

OPEN ACCESS

The DAQ and control system for the CMS Phase-1 pixel detector upgrade

To cite this article: W. Adam *et al* 2019 *JINST* **14** P10017

View the [article online](#) for updates and enhancements.




IOP | ebooks™

Bringing together innovative digital publishing with leading authors from the global scientific community.

Start exploring the collection—download the first chapter of every title for free.

The DAQ and control system for the CMS Phase-1 pixel detector upgrade



Tracker Group of the CMS collaboration

E-mail: bora.akgun@cern.ch

ABSTRACT: In 2017 a new pixel detector was installed in the CMS detector. This so-called Phase-1 pixel detector features four barrel layers in the central region and three disks per end in the forward regions. The upgraded pixel detector requires an upgraded data acquisition (DAQ) system to accept a new data format and larger event sizes. A new DAQ and control system has been developed based on a combination of custom and commercial microTCA parts. Custom mezzanine cards on standard carrier cards provide a front-end driver for readout, and two types of front-end controller for configuration and the distribution of clock and trigger signals. Before the installation of the detector the DAQ system underwent a series of integration tests, including readout of the pilot pixel detector, which was constructed with prototype Phase-1 electronics and operated in CMS from 2015 to 2016, quality assurance of the CMS Phase-1 detector during its assembly, and testing with the CMS Central DAQ. This paper describes the Phase-1 pixel DAQ and control system, along with the integration tests and results. A description of the operational experience and performance in data taking is included.

KEYWORDS: Data acquisition concepts; Detector control systems (detector and experiment monitoring and slow-control systems, architecture, hardware, algorithms, databases); Modular electronics; Optical detector readout concepts



Contents

1	Introduction	2
2	System overview	2
3	Front-end ASICs	5
3.1	Readout chips	5
3.2	Token-Bit Manager chip	6
4	Optical components	7
4.1	Pixel Opto Hybrid (POH)	7
4.2	Digital receiver	8
4.3	Control	9
5	Back-end implementation	9
5.1	Phase-1 Tracker FEC	9
5.2	Phase-1 Pixel FEC	11
5.3	Phase-1 Pixel FED	13
5.3.1	DECODE Pixel FED firmware	14
5.3.2	BUILD Pixel FED firmware	14
5.3.3	Pixel FED data payload	17
6	System tests in the CMS Detector — Phase-1 Pixel Detector Pilot System	17
7	System tests in the laboratory	19
7.1	FED tester setup	20
7.2	The DAQ setup for high data rate tests	20
8	Pixel online software	21
8.1	Hardware access and supervisors	22
8.2	Distributed software architecture	23
8.3	Interface to the Detector Control System	23
9	Operation performance	23
9.1	Software recovery mechanisms	24
9.2	Periodic ROC resets	25
10	Conclusion	26
	Tracker group of the CMS collaboration	29

1 Introduction

The CMS collaboration has adopted the approach to rely on a highly granular pixel detector as key element for the reconstruction of charged particle tracks and interaction vertices. A detailed description of the CMS detector, together with a definition of the coordinate system used and the relevant kinematic variables can be found in ref. [1]. The description of track reconstruction with the CMS tracker can be found in ref. [2].

The original CMS pixel detector [3] featured three barrel layers and two forward disks on each end. It was operated during LHC Run 1 (2010–2012) and the first part of Run 2 (2015–2016), and was designed to record efficiently and with high precision the first three space-points of a charged particle track near the interaction region up to an instantaneous luminosity of $1.0 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, with colliding bunch crossings (BX) at a spacing of 25 ns. The original pixel detector would not have sustained a satisfactory performance given the luminosity conditions expected in LHC running after 2017 due to inefficiencies in the front-end readout chip (ROC), and because the maximum throughput rate for the data links of the innermost layer would have been exceeded.

The goal of the Phase-1 pixel detector upgrade project [4] was to perform an evolutionary upgrade with minimal disruption of data taking and reconstruction by keeping the pixel size, sensor, and readout architecture the same, while improving the performance through a higher rate capability of the ROCs and larger data transmission rate, more robust track reconstruction through the addition of a fourth barrel layer, and a third disk per endcap, as well as a reduced material budget. The Phase-1 pixel detector was designed to maintain a high tracking performance at luminosities up to $2.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, corresponding to an average of 80 simultaneous inelastic interactions per 25 ns spaced BX, (these interactions are referred to as ‘pileup’). The Phase-1 pixel detector, with modified data acquisition (DAQ) and control system, was installed during an extended year-end technical stop at the beginning of 2017. It is expected to deliver high quality data in the high luminosity environment of the LHC up to Long Shutdown (LS) 3, which is scheduled to start in 2024. Roughly 350 fb^{-1} of data will be collected until LS 3.

The Phase-1 pixel DAQ and control system has been developed using a combination of custom and commercial microTCA parts. Custom mezzanine cards on CMS-developed carrier cards provide a Front-End Driver (FED) for readout, as well as a Pixel Front-End Controller (FEC) for configuration, the distribution of clock, fast commands, and trigger signals, and a Tracker FEC for programming auxiliary electronics. The Tracker and Pixel FECs use the same hardware and differ in the firmware.

This paper describes the Phase-1 pixel detector DAQ and control system. Section 2 gives a system overview, section 3 describes the front-end ASICs, section 4 outlines the optical components and section 5 explains the back-end implementation. Sections 6 and 7 describe the Phase-1 pixel pilot system and laboratory tests, respectively. Section 8 explains the software used for the pixel detector operation. Section 9 provides an overview of the performance during operation.

2 System overview

The CMS Phase-1 pixel detector has three disks on both ends of the forward regions (FPIX) and four barrel layers in the central region (BPIX). An overview of the Phase-1 pixel detector DAQ and

control system architecture, including auxiliary components required to interface with the central CMS services, is shown in figure 1. The CMS Phase-1 pixel detector contains one type of sensor, bump bonded to 16 ROCs [5]. The active area of the module is $16.2 \times 64.8 \text{ mm}^2$. The pixel size has remained the same as in the original detector, $100 \times 150 \mu\text{m}^2$. The same $\text{n}^+\text{-in-n}$ technology as for the original detector is used for the silicon sensors. A high density interconnect (HDI) is glued on top of the sensor. The HDI provides signal and power distribution for the ROCs, and it carries the token-bit manager chip (TBM) and decoupling capacitors. The TBM chips are glued onto and wire-bonded to the HDI. The TBM chip is described in detail in section 3.2. They orchestrate the transmission of the data from the ROCs to the back-end electronics. The Phase-1 pixel detector features a fully digital readout system including new back-end electronics. The new ROCs with digital readout operate on a 40 MHz [6] clock¹ and have a 160 Mb/s serial output data stream. This stream is encoded and multiplexed by the TBM using a 4b/5b encoding scheme, to reduce the impact of bit-errors during transmission [7] and for DC balancing. The TBM outputs one or two 400 Mb/s data streams. A dedicated ROC was designed for the sensor modules of the innermost layer of BPIX (layer 1) to cope with the higher hit rates. Layer 1 sensor modules require two TBMs to manage the higher data rates, while all other modules have one TBM.

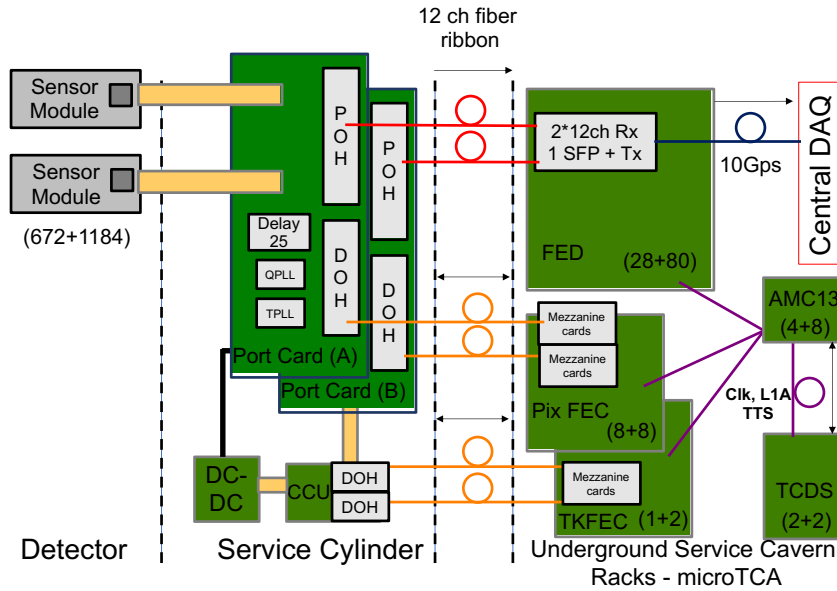


Figure 1. Overview of the microTCA DAQ and control system of the Phase-1 pixel detector. The numbers in parentheses indicate the numbers of the respective installed devices (FPIX and BPIX). Details are explained in the text.

The sensor modules are connected to the on-detector auxiliary electronics (portcards) via flex (FPIX) or twisted pair (BPIX) cables. The portcards are located in the 3 m long service cylinders, which also serve as mechanical support for all detector services (power, cooling and optical links) routed to and from the outside of the 6 m long pixel support tube. There are two different types

¹The LHC frequency and clocks derived from the LHC frequency are referred to as 40 MHz (and multiples of 40 MHz) in this paper.

of optical hybrids on the portcards: the Pixel-Opto-Hybrid (POH) and the Digital-Opto-Hybrid (DOH), which facilitate communication with back-end electronics via optical links.

The POH converts the electrical signal from the TBM to an optical signal and delivers it to the FED. The FED handles decoding and deserialization, and builds event fragments, which are sent to the Central DAQ Front-End Readout Optical Link-40 (FEROL40) card [8], the first stage of the CMS Central DAQ chain, via a small form-factor pluggable (SFP+) 10 Gb/s S-Link Express transceiver (Tx). There are 24 input channels per FED card; two receivers (Rx) with twelve channels each receive the data from the sensor modules. The FED receives clock, trigger and fast commands (called TTC [9] for timing, trigger, and control) from the CMS Trigger Control and Distribution System (TCDS) [10] via a CMS-custom module called AMC13 [11] and the microTCA backplane. The clock runs at the LHC frequency of 40 MHz. The FED also provides a trigger-throttle system [9] (TTS) signal, which is a 4-bit status word as defined in ref. [12], to the AMC13. The AMC13 forwards the TTS signals from all the FEDs in a crate to TCDS. The TTS signal indicates whether FEDs are ready to accept triggers or not, and if the event counters are still synchronized. The overall TTS state depends on the status of each FED. At a given moment a FED should either accept or block CMS level-1 triggers (L1A) [13]. The pixel detector DAQ is able to maintain event synchronization across all FEDs with this back-pressure system. A total of 108 microTCA Pixel FEDs are required to read out the Phase-1 pixel detector.

The AMC13 also propagates the received signals to the Pixel FECs, which distribute them to the sensor modules via the portcards. On the portcards these signals are decoded, after the opto-electrical conversion in the DOHs, by the Tracker Phase Locked Loop (TPLL) [14] and Quartz Phase Locked Loop (QPLL) [15] chips. The signals are then forwarded to the sensor modules on dedicated lines passing through Delay25 chips [16], which provide functionality to delay trigger signals, sent and received clock and data signals with a granularity of 0.5 ns. Each sensor module connected to a pixel-control link is identified by a unique, hardwired 5-bit hub address. The Pixel FEC is also responsible for programming the TBM and the digital-to-analog-converter (DAC) registers of the ROCs. A total of 16 microTCA Pixel FECs are required to operate the Phase-1 pixel detector.

Registers on the portcards, including Delay25 chips, and DC-DC converters [17], used for powering, are programmed by the Tracker FEC via the Inter-Integrated Circuit (I²C) interface and Parallel Interface Adapter (PIA) port, respectively, of a Control & Communication Unit (CCU) [18]. Several CCUs are arranged in a ring-like topology (referred to as *control* or *token ring*) via semi-redundant connections that carry clock and data signals. A total of 3 microTCA Tracker FECs and 12 token rings are required to control the Phase-1 pixel detector auxiliary electronics.

The number of optical readout links has increased with respect to the original detector from 448 to 672 for FPIX and from 1152 to 1696 for BPIX, resulting in a total of 2368 readout links. The first and second layer of BPIX use four and two links per sensor module, respectively, to cope with the higher occupancy and data rate. The third and fourth layer of BPIX, as well as the FPIX disks, use one link per sensor module.

