Available on CMS information server **CMS CR -2022/201**

**The Compact Muon Solenoid Experiment**

**Conference Report**

Mailing address: CMS CERN, CH-1211 GENEVA 23, Switzerland

**18 October 2022 (v4, 21 October 2022)**

# The CMS Muon Drift Tubes HL-LHC Upgrade Demonstrator at Run 3 Start

Ignacio Redondo Fernandez for the CMS Collaboration

## Abstract

The electronic system of the CMS Drift Tubes (DT) chambers will be replaced to operate during High Luminosity (HL-LHC). The upgraded architecture ships all signals to the backend, where complex logic will be performed with a precision matching the maximum chamber resolution, now only available offline. A demonstrator has been installed during Long Shutdown 2 (LS2) in one of the sixty sectors of the detector. Over LS2 we have integrated this system in CMS operations environment and tested its stability over extended cosmics data-taking campaigns, also with the magnetic field on. This paper describes the time synchronization achieved and early performance in collisions at Run3 start.

Presented at *TWEPP2022 Topical Workshop of Electronics for Particle Physics*

# The CMS Muon Drift Tubes HL-LHC Upgrade Demonstrator at Run 3 Start

**I. Redondo,**[a] **on behalf of the CMS Collaboration**
[a] *Centro de Investigaciones Mediambentales y Tecnlógicas,*
*Av. Complutense 40 Madrid 28040, Spain*
*E-mail*: redondo@cern.ch

ABSTRACT: The electronic system of the CMS Drift Tubes (DT) chambers will be replaced to continue operation during high luminosity running of the LHC (HL-LHC). The upgraded architecture sends all signals to the backend, where complex logic will be performed with a precision matching the maximum chamber resolution, now only available offline. A demonstrator has been installed during Long Shutdown 2 (LS2) in one of the sixty sectors of the detector. Over LS2 we have integrated this system in the CMS operations environment and tested its stability over extended cosmics data-taking campaigns, also with the magnetic field on. This paper describes the time synchronization achieved and early performance in collisions at Run 3 start.

KEYWORDS: CMS Muon; System Calibration; High Luminosity LHC; Electronics.

## Contents

## 1. The CMS Muon Drift Tubes HL-LHC Upgrade

The CMS [1] DT chambers instrument the return yoke of CMS, being responsible for identifying, measuring, and triggering on muons in the barrel acceptance region. The electronic system will be replaced during Long Shutdown 3 (LS3) to be able to cope with the stringent data-taking conditions during HL-LHC [2]. The new architecture sends the full time digitized data through high-speed optical links to the backend, where trigger primitives and event matching logic will be performed. This is in contrast with the legacy system which was forced to perform the above-mentioned functionalities locally in the front-end with lower precision due to bandwidth limitations of the available copper links. The time signals coming from the chambers will be digitized by 830 On detector Boards for the Drift Tube chambers (OBDTs) with <1 ns precision. The OBDTs are built around a Microsemi PolarFire FPGA and will use CERN's lpGBT and VTRX+ optics in its production version, now under validation.

### 1.1 Project Challenges

The Phase-2 electronics system is designed to match the maximally achievable chamber precision (~<1 ns), which now is only possible offline. Demonstrating that offline algorithms can be ported to a state of the art FPGA within CMS L1 trigger constraints is one of the main challenges of the project. The DT chambers reconstruct track coordinates through time measurements, thus the clock distribution needs to be kept stable to < 1 ns over active components (serial links & FPGAs).

The upgrade needs to comply with integration constraints related to installation in an already operational detector like CMS such as fitting into the power and cooling envelopes. In addition, it is desirable to reuse the optical fibers from the CMS towers to the counting room that were installed in the cable chains and trigger tunnels during LS1. Moreover, access time and service availability for commissioning in LS3 imposes severe constraints that need to be considered. Therefore during LS2 several mechanical aspects were investigated aiming to streamline installation and commissioning. The FE copper cables of the 2 DT $\varphi$ Superlayers per chamber are accessible from the face of the wheel, which makes possible replacing them by less bulky individually shielded cables. It was identified that using cables of variable length, in contrast to the constant length used in the legacy system, would speed up installation reducing the spare cable management. This leverages on the possibility to have individual channel delays in the backend. Flexibility to recover from mapping mistakes after access is lost also relaxes the time needed for commissioning with access to the detector.

