ATCA crate:

- A) As can be seen in the block diagram of figure 4, in the BCP we used two Si5345 jitter cleaners. The Source-Si5345 (yellow) received a 40.07857 MHz clock from a local clock generator (scope channel 1). The Sync-Si5345 (pink) received one input from the Source-Si5345 (local clock) and another clock input from the recovered TCDS link. That link was the global clock driven by the DTH card (scope channel 4). The Sync-Si5345 sent two clocks to the MGT (multi-gigabit transceiver) serializers of the FPGA (scope channel 2). One MGT was driving one ECAL Front-End (FE) tower (scope channel 3) and another one was in an optical loopback. The lookback SERDES was running a 10 Gb/s Bit-Error-Rate test using PRBS31 patterns. The




Figure 4. Testing the seamless clock switch.

FE tower was configured to send data with CRC checksums to the BCP. During the test, we continuously monitored all 4 10-Gbps lpGBT up-links of the FE Tower for errors. The test consisted of switching the Sync-Si5345 between local and global clock inputs and verifying that the clock output followed the frequency of the switched input on the oscilloscope. We applied the clock switching 1000 times and we got zero errors from the BER test during the entire test. The equivalent BER was in the order of 10^{-15} . Note that the difference between the two clock sources was only 11 Hz and that the PLL on the Sync-Si5345 always remained locked.

B) For this test, we wanted to increase the number of simultaneous channels in order to gather more statistics and to switch the clock inside of the FPGA, where it is more flexible, instead of the Si53XX. The FPGA switched clock output was still fed through a Si53XX. Due to a lack of available prototype hardware, we replaced the FE links with optical loopbacks and used an APxF ATCA card [13], which has a similar architecture to the BCP. We employed 48 10.24 Gbps links, and we switched the clocks 100 times. No errors were observed on any link resulting in a BER better than 10^{-14} . It is worth mentioning that the frequency difference between the two clock sources for this test was much higher (40.078 MHz–40.000 MHz = 78 KHz) than in test A). This caused the PLL of the Sync-Si5345 to go in free-running-mode during the switch between clocks but the links were still able to continually operate error-free.

The positive results of these proof-of-concept tests show that it is possible to maintain a continuous and error-free communication link between the detector front-end electronics and their back-end FEDs by having the FED monitor the global clock and automatically switch to a local free-running clock if the global clock ever goes out of specification. By testing with the prototype electronics for the ECAL HL-LHC upgrade, we are confident that this can be implemented with our final hardware that is currently being produced.

5 Conclusions

As we see in section 2, a clock phase jump up to 10 ps can occur on the MGT deserializers of FPGAs when they reset. The phase jump accumulates to the number of FPGAs used in series, depends on temperature and propagates to the next nodes of the system resulting to an estimation of 30 ps accuracy

error for ECAL Barrel. The Clock Switching Approach that we introduce in this paper is shown to be non-disruptive and prevents phase jumps from occurring in the downstream FPGA nodes as the MGT serializers do not unlock. In addition, the front-end electronics are not forced to be reconfigured, which is time consuming and can cause a loss in physics data.

Acknowledgments

The authors would like to acknowledge the National Science Foundation (NSF) and the Department of Energy (DOE) for funding the US contribution to this research.

References

- [1] CMS collaboration, *The Phase-2 Upgrade of the CMS Barrel Calorimeters*, [CERN-LHCC-2017-011](https://cds.cern.ch/record/2274211), CERN, Geneva (2017).
- [2] CMS collaboration, *The CMS Experiment at the CERN LHC*, [2008 JINST **3** S08004](https://cds.cern.ch/record/1185313).
- [3] CMS collaboration, *Development of the CMS detector for the CERN LHC Run 3*, [2024 JINST **19** P05064](https://cds.cern.ch/record/2404514) [[arXiv:2309.05466](https://arxiv.org/abs/2309.05466)].
- [4] N. Loukas et al., *The CMS Barrel Calorimeter Processor demonstrator (BCPv1) board evaluation*, [2022 JINST **17** C08005](https://cds.cern.ch/record/2604000).
- [5] A. Dolgopolov, C. Jessop, N. Loukas and A. Singovski, *CMS ECAL Upgrade Front End card: design and prototype test results*, [PoS TWEPP2018](https://cds.cern.ch/record/2604001) (2019) 044.
- [6] W. Lustermann et al., *CMS ECAL VFE design, production and testing*, [2024 JINST **19** C05034](https://cds.cern.ch/record/2604002) [[arXiv:2311.02021](https://arxiv.org/abs/2311.02021)].
- [7] G. Mazza et al., *The LiTE-DTU: A Data Conversion and Compression ASIC for the Readout of the CMS Electromagnetic Calorimeter*, [IEEE Trans. Nucl. Sci. **70**](https://cds.cern.ch/record/2604003) (2023) 1215.
- [8] Xilinx, *Large FPGA Methodology Guide*, document UG578 (v14.3), (2012), available: https://www.xilinx.com/support/documents/sw_manuals/xilinx2012_3/ug872_largefpga.pdf.
- [9] AMD, *UltraScale Architecture GTY Transceivers User Guide*, document UG578 (v1.3.1) (2021), available: https://www.amd.com/content/dam/xilinx/support/documents/user_guides/ug578-ultrascale-gty-transceivers.pdf.
- [10] Skyworks, *Si5345/44/42 Rev D Data Sheet*, Tech. Rep. [Rev. 1.2](https://www.skyworksolutions.com/documents/si5345-44-42-rev-d-data-sheet), Skyworks Solutions, Inc., Irvine, CA, U.S.A. (2021).
- [11] Skyworks, *Si5395/94/92 Rev D Data Sheet*, Tech. Rep. [Rev. 1.2](https://www.skyworksolutions.com/documents/si5395-94-92-rev-d-data-sheet), Skyworks Solutions, Inc., Irvine, CA, U.S.A. (2021).
- [12] J. Hegeman et al., *First measurements with the CMS DAQ and Timing Hub prototype-1*, [PoS TWEPP2019](https://cds.cern.ch/record/2604004) (2020) 111.
- [13] C. Herwig, *Particle flow reconstruction for the CMS Phase-II Level-1 Trigger*, [2023 JINST **18** C01037](https://cds.cern.ch/record/2604005).