3 Front-end ASICs

3.1 Readout chips

The ROC used in the original CMS pixel detector, PSI46 [19], was designed for hit rates of a few tens of MHz/cm², encountered at BPIX layer 1 for an LHC instantaneous luminosity of $1.0 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ with a 25 ns bunch spacing. This readout chip performed well during the data taking periods from 2008 to 2016. However, it showed expected inefficiencies when operated at higher data rates when the LHC started operating at instantaneous luminosities above the design value. In addition the innermost layer was moved closer to the beam line. Therefore, new pixel ROCs had to be designed: an evolutionary update of the original PSI46, the PSI46dig used in FPIX and BPIX layers 2, 3, and 4, and a dedicated ROC for BPIX layer 1, the PROC600, to cope with the exceptionally high rates of up to 600 MHz/cm².

The new readout chip evolved from the PSI46 ROC, keeping most of its characteristics: pulse-height readout, and 52×80 pixels organized in 26 double-columns of 2×80 pixels with common data transfer to latency buffers in the periphery outside the active pixel region. The digital Phase-1 pixel ROC (PSI46dig) is manufactured in the same 0.25 μm CMOS technology as the PSI46, and the overall layout and many building blocks remained unchanged. The two main improvements needed for the upgrade were larger data buffers and higher readout speed.

The double-column buffer sizes have been increased from 32 to 80 cells for the hits and from 12 to 24 cells for the time-stamps. Unlike the analog PSI46 ROC, the PSI46dig ROC outputs digital data for which an analog-to-digital converter (ADC) has been implemented in the chip. It is an 8-bit successive approximation register ADC running at 80 MHz. Digitized data are stored in a 64×23 bit First In First Out (FIFO), which is read out serially at 160 Mb/s. The 80 and 160 MHz clocks needed for the ROC operation are generated from the external LHC clock using a Phase Locked Loop (PLL) circuit.

During the trigger latency of the CMS experiment, currently 4.15 μs , the pixel hit data must be stored inside the ROC, and only data corresponding to triggered events are read out through the serial links. The internal transfer and buffer capacities of the ROC were designed to cope with rates up to 200 MHz/cm². The rates of data loss have been measured with high-flux X-ray tubes for pixel hit rates of up to 300 MHz/cm², and were found to be in excellent agreement with expectations based on detailed architecture simulations [20].

In addition to the higher rate capacity of the ROC, several other improvements with respect to the PSI46 have been implemented. An additional metal layer for power distribution was added, which allows a better decoupling of the power lines from the signal lines, resulting in an improved pixel response uniformity as well as lower noise and cross-talk. An optimized comparator reduces the time-walk from about 35 ns [5] to 15 ns [19], resulting in a reduction of the difference between the in-time threshold (within a time window of one clock cycle) and the time-walk independent absolute threshold from about 800 to 150 electrons. The above improvements reduce the effective operational threshold of the ROC from 3400 electrons in the original detector to 1700 electrons for the upgraded one. This is important when the amount of charge per hit starts to decrease after radiation damage to the sensors: a highly irradiated detector will slowly degrade in signal induced charge. With a lower threshold, the charge sharing among neighboring pixels can be exploited for position interpolation up to a higher integrated luminosity leading to a higher resolution.

Based on operational experience with the PSI46 ROC and irradiation tests, further optimizations of the internal biasing were made that extend the range of ionizing dose tolerated by the PSI46dig ROC, reducing the need to re-adjust DAC settings with increasing accumulated dose. The PSI46dig ROC performed well and without significant performance degradation after irradiation to up to 120 Mrad ($4 \times 10^{14} \text{ cm}^{-2}$ 24 MeV protons, at the irradiation facility in Karlsruhe), which is the maximal dose expected during LHC operations for FPIX and BPIX layers 2, 3 and 4. A detailed study on radiation tolerance of the PSI46dig ROC can be found in ref. [21].

Despite the improved performance of the PSI46dig, its architecture would lead to unacceptable data loss rates for the innermost BPIX layer, where pixel hit rates up to 600 MHz/cm² may be encountered. A dedicated chip (PROC600) was designed for layer 1, with a complete re-design of the double-column periphery. The PROC600 features a four times higher hit transfer rate of pixels to the end-of-column buffers, and dead-time-free buffer management. The former is achieved by changing from single pixel to 2×2 pixel cluster transfers and the implementation of a simpler, handshake-free protocol. A faster and more power efficient analog bus was developed for the pulse height transfers. The data buffer was modified considerably: PROC600 has a ring buffer with 56 buffer units, each containing a cluster base address plus four analog storage cells for the charge pulse heights. The readout is zero-suppressed in order to remove pixels in the cluster with zero measured signal amplitude. Only hits which are validated by a L1A are read out without stopping the acquisition of new hits into the buffer. This avoids an interruption of the data acquisition process in the double-column or overwriting of data, as is the case in the PSI46dig ROC.

Both ROCs have performed well in the 2017 and 2018 data taking. For the PSI46dig ROC all targeted improvements, i.e. low noise, lower threshold, and lower inefficiency at high rates, have been confirmed during data-taking. Some shortcomings have been observed for the PROC600, like a higher than expected noise hit rate and the rare loss of data synchronization in double-columns. These can be mitigated by operational procedures: the former by an increase of the in-time charge threshold for layer 1 to 3500 electrons, as compared to 1700 electrons used for other layers. These issues could partially be mitigated by operational procedures and have been addressed in a revised design of the PROC600, which will be used in the planned replacement of the innermost BPIX layer in 2020 during LS2.

Details on the operational performance of both ROCs are shown in section 9.2.

3.2 Token-Bit Manager chip

The Phase-1 pixel detector TBM is a radiation-tolerant integrated circuit that controls the readout of groups of ROCs. It replaces the original TBM [22]. The TBM chip is mounted as a bare die, wire bonded to the HDI that is glued on the sensor modules.

The principal functions of the TBM include the distribution of clock, L1As and fast commands as well as configuration data from the Pixel FECs to the ROCs. The TBM passes a token around the group of ROCs it controls to orchestrate the readout of data associated to a given L1A. The TBM keeps each arriving L1A on a 32-deep stack while waiting for the token to return if the token has not returned before the next L1A arrives. The TBM adds a header and a trailer to the data stream on each token pass.

The TBM has one or two cores that output serial data at 160 Mb/s. Two output data streams are non-return-to-zero-inverted (NRZI) and encoded with a 4b/5b scheme and multiplexed by a

block called the DataKeeper into a 400 Mb/s stream which is transferred to the FED. There are three versions of the Phase-1 pixel detector TBM (summarized in table 1). The TBM08, used in FPIX disks and BPIX layers 3 and 4, combines two groups of ROC data, while the TBM09 and TBM10, used in BPIX layer 2 and layer 1, respectively, combine the output of four groups of ROCs into two 400 Mb/s data streams. The TBM09 and TBM10 differ in their timing settings, which are optimized to match the PSI46dig and PROC600, respectively.

Table 1. Different TBM types.

	Groups of ROCs	Number of ROCs in each group	Number of 400 Mb/s channels	Detector part
TBM08	2	8	1	FPIX + BPIX L3, L4
TBM09	4	4	2	BPIX L2
TBM10	4	2	2	BPIX L1

The data format for Phase-1 sensor modules is as follows: TBM Header, followed by ROC Headers and pixel-level event information, followed by TBM Trailer. Event number and stack count are included in the TBM Header, ROC Headers are followed by column and row addresses of the pixels with hits and hit amplitudes, and the TBM Trailer includes the error information.

4 Optical components

The optical readout link starts at the electro-optic POH interface inside the detector and ends at the opto-electric receiver module interface on the FED. The data coming from TBMs are sent by the POH at a rate of 400 Mb/s. The control optical link system is based on the same components as used in the original pixel system: a DOH communicating bi-directionally with a FEC, which uses standard SFP transceivers. The readout and control links are shown in figure 1.

4.1 Pixel Opto Hybrid (POH)

The POH is a printed circuit board (PCB) mounted on the detector service cylinders. Figure 2 shows the POH4 (left) used in BPIX and the POH7 (right) used in FPIX. The optical characteristics of the two variants are the same. The overall system requires 424 POH4 and 96 POH7.

The design of the POHs uses the Transmitter Optical Sub-Assembly (TOSA) component provided by the Versatile Link project [23]. The POH receives electrical signals from the TBM and converts them into optical signals to be transmitted to the back-end receiver on the FED installed in the counting room, about 65 m away from the detector. Each POH houses single-mode Fabry-Perot laser TOSAs operating at 1310 nm, as well as Digital Level Translators (DLT) and Linear Laser Drivers (LLD) [24]. The DLT chips convert the signals received from the TBM to levels compatible with the LLD and introduce a gain and an offset to the input signal. The LLD chips drive the laser TOSAs. They pre-bias the lasers at their working point and modulate them with a current proportional to the input signal. The modulation gain and pre-bias currents at the LLD are controlled through an I²C interface. The POHs are used to transmit balanced digital signals at a maximum bit rate of 400 Mb/s. A typical output optical eye diagram is shown in figure 3.




Figure 2. Photographs of a fully assembled POH4 (left) and a POH7 (right). The differences are the number of transmitter channels, four in the case of the POH4 and seven in the case of the POH7, and the input matching that adapts to the signal cables of the BPIX and FPIX system, respectively.

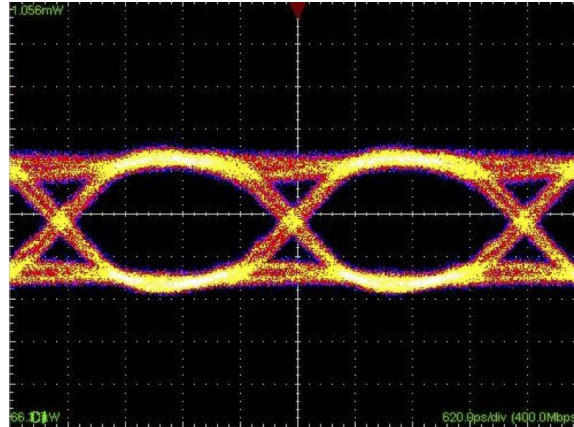


Figure 3. Typical output eye diagram (open and symmetric) measured on a POH. The horizontal scale is 620 ps/div and the vertical scale is 110 μ W/div.

4.2 Digital receiver

The digital receiver module used on the upgraded microTCA FEDs is a commercial component. Since the lasers mounted on the POHs emit light at a wavelength of 1310 nm it was critical to identify a receiver module based on an InGaAs photodiode. Typically, high-density multi-channel receivers are based on GaAs photodiodes that operate with light at a wavelength of around 850 nm and are not sensitive to a wavelength of 1310 nm. One manufacturer was identified being able to produce fully qualified receiver modules [25] with 12-way arrays of InGaAs photodiodes. These are integrated in pairs on an Field Programmable Gate Array (FPGA) Mezzanine Card (FMC) board to be mounted on the FEDs. The receiver modules have a diagnostic feature that allows the DC photocurrent to be measured on each input channel individually. This was used during

initial detector checkout to spot problematic fiber connections. Figure 4 (left) shows a picture of a Receiver-FMC (Rx-FMC), with an SFP+ transceiver attached to it for the Central DAQ link.




Figure 4. (Left) An Rx-FMC with 24 optical input channels feeding two 12-channel optical receivers that are optimized for the 1310 nm 400 Mb/s signal, and a SFP+ 10 Gb/s transceiver for S-Link Express to send data to the CMS Central DAQ. (Right) An FMC equipped with low-speed compatible (80 Mb/s) optical transceivers, with 8 SFPs per FMC.

4.3 Control

The optical link system used to control the Phase-1 pixel detector uses the same components as the previous detector system [26] at the front-end. The DOHs, located in the service cylinders, transmit the control signals between the Pixel and Tracker FECs and the detector front-end. The back-end components that are housed in the FECs are standard single-mode SFP modules rated for 1–2 Gb/s data rates. These SFPs plug into custom-designed FMC boards, shown in figure 4 (right), that are mounted on the FECs.

5 Back-end implementation

The design of the back-end electronics for the Phase-1 pixel detector is based on microTCA modular electronics [27]. A microTCA carrier hub (MCH) card is used as communication interface between the microTCA electronics and the local area network (LAN). The microTCA backplane is used to distribute clock, trigger and fast commands that are received from the TCDS via the AMC13.

The FC7 microTCA FMC carrier [28, 29], was selected as the platform for the new digital FED and the Pixel and Tracker FECs. As shown in figure 5, the FC7 is a full-size, double-width Advanced Mezzanine Card (AMC) holding a Xilinx Kintex 7 FPGA [30] and offering two low-pin-count compatible (LPCC) FMC slots.