Another aspect known from the start was that the FE copper cables from the DT $\theta$ Superlayers cannot be realistically replaced because this would involve extracting all the chambers from the iron yoke. This motivated having two flavours of OBDT boards, one with connectors compatible with the legacy cables. The OBDT boards have to comply with moderate radiation hardness requirements (<0.5 Gy), which is being verified with final firmware optics (see C.F. Bedoya contribution, this workshop).

Finally, HL-LHC background levels will be larger and the chambers need to be operated over a period of time with irradiation larger than what they were originally designed to tolerate. Thus, the Phase-2 design needs to be resilient to aging inefficiencies profiting from the redundancy provided by the large number of detection layers present in the system.

## 2. Slice Test Installation and Commissioning

During LS2 we have instrumented one of the sixty sectors of the CMS DT subdetector with 13 prototypes of the first version of the On detector Board for the Drift Tube chambers (OBDTv1 [3]) , which differs from the final version in using GBTx links and commercial optics, but still hosts a PolarFire that digitizes 240 channels. One OBDTv1 covers one Superlayer, except in the MB4 chamber where 2 are needed per Superlayer. Clock distribution, slow control, readout, and trigger primitive generation are performed in the demonstrator backend, which is based on legacy hardware (DT uTCA boards, TM7 [4]) with firmware performing the required Phase-2 functionality. One TM7 board delivers clock and slow control to 12 OBDTs through its minipod RX/TX links; 2 boards with this firmware (MOCO) are needed. To receive the TDC data, 5 TM7s have different firmware (AB7) that hosts the Analytical Method (AM) [5] trigger algorithm logic and the readout for up to three OBDTs.

In this demonstrator, called the DT Phase-2 Slice Test, the time pulses coming from the chambers are split allowing standard operation of the legacy system and full event-by-event validation of the prototype electronics and the firmware logic. This is greatly helped by fully integrating the Phase-2 chain data into the CMS DAQ environment as a standard CMS DAQ unit (FED). Detector Control System and DQM were also integrated into central CMS systems. During extended cosmics data-taking campaigns in LS2, with and without magnetic field, the long-term stability of the system and the synchronization methods of the ~3000 channels were thoroughly tested. The reconnection of all channels through splitters required verifying the geometrical mapping of all channels that are logically asigned at the backend level, thus keeping the flexibility to recover from cabling errors.

The time scale stability of clock phases used in the TDC firmware in the polarfire over electronics power cycles or resets, called "phase determinism", was identified in the Slice Test as the dominant problem by comparing the time offset between Phase-1 and Phase-2 and was fixed in the OBDT firmware. The problem was traced back to issues with the phases of the resets of the internal PLLs of the Polarfire and the derivation of the fast clocks from the 40 MHz clock in the same FPGA. The main cause of data inefficiency has been an intermittent unlock of GBT optical links, although this did not prevent data taking with the 13 OBDTs for most of the time. The unlocks were typically triggered by TCDS clock reset, which happens regularly during CMS operations cycle. The issue has been mitigated through automatic recovery procedures, but was not investigated thoroughly under the consideration that the final system uses lpGBT.

Before the installation in the detector, all parts of the Slice Test were exercised in an integration setup on the surface at P5, where we have a fully functional DT chamber and access to CMS TCDS. There we have already tested more advanced prototypes (the lpGBT OBDTv2 θ) and the final backend board (BMTL1), which will port one VU13P and firefly optics (see I. Bestinazos contribution, this workshop), achieving full vertical integration in firmware from a DT chamber through the BMTL1 backend to GMT with Phase-2 backend prototypes.