5.1 Phase-1 Tracker FEC

The Tracker FEC is responsible for programming the auxiliary detector electronics, like the DC-DC converters, which is independent from the control of the sensor modules. Each Tracker FEC controls a control/token ring via semi-redundant connections that carry clock and data signals. The



Figure 5. Front (left) and back (right) side of the FC7 with the Xilinx Kintex 7 FPGA and two LPCC FMC connectors.

control is done via a token-ring protocol. For the CMS Phase-1 pixel detector there are four control rings each for FPIX and BPIX.

The Tracker FEC firmware is designed to implement four control ring firmware blocks (CTRL_RING) independently from each other. Each CCU control ring is addressed by one control ring firmware block. The firmware is link compliant with the CCU communication protocol specified for the original detector [18, 32] and access compliant with the control software of the original detector. The firmware can be controlled and monitored over an 1-Gb/s Ethernet/IPBus [33] link via the AMC backplane and, unlike the other parts of the back-end electronics, the firmware does not need to be synchronized with the LHC clock. Four signals are used in every ring: two for data transmission from the FMC via the DOH to the CCU ring, transmitting clock and data, and two for data reception of the FMC via the DOH from the CCU ring, returning clock and data. Two DOHs are available per ring. Four SFPs and eight optical fibers need to be plugged in the FMC in order to connect a CCU ring.

An example topology with all the connections is shown in figure 6, which considers a CCU ring composed of two DOHs and five CCUs. The last CCU is a spare/dummy, which is needed in order to close the redundant path to output B of the control ring.

The commands are transmitted from the control ring firmware block (the master) via the TX line (A or B) to the appropriate CCU of the control ring (the slave). The CCUs are distinguishable by their own defined addresses. A ring-type topology is configured as a standard computer LAN connecting the control ring firmware block to CCUs and the CCUs between themselves. Two types of commands can be executed from the control ring firmware block: register write commands and register read commands. The CCU executes the I²C transactions addressed to the appropriate device from the initial command received.

By default, an idle pattern is sent to the ring on the TX line by the control ring firmware block. The control ring firmware block also verifies that the ring is well initialized at startup and just before transmitting a command, by injecting a token frame to the ring. The ring is well established if the returned token frame matches the token frame injected. In any case, a status register is updated so that the control software (section 8.1) knows the status of the ring in real-time.

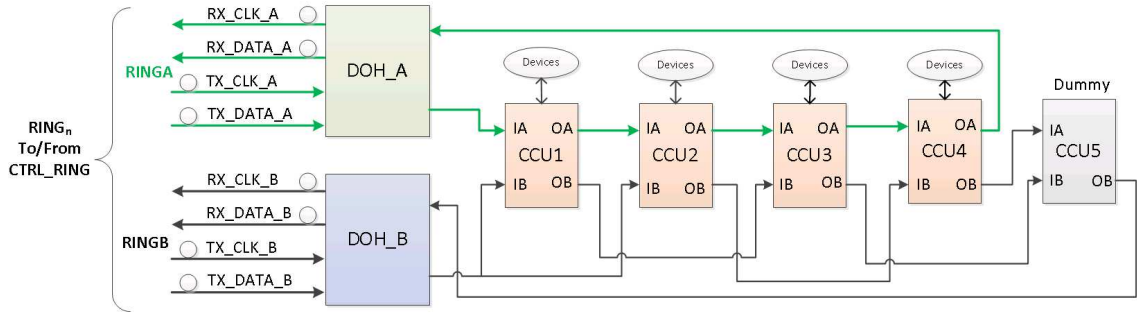


Figure 6. An example topology with two DOHs and five CCUs (the CCU5 is a spare/dummy). Ring A is the primary ring, used by default. In case of a failure, either of DOH_A or of any single CCU, the device can be bypassed by switching to Ring B.

5.2 Phase-1 Pixel FEC

The Pixel FEC is responsible for distributing clock, trigger, and fast commands to the sensor modules, as well as for programming the DAC registers of the ROCs and registers of the TBM chips on the sensor modules. A total of eight DOHs are connected to a Pixel FEC.

A block diagram of the Pixel FEC at the board level is shown in figure 7. The firmware was designed to provide 1-Gb/s Ethernet/IPBus services via the AMC backplane. Pixel FEC registers and channel input and output FIFOs are interfaced via Ethernet through the IPBus.

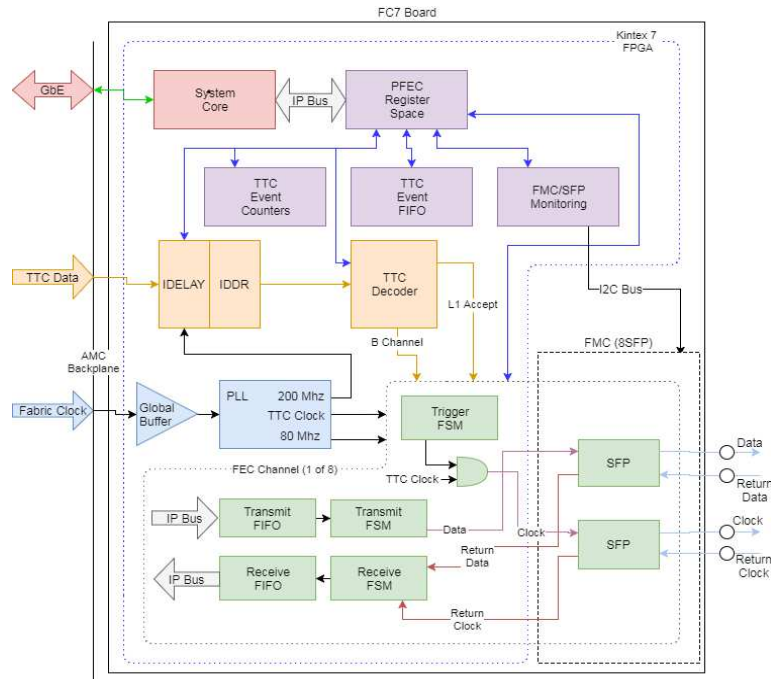


Figure 7. Block diagram of the Pixel FEC. GbE represents gigabit Ethernet, PFEC represents Pixel FEC and B channel is an interpreted fast command. The rest of the abbreviations are explained in the main text.

The LHC 40 MHz input clock (fabric clock) is sent to a PLL to produce a TTC Clock at the same frequency, as well as 80 MHz and 200 MHz clocks for various other subsystems. TTC information is received through IDELAY and IDDR logic blocks and then processed in a hamming decoder block. The outputs of the TTC decoder block are the L1A and fast commands on an 8-bit bus. Pixel-related fast commands include ROC reset and TBM reset and other commands to e.g. reset event counters or clear buffers. Registers in the Pixel FEC register space count how many Pixel-related fast commands are decoded and a FIFO can capture all the TTC events. A trigger finite state machine (FSM) receives the fast commands and encodes the appropriate bit pattern into the TTC clock for transmission to the SFP as the module clock. The L1A, ROC reset, and TBM reset signals can also be encoded in the TTC clock by setting appropriate bits in the Pixel FEC register space.

Eight Pixel FEC channels are instantiated in the FC7's Kintex 7 FPGA. Programming data are loaded into a 16 kB transmit FIFO to be used by the transmit FSM. Either a *Send Data* bit is set in the Pixel FEC register space or the TTC Send Data command makes the transmit FSM undergo a transition using the configuration data stored in the transmit FIFO. An 8b/10b encoded data stream is generated and transmitted to the SFP. The TBM's hub and port addresses along with the number of bytes transmitted in the command are stored in the Pixel FEC register space.

During lab tests and during the initial phase of the detector operation, commands to program the sensor modules were composed by fetching configuration data stored on a remote server, and were loaded into the transmit FIFO, to be sent sequentially for all sensor modules included in the detector configuration. This procedure was relatively time consuming, and not practical in normal operation. Since Spring 2018 a new way of programming the sensor modules has been implemented using the feature of storing the configuration data in the FC7 DDR3 SDRAM. This has two benefits. First, it allows the storing of configuration data locally on the FC7 cards, so during re-configuration there is no need to form the commands again by fetching detector configuration data from the remote server. Secondly, it allows sending configuration commands in parallel, reducing the total time to program the sensor modules by Pixel FECs from 30 seconds to 2 seconds.

The DDR3 memory is partitioned into segments for each of the Pixel FEC channels. One segment is for general calibration purposes, and groups of 4 segments, each used for 28 sensor modules, are used to store TBM settings, two sets of DAC settings for individual ROCs, and settings to trim and mask individual pixels. Each memory segment is assigned a bit used to steer which memory segments are addressed to transmit their commands during a send command.

The data stream returning from the sensor module is parsed by the receive FSM. The clock for the receive FSM is the returned clock from the sensor module. Data reception begins with a start condition ('1's for eight clock cycles).

Data to the same hub/port address can be continuously transmitted until the data are exhausted. Once transmission to a hub/port address is complete the transmit FSM waits for the receive state machine to confirm reception of the command before proceeding to the next hub/port command.

Because the exact optical fiber lengths are unknown to the Pixel FEC and fiber lengths can add delays of several hundred nanoseconds between the data and clock leaving the Pixel FEC and the received data and clock at the sensor modules, a simple handshaking between transmit FSM and receive FSM is implemented to prevent metastability issues that might arise if the data and clock lines are not synchronized. Start of transmission is indicated to the receive FSM so a timeout

counter can be started, limiting the time the receive FSM waits for the start of a transmission to 100 BXs (2495 ns)




Figure 8. Efficiency of transmission as a function of the delay of returned data (RDa) and sent data (SDa). The size of the blue dots is proportional to the transmission efficiency, the largest dots being 100% efficient. The red square is chosen as the working point since it is at the center of the area with 100% transmission efficiency. The step size is 0.5 ns.

The (data and clock) transmission paths between the Pixel FEC and the sensor modules are synchronized by cycling through Delay25 phases of the sent and received data at the portcard and plotting the successful transmissions, as shown in figure 8. The center of the resulting area is used as the calibrated delay for the sent and returned data. The Pixel FEC has been shown to have ± 7.5 ns of phase margin between sent clock and sent data, and ± 6 ns of phase margin between returned clock and returned data signals.

5.3 Phase-1 Pixel FED

The Phase-1 Pixel FED consists of an FC7 board with a Rx-FMC. The Rx-FMC is a mezzanine containing two 12-channel optical receivers that collect signals from the sensor modules, and one SFP+ for data transmission to the CMS Central DAQ, as shown in figure 4 (left). One FED can read out 24 data streams of 400 Mb/s, and transmit output data at 10 Gb/s. The FED can also emulate and transmit data and run without a detector.

The FED firmware consists of two parts. The first part (DECODE) handles decoding of the incoming data. The second part (BUILD) builds pixel events and sends them to the CMS Central DAQ.

5.3.1 DECODE Pixel FED firmware

The main task of the DECODE part of the Phase-1 pixel FED firmware is to decode the NRZI and 4b/5b encoded 400 Mb/s input signals, and to split the multiplexed TBM channels into two data streams. The decoding converts 24 multiplexed data streams of 400 Mb/s to 48 TBM core data streams of 160 Mb/s. A TBM core data stream starts with a TBM Header, followed by an 8-bit event number. This is followed by ROC Headers which indicate the beginning of pixel data. A TBM Trailer followed by 16-bits of status information terminates the data stream.

The DECODE part of the Phase-1 pixel FED firmware (figure 9) was designed to automatically find the best sampling point for the incoming 400 Mb/s signal and to do continuous sampling phase finding without disturbing data integrity. The optical receiver output, which carries the 400 Mb/s data stream, drives a differential input buffer of the FPGA. The negative output of this buffer is used as a copy of the incoming data stream to perform sampling phase finding and phase correction calculations.

The DECODE firmware detects TBM Header, TBM Trailer, and ROC Headers and writes pixel data in so-called TBM FIFOs. Several checks are included to keep data integrity as high as possible. Therefore, not only does the TBM Header marker have to be identified to start a data packet, but also the beginning of the next marker (ROC Header marker or TBM Trailer marker depending on the data packet) is included to validate the start sequence. ROC Headers are only allowed within the expected delay after the arrival of a TBM Header. The number of ROCs is counted and an error is reported if the count does not match the expected number from the TBM type. Furthermore, veto conditions for Header and Trailer arrival times and controllers that check the sequence of Headers, Trailer, and pixel data are added to avoid corrupted data packets. At this stage the TBM 4-bit words, which are the outcome of NRZI decoding, are combined with a 4-bit qualifier marker. This allows the following stage to identify these words as Header, Trailer or pixel information.

In order to keep the event throughput per unit time roughly constant the DECODE firmware truncates the data volume by terminating the TBM core data stream when the TBM FIFOs get filled to a programmable value or too long payloads are sent from layer 1 modules. The data stream from the DECODE firmware block to the BUILD firmware block contains an overflow error type in this case.

The bitwise signal of the events might be altered which causes errors in the FED. The FED has multiple error counters, which can count independently for each channel.

The data streams are forwarded to the BUILD firmware block using a 36-bit wide interface with the possibility of clocking data out at 40, 80, or 160 MHz. The default frequency for clocking out data is 160 MHz.

For debugging purposes, the DECODE part of the firmware has fiber-specific spy FIFOs which store the incoming 5-bit symbols and the decoded 4-bit data words. To monitor the data transfer to the BUILD firmware part, additional spy FIFOs for every TBM channel are implemented.

5.3.2 BUILD Pixel FED firmware

The BUILD FED firmware was designed to handle the data readout of 48 data streams coming from the DECODE FED firmware, transmit data to the CMS Central DAQ via the S-Link Express interface, and communicate with the TCDS system for the TTC and TTS interfaces for synchro-

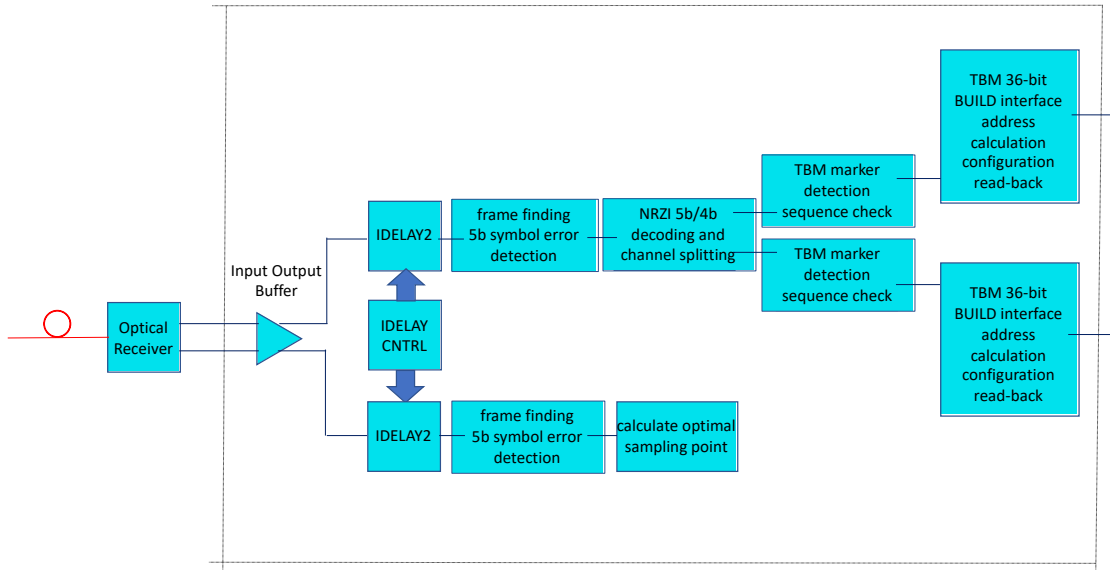


Figure 9. Block diagram of the Phase-1 pixel FED DECODE firmware.

nization. The major challenges are the high data rate, the asynchronous readout with large payload variations between streams and from event to event, maintaining the synchronization, as well as the exception and error handling. Exceptions can occur due to corrupted data provoked by an SEU or sensor modules not sending coherent data. The 1-Gb/s Ethernet/IPBus communication is used to read out information about these exceptions during physics data taking. The block diagram of the BUILD FED firmware is shown in figure 10.

The READOUT firmware block encodes the data coming from the DECODE block, merges all the data fragments, detects and marks exceptions and transmits the merged data to the CMS Central DAQ. The firmware block uses separate FIFOs for L1As and pixel data and, in order to increase the data throughput, it drains the pixel data FIFOs in parallel. The parallel data draining structure allowed to exploit almost the complete potential bandwidth of the used hardware, as discussed in section 7.2, and ensured safe running for the conditions expected during LHC Run 2 and Run 3. The READOUT firmware block also computes and transmits its own individual TTS state which is computed based on the filling level of the L1A and pixel data FIFOs, synchronization loss detection and received TTC commands.

The TTS states used in the READOUT firmware block in the BUILD firmware are ready (RDY), busy (BSY), and out-of-sync (OOS). The other TTS states defined in [12] are not used. The TTS firmware block handles the transitions of the TTS states. The TTS states and transitions are shown in figure 11. After the system is configured, the TTS state is RDY. When FIFOs are almost full, back-pressure is applied to avoid an overflow and a subsequent loss of synchronization. The goal of the BSY state is to rapidly veto the arrival of new triggers. The veto is not instant because of non-negligible propagation time; new triggers are still accepted before the back-pressure is effective. An OOS condition due to either consecutive timeouts (a timeout occurs when a FED channel does not receive data for a programmable time) or consecutive event number mismatches

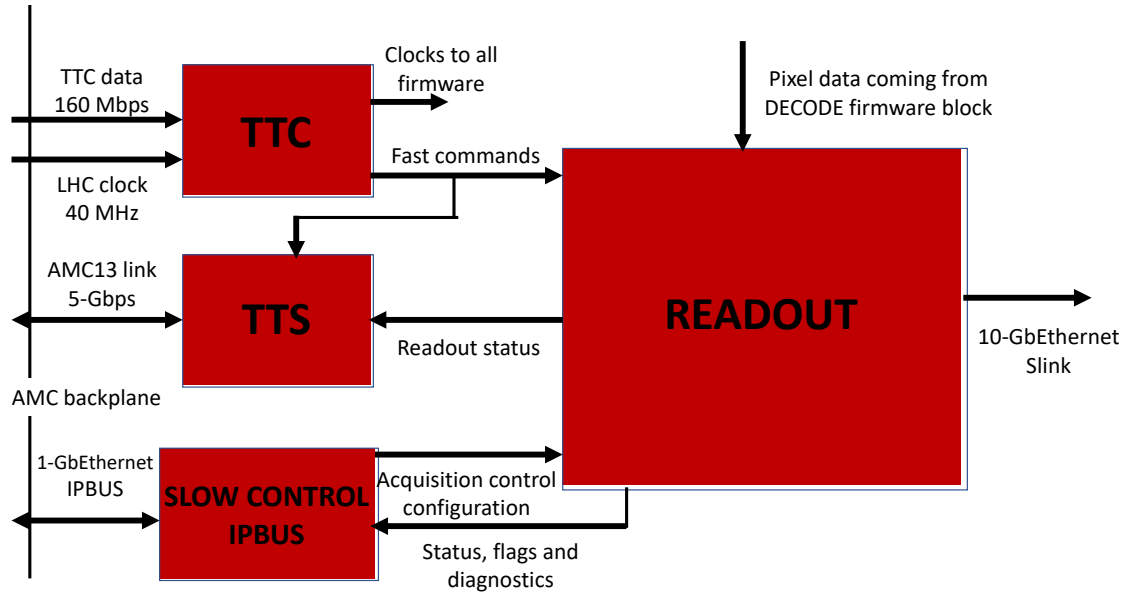


Figure 10. Block diagram of the BUILD FED firmware.

can be triggered at any moment in any TTS state. In order to reestablish synchronization and empty FIFOs, a resynchronization command is propagated from TCDS to all FEDs. The TCDS command is interpreted as a resynchronization sequence (RESYNC). The firmware was written to accept two types of RESYNC commands: global or private. In the global case, which is CMS-wide, both front-end and back-end electronics receive the same command. In the private case, only the back-end electronics receive the command without the event number reset (EC0).

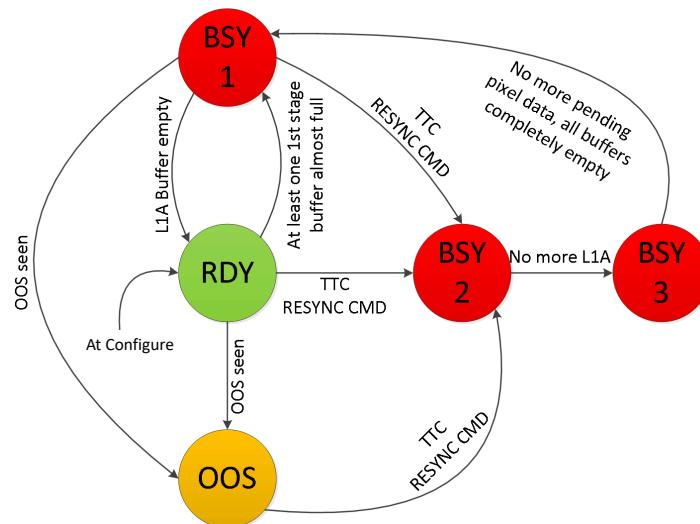


Figure 11. Block diagram of the FED TTS states and transitions. The BSY1, BSY2 and BSY3 are different FSM nodes which all have TTS state BSY.

5.3.3 Pixel FED data payload

The sensor modules transfer zero-suppressed data to the Pixel Detector DAQ system via 2368 optical fibers. The average number of hits in a sensor module decreases with its radial distance from the interaction point. Figure 12 (left) shows the average number of pixel hits per event for all channels. The distribution is uniform for the outer layers in BPIX and in FPIX, while the average number of received hits per event has a large spread for the innermost layer. The twelve channels of each receiver have been chosen from sensor modules of different layers and at different z-coordinates in order to balance the data processing load on the FEDs. Most FEDs take two of these fiber bundles as inputs. Figure 12 (right) shows the distribution of average hits per fiber bundle (ribbon) and per FED. The data load is distributed equally across the fiber bundles and thus also across the FEDs, with the FPIX FEDs (FED number > 96) receiving somewhat higher data rates on average than the BPIX FEDs (FED number ≤ 96). At a pileup of 46, for BPIX the average data payload is 10 hits per event per fiber, and thus lowest, for fibers connected to layer 4 modules. Layer 2 and 3 modules have an average data payload of 15–20 hits per event per fiber. Data payload for layer 1 modules fluctuates between 20 and 40 hits per event per fiber. For FPIX the data payload is 16–18 hits per event per fiber for outer ring modules and 35–40 hits per event per fiber for inner ring modules. The average throughput for a BPIX FED at a pileup of 46 at 100 kHz of triggers is approximately 1.7 Gb/s. As discussed in section 7.2, the data throughput was not a limiting factor during LHC Run 2 and will not be a limiting factor for Run 3.

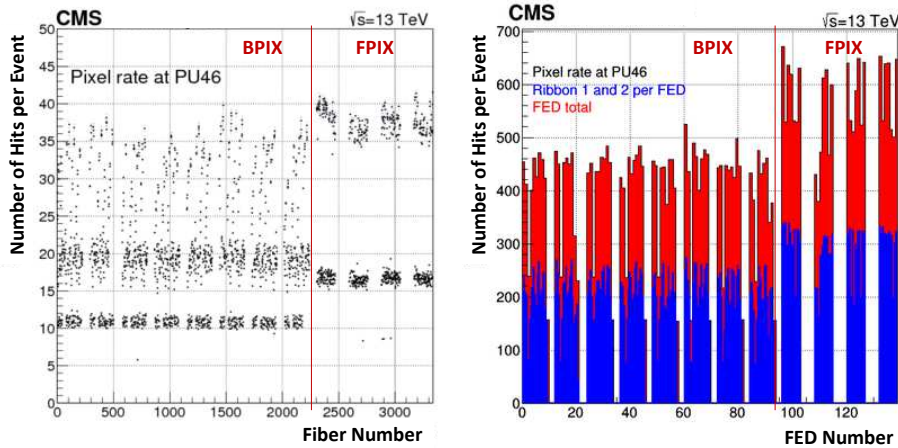


Figure 12. (Left) Data payload versus fiber number and (right) data payload versus FED number for an example CMS run in 2017 with a pileup of 46. For each sub-detector (BPIX and FPIX), the fibers are mapped to FEDs so that the data payload is distributed as evenly as possible, as shown in the right plot.

6 System tests in the CMS Detector — Phase-1 Pixel Detector Pilot System

In order to be well prepared for a short commissioning period during the extended year-end technical stop at the beginning of 2017 and to take advantage of the lengthy access to the original detector possible during LS1, a pilot system [34] was built with eight prototype Phase-1 sensor modules.

The pilot system was installed in 2014 in the available space in the original FPIX half cylinders (figure 13). A prototype microTCA FED system was used to read out the pilot system. The motivation for installing the pilot system was to learn how the readout, control, and offline systems perform in the CMS environment and to start integration with the CMS DAQ.




Figure 13. (Left) Two pilot sensor modules. (Right) Pixel half cylinder with pilot sensor modules installed on the third half disk, and pixel plaquettes of the original pixel detector installed on the first two half disks.