### 2.1 Slice Test Synchronization Strategy

Synchronization of all channels in the detector at all granularity levels (Superlayer, OBDT, chamber, sector) is the main calibration requirement for optimal detector performance. To achieve that, the firmware and software infrastructure was deployed to allow pre-setting of logical delays per channel such that all digitized signals can be aligned based on the time digitized value at the trigger processors to perform segment and bunch-crossing (BX) matching, even if signal arrival times to the backend vary. The first step of calibration is the inter-channel equalisation to homogenize each superlayer: it consists in the determination of individual delays for every channel (each anode wire). The inter-channel calibration delays were obtained using Test Pulse calibration data taken during LS2 and unchanged since, once the stability was monitored for over a month. In these calibration runs, signal pulses are injected at the frontend in all channels simultaneously by the legacy electronics (CCB board) and the resulting TDC times, readout from the Phase-2 electronics, are recorded. Details can be found in [6].

At a second stage, to synchronize the sector, each Superlayer has to be corrected by time-of-flight (TOF) effects and delays generated in the clock distribution. The optimal delay of each Superlayer is calculated by setting to zero the average time of passage of the 4-hit primitives in each φ Superlayer, profiting from the rather good measurement of the arrival time by 4-hit segments. θ Superlayer delays are obtained by shifting the timebox to match the φ Superlayers, but the TM7 resources were too limited to deploy the AM on these Superlayers. A per channel synchronization to take care of TOF effects, relevant in these Superlayers, is possible and a potential improvement in the future.

Optimal delays between Superlayers depend on track direction and timing, which with cosmic rays, can be selected to a certain extent using trigger and offline cuts. Delays were measured and applied for different configurations throughout LS2, and in particular the trigger was configured to select signals from sectors opposite to the ST one in cosmics to i) enhance the fraction of tracks pointing to the interaction point (IP), ii) have similar TOF among Superlayers and latency as in collisions, and iii) avoid trigger bias, because the chambers participating in the trigger were not part of the Slice Test. Fully calibrated results from cosmics can be found here [7,8]. Calibration of the optimal fully configurable delays during LS2 with cosmic-ray muons and the obtained channel-by-channel delays were used as the starting calibration for early Run 3 collisions.

## 2.2 Selection

The information extracted from Phase-2 primitives (generated by the AB7 boards) is compared with Phase-1 segments. Phase-1 segments are reconstructed offline from legacy hits and selected to have i) an inclination with respect to the direction orthogonal to the chamber layer plane < 30°, ii) at least 4 hits, and iii) a $|t0| < 50$ ns, where t0 stands for the time associated to the segment as reconstructed by a 3-parameter fit. The start-up calibration was generated from a cosmic muon sample with 3.8 T collected with the Slice Test setup and triggered by opposite sectors in a global run with the Barrel Muon Track Finder (BMTF). The results shown in the figures are based on a collision data sample of 1.9 $\text{fb}^{-1}$ that was collected during LHC Run 3 startup in July and August 2022.

In collisions, background contributions bias the distributions significantly by up to 4 ns. Thus a sample of events with tight muons (see [9] for the definition) in the Slice Test acceptance was selected requiring $p_T > 25$ GeV. Delays obtained were very similar if the tight muon requirement was released, keeping the selection of a muon in the acceptance of the Slice Test (ST).