The pilot system was commissioned before installation at CERN using a test stand running a standalone test software. Calibration procedures for the pilot detector implemented in the online software were validated after the installation in CMS. During the pilot system tests before installation it was observed that the prototype FED was not able to correctly decode the data at high trigger rates. The issue was traced back to two separate sources: an asymmetric eye diagram due to the TBM design, and jitter on the Phase-1 portcard. While an asymmetric eye diagram could be accepted for the pilot system, new versions of the TBM were designed for the final Phase-1 sensor modules that were built in 2015 and 2016. In order to address the jitter on the portcard, an external QPLL chip had to be added in between the TPLL and Delay25 chips on the pilot portcard. Figure 14 shows asymmetric eye diagrams for one of the pilot modules before and after QPLL installation. For the final Phase-1 portcards, the design incorporated the QPLL chip directly on the PCB.

After the installation in CMS, the pilot system was used to validate the FED firmware. Both the DECODE and BUILD firmware blocks were iteratively improved during pilot system operations within CMS and the DECODE firmware block was finalized.

Six out of eight pilot sensor modules were successfully used for data taking. Clusters are formed from signals from individual pixels by the reconstruction software and the hit position is determined by their barycenter. Figure 15 (left) shows the measured cluster positions projected onto the transverse plane. Figure 15 (right) shows the expected hits, which are derived from extrapolated tracks that are reconstructed in the FPIX detector (without pilot detector hits). As the clusters from the pilot system were not included in the track reconstruction there are no expected hits close to the center. This effect is visible in figure 15 (right). In order to discard the fringes of ROCs, where the uncertainty in the track extrapolation is large, fiducial regions are defined. These are visible as rectangular shapes in figure 15 (right).

Operating the pixel pilot system during the years 2015–16 within CMS, reading out events synchronously with other CMS sub-detectors and reconstructing tracks, provided valuable experience and enabled an early start for the modifications that were required for the integration of the Phase-1 pixel DAQ.

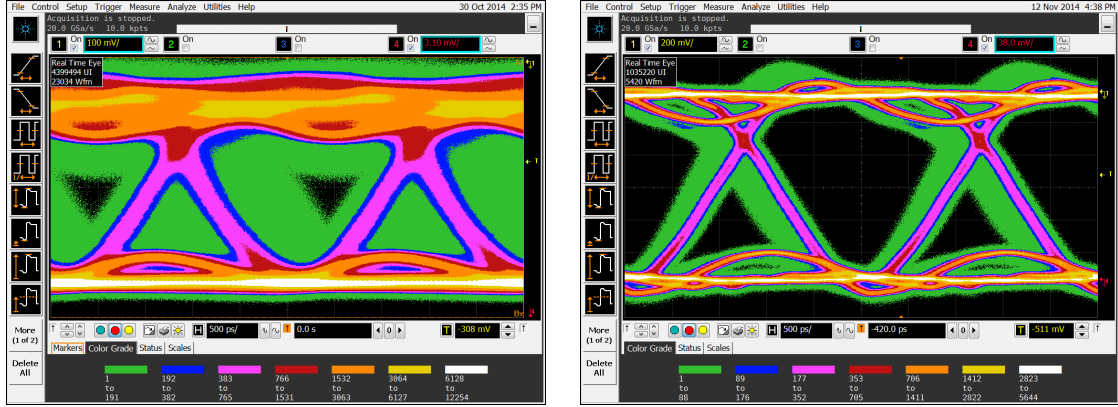


Figure 14. Eye diagrams for one of the pilot modules before (left) and after (right) QPLL installation. The amount of jitter decreased significantly after QPLL installation.




Figure 15. (Left) Cluster positions in each pilot sensor module used in data taking and (right) expected hits in the transverse plane in the CMS coordinate system. Color coding identifies individual sensor modules.

7 System tests in the laboratory

Small scale systems were used for development and testing of final detector parts, which advanced development and uncovered errors and issues well ahead of the final system installation. There were three integration centers using microTCA back-ends: Fermilab, the University of Zurich (UZH), and CERN. In addition, there were test stands at HEPHY in Vienna, IPHC in Strasbourg, and Cornell University for firmware and software development and testing. At Fermilab, qualification of the FPIX detector was performed [35] before shipping it to CERN for installation in CMS. At UZH [36] the focus was on testing the optical components and electronics on the BPIX service cylinders, and on the integration tests for the BPIX detector. At CERN, emphasis was on DAQ hardware testing and integration and on testing firmware before deployment for the Phase-1 pixel detector [37]. Functionality tests were also performed on detector components upon arrival at CERN. A so-called “soak test” facility was set up at CERN to validate all DAQ back-end components before installation

in the CMS service cavern. A rack layout identical to the final setup in the cavern containing all of the production parts, power modules, AC-DC converters, crates, service boards, as well as FEDs and FECs was operated for several weeks before installation. The soak test included regular firmware upload, and power cycling of FEDs and FECs.

To facilitate the development and the validation of the pixel FED firmware a dedicated tool (FED tester) and a setup for high data rate tests were developed. These are described in the next two sections.

7.1 FED tester setup

The FED tester, a data emulator, was designed for the Phase-1 pixel detector upgrade based on the gigabit link interface boards (GLIBs) [38] combined with the same FMC as used on the FECs to transmit the emulated data. FED tester customization enables consistent tests of the FED firmware from version to version. Custom firmware and software was developed to emulate the data bit stream from sensor modules with different TBM and ROC types. Once it receives a trigger emulated data streams are generated and sent: a TBM Header, ROC Headers, pixel hit data, a TBM Trailer, and 16 bits of status information in every data stream. The software is able to generate data patterns in the FED tester framework and can validate the output from the FED the data are sent to.

The GLIB firmware is able to independently emulate 16 TBM core data streams. Three GLIB boards are used to completely fill one FED, and optical splitters can be added in order to feed multiple FEDs in parallel. Each group of two TBM core data streams is then multiplexed, and NRZI and 4b/5b encoded before being transmitted.

The content of the FED error counters are readout in the FED tester framework to confirm that the count for each error is accurate. Event generation can be done with a fixed data size, where every event is the same, or in SRAM mode, where the event size and pixel location can be programmed via software.

The SRAM of the GLIBs is a software loadable memory that can be accessed by the event readout framework. There are two separate memory locations that each hold approximately 8.4 MB of data. The first SRAM is designated to hold the distributions of hits per ROC. The second is designated to hold the emulated pixel locations for each hit. The reading of SRAM memory is driven by a 160 MHz clock. Since each GLIB can emulate 16 independent channels, there are 32 different processes which must occur to emulate an event. The first SRAM only needs to be read once per event, the second SRAM needs to be read any time a pixel hit information needs to be sent. The SRAM readout can be performed at trigger rates expected in CMS.

7.2 The DAQ setup for high data rate tests

Prior to installation the DAQ system was qualified at small scale but with a complete chain of DAQ hardware, for the highest expected data rates from the sensor modules. A microTCA crate with five FEDs was connected to the CMS Central DAQ, with the FEDs emulating data patterns. A FED tester was also installed in the crate, six optical splitters were used to feed the FEDs with the FED tester output. The FED tester output and internally emulated FED data were used at high trigger rates to qualify the pixel DAQ system and the interface to CMS Central DAQ. Clock and trigger signals were supplied by the TCDS system, as in the production DAQ system. The DAQ setup was used to




Figure 16. Data throughput (solid line and left y-axis) and trigger rate (dotted line and right y-axis), as measured when the system is driven by the FED tester. The FED can handle a trigger rate of 100 kHz for up to 48 emulated hits per TBM core data stream. The throttling of the trigger rate is caused by back-pressure in the FED. The blue (red) triangle is the throughput for a simulation with a pileup of 70 (130).

develop configurations to interface with the TCDS system, and to study the robustness of the TTS state transitions and the time spent in TTS states BSY and OOS under different conditions. It was also used to optimize the AMC13 configuration for the pixel use-case, and to study the propagation of the TTC commands from TCDS via the AMC13 to the FEDs. It continues to be used as a test bench to test new FED firmware releases, before their deployment in the production system.

It is possible to check the data sizes through the S-Link Express link of the FED using the high rate test setup. When the event sizes are too large the trigger throttling limits the maximum data throughput. The throughput is tested by using fixed-size data and SRAM data, to allow for more realistic conditions. The maximum pileup during LHC Run 2 was approximately 60 (less than 16 hits/TBM core data stream for an average BPIX FED).² For up to 48 hits/TBM core data stream the FED can run at 100 kHz. Throttling of triggers starts at 56 hits/TBM core data stream and data throughput reaches approximately 7.5 Gb/s (at a trigger rate of 72 kHz). This shows that the data throughput was not a bottleneck during LHC Run 2 and is not expected to be a bottleneck during LHC Run 3. Figure 16 shows the throughput and trigger rates the FED can handle when different numbers of hits/TBM core data stream are generated by the FED tester.

8 Pixel online software

The Pixel Online Software (POS) is a collection of applications that control the front-end and back-end hardware of the CMS pixel detector. The software collection is written in C++ and is based on the CMS online software framework XDAQ [39].

²The average pileup during LHC Run 2 was approximately 35.

8.1 Hardware access and supervisors

For each type of back-end electronics board in the pixel system a corresponding controller class exists. The controller provides the software interface to the hardware and allows access to the hardware functionality within the POS. It uses the CACTUS framework [40], which provides a hardware abstraction layer (HAL) for microTCA hardware. For the actual communication with the hardware the IPBus protocol is used. This protocol is transported via IP and Ethernet. The hierarchical structure of the POS is shown in figure 17.

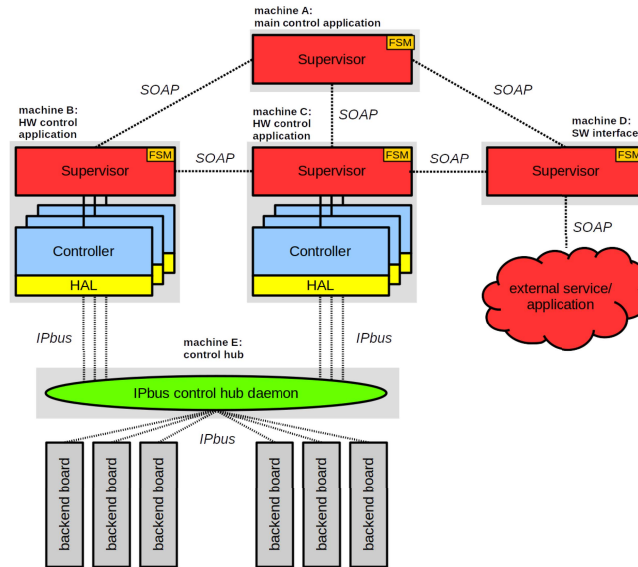


Figure 17. The hierarchical structure of the POS.

In order to prevent conflicts due to potential concurrent hardware accesses, the connection is established via a so-called control hub, which is a service daemon running on a separate computer (machine) that queues incoming requests from different applications and distributes them to the actual hardware.

CMS adopted a state machine-based approach for the control of its DAQ systems. The software must always reflect the current hardware state and must be able to perform well-defined transitions between these states when instructed from a higher control level. For this reason, an additional application layer is built on top of the controller layer, the so-called hardware supervisors. All supervisors implement a common FSM and define interfaces to the other supervisors for state changes. The cross-communication between the supervisors in POS is realized using the SOAP protocol [41]. The message format is XML. Supervisors also provide a graphical user interface using a simple web server. One single supervisor can hold a set of instances of the controller classes allowing control of several hardware boards at the same time.

A second type of supervisor exists (service supervisors), which does not control hardware, but establishes the connection to other services, like the detector control system (DCS) which is used for controlling and monitoring the detector power distribution. This interconnection between the DAQ and DCS system is described in section 8.3.

In order to operate the POS, there has to always be a main (service) supervisor, which orchestrates all the other hardware and service supervisors. This supervisor processes all commands received from the CMS Central Run Control System during global data taking and provides the main user interface during local detector calibrations.

8.2 Distributed software architecture

One advantage of the described software infrastructure is that it is scalable and can be distributed on many different computing nodes. The overall software infrastructure is defined in one common XML configuration file, such that each running process is aware of all the other existing processes in its environment. The current CMS Phase-1 pixel detector software configuration consists of 38 instances of different supervisors: 1 main supervisor (PixelSupervisor), 1 DCS service supervisor (PixelDCSFSMInterface), 1 AMC13 hardware supervisor (PixelAMC13Supervisor) controlling 12 AMC13 boards, 12 FED hardware supervisors (PixelFEDSupervisor) controlling 108 FEDs, 12 Pixel FEC hardware supervisors (PixelFECSupervisor) controlling 16 Pixel FECs, 3 Tracker FEC hardware supervisors (PixelTKFECSupervisor) controlling 3 Tracker FECs, and 8 TCDS hardware supervisors (PixelTCDSSupervisor) controlling 8 TCDS boards.

These software instances are distributed over twelve worker nodes featuring 20 cores and 32 GB RAM each. The number of computers has been chosen in order to follow the organization of the pixel detector back-end hardware in twelve microTCA crates. In addition to the twelve worker nodes, twelve additional computers act as control hubs, defining the gateways to the individual microTCA crates.

8.3 Interface to the Detector Control System

The front-end needs to be configured differently depending on the power status of the detector. For example, a sensor module becomes noisy in case of no external bias voltage. For this reason the different PixelFECSupervisors need to be informed of any state change of the power system of the pixel detector in DCS. The PixelDCSFSMInterface subscribes to the state of individual power supply channels in the DCS and evaluates the power state of a group of power supplies that power parts of the detector controlled by one PixelFECSupervisor. The summary power state is based on a single majority voting, i.e. one single different power supply state is enough to change the summary power state. The new summary state is transmitted to the corresponding supervisors and is considered in the next front-end configuration.

9 Operation performance

The CMS Phase-1 pixel detector has collected data in 2017 and 2018 with 95.5% and 94.4% functional channels, respectively. The non-functional fraction was due to infrastructure issues, caused for example by connector failures, and damaged sensor modules. Software recovery mechanisms and periodic ROC resets are implemented to reduce dead-time, ensure smooth running and maintain a high level of hit efficiency for BPIX layer 1.

9.1 Software recovery mechanisms

One important aspect of an online control and readout system is the ability to react to unexpected hardware states and guarantee the best performance of the hardware. While many of the simple problems, such as a too high trigger rate, are handled by the FEDs themselves, more subtle problems are easier to analyze and handle in software. Within the POS framework several higher level problem recovery systems are implemented, three of which will be discussed as examples here: the recovery from an SEU in the TBM, a non-responsive TBM, and a non-responsive portcard.

For the recovery of an SEU in a TBM the affected sensor module must be reprogrammed. This is handled by the Pixel FECs. The Pixel FED interface corresponding to a group of Pixel FEDs keeps the list of channels that do not send data, and if the number of channels in the list reaches a programmable value, it reports this to the FEDSupervisor. In order to have a full overview of the system the information is then sent from the FEDSupervisor to the PixelSupervisor. In the PixelSupervisor a new thread is started periodically requesting the SEU status count from all FEDSupervisors. When a programmable threshold is reached, which can differ between different parts of the detector for their impact on the data quality, a request to stop the triggers is sent to the CMS Central Run Control System. When the triggers are paused the Pixel FECs are notified to reprogram all the TBM settings. When the TBM is in a controlled state the ROC settings are reprogrammed. In order to use the time of the paused triggers effectively almost all settings are reprogrammed for the whole detector. Only the trim and mask settings are programmed specifically for the sensor modules affected by the SEU. The PixelSupervisor waits until the Pixel FECs have finished this operation and then signals to the CMS Central Run Control System to restart triggers. This procedure takes approximately five seconds.

One problem of the current version of the TBMs is that some SEUs result in a state where the TBM no longer processes triggers. The mechanism for this is understood and has been solved in the revised version of the TBM that will be used for the replacement of the innermost BPIX layer during LS2. For the TBMs currently used in the detector the only solution to revive the TBMs is a power-on reset of the TBM, which means that the low voltage supply of the TBM needs to be switched off and on again. Because of the design of the pixel detector this can be done by disabling the corresponding DC-DC converter; alternatively if a complete low voltage channel is disabled between 14 to 22 sensor modules are turned off. Using the DC-DC converters reduces the number of sensor modules being power cycled to between one and four, depending on their position in the detector. After the sensor modules are turned on, the same procedure as described above is followed to program the TBM and ROC settings.

In rare cases channels connected to a portcard stop sending data. Because of the load balancing among FEDs the readout of the modules served by one portcard is distributed. As a consequence the readout of one portcard is also distributed over several FEDs. This makes the detection of a missing portcard only possible in the PixelSupervisor, where the information from all FEDs is combined. The previously described report chain is used and if a complete portcard is recognized as having stopped to send data, the portcard is reprogrammed by the Tracker FEC. This is followed by the programming of all the affected sensor modules using Pixel FECs. If after this recovery the channels still do not send data to the FEDs, the corresponding channels are masked in the FEDs.

Figure 18 shows the number of ROCs that do not send data as a function of time in BPIX layer 1 during data taking for an LHC fill in 2017. More ROCs become inactive over time due to SEUs in the TBM ($0.7\%/100\text{ fb}^{-1}$). After a programmable threshold is reached the SEU recovery mechanism is activated as described above, during which triggers are paused. Once the triggers are resumed the number of inactive ROCs is again at the baseline value (roughly 1% of BPIX layer 1 ROCs are not functional).




Figure 18. The number of inactive ROCs as a function of time in BPIX layer 1 during a typical LHC fill in 2017. The number of inactive ROCs increases until a programmable threshold is reached, at which point the SEU recovery mechanism is activated and the ROCs are recovered. The SEU recovery mechanism can be activated several times during an LHC fill. The SEU rate depends on the instantaneous luminosity, which decreases over the time of the fill. In the fill used for this plot, the peak luminosity was around $1.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. The typical rate for inactive ROCs is one in five minutes at an instantaneous luminosity of $1.0 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ for BPIX layer 1 modules.

9.2 Periodic ROC resets

As discussed in section 3.1, the PROC600, the readout chip for BPIX layer 1, has rare data synchronization losses in double-columns that lead to lower hit efficiencies. Both at low and high trigger rates, inefficiency is caused by a timing error in the time-stamp buffer of a double-column. Here a coincidence between a new hit and an expiring hit, i.e. a recorded hit exceeding the maximum allowed latency, can generate a spurious column drain and therefore the loss of synchronization of the double-column. This desynchronizes the readout mechanism, and the next hits are not assigned to the correct event. It is more probable to observe this effect at high trigger rates. At low trigger rates another timing error can generate a spurious buffer-full signal. This happens when a buffer is empty, a hit is registered, no other hit arrives within the trigger latency and two hits are registered at exactly the trigger latency in two consecutive clock cycles. The spurious buffer-full signal by itself would not be a problem, but can lead in combination with the problem described above to a loss of

synchronization. In both cases the synchronization is restored by a reset. Both problems have been fixed in the new version of the PROC600.

In order to address these data synchronization losses, periodic ROC resets at 70 Hz are issued by TCDS. Figure 19 shows the BPIX layer 1 hit efficiency versus instantaneous luminosity with and without periodic ROC resets. Issuing resets for the ROCs recovers the hit efficiencies at low and high instantaneous luminosities. If there were no periodic resets, the sharp efficiency drop above an instantaneous luminosity of $1.3 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ would have affected the layer 1 hit efficiency drastically.

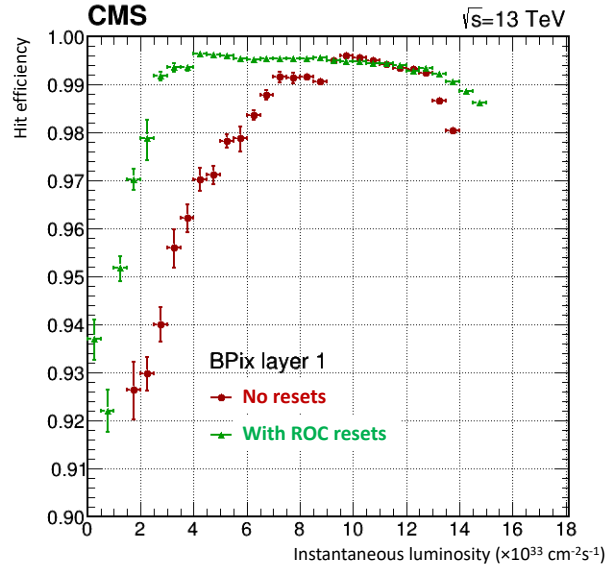


Figure 19. BPIX layer 1 hit efficiency with (green) and without (red) periodic ROC resets at 70 Hz versus instantaneous luminosity.

10 Conclusion

The CMS Phase-1 pixel detector DAQ system was developed based on a combination of custom and standard microTCA parts to satisfy the higher bandwidth requirement of the new pixel detector and to interface correctly to the upgraded front-end electronics and optical links. The DAQ system underwent a series of integration tests, including readout of the pilot pixel detector, quality assurance of the Phase-1 detector during its assembly, and testing with the CMS Central DAQ. It was tested with realistic data streams at high trigger rates (up to 100 kHz) expected during LHC running. The Phase-1 pilot detector system proved to be valuable, leading to new designs for the TBM and the portcard to address an asymmetric eye diagram and excessive clock jitter, and helping with FED firmware development. The CMS Phase-1 pixel detector achieved the required performance improvements compared to the original pixel detector, and the pixel DAQ system performed well during 2017–2018 running delivering high quality data with low dead-time consistently for CMS, without failure of any back-end parts.

Acknowledgments

The tracker groups gratefully acknowledge financial support from the following funding agencies: BMWFW and FWF (Austria); FNRS and FWO (Belgium); CERN; MSE and CSF (Croatia); Academy of Finland, MEC, and HIP (Finland); CEA and CNRS/IN2P3 (France); BMBF, DFG, and HGF (Germany); GSRT (Greece); OTKA K124850, NIH, and Bolyai Fellowship of the Hungarian Academy of Sciences (Hungary); DAE and DST (India); IPM (Iran); INFN (Italy); PAEC (Pakistan); SEIDI, CPAN, PCTI and FEDER (Spain); Swiss Funding Agencies (Switzerland); MST (Taipei); STFC (United Kingdom); DOE and NSF (U.S.A.). Individuals have received support from HFRI (Greece).

References

- [1] CMS collaboration, *The CMS experiment at the CERN LHC*, 2008 *JINST* **3** S08004.
- [2] CMS collaboration, *Description and performance of track and primary-vertex reconstruction with the CMS tracker*, 2014 *JINST* **9** P10009 [[arXiv:1405.6569](https://arxiv.org/abs/1405.6569)].
- [3] CMS collaboration, *The CMS tracker system project: Technical Design Report*, CERN-LHCC-98-006 (1997).
- [4] CMS collaboration, *CMS Technical Design Report for the Pixel Detector Upgrade*, CERN-LHCC-2012-016 (2016).
- [5] H.C. Kästli, *Frontend electronics development for the CMS pixel detector upgrade*, *Nucl. Instrum. Meth. A* **731** (2013) 88.
- [6] J. Varela, *Timing and Synchronization in the LHC Experiments*, at 6th Workshop on Electronics for LHC Experiments, Krakow Poland (2000), CERN-2000-010.
- [7] CMS collaboration, *A digital readout system for the CMS Phase I Pixel Upgrade*, 2015 *JINST* **10** C04037.
- [8] D. Gigi et al., *The FEROL40, a microTCA card interfacing custom point-to-point links and standard TCP/IP*, *PoS(TWEPP-17)* 075.
- [9] S. Baron, *Timing, Trigger and Control (TTC) systems for the LHC*, CERN, Geneva Switzerland (2014); Online: <http://ttc.web.cern.ch/TTC/intro.html>.
- [10] J. Hegeman et al., *The CMS Timing and Control Distribution System*, in *IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, San Diego U.S.A. (2015), pg. 1.
- [11] E. Hazen, A. Heister, C. Hill, J. Rohlf, S.X. Wu and D. Zou, *The AMC13XG: a new generation clock/timing/DAQ module for CMS MicroTCA*, 2013 *JINST* **8** C12036.
- [12] V. Joao, *CMS L1 Trigger Control System*, CMS-NOTE-2002-033 (2002).
- [13] CMS collaboration, *The CMS Level-1 trigger system for LHC Run II*, 2017 *JINST* **12** C03021.
- [14] P. Placidi, A. Marchioro and P. Moreira, *CMS Tracker PLL Reference Manual*, CERN, Geneva Switzerland (2000), <http://cds.cern.ch/record/1069705>.
- [15] P. Moreira and A. Marchioro, *QPLL — a quartz crystal based PLL for jitter filtering applications in LHC*, at 9th Workshop on Electronics for LHC Experiments, Amsterdam The Netherlands (2003).
- [16] H. Furtado et al., *Delay25, an ASIC for timing adjustment in LHC*, at 11th Workshop on Electronics for LHC and future Experiments, Heidelberg Germany (2005).
- [17] L. Feld et al., *The DC-DC conversion power system of the CMS Phase-I pixel upgrade*, 2015 *JINST* **10** C01052.