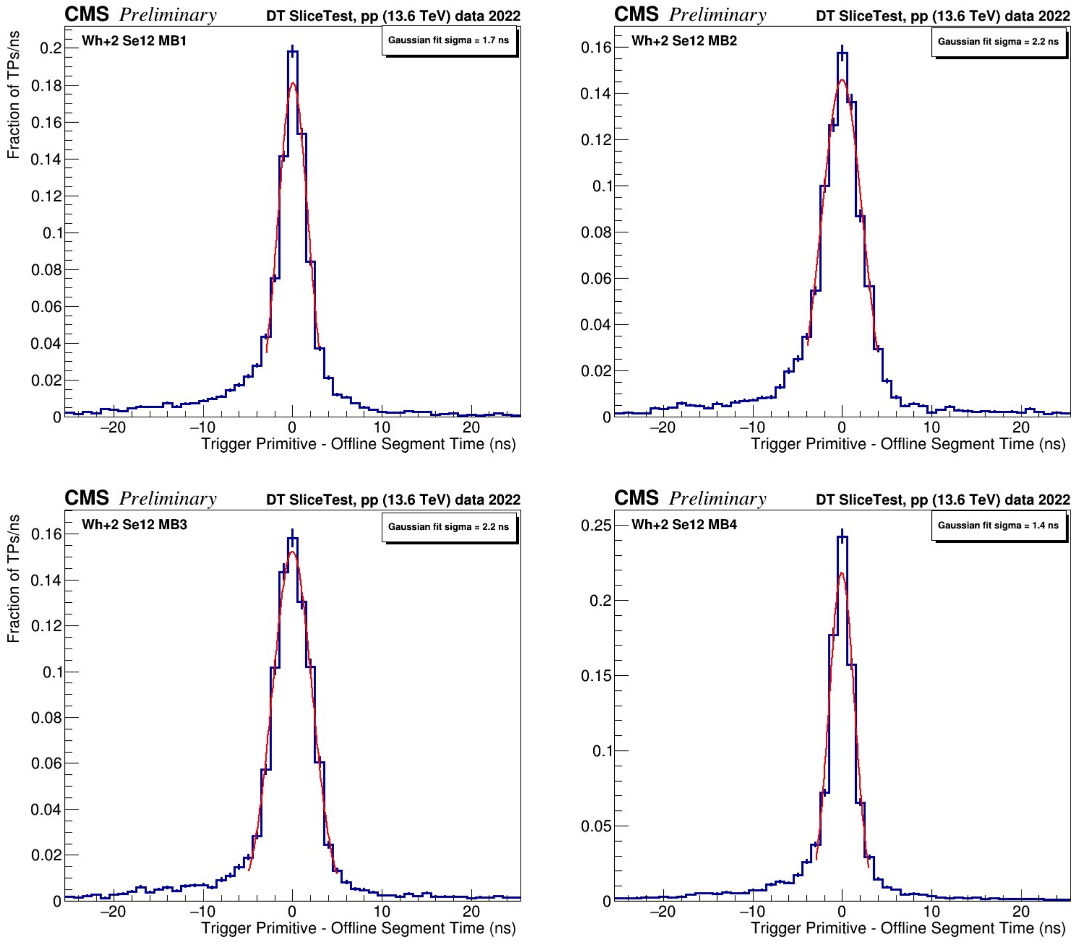

**Figure 1.** Difference between the Phase-2 trigger primitive's time and the offline reconstructed segment time in the four stations in sector Wh+2 S12 for a collision muon sample collected in the DT Slice Test.

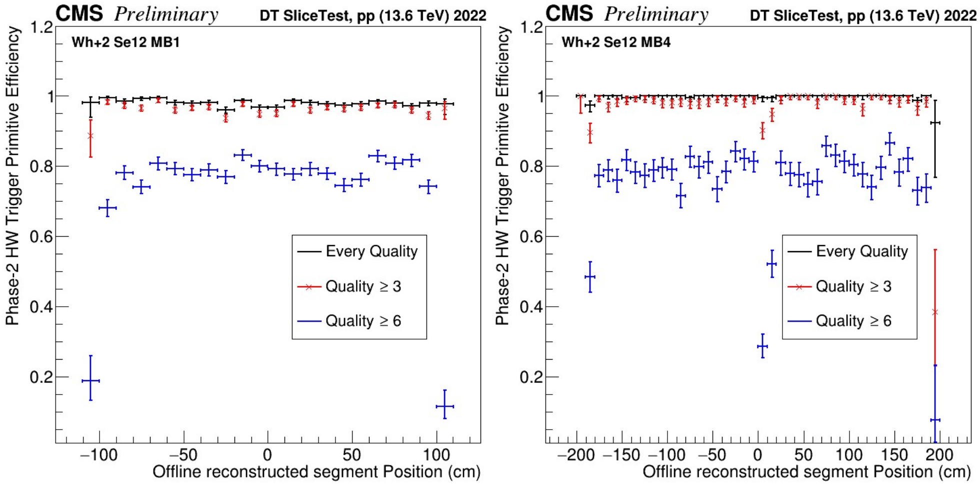

**Figure 2.** Efficiency of finding a Phase-2 Trigger Primitive in any BX with respect to the local position of the offline segment reconstructed out of hits detected by the Phase-1 system considering every primitive (black), primitives built with more than 4 hits in a Superlayer(red), and primitives with more than 6 hits, i.e. 3 or more hits per Superlayer (blue), for the DT Slice Test Wh+2 Se12 MB1 (left) and MB4 (right) data in 2022 collisions.