- [18] C. Paillard, C. Ljuslin and A. Marchioro, *The CCU25: A network oriented communication and control unit integrated circuit in a 0.25 μm CMOS technology*, [CERN-2002-003](#) (2002).
- [19] H.C. Kastli et al., *Design and performance of the CMS pixel detector readout chip*, [Nucl. Instrum. Meth. A **565** \(2006\) 188](#) [[physics/0511166](#)].
- [20] M. Rossini, *Module Prototype Qualification for the CMS Pixel Detector Upgrade*, Ph.D. Thesis, ETH Zurich, Zurich Switzerland (2015), [No. 22934](#).
- [21] J. Hoss, *Search for Supersymmetry with Multiple Charged Leptons at $\sqrt{s} = 13$ TeV with CMS and Radiation Tolerance of the Readout Chip for the Phase I Upgrade of the Pixel Detector*, Ph.D. Thesis, ETH Zurich, Zurich Switzerland (2017), [No. 24286](#).
- [22] E. Bartz, *The 0.25 μm Token Bit Manager Chip for the CMS Pixel Readout*, [CERN-2003-006](#) (2003).
- [23] CMS collaboration, *The versatile link, a common project for super-LHC*, [2009 JINST **4** P12003](#).
- [24] G. Cervelli, A. Marchioro, P. Moreira and F. Vasey, *A linear laser-driver array for optical transmission in the LHC experiments*, [Nucl. Sci. Symp. Conf. Rec. **2** \(2000\) 1](#).
- [25] T. Uemura et al., *1060-nm 10-Gb/s* 12-channel parallel-optical modules for optical interconnects*, [IEEE CPMT Symp. Japan **2010** \(2010\) 1](#).
- [26] J. Troska et al., *Optical readout and control systems for the CMS tracker*, [IEEE Trans. Nucl. Sci. **50** \(2003\) 1067](#).
- [27] <https://www.picmg.org/openstandards/microtca/>.
- [28] M. Pesaresi et al., *The FC7 AMC for generic DAQ & control applications in CMS*, [2015 JINST **10** C03036](#).
- [29] CMS collaboration, *Deployment of the CMS Tracker AMC as backend for the CMS pixel detector*, [2016 JINST **11** C01056](#).
- [30] *Kintex-7 FPGAs Data Sheet*, XILINX DS182.
- [31] <https://www.xilinx.com/products/design-tools/vivado.html>.
- [32] K. Kloukinas et al., *FEC-CCS: A common Front-End Controller card for the CMS detector electronics*, [CERN 2007-001](#) (2006).
- [33] C.Ghabrous Larrea et al., *IPbus: a flexible Ethernet-based control system for xTCA hardware*, [2015 JINST **10** C02019](#).
- [34] B. Akgün, *Pilot system for the Phase-1 pixel upgrade of CMS*, [PoS\(VERTEX2015\)018](#).
- [35] CMS TRACKER collaboration, *Commissioning of the Phase-1 Upgrade of the CMS Pixel Detector*, [Springer Proc. Phys. **213** \(2018\) 385](#).
- [36] J. Ngadiuba, *Testing and Integration of the Service Cylinders for the CMS Phase-1 Pixel Detector*, *Proceedings of TWEPP-16*, Karlsruhe Germany (2016), [CMS-CR-2016/357](#) (2016).
- [37] CMS collaboration, *Integration and testing of the DAQ system for the CMS Phase 1 pixel upgrade*, [2017 JINST **12** C02078](#).
- [38] M. Barros Marin et al., *A GLIB-based uTCA demonstration system for HEP experiments*, [2013 JINST **8** C12011](#).
- [39] V. Brigljevic et al., *Using XDAQ in application scenarios of the CMS experiment*, FERMILAB-CONF-03-293, CHEP-2003-MOGT008 (2003).
- [40] T. Goodale et al., *The Cactus Framework and Toolkit: Design and Applications*, High Performance Computing for Computational Science (VECPAR 2002), at *5th International Conference*, Porto Portugal (2002).
- [41] D. Box et al., *Simple Object Access Protocol (SOAP) 1.1*, W3C Note (2000).

Tracker group of the CMS collaboration

Institut für Hochenergiephysik, Wien, Austria

W. Adam, T. Bergauer, D. Blöchl, E. Brondolin¹, M. Dragicevic, R. Frühwirth², V. Hinger, H. Steininger

Universiteit Antwerpen, Antwerpen, Belgium

W. Beaumont, D. Di Croce, X. Janssen, J. Lauwers, P. Van Mechelen, N. Van Remortel

Vrije Universiteit Brussel, Brussel, Belgium

F. Blekman, S.S. Chhibra, J. De Clercq, J. D'Hondt, S. Lowette, I. Marchesini, S. Moortgat, Q. Python, K. Skovpen, E. Sørensen Bols, P. Van Mulders

Université Libre de Bruxelles, Bruxelles, Belgium

Y. Allard, D. Beghin, B. Bilin, H. Brun, B. Clerbaux, G. De Lentdecker, H. Delannoy, W. Deng, L. Favart, R. Goldouzian, A. Grebenyuk, A. Kalsi, I. Makarenko, L. Moureaux, A. Popov, N. Postiau, F. Robert, Z. Song, L. Thomas, P. Vanlaer, D. Vannerom, Q. Wang, H. Wang, Y. Yang

Université Catholique de Louvain, Louvain-la-Neuve, Belgium

O. Bondu, G. Bruno, C. Caputo, P. David, C. Delaere, M. Delcourt, A. Giammanco, G. Krintiras, V. Lemaître, A. Magitteri, K. Piotrkowski, A. Saggio, N. Szilasi, M. Vidal Marono, P. Vischia, J. Zobeć

Institut Ruđer Bošković, Zagreb, Croatia

V. Brigljević, S. Ceci, D. Ferencsek, M. Roguljić, A. Starodumov³, T. Šuša

Department of Physics, University of Helsinki, Helsinki, Finland

P. Eerola, J. Heikkilä

Helsinki Institute of Physics, Helsinki, Finland

E. Brücken, T. Lampén, P. Luukka, L. Martikainen, E. Tuominen

Lappeenranta University of Technology, Lappeenranta, Finland

T. Tuuva

Université de Strasbourg, CNRS, IPHC UMR 7178, Strasbourg, France

J.-L. Agram⁴, J. Andrea, D. Bloch, C. Bonnin, G. Bourgatte, J.-M. Brom, E. Chabert, L. Charles, E. Dangelser, D. Gelé, U. Goerlach, L. Gross, J. Hosselet, M. Krauth, N. Tonon

Université de Lyon, Université Claude Bernard Lyon 1, CNRS-IN2P3,

Institut de Physique Nucléaire de Lyon, Villeurbanne, France

G. Baulieu, G. Boudoul, L. Caponetto, N. Chanon, D. Contardo, P. Dené, T. Dupasquier, G. Galbit, N. Lumb, L. Mirabito, B. Nodari, S. Perries, M. Vander Donckt, S. Viret

RWTH Aachen University, I. Physikalisches Institut, Aachen, Germany

C. Autermann, L. Feld, W. Karpinski, M.K. Kiesel, K. Klein, M. Lipinski, D. Meuser, A. Ostapchuk, A. Pauls, G. Pierschel, M. Preuten, M. Rauch, N. Röwert, S. Schael, J. Schulz, G. Schwering, M. Teroerde, M. Wlochal, V. Zhukov

RWTH Aachen University, III. Physikalisches Institut B, Aachen, Germany

C. Dziwok, G. Fluegge, T. Müller, O. Pooth, A. Stahl, T. Ziemons

Deutsches Elektronen-Synchrotron, Hamburg, Germany

M. Aldaya, C. Asawatangtrakuldee, G. Eckerlin, D. Eckstein, T. Eichhorn, E. Gallo, M. Guthoff, M. Haranko, A. Harb, J. Keaveney, C. Kleinwort, R. Mankel, H. Maser, M. Meyer, M. Missiroli, C. Muhl, A. Mussgiller, D. Pitzl, O. Reichelt, M. Savitskyi, P. Schuetze, R. Stever, R. Walsh, A. Zuber

University of Hamburg, Hamburg, Germany

A. Benecke, H. Biskop, P. Buhmann, A. Ebrahimi, M. Eich, F. Feindt, A. Froehlich, E. Garutti, P. Gunnellini, J. Haller, A. Hinemann, G. Kasieczka, R. Klanner, V. Kutzner, T. Lange, M. Matysek, M. Mrowietz, C. Niemeyer, Y. Nissan, K. Pena, A. Perieanu, O. Rieger, P. Schleper, J. Schwandt, D. Schwarz, J. Sonneveld, G. Steinbrück, A. Tews, B. Vormwald, J. Wellhausen, I. Zoi

Institut für Experimentelle Kernphysik, Karlsruhe, Germany

M. Abbas, L. Ardila, M. Balzer, C. Barth, T. Barvich, M. Baselga, T. Blank, F. Bögelspacher, E. Butz, M. Caselle, W. De Boer, A. Dierlamm, K. El Morabit, J.-O. Gosewisch, F. Hartmann, U. Husemann, R. Koppenhöfer, S. Kudella, S. Maier, S. Mallows, M. Metzler, Th. Muller, M. Neufeld, A. Nürnberg, O. Sander, D. Schell, M. Schröder, T. Schuh, I. Shvetsov, H.-J. Simonis, P. Steck, M. Wassmer, M. Weber, A. Weddigen

Institute of Nuclear and Particle Physics (INPP), NCSR Demokritos, Aghia Paraskevi, Greece

G. Anagnostou, P. Asenov, P. Assiouras, G. Daskalakis, A. Kyriakis, D. Loukas, L. Paspalaki

Wigner Research Centre for Physics, Budapest, Hungary

T. Balázs, F. Siklér, T. Vámi, V. Veszprémi

University of Delhi, Delhi, India

A. Bhardwaj, C. Jain, G. Jain, K. Ranjan

Saha Institute of Nuclear Physics, Kolkata, India

R. Bhattacharya, S. Dutta, S. Roy Chowdhury, G. Saha, S. Sarkar

INFN Sezione di Bari^a, Università di Bari^b, Politecnico di Bari^c, Bari, Italy

P. Cariola^a, D. Creanza^{a,c}, M. de Palma^{a,b}, G. De Robertis^a, L. Fiore^a, M. Ince^{a,b}, F. Loddo^a, G. Maggi^{a,c}, S. Martiradonna^a, M. Mongelli^a, S. My^{a,b}, G. Selvaggi^{a,b}, L. Silvestris^a

INFN Sezione di Catania^a, Università di Catania^b, Catania, Italy

S. Albergo^{a,b}, S. Costa^{a,b}, A. Di Mattia^a, R. Potenza^{a,b}, M.A. Saizu^{a,5}, A. Tricomi^{a,b}, C. Tuve^{a,b}

INFN Sezione di Firenze^a, Università di Firenze^b, Firenze, Italy

G. Barbagli^a, M. Brianzi^a, A. Cassese^a, R. Ceccarelli^{a,b}, R. Ciaranfi^a, V. Ciulli^{a,b}, C. Civinini^a, R. D'Alessandro^{a,b}, E. Focardi^{a,b}, G. Latino^{a,b}, P. Lenzi^{a,b}, M. Meschini^a, S. Paoletti^a, L. Russo^{a,b}, E. Scarlini^{a,b}, G. Sguazzoni^a, L. Viliani^{a,b}

INFN Sezione di Genova^a, Università di Genova^b, Genova, Italy

S. Cerchi^a, F. Ferro^a, R. Mulargia^{a,b}, E. Robutti^a

INFN Sezione di Milano-Bicocca^a, Università di Milano-Bicocca^b, Milano, Italy

F. Brivio^{a,b}, M.E. Dinardo^{a,b}, P. Dini^a, S. Gennai^a, L. Guzzi, S. Malvezzi^a, D. Menasce^a,
L. Moroni^a, D. Pedrini^a, D. Zuolo^{a,b}

INFN Sezione di Padova^a, Università di Padova^b, Padova, Italy

P. Azzi^a, N. Bacchetta^a, D. Bisello^a, T. Dorigo^a, N. Pozzobon^{a,b}, M. Tosi^{a,b}

INFN Sezione di Pavia^a, Università di Bergamo^b, Bergamo, Italy

F. De Canio^{a,b}, L. Gaioni^{a,b}, M. Manghisoni^{a,b}, L. Ratti^a, V. Re^{a,b}, E. Riceputi^{a,b},
G. Traversi^{a,b}

INFN Sezione di Perugia^a, Università di Perugia^b, CNR-IOM Perugia^c, Perugia, Italy