## 3. Performance Results

These results focus on the commissioning of the ST during early collisions in Run 3. Figure 1 shows the difference between the Phase-2 trigger primitive's time and the offline reconstructed segment time for a collision muon sample collected in the DT Slice Test setup. This distribution is a good metric of how well the AM trigger algorithm delivers online the performance now only available offline. The figure includes every quality (3-hit primitives in one Superlayer were included). A Gaussian fit to the core of the distributions, used as first order assesment of the performance, gives values < 2.2 ns. The ultimate performance was not achieved because the data were taken before the Phase-2 fine synchronization and Phase-1 full-precision offline calibration of the DT system were deployed. Moreover, the effect of the time of propagation along the wire is included offline in the three internal chambers, which is the reason for the better agreement in MB4, where this correction is not yet included ($\sigma$=1.4 ns), and in MB1 ($\sigma$=1.7 ns), where this correction is partially cancelled by the muon TOF from the IP.

The efficiency of the AM algorithm for finding a Phase-2 primitive in any BX is shown with respect to the local position of the Phase-1 offline segment in Figure 2 for several quality levels. No geometrical matching between the offline segment and the Trigger Primitive is required. Since one prototype backend board only generates primitives from one-half of the MB4 chamber, an efficiency drop is visible at the boundary (segment position=0) between the two regions due to edge effects. In the final system all channels from a chamber, in fact from more than a sector, will be processed by one backend board, so this effect will disappear.

## 4. Summary & Plans

We have run a Phase-2 parallel demonstrator system, a 1/60 sector-scale Slice Test, in collisions during Run 3 of LHC, which allows us to refine trigger algorithms under realistic conditions (radiation, magnetic field, background rate) and test main components also used in the final board, such as the PolarFire. With this setup we have measured the performance of the synchronization of the Phase-2 Slice Test chain, largely based on previous cosmics data taking. The Analytical Method algorithm deployed in the backend FPGAs is reaching a time resolution within 2 ns of that of the offline reconstruction, as promised by the new architecture. This rather satisfactory performance of OBDTv1-based Slice Test in collisions proves the feasibility of the chosen architecture.

We plan to deploy in CMS backend prototypes based on the ATCA standard that are now being tested on the surface. Final OBDT boards, now ongoing validation, will be installed in an appropriate technical stop during Run 3.

## Acknowledgments

These results are possible thanks to work by many colleagues from the CMS Muon DT group.

## References

[1] CMS Collaboration, *The CMS experiment at the CERN LHC,* JINST 3 (2008) S08004, doi:10.1088/1748-0221/3/08/S08004

[2] CMS Collaboration, *The Phase-2 Upgrade of the CMS Muon Detectors*, CERN-LHCC-2017-012, CMS-TDR-016,2017.

[3] A. Triossi et al., *Electronics Developments for Phase-2 Upgrade of CMS Drift Tubes* , TWEPP2018

[4] A.Triossi et al., *A New Data Concentrator for the CMS Muon Barrel Track Finder,* PoS(TIPP2014)412

[5] The CMS Muon group, *A Firmware-oriented Analytical Algorithm for Trigger Primitive Generation in the CMS Drift Tubes Detector for the High Luminosity LHC*. CMS NOTE-2022/009 submitted to Nucl. Instr. Meth. A.

[6] CMS Collaboration*, DT Slice Test 2021 Calibration Stability plots* DP/2022-01

[7] *Analytical Method (AM) for DT Trigger Primitive Generation in Phase 2. Performance results on simulation* DP/2021-28

[8] CMS Collaboration *Performance of the DT Phase-2 readout and local reconstruction with cosmic ray data from the DT slice-test* DP/2021-36.

[9] CMS Collaboration *Performance of CMS muon reconstruction in pp collision events at sqrt(s) = 7 TeV,* 2012 JINST 7 P10002