G. Baldinelli^{a,b}, F. Bianchi^{a,b}, M. Biasini^{a,b}, G.M. Bilei^a, S. Bizzaglia^a, M. Caprai^a,
C. Cecchi^{a,b}, B. Checcucci^a, D. Ciangottini^a, L. Fanò^{a,b}, L. Farnesini^a, M. Ionica^a,
R. Leonardi^{a,b}, E. Manoni^a, G. Mantovani^{a,b}, V. Mariani^{a,b}, M. Menichelli^a, A. Morozzi^a,
F. Moscatelli^{a,c}, D. Passeri^{a,b}, P. Placidi^{a,b}, A. Rossi^{a,b}, A. Santocchia^{a,b}, D. Spiga^a,
L. Storchi^a, C. Turrioni^{a,b}

INFN Sezione di Pisa^a, Università di Pisa^b, Scuola Normale Superiore di Pisa^c, Pisa, Italy

K. Androsov^a, P. Azzurri^a, G. Bagliesi^a, A. Basti^a, R. Beccherle^a, V. Bertacchi^{a,c}, L. Bianchini^a,
T. Boccali^a, L. Borrello^a, F. Bosi^a, R. Castaldi^a, M.A. Ciocci^{a,b}, R. Dell'Orso^a, G. Fedì^a,
F. Fiori^{a,c}, L. Giannini^{a,c}, A. Giassi^a, M.T. Grippo^{a,b}, F. Ligabue^{a,c}, G. Magazzu^a, E. Manca^{a,c},
G. Mandorli^{a,c}, E. Mazzoni^a, A. Messineo^{a,b}, A. Moggi^a, F. Morsani^a, F. Palla^a, F. Palmonari^a,
F. Raffaelli^a, A. Rizzi^{a,b}, P. Spagnolo^a, R. Tenchini^a, G. Tonelli^{a,b}, A. Venturi^a, P.G. Verdini^a

INFN Sezione di Torino^a, Università di Torino^b, Politecnico di Torino^c, Torino, Italy

R. Bellan^{a,b}, M. Costa^{a,b}, R. Covarelli^{a,b}, G. Dellacasa^a, N. Demaria^a, A. Di Salvo^{a,c},
G. Mazza^a, E. Migliore^{a,b}, E. Monteil^{a,b}, L. Pacher^a, A. Paterno^{a,c}, A. Rivetti^a, A. Solano^{a,b}

Vilnius University, Vilnius, Lithuania

D. Simelevicius

Instituto de Física de Cantabria (IFCA), CSIC-Universidad de Cantabria, Santander, Spain

E. Curras Rivera, J. Duarte Campderros, M. Fernandez, G. Gomez, F.J. Gonzalez Sanchez,
R. Jaramillo Echeverria, D. Moya, E. Silva Jimenez, I. Vila, A.L. Virto

CERN, European Organization for Nuclear Research, Geneva, Switzerland

D. Abbaneo, I. Ahmed, B. Akgun, E. Albert, G. Auzinger, J. Bendotti, G. Berruti, G. Blanchot,
F. Boyer, A. Caratelli, D. Ceresa, J. Christiansen, K. Cichy, J. Daguin, N. Deelen⁶, S. Detraz,
D. Deyrail, M. Dobson, N. Emriskova⁷, B. Engegaard, F. Faccio, A. Filenius, N. Frank, T. French,
J. Fulcher, R. Gajanec, D. Gigi, F. Glege, M. Hansen, J. Hegeman, A. Honma, G. Hugo, W. Hulek,
L.M. Jara Casas, J. Kaplon, K. Kloukinas, A. Kornmayer, N. Koss, L. Kottelat, D. Koukola,
M. Kovacs, A. La Rosa, P. Lenoir, R. Loos, A. Marchioro, S. Marconi, F. Meijers, S. Mersi,
E. Meschi, S. Michelis, C. Nieto Martin, A. Onnela, S. Orfanelli, L. Orsini, T. Pakulski, A. Perez,
F. Perez Gomez, J.-F. Pernot, P. Petagna, Q. Piazza, H. Postema, T. Prousalidi, R. Puente Rico⁸,

A. Racz, A. Remigiusz Labaza, H. Sakulin, S. Scarff⁹, S. Spathopoulos, S. Sroka, P. Tropea, J. Troska, A. Tsirou, F. Vasey, P. Vichoudis

Paul Scherrer Institut, Villigen, Switzerland

W. Bertl[†], L. Caminada¹⁰, K. Deiters, W. Erdmann, R. Horisberger, H.-C. Kaestli, D. Kotlinski, U. Langenegger, B. Meier, T. Rohe, S. Streuli

Institute for Particle Physics, ETH Zurich, Zurich, Switzerland

F. Bachmair, M. Backhaus, R. Becker, P. Berger, D. di Calafiori, L. Djambazov, M. Donega, C. Grab, D. Hits, J. Hoss, W. Lustermann, M. Masciovecchio, M. Meinhard, V. Perovic, L. Perozzi, B. Ristic, U. Roeser, D. Ruini, V. Tavoraro, R. Wallny, D. Zhu

Universität Zürich, Zurich, Switzerland

T. Aarrestad, C. Amsler¹¹, K. Bösigler, F. Canelli, V. Chiochia, A. De Cosa, R. Del Burgo, C. Galloni, B. Kilminster, S. Leontsinis, R. Maier, G. Rauco, P. Robmann, Y. Takahashi, A. Zucchetta

National Taiwan University (NTU), Taipei, Taiwan

P.-H. Chen, W.-S. Hou, R.-S. Lu, M. Moya, J.F. Tsai

University of Bristol, Bristol, United Kingdom

D. Burns, E. Clement, D. Cussans, J. Goldstein, S. Seif El Nasr-Storey

Rutherford Appleton Laboratory, Didcot, United Kingdom

J.A. Coughlan, K. Harder, K. Manolopoulos, I.R. Tomalin

Imperial College, London, United Kingdom

R. Bainbridge, J. Borg, G. Hall, T. James, M. Pesaresi, S. Summers, K. Uchida

Brunel University, Uxbridge, United Kingdom

J. Cole, C. Hoad, P. Hobson, I.D. Reid

The Catholic University of America, Washington DC, U.S.A.

R. Bartek, A. Dominguez, R. Uniyal

Boston University, Boston, U.S.A.

Z. Demiragli, E. Hazen, J. Rohlf

Brown University, Providence, U.S.A.

G. Altopp, B. Burkle, C. Chen, X. Coubez, Y.-T. Duh, M. Hadley, U. Heintz, N. Hinton, J. Hogan¹², A. Korotkov, J. Lee, M. Narain, S. Sagir¹³, E. Spencer, R. Syarif, V. Truong, E. Usai, J. Voelker

University of California, Davis, Davis, U.S.A.

M. Chertok, J. Conway, G. Funk, F. Jensen, R. Lander, S. Macaudo, D. Pellett, J. Thomson, R. Yohay¹⁴, F. Zhang

University of California, Los Angeles, U.S.A.

S. Erhan

University of California, Riverside, Riverside, U.S.A.

G. Hanson, W. Si

University of California, San Diego, La Jolla, U.S.A.

R. Gerosa, A. Holzner, S. Krutelyov, V. Sharma, A. Yagil

University of California, Santa Barbara - Department of Physics, Santa Barbara, U.S.A.

O. Colegrove, V. Dutta, L. Gouskos, J. Incandela, S. Kyre, H. Qu, M. Quinnan, D. White

University of Colorado Boulder, Boulder, U.S.A.

J.P. Cumalat, W.T. Ford, E. MacDonald, A. Perloff, K. Stenson, K.A. Ulmer, S.R. Wagner

Cornell University, Ithaca, U.S.A.

J. Alexander, Y. Cheng, J. Chu, J. Conway, D. Cranshaw, A. Datta, K. McDermott, J. Monroy, Y. Bordlemay Padilla, D. Quach, A. Rinkevicius¹⁵, A. Ryd, L. Skinnari, L. Soffi, C. Strohman, Z. Tao, J. Thom, J. Tucker, P. Wittich, M. Zientek

Fermi National Accelerator Laboratory, Batavia, U.S.A.

M. Alyari, A. Bakshi, G. Bolla[†], K. Burkett, J.N. Butler, A. Canepa, H.W.K. Cheung, J. Chramowicz, G. Derylo, A. Ghosh, C. Gingu, H. Gonzalez, S. Grünendahl, S. Hasegawa, J. Hoff, M. Johnson, C.M. Lei, R. Lipton, M. Liu, S. Los, M. Matulik, P. Merkel, R. Mommsen, S. Nahn, A. Prosser, F. Ravera, R. Rivera, B. Schneider, W.J. Spalding, L. Spiegel, S. Timpone, L. Uplegger, E. Voirin, H.A. Weber, P. Zejdl

University of Illinois at Chicago (UIC), Chicago, U.S.A.

D.R. Berry, X. Chen, S. Dittmer, A. Evdokimov, O. Evdokimov, C.E. Gerber, D.J. Hofman, C. Mills

The University of Iowa, Iowa City, U.S.A.

M. Alhusseini, S. Durgut, J. Nachtman, Y. Onel, C. Rude, C. Snyder, K. Yi¹⁶

Johns Hopkins University, Baltimore, U.S.A.

N. Eminizer, A. Gritsan, P. Maksimovic, J. Roskes, M. Swartz, M. Xiao

The University of Kansas, Lawrence, U.S.A.

P. Baringer, A. Bean, S. Khalil, A. Kropivnitskaya, D. Majumder, E. Schmitz, G. Wilson

Kansas State University, Manhattan, U.S.A.

A. Ivanov, R. Mendis, T. Mitchell, A. Modak, N. Skhirladze, R. Taylor

University of Mississippi, Oxford, U.S.A.

J.G. Acosta, L.M. Cremaldi, S. Oliveros, L. Perera, D. Summers

University of Nebraska-Lincoln, Lincoln, U.S.A.

K. Bloom, D.R. Claes, C. Fangmeier, F. Golf, I. Kravchenko, J. Siado

State University of New York at Buffalo, Buffalo, U.S.A.

I. Iashvili, A. Kharchilava, C. McLean, D. Nguyen, A. Parker, J. Pekkanen, S. Rappoccio

Northwestern University, Evanston, U.S.A.

K. Hahn, Y. Liu, K. Sung

The Ohio State University, Columbus, U.S.A.

J. Alimena, B. Cardwell, B. Francis, C.S. Hill

University of Puerto Rico, Mayaguez, U.S.A.

S. Malik, S. Norberg, J.E. Ramirez Vargas

Purdue University, West Lafayette, U.S.A.

S. Das, M. Jones, A. Jung, A. Khaliwada, G. Negro, J. Thieman

Purdue University Northwest, Hammond, U.S.A.

T. Cheng, J. Dolen, N. Parashar

Rice University, Houston, U.S.A.

K.M. Ecklund, S. Freed, M. Kilpatrick, T. Nussbaum

University of Rochester, Rochester, U.S.A.

R. Demina, J. Dulemba, O. Hindrichs

Rutgers, The State University of New Jersey, Piscataway, U.S.A.

E. Bartz, A. Gandrakotra, Y. Gershtein, E. Halkiadakis, A. Hart, S. Kyriacou, A. Lath, K. Nash, M. Osherson, S. Schnetzer, R. Stone

Texas A&M University, College Station, U.S.A.

R. Eusebi

Vanderbilt University, Nashville, U.S.A.

P. D'Angelo, W. Johns, K.O. Padeken

Visitor from Université Saint-Joseph de Beyrouth, Beirut, Lebanon

W. Karimeh

†: Deceased

1: Now at CERN, European Organization for Nuclear Research, Geneva, Switzerland

2: Also at Vienna University of Technology, Vienna, Austria

3: Also at Institute for Theoretical and Experimental Physics, Moscow, Russia

4: Also at Université de Haute-Alsace, Mulhouse, France

5: Also at Horia Hulubei National Institute of Physics and Nuclear Engineering (IFIN-HH), Bucharest, Romania

6: Also at Institut für Experimentelle Kernphysik, Karlsruhe, Germany

7: Also at Université de Strasbourg, CNRS, IPHC UMR 7178, Strasbourg, France

8: Also at Instituto de Física de Cantabria (IFCA), CSIC-Universidad de Cantabria, Santander, Spain

9: Also at École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland

10: Also at Universität Zürich, Zurich, Switzerland

11: Also at Albert Einstein Center for Fundamental Physics, Bern, Switzerland

12: Now at Bethel University, St. Paul, Minnesota, U.S.A.

13: Now at Karamanoglu Mehmetbey University, Karaman, Turkey

14: Now at Florida State University, Tallahassee, U.S.A.

15: Now at Vilnius University, Vilnius, Lithuania

16: Also at Nanjing Normal University, Nanjing, China